

“All-in-One Is All You Need.”

ALL-IN-ONE

SSCP[®]

Systems Security Certified Practitioner

EXAM GUIDE

THIRD EDITION

*Complete coverage
of all seven domains of
the SSCP exam*

■
*Includes updated
exam objectives effective
November 1, 2018*

■
*Ideal as both a study tool
and an on-the-job reference*

■
*Filled with practice exam
questions and
in-depth explanations*

DIGITAL
CONTENT
INCLUDES:

- Two practice exams

Mc
Graw
Hill
Education

DARRIL GIBSON

SSCP, CISSP[®], CompTIA Security+[®], CASP, MCSE, MCITP, ITIL[®]

3/28/2020

ALL ■ IN ■ ONE

SSCP[®]

Systems Security Certified Practitioner

EXAM GUIDE

Third Edition

Darril Gibson



New York Chicago San Francisco
Athens London Madrid Mexico City
Milan New Delhi Singapore Sydney Toronto

McGraw-Hill Education is an independent entity from (ISC)² and is not affiliated with (ISC)² in any manner. This study/training guide and/or material is not sponsored by, endorsed by, or affiliated with (ISC)² in any manner. This publication and accompanying media may be used in assisting students to prepare for the SSCP[®] exam. Neither (ISC)² nor McGraw-Hill Education warrant that use of this publication and accompanying media will ensure passing any exam. (ISC)², CISSP[®], CAP[®], ISSAP[®], ISSEP[®], ISSMP[®], SSCP[®], CSSLP[®], CCSP[®], and CBK[®] are trademarks or registered trademarks of (ISC)² in the United States and certain other countries. All other trademarks are trademarks of their respective owners.

CONTENTS AT A GLANCE

Chapter 1 Security Fundamentals

Chapter 2 Access Controls

Chapter 3 Basic Networking and Communications

Chapter 4 Advanced Networking and Communications

Chapter 5 Attacks

Chapter 6 Malicious Code and Activity

Chapter 7 Risk, Response, and Recovery

Chapter 8 Monitoring and Analysis

Chapter 9 Controls and Countermeasures

Chapter 10 Auditing and Management Processes

Chapter 11 Security Operations

Chapter 12 Security Administration and Planning

Chapter 13 Legal Issues

Chapter 14 Cryptography

Appendix About the Online Content

Glossary

Index

Exam Objective Map

The SSCP exam is composed of questions from seven domains:

- Access Controls (16%)
- Security Operations and Administration (15%)
- Risk Identification, Monitoring, and Analysis (15%)
- Incident Response and Recovery (13%)
- Cryptography (10%)
- Network and Communications Security (16%)
- Systems and Application Security (15%)

CHAPTER 1

Security Fundamentals

In this chapter, you will learn about

- Requirements to earn the (ISC)² Systems Security Certified Practitioner (SSCP) certification
- Primary goals of security related to confidentiality, integrity, and availability (CIA)
- Fundamental security terminology

Reviewing the Requirements for SSCP

The Systems Security Certified Practitioner (SSCP) certification is one of the certifications sponsored by the International Information Systems Security Certification Consortium, Inc., more commonly known as (ISC)².

You must complete the following requirements to earn the certification:

- Have at least one year of cumulative work experience in one or more of the (ISC)² SSCP domains
- Register for the exam
- Pass the exam
- Subscribe to the Code of Ethics and get endorsed



EXAM TIP Earning the certification is more than just passing the exam. You must also have one year of experience and commit to the Code of Ethics.

If you've earned certifications from other vendors (such as CompTIA), you'll find that the (ISC)² process is different. It is often confusing to people the first time they earn an (ISC)² certification. The following sections explain the process of registering for the exam, some key information about the exam, the requirement to submit paperwork to validate your experience, and requirements to maintain the SSCP certification.

Have One Year of Experience

To earn the SSCP certification, you need to have a minimum of at least one year of cumulative paid full-time work experience in one or more of the seven (ISC)² SSCP domains. The domains are as follows:

- Access Controls
- Security Operations and Administration
- Risk Identification, Monitoring, and Analysis
- Incident Response and Recovery
- Cryptography
- Network and Communications Security
- Systems and Application Security



TIP These domains represent the seven major categories of information in the SSCP Common Body of Knowledge (CBK). The CBK is a group of topics updated periodically by subject-matter experts.

(ISC)² has added an alternative to the one year of experience requirement. If you don't have the experience but do have a bachelors or master's degree in a cybersecurity program, you can use that instead. The degree must be from an accredited college or university and either focused on cybersecurity or in one of the preapproved degree programs listed by (ISC)².

When you register for the exam, you identify the number of years' experience that you have in any of the domains. After you take and pass the exam, you're required to submit a résumé that documents this experience. Additionally, you must find an (ISC)² certified professional (in good standing) to validate your experience and submit an endorsement form on your behalf. See the section "Getting Endorsed" later in the chapter for more details.

If you don't have one year of experience, you can still earn the Associate of (ISC)² designation. You still need to complete the other requirements, including passing the exam and subscribing to the Code of Ethics. You'll then have up to two years to obtain the required experience and submit the endorsement form to convert your status from Associate of (ISC)² to become a fully certified SSCP.

Registering for the Exam

The exams are computer based and administered at Pearson VUE test centers. You can register for exams through the Pearson VUE website (www.pearsonvue.com/isc2/). If you register online, you'll need to have or create a Pearson VUE account. The Pearson VUE website also includes a search feature so that you can locate a test center near you.



NOTE The SSCP exam was previously paper based and proctored in large conference rooms. While (ISC)² has moved to Computer-Based Testing (CBT), it does occasionally authorize paper-based exams on a limited basis, such as after an official training seminar. The primary portal for all SSCP information (www.isc2.org/sscp) includes up-to-date information on the exam, including current exam prices.

Registering for the exam includes three steps:

1. Complete the examination agreement.
2. Answer questions related to criminal history and background.
3. Submit the exam fee.

Complete the Examination Agreement

The examination agreement requires that you confirm your assertions regarding your minimum one year of professional experience and requires that you legally commit to the adherence of the (ISC)² Code of Ethics. Refer to the section “Subscribing to the (ISC)² Code of Ethics and Getting Endorsed” later in this chapter for an overview of the Code of Ethics.

Answer Questions Related to Criminal History and Background

The registration process requires you to answer several questions related to your history and background. These questions ask about the following topics:

- Felony convictions
- Involvement with hackers or hacking
- Revocation of any professional licenses, certifications, etc. or disciplinary actions
- Any use of aliases or pseudonyms

Answering yes to any of these questions doesn't necessarily disqualify you. However, it's best to resolve a potential problem before taking the exam. You can e-mail (ISC)² at legal@isc2.org to discuss your situation. If appropriate, (ISC)² will declare you eligible to take the exam and pursue the certification.

Submit the Exam Fee

You submit the exam fee through Pearson Vue, the testing provider. The page www.vue.com/isc2/includes links you can use to create a Pearson Vue account, find a testing center close to you, and register for the exam. Some organizations purchase vouchers in bulk and give them to their employees. If you have a voucher, you can use it instead of submitting a fee.

Passing the Exam

The exam traditionally includes 125 multiple-choice questions and you have three hours to complete it. Most questions are stand-alone multiple-choice questions. However, you may encounter some scenario-type questions that start with a paragraph or two describing the scenario, followed by two or more multiple-choice questions about the scenario. For example, you may be given a scenario that explains an organization's security goals, followed by questions that ask how best to implement those goals.

(ISC)² maintains a large pool of questions and regularly adds new questions to the pool. However, (ISC)² typically tests the new questions on actual exams before using them in the exam score. Up to 25 questions on any exam might be ungraded test questions that (ISC)² has included for research purposes. You can think of these as beta questions that (ISC)² analyzes for the questions' effectiveness. For example, if everyone answers the beta question correctly, (ISC)² may conclude that the question is too easy and needs to be revised. If everyone answers the question incorrectly, (ISC)² will likely conclude that something is wrong with the question itself. Through this type of test question analysis, (ISC)² attempts to identify problematic questions and revise them before including them in the questions that are graded. You won't know which questions are graded questions and which questions are beta-type questions. In other words, you must treat each question as though it's a valid question.

A score of 700 out of a possible 1,000 points is required to pass the exam. However, that doesn't necessarily mean that you'll pass the exam if you answer 70 out of 100 graded questions correctly, because questions aren't weighted the same. Out of the 100 valid questions, some questions may be worth 10 points, more difficult questions may be worth more than 10 points, and easier questions may be worth less. Just as (ISC)² doesn't tell you which questions are graded; it doesn't advertise the actual value of any question.

(ISC)² derives the examination questions from the SSCP CBK. The SSCP Exam Outline is an excellent source to see which topics are tested. The last page includes a short section on Supplementary References, along with a link to list of references. The goal of this book is to compile the relevant knowledge from these references into a single source. However, I strongly encourage you to download and review a copy of the current Exam Outline to ensure the document hasn't changed. The Exam Outline includes a listing of objectives and subobjectives

for each of the domains. You'll find an Exam Objective Map in the Introduction of this book. It lists the objectives that went live on April 15, 2018, and maps them to the chapter that covers the objectives.



TIP You can access the SSCP Exam Outline from here:
<https://www.isc2.org/Certifications/SSCP>.

SSCP to CISSP

The (ISC)² Certified Information Systems Security Professional (CISSP) certification is one of the top security certifications. It requires five years of experience in two or more of the eight (ISC)² CISSP domains. It is a highly respected certification and opens many doors of opportunity for those who earn it.

However, not everyone has five years of experience. The good news is that you can request a waiver of one year of experience if you have the (ISC)² SSCP certification. In other words, you only need four years of experience to earn the CISSP if you get the SSCP certification first and receive a waiver.

Many people use the (ISC)² SSCP as a stepping stone to the CISSP. By first achieving the SSCP, you gain an understanding of the (ISC)² certification process. Further, you'll find that the knowledge you gain studying for the SSCP certification helps you to build a solid foundation for the CISSP exam.

Predicting Future Question Types

It's worth noting that (ISC)² added what it calls "innovative Drag & Drop and Hotspot" questions to the CISSP exam. (ISC)² started using these new types of questions in January 2014, which was two years after the release of the CISSP Common Body of Knowledge (CBK) for that version of the exam. (ISC)² didn't change the objectives, but instead changed the way some of the questions are presented. (ISC)² announced the changes a couple of months before implementing them.

(ISC)² has not given any indication that it will implement these types of questions in the SSCP exam. However, it's entirely possible that (ISC)² will start adding them at some point, just as it did with the CISSP exam. If it does so prior to your exam, don't let the questions scare you. If you understand the content, you should be able to answer the questions without too much trouble. The following sections give you an idea of how to handle these types of questions.



TIP I regularly post blog articles on the Get Certified Get Ahead website (<http://blogs.getcertifiedgetahead.com/>). If (ISC)² does announce its intent to add these types of questions, I will post information there to help keep readers informed.

Drag & Drop Questions In a Drag & Drop question, you simply drag an object from one area of the screen to another. As an example, the following question asks you to identify symmetric encryption algorithms from a list:

Question: Which of the following algorithms are examples of symmetric encryption? Drag and drop the correct answers from the list on the left to the box on the right.

[Figure 1-1](#) shows the initial display you might see, along with how the display looks after you have correctly answered the question.

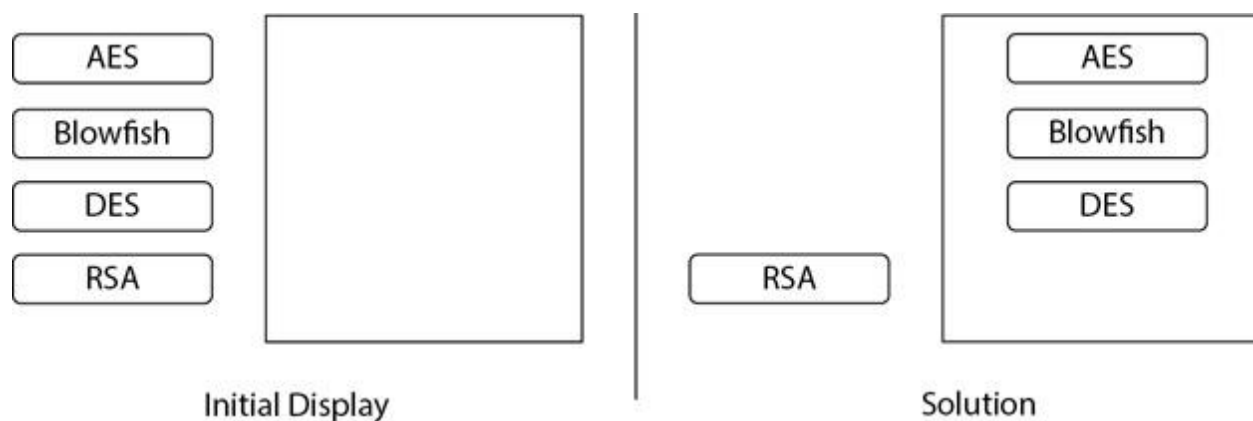


Figure 1-1 Example Drag & Drop question

This is really just another way of presenting a multiple-choice question that asks you to select all the correct answers. Here's an example of a multiple-choice question that tests the same knowledge:

Question: Which of the following algorithms are examples of symmetric encryption? (Select all that apply.)

- A. AES
- B. Blowfish
- C. 3DES
- D. RSA

While you won't see the correct answer and an explanation when you take the exam, you might like this knowledge. For clarity, here is the correct answer and an explanation.

Answer: The correct answers are A, B, and C. Advanced Encryption Standard (AES), Blowfish, and Triple Data Encryption Standard (3DES) are all symmetric encryption algorithms.

D is incorrect. RSA (named after its authors Rivest, Shamir, and Adleman) is an asymmetric encryption algorithm.

Hotspot Questions In a Hotspot question, you simply click on an area of the screen to select the correct answer. As an example, the following question asks you to identify the best location for a public-facing web server:

Question: Your organization is planning to deploy a web server. It needs to be accessible via the Internet and will access a database on the internal network. Where should they locate the new web server? (Click on the area of the diagram to indicate your choice.)

[Figure 1-2](#) shows the graphics for this question. The numbers (1 through 4) are for the explanation only and probably would not be on an actual Hotspot question.

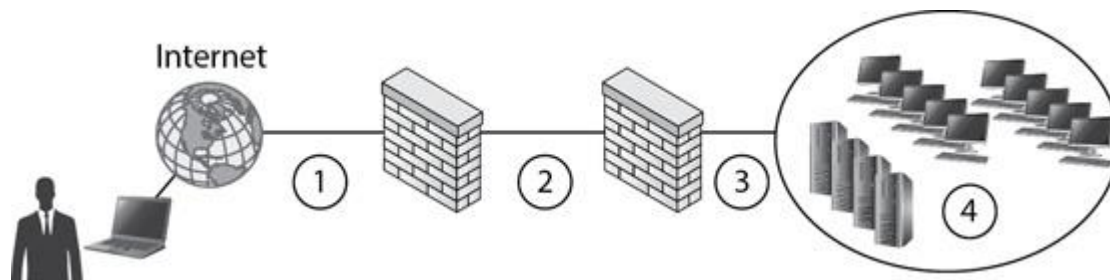


Figure 1-2 Example Hotspot question

The Hotspot format is also just another way of asking a question to test your knowledge. Here's an example of a similar question worded as a multiple-choice question:

Question: Your organization is planning to deploy a web server. It needs to be accessible via the Internet and will access a database on the internal network. Where should they locate the new web server?

- A. On the Internet side of the firewalls (the area marked with a 1)
- B. Between the firewalls (the area marked with a 2)
- C. On the private side of both firewalls (the area marked with a 3)
- D. With the database server (the area marked with a 4)

Answer: The correct answer is B. This area is a perimeter network or demilitarized zone (DMZ). It provides a layer of protection against external attacks, and helps prevent a compromised server from accessing internal resources. A is incorrect. Placing the web server directly on the Internet

leaves it vulnerable to more attacks. C and D are incorrect. Placing a public-facing server on an internal network increases the risk of attacks for other resources on the internal network.

Predicting the Future Exam Type

(ISC)² has implemented adaptive testing for the Certified Information Systems Security Professional (CISSP) exam. An adaptive exam uses algorithms to determine the candidate's knowledge level and gives different questions based on how the candidate answered previous questions. Instead of answering 250 questions within six hours, a CISSP exam candidate will see between 100 and 150 questions in up to three hours.

If (ISC)² follows the same path with the SSCP exam, you might see fewer questions and have less time to take it. However, implementing adaptive testing is more expensive so it's entirely possible (ISC)² won't implement adaptive testing for the SSCP exam.

Subscribing to the (ISC)² Code of Ethics and Getting Endorsed

After passing the exam, you need to subscribe to the (ISC)² Code of Ethics and go through an endorsement process. Both are described in the following sections.

Subscribing to the Code of Ethics

The Code of Ethics includes a preamble and four canons describing (ISC)²'s ethical expectations of its certified practitioners. Candidates must commit to and abide by the Code of Ethics to earn and keep the SSCP certification. Members who violate any provision of the Code of Ethics may have their certification revoked based on recommendations from a peer review panel.

The following sections quote the preamble and canons exactly as they appear on the (ISC)² page (<https://www.isc2.org/ethics>).

Code of Ethics Preamble The preamble consists of two points:

- The safety and welfare of society and the common good, duty to our principals, and to each other, requires that we adhere, and be seen to adhere, to the highest ethical standards of behavior.
- Therefore, strict adherence to this Code is a condition of certification.

Code of Ethics Canons The four canons are as follows:

- Protect society, the common good, necessary public trust and confidence, and the infrastructure.
- Act honorably, honestly, justly, responsibly, and legally.
- Provide diligent and competent service to principals.

- Advance and protect the profession.

If a situation arises resulting in a conflict between the canons, the conflict should be resolved in the order in which the canons are listed. In other words, the first canon is more important than the second one, and so forth.



EXAM TIP The (ISC)² Code of Ethics is included in the Security Operations and Administration domain and you can expect to be tested on it.

Getting Endorsed

As previously mentioned, when you register for the exam, you identify the number of years' experience that you have in any of the domains. After you take and pass the exam, you're required to submit a résumé that documents this experience. Additionally, you must find an (ISC)² certified professional (in good standing) to validate your experience and submit an endorsement form on your behalf. You need to complete this endorsement process within nine months of learning that you've passed the exam. (ISC)² completes periodic random audits to ensure the integrity of these documents.

Don't worry if you currently don't know an (ISC)² certified professional. During the course of your studies, you are very likely to meet one or more people who can endorse you. For example, if you join a local security group, you're sure to run into someone with an (ISC)² certification. (ISC)² chapters exist around the world, and you don't have to be certified to join. This page describes the chapters in more detail: <https://www.isc2.org/chapters/default.aspx>. Even if you don't meet an (ISC)² certified professional, you can still get an endorsement through (ISC)². You have to submit a special form to request (ISC)² endorsement (see <https://isc2.org/endorsement.aspx>), and the process takes much longer to get an endorsement than it would take using the normal process.

Maintaining Your SSCP Certification

After you've earned the SSCP certification, you're required to recertify every three years. The primary method of doing this is by acquiring 60 continuing professional education (CPE) credits every three years, with a minimum of 10 CPEs earned each year. Security constantly changes, and earning CPEs is one of the ways security professionals keep abreast of current security trends.

The CPE requirement is a surprise to some people, but many professions use the same concept. As an example, medical doctors are required to complete a minimum number of continuing medical education (CME) credits to maintain their medical licenses.

(ISC)² categorizes CPE credits as Group A credits and Group B credits. Group A credits are for activities directly related to one of the domains in the CBK. Group B credits are optional and are earned for activities that are outside of the domain, but can still enhance a member's general professional skills and competencies.

You typically earn one CPE credit for each hour you spend in a related activity. Group A credits can be earned by attending educational/training courses and seminars, attending conferences, attending vendor presentations, completing some academic courses, and preparing for a presentation, lecture, or training event. Some examples of Group B credits include attending a management course and participating in project planning activities. This is not an exhaustive list, but instead just a few examples.

To maintain your SSCP certification, you must earn the following CPE credits:

- At least 10 Group A CPEs annually
- At least 60 CPE credits during a three-year certification period
- At least 40 Group A CPE credits during a three-year certification period
- As many as 20 Group B CPE credits during a three-year certification period



TIP You have a lot of flexibility with Group B credits. Just about anything that enhances your career can be applied as a Group B credit. The only restriction is that you can't use more than 20.

Understanding Basic Security Concepts

Three primary goals of information security are to prevent the loss of confidentiality, the loss of integrity, and the loss of availability for information technology (IT) systems and data. You'll find that most security practices and security controls are designed to help prevent losses in one or more of these areas. The SSCP objectives identify these as the CIA triad, using the initials for confidentiality, integrity, and availability. [Figure 1-3](#) depicts the CIA security triad.



Figure 1-3 CIA security triad protecting information security



NOTE The CIA security triad is sometimes called the AIC security triad. Both are correct because the order of the initials doesn't matter. What you really need to know is what each letter represents (confidentiality, integrity, and availability) and what it means to prevent losses in these areas.

System and information owners are responsible for ensuring that security controls are in place to protect their systems and data. For example, managers that own proprietary data need to ensure that security controls are in place to prevent the unauthorized disclosure of that data. IT security personnel implement and maintain these security controls.



EXAM TIP The three primary goals of an information security program are to prevent the loss of confidentiality, the loss of integrity, and the loss of availability for any IT systems and data. System and information owners are responsible for ensuring that security controls address the confidentiality, integrity, and availability of their systems and data.

Confidentiality

You protect against the loss of confidentiality by ensuring that unauthorized users cannot access data. This starts by identifying and authenticating users, and then implementing access controls to restrict access. For example, you can use permissions to ensure that only authorized users can access the data.



NOTE [Chapter 2](#) covers authentication and access controls in more depth, and [Chapter 14](#) covers encryption algorithms.

Encryption provides another layer of protection for confidentiality. [Figure 1-4](#) shows the overall process of encryption, where data starts as plaintext, is ciphered using an encryption algorithm, and then becomes ciphered text. For example, if you've ever ordered anything over the Internet using a credit card, you've probably used a Hypertext Transfer Protocol Secure (HTTPS) connection. HTTPS encrypted your credit card information to prevent unauthorized individuals from intercepting it and using it without your permission. If criminals had managed to capture this transmission, they wouldn't have been able to read and use your credit card information.

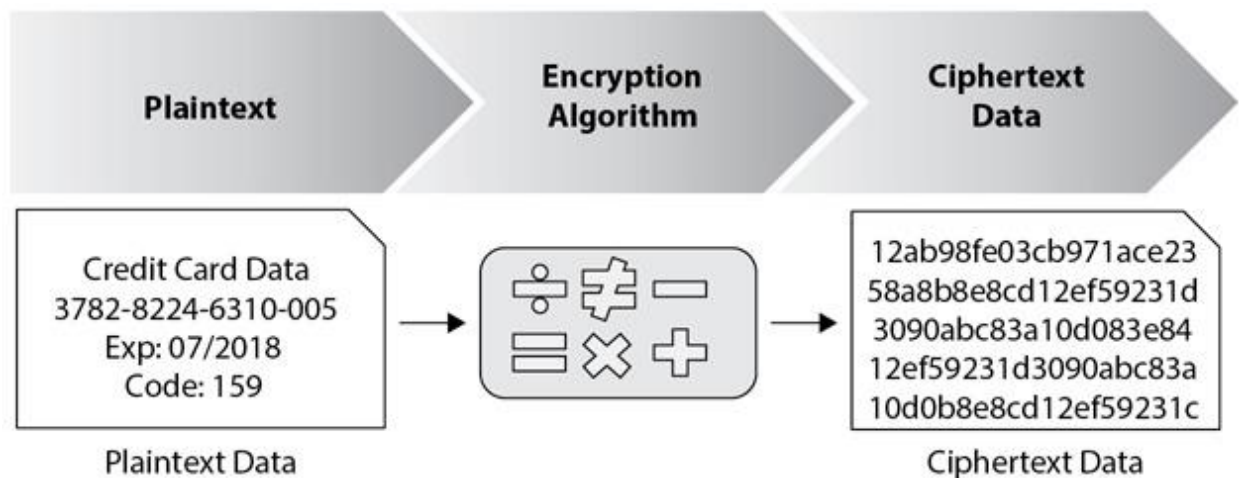


Figure 1-4 Encryption used to provide confidentiality

Confidentiality only works when organizations implement secure encryption algorithms and practice sound security practices. With this in mind, it's important to know which algorithms are secure and which algorithms have been compromised and shouldn't be used anymore.

As an example, the Data Encryption Standard (DES) is an older encryption algorithm that was used to encrypt data. It is now considered compromised and it shouldn't be used. The common replacement used today is the Advanced Encryption Standard (AES).



EXAM TIP Confidentiality controls help prevent the unauthorized disclosure of data. You ensure confidentiality by authenticating users and implementing access controls to ensure that

only authorized users can access the data. You can also encrypt data to ensure that even if the data falls into the wrong hands, it's less likely that unauthorized users can read it.

Integrity

Integrity controls prevent any unauthorized or unwanted modification of data or systems. Several different methods are used to protect integrity, including hashing and audit logging. Hashing methods will detect the alteration of data and alert personnel that it has lost its integrity.



NOTE [Chapter 10](#) covers auditing and logs in more depth, and [Chapter 14](#) covers hashing, including specific hashing algorithms.

A *hash* is simply a number created by performing a mathematical operation against data, such as a file or message. As long as the data stays the same, the hash (the number) will always be the same.

As an example, imagine that Sally needs to inform Joe of the price of a particular company's stock. She creates an e-mail message stating "The price is \$99" and creates a hash of the message. To keep the example simple, the hash for this message is 1234, though an actual hash would be much longer. Sally could re-create the hash on this message 100 times and each time it would be 1234. Here's how she can use the hash to validate the integrity of the message after she sends it:

1. Sally sends her message to Joe along with the calculated hash, as shown in [Figure 1-5](#).

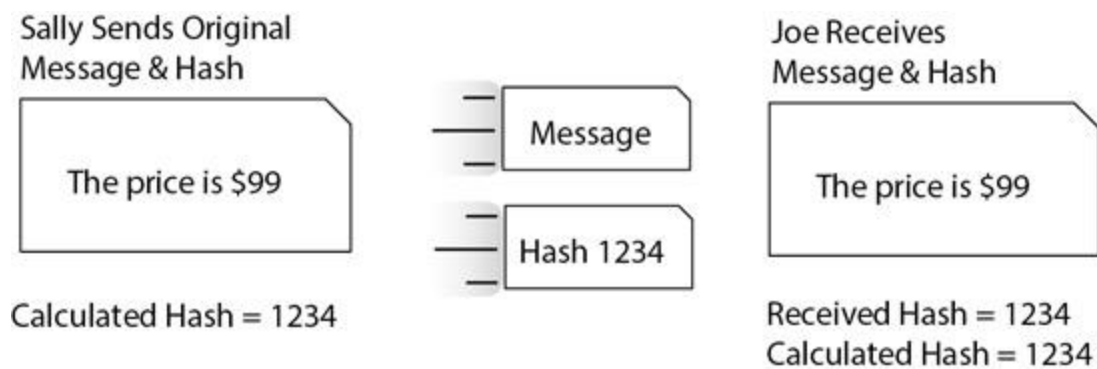


Figure 1-5 Hashing used to check integrity

2. When Joe's system receives the message, it calculates the hash again.
3. The received hash is the same as the calculated hash, providing assurances that the received message is the same as the sent message.

In contrast, imagine that someone intercepted the message and modified it before resending it to Joe, as shown in [Figure 1-6](#). Joe’s system calculates the hash on the received message, “The price is .99,” and determines it is 9876. It compares the calculated hash (9876) against the received hash (1234) and discovers that the message is different. In other words, the message has lost integrity.

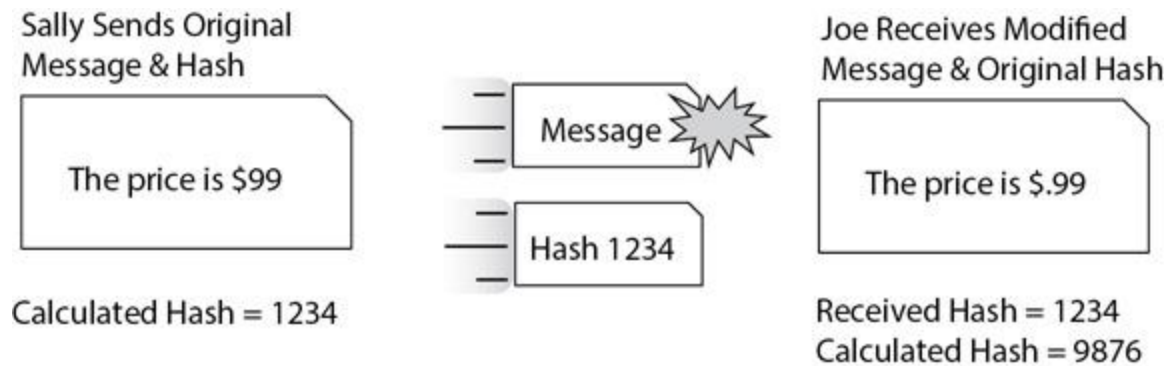


Figure 1-6 Hashing identifies altered data.



TIP If attackers can change the message, they can also change the hash. Some cryptography systems protect the hash by encrypting it, which prevents attackers from modifying the hash.

Using hashing alone, you can’t determine what altered the message. You only know that it was changed. However, this is valuable information. If the message is different, it shouldn’t be trusted.

While it’s important to be able to verify the integrity of data, integrity also applies to system configuration. As an example, organizations often use change and configuration management processes to prevent unexpected system outages resulting from changes to the system. If a technician makes an unauthorized change, it results in a loss of integrity for the system. Worse, many unauthorized changes have caused unexpected system outages.

You can also use audit logging for system integrity. An audit log tracks changes to a resource, including what was changed, who changed it, and when. A set of one or more audit logs creates an audit trail that you can use to verify whether the configuration of a system is the same or has been modified. If someone did make a change that caused an outage, investigators can use audit logs to identify what was changed and who made the change.



EXAM TIP Integrity ensures that data or systems have not been altered. Two common methods used to ensure integrity are hashing and audit logs.

Availability

Preventing the loss of availability ensures that IT systems and data are available when needed. Note that there isn't a timeframe here. Some organizations operate only during the daytime from Monday to Friday, so this is the only time when the systems are needed. Other organizations are operational 24/7, so the systems and data must also be available 24/7.

If users need to access data on a server and they can access it, then the data is available. However, if the data becomes corrupt or the server fails, the result is a loss of availability.

Organizations protect against loss of availability using a variety of different technologies. These include the following:

- **Backups** Regular backups capture a copy of the data. If something happens to the original data, administrators can restore the data from backups. It's important to keep a copy of backup data in an offsite location, so that if a fire or other catastrophe destroys the entire building, the data is still available.
- **Redundant disks** Many Redundant Array of Independent Disks (RAID) systems will continue to operate even if a disk fails. A mirror (RAID 1) is one example of redundant disks. The RAID system stores identical data on two disks, and if one disk drive fails, a copy of the data is still available.
- **Redundant servers** If a service provided by a server is critically important to an organization, the organization can add redundant servers. For example, failover clustering uses multiple servers and ensures that a service will remain available, even if a server fails.



NOTE [Chapter 9](#) covers backups, redundancy, and fault-tolerant techniques in greater depth. [Chapter 10](#) discusses alternate locations used as redundant sites.

- **Redundant connections** Organizations often need to stay connected to the Internet or stay connected between buildings in separate locations via an intranet. When this connectivity is critical to the operation of the organization, two or more connections are used so that even if one fails, the organization still has connectivity.
- **Redundant sites** Many organizations must stay operational even if a catastrophic event destroys their building or makes it uninhabitable. For example, many locations are susceptible to earthquakes, tornadoes, floods, and hurricanes. An organization can plan for these catastrophes by establishing a separate location. Redundant sites are known as *hot sites* (ready at a moment's

notice), *cold sites* (an empty building with electricity and running water), and *warm sites* (a cross between a hot site and a cold site).

In addition to using fault-tolerant and redundant technologies, organizations create business continuity plans and disaster recovery plans. These help the organizations maintain the availability of critical systems even after a disaster.



EXAM TIP Availability ensures that authorized users can access any resource when it's needed. Fault-tolerant and redundant technologies ensure that availability is not lost even if a system suffers a failure.

Exploring Fundamentals of Security

In addition to knowing the main goals of security (the CIA security triad), you need to understand some basic terms and concepts for the SSCP exam. The following sections introduce these concepts and some of the common terminology.

Least Privilege

An important security principle is the *principle of least privilege*. In short, this means that you grant users access to what they need to perform their jobs, and no more. This includes granting permissions to access resources such as files and granting rights to perform actions such as modifying system configurations.

For example, consider a group of project managers and project team members who all need access to a folder named Project Data. The team members need to be able to read the data but not modify it, while the managers need full control over the folder. [Table 1-1](#) shows a simple matrix identifying appropriate permissions. Notice that regular users are not granted any permissions because they don't need access to any of this data to perform their jobs.

| User Group | Project Data Folder Permissions |
|----------------------|---------------------------------|
| Project Managers | Full control |
| Project Team Members | Read access (cannot modify) |
| Regular Users | None |

Table 1-1 Permission Matrix

If you instead gave the project team members full control permission on the folder, they would still be able to read the files, but they could also modify them or even delete them. Consider what

could happen if the team members had full control and one of the team members became a disgruntled employee. That user could deliberately modify or delete data on the server. Even the most loyal employees can accidentally modify or delete files, but not if they don't have permissions to do so. By assigning only the required permissions, you reduce the risk of anyone modifying or deleting the data.

Similarly, network administrators need elevated privileges to modify network and server configuration settings. Regular users don't need these privileges to perform their jobs, so administrators ensure that regular users don't have them. This reduces the possibility of a regular user accidentally making a change that affects the availability of a system.



EXAM TIP The principle of least privilege ensures that users are granted only the rights and permissions needed to perform their jobs, and no more.

Separation of Duties

Separation of duties is a security principle that ensures that no single person has complete control over a process. When properly implemented, separation of duties significantly reduces the risk of fraud within an organization.

Consider the process of approving and paying invoices. If Joe controlled the entire process, he could create an invoice for his own fictitious company, approve the invoice, and then make a payment to his own bank account. Of course, the loser in this scenario is the company that is employing Joe. Separating the payment process into two steps and assigning different people to handle each step reduces the risk of fraud. [Figure 1-7](#) shows how one person approves the invoice, while another person pays the invoice. Because neither person has full control of the process, neither person can defraud the company without involving the other person.

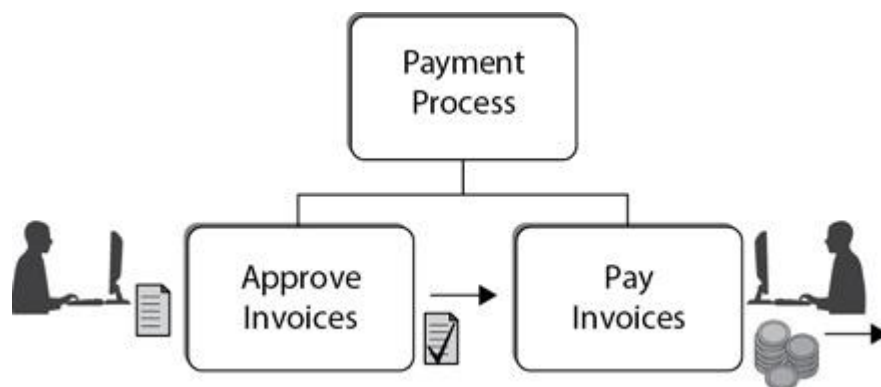


Figure 1-7 Separation of duties

Implementing separation of duties policies doesn't eliminate the possibility of fraud, because the two employees could choose to collude to defraud the company. Therefore, many companies also use job rotation and mandatory vacations to reduce the risks of collusion. [Chapter 13](#) covers both of these concepts in more depth.



EXAM TIP Separation of duties helps prevent fraud by ensuring that no single person has complete control over a process.

Privacy and Data Sensitivity

Data is often an extremely valuable asset for an organization. Three types of data that are sensitive and require extra steps to protect are PII, PHI, and intellectual property.

Personally identifiable information (PII) is information that identifies an individual and includes items such as the person's name, national identification number such as the U.S. Social Security number, and birthdate. *Protected health information (PHI)* is any information about an individual's medical and health history.

National Institute of Standards and Technology (NIST) Special Publication (SP) 800-122, "Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)," includes details on identifying and protecting PII. NIST SP 800-66, Revision 1, "An Introductory Resource Guide for Implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule," discusses various methods of protecting PHI.

Intellectual property refers to creations of the mind or intellect, such as literary and artistic works, and inventions. The World Intellectual Property Organization (WIPO) divides intellectual property into two categories: industrial property and copyrights. Industrial property applies to industrial designs, trademarks, and patents for inventions. Copyrights apply to literary works (such as books and screenplays), music, films, and other artistic works.

Organizations often apply confidentiality principles to protect sensitive data such as PII, PHI, and intellectual property. This includes using strong access controls to restrict access to the data. It also includes encrypting privacy data as an added layer of protection.



TIP Several laws mandate the protection of an individual's personally identifiable information (PII) and protected health information (PHI). Organizations have a requirement to exercise due care to protect PII, PHI, and intellectual property.

Defense in Depth

One of the primary tenets of security is that you're never done. You can't just write a security policy, install antivirus software or enable firewalls, and say, "There. We're safe and secure now." Instead, IT security uses the principle of *defense in depth* to implement several layers of security.

Consider [Figure 1-8](#). It shows network resources protected through several layers of security. [Chapter 9](#) covers security controls in greater depth, but in short, a security control attempts to reduce risk by either reducing vulnerabilities or minimizing the impact of a threat. One of the primary benefits of a defense in depth strategy is that even if a single control fails, other controls still provide protection.

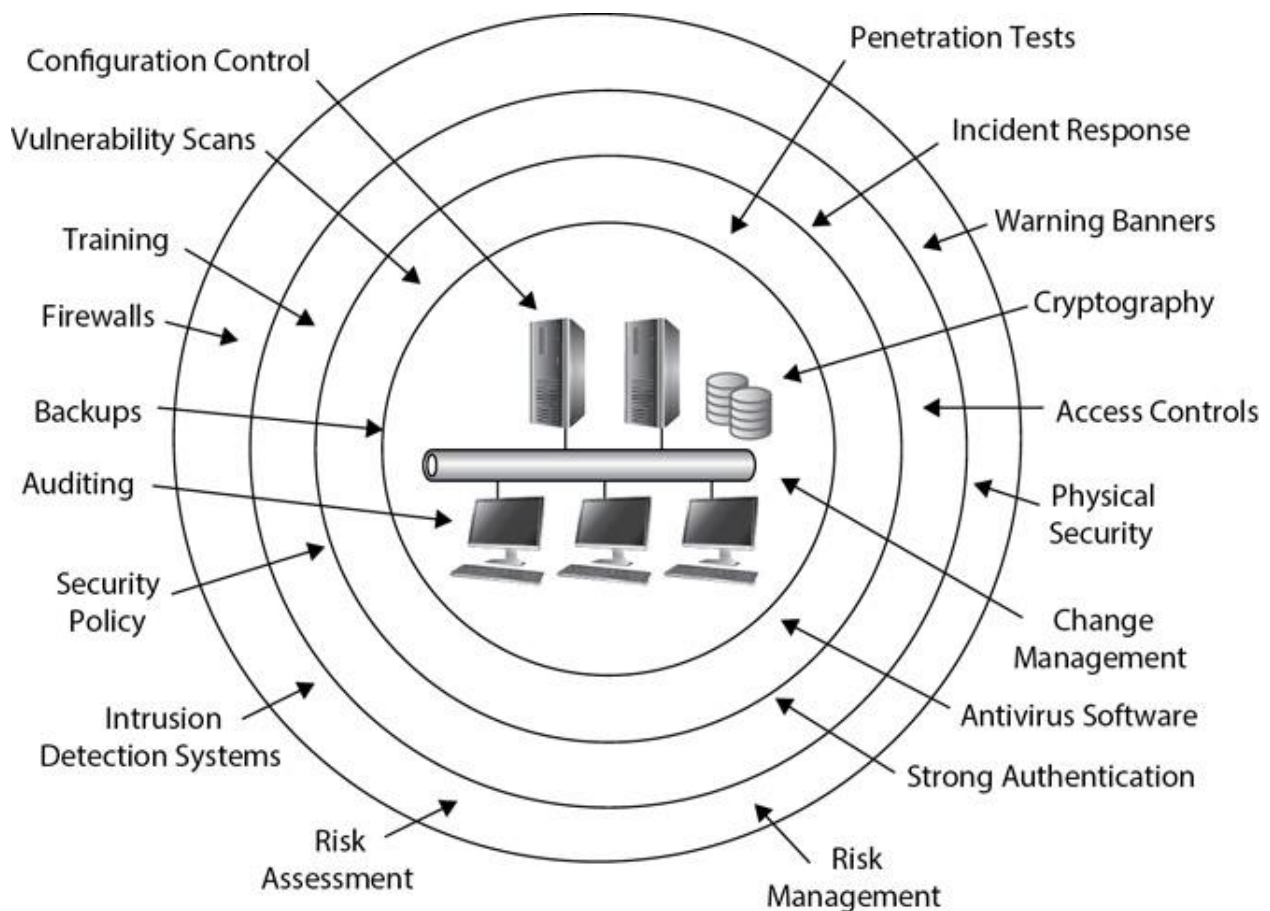


Figure 1-8 Defense in depth includes several layers of security.



NOTE The arrows in [Figure 1-8](#) are not meant to imply at which layer any of the controls are implemented. Instead, the message is that IT and security personnel apply multiple methods of security at multiple layers.

For example, you may combine access controls with the principle of least privilege to restrict access to data within your organization. You may also have some research and development data that you want to ensure remains confidential. In addition to access controls and least privilege, you can use cryptography methods to add an extra layer of security for this research and development data. Even if someone is able to bypass the access controls, he or she will not be able to decrypt the data easily.



EXAM TIP A defense in depth strategy provides a layered approach to security by implementing multiple controls at different layers.

Nonrepudiation

Nonrepudiation ensures that a party cannot believably deny (or repudiate) taking an action. Audit logging and digital signatures are two common methods used to enforce nonrepudiation.

Consider a system that has audit logging enabled for a specific folder. If any user reads, modifies, or deletes data in the folder, the system logs the event in an audit log. The log includes who performed the activity, when they did it, and what they did. If Joe logs on to a computer using his credentials and he deletes a file, the audit log holds a record of his actions. Because the log recorded information from Joe's credentials, you know that Joe did it.



NOTE A remote possibility is that someone else is using Joe's credentials. This possibility increases if the organization uses weak authentication or has poor security practices. However, if Joe logs on with strong authentication (such as with a smart card or biometrics), it's highly unlikely someone is impersonating him.

Digital signatures also provide nonrepudiation. For example, if Sally sends an e-mail to Bob and signs it with a digital signature, Sally can't later deny that she sent the e-mail. Digital signatures use certificates and public/private key encryption. They also provide authentication, giving assurances of who sent the e-mail.

Another example of nonrepudiation is related to commerce and e-commerce transactions. If you use a credit card to purchase a product and you sign the credit card bill, the company can use your signature to prove you are the person who made the purchase. You couldn't later deny it,

because your signature verifies that you purchased it. Similarly, e-commerce transactions require you to enter additional information such as the expiration date and the security code on the card. The idea is that only someone with the card in his or her possession knows this additional information.



EXAM TIP Nonrepudiation prevents a party from denying that he or she took an action. The sender of a digitally signed e-mail cannot believably deny sending it. If a system has accountability and an audit trail shows the user took an action, the user cannot believably repudiate it.

AAAs of Security

The AAAs of security are *authentication, authorization, and accounting*. Combined, they help to ensure that only authorized entities have access to resources and that their access is recorded. [Figure 1-9](#) shows the AAAs of security, and the following list explains them. While reading the explanations, imagine that Dawn is a project manager and she needs to modify project files stored on a server.

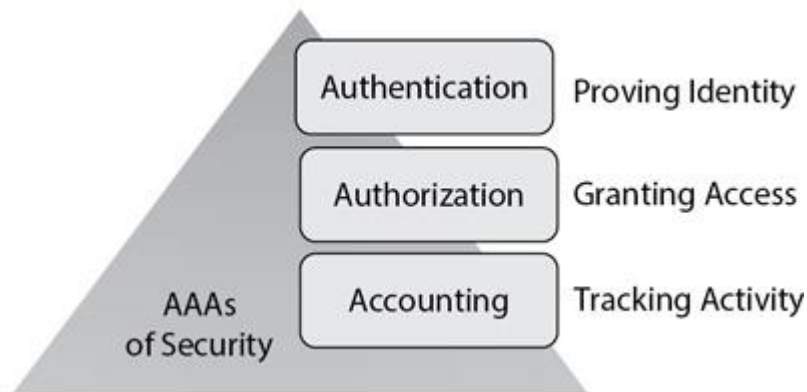


Figure 1-9 AAAs of security

- **Authentication** A user provides credentials (such as a username and password) that are checked against a database to prove the user's identity. The authentication system verifies the credentials. In the scenario, Dawn's username is her identity, and she provides the correct password to prove her identity. If she enters the correct username and password, the system authenticates her.



EXAM TIP There are three types or factors of authentication, known as (1) something you know (such as a username and password), (2) something you have (such as a smart card), and (3) something you are (using biometrics). [Chapter 2](#) covers authentication in greater depth.

- **Authorization** Administrators assign rights and permissions to resources, which authorize users to access the resources. In the scenario, when Dawn attempts to access files on the server, the system checks to see if Dawn's user account has been granted appropriate permissions and is authorized to access the files. If her account has appropriate permissions, the system grants her access. If her account doesn't have appropriate permissions, the system blocks her access.
- **Accounting** Logging tracks activity of a user through monitoring. A basic accounting mechanism is an audit log, such as the Security log in Windows systems, and audit logs create an audit trail. Imagine that administrators have configured the server storing the project files to log events anytime someone accesses the files. When Dawn opens a file, the system logs an event in the Security log. Similarly, if she saves changes that modify a file, or deletes a file, the system records these events with log entries.

Note that you cannot restrict authorization without proper identification and authentication. If all users have the same account, you can either authorize access to everyone or block access to everyone. Similarly, without authentication, you can't have accounting. If everyone used the same account, event logs could record when users accessed a resource, but they wouldn't be able to record who accessed the resource.



EXAM TIP The AAAs of security are authentication, authorization, and accounting.

Accountability

One of the underlying goals of the AAAs of security is *accountability*. If a system can identify individual users, track their actions, and monitor their behavior, it provides accountability. Authentication provides identification for users, and accounting tracks their activities in audit logs. If users are not required to authenticate or if audit trails are not created, then a system does not provide accountability.

It's worthwhile noting that you don't have to track every single action of a user to provide accountability. For example, your network may have proprietary data stored in a folder named Research and publicly available information stored in a folder named Public. You may want to track each time any single user accesses any single file within the Research folder. This includes any time someone views, modifies, or deletes a file. However, it's not important to know who viewed any of the data in the Public folder; instead, you only want to know who modified it. In this situation, you would track all activity in the Research folder, but only some of the activity in the Public folder.



EXAM TIP If a system can track activity of an individual on a system, it provides accountability.

Due Diligence

Due diligence refers to the investigative steps that an organization takes prior to taking on something new, such as signing a contract or making a major purchase. In the IT world, an organization has an obligation to exercise due diligence to discover risks associated with a large purchase.

For example, if an organization is planning to purchase a software development company, that organization is obligated to exercise due diligence to determine as much as it can about the company and whether the purchase is a sound decision. Suppose the software company claims it earns about \$100 million in revenue a year, but in actuality only earns \$10,000. If the organization ignores due diligence and simply believes the company's statements without verifying them, they may purchase the company based on the \$100 million claim and pay way too much for it. By exercising due diligence and investigating the claims, the organization will likely discover the discrepancies and avoid the losses.

From an IT security perspective, an organization has a responsibility to exercise due diligence to discover risks. The goal is to identify risks that can result in the loss of availability, integrity, or confidentiality of any data or IT systems. Risks can't be eliminated, but an organization can take steps to mitigate or reduce the risks by exercising due care.



EXAM TIP You cannot eliminate risk. Management decides which risks to mitigate, and the risk that remains is residual risk. Management is responsible for any losses that occur because of residual risk.

Due Care

Due care is the practice of implementing security policies and practices to protect resources. It ensures that a certain level of protection is applied to protect against losses from known risks. The goal is to reduce the risk to the resources to a manageable level.

Because risks can't be eliminated, an organization is likely to experience losses. If these losses are due to negligence, then the organization may face legal action against it. However, if the organization took due care to protect the resources but still suffered the loss, it's less likely that a court will find the organization negligent.

For example, imagine that a company holds customer data—including names, addresses, birth dates, and credit card data—in cleartext in a database hosted on a web server. The company uses this information when customers make purchases through a website. A hacker checks out the website, discovers the database, and realizes that he can easily retrieve all the data. He steals the data and sells it to identity thieves, who proceed to steal millions of dollars.

Did the company take due care? Many would say no.

A web server accessible by users on the Internet will be attacked. It's not a matter of *if* it will happen, but *when*. Further, if valuable data is on the server, attackers will find it. Even if it does have some protection, such as encryption, it's still at risk if it's accessible from the Internet.

Taking due care, the organization would implement security controls to protect the data. For example, administrators could store the customer information on a different server within the private network that isn't accessible from the Internet but is accessible from the web server. Additionally, they could encrypt sensitive data on the server to protect against the loss of confidentiality. They don't need to encrypt the entire database, but instead just encrypt columns holding sensitive data such as credit card data.



EXAM TIP If an organization doesn't take due care to protect private data, it might be sued or forced to pay fines for privacy violations.

Chapter Review

There are several requirements to earn the (ISC)² System Security Certified Practitioner (SSCP) certification. Candidates must have at least one year of experience in one of the seven domains defined by the (ISC)² for the SSCP exam. During the registration process, candidates must agree to comply with the (ISC)² Code of Ethics.

The SSCP exam is computer based and is administered at Pearson VUE test centers. The exam includes 125 questions, but up to 25 of these are ungraded beta questions. You need to score a minimum of 700 on a scale of 1,000 to pass the exam.

The CIA security triad includes confidentiality, integrity, and availability. System and information owners are responsible for ensuring that security and access controls are in place to prevent losses in these areas.

- Confidentiality protects against the unauthorized disclosure of information. You can ensure confidentiality with access controls and encryption.

- Integrity ensures that unauthorized users have not altered data. Hashing provides assurances that data, such as messages or files, has not been altered. Audit logs provide an audit trail identifying whether data has been modified, who modified it, and when.
- Preventing loss of availability ensures that systems and data are operational and available when needed. Organizations prevent the loss of availability with backups, fault tolerance, and redundant systems.

This chapter also presented several fundamental security terms:

- The principle of least privilege ensures that users are granted only the rights and permissions they need to perform their jobs, and no more.
- Separation of duties helps prevent fraud by ensuring that no single individual has complete control over a process. Instead, tasks are divided between multiple people.
- Organizations are responsible for protecting the confidentiality of sensitive data. Many laws mandate the protection of both PII and PHI. Organizations also have a responsibility for protecting intellectual property.
- Defense in depth provides a layered approach to security and protects an organization even if one or more security elements fail.
- Nonrepudiation prevents an individual from denying that he or she took an action. For example, if a user sends a digitally signed message, the user cannot later deny sending it, because a digital signature provides nonrepudiation.
- The AAAs of security are authentication, authorization, and accounting. Authentication identifies users, authorization defines what the users can access, and accounting tracks a user's activities.
- If a system has accountability, it can identify and track the activity of a user. Accountability starts with strong authentication practices.
- Due diligence refers to the investigative steps that an organization takes prior to making a major purchase or taking on an obligation.
- Due care refers to the steps that an organization takes to protect resources. If an organization doesn't take due care to protect resources, such as any individual's private data, it can be susceptible to legal action or fines.

Questions

1. How many years of experience are required to earn the Associate of (ISC)² designation?

A. Zero

B. One

C. Two

D. Five

2. What are the three elements of the security triad?

A. Authentication, authorization, and accounting

B. Confidentiality, integrity, and availability

C. Identification, authentication, and authorization

D. Confidentiality, integrity, and authorization

3. Who is responsible for ensuring that security controls are in place to protect against the loss of confidentiality, integrity, or availability of their systems and data?

A. IT administrators

B. System and information owners

C. CFO

D. Everyone

4. You are sending an e-mail to a business partner that includes proprietary data. You want to ensure that the partner can access the data but that no one else can. What security principle should you apply?

A. Authentication

B. Availability

C. Confidentiality

D. Integrity

5. Your organization wants to ensure that attackers are unable to modify data within a database. What security principle is the organization trying to enforce?

A. Accountability

B. Availability

C. Confidentiality

D. Integrity

6. An organization wants to ensure that authorized employees are able to access resources during normal business hours. What security principle is the organization trying to enforce?

A. Accountability

B. Availability

C. Integrity

D. Confidentiality

7. An organization has created a disaster recovery plan. What security principle is the organization trying to enforce?

A. Authentication

B. Availability

C. Integrity

D. Confidentiality

8. Your organization has implemented a least privilege policy. Which of the following choices describes the most likely result of this policy?

A. It adds multiple layers of security.

B. No single user has full control over any process.

C. Users can only access data they need to perform their jobs.

D. It prevents users from denying they took an action.

9. Your organization wants to implement policies that will deter fraud by dividing job responsibilities. Which of the following policies should they implement?

A. Nonrepudiation

B. Least privilege

C. Defense in depth

D. Separation of duties

10. Which one of the following concepts provides the strongest security?

A. Defense in depth

B. Nonrepudiation

C. Security triad

D. AAAs of security

11. Which of the following would a financial institution use to validate an e-commerce transaction?

A. Nonrepudiation

B. Least privilege

C. Authentication

D. Signature

12. What are the AAAs of information security?

A. Authentication, availability, and authorization

B. Accounting, authentication, and availability

C. Authentication, authorization, and accounting

D. Availability, accountability, and authorization

13. You want to ensure that a system can identify individual users, track their activity, and log their actions. What does this provide?

A. Accountability

B. Availability

C. Authentication

D. Authorization

14. Which of the following is required to support accountability?

- A. Encryption
- B. Authentication
- C. Hashing
- D. Redundant systems

15. Which of the following statements accurately describes due care?

- A. It is the practice of implementing security policies and procedures to protect resources.
- B. Due care eliminates risk.
- C. A company is not responsible for exercising due care over PII.
- D. Organizations cannot be sued if they fail to exercise due care over resources such as customer data.

Answers

1. A. You don't need to meet the experience requirement to earn the Associate of (ISC)² designation, so zero years of experience are acceptable. The SSCP certification requires one year of direct full-time security work experience. If you earn the Associate of (ISC)² designation, you have two years from the date (ISC)² notifies you that you have passed the SSCP exam to obtain the required experience and apply to become a fully certified SSCP (which includes submitting the required endorsement form). The CISSP certification requires five years of experience.

2. B. The CIA security triad includes three fundamental principles of security designed to prevent losses in confidentiality, integrity, and availability. Authentication, authorization, and accounting are the AAAs of security, and identification, authentication, and authorization are required for accountability, but these are not part of the CIA security triad.

3. B. System and information owners are responsible for ensuring that these security controls are in place. IT administrators or other IT security personnel might implement and maintain them. While it can be argued that the Chief Executive Officer (CEO) is ultimately responsible for all security, the Chief Financial Officer is responsible for finances, not IT security. Assigning responsibility to everyone results in no one taking responsibility.

4. C. Confidentiality helps prevent the unauthorized disclosure of data to unauthorized personnel, and you can enforce it with encryption in this scenario. Authentication allows a user to claim an identity (such as with a username) and prove the identity (such as with a password). Availability ensures that data is available when needed. Integrity ensures that the data hasn't been modified.

5. D. Integrity ensures that data is not modified, and this includes data within a database. Accountability ensures that systems identify users, track their actions, and monitor their behavior. Availability ensures that IT systems and data are available when needed. Confidentiality protects against the unauthorized disclosure of data.

6. B. Availability ensures that IT systems and data are available when needed, such as during normal business hours. Accountability ensures that users are accurately identified and authenticated, and their actions are tracked with logs. Integrity ensures that data is not modified. Confidentiality protects the unauthorized disclosure of data to unauthorized users.

7. B. Availability ensures that IT systems and data are available when needed. Disaster recovery plans help an organization ensure availability of critical systems after a disaster. Users prove their identity with authentication. Integrity provides assurances that data and systems have not been modified. Confidentiality protects against the unauthorized disclosure of data.

8. C. The principle of least privilege ensures that users have access to the data they need to perform their jobs, but no more. Defense in depth ensures an organization has multiple layers of security. Separation of duties ensures that no single user has full control over any process. Nonrepudiation prevents users from denying they took an action.

9. D. Separation of duties helps prevent fraud by dividing job responsibilities and ensuring that no single person has complete control over an entire process. Nonrepudiation ensures that parties are not able to deny taking an action. The principle of least privilege ensures that users have only the rights and permissions they need to perform their jobs, but no more. Defense in depth provides a layered approach to security.

10. A. Defense in depth provides a layered approach to security by implementing several different security practices simultaneously and is the best choice of the available answers to provide the strongest security. The security triad (confidentiality, integrity, and availability) identifies the main goals of security. Nonrepudiation prevents an individual from denying that he or she took an action. The AAAs of security are authentication, authorization, and accounting.

11. A. Digital signatures used by some online institutions to validate transactions and provide nonrepudiation. Least privilege ensures that users have only the rights and permissions they need to perform their jobs, and no more. Authentication verifies a user's identity. A written signature is not used in e-commerce.

12. C. The AAAs of information security are authentication, authorization, and accounting. Availability is part of the CIA security triad (confidentiality, integrity, and availability), but it is not part of the AAAs of information security.

13. A. If a system can identify individual users, track their activity, and log their actions, it provides accountability. Availability ensures the system is operational when needed. Authentication identifies the individual using credentials. Authorization identifies resources that a user can access.

14. B. Users prove their identity with authentication, and strong authentication mechanisms are required to support accountability. Encryption helps provide confidentiality. Hashing helps provide integrity. Redundant systems help provide availability.

15. A. Due care is the practice of implementing security policies and procedures to protect resources. You cannot eliminate risk. A company is responsible for exercising due care over PII and can be sued if it fails to exercise due care.

NOTES

CHAPTER 2

Access Controls

In this chapter, you will learn about

- Identification, authentication, and authorization
- Three factors of authentication
- Single sign-on
- One-time passwords
- Access controls with subjects and objects
- Access control models (DAC, non-DAC, MAC, RBAC, and ABAC)
- Bell-LaPadula, Biba, Clark-Wilson, and Brewer-Nash architectures
- Identity and access management (IAM) systems
- Physical security controls

Comparing Identification, Authentication, and Authorization

A key part of any access control system is the identification and authentication of individuals. If you can't identify individuals, everyone is anonymous. If everyone is anonymous, there is no way you can control access to different resources. Either everyone has access or no one does.

However, if you are able to differentiate between different users, you can grant access to some users while denying access to other users. This process starts with identification. In many authentication systems, the identity of a user is simply the user's name.

[Figure 2-1](#) shows the overall process of identification, authentication, and authorization. The user professes an identity with a username and validates the identity by providing authentication information. An authentication system verifies the credentials. Then access controls (such as permissions) authorize the user to access resources such as files stored on a server. If these three steps don't come together, the user is not able to access resources protected with access controls.

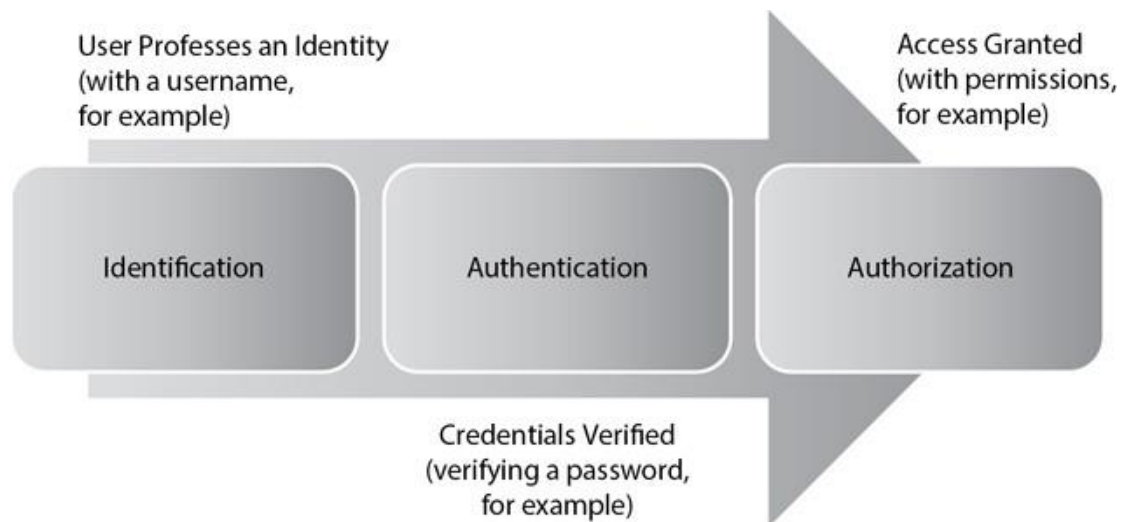


Figure 2-1 Identification, authentication, and authorization

Just because someone knows a username doesn't mean they are authenticated. They must also prove their identity by providing authentication data, such as a password or fingerprint. Similarly, just because a user can log on with a username and password doesn't mean they can access everything on a system or a network. Administrators authorize access by configuring rights and permissions to users based on their proven identities.



EXAM TIP Identification and authentication are two primary controls used by access control systems. The user professes an identity and then proves his or her identity with authentication. Administrators grant a user rights and permissions based on the user's proven identity.

Exploring Authentication

Authentication is the process of users proving they are who they claim to be. A common method is by having a user provide a logon name to profess an identity and then using a password for authentication.

For example, [Figure 2-2](#) shows a challenge screen where a user must provide account credentials. The username provides the identification, and the password provides the authentication.

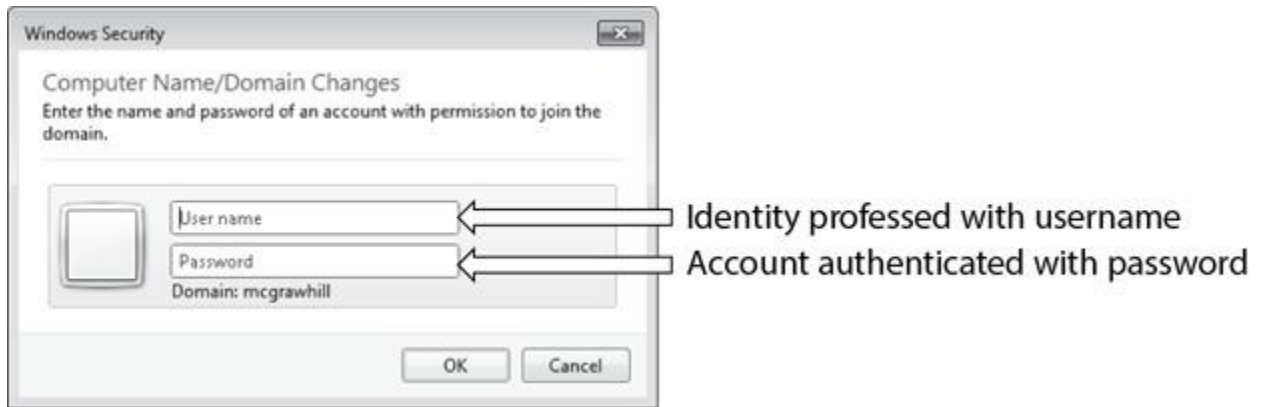


Figure 2-2 Username and password for identification and authentication

While you've probably used a password to authenticate to a system, you might not have used all possible authentication mechanisms. The following sections cover the different authentication methods.

Three Factors of Authentication

As shown in [Figure 2-3](#), most authentication falls into one of three primary types or factors:

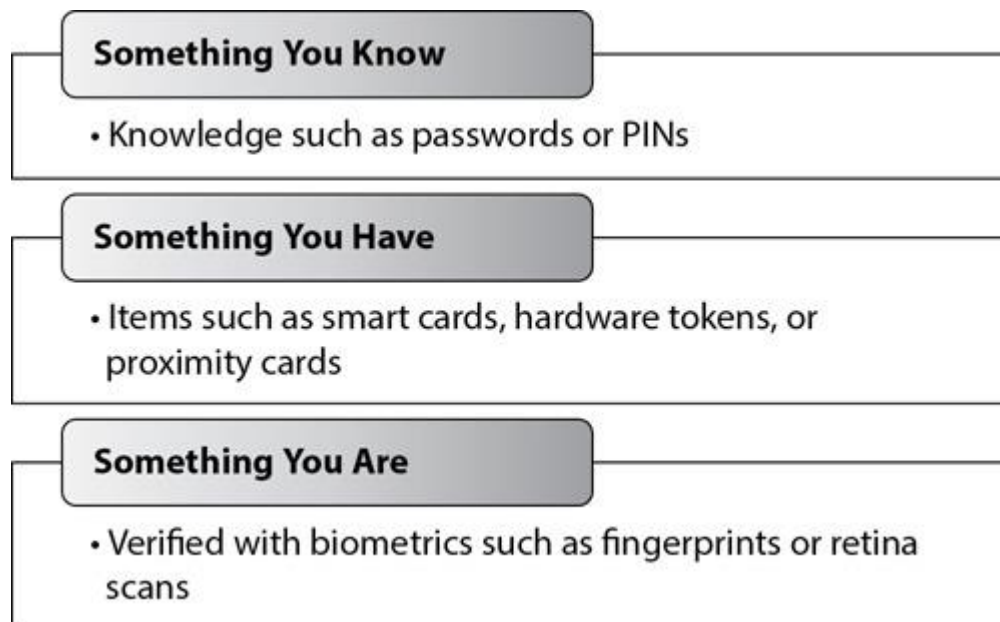


Figure 2-3 Three factors of authentication

- **Something you know (type 1)** This includes knowledge, such as passwords, personal identification numbers (PINs), your mother's maiden name, or even personal information such as the name of your first pet.

- **Something you have (type 2)** This includes items such as smart cards, hardware and software tokens, and proximity cards.
- **Something you are (type 3)** This includes the use of biometrics, such as fingerprints and retina scans.

Something You Know

The simplest type of authentication is *something you know* and includes passwords, PINs, and other information known by an individual. Passwords can typically include multiple character types, such as uppercase letters, lowercase letters, numbers, and special characters. A PIN is typically only numbers.

Passwords have the following classifications:

- **Static password** A static password stays the same over a period of time. Each time a user logs on to a session, they use the same password.
- **One-time or dynamic password** One-time passwords are used only once per session. Many hardware tokens generate new passwords every 60 seconds that users can enter as a one-time password. The “Something You Have” section, later in this chapter, explains hardware tokens in more depth.
- **Cognitive password** Information that a user would know, such as the name of a first pet, a first employer, and a favorite color, is combined to create a cognitive password. Users enter this information during a registration process. Later, the authentication system challenges the user with these questions in specific situations. For example, many authentication systems register information about the user’s computer, and if they detect the user is accessing the account from a different system, they challenge the user with cognitive password questions. If the user can give the same information when accessing the account at a different computer, the system authenticates the user.
- **Passphrase** A passphrase is a long string of characters that has meaning to the user. Users typically code these to make them more complex. For example, you can create a passphrase of *IWillPassTheSSCPEXam* from the statement *I will pass the SSCP exam*. If a system forces a user to use all four character types (uppercase letters, lowercase letters, numbers, and special characters), a user can change it to *IWi11P@\$\$The\$\$CPEXam*. Both are long passwords of 20 characters, but relatively easy to remember. In the second example, the letter *l* is replaced by the number 1 in each instance, the *a* is replaced by @, each *s* or *S* is replaced by \$, and the first letter in every word is capitalized. Even as complex as it is, most people could easily commit this 20-character passphrase to memory.

Password Policies Many organizations implement written password policies providing rules for users on passwords. The policy identifies the minimum requirements. Administrators enforce the password policy using a technical control. For example, Windows Group Policy includes Password Policy settings that administrators can configure to apply to all users in a domain.

[Figure 2-4](#) shows the Windows Local Group Policy Editor with the Password Policy selected. The following list describes the Password Policy settings:

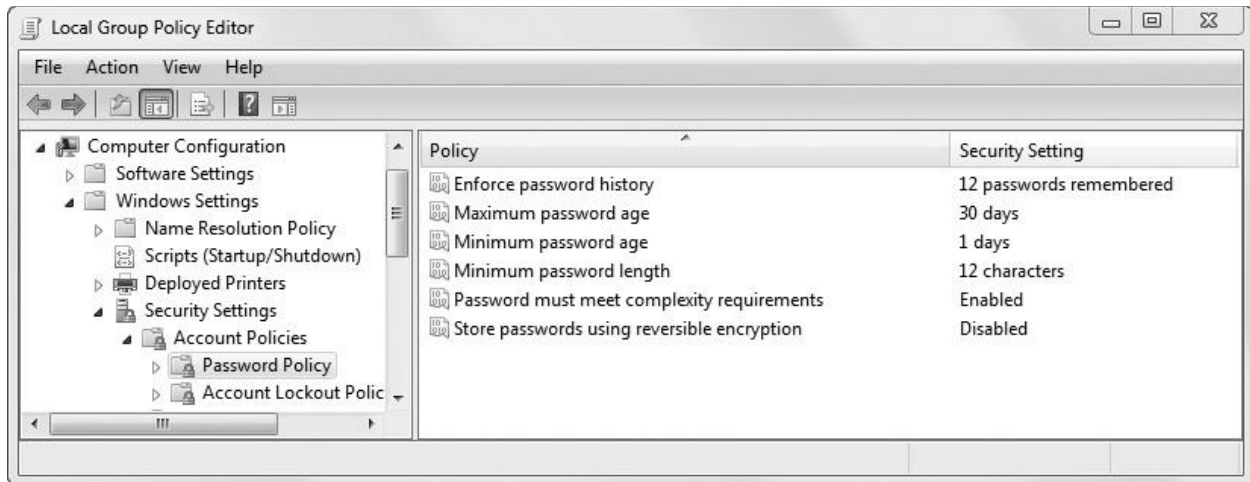


Figure 2-4 Password Policy within Windows Local Group Policy

- **Enforce password history** This setting remembers previous passwords that a user has created and prevents the user from reusing the same passwords. As configured, it prevents a user from reusing any of the last 12 passwords they've created.
- **Maximum password age** This setting defines when users must change their password. As configured, it requires users to reset their password every 30 days.
- **Minimum password age** This setting defines how long users must wait before changing their password again. As configured, it prevents users from changing their password until 1 day has passed. This setting works with the password history setting. If set to 0, users can change their password 12 times in a couple of minutes and get back to their original password.
- **Minimum password length** This setting enforces the character length of the password. As configured, it requires users to create passwords at least 12 characters long.
- **Password must meet complexity requirements** When enabled, this setting requires users to create complex passwords that include at least three of the four character types—uppercase letters, lowercase letters, numbers, and special characters.
- **Store passwords using reversible encryption** Some legacy systems require storing passwords with reversible encryption. However, this is not recommended because it makes it easier for attackers to discover the password. This setting should be disabled unless otherwise required.



EXAM TIP A written password policy often identifies requirements for strong passwords, the minimum amount of time before a password should be changed, and how passwords are audited. Technical controls enforce the policy. Many technical password policies can remember previous passwords and prevent users from reusing the same passwords again.

A popular method to create strong passwords is to use a password manager such as KeePassXC. It will create and store complex passwords within a database. As long as you use a strong master password to open the application, it keeps your passwords relatively secure. This is much better than storing all your passwords in a file named passwords.

Passwords are the least secure method of authentication; however, you can increase password security by following some basic guidelines:

- **Use strong passwords** A strong password includes a combination of different character types, such as uppercase letters, lowercase letters, numbers, and symbols. A strong password also needs to be sufficiently long. However, sufficiently long is a moving target. They should be at least 8 characters long, but many security professionals recommend passwords of 12 characters, 15 characters, or more. Strong passwords do not include usernames, logon names, children names, pet names, or any words found in a dictionary (of any language).



NOTE Some systems create random passwords for users. However, these are often much more difficult for users to remember, causing them to write them down.

- **Don't write down passwords** If a password is only in a person's head, it can't be read. On the other hand, if a user writes it down on a yellow sticky note and sticks it beneath a keyboard, an attacker can find it.
- **Change passwords often** Many security professionals have historically recommended changing passwords at least every 90 days. Some organizations require users to change their passwords more often, such as every 30, 45, or 60 days. However, newer recommendations from the National Institute of Standards and Technology (NIST) indicate this isn't useful and should be avoided, as discussed in the following section.
- **Don't use the same password on multiple systems** Unfortunately, attackers have been successful at hacking into systems and downloading huge databases that include user credentials. Attackers then try these credentials on other sites. If users have the same credentials on multiple sites, it becomes easier for attackers to compromise their accounts.
- **Never give out your password** Many social engineering attacks can easily be avoided if users understand and follow this simple rule. Because social engineers frequently try to get passwords simply by requesting them, security professionals often repeat this password policy rule to end users.



NOTE [Chapter 5](#) covers social engineering and many types of social engineering attacks.

- **Audit passwords** Many organizations use technology such as built-in operating system tools or third-party applications to verify that passwords are strong and that users change them regularly. A password audit ensures that passwords are strong (never blank), are of a minimum length, and have been changed within a given time period (such as 90 days). This is much more effective than a security professional trying to verify that all user passwords are strong and changed regularly.
- **Use a credential management system** Credential management systems provide a storage space for users to keep their credentials. For example, Windows systems include the Credential Manager tool, which allows users to store different username/password pairs and certificates used for authentication. Similarly, there are many password wallet or password vault apps that allow users to store multiple passwords within a single application. For example, KeePassXC is a free open source utility that allows users to store usernames and passwords within a secure encrypted database. Users open the database by entering a single master password or key file.



EXAM TIP Some password audit tools only enforce the use of strong passwords when users create the password. Other tools can also check passwords periodically to ensure that password policies are followed, such as ensuring that they have been changed within a given time period.

NIST Password Recommendations NIST Special Publication (SP) 800-63B, titled “Digital Identity Guidelines,” provides some newer recommendations related to passwords (also known as “memorized secrets” in NIST lingo) that are counter to what conventional wisdom has told us in the past. While SSCP test item writers may take a while to catch up with these recommendations, many organizations are starting to implement them.

Following conventional wisdom, a strong password is of sufficient length and complex, using multiple character types. However, when password policies force users to create complex passwords, users have responded in predictable ways. For example, imagine Joe typically uses a password of “password” for online sites. However, a password policy forces him to use all four character types. If Joe is a typical user, he will likely create a password of “Password1!”.

Instead of forcing users to create overly complex passwords that are difficult to remember, SP 800-63B recommends the following actions:

- **Users should be encouraged to create passphrases** As an example, a passphrase of ICanCountTo6 is a simple password to remember. However, because it has 12 characters and uses a variety of character types, it is difficult to crack.
- **Passwords should be a minimum of 8 characters** While longer passwords can be required, NIST recommends using multifactor authentication (described later in this chapter) for accounts that need more security.
- **Passwords should never be stored in cleartext** Passwords should be salted and hashed with a strong hashing algorithm such as Secure Hash Algorithm 3 (SHA-3). Salts are random bits added to a password, and NIST recommends a salt of at least 32 bits. A hashing algorithm reduces the salted password to a string of hexadecimal characters. For example, SHA-3(256) creates a hash 256 bits long.
- **Users should be able to create passwords with as many as 64 characters** Because passwords should never be stored, users should be able to create long passwords. For example, *ILoveTheHolidaySeasonEspeciallyWhenItSnows* is an easy password to remember (especially for someone who loves the holiday season), but with so many characters, it is unlikely to be cracked by any method. Additionally, the password length is not relevant if it is salted and hashed.
- **Blacklisted passwords should be blocked** A password system should check user passwords against a blacklist of known simple passwords such as 12345678, password, password1, qwertyui, and so on. Blacklists should also include words in a dictionary.
- **No other complexity requirements should be imposed on users** Users should be required to create passwords of at least 8 characters and not use passwords in the blacklist. However, they should not be required to create overly complex passwords. Instead, the salted hash should provide the complexity.
- **Passwords should not be set to expire** Users should not be required to change their passwords periodically. The only reason they should be required to change their password is if there is evidence of compromise.

How Strong Is Your Password?

How Strong Is Your Password? Passwords take significantly longer to crack as you add more characters and as you add complexity. One site (<https://howsecureismypassword.net/>) calculates how secure your password is by letting you know how long it would take a system to crack it. When you enter a password, the site tells you about how long it takes a desktop PC to crack your password. You shouldn't put an actual password that you're using into a site such as this, but you can use the site to see how easy, or difficult, it is to crack a similar password.

Although the times may not be exact, you can see the big difference in times when you add more characters to the password:

- Five characters takes less than a second to crack.

- Six characters with a mixture of three character types takes about 1 second.
- Six characters with a mixture of four character types takes about 5 seconds.
- Seven characters with a mixture of four character types takes about 7 minutes.
- Eight characters with a mixture of four character types takes about 9 hours.
- Nine characters with a mixture of four character types takes about 4 weeks.
- Ten characters with a mixture of four character types takes about 6 years.
- Eleven characters with a mixture of four character types takes about 400 years.
- Twelve characters with a mixture of four character types takes about 34,000 years.

Note that these timeframes are a moving target. Computer processing power continues to improve. This allows computers to crack these passwords quicker and quicker.

Something You Have

Organizations that want to increase the security of authentication often use the next factor: *something you have*. Smart cards, hardware and software tokens, and proximity cards are some examples.

Smart Cards A smart card is a small card about the size of a credit card that has an embedded certificate used to identify the user. The user inserts the card into a smart card reader that reads the data on the card to identify and authenticate the user. Because it's easy for a user to lose a smart card, most smart card authentication systems require users to enter a PIN or a password in addition to using the smart card.



EXAM TIP Combining the something you have factor with the something you know factor provides two-factor authentication. For example, when a user logs on with both a smart card and a PIN to authenticate, the user is authenticating with two-factor authentication. Two-factor authentication is more secure than just a single factor of authentication.

Hardware Tokens and One-time Passwords A hardware token is a small device that displays a number. The number changes periodically, such as every 30 seconds, and the user enters the number for authentication. This number is a one-time password because the user can only enter it once. You'll often see this called a *dynamic password* to contrast it with a static password that stays the same. A one-time password reduces the risk of an attacker capturing the password and reusing it.

Authentication servers know the number displayed on the token, as shown in [Figure 2-5](#). A user typically authenticates via a website by entering the username, password, and the number displayed in the token, which is sent to the authentication server. If the user enters the correct number, it proves that the user has the token. As long as the user entered the correct username and password, this provides authentication. As with smart cards, hardware tokens are typically used with another factor of authentication. In this example, the user entered a password and the number displayed on the token.

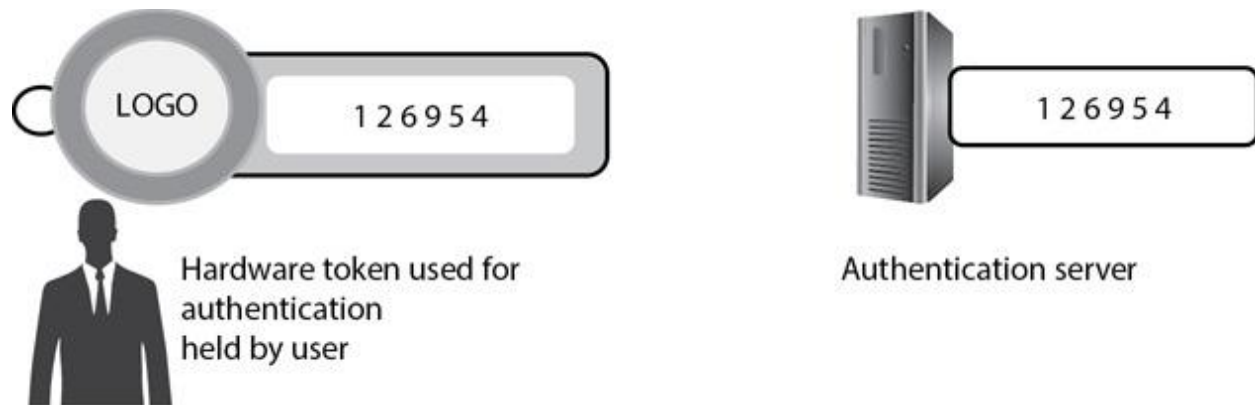


Figure 2-5 Token used for token-based authentication

Hardware tokens can use either a synchronous dynamic password or an asynchronous dynamic password:

- **Synchronous dynamic password** A token using a *synchronous dynamic password* changes the password at specific times, such as every 60 seconds. This requires the token and the server to be synchronized with the same time.
- **Asynchronous dynamic password** A token using an *asynchronous dynamic password* doesn't require exact clock synchronization. Instead, it uses another method to create one-time dynamic passwords. As an example, the user can enter a PIN into the device, which prompts the token to generate a one-time password.

Software Tokens and One-time Passwords Software tokens are being used more and more today. These use open source protocols such as the open-standard Time-based One-Time Password (TOTP) protocol or the HMAC-based One-Time Password (HOTP) protocol to create one-time passwords. Google Authenticator and Symantec's VIP Access are two examples of software tokens.

The TOTP protocol creates synchronous dynamic passwords based on time. The HOTP protocol creates asynchronous dynamic passwords. Instead of using a chronological time, HOTP uses an incrementing counter combined with a secret key, known by both the server and the token. HOTP creates a hash using the hash message authentication code (HMAC) and then reduces this hash to a six- to eight-digit HOTP value. Because HOTP uses an incrementing counter, these one-time passwords remain valid until used.

One-time Password In Everything (OPIE) is another type of one-time password used in networks. OPIE is based on S/KEY, a one-time password system used on some UNIX systems. The goal of both OPIE and S/KEY is to generate an asynchronous password known to both a client and a server and used in a single session. Both S/KEY and OPIE combine the user's real password with other data, and use a hashing algorithm such as Message Digest 5 (MD5) to create a one-time password.



EXAM TIP Tokens synchronized with a server based on time create synchronous dynamic one-time passwords. Asynchronous dynamic passwords are not based on time. HOTP, OPIE, and S/KEY are some examples of protocols that create asynchronous dynamic passwords.

Proximity Cards A proximity card includes data electronically embedded within the card. It doesn't have power, but does have an inductor that can be electrically excited when it passes through a magnetic field. This provides enough power for the proximity card to transmit data to a nearby receiver. For example, some credit cards have this technology embedded within them. When you pass the card over the proximity card reader, it activates the card's electronics and transmits your data to the card reader.

Many organizations issue identification badges to employees, which they show to security guards for identification, and some identification badges include proximity card capabilities. For example, an employee can wear the badge for identification and also use it to open doors by passing the badge in front of a proximity card reader.

Something You Are

Biometrics are used to provide authentication as *something you are* by identifying unique characteristics of a person, such as a fingerprint. Users register the biometric value and associate it with an identity such as a username. Later, users authenticate by entering the username and presenting the biometric value. For example, a user can place their finger in a fingerprint scanner. The following list shows some of the biometric methods used for authentication:

- **Fingerprints and thumbprints** Fingers and thumbs have unique patterns that law enforcement personnel have used for decades to identify individuals. Many IT systems now use fingerprints and thumbprints for authentication. As an example, many laptops and small devices include a fingerprint scanner. However, some fingerprint scanners can be tricked with a picture of a fingerprint or even a duplicate fingerprint transferred to and from a gummy bear.
- **Palms** A palm scanner can measure the vein pattern in a person's palm using an infrared scanner. Users don't need to touch the scanner, but instead just hover their hand over the scanner.

- **Retina** A person's eyes have a pattern of blood vessels at the back of the eye, and a retina scan uses an infrared light to measure that pattern. While retina scans are very accurate, people object to them because they can reveal medical conditions. Also, retina scans typically require physical contact with the scanner.
- **Iris** The area surrounding the eye's pupil is the iris, which is almost as unique as the retina. Iris scans are more acceptable to users because they don't require physical contact and cameras can take pictures of the iris from a distance. However, lighting can affect the accuracy of an iris scan, and some iris scanners can be tricked with a high-quality picture.

Behavioral biometrics identify behavioral traits of an individual. For example, keystroke dynamics measures the pattern and rhythm as a user types on a keyboard. Similarly, handwriting analysis can sometimes verify that a specific person wrote a note. However, these methods aren't as reliable for authentication as biometric methods that focus on a person's physical characteristics.



NOTE While this section focuses on using biometrics for authentication, biometrics can also be used for identification. For example, law enforcement has used fingerprint technologies for more than 100 years for accurate identification of criminals because fingerprints are unique for each person.

One of the challenges with biometrics is the potential for errors. When considering a biometric authentication system, it's important to understand the different error types and the accuracy of the system. Refer to [Figure 2-6](#) while reading the following error rate definitions:

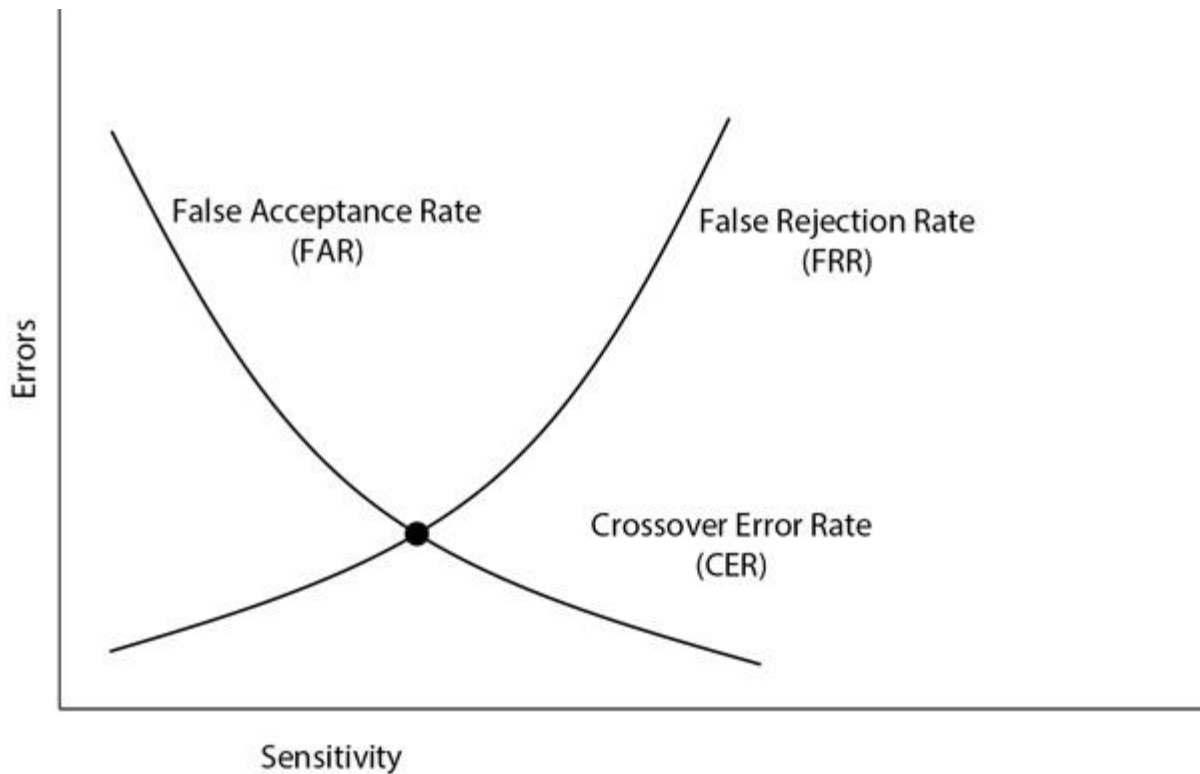


Figure 2-6 Biometric error rates

- **False Rejection Rate (FRR)** Also called a type 1 error, FRR refers to the percentage of times a biometric system falsely rejects a known user. Instead, the system indicates that the user is unknown.
- **False Acceptance Rate (FAR)** Also called a type 2 error, FAR refers to the percentage of times a biometric system falsely identifies an unknown user. Instead, the system indicates the user is a known user.
- **Crossover Error Rate (CER)** Also called the Equal Error Rate (EER), CER identifies the point where the FAR and FRR of a biometric system are equal or cross over each other on the chart. A lower CER indicates a better-performing biometric system.



EXAM TIP The three primary performance measurements used in biometrics are the False Acceptance Rate (FAR), the False Rejection Rate (FRR), and the Crossover Error Rate (CER). Systems with low CERs are more accurate than systems with high CERs.

Biometric Errors

Ideally, a biometric system won't have any errors at all, but the cost of such a system is prohibitive for most organizations. Instead, an organization often weighs what it considers more important and adjusts the system to accept more of one type of an error than another error.

For example, consider a Top Secret military facility that chooses to use biometrics for authentication. The most important result here is that the system does not authenticate unauthorized personnel. In other words, the FAR must be as close to zero as possible. However, it is acceptable to have a higher FRR where the system occasionally rejects known users.

Multifactor Authentication

Multifactor authentication uses two or more factors of authentication. This can be any combination of two of the three factors (something you know, something you have, and something you are) or a combination of all three factors.

Some examples of multifactor authentication include the following:

- A smart card (something you have) and a PIN (something you know)
- A fingerprint (something you are) and a password (something you know)
- A hardware token (something you have) with a username and password (something you know)

A key point is that the different authentication methods must use different factors. For example, using a password and a PIN is not multifactor authentication because it uses two authentication methods of the same type, something you know.



EXAM TIP Multifactor authentication requires users to use more than one factor of authentication. When two factors are required, it's also called two-factor authentication. Multifactor authentication is more secure than any single authentication type used by itself.

Two-Step Verification

Many organizations have implemented two-step verification using HOTP and/or TOTP. Google 2-Step Verification is one example. If you have a Google account, you'll often be asked for your username, password, and something else, such as a six-digit code, as a second step. Google sends this code to your phone via text, a voice phone call, or a Google mobile app.

Imagine that you registered your smartphone number with Google. Later, when you log on to your Google account with your tablet device, Google notifies you that it has sent a code to your smartphone and prompts you to enter the code. You then look at your phone to get the code and enter the code onto your tablet. If you use your tablet frequently to log in to Google, this can get

tedious quickly. However, Google 2-Step Verification allows you to set an option to skip the verification on that device or computer. It records information about your device so that it can identify it later. The next time you log on with the tablet, it won't challenge you for the code again.

However, if Hacker Harry somehow gets your password and tries to log in to your Google account, Google 2-Step Verification will challenge him to provide the code. Google will send you the verification code telling you to use this to log on. Hacker Harry won't have your phone, so he won't see the code and he won't be able to log on. This also gives you a strong indication that your password may have been compromised. Change it as quickly as you can.

When using two-step authentication in this scenario, your smartphone is using a software token and operating as a hardware token in the something you have factor. Unless you have your smartphone with you, you won't be able to log on.

Other vendors have similar processes but name them slightly differently. Amazon uses Two-Step Verification (spelling out Two and using initial capitalization), while Google uses 2-Step Verification (using the number 2 and initial capitalization). Microsoft names its process generically as two-step verification.

Reviewing Identification

A username is a primary method of identification, but it's important to realize that there are more methods. Further, many of these methods are used for both identification and authentication.

Biometric systems are effective at accurately identifying individuals using a one-to-many search. As an example, the TV show *Person of Interest* showed many security technologies, including face-recognition technologies, which use cameras to capture the facial characteristics of a person and then search a database of many faces to find a match. In this fictional show, they used face-recognition technologies to track activity of different people. Similarly, many passport-free border crossings use iris scans to identify individuals.

In contrast, biometric systems used for authentication use a one-to-one match. For example, if a user claims an identity with a username and then authenticates with a fingerprint, the authentication system compares the user's fingerprint against the fingerprint associated with the user.

Smart cards can also be used for identification. A user inserts the smart card into a reader for identification and then enters a PIN for authentication. For example, the U.S. government uses Common Access Cards (CACs) and Personal Identity Verification (PIV) cards, which are specialized types of smart cards. They include a photo, which users can also present for identification.

Single Sign-on Authentication

With single sign-on (SSO), a user authenticates once and then the system uses the same credentials for the entire session. This includes when the user accesses any resource in the organization, or with a trusted partner. This increases security because users have to remember only a single set of credentials and are less likely to write them down.

Consider the alternative in an organization that has seven servers that any user can access but where each server requires a different username and password. Users have to memorize seven sets of credentials. Even if the usernames are the same on each server, it is almost impossible to ensure that the passwords remain the same on all systems unless the passwords are never changed.

However, if the organization implements an SSO system, users would need to remember only a single username and password. This provides convenience for the users. Additionally, with fewer centralized user accounts, it decreases the administrative workload. For example, it is easier for administrators to implement password policies with a centralized SSO system. Of course, the major drawback of an SSO system is that if an attacker obtains a user's credentials, the attacker can access multiple systems. The following sections discuss some commonly used SSO systems.

Kerberos

Windows domains and Linux/UNIX realms use Kerberos as the underlying protocol for SSO. A user can log on to the domain at the beginning of the workday, and the system uses the same credentials to access any resources in the domain.

Kerberos is a vendor-neutral authentication protocol originally developed at the Massachusetts Institute of Technology (MIT) for UNIX realms. It uses a complex process of issuing tickets to authenticated accounts and then uses the tickets to access resources. Microsoft has adapted Kerberos for its Windows domains since Windows 2000.

The Kerberos server issues different types of tickets used for authentication, and is often referred to as a ticket-granting server. Kerberos uses symmetric keys and symmetric encryption (also called secret key cryptography) to protect the confidentiality of the Kerberos transmissions, and you'll also see the Kerberos server referred to as a Key Distribution Center (KDC).

The following steps and [Figure 2-7](#) provide an overview of how Kerberos works:

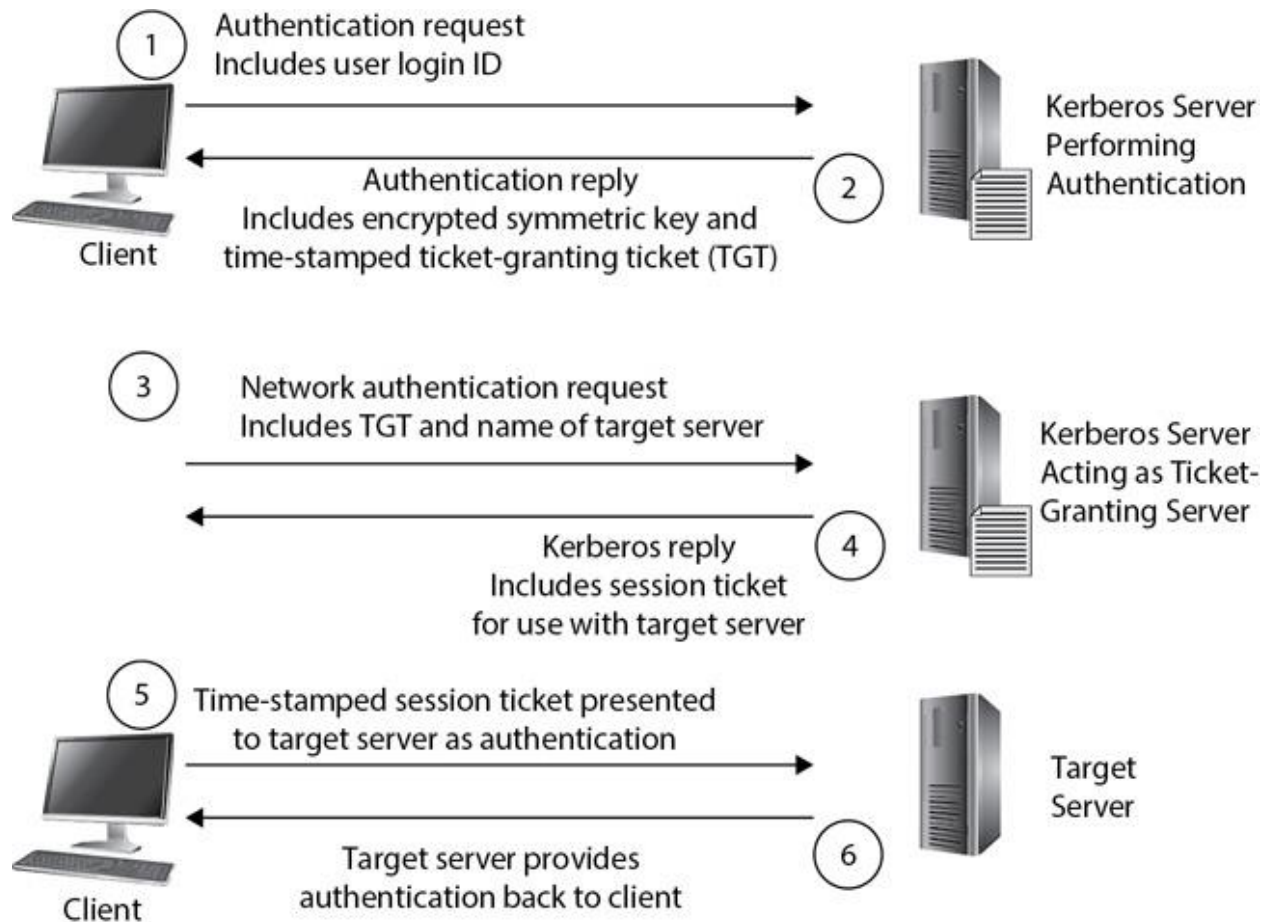


Figure 2-7 Kerberos ticket-granting process

1. The client (the user's computer) sends an authentication request to the Kerberos server. This includes the user's login ID or username, but not the user's password.
2. After verifying the user's login ID with a database (such as Microsoft Active Directory), the Kerberos server creates a symmetric key that will be known only to the client and the Kerberos server. It encrypts this symmetric key using a hash of the user's password. It also encrypts a time-stamped ticket-granting ticket (TGT) with a separate key, and sends both the encrypted symmetric key and encrypted TGT to the client.

The client decrypts the symmetric key with a hash of its password. It can then use this symmetric key to encrypt traffic between the client and server. Note that this verifies the password too. If the client doesn't know the correct password, it isn't able to decrypt the symmetric key.

3. When the client wants to access resources on a target server (such as a file server on the network), it sends the encrypted TGT to the ticket-granting server with the name of the target server.

4. The ticket-granting server replies with an encrypted ticket (also called a session ticket) that the client uses to communicate with the target server. The ticket-granting server also provides an encrypted symmetric key that systems use to encrypt information between the client and the target server.

5. The client presents the encrypted time-stamped session ticket to the target server for authentication.

6. The target server decrypts the session ticket and verifies it. It then replies with a message, which includes the session ticket's time stamp. This authenticates the server to the client (providing mutual authentication), because other servers would not be able to decrypt the session ticket.

Kerberos also provides a level of confidentiality and integrity for information. It helps prevent the unauthorized access to data by ensuring that only authenticated users can access data on target servers. This also contributes to integrity by preventing unauthorized modifications of data on target servers.



EXAM TIP Kerberos provides authentication on a network and contributes to the confidentiality and integrity of information. It uses symmetric encryption (also called secret key cryptography) to encrypt tickets in a secure format between systems. Kerberos requires a database of accounts and time synchronization for all systems.

Federated Access

Many SSO systems use federated access technologies. Federated access allows users in different networks to log on only once, even if they are accessing multiple systems. The systems can be different operating systems, owned and managed by different organizations.

For example, a company can have an intranet website used by employees that also includes links to outside websites. The outside websites can be for employee financial services, such as 401K information, or insurance company websites used to provide information on an employee's health plan.

Users log in to their work account using their regular account. When they access the intranet website, it uses typical SSO technologies such as Kerberos to grant them access. Then if they click on any links to the external sites, the federated access SSO system verifies the user's identity and provides access without additional user interaction.

Each of the different websites has its own independence or autonomy. In other words, the owners of the financial website are not dependent on information on the intranet website or the health

insurance website. However, the different websites do provide information on the identity of users and provide assurances to the other sites that the user has authenticated.

Many federated SSO systems share information via a federated database. A federated database does not include passwords of users, but it does include identity information needed by each of the sites. It provides the users with a federated identity, which each of the sites can use.



EXAM TIP Federated access provides SSO to different operating systems or networks. A federated database provides central SSO authentication.

SAML

Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML)–based data format used for SSO on the Internet. As an example, consider two websites hosted by two separate organizations. Normally, a user would have to log on to each site separately. However, the organizations can use SAML as a federated identity management system. Users authenticate once with the first website, and they are not required to authenticate again when they access the second website.

Many online banking sites use SAML for SSO. For example, the banking site might have one service for accessing checking and savings accounts, another service for online bill paying, and another service for handling mortgages. With SSO, the user is able to log on to the primary banking site one time, and then access all the services without logging on again.

SAML includes three roles:

- **Principal** The principal is typically the user that logs on to the system. If necessary, the user might request a principal identity from the identity provider.
- **Identity provider** The identity provider creates, maintains, and manages the identity information for principals.
- **Service provider** A service provider is the entity that provides services to principals. For example, a banking institution that hosts different banking services is the service provider.

The primary purpose of SSO is for identification and authentication of users. Users claim an identity and prove that identity with credentials, such as a password. SSO does not provide authorization. For example, if a user logs on to the banking site, SSO doesn't control what services the user can access. However, SAML includes the ability to transfer authorization data between the bank's systems. In other words, it's possible to use SAML for SSO authentication and for authorization.



EXAM TIP SAML provides SSO for web-based application servers on the Internet. It is an XML-based standard used to exchange authentication and authorization information between different parties.

Other SSO Technologies

Two newer SSO technologies you should be aware of are OAuth and OpenID Connect. Many online websites are using them to allow users to log on to their sites with an account they've created elsewhere. Many popular websites such as Google, Facebook, PayPal, and Microsoft support these technologies.

As an example, imagine you have a Google account. You want to log on to a website to post a comment, but doing so requires you to have an account. Instead of creating another account, the site gives you the option of logging on with your Google account.

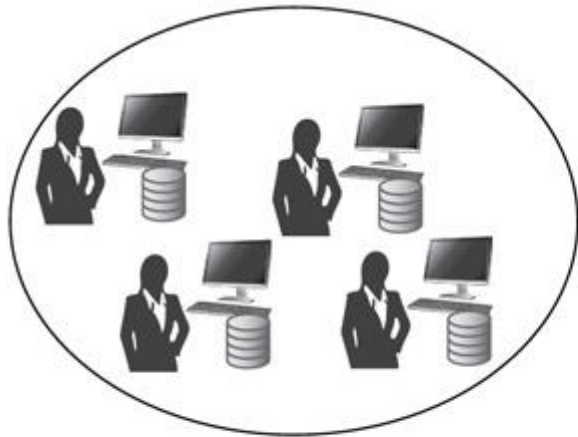
Most organizations support OAuth 2.0 but not OAuth 1.0. OAuth 2.0 provides an open source authentication framework that organizations can build on. OpenID Connect is an authentication layer that can be used with OAuth 2.0 to extend its capabilities. Combined, OAuth and OpenID Connect allow websites and online web applications to verify the identity of end users without managing the users' credentials.



EXAM TIP Kerberos, SAML, OAuth, and OpenIDConnect each provide SSO capabilities. With SSO, users only have to log on once and then use the same credentials to access multiple resources.

Centralized vs. Decentralized Authentication

Authentication can be decentralized or centralized. Consider [Figure 2-8](#). The group of users on the left is in a decentralized environment. Every computer has a separate database that stores credentials. If Sally needed to log on to all four computers in this network, she would need to have four separate sets of credentials—one for each system.



Decentralized Authentication
Four Passwords Needed for Four Systems



Centralized Authentication
One Password Needed

Figure 2-8 Centralized vs. decentralized authentication

On the other hand, the computers on the right are using centralized authentication. Credentials for the users are stored on a central server. Any user can log on to the network once and then access any computer in the network (as long as the user has permissions). For example, if a computer is part of a Microsoft domain, the central server will be a domain controller and hold accounts for all users in the domain.

Centralized authentication is used for single sign-on networks. It's not possible to use SSO without some type of centralized authentication.

Offline Authentication

Offline authentication allows users who have logged on to the system at one time to still log on even when they are disconnected from a network.

For example, consider a consultant named Steve with an account in a Windows domain named mcgrawhill. He uses his laptop computer to log on to the domain while he's at work. When his company deploys him to a customer site to provide consulting services, he's no longer connected to the mcgrawhill domain, so there's no way that a Kerberos server in the domain can authenticate him if he is not connected to the domain.

However, he is still able to log on while he's offline. In a Windows environment, the system uses cached credentials. In other words, when Steve logged on to his laptop, encrypted data about his credentials were stored (cached) on his system. If he's disconnected when he logs on, these cached credentials are compared to the credentials he's entering, and if they match, he is authenticated.



TIP A user will not be able to access network resources while using cached credentials. The user can only access resources on the local system.

Linux and UNIX users have a similar capability when they are interacting with a Microsoft domain. As long as offline authentication is selected when the computer joins the domain, they will be able to use domain credentials even when they are disconnected.

Device Authentication

Device authentication methods focus on hardware instead of the user. The goal is to prevent unauthorized devices from accessing the network.

Windows domains require users to log on from known computers that have joined the domain. This requires the computer and the user to authenticate before the user is able to access domain resources. However, there are times when you might want to allow network access without joining the device to the domain.

One method used for device authentication allows or blocks devices with media access control (MAC) address filtering, described in [Chapter 3](#). However, this is very tedious and users can easily bypass MAC address filtering if they know how to change their MAC address.

Another method is device fingerprinting that identifies devices based on several characteristics. Many of the following characteristics are easy to capture when a user attempts to connect to a network with a web page: operating system and version, IP addresses, browser, browser fonts, browser plug-ins, time zone, data storage, screen resolution, cookie settings, and HTTP headers.

As an example, SecureAuth created SecureAuth Identity Provider (IdP), which provides mobile device management (MDM), including device authentication. Authorized devices are registered with the system. Later, when a user attempts to access the network, SecureAuth IdP verifies the device is recognized. If not, it blocks access. [Chapter 9](#) discusses MDM in more depth.



TIP Marketers and web analytics services have used similar device fingerprinting services for years. Canvas fingerprinting is a common method that allows companies to track user activity.

Implementing Access Controls

An access control provides a mechanism to restrict or control access to resources. These resources can be logical, such as files and folders hosted within a network, or physical, such as

servers within a locked server room. Access control starts with identification and authentication. Once individuals have been reliability authenticated, you can control their access.

Comparing Subjects and Objects

Two primary terms related to access controls are *subjects* and *objects*. In short, a subject accesses a resource, and an object is the resource being accessed. As a simple example, if a user accesses a project file stored on a server, the user is the subject and the project file is the object.

With this in mind, you would first ensure that subjects are authenticated to validate their identity and then use access controls (such as permissions) on the objects to control the access. The access controls ensure that only certain subjects have access to the objects.

In some examples, an entity can be both a subject and an object. Consider [Figure 2-9](#). A user requesting a web page from a server is clearly a subject. Similarly, the web pages served by the web server are the objects. However, this web server is retrieving data from a back-end database (protected behind a firewall) to build the web page. When the web server retrieves the data, it is a subject and the database is the object.

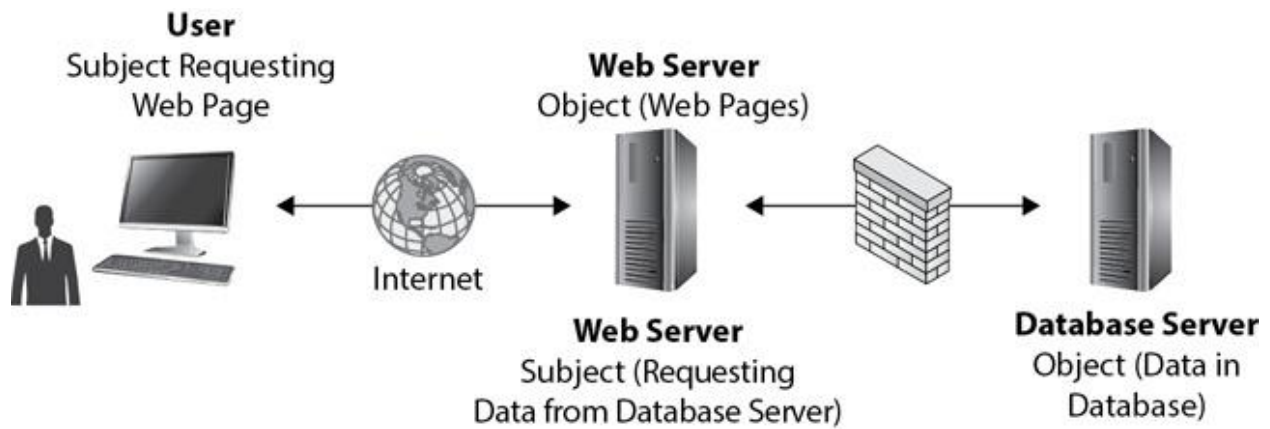


Figure 2-9 Comparing subjects and objects



TIP It may not be apparent, but web servers commonly access back-end databases. For example, [Amazon.com](#) hosts web servers to serve the web pages and implements powerful databases to host the data. When you search for a product, the web server sends the query to the database and then formats the reply into the web page. This separation provides security for the database and performance gains for the web server.

While there aren't any formal access control models called subject-based access control or object-based access control, it is possible to control access based on subjects or objects. The following sections dig into subjects and objects a little deeper, including how they can be used

for access control. The upcoming “Comparing Access Control Models” section describes how different formally named models control access based on subjects or objects.

Subjects and Subject-based Access Control

Access control systems can treat any of the following as subjects:

- Users
- Computers
- Applications
- Networks

You can create accounts to represent most of the subjects. For example, in most networks, each user has at least one user account and each computer has a computer account. Similarly, many applications use service accounts to access resources. You can then assign specific permissions to these accounts for any object they need to access.

Additionally, an access control system can often apply different attributes to the subject to control access further. The following are some subject attributes:

- **Time (temporal)** For example, it’s possible to prevent a user from logging on to a system outside of normal working hours.
- **Remote access attribute** [Figure 2-10](#) shows how the Network Access Permission is set to Allow Access for a user in a Microsoft Active Directory domain. This setting allows a user to access a remote access server, such as a virtual private network (VPN) server, from any location that has Internet access.

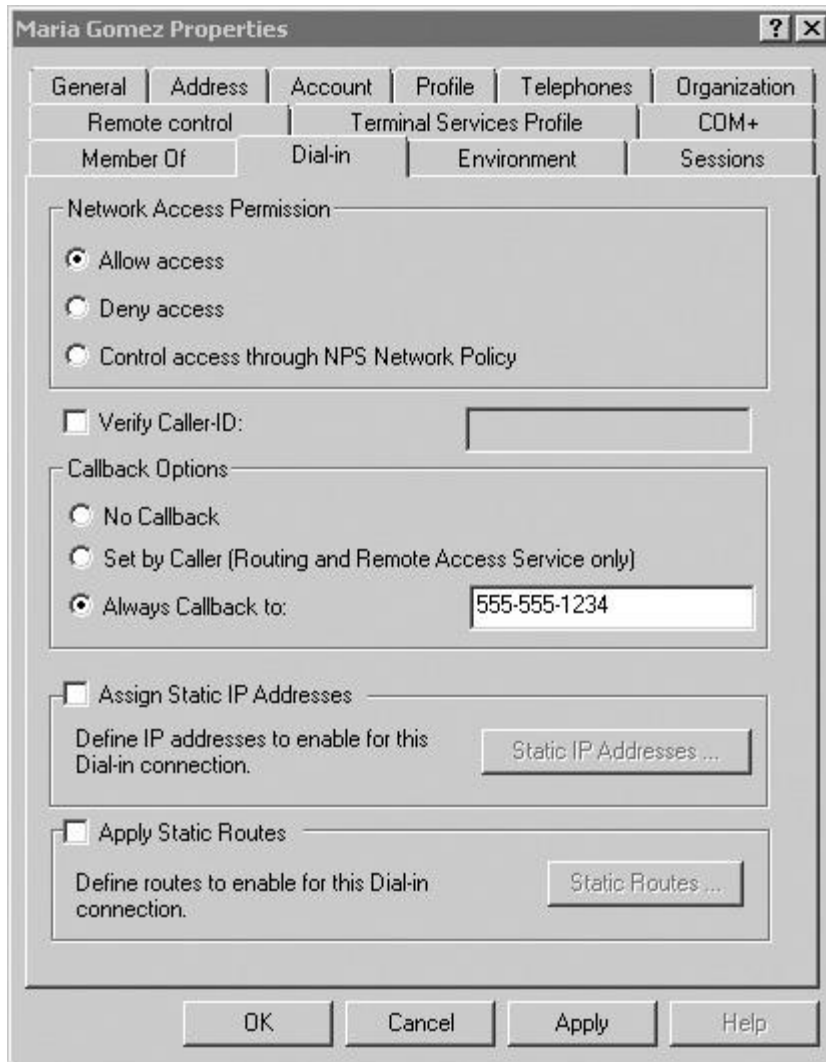


Figure 2-10 Viewing the Network Access Privilege for a user

- **Location** Within a network, it's possible to restrict users to only log on to specific computers, based on the computer name or MAC address. Many Internet-based systems analyze the IP address of the user's device to determine the user's location. A global positioning system (GPS) can identify a user's physical location and restrict access based on the GPS coordinates. A geofence can create a virtual boundary (such as within an organization's building) and block access to users outside of the geofence.
- **Role or group membership** Subjects (such as users) are often assigned roles or placed into groups to grant them access to objects. The "Role-based Access Control" section a bit later discusses this model in more detail.

Many organizations assign administrators two accounts. Administrators log on with a regular user account for day-to-day work, and this account gives them limited access. When they need administrative privileges, they use their administrator account. As an example, [Figure 2-11](#)

shows Active Directory Users and Computers (ADUC) in a Microsoft domain. It shows four accounts in the IT Department Organizational Unit. Each of the users has a regular account for day-to-day work and an administrator account used to perform work requiring elevated privilege.

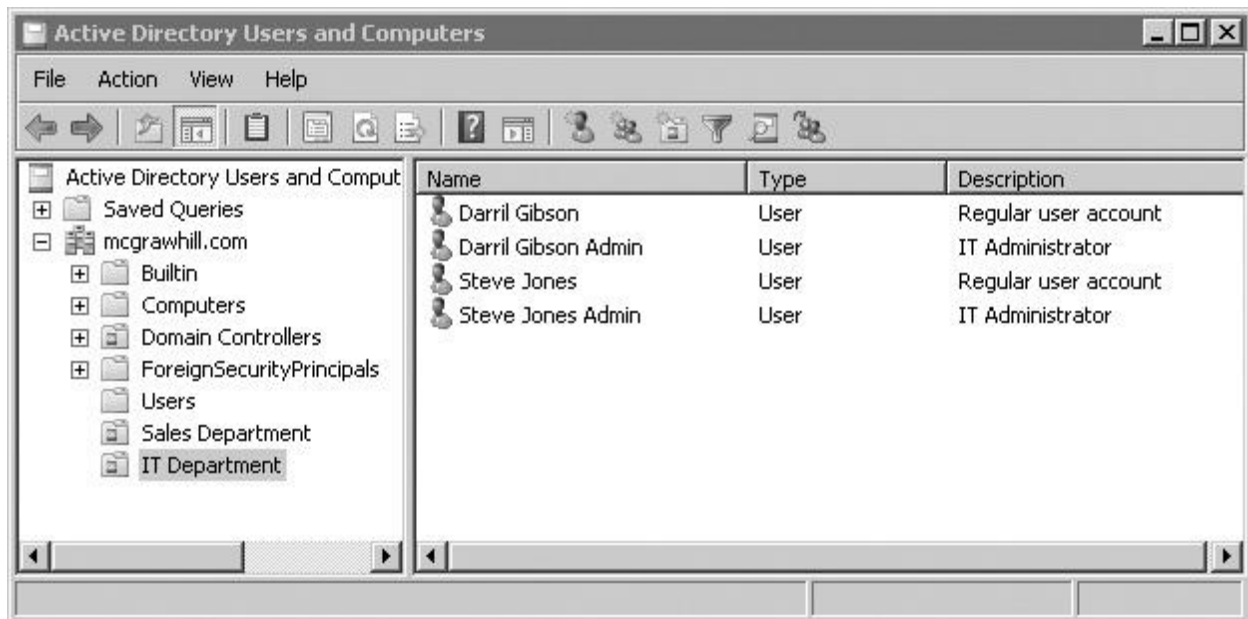


Figure 2-11 Viewing accounts in a domain



NOTE [Chapter 1](#) covers privileges within the context of least privilege. As a reminder, privileges include both permissions and rights. Users are granted permissions to access resources such as files and they are granted rights to perform actions such as modifying the time on a system.

Objects and Object-based Access Control

Access control systems can treat any of the following as objects:

- Data (stored in files, folders, and shares)
- Hardware (such as desktop computers, servers, and printers)
- Applications (such as a web server application)
- Networks (such as an Internet connection or an internal connection)
- Facilities (controlled with physical security)

Most objects allow various levels of access. For example, a user might be able to view a web page but not have permission to modify it. Similarly, regular employees can get into the company building but they don't have access to secure server rooms.

It's common to store groups of objects with the same access requirements together. For example, if salespeople need identical access to 100 files, you can store all 100 files in a folder and then assign the permissions to the folder. This simplifies administration. Permissions are the most common type of access control applied to objects.

Logical Access Controls

Logical access controls (also called technical access controls) are implemented with technologies. Many logical access controls use access control lists (ACLs). For example, a router has an ACL with multiple rules that identify the traffic allowed in or out of a network. Similarly, files and folders often use ACLs to identify who can access the resource and what level of access each user has.



EXAM TIP Within IT systems, the security kernel does much of this work. The security kernel is the central part of the operating system that controls access to the system's resources.

Comparing Access Control Models

There are several formally defined access control models available. At this stage of your IT career, you aren't expected to design access based on any of the models. However, you should be aware of the models that exist and have a basic understanding of how they work. The following sections summarize some basic characteristics of these models.

Discretionary Access Control

The Discretionary Access Control (DAC) model provides the most granular level of access control. It is an identity-based (or subject-based) model and allows data owners to assign permissions to subjects at the most basic level. For example, you can assign read permission to a single user. File systems such as New Technology File System (NTFS, used by Microsoft) and Network File System (NFS, used on UNIX-based systems such as Solaris and Linux) use the DAC model.



EXAM TIP In the DAC model, users have ownership of the data and can exercise full control over it, including assigning permissions to others. The DAC model allows data owners to assign permissions at the granular level.

DAC uses ACLs, or, more specifically, discretionary access control lists (DACs). As an example, consider [Figure 2-12](#). It shows the permissions assigned to different subjects for the Sales Data object (a folder on a server).

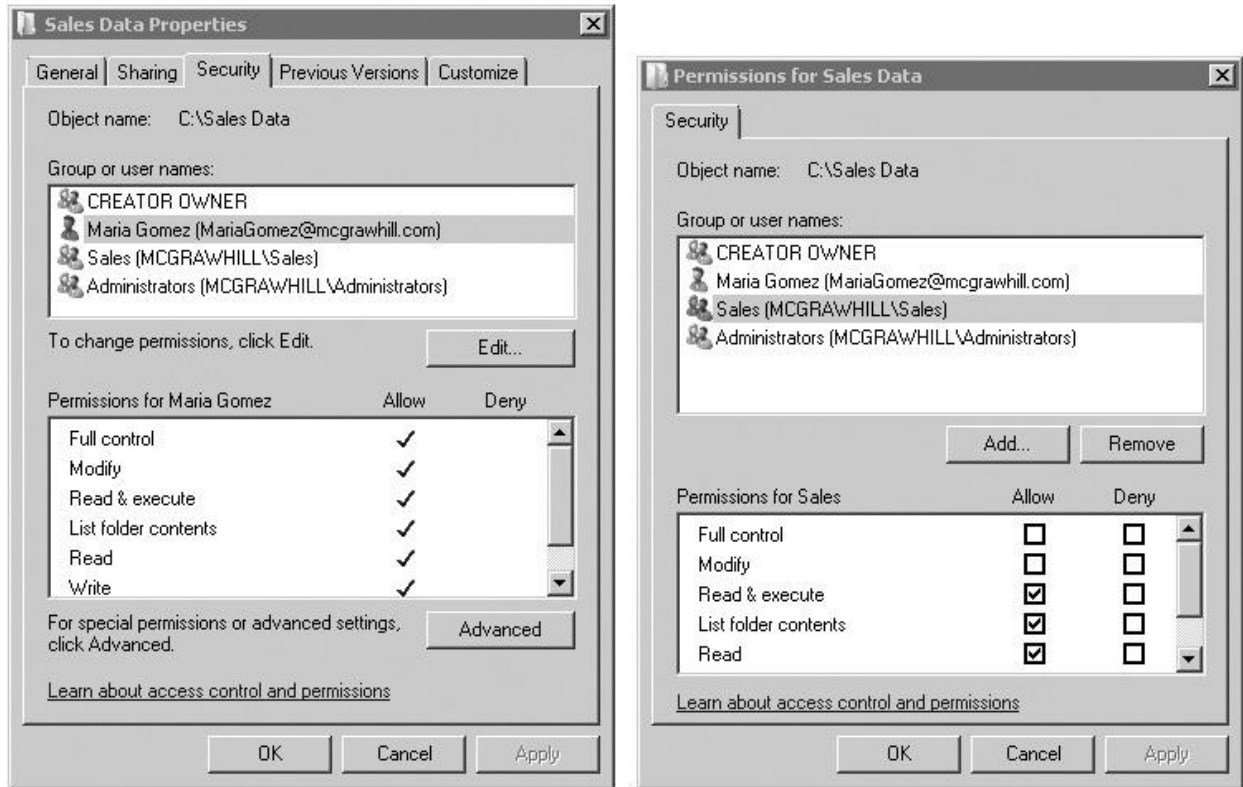


Figure 2-12 Viewing permissions as a DACL



TIP Windows systems identify users and groups with a security identifier (SID) that you'll rarely see. Each entry in the DACL includes the SID and the permission assigned to that SID. However, interfaces such as the one shown in [Figure 2-12](#) perform a lookup to identify the user-friendly name for the SID. They display the name instead of the SID.

A DACL is composed of several access control entries (ACEs), and [Table 2-1](#) shows the underlying ACEs for [Figure 2-12](#). A distinguishing feature of the DAC model is that users have full control over any objects they create. The CREATOR OWNER group is a special group that ensures that users are automatically assigned full control permission to any file or folder they create. Maria Gomez (the sales manager) has full control permission on the folder. The Sales

group would include all personnel in the Sales department, and it has read & execute, list folder contents, and read permissions. The Administrators group has full control permission.

| | Subject | Permission(s) |
|-----|----------------|--|
| ACE | CREATOR OWNER | Special |
| ACE | Maria Gomez | Full control |
| ACE | Sales | Read & execute, list folder contents, read |
| ACE | Administrators | Full control |

Table 2-1 Individual ACEs Creating a DACL



EXAM TIP The DAC model uses a discretionary access control list as shown in [Table 2-1](#). This is a list of subjects along with their rights and/or permissions.

Non-Discretionary Access Control

In non-Discretionary Access Control (non-DAC) models, security administrators control the access granted to users. In contrast, in a DAC model, users have ownership of their resources and users have full control of the resources they own. In general, any model that is not a DAC model is a non-DAC model.

Some operating systems implement non-DAC models for system file access. This prevents malware from taking ownership of any critical or sensitive system files or modifying permissions on any of these files. Users still own and manage their own files using DAC, but the non-DAC model methods protect system files. In other cases, the operating system uses a non-DAC model exclusively, such as Mandatory Access Control (described later in this chapter). Other non-DAC models include role-based, rule-based, and attribute-based models.

Role-based Access Control

The Role-based Access Control (Role-BAC or RBAC) model uses roles or groups to determine access. Subjects are placed into specific roles and object permissions are granted to the roles. Although the Role-BAC model doesn't provide the granularity offered by DAC, it is easier to implement for large groups of people.



TIP Many operating systems support the use of groups that function as roles. Administrators place user accounts into groups and assign appropriate privileges to the groups.

Consider [Figure 2-13](#). An organization has a group of salespeople in the Sales department. The organization wants each of these users to have access to a printer, a file server, and several folders and files. An administrator adds the accounts for each salesperson to the Sales Role and then assigns appropriate permissions to the Sales Role. These permissions apply to all users in the Sales Role without having to assign the permissions individually.

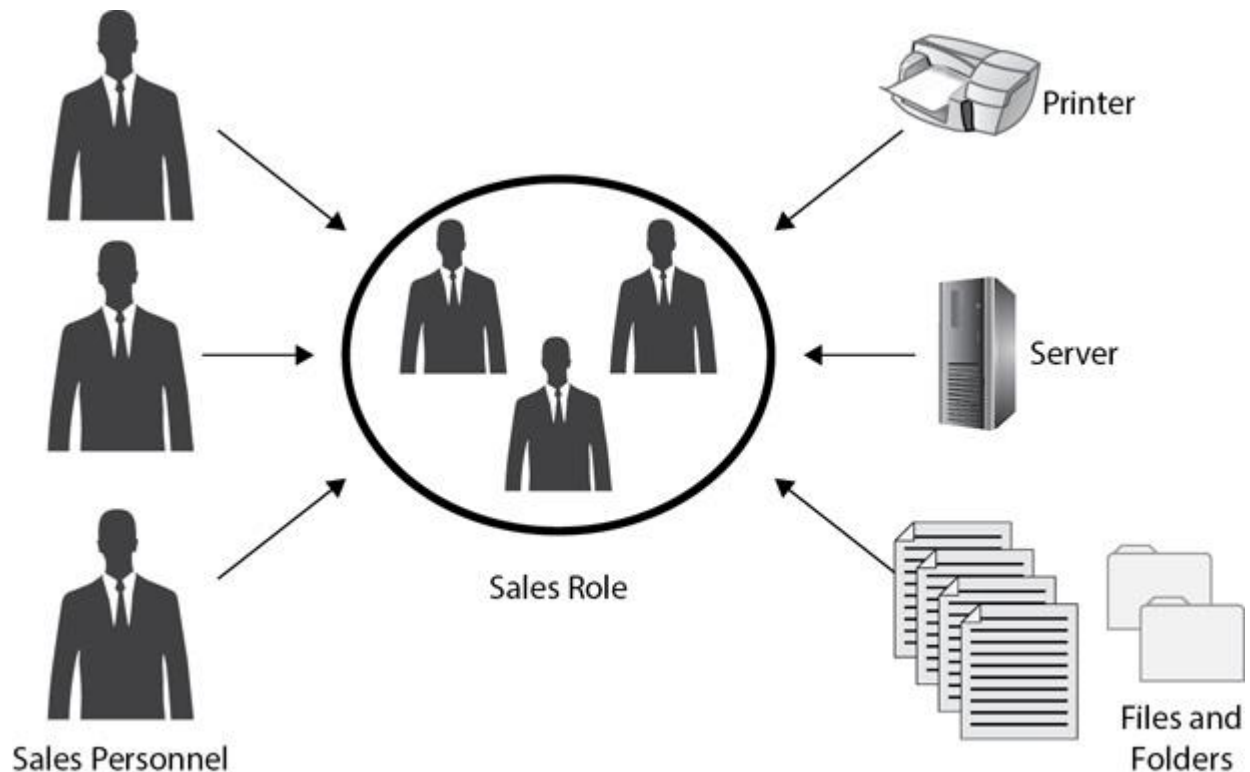


Figure 2-13 The Role-BAC model

If the organization hires another salesperson, administrators place the new salesperson's user account into the Sales Role. This gives the new employee all the appropriate access in a single step. If a salesperson transfers to a different department, an administrator would remove the account from the Sales Role. When you remove a user from a group or role, it revokes all the permissions they had as a member of the old group or role. An administrator would add the user account to appropriate roles in the new department. This reduces the administrative workload and improves security.



EXAM TIP Role-based models reduce the administrative workload and are very useful in organizations with high employee turnover. (ISC)² has been using more and more acronyms within their exam questions and you may see Role-based Access Control represented as either RBAC or Role-BAC.

In contrast, if the organization doesn't use roles, but instead forces administrators to assign permissions individually, it becomes very difficult to manage. Looking at [Figure 2-13](#), administrators would have to assign permissions individually for each salesperson at each resource. When a new salesperson is hired, administrators would have to remember all the resources that salespeople need to access and assign permissions individually. When employees transfer, administrators should remove their permissions to resources associated with their previous job. However, without roles or groups, this step is very difficult and often skipped. This results in *permission creep*, where employees just collect more and more permissions. All of this can be challenging with just a few people, but it is unmanageable if your organization has hundreds of users.

Rule-based Access Control

The Rule-based Access Control model is another non-DAC model. Administrators create rules that determine access to resources. As an example, routers have rules within an ACL. These rules identify what traffic the router will pass based on IP addresses, ports, and protocols.



TIP The last rule in a router is an implicit deny rule. It blocks all traffic that isn't explicitly allowed by previous rules. Permissions assigned in DACLs use a similar concept. For example, you explicitly grant permissions to users for a folder. If you don't assign permissions to a specific user, the system blocks that user from accessing the folder.

Attribute-based Access Control

Attribute-based Access Control (ABAC) is a more sophisticated type of Rule-based Access Control. It evaluates subject and object attributes, and grants access based on the value of these attributes. Attributes can be almost any characteristic of a user, the environment, or the resource.

Some systems implement ABAC with plain-language statements to build policies. These policies work as rules, but provide you with a lot more granularity than you can achieve with traditional Rule-based Access Control models. Software-defined networking is becoming very popular and often uses these types of policies. For example, you could create a policy rule that states, "Allow logged-on users to access YouTube videos using smartphones or tablets via the guest network."

Policy statements typically include four elements:

- **Subject** The user accessing a resource. You can use any user property as an attribute. This includes roles, group memberships, management level, and assigned department. The preceding example uses the status of the user—logged on.
- **Object** The resource that the user is trying to access. It can be a file, a database, or an application. The example uses the YouTube application. Notice that this allows the rule to cover much more than just Internet access.
- **Action** What the user is attempting to do. It could be reading or manipulating a file, accessing specific websites, or accessing website applications. The example allows access to the YouTube website via the guest network. This could prevent video streaming from overloading the primary network, while allowing it on the guest network.
- **Environment** Includes everything outside of the subject and object attributes. This is often referred to as the context of the access request. It can include the time, location, protocols, encryption, devices, and communication method. In the example, it specifies smartphones, tablets, and the guest network as environmental attributes.

While the policy statement is stated in plain language, it would typically be stored in a database using an XML format. The eXtensible Access Control Markup Language (XACML) is a specialized XML format used specifically for ABAC.

Mandatory Access Control

The Mandatory Access Control (MAC) model is a non-DAC model that uses labels to identify both subjects and objects. It provides the highest level of security among the models (MAC, DAC, Role-BAC, Rule-BAC, and ABAC) and is commonly used by the U.S. military to ensure that data is protected in mission-critical systems.

The U.S. government uses the following classifications for data, from highest to lowest:

- Top Secret
- Secret
- Confidential
- Unclassified

Just because someone has a Top Secret classification, they don't automatically have access to all Top Secret data. Instead, data owners authorize access based first on their Top Secret security clearance, and second based on their need to know the information for their job. Additionally, it's possible to create subclassifications or compartments within each classification level.

As an example, [Figure 2-14](#) shows the three classification levels of Top Secret, Secret, and Confidential. Within each level, there are additional labels identifying compartments within the

classification level. The Top Secret level includes compartments labeled Diplomat, Lucid, and Foil (names from World War II military operations). The Secret level includes compartments labeled Teardrop, Alacrity, and Marmalade, and the Confidential level includes compartments labeled Stonewall, Cockpit, and Jaywick.



| | | | | |
|-----------|----------|-----------|--|--------------|
| Diplomat | Lucid | Foil | | Top Secret |
| Teardrop | Alacrity | Marmalade | | Secret |
| Stonewall | Cockpit | Jaywick | | Confidential |

Figure 2-14 Using labels with the MAC model

Lisa has both the Top Secret label and the Lucid label. She can be granted access to any data labeled Top Secret and below as long as it doesn't have an additional label. Additionally, she can be granted access to compartmentalized data with the Lucid label. Maggie has the Secret and Marmalade labels. She can be granted access to any data labeled Secret and below as long as it doesn't have an additional label. Additionally, she can be granted access to compartmentalized data with the Marmalade label. Bart has only the Confidential label so he can only be granted access to Confidential data that doesn't have an additional label.



EXAM TIP The MAC model is a non-DAC model that uses labels to control access to data. It is the most secure model when compared to other access control models.

Some of the MAC-based architecture models are Biba, Bell-LaPadula, Clark-Wilson, and Brewer-Nash. [Table 2-2](#) provides an overview of these models, including their primary goals and some of the rules they use. The following sections explain them in a little more depth.

| Model | Goal | Rules |
|---------------|------------------------------|--|
| Bell-LaPadula | Confidentiality | No read up, no write down |
| Biba | Integrity | No read down, no write up |
| Clark-Wilson | Integrity | Certification (C) rules and enforcement (E) rules |
| Brewer-Nash | Prevent conflict of interest | Access governed by membership in groups to prevent conflicts of interest |

Table 2-2 Goals and Rules of Some Access Control Models

Bell-LaPadula David Elliott Bell and Leonard J. LaPadula designed the Bell-LaPadula model with a primary goal of ensuring confidentiality. It enforces security through two primary rules commonly called *no read up* and *no write down*. Each of these rules compares the subject’s clearance with the object’s classification.

- **Simple security property rule—no read up** Subjects granted access to any security level may not *read* an object at a higher security level. For example, if Joe is granted Secret access, he cannot read materials classified as Top Secret.
- **The * property (read as “star-property”) rule—no write down** Subjects granted access to any security level may not *write* to any object at a lower security level. For example, if Sally is granted Top Secret access, she cannot create or modify documents classified as Secret.



EXAM TIP Bell-LaPadula is a MAC model that uses the no read up and no write down rules to enforce confidentiality. These rules compare the subject’s clearance with the object’s classification. Confidentiality ensures that unauthorized personnel cannot access data.

Biba The Biba model, another MAC-based model, enforces integrity (unlike the Bell-LaPadula model, which enforces confidentiality). Biba includes two rules that are reversed from the Bell-LaPadula model:

- **Simple Integrity Axiom—no read down** Subjects granted access to any security level may not *read* an object at a lower security level, at least not as the authoritative source. For example, a captain of a ship can read orders from an admiral and consider them authoritative and actionable. However, if a seaman recruit tries to issue orders to the captain, the captain will not read them as authoritative.

- **The * Integrity Axiom (read as “star Integrity Axiom”)**—no write up Subjects granted access to any security level may not *write* to any object at a higher security level. For example, a seaman recruit cannot write orders for the captain of the ship. Similarly, the captain cannot write orders for the admiral.



EXAM TIP Biba is a MAC model that uses the no read down and no write up rules to enforce integrity. Integrity protects against unauthorized data modifications.

Clark-Wilson David Clark and David Wilson created the Clark-Wilson model, and its primary goal is information integrity, although it is more stringent than the Biba model. It also helps enforce the separation of duties principle. [Chapter 1](#) described separation of duties. As a reminder, separation of duties is a security principle that ensures that no single person has complete control over a process.

The Clark-Wilson model uses certification rules (identified as C1 through C5) and enforcement rules (identified as E1 through E4) to enforce separation of duties. The certification rules are integrity-monitoring rules, and the enforcement rules are integrity-preserving rules.

These rules are complex, but they work together to ensure there is adequate separation of the elements of any transaction. For example, consider a company that purchases products. The transaction includes someone placing an order, someone receiving the products (and verifying that they were received), and someone paying for the products (after verification that the products were received).

Imagine that Fred has sole responsibility for doing all three tasks in the transaction. This increases the potential for fraud, because he could place an order through a fictitious company, acknowledge receipt (even if no products were received), and deposit the funds into his own bank account. The Clark-Wilson model ensures that different people perform the separate tasks independently of each other.

Brewer-Nash The Brewer-Nash model (sometimes called the Chinese Wall model) is documented by Dr. David F. C. Brewer and Dr. Michael J. Nash. Financial services organizations often implement this model to help prevent a conflict of interest, and it helps enforce the separation of duties principle.



NOTE The Brewer-Nash model is often referred to as the Chinese Wall model. If you consider The Great Wall of China, one of the most impressive architectural creations in history, it evokes an image of a strong barrier. This image helps provide a visual of how the Brewer-Nash model

creates a logical barrier to prevent conflicts of interest. However, many people consider the use of the term Chinese Wall as culturally insensitive and it has been challenged for decades. Some alternatives are ethical wall or ethical screen, but you may still see the term Chinese Wall on the exam.

Data is classified in the ethical screen model using different conflict-of-interest classes. If a subject has access to data in one class, he or she cannot access data in a conflicting class.

As an example, consider traders in an investment bank who are selling financial securities to investors. These traders should all have access to the same information. In another part of the investment bank, advisors are helping a separate organization with a major merger. If the merger is finalized, it will affect investments sold by the traders. If the traders have access to this information, they could use it to increase investment returns for themselves or their clients. This type of insider trading is illegal, but that isn't enough to stop some people.

A Brewer-Nash model provides a barrier between these two groups of employees by classifying data. It can be as simple as classifying data as Trader Data and Advisor Data and indicating the two classes have a conflict of interest. Traders have access to Trader Data, and the Brewer-Nash model blocks their access to the Advisor Data.



EXAM TIP Both the Clark-Wilson model and the Brewer-Nash model can be used to enforce the principle of separation of duties. The Brewer-Nash model also helps prevent a conflict of interest.

Access Control Matrix vs. Capability Table

Administrators use various tools when implementing access control models. Two common tools are access control matrixes and capability tables. A key difference between the two is that access control matrixes are object based and capability tables are subject based.

An *access control matrix* is a list of objects along with the permissions granted for each object. You can think of an access control matrix as a group of ACLs. Each ACL represents a single object and lists all the permissions for that object.

A *capability table* is a list of subjects, along with the capabilities granted to the subjects. These capabilities include rights and permissions. For example, a capability table can list several groups such as Project Managers, Project Team Leads, and Project Members. It then lists the rights and permissions granted to each of these groups.

Participating in the Identity and Access Management Lifecycle

Identity management lifecycle refers to managing accounts through their lifetime. *Identity and access management (IAM)* refers to granting accounts access to resources. Accounts need to be configured properly when they are first created, managed while in use, and disabled or deleted when they are no longer needed.

Privileges include both rights and permissions. Rights are actions that an account can take on a system, such as backing up files, changing the time, or rebooting the computer. Permissions identify what a user can do with resources, such as reading and writing to a file or printing to a printer.

The principle of *least privilege* dictates that users should have only the rights and permissions they need to perform their jobs and no more. However, enforcing the principle of least privilege isn't a one-time event. Instead, administrators take steps to enforce this principle throughout the lifecycle of any identity.

Identity Proofing

It's important to verify a person's identity before creating a user account. Within an organization, Human Resources (HR) departments typically have processes they use for identity verification. Many HR departments require documents such as a driver's license.

Online banks and retailers often use identity-proofing methods to prevent fraud. One common method to prevent credit card fraud is the use of a credit card verification value (CVV). While criminals might steal databases with credit card numbers, these databases rarely include the CVVs too.

Cognitive passwords (described earlier in this chapter) provide identity proofing after an account has been created. For example, if a user claims to have forgotten a password, the system will challenge the user with cognitive password questions to verify the person's identity.

Provisioning and Authorization

Provisioning refers to creating accounts for users and granting them access to appropriate resources. *Authorization* refers to granting appropriate privileges for the accounts. Many organizations use groups, as described in the "Role-based Access Control" section of this chapter, to streamline the process of granting privileges. For example, administrators can create a group named Sales to represent the Sales department. They add the user accounts for people in the Sales department to the Sales group, and they assign appropriate privileges to the Sales group.

It's common to use template user accounts for different departments. For example, administrators can create an account named SalesTemplate and configure this account with appropriate group membership in the Sales department. When creating an account for a new

employee, they make a copy of the template, and it is automatically provisioned with the appropriate group membership.

Some organizations automate account creation. For example, when HR personnel process a new hire, they can enter information about the new employee into an application. This includes basics such as name, address, and the department in which the employee will be working. After the HR person enters the information, the application automatically creates the account and adds the user to the appropriate groups or roles based on the employee's job.

This type of application helps ensure that account-naming conventions are always followed. For example, if the organization wants to create accounts based on a <first name><dot><last name> format (such as darril.gibson), the application will always create it using this convention.

It's important to consider accounts that have elevated permissions, such as administrator accounts. Many organizations require administrators to maintain two accounts. They use one account for day-to-day tasks that have limited privileges, similar to any regular user. They have a second administrator account that they use only when performing administrative work. This helps limit the risks to the administrator account.

For example, if malware infects a system while a user is logged on, the malware can assume the privileges of the logged-on user. If the user is logged on with administrative privileges, the malware assumes the same level of privileges and can then take actions as an administrator.

In contrast, if malware infects a system while a regular user is logged on, the malware must use more technical methods to escalate its privileges. If an administrator is logged on as a regular user 90 percent of the time and uses the administrator account only 10 percent of the time, it reduces this risk.

Maintenance and Entitlement

Many organizations have account management policies that dictate how accounts are created and managed through their lifetime. These policies address password management, account lockouts, and procedures to add and remove privileges.

The "Something You Know" section earlier in this chapter discussed password policies. They ensure that users create strong passwords, change them regularly, and don't reuse the same passwords. It's common to enforce password policies with a technical policy such as with Microsoft's Group Policy. In other words, the password policy starts as a written policy and a technical policy enforces it.

Account Lockout Policies

An account lockout policy locks out an account after too many failed logon attempts. The goal is to prevent attackers from guessing passwords in an online brute-force attack. For example, a policy may dictate that an account is locked out for 30 minutes if someone enters an incorrect password five times in a row.

There are two common settings for account lockout policies:

- **Threshold** This identifies how many incorrect passwords are allowed. For example, if it is set to five, an account is locked out after someone enters the wrong password five times in a row.
- **Duration** This identifies how long the account remains locked out. For example, if the duration is set to 30 minutes, the account will remain locked out for 30 minutes. After 30 minutes, the user will be able to log on with the correct password. In some cases, the duration is configured so that it never expires. The account remains locked out until an administrator unlocks it.



EXAM TIP Account lockout policies protect against online brute-force attacks where the attacker attempts to guess the password. They are not effective against offline password attacks where the attacker has obtained a database of encrypted passwords and attempts to discover the actual passwords from the database.

Entitlement and the Principle of Least Privilege

Entitlement refers to the privileges granted to users, and following the principle of least privilege, it is an important part of entitlement. As mentioned in the “Role-based Access Control” section, using roles (or groups) makes it easier to manage privileges.

Roles and groups provide significant security benefits related to entitlement. For example, suppose that Maria is working in the Sales department. She has appropriate privileges to perform her job there. After earning her SSCP certification, she transfers to the IT department and starts working in IT security. Administrators add her account to IT groups related to her new IT security job and remove her account from groups related to her job in sales.

Removing a user from a group immediately removes all the privileges assigned to the user as a member of that group. In other words, removing Maria from the Sales group removes the privileges that she had while working in the Sales department. The key here is that it’s important to remove users from groups when they transfer. If processes aren’t in place to monitor entitlement, users end up having more privileges than they need to perform their jobs.

De-provisioning

Account de-provisioning refers to disabling and deleting inactive accounts. At a minimum, inactive accounts need to be disabled. When it’s determined that the account is no longer needed, it should be deleted.

As an example, if Sally leaves the company to take an executive position in another company, her account needs to be disabled. If it is not, she may be able to access it. If she was able to work

from home using a VPN connection, she could still use this VPN connection after she leaves the company. If the account is not disabled and another employee discovers the password, the other employee can log in to the account. Auditing will record any malicious action taken by the other employee as if Sally did it.

Most organizations require that someone disables user accounts as soon as possible after employees leave the company. This includes employees who are terminated and employees who resign. After a specified time, such as 90 days, administrators delete the account. This allows managers to verify the account is not needed prior to deleting it.

In addition to disabling accounts for employees who leave the company, many organizations disable accounts when employees leave for an extended period. For example, if an employee must take an extended leave of absence, the organization disables the account.



EXAM TIP Accounts should be disabled when a user leaves an organization. This includes users who leave an organization amicably or are terminated due to poor performance or other issues.

Participating in Physical Security Operations

While the majority of this chapter focused on logical controls, organizations also use physical access controls to protect assets. Physical access controls include items that you can physically touch. The following list identifies many physical access controls:

- Guards
- Mantraps
- Locked doors
- Alarm systems
- Cameras and closed circuit TVs (CCTVs)
- Facilities (controlled with physical security)

Personnel in the organization need to understand and follow appropriate policies related to physical security. As an example, personnel should help prevent tailgating. Tailgating occurs when someone follows someone through a controlled entry point without providing credentials. Imagine that Bob enters the numbers into a cipher lock to open a door. Suzie follows right behind him without entering the numbers into the cipher lock. In this example, Suzie is tailgating. A mantrap would prevent this activity.

[Chapter 9](#) covers security controls in more depth including physical security controls. In the context of physical security operations, it's worth discussing badging and data center assessments.

Badging

Organizations often issue badges to employees to help identify them. These badges typically include information about the employee, such as the employee's name, along with a picture of the employee.

Many times these badges include additional capabilities. For example, the badge can also be a proximity card used to open doors. Similarly, the badge can also be a smart card used for authentication with computers. As mentioned previously, PIV cards and CACs are specialized types of smart cards that include a photo and can be used for identification as well as authentication.

Data Center Assessment

A data center assessment can evaluate the effectiveness of physical security for a data center. Data centers can house hundreds of physical and/or virtual servers, and these servers often hold valuable data that needs to be protected.

Security professionals and/or security consultants look at various elements of physical security. As an example, a security consultant can test physical security by attempting to gain access to secure areas by tailgating or via other methods. If a security consultant is successful, an attacker could also be successful.

Similarly, an assessment can evaluate cameras and CCTV systems. A primary concern is ensuring that the systems meet the needs of written security policies. For example, a security policy may state that both entrances and exits need to be monitored, which requires at least two cameras. [Chapter 8](#) covers different types of security assessments.

Chapter Review

Access control systems use identification and authentication. Identification occurs when a subject professes an identity, and authentication occurs when the user proves the identity. Administrators authorize authenticated users access to resources by granting permissions to the resource.

The three factors of authentication are something you know (such as passwords), something you have (such as smart cards), and something you are (using biometrics). Most organizations use written password policies to define requirements. Technical controls implement and enforce password policies, and auditing tools verify compliance with the policies.

Passwords are the weakest authentication mechanism. Static passwords stay the same over a period of time. Dynamic one-time passwords use a different password for each user session and

provide much more security than static passwords. Cognitive passwords are based on information that a user knows but an attacker is unlikely to know.

Smart cards, hardware tokens, and proximity cards are examples of something you have. It's common to combine these methods with an additional factor of authentication, providing multifactor authentication. For example, users often enter a PIN along with using a smart card.

Hardware tokens use one-time passwords. A synchronous one-time password is time-based and requires the token and a server to be synchronized. An asynchronous one-time password doesn't require time synchronization but instead uses another method to create the password. Software tokens use protocols such as Time-based One-Time Password (TOTP) and HMAC-based One-Time Password (HOTP). TOTP creates synchronous time-based one-time passwords. HOTP, One-time Password In Everything (OPIE), and S/KEY create asynchronous one-time passwords.

Many organizations implement two-step authentication using HOTP and/or TOTP. After a user attempts to log on, the authenticating system sends a numeric code to the user and requires the user to enter this code to complete the log on process.

Biometrics methods measure physical characteristics of individuals, such as their fingerprints, palms, retinas, and irises. These characteristics are in the something you are factor of authentication. When evaluating the effectiveness of biometrics, you should consider the type 1 and type 2 errors. The False Rejection Rate (FRR, or type 1 error) refers to the percentage of times a system falsely rejects a known user. The False Acceptance Rate (FAR, or type 2 error) refers to the percentage of times a biometric system falsely identifies an unknown user as a known user. The Crossover Error Rate (CER) indicates the point where the FAR and FRR are equal. Lower CERs indicate a better biometric system.

Single sign-on (SSO) allows a user to authenticate once for a system. The user can then access any resources in the system (as long as the user has access) without authenticating again. Kerberos is a popular SSO protocol used within networks. A Kerberos server issues symmetric encryption keys as a Key Distribution Center (KDC) and issues tickets as a ticket-granting server. Kerberos provides authentication on a network and contributes to the confidentiality and integrity of information. Many SSO systems use federated access, providing centralized authentication for different systems. Security Assertion Markup Language (SAML) uses an XML-based data format to share SSO information between servers within a federated identity management system. While SAML is primarily used for authentication, it can also exchange authorization information between different entities.

Two newer SSO technologies are OAuth 2.0 and OpenID Connect. OAuth 2.0 is an authentication framework. OpenID Connect is an authentication layer that is commonly used to extend the capabilities of OAuth. Combined, the two allow websites and online web applications to verify the identity of end users without managing the users' credentials.

Access controls enforce security. For example, a security kernel (a central part of an operating system) enforces security for the operating system by monitoring subjects and objects. Examples of subjects are users, computers, and applications. Examples of objects include data, hardware,

and facilities. Access controls can be logical (implemented with technology such as a security kernel) or physical (such as locked doors).

Several access control models are in use, including discretionary (DAC), non-discretionary (non-DAC), mandatory (MAC), role-based (Role-BAC or RBAC), rule-based, and attribute-based (ABAC) models. The DAC model provides the most granular control and is a subject-based or identity-based model. Individual users own the objects (such as files), and as the data owners, they can assign permissions to subjects (such as other users) as desired. DAC is used with file systems such as NTFS with Microsoft systems and NFS with Linux systems.

The RBAC model uses roles. Subjects (such as users) are placed into roles, and permissions to objects are assigned directly to the role. In non-DAC models, security administrators control access. An ABAC model is a sophisticated type of Rule-based Access Control model. Policies identify attributes of subjects, objects, and the environment to control access.

The MAC model provides the highest level of security when compared to DAC and RBAC models. Both subjects and objects are assigned labels, and when the labels match, access is granted. Several architectures are based on the MAC model. The Bell-LaPadula model has a primary goal of ensuring confidentiality and uses rules of no read up and no write down. The Biba model has a primary goal of ensuring integrity and uses rules of no read down and no write up. The Clark-Wilson model provides integrity by using certification and enforcement rules to enforce integrity and separation of duties. The Brewer-Nash model helps prevent conflicts of interest by preventing access to data organized in conflict-of-interest classes.

Many access control models use an access control matrix and/or a capability table. An access control matrix is object based, with a list of objects along with the permissions granted for each object. A capability table is subject based, with a list of subjects along with the capabilities granted to the subjects.

Identity and access management includes provisioning, maintenance, entitlement, and de-provisioning. Provisioning includes creating accounts and providing appropriate access. Many organizations use roles or groups to manage access to resources. Account management ensures that accounts use appropriate password policies and account lockout policies to help thwart online brute-force password attacks. Entitlement refers to privileges granted to users and helps ensure that the principle of least privilege is followed. Proper de-provisioning practices ensure that accounts are disabled when the employee is no longer working for the company.

Physical security controls include any controls that you can physically touch. Badges are one type of physical security control and typically include a picture of the employee. Badges can also include other capabilities such as a proximity card and/or a smart card. Data center assessments evaluate the physical security methods used to protect data centers.

Questions

1. A user professes an identity by entering a user logon name and then enters a password. What is the purpose of the logon name?

A. Authentication

B. Accountability

C. Identification

D. Accounting

2. Access controls protect assets such as files by preventing unauthorized access. What must occur before a system can implement access controls to restrict access to these types of assets?

A. Identification and authentication

B. Identification and accountability

C. Authentication and accounting

D. Accountability and availability

3. Users are required to enter a different password each time they log on. What type of password is this?

A. Static password

B. Cognitive password

C. Passphrase

D. Dynamic password

4. Authentication includes three types, or factors. Which of the following best describes these authentication methods?

A. Something you say, something you think, and something you are

B. Something you know, something you have, and something you type

C. Something you know, something you say, and something you are

D. Something you know, something you have, and something you are

5. Which of the following choices does NOT ensure that a password is strong?

A. Ensuring that the password is of a sufficient length

B. Ensuring that the password is changed frequently

- C. Ensuring that the password has a mixture of different character types
- D. Ensuring that the password does not include any part of the user's name

6. What can be used to prevent a user from reusing the same password?

- A. Minimum password age
- B. Maximum password age
- C. Password length
- D. Password history

7. A user tries to log on to his bank account via the Internet with his username and password. The webpage then displays a message indicating a code was sent to his smartphone and prompts him to enter a six-digit code. What type of authentication is this?

- A. One-factor authentication
- B. Two-step authentication
- C. Three-factor authentication
- D. TOTP-based authentication

8. What form(s) of authentication are individuals using when they authenticate with a hardware token and a password?

- A. Something they have only
- B. Something they know only
- C. Something they have and something they know
- D. Something they have and something they are

9. An organization uses a biometric system with a one-to-many search method. What does this system provide for the organization?

- A. Authentication
- B. Accountability
- C. Authorization

D. Identification

10. An organization has been using an iris scanner for authentication but has noticed a significant number of errors. Assuming the iris scanner is a high-quality scanner, which of the following is MOST LIKELY affecting its accuracy?

- A. False Acceptance Rate (FAR)
- B. False Rejection Rate (FRR)
- C. Sunlight shining into the scanner
- D. Faulty laser beam

11. Which of the following metrics identifies the number of valid users that a biometric authentication system falsely rejects?

- A. FAR
- B. FRR
- C. CER
- D. AAA

12. What is SSO?

- A. A system that requires user credentials once and uses the same credentials for the entire session
- B. An authentication system that requires users to use different credentials for each resource they access
- C. A secure system used for operations
- D. Any network that employs secure access controls

13. What type of service does Kerberos provide?

- A. Authentication
- B. Accounting
- C. Availability
- D. Accountability

14. Of the following choices, what most accurately identifies the major drawback of SSO systems?

- A.** It allows users to access multiple systems after logging on once.
- B.** It increases the difficulty for users to log on.
- C.** It increases the administrative workload.
- D.** It risks maximum unauthorized access with compromised accounts.

15. What type of access control is subject based?

- A.** Discretionary
- B.** Non-discretionary
- C.** ABAC
- D.** Biba

16. What is the primary goal of the Bell-LaPadula model?

- A.** Enforce separation of duties
- B.** Enforce two-factor authentication
- C.** Enforce confidentiality
- D.** Enforce integrity

17. Which of the following models helps enforce the principle of separation of duties?

- A.** Brewer-Nash and Clark-Wilson
- B.** Brewer-Nash and Biba
- C.** Clark-Wilson and Bell-LaPadula
- D.** Biba and Bell-LaPadula

18. Which of the following statements is true?

- A.** An access control matrix is object based and a capability table is object based.
- B.** An access control matrix is subject based and a capability table is object based.

- C. An access control matrix is object based and a capability table is subject based.
- D. An access control matrix is subject based and a capability table is subject based.

19. Which of the following will disable an account if an attacker tries to guess the password multiple times?

- A. A password policy
- B. An account lockout policy
- C. A password history
- D. De-provisioning accounts

20. Which of the following actions is most appropriate if an employee leaves the company?

- A. Delete the user's account as soon as possible.
- B. Disable the user's account as soon as possible.
- C. Change the user's password as soon as possible.
- D. Change the user's permissions as soon as possible.

Answers

1. C. The logon name provides identification for the user. When combined with the username, the password provides authentication. Accountability is possible if a system can identify users and track their activities. Accounting is provided by logging after proper identification and authentication.

2. A. Identification and authentication must occur before a system can implement access controls. Identification is the act of a user professing an identity, and authentication occurs when an authentication system verifies the user's credentials (such as a username and password). Without proper identification and authentication, it isn't possible to restrict access to specific subjects. Accountability is not possible if an authentication system does not identify and authenticate users. Similarly, you can't provide accurate accounting if users haven't been identified and authenticated.

3. D. A dynamic password is a one-time password that changes for each session. A static password stays the same over a period of time, such as 30 days. A cognitive password includes information that a user would know but isn't necessarily public knowledge. A passphrase is a static password using a long string of characters that has meaning to the user.

4. D. The three factors of authentication are something you know, something you have, and something you are. Something you think, something you type, and something you say are not authentication factors.

5. B. A password should be changed regularly, but doing so doesn't ensure the password is strong. For example, if a user changes a password from "1234" to "4321," it is not strong. The other options all contribute to the strength of a password. Note that this is an example of a question that can be switched to a multiple-correct-answer question. For example, it could be "Which of the following choices ensures that a password is strong? (Choose THREE)."

6. D. Password history remembers users' previous passwords and prevents them from reusing passwords. The minimum password age is used with the password history to prevent users from changing their password repeatedly to get back to the original password. It is often set to one day. The maximum password age identifies when users must change their passwords. The password length identifies the minimum number of characters in the password.

7. B. This is two-step authentication. It is not one-factor authentication because it requires the user to know something (the username and password) and have something (the smartphone) to retrieve the code. It is not three-factor authentication because it is not using biometric authentication (something you are). While two-step authentication might use the Time-based One-Time Password (TOTP) protocol, all two-step authentication methods do not use TOTP. Some use the HMAC-based One-Time Password (HOTP) protocol.

8. C. The two factors of authentication are something they have (the hardware token) and something they know (the password). The third factor of authentication is something you are (using biometrics), but neither a hardware token nor a PIN uses biometrics.

9. D. Biometric systems used for identification use a one-to-many search method. They use the one biometric value (such as a fingerprint) provided by the user and search a database of many similar values (fingerprints) to find a match. Biometric systems used for authentication use a one-to-one search method. They use the biometric value provided by the user, along with the user's claimed identity, and check to see if the value matches the value stored with the user's account. Once a system identifies and authenticates a user, biometric systems are not used for accountability or authorization.

10. C. Lighting affects the accuracy of an iris scanner, so sunlight shining into the scanner's aperture will affect the accuracy. The FAR and FRR indicate specific types of errors. An iris scanner does not use a laser beam.

11. B. The False Rejection Rate (FRR, also called a type 1 error) refers to the percentage of times a biometric system falsely rejects a known user. The False Acceptance Rate (FAR, also called a type 2 error) refers to the percentage of times a biometric system falsely identifies an unknown user. The Crossover Error Rate (CER) identifies where the FAR matches the FRR. AAA refers to the three As of security—authentication, authorization, and accounting.

12. A. Single sign-on (SSO) requires users to log on once, and it uses the same credentials for any other resources accessed during the session. Users are not required to use different credentials for each resource with SSO. SSO is not a network.

13. A. Kerberos provides authentication. Accounting and accountability are possible if a system can identify users and track their activities, but Kerberos doesn't provide logging of user activities. Availability ensures systems are operational when needed. While not listed as one of the answers, Kerberos also contributes to confidentiality and integrity.

14. D. A major concern with SSO systems is that if any single account is compromised, it maximizes the potential unauthorized access. A primary benefit is that SSO allows users to access multiple systems after logging on once. It decreases the difficulty for users to log on and decreases the administrative workload.

15. A. A Discretionary Access Control (DAC) model assigns permissions to identities (subjects), making it a subject-based model. Not all non-discretionary models assign permissions to identities, so B is not the best answer. Attribute-based Access Control (ABAC) focuses on attributes of subjects, objects, and/or the environment. Biba is a MAC-based model that uses labels assigned to both subjects and objects to enforce integrity.

16. C. The Bell-LaPadula model has a primary goal of ensuring confidentiality. The Clark-Wilson and Brewer-Nash access control models help enforce the principle of separation of duties. The Biba model enforces integrity. Access control models do not enforce authentication factors.

17. A. Both the Clark-Wilson model and the Brewer-Nash model enforce the principle of separation of duties. The Clark-Wilson model also enforces integrity, and the Brewer-Nash model also helps prevent conflicts of interest. Biba enforces integrity. Bell-LaPadula enforces confidentiality.

18. C. An access control matrix is object based and a capability table is subject based. An access control matrix is a list of objects along with the permissions granted for each object. A capability table is a list of subjects along with the capabilities granted to the subjects.

19. B. An account lockout policy can disable an account if an attacker (or a user) enters the wrong password too many times. The threshold is often set to three or five, causing an account to be locked out after a user enters the wrong password three or five times, respectively. A password policy ensures that users create strong passwords and regularly change their password. Password history prevents users from reusing the same password. De-provisioning refers to disabling and deleting inactive accounts.

20. B. User accounts should be disabled as soon as possible after the user leaves the company under any circumstances. The account should not be deleted until it's determined to be no longer needed. Changing the password without disabling the account stills allow the account to be used. Disabling the account removes the access and is more direct than changing permissions.

CHAPTER 3

Basic Networking and Communications

In this chapter, you will learn about

- Use cases
- The OSI and TCP/IP models
- Bus, star, tree, token ring, and mesh network configurations
- Common protocols in the TCP/IP suite
- Ports and protocol numbers used with common protocols
- Different network and trust architectures such as the Internet, DMZs, and transitive trusts
- Wireless security protocols such as WPA, WPA2/802.11i, AES, and TKIP
- Wireless technologies such as Bluetooth, NFC, RFID, 802.11, WiMAX, GSM, and 3G

Understanding Use Cases

The 2018 version of the (ISC)² Systems Security Certified Practitioner (SSCP) exam outline added the phrase “Common use cases” for secure protocols. This phrase might be unfamiliar, but don’t let it throw you. You can simply think of this as a goal that an organization wants to achieve.

As an example, an organization may decide that it wants to allow employees to access the internal network from remote locations in a secure manner. A use case to support this goal might be “create a secure virtual private network (VPN).”

Use cases are implemented differently by different professionals. For example, developers often create use cases to help them define the requirements for an application. Marketing and sales people sometimes create use cases to help them define sales goals. Again though, the SSCP exam limits use cases to secure protocols. While reading through this chapter, you’ll see explanations of how various protocols are used. These implementations can also be described as use cases.

Reviewing the OSI Model

The International Organization for Standardization (ISO) developed the Open Systems Interconnection (OSI) Model as a standard model for network communications. It describes a framework for connecting multiple computers to each other and allows dissimilar networks to

communicate with each other. It includes seven layers, and each layer focuses on a set of specific tasks or functions.

Data is packaged, or encapsulated, at each layer with additional data that the layer requires. Generically, the encapsulated data is a *datagram*. However, datagrams have specific names at each layer. For example, at the lowest layer (the Physical layer), data arrives as *bits* (ones and zeros). The next layer (the Data Link layer) adds more data, and packages the data into *frames*. [Figure 3-1](#) shows a diagram of the OSI Model identifying the layers and the data encapsulation names. The following subsections describe these different layers.

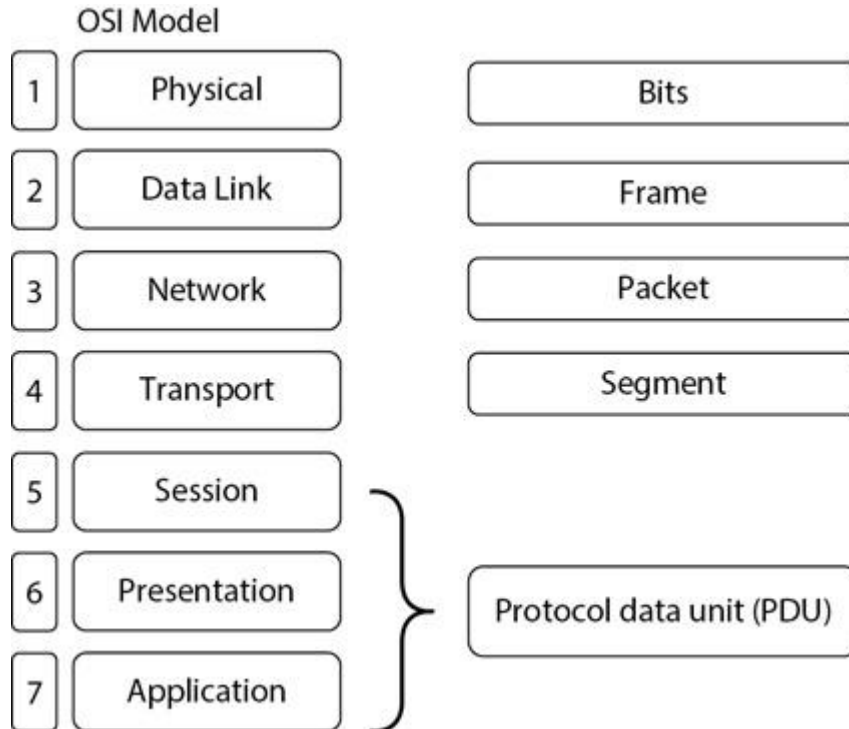


Figure 3-1 OSI Model



TIP Many people memorize the layers of the OSI Model with a mnemonic. For example, All People Seem To Need Data Processing (for Application, Presentation, Session, Transport, Network, Data Link, and Physical) is one. Please Do Not Throw Sausage Pizza Away (for Physical, Data Link, Network, Transport, Session, Presentation, and Application) is another.

For the SSCP exam, you don't need to know the OSI Model as in-depth as you would for a networking certification such as Cisco's CCNA or CompTIA's Network+. However, you should have a basic understanding of the model, including what occurs at each layer. This section introduces each of the OSI Model layers. The "Reviewing Basic Protocols and Ports" section later in this chapter includes a figure that maps protocols to specific layers.

The Physical Layer (Layer 1)

The Physical layer provides a direct connection to the network. It defines how data is passed onto the network, including the transmission method, such as via wired or wireless techniques. The Physical layer includes physical items such as cabling, hubs, repeaters, and wireless radio transmissions required to provide a connection to the network.

Security at this layer focuses on protecting access to the network connections. If attackers are able to connect a protocol analyzer or sniffer via unprotected cabling, they can intercept data going across the network.



TIP A sniffer can capture and analyze packets sent over both wired and wireless networks. For example, if systems send credentials such as a username and password over a network in clear, readable text, the attacker can learn the credentials by capturing and analyzing the traffic. In contrast, attackers are unable to read the data if systems encrypt it prior to transmission, making the text unreadable.

The following list describes the different transmission media types commonly used to transmit data:

- **Coaxial cable** Coaxial cable is similar to the cable used for cable television. It's an older technology that is rarely used on networks today, though it was commonly used in early bus networks (described later in this chapter). It is relatively easy for an attacker to tap into a coaxial cable connection.
- **Twisted pair** This is the most commonly used media for data transmissions. Unshielded twisted pair (UTP) does not have shielding and is susceptible to electromagnetic interference (EMI) and radio frequency interference (RFI). Shielded twisted pair (STP) provides protection against EMI and RFI. The EIA/TIA-568 standards define several categories of twisted pair cables. Common standards are Cat 5e and Cat 6. It is relatively easy for an attacker to tap into a twisted pair connection.
- **Fiber-optic cabling** Fiber provides the highest bandwidth and can support cable runs up to 40 kilometers. Although more expensive than copper cable such as twisted pair, fiber is becoming more and more popular. It is free from EMI and RFI issues and it is harder to tap into than wireless, twisted pair, and coaxial.
- **Wireless transmissions** Wireless transmissions use radio frequency transmissions to broadcast signals over the air. Because they travel over the air, they are the easiest to intercept.



EXAM TIP Cable standards are defined on the Physical layer. Wireless transmissions are the easiest to intercept, and fiber connections are the most difficult to tap into.

Data traveling on the Physical layer is packaged and transmitted as bits (ones and zeros) and bit streams. The Physical layer converts the ones and zeros from the media and passes a bit stream to the Data Link layer, which packages the data as a frame.

The Data Link Layer (Layer 2)

The Data Link layer provides reliable delivery of data across the network and establishes communication links between devices over a network. It includes two sublayers. The Media Access Control (MAC) sublayer interacts with the Physical layer and includes the MAC address (also called the physical address or hardware address). The Logical Link Control sublayer interacts with the Network layer.

Although manufacturers assign a MAC address to network interface cards (NICs), it's worth noting that it's possible to assign a different MAC address to a NIC by using software. This allows an attacker to spoof the source MAC address so that traffic appears as though it's coming from a different computer.

Basic layer 2 switches work at this layer. As traffic comes into a switch on a physical port, the switch identifies the MAC address of the computer sending the data. It then creates a table mapping the MAC address to the physical port. When the switch receives unicast traffic for that computer, the switch sends the data to that port only.

Data traveling on the Data Link layer is packaged and transmitted as frames. The Data Link layer includes orderly delivery of frames and error notification if frames are not delivered properly.



EXAM TIP The Data Link layer is the only layer that includes sublayers. It includes the Media Access Control (MAC) sublayer, which defines MAC addresses (or physical addresses), and the Logical Link Control sublayer. Data is packaged in frames, and the Data Link layer ensures orderly delivery of frames.

The Network Layer (Layer 3)

The Network layer provides routing for data across a network by analyzing the Internet Protocol (IP) address and determining the best route to the destination computer. This layer also controls the flow of data across the network.

Routers and layer 3 switches are the primary hardware devices used at this layer. Routers use access control lists (ACLs) for basic packet filtering to control traffic. An ACL can block or allow traffic based on IP addresses, logical ports, and even some protocols such as Internet Protocol security (IPsec, identified with protocol numbers 50 and 51). Data traveling on the Network layer is packaged and transmitted as packets.

IP is the protocol used for addressing. Currently both IPv4 and IPv6 addresses are in use on the Internet and in many internal networks.



EXAM TIP Routers and layer 3 switches operate on the Network layer and determine the best route based on the destination IP address (also known as the logical address). Both IPv4 and IPv6 addresses are defined on this layer.

The Transport Layer (Layer 4)

The Transport layer provides reliable end-to-end communication services for applications, and includes error detection and recovery mechanisms. Data traveling on the Transport layer is packaged and transmitted as segments. Many of the logical port assignments for upper-layer protocols occur on this layer. Two primary protocols used at this layer are Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).



TIP Well-known ports are logical numbers matched to specific protocols. For example, Hypertext Transfer Protocol (HTTP) uses the well-known port of 80. By blocking port 80, the router's ACL blocks HTTP traffic that is using port 80. The port number is embedded in the datagram.

TCP

TCP is a connection-oriented protocol that provides guaranteed, reliable communication for devices on a network. It uses packet sequencing, and the destination acknowledges every packet that it receives. A three-way handshake establishes a TCP connection.

In [Figure 3-2](#), a client is establishing a TCP session with a server. In the first step of the three-way handshake, the client sends a packet with the synchronize (SYN) flag set. The server responds with both the acknowledge (ACK) and SYN flags set. Last, the client sends back a packet with the ACK flag set. At this point, the session is established.

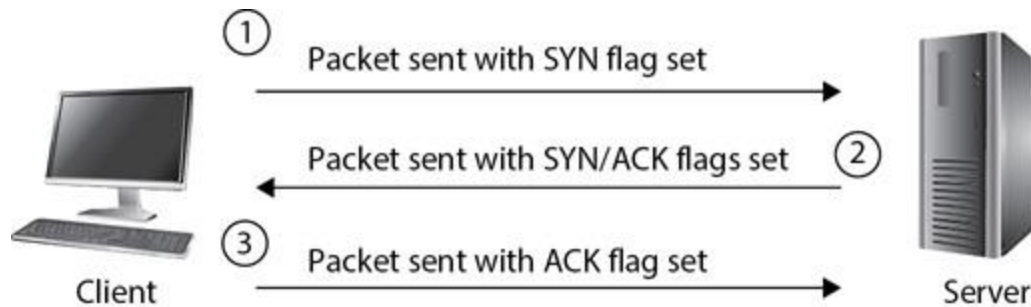


Figure 3-2 TCP three-way handshake



EXAM TIP In a common denial of service (DoS) attack known as a SYN flood attack, the client withholds the third packet. The server is left with a half-open session consuming system resources while it waits for the third packet. However, instead of sending the third packet, the attacker floods the system with more SYN packets. [Chapter 5](#) covers DoS attacks in more depth.

If another protocol requires guaranteed delivery, it uses TCP. For example, IP doesn't include the ability to ensure guaranteed delivery. Instead, IP uses TCP to provide the guaranteed delivery. If the protocol does not require guaranteed delivery, it uses UDP.

UDP

UDP is a connectionless protocol. Instead of checking to see whether a connection exists with another system before sending data, UDP simply sends it. Another way of thinking of this is that UDP gives a "best effort" to send the data, but it does not verify delivery. This could be like you sending a text or an e-mail to a friend to pass on some information. There's a good chance your friend will receive the message, but you don't have any automatic verification.

Many diagnostic protocols use UDP because they don't need to establish a session when performing diagnostic tests. You'll also see UDP used in simpler protocols such as Trivial File Transfer Protocol (TFTP), a simple, streamlined version of File Transfer Protocol (FTP). TFTP sends small files such as configuration files to network devices. FTP uses TCP, but TFTP uses UDP. UDP is also common with streaming data such as streaming audio or video. With streaming data, it's acceptable to lose some bits occasionally here and there. For example, a streaming video file may hiccup occasionally when some bits are lost, but the overall file still plays.



TIP The terms *connection-oriented* and *connectionless* have nothing to do with a wired or wireless connection. Instead, they refer to the logical nature of a connection. TCP ensures both sides can communicate before starting a session. UDP simply sends the data without checking to see whether a logical connection exists.

The Session Layer (Layer 5)

The Session layer establishes and maintains sessions between applications or software processes on the local and remote systems. It keeps the session information for one application separate from the session information from other applications. Remote Procedure Call (RPC) is one of the few protocols that operates directly on this layer. Applications can use RPC to request a service from an application running on a different computer.

Data traveling on the application, presentation, and session layers is packaged and named protocol data units (PDUs).

The Presentation Layer (Layer 6)

The Presentation layer standardizes data presentation for the application layer. It translates data using standards such as American Standard Code for Information Interchange (ASCII), Extended Binary Coded Decimal Interchange Code (EBCDIC), Joint Photographic Experts Group (JPEG), Moving Picture Experts Group (MPEG), and so on. Some data encryption and decryption occurs at this layer. Additionally, data can be compressed and decompressed at this layer for better network performance.

The Application Layer (Layer 7)

Network applications use the Application layer to implement specific user applications. This layer provides the services needed for applications that communicate across a network.

The Application layer supports many additional security features. These include authentication, access control with mechanisms such as permissions, encryption for confidentiality, hashing for integrity, and the use of digital signatures for nonrepudiation.



TIP [Chapter 14](#) covers many cryptographic topics, including encryption, hashing, and digital signatures.

Comparing the OSI and TCP/IP Models

The Defense Advanced Research Projects Agency (DARPA) created the TCP/IP Model. It has a similar goal as the OSI Model—to provide a common framework for connecting multiple computers together. Request for Comments (RFC) 1122 and 1123 identify the TCP/IP Model

with four layers that map to the OSI layers. Some documentation identifies the TCP/IP Model with five layers, but the RFCs showing four layers are authoritative. [Figure 3-3](#) compares the OSI Model to both four- and five-layer TCP/IP Models.

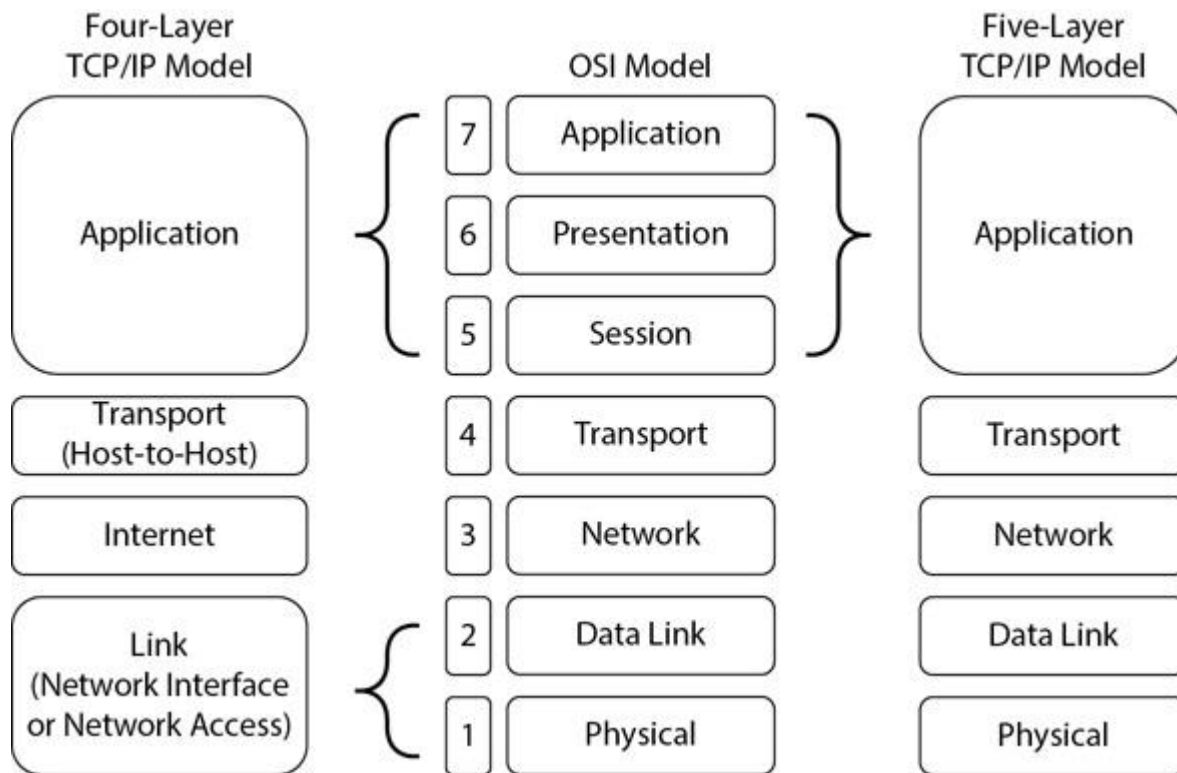


Figure 3-3 TCP/IP Models compared to the OSI Model

Looking at the four-layer TCP/IP Model, you can see that the upper-three layers of the OSI Model map directly to the Application layer of the TCP/IP Model. The OSI Transport layer maps to the TCP/IP Transport layer (also known as the Host or Host-to-Host layer). The OSI Network layer maps to the TCP/IP Internet layer. Finally, the OSI Data Link and Physical layers map to the TCP/IP Link layer (also known as the Network Interface or Network Access layer).

Documentation that shows five layers typically names the bottom-four layers the same as they are on the OSI Model. In essence, they perform the same tasks.



EXAM TIP Although RFCs 1122 and 1123 identify the TCP/IP Model with four layers, some references identify it with five layers. The TCP/IP Model is also called the TCP Model and the DoD Model. DoD is an acronym for the Department of Defense.

Understanding Network Topologies and Relationships

The network topology refers to the layout of a network. It includes items you can touch such as cables and network devices. The primary configuration you'll see in most networks today is star, but you may also see bus, tree, token ring, and mesh configurations. In addition to understanding the topologies, it's important to understand how Ethernet fits into the mix.

Ethernet

The IEEE 802.3 family of standards defines Ethernet. It is the most widely used standard for local area networks (LANs). A working group drafted the first standard in 1973, and the Institute of Electrical and Electronics Engineers (IEEE) has steadily released updates since then. For example, Gigabit Ethernet over twisted pair (802.3bp) and 10 Gigabit Ethernet over fiber (802.3ae) define the standards used for sending data over a network with progressively higher bandwidths.



TIP Modifications to the 802.3 standard are identified with lowercase letters such as 802.3a, 802.3b, and so on. While the standard alphabet has only 26 letters, there have been more than 26 modifications to the 802.3 standard. The modification after 802.3z is 802.3aa. Similarly, the modification after 802.3az is 802.3ba.

Topology vs. Topography

Network concepts such as bus, star, and token ring were included in previous versions of the SSCP domain objectives as both topologies and topographies. The current version uses only the term topographies.

Two questions came to mind for me: 1) What's the difference between topology and topography? 2) Does the change imply a different focus on these concepts for a network?

In the context of networking, topology refers to the arrangement of elements in a network. The logical topology identifies how data flows through the network (such as over a single-line bus or in a ring). The physical topology identifies the actual placement of devices in the network.

Topography usually refers to maps that show physical characteristics of a land mass, such as the elevations of mountains or the physical features of an island. It can also refer to a detailed description of map features, but you'll find it very difficult to find any documentation referring to a network topography.

I've concluded that some people are just merging these terms. While some people might want to debate the correctness of the terms, the good news is that when you're taking the SSCP exam, it

doesn't matter. Exam questions might use the term topology or topography, but what you really need to know is the difference between concepts such as bus, star, tree, token ring, and mesh.

When discussing Ethernet, it's important to understand Carrier Sense Multiple Access (CSMA) and collisions. A collision occurs when two devices try to communicate at the same time over the same medium. Because both devices are sending traffic, the data from both systems collides with each other and none of the data is readable. There are two ways of dealing with collisions if the systems are using the same wire for sending and receiving: detecting them or avoiding them.

CSMA/CD

Ethernet networks support Carrier Sense Multiple Access/Collision Detection (CSMA/CD). In CSMA/CD, each of the devices on the network can detect collisions. As soon as a device detects a collision, it sends a signal on the wire letting all other systems know that a collision occurred. Each device then waits a random amount of time before resending.

However, most switches and computers today use full-duplex connections. In other words, the cable connecting a switch and a computer includes two sets of wires, with one set of wires used for receiving and another set of wires used for transmitting. If full duplex is used, and each device is connected directly with only one other device, collisions don't occur on this cable.

Unfortunately, some older devices still use half duplex. In a half-duplex connection, a single set of wires is used for both transmitting and receiving. It's possible for one system to transmit at the same time as another device is transmitting, resulting in a collision. Ethernet supports CSMA/CD for backward compatibility.

CSMA/CA

In Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA), systems listen before transmitting data. If they sense that other systems are transmitting data, they wait a random amount of time to check again before sending traffic. Wireless networks such as 802.11 networks use CSMA/CA.

Wireless networks also support Request To Send (RTS) and Clear To Send (CTS) to negotiate traffic. When a computer wants to transmit data, it sends out the RTS signal. The destination replies with a CTS signal and waits for the transmission.

Bus Topology

In a bus configuration, all the hosts are connected together via a single shared cable connection as shown in [Figure 3-4](#). Earlier networks commonly used the bus configuration because it was easy to set up. However, it is very difficult to troubleshoot, so you won't see it used to connect entire networks together today.

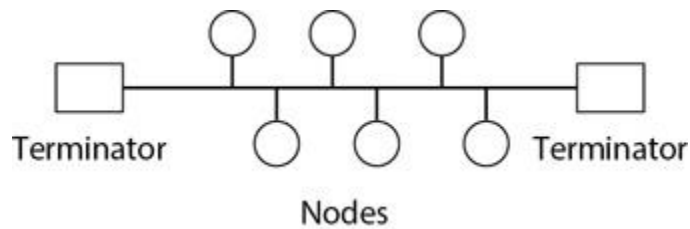


Figure 3-4 Bus configuration

The cable used in a bus configuration is typically coaxial cable. This is similar to a cable used for cable television. T connectors connect each computer on the network. The cable connects to one end of the T, the base of the T connects to the computer, and the other end of the T goes to the next computer.

Notice that each end of the bus includes a terminator. When an electrical signal travels down a wire, it will reflect back when it reaches the end unless the wire has a terminator. This reflected signal corrupts other data transmissions on the bus. The terminator is a connector that prevents this signal reflection. Removing the terminator effectively takes the entire bus network down.

Consider how the computers are connected with a T connector to the main bus. If any one of these T connections is disconnected, you no longer have a single bus with two terminators. Instead, you have two buses, and each one only has a single terminator instead of two terminators. In other words, none of the computers will be able to communicate on the network.

If you only have three or four computers on your network, you can probably locate the problem relatively easily. However, if the network includes several hundred computers, it can take quite a long time to discover the source of the problem.

Older Ethernet standards for bus configurations include 10Base2 (ThinNet) and 10Base5 (ThickNet). These use thin and thick coaxial cables, respectively, for bus configurations.



EXAM TIP Bus networks are expensive to maintain because of how difficult it is to troubleshoot problems. A single break in the cable takes down the entire network, reducing availability.

Star Topology

In a star configuration, all the devices are connected to a central device such as a hub or switch, as shown in [Figure 3-5](#). When drawn this way, it is reminiscent of a star, which is how it gets its name.

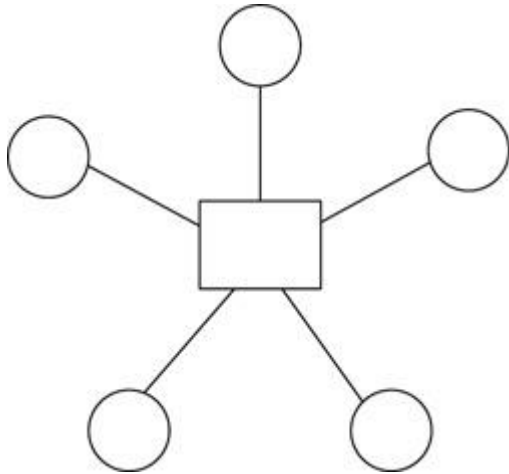


Figure 3-5 Star configuration



NOTE Many network diagrams are drawn like a bus configuration diagram but are actually configured in a star. In other words, network diagrams don't indicate whether the network is configured as a bus, star, or token ring. Instead, they show the relevant components, such as the routers and networks.

A star configuration is much more common in typical computer networks today. Additionally, the devices typically connect with UTP or STP cable. UTP and STP cables use RJ-45 connectors. Common twisted pair cable categories are Cat 5e, 6, and 6a. Cat 5e cable supports speeds up to 1 gigabit and Cat 6 and 6a cables support speeds up to 10 gigabit. The maximum length of these cables is limited to 100 meters (about 330 feet).

A typical star configuration uses a switch as the central component, and computers connect to the switch as nodes in the star with UTP or STP cable. Each of the computers communicates with other computers on the network via the switch. Network devices include components to terminate connections automatically, so external terminators are not required in a star configuration as they are in a bus configuration.

Networks configured in a star are much easier to troubleshoot than a bus network. If someone disconnects a cable from a single computer, only that computer is affected. All the other computers still work. However, compared to the bus configuration, you need much more cable. Each device uses a separate cable from it to the network device, which is typically stored in a secure location such as a locked wiring closet or a server room.

The central device in a star configuration could be a hub, but most organizations have replaced hubs with switches. When a hub is used, all traffic that goes in one port goes out all the other ports of the hub. An attacker could install a protocol analyzer (sniffer) on any computer

connected to the hub and capture all the data going through the hub. However, when organizations replace hubs with switches, the only unicast traffic that will reach the attacker's computer is traffic that is directly addressed to or from the attacker's computer. This improves network efficiency by limiting traffic to computers, and it reduces the potential success of sniffing attacks.

It is possible to configure most switches with a mirrored port. This sends a copy of all traffic through the switch to the mirrored port for monitoring. If an attacker has physical access to the switch, the attacker can connect to the monitoring or mirrored port and sniff all the traffic. Because of this, organizations keep switches in secure areas, such as in a locked wiring closet or a locked server room. Securing a switch in a wiring closet or server room provides physical security.

Ethernet standards for star configurations use twisted pair cable such as 10Base-T, 100Base-TX, 1000Base-T, and 10GBase-T supporting bandwidths of 10 Mbps, 100 Mbps, 1 Gbps, and 10 Gbps, respectively.

Tree Topology

A tree topology is a combination of a bus topology and a star topology. Instead of connecting multiple computer nodes in the bus configuration, it connects multiple star networks along a type of a bus network.

Because all of the star networks are connected to each other via the bus network, a break in the bus network disconnects the star networks from each other. However, a break in the bus does not affect communication within any of the individual star networks.

Token Ring Topology

IEEE 802.5 defines token ring networks. In a token ring network, each device is connected together in a ring. Additionally, the ring includes a logical token that controls when computers can communicate. This prevents collisions because computers can send traffic on the network only when they have the token.

For example, consider a group of five people in a room sitting in a circle. Each person can talk, but only while holding a piece of paper identified as a token. Each person takes turns talking, and when they are done, they pass the token to the next person in the ring.

[Figure 3-6](#) shows two diagrams of a token ring. The diagram on the left shows what the original token ring networks looked like. Each computer is connected to two other computers in the ring, and the computers pass the token around the network. If a computer has the token, it can send traffic on the network. However, if any of the computers fails or is disconnected from the network, it breaks the ring. No other computers will be able to communicate.

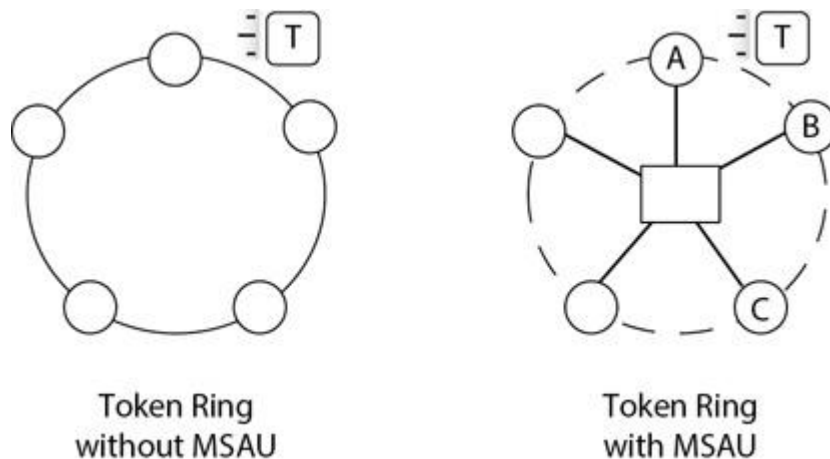


Figure 3-6 Token ring configuration

The configuration on the right includes a multistation access unit (MSAU), also called a media access unit (MAU). The traffic still flows logically in a circle, but it goes through the MSAU. For example, Computers A, B, and C aren't connected directly to each other, but instead they are connected to the MSAU. Traffic passes from Computer A to the MSAU, then to Computer B, back to the MSAU, and then to Computer C. If a computer fails, or a connection between the MSAU and a computer breaks, the MSAU senses the failure and bypasses the computer. In other words, a single failure does not take down the entire network.

Token ring networks don't scale well. Adding more computers significantly reduces performance. If a token ring has five computers, each computer has the token for about 20 percent of the time. However, if the token ring has 100 computers, each computer has the token for only about 1 percent of the time. Traffic can slow to a crawl.

An extension of token ring networks is the Fiber Distributed Data Interface (FDDI) standard. It uses fiber-optic connections instead of copper connections and thus supports much higher bandwidth. It also uses dual rings. The second ring is used only when the system detects an error in the primary ring. Conventional token ring networks support speeds up to 16 Mbps. However, FDDI token rings support speeds up to 100 Mbps over distances as far away as 120 miles.

Mesh Topology

Mesh topologies provide redundancy with multiple connections, providing the highest availability of any of the topology methods. A full mesh topology connects every computer or node in a network with every other node. A partial mesh topology provides multiple redundant connections but doesn't connect every node with every other node in the network. In either a full mesh or a partial mesh, computers in the network are still able to communicate with each other, even if some connections fail.

[Figure 3-7](#) shows two examples of full mesh topologies. The one on the left has four nodes and six connections. The one on the right has five nodes and ten connections.

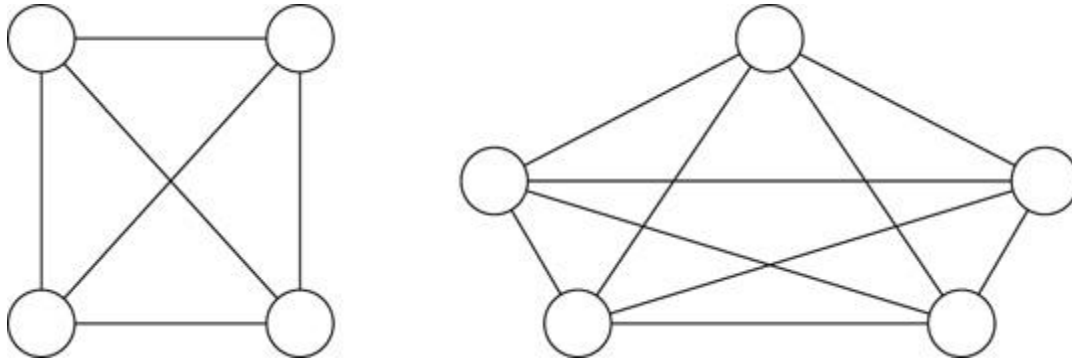


Figure 3-7 Mesh topology

Notice how adding a single node requires the addition of multiple new connections. You can calculate the number of connections needed in a mesh topology with the formula $n(n - 1)/2$, where n is the number of computers. For example, with four computers, $n=4$ and the formula is

- $4(4 - 1)/2$
- $4 \times 3/2$
- $12/2 = 6$

Add another computer and the calculation is $5(5 - 1)/2$, or 10 connections. Six computers require 15 connections. Each of these connections adds costs, so you'll rarely see full mesh topologies in a regular network. However, there are many instances where designers combine mesh topologies with another topology to create a partial mesh. This partial mesh topology has multiple connections to provide a high level of redundancy, but it doesn't connect every single computer with every other computer in the network.

Network Relationships

Network relationships refers to how computers connect with each other. The two common network relationship types are peer-to-peer and client-server.

In a *peer-to-peer network*, systems are independent of each other and authentication is decentralized. As an example, you can have a home network with four computers and a printer connected with both wired and wireless connections. You can log on to your computer, other people in your home can log on to their computers, and everyone can print to the networked printer. In a Microsoft environment, this is referred to as a *workgroup*. Workgroups are typically used for smaller networks.

It's also possible to share folders on individual computers in a peer-to-peer network. However, these are typically accessible to anyone and everyone that has access to the workgroup. If you share a folder on your computer, you can limit access. You'll need to create another account on your computer and limit the access to this new account. Anyone that knows the credentials for

this new account can access the share, but this requires other users to know multiple usernames and passwords.

In a *client-server network*, a server holds accounts and provides centralized authentication. In a Microsoft network, this is called a *domain*. A domain controller holds user accounts, and when users log on, they log on to their computer (and the domain) using their domain account. Within a domain, it's common to create groups to organize users. For example, a group named HelpDesk can be used to organize employees working on the help desk. It's then possible to share folders and grant access to the folders to the HelpDesk group.

Reviewing Basic Protocols and Ports

The TCP/IP protocol suite consists of a full range of protocols used to connect computers on the Internet and on most internal networks. You don't need to know the details of all the protocols, but you should understand the basic function of many of them. The following sections list some common protocols along with their acronyms and a brief explanation of the protocol.

It's also important to know the logical port numbers for many of the protocols. The "Mapping Well-Known Ports to Protocols" section later in this chapter describes the use of these ports in more detail and includes a table mapping some well-known ports to protocols. However, I've also listed the relevant port numbers in the protocol descriptions in case you turn to only that section for a description of the protocol.

A core step in hardening a server, or making it more secure from the default configuration, is to remove all unneeded protocols and services. It's important to know what protocols are required and what protocols are optional when securing any system. Consider a server that is *not* running FTP. If an attacker tries to attack this server using a known FTP exploit, it simply won't work because FTP is not installed on the server.

On the other hand, consider a server that is running 100 different protocols. It is now susceptible to attacks on each of these protocols. By removing unneeded protocols, you reduce the attack surface of the server. Once you understand some basics about protocols, you can determine if a server needs specific protocols or if you can disable or remove them.

[Figure 3-8](#) identifies the layers of the TCP/IP and OSI models where several protocols reside. The following pages include descriptions of these protocols.

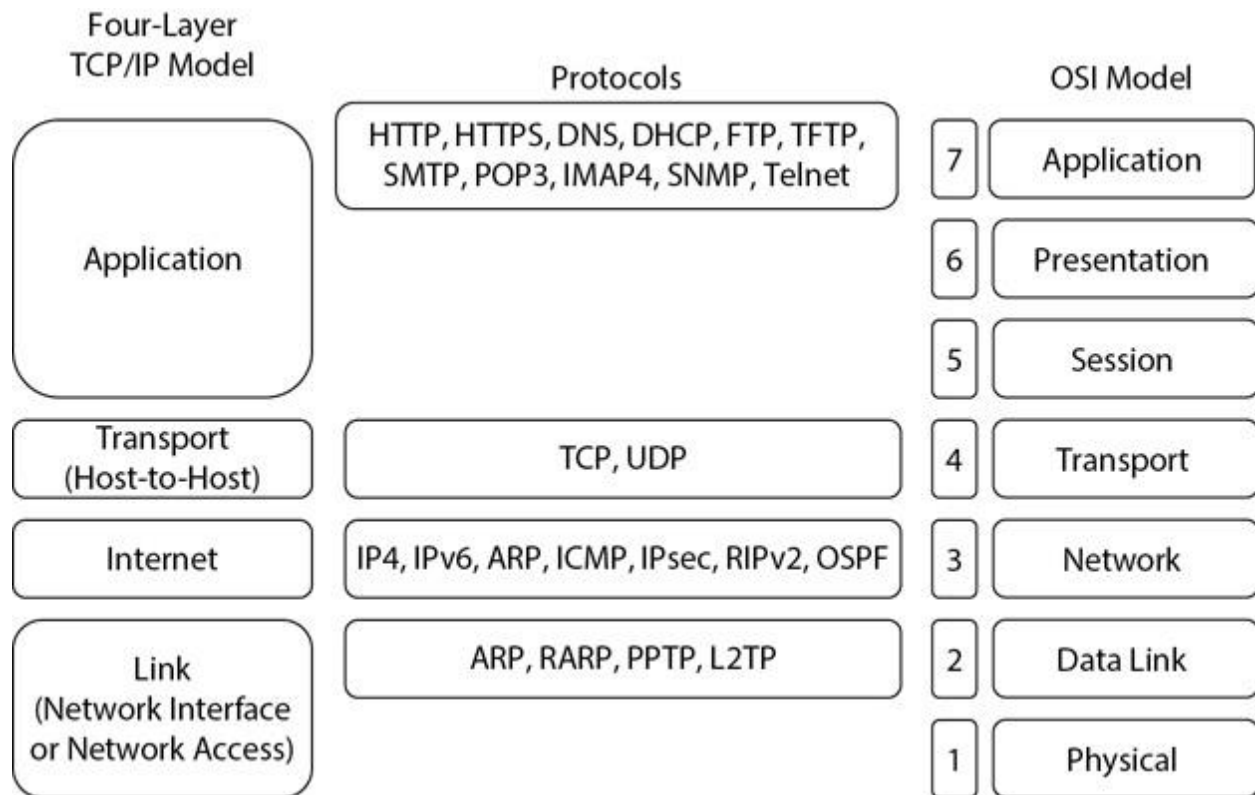


Figure 3-8 Mapping protocols to the OSI and TCP/IP models

Comparing IPv4 and IPv6

IP is the protocol used for addressing. Currently, both IPv4 and IPv6 addresses are in use on the Internet and in many internal networks. An IPv4 address uses 32 bits, but is typically displayed in dotted decimal format (such as 192.168.1.1), where each decimal number represents 8 bits.

An IPv6 address uses 128 bits, and is displayed with hexadecimal characters instead of decimal numbers. When the full IPv6 address is displayed, it includes eight groups of four hexadecimal characters separated by seven colons like this:

fe80:0000:0000:0000:045d:2376:095f:0075



NOTE RFC 5952 states that IPv6 hexadecimal characters such as a, b, c, d, e, and f must be represented in lowercase. This eliminates potential problems when systems search for an IPv6 address.

Two methods used to shorten the displayed address include omitting leading zeros and using zero compression. When omitting leading zeros, the address looks like this:

```
fe80:0:0:0:45d:2376:95f:75
```

Even though this omits leading zeros, it's just for display. The actual IPv6 address still includes 128 bits.

Zero compression replaces a contiguous string of zeros with a double colon like this:

```
fe80::45d:2376:95f:75
```

Even though you don't see seven colons, or eight groups, this is also just for display. The actual IPv6 address still includes 128 bits.

Dynamic Host Configuration Protocol

A Dynamic Host Configuration Protocol (DHCP) server uses DHCP to assign TCP/IP configuration information to DHCP clients. This includes IP addresses, subnet masks, addresses of DNS servers, and much more. Many people shorten the definition of DHCP to say that it simply assigns IP addresses, but that is not completely accurate. There are close to 70 different options that a DHCP server can assign.

DHCP uses four primary messages between the client and the DHCP server:

- **DHCPDiscover** The DHCP client broadcasts this message when it is looking for a DHCP server that can provide the IP address and other configuration information.
- **DHCPOffer** DHCP servers that have IP addresses to offer for the client's subnet respond by broadcasting an offer. The message includes information such as the offered IP address, subnet mask, the IP address of the DHCP server making the offer, and any additional information configured on the DHCP server.
- **DHCPRequest** The DHCP client responds to the first offer it receives with a broadcast request to use the offer. This message includes the IP address of the DHCP server that made the offer.
- **DHCPAck** The DHCP server responds with a DHCP acknowledge message. At this point, the DHCP server re-serves that IP address to the DHCP client. When the DHCP client receives the information, it configures itself with the IP address and other TCP/IP configuration information.

DHCP uses UDP ports 67 and 68. The DHCP client sends traffic out on port 67 and receives answers on port 68. The DHCP server sends traffic on port 68 and receives traffic on port 67.



NOTE DHCP isn't limited to only running on a server. Many wireless routers include the DHCP service, which greatly simplifies the configuration of a wireless network.

Address Resolution Protocol

The Address Resolution Protocol (ARP) resolves IP addresses to MAC (or physical) addresses. The MAC address is 48 bits displayed as 12 hexadecimal numbers (such as 6c-62-6d-ba-73-6c). It is assigned to the NIC and uniquely identifies the computer. Systems use the MAC address when connecting to other computers on the same subnet.

From a big picture perspective, systems use the IP address to get packets from one subnet to another on the Transport layer. Once the packet arrives on the destination subnet, systems use the MAC address to get the frame to the destination computer on the Data Link layer.

ARP broadcasts a datagram that essentially asks all the computers on the subnet, "Who has this IP? If it's you, tell me your MAC address." The computer with that IP address responds with its MAC address.

Reverse Address Resolution Protocol (RARP) works just the opposite of ARP. A client has a MAC address and uses RARP to get an IP address. You can think of this as a simplified version of DHCP. It provides the client with some TCP/IP configuration information, but doesn't have the full capabilities of DHCP.

The Bootstrap Protocol (BootP) is similar to RARP in that it provides clients with IP addresses. However, it goes beyond just getting an IP address and can be used to retrieve a bootable image for diskless clients.



EXAM TIP ARP resolves IP addresses to physical or hardware addresses (MAC addresses). RARP allows a client with a MAC address to get an IP address. BootP allows a diskless client with a MAC address to get an IP address and retrieve a bootable operating system image.

Network Discovery Protocol

IPv6 uses the Network Discovery Protocol (NDP) instead of ARP to resolve IPv6 addresses to the IPv6 version of MAC addresses. IPv4 uses a 12-hexadecimal number (48 bits) for MAC addresses. IPv6 uses a 64-bit modified EUI-64 address instead of the IPv4 MAC address. EUI is an acronym for Extended Unique Identifier.

NDP does more than just resolve IPv6 addresses to EUI-64 addresses. Some of the tasks include

- Automatically configures nodes
- Discovers other nodes on the network
- Locates available routers on the network
- Detects duplicate addresses on the network

Domain Name System

A Domain Name System (DNS) server provides name resolution services. More specifically, DNS resolves host names to IP addresses as its primary function. For example, a DNS client can query the DNS server with the host name of a computer, and the DNS server replies with the host's IP address.

DNS uses a hierarchical naming system with a distributed database to provide responses to clients. Any single DNS server doesn't know the IP address of all the systems on the Internet. However, by querying different servers in the hierarchy, it's possible to discover the IP address of any system with a public IP address.

Thirteen root servers at the top of the hierarchy know the IP addresses for top-level domain DNS servers. Top-level domains include .com, .gov, .edu, .net, and so on. These top-level domain DNS servers know the IP address of second-level DNS servers within the same domain. For example, a .com DNS server knows the IP address of authoritative DNS servers in the mcgraw-hill.com domain. Second-level domain DNS servers know the IP address of third-level DNS servers within the same domain.

The distributed database includes several different types of records to provide different responses. Some of the records in a DNS database are

- **A (host) record** Resolves a host name to an IPv4 address
- **AAAA (host) record** Resolves a host name to an IPv6 address
- **PTR (pointer) record** Resolves an IP address to a host name
- **MX (mail exchange) record** Identifies a mail server
- **CNAME (canonical name or alias) record** Allows a system to be known by different names, or aliases. For example, a web server can be named Web1 but also respond to a name of www.

Berkeley Internet Name Domain (BIND) is a version of DNS software that runs on UNIX systems. It is freely available and runs on many DNS servers on the Internet. Microsoft networks running in internal organizations commonly use Microsoft's version of DNS.

When clients query DNS for name resolution, DNS uses UDP port 53. However, when DNS servers transfer data between each other, they use TCP port 53.

Domain Name System Security Extensions (DNSSEC) adds security to DNS systems. DNSSEC is a suite of extensions that validates DNS responses with a Resource Record Signature (RRSIG), which is similar to a digital signature. The digital signature provides authentication of the response so that DNS clients know that the data hasn't lost integrity. This helps prevent DNS cache poisoning attacks. [Chapter 14](#) provides more information on digital signatures.



EXAM TIP DNS uses a hierarchical structure with a distributed database to provide name resolution services. BIND is commonly used on the Internet, while Microsoft DNS is often used within a Microsoft internal network. DNSSEC helps prevent against DNS cache poisoning attacks.

Internet Control Message Protocol

The Internet Control Message Protocol (ICMP) is used with many diagnostic protocols such as ping, pathping, and tracert to check or verify the health of a network. For example, if you enter **ping** and an IP address (such as **ping 192.168.1.10**) on a Windows system, ping sends four packets to the IP address. If firewalls aren't blocking ICMP, the host with that IP address will reply with four ping packets. Other operating systems, such as Linux and UNIX, will continue to send ping requests until the user presses CTRL-C.



TIP UNIX and Linux systems use traceroute instead of tracert. While there are many similarities between the two commands, traceroute uses UDP instead of ICMP.

Attackers often use ICMP in DoS attacks, so firewalls frequently block ICMP traffic. In other words, a server may be operational, but a ping to that server will fail due to a firewall blocking ICMP traffic.

Internet Group Message Protocol

The Internet Group Message Protocol (IGMP) is used for IPv4 multicasting, where data is sent from one computer to multiple computers at the same time. Computers join a multicast group

with a multicast address (in the range of 224.0.0.0 through 239.255.255.255) to receive the multicast transmissions.



TIP IPv6 does multicasting differently than IPv4. IPv6 uses Multicast Listener Discovery (MLD) with ICMPv6 messaging instead of IGMP.

You can compare multicasting used with IGMP to unicast and broadcast messages:

- **Unicast** Systems send traffic from one computer to one other computer.
- **Broadcast** Systems send traffic from one computer to all other computers on the subnet. Switches pass broadcast traffic, but routers do not pass broadcast traffic, so broadcast traffic only goes to computers on the same subnet.
- **Multicast** Systems send traffic from one computer to multiple computers in the network. Routers pass multicast traffic, so multicast transmissions can reach computers on multiple subnets within the network. Many audio and video streaming applications use multicast, and multicast uses UDP.

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) is commonly used to manage network devices (such as routers and layer 3 switches) on a network. SNMP agents on the network devices send traps (errors) and other notifications back to a central server running SNMP management software. Additionally, a server hosting SNMP management software can use SNMP to configure the network devices remotely.

SNMPv1 and SNMPv2 had several security flaws, but SNMPv3 provides significant improvements. For example, client authentication in SNMPv1 was handled with a community string, and the default community was “public” in SNMPv1 and SNMPv2. However, SNMPv3 introduced several security enhancements. SNMPv3 uses encryption to achieve confidentiality, verifies the integrity of the packets, and includes authentication to verify that the data is from a known entity. SNMP agents receive data on UDP 161 and send traps and notifications on UDP port 162. When using Transport Layer Security (TLS) to encrypt data, Secure SNMP uses TCP ports 10161 and 10162.

When SNMP is active on a network, internal routers will pass SNMP traffic. However, it’s common to block all SNMP traffic at the boundary of the private network with the Internet. This prevents attackers from accessing internal devices using SNMP.

File Transfer Protocol

File Transfer Protocol (FTP) is used to transfer files from one system to another. Some organizations host FTP servers that allow individuals to upload and download data. The FTP server provides a simple storage location for people located in different areas of the world. One benefit of FTP is that it supports authentication in the form of a username or password, but administrators can also configure FTP to support anonymous connections. FTP uses TCP ports 20 and 21.

Trivial FTP (TFTP) was mentioned briefly earlier in the chapter in the “UDP” section. TFTP is a stripped-down version of FTP that provides very basic capabilities. It doesn’t support authentication and doesn’t use TCP to establish a session. However, it is very useful when transferring configuration files to and from network devices. TFTP uses UDP port 69.



EXAM TIP FTP supports authentication and is commonly used to transfer large files across a network. FTP uses TCP ports 20 and 21. TFTP does not support authentication and is commonly used to transfer small configuration files to and from network devices. TFTP uses UDP port 69. Both transfer data across a network in cleartext.

Both FTP and TFTP have significant security issues. Data sent across the network using FTP and TFTP is sent in cleartext. Anyone with a sniffer can easily capture the data and read the contents of the FTP packets. This includes the username and password used to log on to an FTP server, and the data.

Secure FTP (SFTP) uses Secure Shell (SSH) to encrypt FTP traffic. SFTP functions similarly to FTP, but the encryption prevents sniffing attacks from capturing data in cleartext. FTP Secure (FTPS) is another protocol that encrypts FTP sessions. It uses Transport Layer Security (TLS) for encryption.

Telnet

Telnet is a legacy protocol used to interact with a remote system using text-based commands. For example, a server with Telnet enabled will accept connections from a system running a Telnet client. When the client connects, the user can enter commands from a Telnet command prompt. Telnet uses TCP port 23.

Although Telnet used to be much more popular, especially in UNIX and UNIX derivatives such as Linux, you will rarely see it in use today. Telnet sends data (including login credentials) across the network in cleartext, which attackers can capture and read using a sniffer. Once attackers have the credentials, they can log on to the Telnet server and launch additional attacks.

Secure Shell

Secure Shell (SSH) is an encryption protocol that can create a secure session between two computers. It is used to encrypt other protocol traffic, such as FTP (called Secure FTP or SFTP) and Secure Copy (SCP), and can replace many protocols such as Telnet, rlogin (remote login), rsh (remote shell), and rexec (remote execute). SSH also supports secure port forwarding. SSH uses TCP port 22.

Many Linux and UNIX systems support rlogin, a utility that allows users to log in to a remote computer. Similarly, many Linux and UNIX systems support rsh and rexec, which allow users to execute commands on remote computers. However, these commands send data across a network in cleartext, representing significant risks.

SSH provides several security services. It uses asymmetric and symmetric encryption to provide confidentiality, and hashing techniques to provide integrity. SSH also supports mutual authentication, where both the client and the server prove their identity within an SSH authenticated tunnel. RFCs 4251, 4252, 4253, and 4254 document the SSH protocol.



EXAM TIP SSH can encrypt other protocols such as FTP (as SFTP). It is a more secure alternative than older protocols such as Telnet, rlogin, rsh, and rexec that send data across a network in cleartext. SSH uses a combination of asymmetric and symmetric encryption for confidentiality, hashing for integrity, and an SSH authenticated tunnel for authentication. SSH uses TCP port 22.

Hypertext Transfer Protocol and Hypertext Transfer Protocol Secure

Hypertext Transfer Protocol (HTTP) is the common protocol used to transfer web pages over the Internet. Most web browsers use HTTP by default. For example, you could type in the address of a site such as www.mcgraw-hill.com and it will default to <http://www.mcgraw-hill.com>. HTTP uses TCP port 80.



NOTE Some web browsers hide “http://” in the address bar so you may not see it. However, if you copied the URL and pasted it into a document, you would see that it is including http:// even if you don’t type it in the address bar.

HTTP Secure (HTTPS) is an encrypted form of HTTP and is commonly used when processing e-commerce transactions over the Internet. Transport Layer Security (TLS) encrypts the HTTPS sessions in almost all Internet applications. Secure Sockets Layer (SSL) encrypted HTTPS

originally, but due to known vulnerabilities, it has been deprecated and is not recommended for use. Data encrypted within HTTPS is not readable by anyone who is not involved in the HTTPS session. HTTPS uses TCP port 443.

Transport Layer Security and Secure Sockets Layer

Transport Layer Security (TLS) is a popular encryption protocol used to encrypt different types of traffic. You most often see TLS used to encrypt HTTP traffic (as HTTPS). However, other implementations also use TLS. For example, some VPNs use TLS to encrypt tunnel traffic.

Transport Layer Security (TLS) is the designated replacement for Secure Sockets Layer (SSL), and is based on SSL. Although the acronym HTTPS historically has meant that SSL is encrypting an HTTP session, current HTTPS sessions use TLS. [Chapter 14](#) describes the process used by TLS when encrypting HTTPS traffic.



TIP SSL is vulnerable to POODLE (Padding Oracle On Downgraded Legacy Encryption) attacks and no longer recommended for use. Personnel from Google documented this attack in late 2014, which prompted organizations still using SSL to switch to TLS.

IT personnel often debate the actual OSI layer where these protocols operate. RFCs typically provide a definitive answer, but both TLS and SSL are layered protocols, meaning that they operate on different layers. RFC 6101 is a historic RFC that describes SSL. RFCs 2246, 4346, 5246, and 6176 formally define TLS. The following list describes operations performed by TLS on different layers:

- **Transport layer** TLS creates the initial connection on the Transport layer (layer 4). With the name Transport Layer Security, it's implied that TLS operates on the Transport layer.
- **Session layer** The Handshake Protocol within TLS operates on the Session layer (layer 5) when negotiating details for the TLS session. This includes details such as what specific encryption algorithm to use.
- **Application layer** The TLS Record Protocol within TLS operates on the Application layer (layer 7). It provides confidentiality with an encryption algorithm and integrity with a hashing algorithm.

Network File System

Network File System (NFS) allows computers running different operating systems to access and share files over the network. NFS is an open standard created by Sun Microsystems. It allows users running UNIX (and UNIX derivatives such as Linux) to access files on Microsoft systems. Similarly, NFS allows users running Microsoft systems to access files on UNIX-based systems.

Routing Protocols

As mentioned previously, routing occurs at the Network layer. Routers identify the best path for a packet to take based on information it knows about the network. It's possible to enter routing information manually in a smaller network, but for larger networks, the routers communicate with each other using a routing protocol. The two primary routing protocols used on internal networks are Routing Information Protocol version 2 (RIPv2) and Open Shortest Path First (OSPF).

Although RIPv2 works well, it can generate a lot of traffic and doesn't scale well. In general, RIPv2 is good for smaller networks, while OSPF works better in larger networks. RIPv2 is very easy to enable and configure on routers, while OSPF takes a little more effort. Both RIPv2 and OSPF operate on the Network layer of the OSI Model.

E-mail Protocols

You should know about three primary e-mail protocols when preparing for the SSCP exam: Simple Mail Transfer Protocol (SMTP), Post Office Protocol version 3 (POP3), and Internet Message Access Protocol version 4 (IMAP4).

Clients use SMTP to send e-mail to an e-mail server for delivery. Some clients receive e-mail from a POP3 server, and other clients access e-mail from an IMAP4 server. POP3 servers simply send the e-mail to the client when the client connects. Some POP3 servers delete the e-mail as soon as the client receives it. Other POP3 servers keep the e-mail for a preset time limit, or until the user deletes it off the server.

IMAP4 servers store e-mail on the server and allow users to organize the e-mail in folders. This allows users to access e-mail on multiple devices. As an example, e-mail downloaded to your smartphone can also be downloaded to your tablet device. Additionally, if you move an e-mail to a folder on one device, it will be moved to the same folder on all other devices configured to sync with the IMAP4 server.

Another advantage of IMAP4 over POP3 is that users can configure it to show only the headers. This allows users to screen some of the e-mail before downloading it to their system. In contrast, POP3 sends the entire contents to the client as soon as the client connects.

SMTP uses TCP port 25, POP3 uses TCP port 110, and IMAP4 uses TCP port 143. Each of these protocols supports encryption with SSL or TLS. SMTP with SSL/TLS uses TCP port 465, IMAP4 with SSL/TLS uses TCP port 993, and POP3 with SSL/TLS uses TCP port 995.

With the overwhelming amount of spam and phishing attacks, IMAP4 is an appealing alternative used by many e-mail servers on the Internet. Users can often tell by the sender e-mail address or the subject line whether they want to open the e-mail or just delete it. For example, I'd love to actually win \$1,500,000 just because I have a Gmail account, but when I see the subject line saying, "Your e-mail ID has won \$1,500,000," I know it's a phishing attempt and I simply delete the message to eliminate the risk.



EXAM TIP Users send e-mail with SMTP over TCP port 25. Users receive e-mail with POP3 over TCP port 110 or IMAP4 using TCP port 143. IMAP4 allows users to manipulate the e-mail on the server, including organizing e-mail in folders.

Internet Protocol Security

Internet Protocol security (IPsec) provides security for IP traffic traveling over a network. It provides mutual authentication of systems, integrity, and confidentiality. [Chapter 4](#) expands on IPsec, but as an introduction, IPsec has two primary components:

- **Authentication Header (AH)** AH provides authentication between the systems and verifies the integrity of the packets, but it does not encrypt the data. The header content is created from a hash of the packet with additional authentication data, and this hash is then encrypted to prevent tampering. The hash provides integrity. Systems identify AH packets with protocol number 51.
- **Encapsulating Security Protocol (ESP)** ESP encrypts the data and includes the same authentication services provided from AH. Data encrypted with ESP cannot be read if captured with a sniffer. Systems identify ESP packets with protocol number 50.



EXAM TIP Note that AH and ESP are identified with protocol numbers (not port numbers). The “Comparing Ports and Protocol Numbers” section later in this chapter identifies the differences. The IP header includes the protocol number field. Other protocol number values that aren’t associated with IPsec are as follows: ICMP is 1, IGMP is 2, TCP is 6, and UDP is 17.

Tunneling Protocols

[Chapter 4](#) explores virtual private networks (VPNs) in greater depth. In general, a VPN provides access to a private network over a public network such as the Internet. Because traffic traverses a public network, potential attackers can access the data using a protocol analyzer or sniffer. A common method of protecting VPN traffic is to encapsulate it in a tunneling protocol and encrypt the encapsulated data.

Two common tunneling protocols are the Point-to-Point Tunneling Protocol (PPTP) and the Layer 2 Tunneling Protocol (L2TP). PPTP provides encryption, while L2TP relies on IPsec for encryption. Both operate on the Data Link layer. PPTP uses TCP port 1723 and L2TP uses UDP port 1701. TLS and IPsec can also be used as tunneling protocols. [Chapter 4](#) explores these protocols in more depth.

Mapping Well-Known Ports to Protocols

There are a total of 65,536 TCP ports and 65,536 UDP ports. They are divided into three sections:

- **Well-known ports (0 to 1023)** These ports are assigned to widely used protocols by the Internet Assigned Numbers Authority (**IANA**).
- **Registered ports (1024 to 49,151)** These ports are also assigned by IANA, although most are not as widely used as the well-known ports.
- **Dynamic (49,152 to 65,535)** These ports are assigned dynamically by internal systems.

As an example of how ports are used, consider [Figure 3-9](#). The client is using a web browser to connect to an organization's internal web server. When the user clicks a link to connect to the web server, several things happen to retrieve a web page. To begin, the client creates a datagram and sends it to the web server.

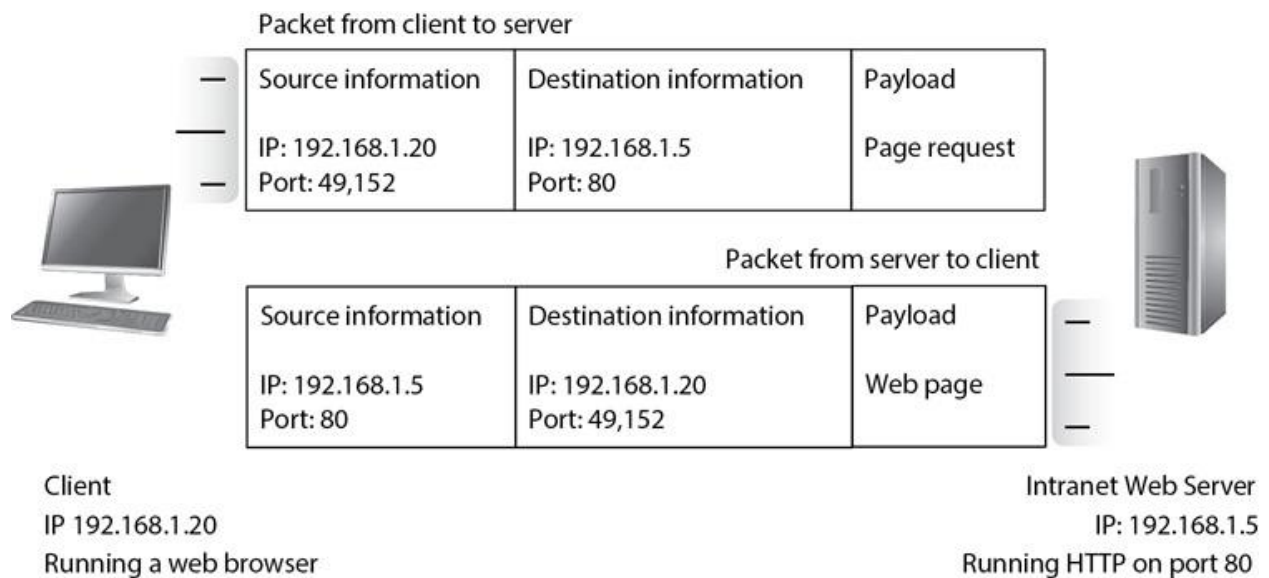


Figure 3-9 Using IP addresses and ports

The datagram includes the source and destination IP addresses and the source and destination ports. A web server has an IP address of 192.168.1.5 and uses TCP port 80 for HTTP, so the destination port is port 80. The client adds its IP address (192.168.1.20) as the source IP address and then identifies an unused port in the dynamic range and assigns it as the source port for this datagram. The user's computer maps the source port to the user's web browser application.

TCP/IP gets the datagram to the server using the destination IP address. Once the datagram reaches the server, the server uses the destination port to get the datagram to the correct protocol, service, or application that is servicing that port. In this example, the destination is port 80,

which is the well-known port for HTTP. The server sends the datagram to the web server application servicing HTTP.

The web server then creates the web page and puts the data together to return the web page to the client. Because the source of the original packet was 192.168.1.20 with a port of 49,152, this now becomes the destination. Again, TCP/IP gets the packet to the client computer using the destination IP address.

When the client receives the packet, it identifies the destination port (49,152), which it previously mapped to the web browser application. The operating system sends the data to the web browser, which then displays it. It's worth noting that some firewalls log the combination of the IP address and the port as a socket. For example, the packet from the client to the server would have a source socket of 192.168.1.20:49152 and a destination socket of 192.168.1.5:80.

[Table 3-1](#) lists many common protocols and ports, along with some basic comments about the protocol.

| Protocol | Port(s) | Comments |
|-------------|------------------|--|
| FTP | TCP 20, 21 | Optional but used on FTP servers. Data is sent in cleartext but can be secured with SSH. |
| TFTP | UDP 69 | Optional and usually disabled. |
| SSH | TCP 22 | Optional. It is used to secure other protocols, such as FTP (SFTP) and Secure Copy (SCP). |
| Telnet | TCP 23 | Optional and usually disabled. Data is sent in cleartext, and Telnet can be used to launch other attacks. SSH is a preferred alternative. |
| DNS | TCP 53, UDP 53 | Required for DNS. Users send DNS queries using UDP 53. DNS servers transfer data using TCP 53. |
| DHCP | UDP 67, 68 | Needed on DHCP servers and for DHCP clients. Clients send DHCP data on UDP 67 and receive DHCP data on UDP 68. |
| SMTP | TCP 25 | Used to send e-mail. E-mail servers and clients with e-mail applications need SMTP enabled. Some web servers also use SMTP. |
| POP3 | TCP 110 | Used to receive e-mail from a POP3 server. |
| IMAP4 | TCP 143 | Used to store, manipulate, and manage e-mail on an e-mail server. |
| HTTP | TCP 80 | Needed on web servers. |
| HTTPS | TCP 443 | Needed on web servers that support secure transmissions. Uses TLS or SSL. |
| TLS | TCP 443 | Used to encrypt many protocols such as HTTPS and FTPS. |
| SNMP | UDP 161, 162 | Used to manage network devices (such as routers and layer 3 switches). Agents receive data on UDP 161 and send traps and notifications on UDP 162. |
| Secure SNMP | TCP 10161, 10162 | Secured with Transport Layer Security. |
| Kerberos | TCP 88 | A network authentication protocol (presented in Chapter 2). |
| PPTP | TCP 1723 | Used with VPNs. |
| L2TP | UDP 1701 | Used with VPNs. |

Table 3-1 Common Protocols and Their Port Numbers

Comparing Ports and Protocol Numbers

While some documentation indicates that ports are the same as protocol numbers, they are completely different. The first octet in TCP and UDP headers includes source and destination ports for TCP and UDP segments, respectively. As a reminder, TCP and UDP operate on the Transport layer of the OSI Model. IPv4 packets include protocol numbers in the Protocol field of

the IPv4 header, and IPv6 packets include protocol numbers in the Next Header field of the IPv6 header. As a reminder, IPv4 and IPv6 operate on the Network layer of the OSI Model.

[Table 3-2](#) lists common protocol numbers.

| Protocol Number | Protocol |
|-----------------|-------------|
| 1 | ICMP (IPv4) |
| 2 | IGMP (IPv4) |
| 6 | TCP |
| 17 | UDP |
| 50 | IPsec ESP |
| 51 | IPsec AH |
| 56 | ICMP (IPv6) |

Table 3-2 Common Protocol Numbers



EXAM TIP Know the common ports and protocol numbers listed in [Tables 3-1](#) and [3-2](#). Routers and firewalls use these to control traffic in and out of a network, so these are important security concepts to know.

Comparing Internetwork Trust Architectures

Computers are connected together using networks, and different types of networks provide different levels of trust. Generically, there are four types of networks: the Internet, an intranet, an extranet, and a demilitarized zone (DMZ, or perimeter network). Each of these can support different types of trust architectures, such as one-way trusts, two-way trusts, and transitive trusts.

Consider [Figure 3-10](#) as you review the definitions of the first three different network types in the following list:

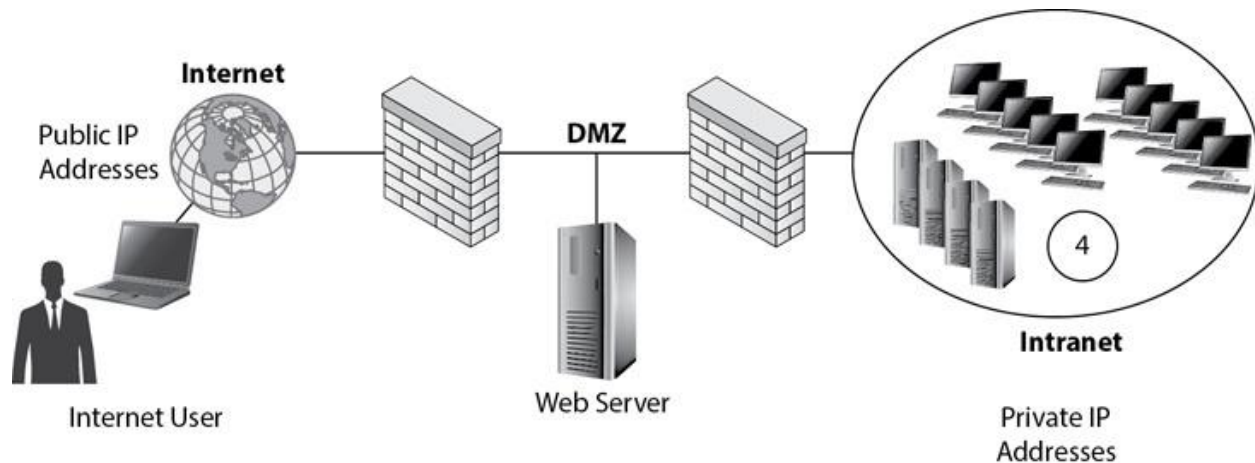


Figure 3-10 Comparing the Internet, an intranet, and a DMZ

- **Internet** The Internet is the public network accessible around the world. Any computer with access to the Internet can reach any other computer with access to the Internet. From a risk perspective, the Internet presents the most risk.
- **Intranet** An organization's internal network is an intranet. From a risk perspective, an intranet presents the least amount of risk. An organization controls the resources on the internal network.
- **DMZ** When an organization wants to host resources on the Internet (such as with a web server, an e-mail server, or an FTP server), a common practice is to place these Internet-facing servers within a DMZ. The DMZ provides a layer of protection for the servers that would not be available if they were placed directly on the Internet. Although there are multiple methods for creating a DMZ, a common method is to use two firewalls, as shown in [Figure 3-10](#). Another method is to create a multihomed firewall, using three or more NICs to create a separate screened subnet.
- **Extranet and third-party connections** An extranet is a network used to host resources via the Internet, but is accessible only to trusted third-party entities, as shown in [Figure 3-11](#). For example, a company may have a partnership with another company and want to share resources with that partner through a web-based portal. Administrators at the first company can create an extranet to provide access to these resources, but restrict access to this network to only the trusted partner. Many organizations use federated access systems to provide single sign-on (SSO) capabilities to systems hosted in an extranet. [Chapter 2](#) describes federated access in more depth.

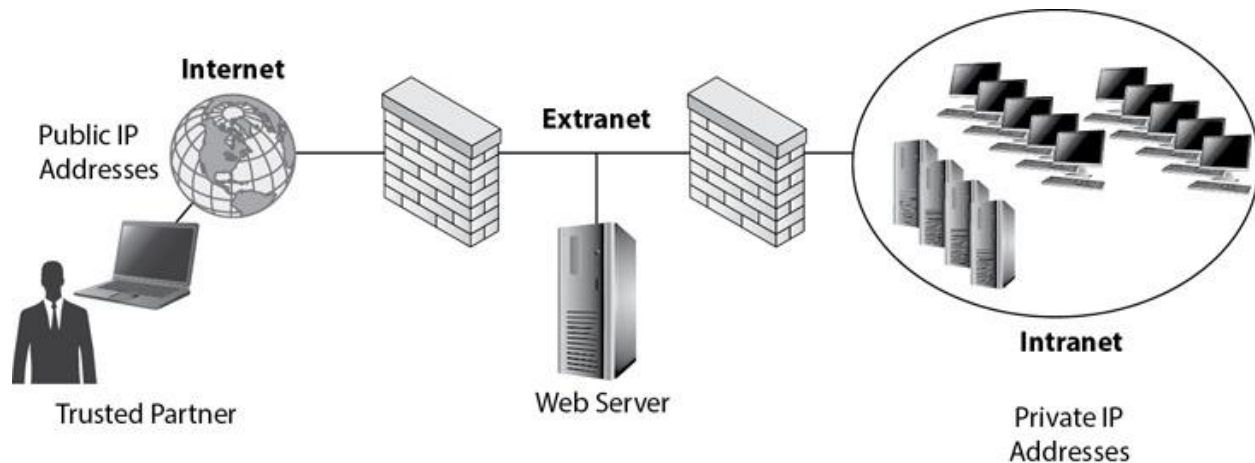


Figure 3-11 Comparing the Internet, an intranet, and an extranet



TIP Notice that the drawing representing an extranet ([Figure 3-11](#)) is very similar to the drawing for a DMZ ([Figure 3-10](#)). The difference is in the intent. Resources in a DMZ are available to any users on the Internet. Administrators configure an extranet so that resources are only available to trusted entities. An organization can host both a DMZ and an extranet.

Additionally, networks are identified with some other terms as follows:

- **Personal area network (PAN)** A small network centered on a person. For example, a person may have a cell phone paired with an earpiece and a microphone within a PAN.
- **Local area network (LAN)** A network that connects computers together in the same geographical location. A LAN is a single, well-connected network. It could be in a single room, a single building, or a group of buildings connected together with links of similar speeds.
- **Metropolitan area network (MAN)** A group of networks that are connected over a large area such as a university campus or even an entire city.
- **Wide area network (WAN)** A large network that connects computers across a large geographical area, such as in different cities. WAN links connect WANs together, and WAN links are typically slower than the links within individual LANs.

Comparing Public and Private IP Addresses

You may have noticed in [Figures 3-10](#) and [3-11](#) that the Internet included public IP addresses, while the intranet included private IP addresses. Public IP addresses are used only on the Internet, and private IP addresses are used within an organization. A public IP address is accessible from any other public IP address. In other words, any computer (or any attacker) with

access to the Internet from anywhere in the world can access another computer on the Internet because they both have public IP addresses.

Every public IP address must be unique. No two computers can be assigned the same public IP address. In contrast, private IP addresses must be unique only within an organization. For example, all the IP addresses used by McGraw-Hill in its network are unique within the McGraw-Hill network. However, (ISC)² can use the same private IP addresses as McGraw-Hill because both companies use these IP addresses internally.

It's common for IP addresses within the DMZ or within an extranet to have public IP addresses. However, it's also possible to configure the firewalls to forward traffic to servers within a DMZ or extranet and configure these servers with private IP addresses.

RFC 1918 lists the IPv4 addresses reserved for private networks. They are as follows:

- **10.0.0.0 through 10.255.255.255** A very large address space with enough addresses for large organizations
- **172.16.0.0 through 172.31.255.255** A medium-sized address space
- **192.168.0.0 through 192.168.255.255** A smaller address space for smaller organizations

Link Local Addresses

IPv4 and IPv6 both have specific IP address ranges that can be assigned to clients automatically. These are called link local addresses. A computer with only a link local address can communicate with other computers on the local link or local subnet, but not with computers on any other subnets. In other words, these addresses cannot pass a router.

The link local address range in IPv4 is 169.254.0.1 through 169.254.255.254, more commonly known as Automatic Private IP Addresses (APIPAs). DHCP clients will automatically assign themselves an address in this range if they do not receive an IP address from a DHCP server.

IPv6 works a little differently. IPv6 clients will typically assign themselves a link local address any time IPv6 is enabled. It doesn't matter if they've received an IPv6 address from a DHCP server, or even if they have a manually assigned IPv6 address. The IPv6 link local address range starts with the fe80::/10 prefix.



EXAM TIP Private IP addresses are not routable on the Internet. If traffic does reach the Internet with a private source or destination IP address, routers on the Internet drop it.

RFC 4193 lists the IPv6 addresses reserved for private networks. They are identified as unique local IPv6 unicast addresses (or local unicast addresses). They have a prefix of fc00::/7, which is a little misleading.

The /7 in fc00::/7 identifies the first seven bits in the address. A hexadecimal f is binary 1111, and a hexadecimal c is binary 1100. If only the first seven characters are used, it is 1111 110x, with x being the eighth bit. However, RFC 4193 specifies that the eighth bit is always a 1 (although this may change in the future), so the first eight characters of a unique local address are 1111 1101 or fd. In other words, it's technically accurate to say that the link local address starts with fc00::/7, but in practice, you'll actually see link local addresses that start with fd (using a hexadecimal d instead of a hexadecimal c).

Using NAT

The Internet has public IPv4 addresses, and internal networks have private IPv4 addresses. However, users with private IPv4 addresses can still access the Internet. Network Address Translation (NAT) provides a mechanism that allows this to happen. NAT is an additional service running on a system at the boundary between the Internet and an internal network.

Consider [Figure 3-12](#). It shows an intranet with private IP addresses as well as the Internet with public IP addresses, separated by a proxy server with NAT installed. Internal clients connect to the Internet via the proxy server, and NAT translates the IP addresses from public to private and from private back to public.

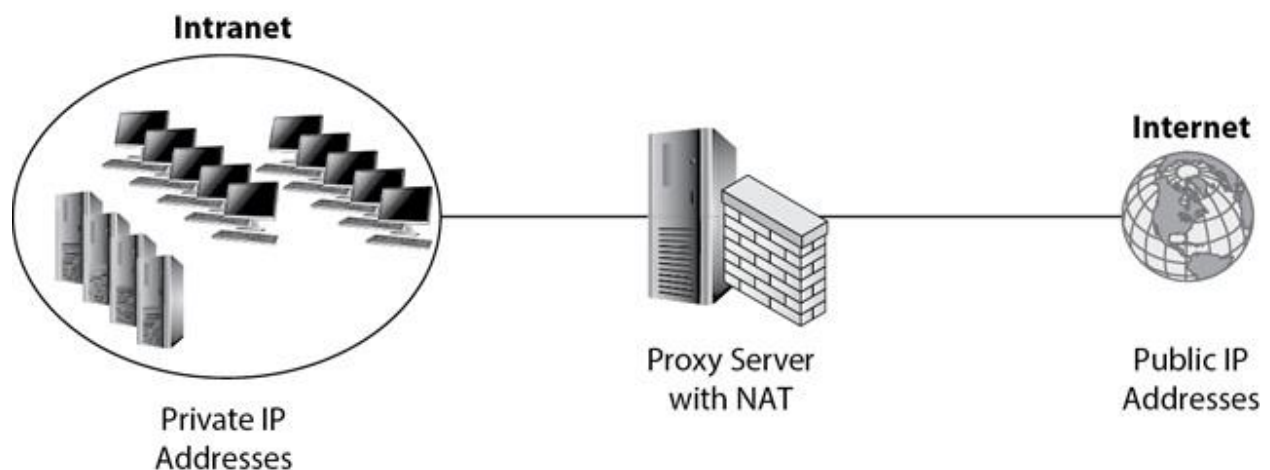


Figure 3-12 Using a proxy server with NAT

One of the benefits of NAT is that internal clients do not need public IP addresses to access the Internet. This provides a layer of protection for these clients because they aren't directly accessible from other Internet clients. It also reduces costs for an organization because the organization needs to purchase public IP addresses only for the proxy server and not for every internal client.

Smaller networks often use NAT on a router instead of a proxy server. For example, wireless routers used in small wireless networks often include NAT. All internal clients use private IP addresses and access the Internet via the wireless router.

NAT is typically identified as either static or dynamic:

- **Static** Maps an unregistered private IP address to a registered public IP address on a one-to-one basis. Static NAT allows an internal client to communicate with devices on the Internet with the same IP address.
- **Dynamic** Uses multiple public IP addresses and maps private IP addresses to one of the public IP addresses based on a first-come, first-served basis.

Port Address Translation (PAT) is an extension of NAT that maps multiple private IP addresses to a single public IP address using different ports. As an example, imagine the proxy server in [Figure 3-12](#) is using PAT instead of NAT and Sally uses her web browser to connect to [mcgraw-hill.com](#) from the private network. The PAT server changes the request using the PAT server's public IP address as the source IP address, and it changes the source port to a unique source port number, such as 49,152. When the PAT server receives the response, it has a destination port of 49,152. The proxy server recognizes this port and modifies the packet using Sally's original private IP address and port number, and forwards it to her computer.



EXAM TIP Static NAT uses a one-to-one mapping with a single IP address, but dynamic NAT uses a one-to-many mapping with multiple IP addresses. PAT is an extension of NAT that uses a many-to-many mapping using port numbers.

Comparing Trust Relationships

Networks utilize trust relationships to simplify access control. The three primary types of trust relationships are one-way, two-way, and transitive. Refer to [Figure 3-13](#) when reviewing the following explanations:

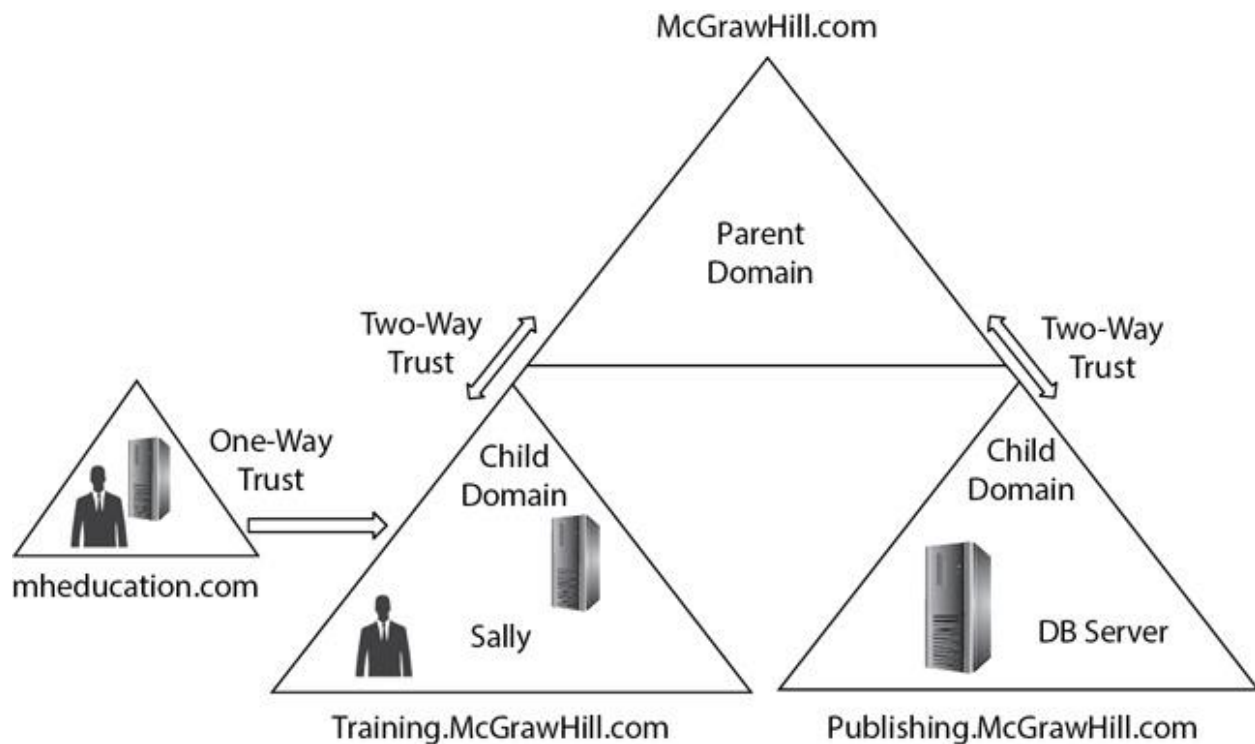


Figure 3-13 Trust relationships

- **One-way trust** A one-way trust allows subjects (such as users) in one domain to access resources in another domain. For example, [Figure 3-13](#) shows a one-way trust relationship where [mheducation.com](#) trusts [Training.McGrawHill.com](#). This would allow users in the [Training.McGrawHill.com](#) domain (such as Sally) to access resources in the [mheducation.com](#) domain. However, users in the [mheducation.com](#) domain would not be able to access resources in the [Training.McGrawHill.com](#) domain.



NOTE A one-way trust does not automatically grant permissions. In other words, everyone in the [Training.McGrawHill.com](#) domain would not automatically have access to all resources in the [mheducation.com](#) domain. However, the one-way trust makes it possible to grant users permissions to the resources.

- **Two-way trust** A two-way trust indicates that both domains trust each other. In [Figure 3-13](#), there is a two-way trust between the parent domain ([McGrawHill.com](#)) and each child domain (Training and Publishing). You can also think of this as two one-way trusts.
- **Transitive trust** A transitive trust is an indirect trust relationship. In [Figure 3-13](#), the Training domain and the Publishing domain do not have a direct trust relationship. However, because they both have a two-way trust relationship with the parent domain ([McGrawHill.com](#)), it creates a

transitive trust relationship between both domains. This allows users in the Training domain to access resources in the Publishing domain, and it allows users in the Publishing domain to access resources in the Training domain. This doesn't mean they automatically have full rights and permissions in the other domain. However, administrators can grant users rights and permissions when the transitive trust relationship exists.

Exploring Wireless Technologies

Wireless technologies are quite popular in both homes and organizations. With a wireless network (sometimes called a WiFi network), computers connect to each other and network devices by transmitting data over the airwaves. There are five primary IEEE 802.11 standards for wireless local area network currently supported by most access points or wireless routers. [Table 3-3](#) outlines the base frequency and theoretical maximum speeds of each.

| Standard | Base Frequency | Maximum Speed |
|----------|----------------|--------------------------------------|
| 802.11a | 5 GHz | 54 Mbps |
| 802.11b | 2.4 GHz | 11 Mbps |
| 802.11g | 2.4 GHz | 54 Mbps |
| 802.11n | 2.4 GHz, 5 GHz | 600 Mbps |
| 802.11ac | 2.4 GHz, 5 GHz | 600 Mbps (2.4 GHz), 2.6 Gbps (5 GHz) |

Table 3-3 802.11 Wireless Standards



NOTE The speed is the maximum speed in ideal conditions. However, many variables affect the actual speed. For example, as the distance increases between two wireless devices, the actual speed decreases. Obstructions such as walls and interference from EMI or RFI also reduce the actual speed.

Other standards have been defined and are being developed, but they are not widely supported as I write this chapter. As an example, 802.11af was approved in February 2014 and supports the use of unused television channels. 802.11ax is slated to be an upgrade to 802.11ac with a goal of providing four times the throughput of 802.11ac.

A wireless network includes a wireless access point (AP) or a wireless router. Wireless devices, such as laptops, tablets, and smartphones, connect to the network through the AP or the wireless router. An AP provides a bridge for wireless clients to a wired network. For example, [Figure 3-14](#) shows an AP for an organization with one user connecting to the AP with a smartphone while another user is connecting with a laptop. Once connected, the wireless clients are able to access any resources in the network, just as other wired clients can.

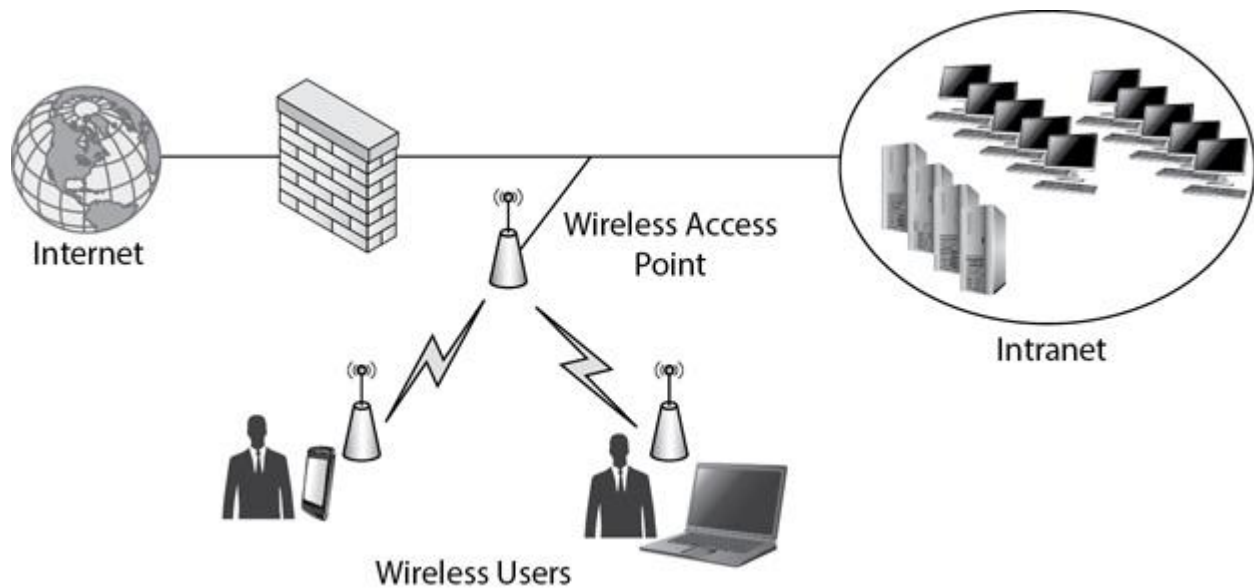


Figure 3-14 Wireless users connecting to a network via an AP

Wireless routers include additional capabilities beyond the AP. Technically, they are access points with routing capabilities, but they are widely marketed as wireless routers. They are used in small offices and home offices (SOHOs), along with many home networks. Some of the common capabilities found in a wireless router include

- **AP** Provides a bridge for wireless clients to the wired network.
- **DHCP** Provides both wired and wireless clients with TCP/IP configuration information, such as IP addresses, subnet masks, the address of a DNS server, and more.
- **NAT** Translates between private IP addresses used internally and public IP addresses used on the Internet. Wireless routers typically use PAT, which is an extension of NAT.
- **Routing** Provides basic routing between the internal network and the Internet. The wireless router is the default gateway for all internal clients.
- **Basic firewall** Provides basic firewall capabilities to filter traffic, but the capabilities can differ substantially from one brand to another.

[Figure 3-15](#) shows an example of a wireless router within a network. This wireless router is providing many additional services for both wired and wireless clients. All clients access the Internet via this wireless router.

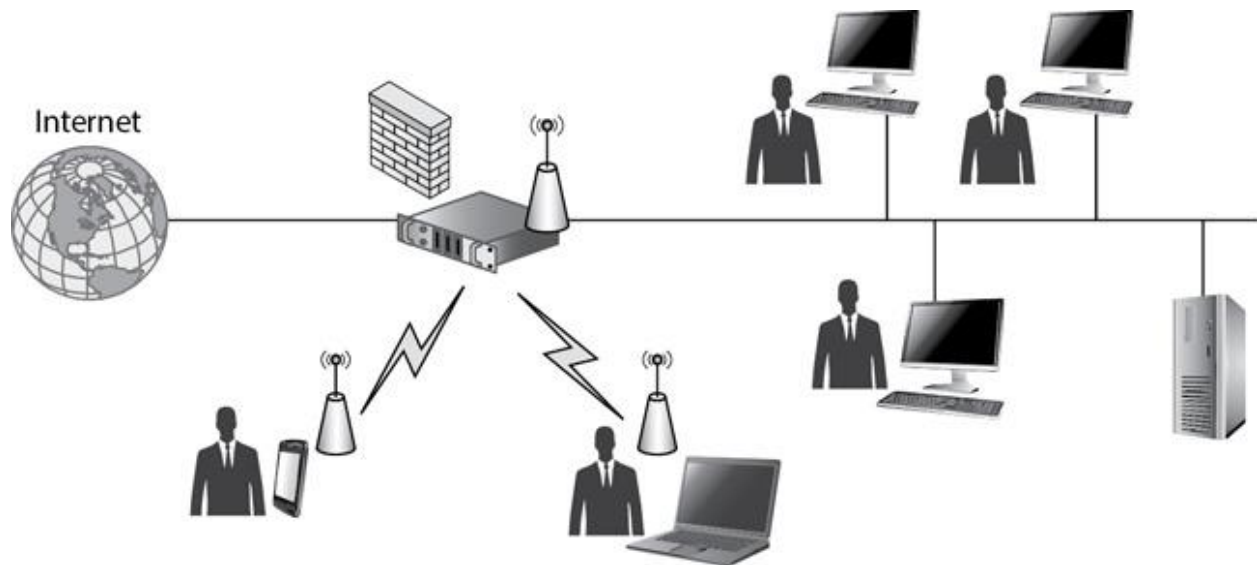


Figure 3-15 Wireless users connecting to a network via a wireless router

Wireless systems typically transmit in multiple directions (omnidirectional) at the same time. The area where wireless systems can receive the signal is the wireless footprint. Within an organization, walls or devices transmitting EMI or RFI reduce the overall footprint.

However, an attacker can extend the footprint of a wireless network by using a receiver with a directional antenna. For example, it's possible to hook up something as simple as a can to a wireless antenna (creating a cantenna) and then point the can toward the wireless transmitter. Some tests have shown that an attacker can capture transmissions of wireless networks more than a mile away by using a cantenna.

Securing Data Transmissions

Because wireless technologies transmit data over the air using radio waves, it's easy for an attacker to capture the data with a wireless receiver. This can be as simple as installing a wireless sniffer on a laptop that includes wireless capabilities. For example, Wireshark is a free sniffer application that includes the ability to capture wireless transmissions. An attacker can download and install it, and then sit in an organization's parking lot and capture wireless transmissions.



TIP Wireshark can capture all transmissions as long as the NIC is in promiscuous mode. If it's not in promiscuous mode, Wireshark only captures traffic addressed to or from the Wireshark client. The challenge with capturing wireless transmissions is that many wireless NICs do not support promiscuous mode. However, some USB-based adapters such as Riverbed AirPcap support promiscuous mode.

Secure wireless transmissions make it more difficult for an attacker to capture any usable information from these transmissions. The three primary security protocols are Wired Equivalent Privacy (WEP), Wi-Fi Protected Access (WPA), and Wi-Fi Protected Access 2 (WPA2).

WEP was the first security algorithm used for 802.11 wireless networks. The goal was to provide the same level of privacy on a wireless network as users could achieve on a wired network. It failed. WEP has several security issues and should not be used.

WPA was the interim replacement for WEP. The goal was to provide better security immediately for wireless devices using existing hardware. In other words, users running WEP could switch over to WPA without purchasing new wireless adapters, wireless access points, and/or wireless routers. Although WPA is much better than WEP, its designers never intended it to be a permanent replacement. WPA has known vulnerabilities and is susceptible to WPA cracking attacks allowing an attacker to discover the WPA preshared key (PSK), which is a password or passphrase of 8 to 63 characters.

WPA2 is the permanent replacement for WEP and WPA. The IEEE formally ratified it in 2004 as 802.11i. WPA2 requires different hardware than WEP and WPA. However, the Wi-Fi Alliance requires that all new hardware support WPA2 and Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP) in order to be Wi-Fi Certified. The Wi-Fi Alliance is a global nonprofit organization that has helped standardize wireless technologies.

WPA originally used only Temporal Key Integrity Protocol (TKIP). TKIP could use the same hardware as WEP, but provided significantly more security than WEP. Later implementations of WPA support Advanced Encryption Standard (AES), a much stronger encryption standard than TKIP.

WPA2 uses CCMP, which is based on AES and is much stronger than WPA using either TKIP or AES. WPA2 with CCMP is the recommended standard for security.

WPA-Personal and WPA2-Personal

WPA-Personal and WPA2-Personal use a preshared key (PSK). WPA/WPA2 then uses this PSK to create an encryption key, which wireless devices use to encrypt and decrypt wireless traffic. It's worth stressing that the PSK isn't the encryption key itself, but instead the PSK is used to create the key.



TIP This encryption key is used by either AES or TKIP for symmetric encryption depending on whether the system is using WPA2 or WPA. [Chapter 14](#) explores symmetric encryption and AES in more depth.

Administrators enter the PSK into the AP or wireless router, and users then enter the same PSK into their wireless devices. For example, [Figure 3-16](#) shows one of the setup pages for a Linksys wireless router. The PSK is IW!lIBe\$\$CPCertified. You can also see that WPA2-Personal is selected, and on this router, WPA2-Personal uses AES-based CCMP.

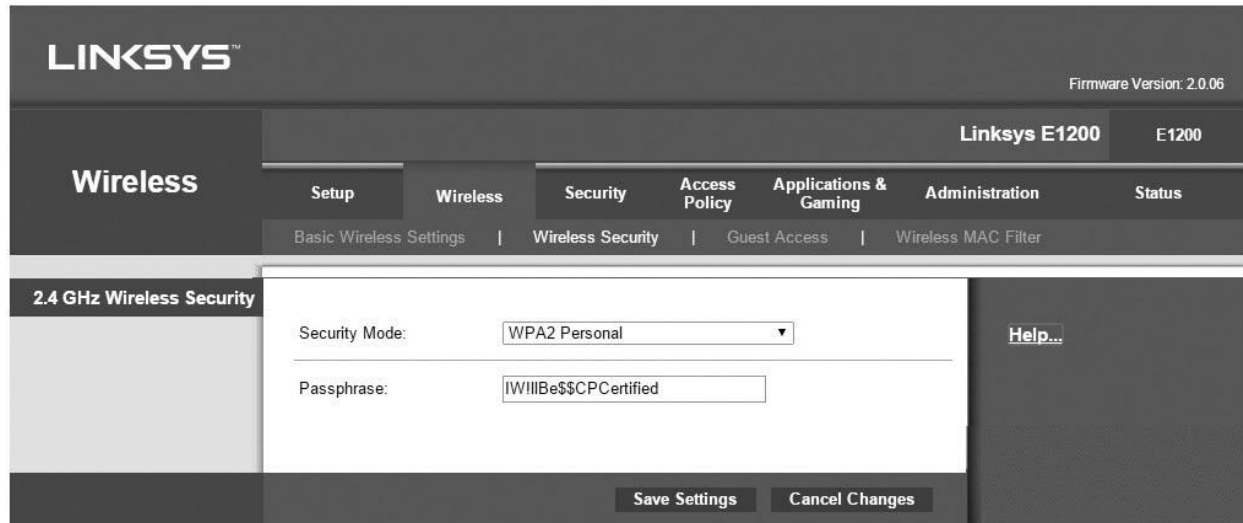


Figure 3-16 Configuring security for a wireless router

WPA-Enterprise and WPA2-Enterprise

An organization can increase security with wireless networks by adding an 802.1x authentication server and using Enterprise mode. A Remote Authentication Dial-In User Service (RADIUS) is commonly implemented as the 802.1x authentication server. Both WPA and WPA2 support Enterprise mode.



EXAM TIP Enterprise mode requires an 802.1x authentication server and increases security by requiring clients to authenticate before granting them access.

When an 802.1x authentication server is added, the clients are required to authenticate (such as with a username and password) prior to being granted access to the network. Additionally, clients don't need to enter a PSK. Instead, the system creates and assigns unique encryption keys after the clients log on. In contrast, Personal mode requires users to know and enter the PSK shared by all users.

An 802.1x authentication server will often be configured to check an authentication database on another server. For example, in a Microsoft domain, the 802.1x authentication server can check a user's credentials via a domain controller hosting Active Directory Domain Services. Users log

on using their domain credentials. If they don't have an Active Directory account, they cannot log on and cannot access the wireless network.

Although the focus in this section is to use an 802.1x authentication server to authenticate wireless clients, an 802.1x authentication server can also be used to authenticate other clients. For example, it can be used to authenticate DHCP clients prior to providing TCP/IP configuration information.



EXAM TIP WEP should not be used because it has several security vulnerabilities. Additionally, WPA has known vulnerabilities and should not be used. WPA2 with CCMP is recommended over WPA2 with TKIP. If legacy hardware doesn't support WPA2, WPA might be the only alternative to WEP. WPA-Personal and WPA2-Personal use a PSK. WPA-Enterprise and WPA2-Enterprise use an 802.1x authentication server, providing a higher level of security.

Wireless Device Administrator Password

Wireless devices such as APs and wireless routers commonly have a web-based administration program that you can access with a basic web browser. You can access this program using the IP address assigned to the router (often either 192.168.0.1 or 192.168.1.1). You enter the IP address as the URL in your browser. When you connect, the wireless router prompts you to provide the name of the administrator account and the administrator password.

You can check the manual (or the Internet) for the default name and password if the device uses defaults. A common default combination used in older APs and wireless routers is **admin** (for the administrator account) and **admin** (as the password), although there are many variations. Most newer devices force you to create the password for the device when you first set it up.

If the device has a default password, an important step in securing the wireless device is to change the password from the default. Ideally, you'll change it to a strong password or passphrase. If you don't, an attacker might see your wireless network, log in using the defaults, and modify your network. An attacker can actually lock you completely out of your network until you reset the wireless device to factory defaults.



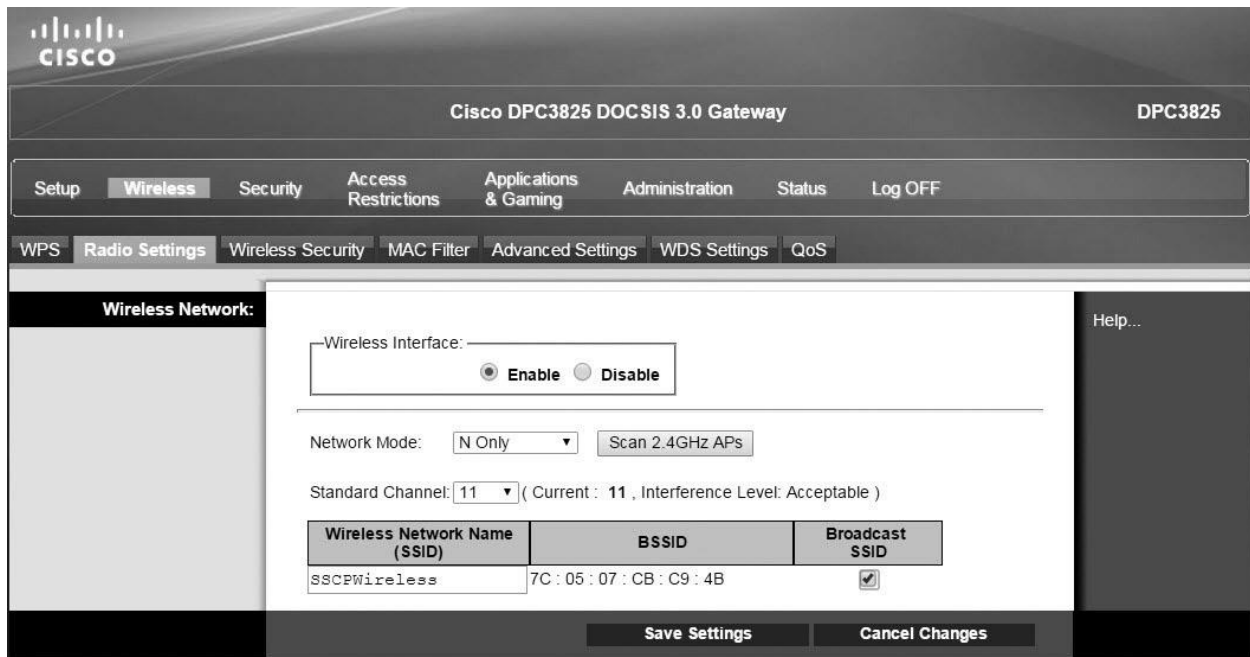
TIP Wardriving is the practice of driving around with a wireless receiver (such as a laptop with a wireless NIC) and looking for wireless networks.

Wireless Service Set Identifier

The service set identifier (SSID) is the name of the wireless network. Many wireless devices have default SSIDs, but newer ones often require the user to enter the name of the SSID as part of the setup process. For example, some common SSIDs for older wireless routers were **Linksys** (for Linksys systems) and **Netgear** (for Netgear devices).

An SSID can be almost any name you want, although it can't be longer than 32 characters. It's a user-friendly name, but you don't necessarily want to make it user friendly to attackers. For example, if your wireless network includes secret financial data, you probably don't want to call it **SecretFinancialData**.

[Figure 3-17](#) shows the configuration page for a Cisco wireless router with the SSID entered as **SSCPWireless**. You can also see that the Broadcast SSID setting is checked, enabling SSID broadcast. When the wireless device is set to enable wireless SSID broadcasting, it regularly broadcasts messages with the SSID name. This enables other wireless devices in range of the broadcast to easily identify the wired network, and if they have security configured to match the wireless router, the SSID broadcast allows these other devices to connect easily.



The screenshot shows the configuration page for a Cisco DPC3825 DOCSIS 3.0 Gateway. The page is titled "Cisco DPC3825 DOCSIS 3.0 Gateway" and "DPC3825". The navigation menu includes Setup, Wireless, Security, Access Restrictions, Applications & Gaming, Administration, Status, and Log OFF. The "Wireless" tab is selected, and the "Radio Settings" sub-tab is active. The "Wireless Network:" section is expanded, showing the following settings:

- Wireless Interface: Enable Disable
- Network Mode: N Only (dropdown) Scan 2.4GHz APs (button)
- Standard Channel: 11 (dropdown) (Current: 11, Interference Level: Acceptable)
- Wireless Network Name (SSID): SSCPWireless
- BSSID: 7C:05:07:CB:C9:4B
- Broadcast SSID:

At the bottom of the page, there are buttons for "Save Settings" and "Cancel Changes".

Figure 3-17 Configuring the SSID



TIP [Figure 3-17](#) shows other settings. It's possible to convert this to a wired router only by setting the Wireless Interface to Disable. The Network Mode is set to N Only using 2.4 GHz, so

it only supports 802.11n devices on 2.4 GHz. It can also be configured to support 802.11b, 802.11g, and 802.11n devices in mixed mode. This page also allows you to select the wireless channel, which is useful to avoid interference on another channel.

An ongoing discussion related to the SSID is whether you should enable or disable SSID broadcasting. Many people originally recommended disabling SSID broadcasting to hide the network. However, this really doesn't provide security. Even though this stops the normal SSID broadcasts used by other wireless devices, it doesn't prevent the device from transmitting the SSID over the air. An attacker with a wireless sniffer can easily capture key packets and discover the SSID, even when SSID broadcasting is disabled.

MAC Address Filtering

Most wireless devices also provide the ability to filter connections based on MAC addresses. For example, [Figure 3-18](#) shows a configuration page for a wireless router with a Wireless MAC Filter set to Enable. As configured, it will permit one computer with the MAC address of 1A:2B:3C:4D:5:6F access to the network. It will block access to all other computers.



Figure 3-18 Enabling a MAC address filter

However, it's quite simple to spoof the MAC address of a system. An attacker with a wireless sniffer can identify the MAC addresses that are connecting to a wireless network. Armed with

this knowledge, the attacker can then use software to change the wireless MAC address of the attacker's system and then connect.

In theory, MAC address filtering as a security precaution sounds good. In practice, it's easily defeated. The best protection you can use is WPA2 with CCMP. If you have a choice and can configure an 802.1x authentication server, use WPA2-Enterprise instead of WPA2-Personal (WPA2-PSK).

Bluetooth

Bluetooth is a wireless technology used for small wireless networks commonly called personal area networks (PANs). As an example, smaller wireless devices such as smartphones along with an earpiece and microphone use Bluetooth. A user can wear a Bluetooth-enabled earpiece while keeping a Bluetooth phone in their pocket.

Devices must be paired with each other before they will communicate. Pairing is done by changing the devices to discovery mode so that they can easily locate each other and connect. Once they are paired, discovery mode should be disabled. If discovery mode is left enabled, an attacker can pair with the devices and take them over. It's relatively easy for a user running Linux to exploit a Bluetooth device left in discovery mode.

The following are three well-known types of Bluetooth attacks:

- **Bluesnarfing** A hacker that is able to pair with a Bluetooth-enabled smartphone can access data on the device. The attacker can access information on the user's phone such as e-mail, contact lists, calendars, and text messages. The attacker must be close to the phone (usually within 30 feet).
- **Bluebugging** In some instances, an attacker is able to access a smartphone and issue specific commands to it. This enables the attacker to make phone calls, enable call forwarding, send messages, read and write data, connect to the Internet, and turn on the speaker to eavesdrop on conversations.
- **Bluejacking** Compared to Bluesnarfing and Bluebugging, Bluejacking is relatively harmless. It allows users to send unsolicited messages to nearby Bluetooth devices, but it does not involve modifying a phone's data.

NFC

Near field communication (NFC) is a technology available on many smartphones that allows users to make purchases with those smartphones. For example, if you have an NFC-enabled smartphone, you can make a purchase at a store by simply tapping your smartphone on the NFC reader or waving it close to the reader. Similarly, users can share information with each other by identifying the information to share and tapping the two devices together.

NFC uses similar technology as a proximity card, allowing devices to communicate with each other when they are in close proximity. The devices transmit and receive data only between short distances (identified as a *near field*). NFC encrypts this data, providing some protection. Additionally, it's often possible to configure the device to require the user to enter a PIN or authenticate with a fingerprint before it transmits any data.

Of course, if you add your credit information to your NFC-enabled smartphone, it does make the risk of losing your phone that much greater, especially if additional security hasn't been configured on the device. If someone steals your phone and can access it, they can wave and tap it during a massive shopping spree that will be billed on your credit card.

GSM

Global System for Mobile Communications (GSM) began as a European standard for 2G, or digital cellular, networks in 1982. As a comparison, 1G refers to the analog communications that came out in the 1980s, and 2G refers to the digital technologies that came out in the 1990s. Industries embraced the GSM standard and it has since grown into a large family of technologies used in mobile communications.

GSM is a standard used in 219 countries (including the United States). The GSM Association (GSMA) includes almost 800 mobile operators and 200 other companies that help standardize, deploy, and promote the GSM standard.

3G, LTE, and 4G

3G refers to the third generation of standards for wireless access used by mobile phones and mobile telecommunications. 3G provides higher transmission speeds for mobile communications, comparable to broadband speeds that wired users have. Speeds up to 3 Mbps are common. Users can surf the Web with a mobile phone and access other web-based services such as e-mail.

Improvements of 3G are sometimes dubbed 3.5G and 3.75G. They use High Speed Packet Access (HSPA), and in many cases, users can get 20 Mbps download times or better. However, not everyone can achieve these high rates of downloads. Worse, when large numbers of cellular users are operating in the same area, all users have reduced speeds because they're competing for the same bandwidth.

Telecommunication companies often market Long-Term Evolution (LTE) as 4G, but it doesn't always match the speed of 4G or meet the technical specifications. 4G (the fourth generation of mobile telecommunications) includes speeds as high as 100 Mbps for highly mobile users in trains and cars, and 1 Gbps for individuals walking or staying still.

New generations of telecommunications appear about every ten years. Fully compliant 4G systems were standardized in 2012. With that in mind, 5G systems are likely to appear sometime after 2020.

WiMAX

Worldwide Interoperability for Microwave Access (WiMAX) is a standard based on IEEE 802.16. IEEE 802.16 is officially known as Wireless MAN, but is commercially known as WiMAX. WiMAX provides broadband wireless coverage to large areas or metropolitan area networks (MANs).

For example, a WiMAX fixed station can provide wireless access for up to 30 miles. In comparison, wireless local networks using 802.11 are limited to about 300 feet. Many providers have created WiMAX networks around the world, and users can rent access to the networks. WiMAX provides speeds up to 1 Gbps for wireless connections.

Radio Frequency Identification

Radio frequency identification (RFID) is used for identification, tracking, asset management, and inventory control. Products are “tagged” with an RFID tag, which is a small electronic device that marks the product. RFID readers can then read the tag to get information about the product.

RFID tags can be either active (with a battery) or passive (without a battery). Active RFID tags are constantly transmitting. A passive RFID tag works similarly to how proximity cards work for authentication. An RFID reader can excite the electronics on the RFID tag and cause the tag to transmit its information.



TIP [Chapter 2](#) describes proximity cards in more depth within the “Something You Have” section.

Protecting Mobile Devices

The primary vulnerabilities with mobile devices occur if they are lost. If someone finds or steals a mobile device, that person may be able to get some valuable information. With this in mind, users can do several things to protect their devices. These include the following:

- **Password protect devices** Most devices support password protection. Users must enter a password or personal identification number (PIN) to unlock the device to use it. If a thief gets the device, this password makes it more difficult for the thief to access the data.
- **Encrypt data** Many devices support data encryption, and some even support full device encryption. If a thief does get an encrypted smartphone, the encryption provides confidentiality by restricting the thief’s access to the data. This encryption is often enabled by default when a passcode (such as a PIN) is enabled.

- **Remote wipe** Some devices support remote wipe. This allows the owner to send a command to a lost device to delete all the data on the device. The command will delete cached data, such as cached passwords, along with other data such as information in the user's contact list. For mobile devices that do not have a built-in remote wipe capability, users can install apps to perform the same function.
- **Enable GPS** Many devices support global positioning system (GPS) location services. If the device is lost, the user can use the GPS service to determine its location.



EXAM TIP Mobile devices are easily lost. Because they can often contain valuable data, it's important to know the different methods available to protect data. Password protection and data encryption are valuable if done before the device is lost. Remote wipe and GPS are valuable after the device is lost.

While this section identified some specific security controls that users can implement to protect their devices, it's also possible for an organization to configure these settings remotely. [Chapter 9](#) covers mobile device management (MDM) in more depth.

Chapter Review

The OSI Model is a framework for connecting multiple computers to each other. The ISO developed it and it has seven layers. The layers, listed from layer 1 to layer 7, are Physical, Data Link, Network, Transport, Session, Presentation, and Application. The TCP/IP Model (also called the TCP Model and the DoD Model) is similar to the OSI Model, but it has fewer layers. RFCs 1122 and 1123 identify four layers of the OSI Model, but some references identify five layers.

Two important protocols operate on the Transport layer (layer 4). They are TCP and UDP. TCP is connection-oriented and establishes sessions with a three-way handshake. UDP is connectionless and doesn't use a three-way handshake to establish a session. Instead, UDP uses its best effort to send the data. Other protocols can work with either TCP or UDP, depending on the protocol's needs. For example, FTP uses TCP for guaranteed delivery, while TFTP uses UDP.

Network configurations include bus, star, tree, token ring, and mesh. In a bus configuration, all devices share a common bus cable connection that has terminators at both ends. A bus is inexpensive to install but very difficult to troubleshoot, increasing the total cost of ownership. A star configuration is much more common and uses a central device such as a switch. All devices connect to each other via the central device in a star. A tree is a combination of bus and star topologies. A token ring configuration uses a logical token, and devices in the ring can transmit only when they have the token. Traditional token ring configurations are slower and don't scale well. FDDI uses a token ring concept but adds a second ring for redundancy and has significantly

higher speeds because it uses fiber-optic cables. Mesh networks have multiple redundant connections, connecting each device with all other devices, and they provide the highest availability.

Network relationship types include peer-to-peer and client-server. Smaller networks use a peer-to-peer configuration and have decentralized authentication. Larger networks typically use a client-server configuration with centralized authentication. In a Microsoft environment, user accounts are maintained and managed on domain controllers.

TCP/IP includes a full suite of protocols used by computers on the Internet and on many internal networks. The SSCP exam doesn't expect you to be an expert on all the protocols within TCP/IP, but you should be familiar with many of them.

DHCP provides IP addresses and other TCP/IP configuration information to hosts. ARP resolves IPv4 addresses to MAC addresses and must be running to transfer data on a network. RARP is similar to DHCP in that it provides IP addresses to clients with MAC addresses (also known as physical or hardware addresses). BootP provides diskless clients with an operating system. NDP is similar to ARP but resolves IPv6 addresses to EUI-64 addresses.

DNS uses a hierarchical naming system with a distributed database to resolve host names to IP addresses. DNSSEC protects against DNS poisoning attacks using Resource Record Signatures (RRSIGs), which are similar to digital signatures. The digital signature provides authentication of the response, validating the integrity of the DNS response.

ICMP is used for diagnostics. While it is useful when it is running, it is often disabled or blocked at firewalls to thwart attacks that use ICMP. IPv4 uses IGMP for multicasting, while IPv6 uses ICMPv6 for multicasting. SNMP is used to manage network devices. Agents receive data on UDP port 161 and send traps and other data over UDP port 162.

FTP is efficient for sending large files, while TFTP is good for sending small files such as configuration files for network devices. FTP uses TCP ports 20 and 21, and supports authentication. TFTP uses UDP port 69 and does not include authentication. Both FTP and TFTP send data across a network in cleartext.

Sending data over a network in cleartext is a security risk because an attacker with a sniffer can capture these transmissions and read the data. Several protocols, such as Telnet and rlogin, send credentials over a network in cleartext. SSH is a more secure alternative as it encrypts the transmissions using a combination of symmetric and asymmetric encryption for confidentiality. SSH uses hashing for integrity, and can use an authenticated tunnel for mutual authentication. SSH can also encrypt other protocols such as FTP (as SFTP) and Secure Copy. SSH uses TCP port 22.

HTTP is the common protocol used to transfer web pages over the Internet. HTTPS encrypts e-commerce transactions with SSL or TLS. TLS is the designated replacement for SSL. HTTP uses TCP port 80 and HTTPS uses TCP port 443.

Routers use routing protocols such as RIPv2 and OSPF to share information with each other. RIPv2 is good for smaller networks, while OSPF works better in larger networks. Both protocols operate on the Network layer.

E-mail protocols include SMTP, POP3, and IMAP4. Clients send e-mail with SMTP and receive e-mail with POP3. IMAP4 allows clients to access and manipulate e-mail on a server, including organizing e-mail within folders. SMTP uses TCP port 25, POP3 uses TCP port 110, and IMAP4 uses TCP port 143.

VPNs use tunneling protocols to protect data. PPTP and L2TP are two tunneling protocols, and both operate on the Data Link layer. PPTP provides encryption, while L2TP depends on IPsec for encryption. IPsec uses AH for authentication and integrity. IPsec ESP includes the functionality of AH and provides confidentiality by encrypting the data. AH uses protocol number 51, and ESP uses protocol number 50. TLS can also be used as a VPN tunneling protocol.

Both TCP and UDP can use any of 65,536 ports (numbered from 0 to 65,535). Ports 0 to 1023 are well-known ports, and IANA maps them to specific protocols. For example, SSH uses port 22, HTTP uses port 80, and SMTP uses port 25. You should be able to identify the port used by many common protocols, in addition to the protocol numbers used by other protocols.

Different network architectures provide different levels of trust. For example, an intranet is internal within an organization and provides a high level of trust. The Internet is external and has the least level of trust. A DMZ hosts Internet-facing servers that are accessible by anyone on the Internet, and an extranet hosts Internet-facing servers that are accessible only by trusted entities such as business partners. Internal clients use private IP addresses, and Internet hosts have public IP addresses. NAT translates IP addresses from public to private and from private back to public.

Wireless networks (sometimes called WiFi networks) allow clients to connect using radio-wave transmissions. These radio-wave transmissions are the easiest to capture and intercept (compared to transmissions on wired networks), but some basic security steps can increase the security of a wireless network. WEP is older and should not be used. WPA was an interim replacement for WEP and was compatible with existing hardware. WPA2 is the permanent replacement and is formally defined by 802.11i. WPA originally used TKIP, but many WPA systems also support AES. WPA2 uses AES-based CCMP. Both WPA and WPA2 can use Personal mode (with a preshared key) or Enterprise mode. Enterprise mode requires an 802.1x authentication server.

Smaller wireless devices use technologies such as Bluetooth, NFC, GSM, 3G, LTE, 4G, and WiMAX. Several methods protect these devices if they are lost. Methods include password-protecting the devices, encrypting data on them, using remote wipe to delete data on lost devices, and enabling GPS to locate lost devices.

Questions

1. Which layer of the OSI Model defines cable standards?

- A. Physical layer
- B. Data Link layer
- C. Network layer
- D. Transport layer

2. Which layer of the OSI Model packages data as a frame?

- A. Physical layer
- B. Data Link layer
- C. Network layer
- D. Transport layer

3. Which layer of the OSI Model provides reliable end-to-end communication services?

- A. Physical layer
- B. Transport layer
- C. Data Link layer
- D. Host layer

4. Which layer of the OSI Model includes TCP and UDP?

- A. Transport layer
- B. Network layer
- C. Data Link layer
- D. Application layer

5. Which of the following protocols is connection-oriented?

- A. IP
- B. RIP
- C. TCP

D. UDP

6. User accounts are located on a central authentication server. What type of network relationship does this represent?

A. Peer-to-peer

B. Client-server

C. Star

D. Tree

7. What protocol would a system use to determine a system's physical address?

A. ARP

B. RARP

C. BootP

D. DNS

8. Your organization plans to place a server in the DMZ to provide name resolution services. However, management wants to ensure that data it sends maintains integrity. Which of the following will BEST meet these goals?

A. DNS

B. TLS

C. DHCP

D. DNSSEC

9. Which of the following is the recommended protocol to secure HTTPS sessions?

A. SSL

B. TLS

C. SSH

D. IPsec

10. Which of these ports does DNS use?

- A. TCP 22
- B. TCP 25
- C. UDP 53
- D. UDP 69

11. Which of the following protocols is commonly used with diagnostic utilities?

- A. TFTP
- B. RARP
- C. IGMP
- D. ICMP

12. Which of the following accurately identifies a difference between FTP and TFTP?

- A. FTP uses UDP, and TFTP uses TCP.
- B. FTP supports authentication, but TFTP does not support authentication.
- C. TFTP sends data across a network in cleartext, but FTP encrypts data.
- D. TFTP is primarily used to transfer large files, and FTP is used to transfer configuration information to and from network devices.

13. Which of the following protocols is the most secure method to remotely administer a server?

- A. Telnet
- B. rlogin
- C. rexec
- D. SSH

14. What port does POP3 use?

- A. 25
- B. 110

C. 143

D. 443

15. You are purchasing a product from a website. Which of the following protocols will your system most likely use to provide confidentiality for this transaction?

A. TLS

B. SSH

C. IPsec

D. HTTP

16. Which of the following statements is correct related to IPsec?

A. IPsec provides confidentiality by encrypting data with AH.

B. IPsec provides confidentiality by encrypting data on the Network layer.

C. IPsec AH uses protocol number 50.

D. IPsec ESP uses protocol number 51.

17. You need to allow IPsec AH through your firewall. Which of the following should you configure to allow this traffic through the firewall?

A. Port 50

B. Port 51

C. Protocol number 50

D. Protocol number 51

18. Where is a DMZ located?

A. Behind the intranet firewall

B. In front of the first intranet-facing firewall

C. In front of the first Internet-facing firewall

D. Behind the first Internet-facing firewall

19. Which of the following is the recommended security mechanism to use with wireless networks?

- A. 802.11a
- B. 802.11g
- C. 802.11i
- D. 802.11n

20. You have created a wireless network for guests and several guests have connected to the network in the past. Today, new guests are unable to see the wireless network. However, guests that connected to it previously are using it without any trouble. What is the MOST likely problem?

- A. The PSK has been changed.
- B. The AP is now using WPA2-Enterprise.
- C. SSID broadcasting has been disabled.
- D. MAC address filtering has been enabled.

Answers

1. A. Cable standards are defined at the Physical layer, layer 1. They are not defined at the Data Link layer (layer 2), the Network layer (layer 3), or the Transport layer (layer 4).

2. B. The Data Link layer packages data as a frame. The Physical layer packages data as bits. The Network layer packages data as a packet. The Transport layer packages data as a segment.

3. B. The Transport layer provides reliable end-to-end communication services. Neither the Physical layer nor the Data Link layer provides this service. The Host layer is on the TCP/IP Model, not the OSI Model.

4. A. The Transport layer includes the TCP and UDP protocols. These protocols are not implemented on the Network layer, the Data Link layer, or the Application layer.

5. C. TCP is connection oriented. IP uses TCP to provide a connection-oriented session but is not connection-oriented itself. RIP is a routing protocol and is not connection-oriented. UDP is connectionless. Instead of establishing a session, it makes a best effort to deliver data.

6. B. In a client-server network relationship, users log on to a central authentication server, such as a domain controller. A peer-to-peer network relationship has decentralized authentication. Star and tree refer to network topologies, not network relationships.

7. A. Systems use the Address Resolution Protocol (ARP) to identify the assigned physical (or MAC) address matching an assigned IP address. Reverse ARP (RARP) allows a system with a MAC address to get an IP address. The Bootstrap Protocol (BootP) allows a diskless system to get an IP address and then download the image of an operating system. Domain Name System (DNS) resolves host names to IP addresses but not physical addresses.

8. D. Domain Name System Security Extensions (DNSSEC) adds security to DNS systems and best meets these goals. DNS responses are signed with a Resource Record Signature (RRSIG) to provide authentication of the response and prove it hasn't lost integrity. While a server that provides name resolution services is a DNS server, a DNS server without DNSSEC will not meet the stated goals. While DNSSEC can use Transport Layer Security (TLS), TLS by itself provides encryption for confidentiality, not authentication for integrity. Dynamic Host Configuration Protocol (DHCP) assigns TCP/IP configuration information for DHCP clients and is unrelated to this question.

9. B. Transport Layer Security (TLS) encrypts Hypertext Transfer Protocol Secure (HTTPS) sessions. Secure Sockets Layer (SSL) has known vulnerabilities and should not be used. Secure Shell (SSH) is used to connect to remote servers, but not to encrypt HTTPS. Internet Protocol security (IPsec) is used to encrypt VPN tunnels, but not for HTTPS.

10. C. Domain Name System (DNS) uses UDP port 53 when clients query the DNS server and TCP port 53 when DNS servers transfer data between each other. Secure Shell (SSH) uses TCP port 22. Simple Mail Transfer Protocol (SMTP) uses TCP port 25. Trivial FTP (TFTP) uses UDP port 69.

11. D. Many diagnostic utilities such as ping, pathping, and tracert use Internet Control Message Protocol (ICMP). Administrators commonly use Trivial FTP (TFTP) to transfer configuration files to and from network devices. Reverse ARP (RARP) allows a system with a MAC address to get an IP address. The Internet Group Message Protocol (IGMP) is used for IPv4 multicasting.

12. B. File Transfer Protocol (FTP) supports authentication, but Trivial FTP (TFTP) does not support authentication. FTP uses TCP ports 20 and 21, while TFTP uses UDP port 69. Both FTP and TFTP send data across a network in cleartext, but it is possible to encrypt FTP with Secure Shell (as SFTP). TFTP is commonly used to transfer configuration files to and from network devices, and FTP is primarily used to transfer large files.

13. D. Secure Shell (SSH) encrypts data sent over a network and is the most secure method for remotely accessing systems of the given choices. Telnet, rlogin (which is remote login), and rexec (remote execute) all send data across a network in cleartext.

14. B. Post Office Protocol version 3 (POP3) uses TCP port 110. Simple Mail Transfer Protocol (SMTP) uses TCP port 25. Internet Message Access Protocol version 4 (IMAP4) uses TCP port 143. Hypertext Transfer Protocol Secure (HTTPS) uses TCP port 443.

15. A. E-commerce transactions use Hypertext Transfer Protocol Secure (HTTPS) for confidentiality, and Transport Layer Security (TLS) is used to encrypt HTTPS. While Secure

Shell (SSH) and Internet Protocol security (IPsec) both provide confidentiality with encryption, HTTPS doesn't use SSH or IPsec. HTTP sends data in cleartext, so it doesn't provide confidentiality.

16. B. Internet Protocol security (IPsec) provides confidentiality by encrypting data on the Network layer. Encapsulating Security Payload (ESP) provides confidentiality by encrypting data, but Authentication Header (AH) only provides authentication and integrity. AH uses protocol number 51, and ESP uses protocol number 50.

17. D. You should enable protocol number 51 because it is the protocol number for Internet Protocol security (IPsec) Authentication Header. The protocol number for IPsec Encapsulating Security Protocol (ESP) is 50. IPsec AH and ESP are identified by protocol numbers, not port numbers.

18. D. A demilitarized zone (DMZ), or perimeter network, is located behind the first Internet-facing firewall. It is not on the private network (behind the intranet firewall) or directly on the Internet (in front of the intranet or Internet-facing firewall).

19. C. The 802.11i standard documents Wi-Fi Protected Access 2 (WPA2), the recommended security mechanism for wireless networks. It uses AES-based CCMP for very strong security. The other standards focus on the base frequency and speed of wireless networks, not security.

20. C. The most likely problem is that service set identifier (SSID) broadcasting has been disabled, primarily because new users are unable to see the wireless network. Clients that previously connected could have configured their systems to connect automatically. If the preshared key (PSK) was changed or the AP was changed to use WPA2-Enterprise, no one would be able to connect. It would be extremely difficult to enable media access control (MAC) filtering for only guests that previously connected, so this is highly unlikely.

NOTES

CHAPTER 4

Advanced Networking and Communications

In this chapter, you will learn about

- The difference between switches and routers
- Physical and logical segmentation of hosts and networks
- Securing the Voice over Internet Protocol (VoIP) with the Secure Real-time Transport Protocol (SRTP)
- Proxy servers and filtering methods
- Packet-filtering, stateful inspection, and application firewalls
- Defense diversity with firewalls
- The differences between network-based and host-based firewalls
- Risks and vulnerabilities related to remote access solutions
- Different tunneling protocols used with remote access
- Authentication methods used with remote access
- Network access control
- Virtualization environments
- Separating the control and data planes
- Cloud service models, shared responsibility, and operation models

Managing LAN-based Security

A local area network (LAN) connects computing devices together, allowing users to share resources such as data and devices. While LANs have helped improve productivity for many people, they have also added many risks. Because of this, it's important for IT professionals to understand networks and some basic LAN-based security methods.

Comparing Switches and Routers

Switches connect devices together in a network and routers connect networks together. [Figure 4-1](#) shows devices connected together via a switch in Network 1, devices connected together via a

switch in Network 2, and both networks connected together via a router. Notice that Port 1 of the router is the gateway for computers in Network 1 to other networks. Similarly, Port 2 is the gateway for computers in Network 2 to other networks. Most networks have only one gateway labeled as the default gateway.

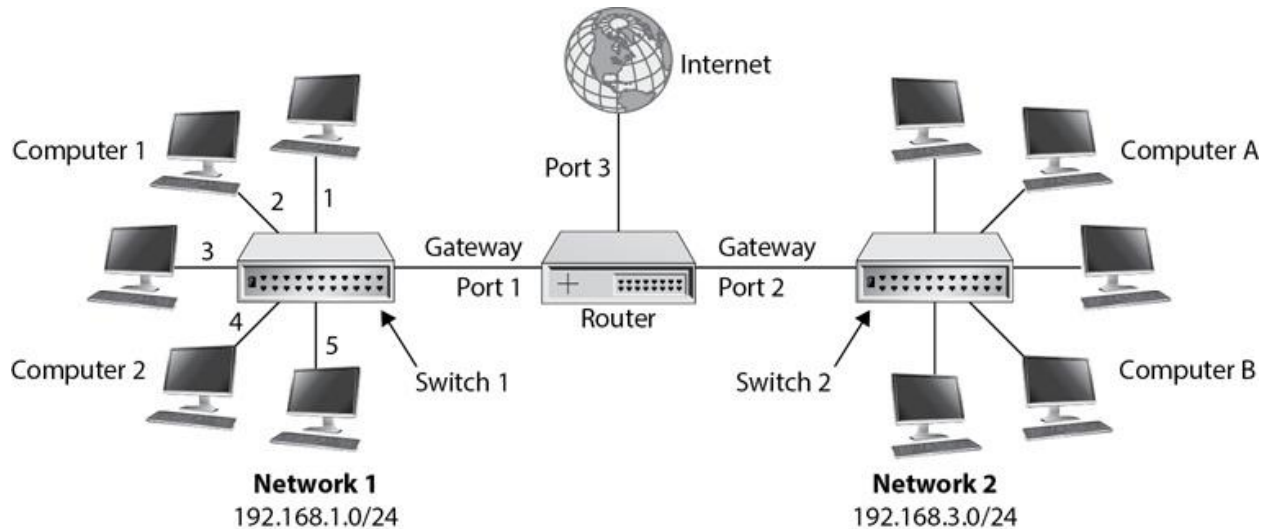


Figure 4-1 Comparing switches and routers



NOTE Networks connected together in the same area, such as within an office building or school, are local area networks (LANs). However, technicians often just refer to a LAN as a network.

[Chapter 3](#) covers the Open Systems Interconnection (OSI) Model and mentions how switches operate on layer 2 (the Data Link layer) and routers operate on layer 3 (the Network layer). Layer 2 uses media access control (MAC) addresses, and switches control traffic within a network using the MAC addresses. Layer 3 uses Internet Protocol (IP) addresses, and routers control traffic between networks using IP addresses.

Switches typically build tables mapping physical ports to the MAC address of the device connected to the port. For example, when Computer 1 sends traffic through the switch, the switch captures its MAC address and records it in the table. While administrators can statically configure the MAC address table, most switches create the table dynamically. [Table 4-1](#) shows a table for the switch in Network 1. Note that Computer 1 is connected to physical Port 2 and Computer 2 is connected to physical Port 4. Other computers are connected to Ports 1, 3, and 5 but are just listed as other devices to simplify the discussion.

| Physical Port | MAC Address |
|--------------------|-------------------|
| 1 (Other Computer) | 76-29-af-e4-33-99 |
| 2 (Computer 1) | 74-29-af-e4-33-9a |
| 3 (Other Computer) | 34-64-a9-13-44-08 |
| 4 (Computer 2) | 34-64-a9-13-44-a8 |
| 5 (Other Computer) | 00-14-22-a3-32-c8 |

Table 4-1 MAC Address Table in a Switch

When Computer 1 sends unicast traffic to Computer 2, it includes the destination MAC address (34-64-a9-13-44-a8). The switch uses the MAC address table and sees that the destination MAC address is connected to Port 4, so it forwards the traffic to Port 4. This traffic doesn't reach any of the other ports.

In contrast, older networks used hubs in place of switches. A hub forwards traffic going into one port to all other ports. This represents a risk because an attacker could configure a protocol analyzer (or sniffer) to capture traffic going through the hub. [Chapter 5](#) covers sniffing attacks in more depth.

Switches can also restrict traffic with MAC address filtering, as discussed in [Chapter 3](#).



NOTE Layer 3 switches are an advanced implementation of switches. They work similarly to how routers work on layer 3.

Routers use tables too, but they use IP addresses and network addresses instead of MAC addresses. As an example, the routing table for the router in [Figure 4-1](#) would include the following information:

- Physical Port 1 is connected to Network 1 (all IP addresses in the 192.168.1.0/24 network).
- Physical Port 2 is connected to Network 2 (all IP addresses in the 192.168.3.0/24 network).
- The path to all unknown networks (such as those on the Internet) is through Port 3.



EXAM TIP The /24 in 192.168.1.0/24 indicates the subnet mask is 255.255.255.0. It is using Classless Inter-Domain Routing (CIDR) notation and indicates the first 24 bits in the subnet mask are 1's. Check out this free mini-tutorial that provides an overview on subnetting and CIDR notation: <http://gcgpremium.com/courses/subnetting>.

Administrators can create routing tables manually or configure routers with routing protocols. Routing protocols such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) allow routers to share information with other routers and learn the best path to a different network.

Physical and Logical Segmentation

Network segmentation refers to breaking up a larger network into smaller interconnected networks. [Figure 4-1](#) shows how a router segments traffic between two networks. The router divides a large network into two smaller networks using *physical segmentation*. Another method of segmenting a network is with virtual local area networks (VLANs) using *logical segmentation*. A VLAN allows an organization to segment traffic with a lot more flexibility when compared to segmenting traffic with a router.

Consider the single-network configuration shown in [Figure 4-2](#). Imagine that Switch 1 connects the computers on a building's first floor, and Switch 2 connects the computers on a building's second floor. Switch 3 is connecting the two switches to the router. Configured like this, it creates a single broadcast domain. In other words, if any computer transmits broadcast traffic, it will reach all of the computers on the first and second floors. Routers do not pass broadcast traffic, so the broadcast domain does not extend beyond the router.

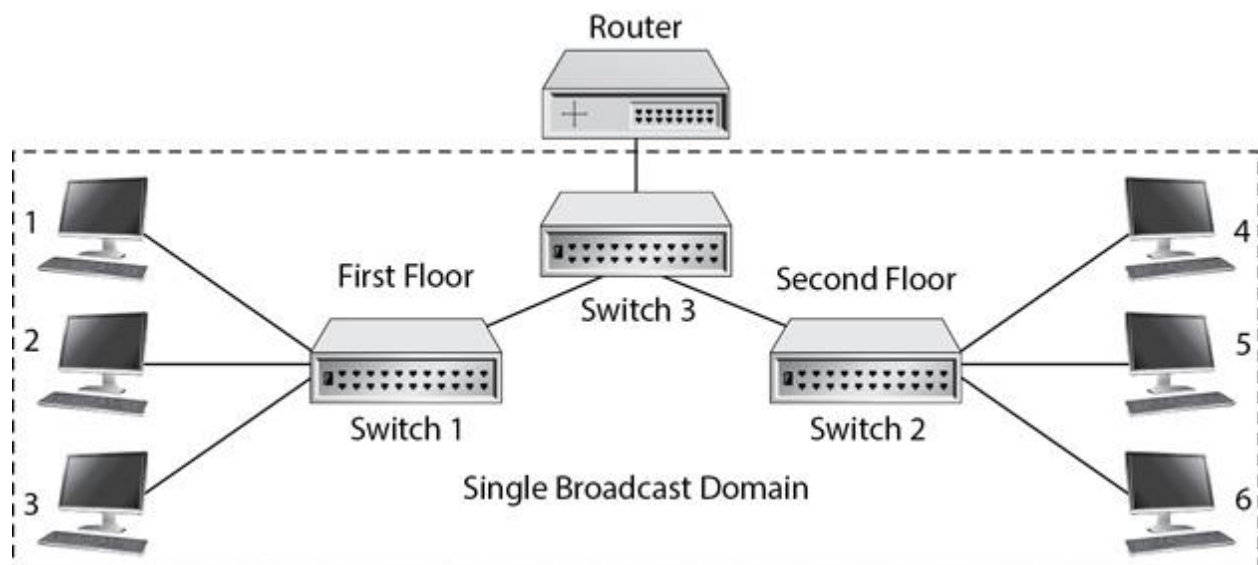


Figure 4-2 Building network on two floors

While this configuration works, it also presents a potential risk. Imagine that employees using computers 1 and 4 are working on a sensitive project. Data broadcasts between these two

computers will go to all computers in the broadcast domain, meaning someone within the broadcast domain can capture this traffic with a sniffer.

One way to segment these two computers is to add another switch on the other side of the router. You can then connect the two computers to the switch as shown in [Figure 4-3](#). However, this is highly inefficient. You have to add a switch and run additional cables from the switch to the computers on two separate floors.

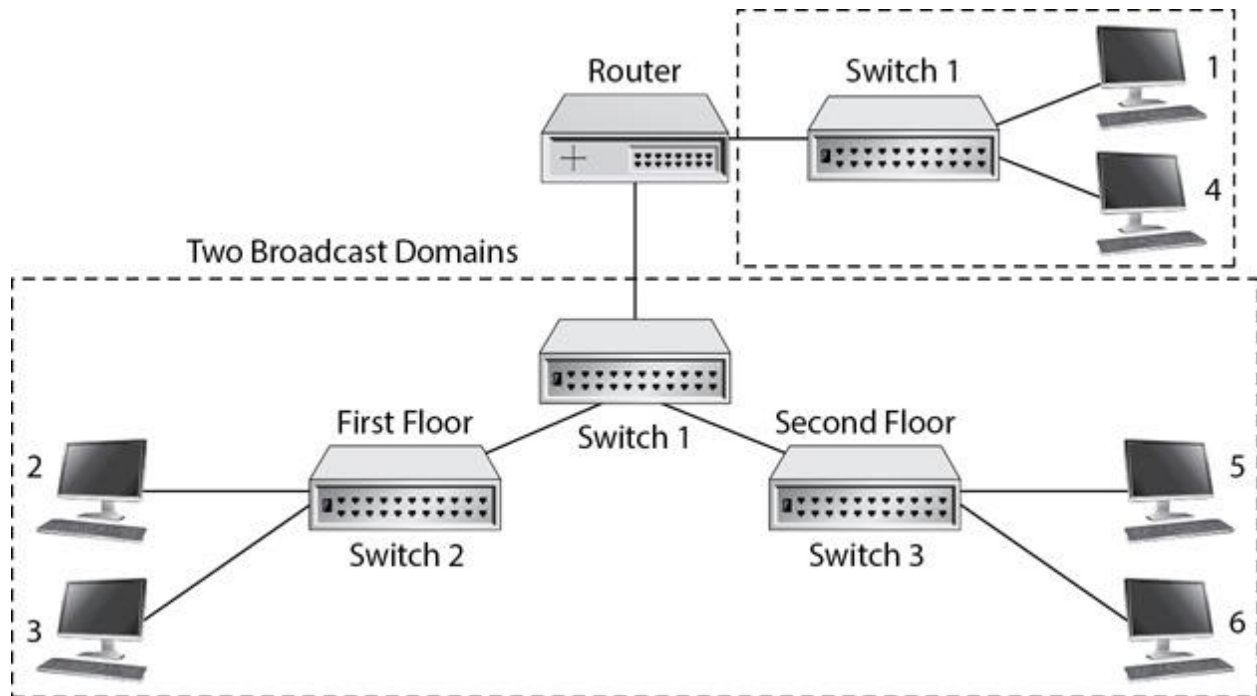


Figure 4-3 Inefficient network segmentation

VLANs provide a much more efficient solution, as shown in [Figure 4-4](#). Computers 1 and 4 are within the VLAN named VLAN 10. This creates a logical broadcast domain for these two computers just as if they were configured with another switch connected to a different router port as shown in [Figure 4-3](#). The switch doesn't pass broadcast traffic within VLAN 10 to other devices on the network. Similarly, the switch won't pass broadcast traffic from outside of VLAN 10 to computers 1 and 4. Less broadcast traffic also improves the performance of the LAN.

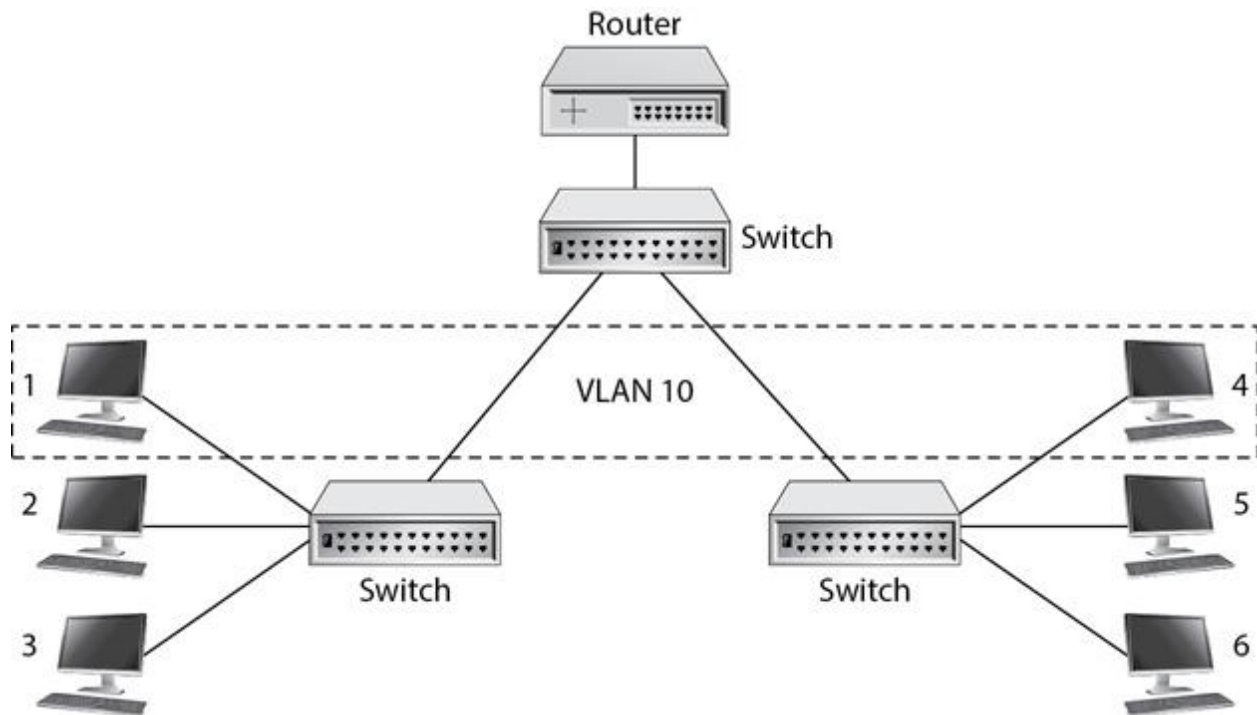


Figure 4-4 VLAN configuration



NOTE When configuring VLANs, it's important to configure all of the switches within the network to recognize the VLAN. This ensures that broadcast traffic within a VLAN does not reach devices connected to other switches.

While [Figure 4-4](#) shows a single VLAN, it's also possible to create multiple VLANs. For example, you could create a VLAN for computers 2 and 5, and another VLAN for computers 3 and 6. VLANs provide several benefits:

- **Security** VLANs separate traffic from one broadcast domain to another. This is especially useful if sensitive data is transmitted over a network.
- **Improved performance** VLANs reduce the size of broadcast domains and improve their performance because there is less traffic within each broadcast domain.
- **Reduced costs** A VLAN doesn't require any additional hardware if the switches support VLANs.



EXAM TIP VLANs increase security by segmenting traffic on a network. Administrators create VLANs with switches.

Secure Device Management

It's important to protect routers and switches with strong physical security. A data center or server room already has strong physical security, so it is a good location for routers and switches. Another option that many organizations use is to store them in wiring closets. However, it's important to ensure that the wiring closet is kept locked to prevent unauthorized access.

If attackers can access the devices, they can hook up a sniffer and capture traffic going through them. As an example, an attacker can connect a rogue access point to a switch. The access point will then transmit captured data wirelessly. The attacker can collect the data using a laptop while sitting in the parking lot outside the organization.

By default, a switch only forwards broadcast traffic to other ports, so the attacker cannot see all traffic going through the switch. However, most switches (and some routers) support port mirroring. Port mirroring sends a copy of all packets going through a switch to a separate port used for monitoring. Cisco switches support port mirroring with a feature called Switched Port Analyzer (SPAN), or the newer Remote SPAN (RSPAN) and Encapsulated RSPAN (ERSPAN). If the attacker can configure port mirroring using one of these features, it's possible to capture all traffic going through the switch.



NOTE While Cisco routers don't support port mirroring with SPAN, RSPAN, or ERSPAN, there are other methods to achieve the same result. For example, many Cisco routers support the IP Traffic Export feature, which will send a copy of all traffic to a monitored port. Similarly, Juniper routers and switches include methods for port mirroring with their devices.

VoIP

Voice over Internet Protocol (VoIP) includes the different technologies used to transmit multimedia and voice communications over IP networks, including the Internet. It is possible to transmit VoIP over the Internet in a plain, unencrypted format, but such transmissions are easily intercepted. Instead, many VoIP implementations provide some type of protection for the transmissions.

RFC 3711 defines Secure Real-time Transport Protocol (SRTP), which provides confidentiality, message authentication, and replay protection for audio and video traffic, including VoIP. [Chapter 14](#) covers cryptography topics in greater depth, but it's worth mentioning that SRTP supports strong cryptographic methods such as Advanced Encryption Standard (AES) for

encryption, and salting to protect against offline attacks. Salting adds bits to passwords or encryption keys to thwart brute-force attacks.



EXAM TIP SRTP can protect VoIP transmissions and provides confidentiality, authentication, and replay protection.

Using Proxy Servers

A proxy server is used as an intermediary, or proxy, for clients. Instead of clients' retrieving web pages directly from the Internet, they send the request to the proxy server. The proxy server retrieves the web page on the client's behalf and returns it to the client.

[Chapter 3](#) briefly introduced proxy servers in the context of Network Address Translation (NAT). Remember that NAT translates private IP addresses to public addresses and public IP addresses back to private. Proxy servers commonly include NAT as an additional service. Additionally, proxy servers provide two more primary services: caching and filtering. [Figure 4-5](#) shows how caching and filtering work.

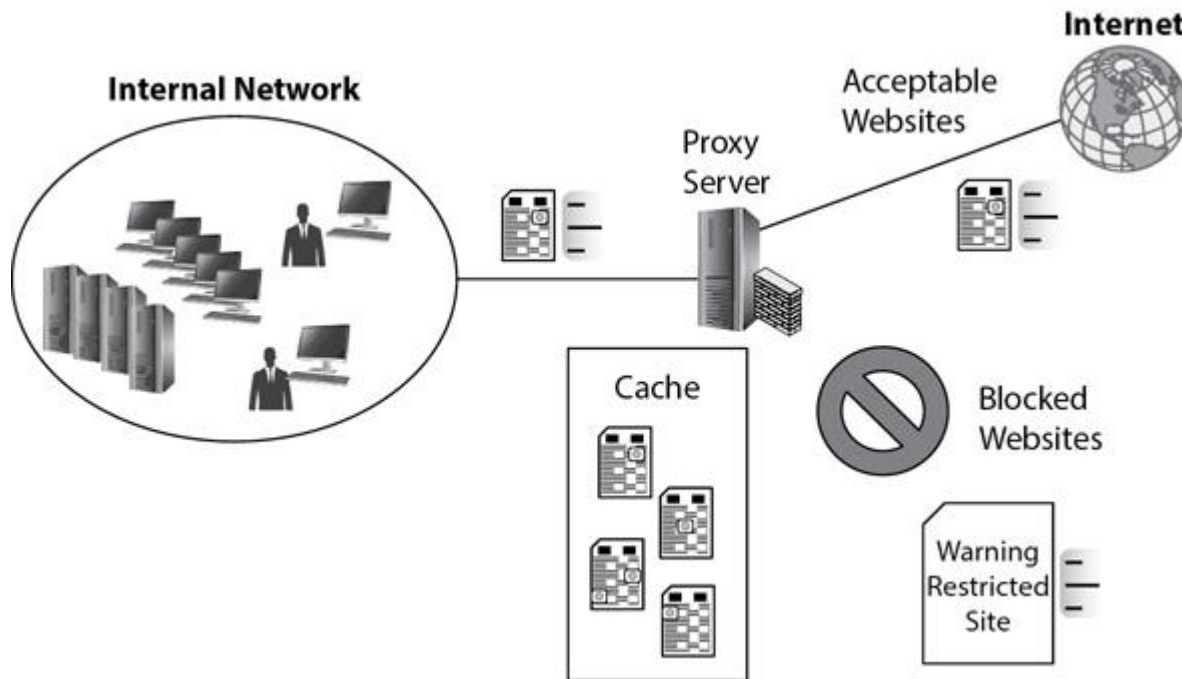


Figure 4-5 Proxy used for caching and filtering

When a user first retrieves a web page, the proxy server retrieves it from the Internet, stores a copy in cache, and sends a copy to the user. If another user requests the same page, the proxy server retrieves it from cache and sends the page to the other user instead of retrieving it from the

Internet again. This improves performance for the clients and also reduces Internet bandwidth usage.



NOTE Many web pages include metadata that can indicate when a web page is no longer valid. If the web page times out based on this metadata, the proxy server flushes it from the cache. If a user requests the page again, the proxy server retrieves another copy from the Internet and places the new copy in cache.

Most proxy servers can filter traffic based on the Uniform Resource Locator (URL). For example, if an organization wants to block employees from accessing websites with certain types of content, such as gambling sites, they can use a filter to block the URLs for these websites. Many organizations purchase lists of websites in certain categories from third-party companies. For example, if an organization wants to block users from accessing web-based e-mail, it can purchase a list of websites in the web-based e-mail category.



TIP Lists of websites are available in multiple categories, including gambling, shopping, pornography, and many more. Organizations subscribe to the lists that they want to block, and the vendor updates these lists regularly as new sites appear.

Administrators install the list of websites on the proxy server, and when users attempt to access a site on the list, the proxy server blocks the request. Instead, the proxy server displays a page indicating that the proxy server blocked their query. Many organizations personalize this web page to remind users about the organization's security policy.

Some organizations have switched to whitelisting websites. For example, instead of blocking specific websites and allowing all others, an organization can list only the sites that are allowed and block all others. Although this does limit a user's ability to surf the Internet, that is sometimes the goal.

Whitelists typically need to include more than just the website's URL. As an example, many websites pull JavaScript and cascading stylesheet (CSS) libraries from other locations on the Internet.

Website developers often use JavaScript libraries to simplify the development process. These libraries include prewritten JavaScript snippets that developers use to create rich web applications. These libraries are typically stored elsewhere on the Internet, and called by web applications when they're needed. When using a whitelist, the whitelist also needs to include the location of the JavaScript library.



TIP A whitelist identifies sites that are authorized and a proxy server blocks all sites not on the whitelist. In contrast, a blacklist identifies unauthorized sites and a proxy server allows access to any other sites.

Proxy servers are placed logically between the user and the Internet. It's common to place the proxy server in the demilitarized zone (DMZ) and configure web browsers to send all web page requests to the Internet via the proxy server.

Understanding Firewalls

[Chapter 3](#) introduced firewalls when describing differences between the Internet, an intranet, an extranet, and a DMZ. Network-based firewalls are placed inline with the traffic and they filter all traffic going into or out of a network. Host-based firewalls are software applications that run on individual systems such as a server or desktop computer. Firewalls have gone through several generations, as described in the following sections.

Packet-Filtering Firewall

The first generation of firewalls is packet-filtering firewalls that filter traffic based on the contents of a packet. More specifically, a packet-filtering firewall can filter traffic based on IP addresses, subnet addresses, ports, some protocols, or any combination of these:



EXAM TIP A packet-filtering firewall uses rules within access control lists (ACLs) to filter the traffic. ACLs filter traffic based on source or destination IP addresses, subnet addresses, entire domains, ports, and/or protocols. The ACLs provide segmentation of a network, by restricting traffic into or out of the network.

- **IP addresses** Each packet contains the source and destination IP addresses. A packet filter can examine this data and either block or allow traffic to or from specific IP addresses. For example, if an attacker is launching an attack from a specific IP address, it's possible to create a rule to block traffic from that IP address.
- **Subnet addresses** In addition to filtering specific IP addresses, you can filter entire subnets. For example, the network ID of 192.168.1.0/24 could be used to block or allow IP addresses from 192.168.1.1 through 192.168.1.254. Similarly, rules can block entire blocks of IP addresses, such as those from another country.

- **Ports** Well-known ports (ports from 0 to 1023) identify many protocols. For example, the well-known port for the Simple Mail Transport Protocol (SMTP) is port 25. By allowing traffic through port 25 (or opening port 25), you are effectively allowing SMTP traffic.
- **Some protocols** Many protocols are identified by a protocol number (or protocol ID) embedded in the IP header. For example, the protocol number for Internet Control Message Protocol (ICMP) is 1. A rule that blocks traffic with a protocol number of 1 blocks all ICMP traffic.



EXAM TIP [Chapter 3](#) covers ports in more depth and includes a table listing many commonly used ports. Additionally, [Chapter 3](#) describes protocol numbers and includes a table listing many commonly used protocol numbers.

- **A combination of IP addresses, ports, or protocols** Any combination of IP addresses, ports, or protocols can be used. For example, consider [Figure 4-6](#), which shows an e-mail server (named mail1) within a DMZ and another e-mail server (named mail2) in the intranet. Administrators would configure Firewall 1 to allow SMTP traffic (using port 25) to the mail1 server (using IP address 192.168.1.25). Similarly, administrators would configure Firewall 2 to allow e-mail traffic between mail1 and mail2. However, unless administrators add more rules to these firewalls, they will block SMTP traffic coming from the Internet and going directly to mail2.

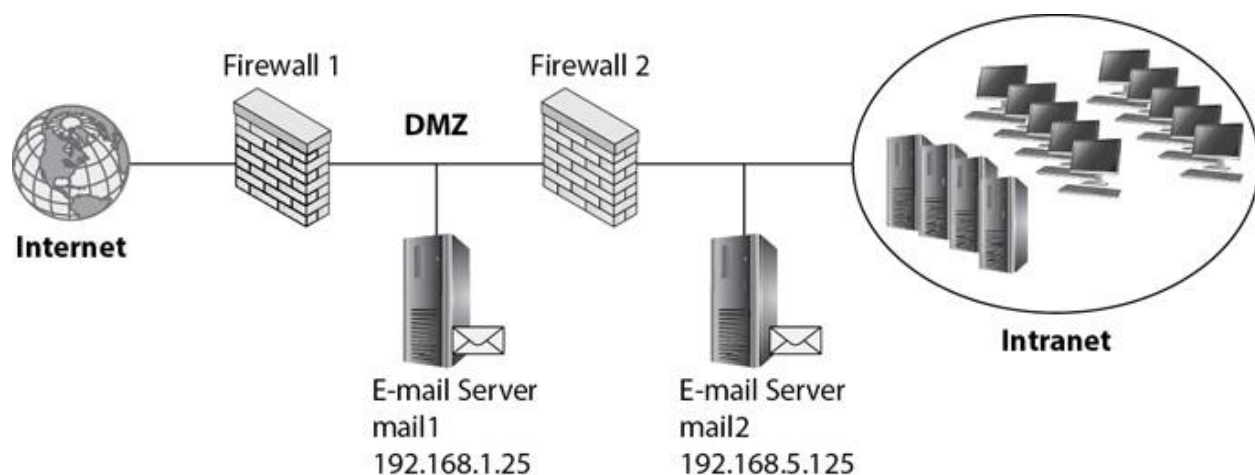


Figure 4-6 E-mail server within a DMZ

A packet-filtering firewall uses an implicit deny philosophy. All traffic is blocked (implicitly denied) unless there is a rule in the ACL that explicitly allows the traffic. Most firewalls automatically apply this rule, while some firewalls require administrators to add it. The implicit deny rule is placed at the end of the ACL and might look like one of the following statements:

- deny any
- deny any any

The actual syntax is dependent on the router brand and model, but the most important point is that it is the last rule in the ACL. If a previous rule doesn't explicitly allow the traffic, this rule blocks the traffic. If you place this rule first in the ACL, the ACL will ignore any other rules in the ACL and deny all traffic.



TIP Most routers and packet-filtering firewalls automatically use a deny any or deny any any rule even if the rule hasn't been typed into the ACL. In other words, you may not see the rule in the ACL, but the firewall is enforcing the rule.

Although a packet-filtering firewall is effective at blocking individual packets, it can be difficult to maintain when the ACL includes a large number of rules. Additionally, a large number of rules can negatively affect performance because the firewall must compare each packet against each rule.

Many organizations use a simple packet-filtering firewall at their boundary to the Internet. They do so by adding rules to the ACL of an external screening router. Although such a firewall doesn't provide all the capabilities of a more advanced firewall appliance, it can be effective for smaller organizations.

Stateful Inspection Firewall

The second generation of firewalls is stateful inspection firewalls. They include the same capabilities of a packet-filtering firewall but can also filter traffic by examining much more than individual packets. Instead, they monitor the complete session. A stateful inspection firewall identifies active connections and monitors the status of these connections in a state table within the firewall.

As an example, a TCP session begins with the three-way handshake. The stateful inspection firewall uses this three-way handshake to identify that a session has started between two systems. The firewall can then monitor the traffic between the systems. Even though UDP connections don't start with a three-way handshake, a stateful inspection firewall can still identify active UDP connections and track the state of the UDP connection.



TIP A stateful inspection firewall allows traffic that is part of an active connection or that is initiating a new connection. It rejects traffic that is not part of an active connection or that is not initiating a new connection. It also includes basic packet-filtering capabilities of a packet-filtering firewall.

A stateful inspection firewall is aware of packets exchanged in many communication sessions. When a session starts, the firewall dynamically opens ports based on the needs of the connections, and when the session completes, it dynamically closes ports.

In comparison, the packet-filtering firewall is a stateless firewall. It cannot determine whether a packet is part of an ongoing session between two systems or is from an attacker trying to probe or attack a system.

Application Firewall

An application firewall (also called an application proxy firewall, an application gateway firewall, and a third-generation firewall) includes different elements for different applications. For example, it would have a Hypertext Transfer Protocol (HTTP) proxy to examine web traffic, an SMTP proxy to examine e-mail traffic, a File Transfer Protocol (FTP) proxy to examine FTP traffic, and so on. Each proxy can analyze specific traffic for the protocol that it is monitoring and understands the individual protocol commands. For example, an HTTP proxy can interpret Get and Put commands used to retrieve and write web page elements, and administrators can configure the proxy to allow or block specific commands.



TIP An application firewall is also known as an application proxy firewall, an application gateway firewall, and a third-generation firewall.

Although the application firewall does provide effective filtering for individual applications or protocols (such as HTTP, SMTP, and FTP), it requires a separate component for each individual protocol. If the firewall must filter many different protocols, this process can be extensive and negatively affect performance. However, if the firewall is being used only to filter a few protocols, it is manageable and can be highly efficient.

Application firewalls are often used with packet filtering. For example, the packet-filtering firewall can drop packets for protocols that aren't desired based on the port, while the application proxy firewall can perform deep inspection of traffic for protocols that are allowed.

Next-Generation Firewall

Next-generation firewalls integrate multiple capabilities into a single security appliance, and these are often referred to as a unified threat management (UTM) device. Additionally, vendors update the UTM devices to adapt to new threats as they emerge.

Some of the additional capabilities that are included in UTM devices include

- **Malware blocking** UTM appliances typically include antivirus capabilities to detect and block malware before it enters the network. This includes malware included as attachments or downloaded from websites. Many UTM appliances can often detect and block malicious scripts embedded in web pages.
- **Spam filtering** Some UTM devices include spam filters to block unwanted e-mail. This can be especially useful because attackers commonly use e-mail to launch various attacks. [Chapter 5](#) discusses many common e-mail attacks such as phishing, spear phishing, and whaling.
- **URL filtering** UTM devices commonly include URL-filtering capabilities, similar to a proxy server. URL filters can block users from accessing external websites.

Defense Diversity

[Chapter 1](#) introduced the concept of defense in depth. A solid security program includes multiple layers of security so that even if one layer breaks down, other layers still provide protection. Similarly, defense diversity can be used when implementing a DMZ by using firewalls from two separate vendors.

For example, consider [Figure 4-7](#). The firewall separating the Internet from the DMZ is from Cisco and the firewall separating the DMZ from the intranet is from Microsoft. Because the DMZ was created by firewalls from two different vendors, it is using defense diversity. Both firewalls are susceptible to vulnerabilities and both are regularly updated as vulnerabilities become known.

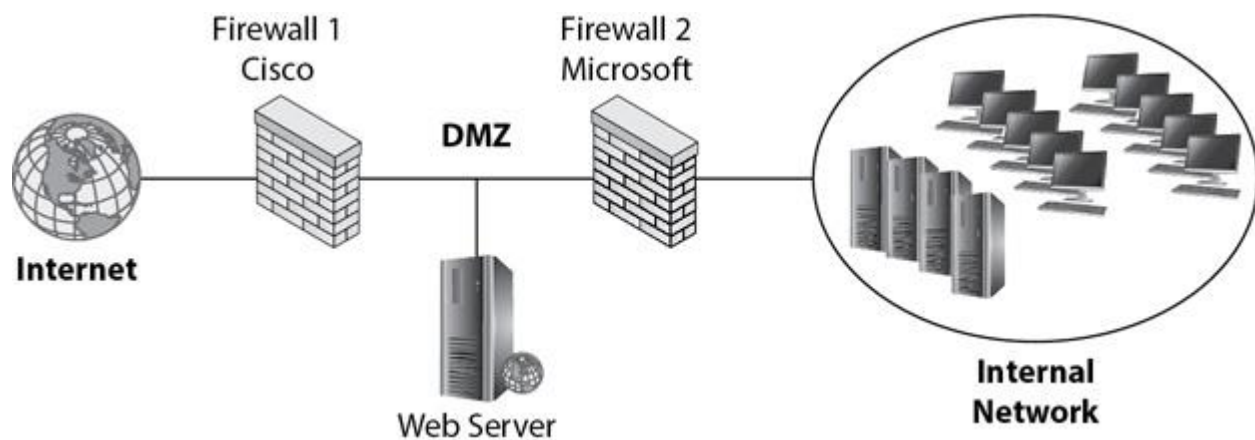


Figure 4-7 Implementing defense diversity with a DMZ

At any time, an attacker could discover a previously unknown vulnerability on either firewall. If both firewalls are the same model from the same vendor, both firewalls will be susceptible to the same vulnerability at the same time. However, if both firewalls are from different vendors, it is

less likely that both will be subject to a vulnerability at the same time. Defense diversity with the DMZ firewalls provides an extra layer of protection for the internal network.



EXAM TIP A DMZ is typically created with two firewalls. Using firewalls from two separate vendors provides defense diversity for the DMZ.

Additionally, defense diversity requires the attacker to have more skills and knowledge in order to exploit both firewalls. For example, an attacker may have extensive expertise with the Cisco router, but very little on the Microsoft router. Fewer attackers will have the same level of expertise on firewalls from different vendors.

Comparing Network-based and Host-based Firewalls

As previously mentioned, firewalls are often classified as either network-based or host-based. A network-based firewall provides protection for a network, while a host-based firewall provides protection for a single host. As an example, consider [Figure 4-8](#). The two firewalls used to create the DMZ are network-based firewalls and are filtering all traffic coming into or out of the network.

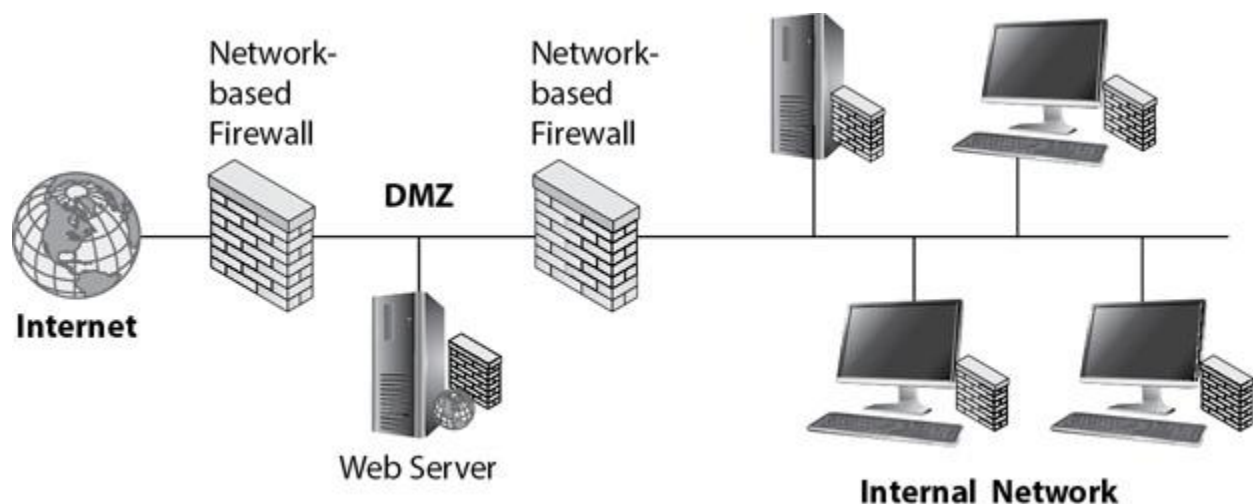


Figure 4-8 Network with host-based and network-based firewalls

Additionally, each internal system is running a host-based firewall that filters traffic coming into or out of the computer's network interface card (NIC). Both servers and desktop operating systems can run a host-based firewall as an additional application. As an example, many antivirus vendors now include firewalls as part of a comprehensive security suite.



TIP Some operating systems include built-in host-based firewalls running as additional software. For example, current Microsoft desktop and server products include the Windows Firewall, which is enabled by default.

Most organizations use both network-based and host-based firewalls. The network-based firewall filters and blocks malicious traffic coming from the Internet. Internal host-based firewalls block malicious traffic that reaches the internal network.

For example, imagine a user's home computer is infected with malware. When transferring files from home to work with a USB flash drive, the user can inadvertently infect the work computer just by inserting the USB flash drive into the work computer. This virus may then have a worm component that tries to locate and infect other systems on the network. However, if other systems have host-based firewalls installed, they can block the worm.

Exploring Remote Access Solutions

Remote access solutions allow users from remote locations to access internal network resources. Remote users can connect via the Internet to a virtual private network (VPN), as shown in [Figure 4-9](#). A VPN provides access to a private network over a public network such as the Internet.

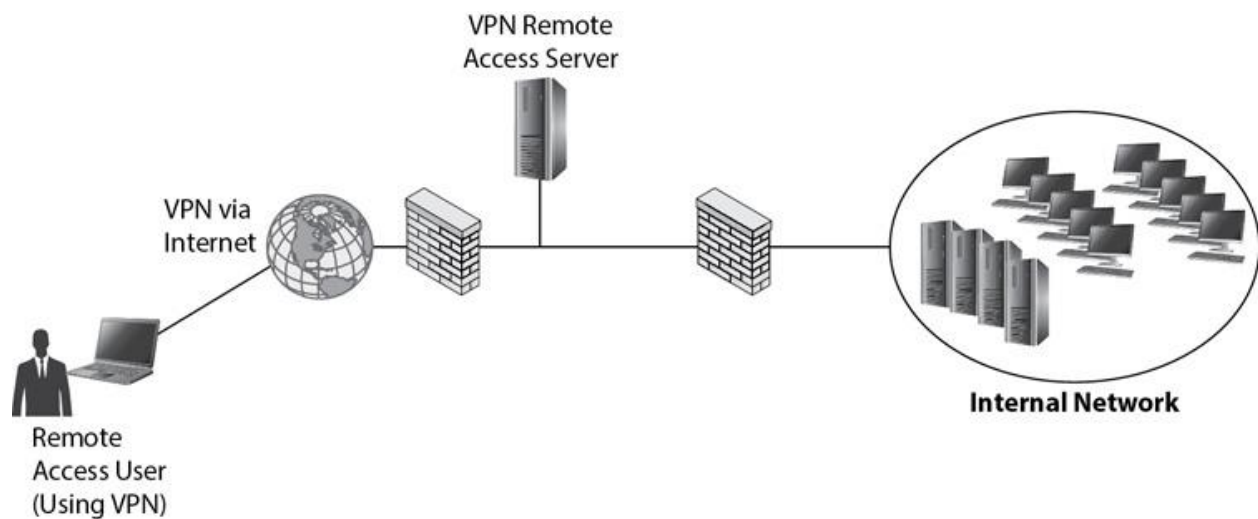


Figure 4-9 Remote access via dial-up or a VPN

It's also possible to connect to a remote access server using a dial-up connection, though the connection is much slower. Employees of organizations hosting remote access servers typically have broadband access to the Internet, so it is much more common for organizations to host VPNs. Many organizations allow employees to work from home. VPNs support telework

(sometimes called telecommute) practices, allowing employees to access resources within the corporate network from home.

Users with an Internet connection have a public IP address when they are connected to the Internet, typically assigned by their Internet service provider (ISP). The VPN remote access server also has a public IP address. The VPN client connects to the VPN server over the Internet, authenticates, and then accesses the private internal network.

Risks and Vulnerabilities

One of the primary risks of a remote access server is that anyone with access to a phone or the Internet can try to connect to the server. However, if strong authentication mechanisms are used, attackers cannot access the internal network because they lack the credentials to authenticate. Several different authentication protocols are available, and they are explored later in this chapter.

VPNs transfer data over the Internet in an encrypted format to prevent attackers from capturing the data. This includes authentication information, such as usernames and passwords, and the actual data. If the data isn't encrypted, attackers can use a sniffer to capture the data transmissions and read the data. VPN tunneling protocols use encryption to prevent unauthorized disclosure of credentials and data.

Due to the inherent risks of using VPNs, many organizations simply don't implement them. Instead, users perform all their work on the organization's network while they are physically at the organization's location.



NOTE Many organizations are recognizing mutual benefits when workers can work from home. For example, there are fewer costs for the worker for the commute, and in some instances, the company can reduce its workspace. A VPN provides a good way for workers to telecommute, but the VPN must be secure.

Tunneling Protocols

A tunneling protocol encapsulates data used on an internal network and transmits it over the public network. For example, an internal network may use the Server Message Block (SMB) protocol with TCP/IP to access files and printers within an internal network. However, SMB is not used on the Internet. The tunneling protocol encapsulates SMB traffic to transmit it over the Internet.



NOTE SMB is used with other protocols within the TCP/IP suite. The supporting TCP/IP protocol information used on the remote network is also encapsulated with SMB.

[Figure 4-10](#) shows the overall process. The client wants to transfer data using SMB. Before sending the data via the VPN, the client software encapsulates the SMB protocol data within a tunneling protocol. This encapsulated SMB data is then transmitted over the Internet to the VPN server. The VPN server strips off the tunneling protocol data, leaving only the SMB and TCP/IP data. The SMB data is then sent to the appropriate server within the internal network using TCP/IP. When the data is returned to the client, the VPN server again encapsulates the data in a tunneling protocol to send it back to the VPN client.

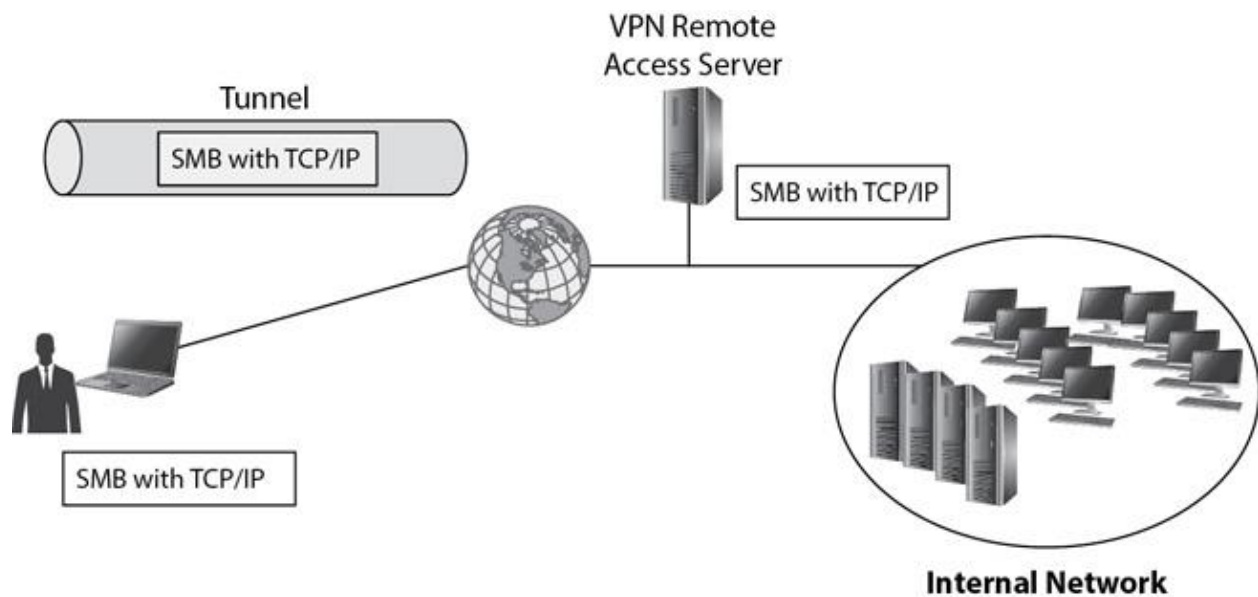


Figure 4-10 Encapsulating SMB data before sending it over the Internet

Although [Figure 4-10](#) may give you the impression that the SMB data is sent in plaintext, that is not the case. Instead, the tunneling protocol encrypts the data before sending it over the Internet. If attackers intercept any of the data, they will not be able to read it.

Some of the tunneling protocols that you may see on the SSCP exam to encapsulate VPN traffic are Secure Shell (SSH), Layer 2 Forwarding (L2F), Point-to-Point Tunneling Protocol (PPTP), Internet Protocol security (IPsec), Layer 2 Tunneling Protocol (L2TP), with or without IPsec, and Transport Layer Security (TLS). The following subsections briefly describe these tunneling protocols.

Secure Shell

The Secure Shell (SSH) protocol can be used to create an encrypted tunnel. Although the concept is similar to that of other tunneling protocols where the underlying protocol is encapsulated, SSH isn't commonly used for VPNs. It is used to encrypt other protocols such as Secure Copy (SCP) and Secure FTP (SFTP). Additionally, it's often used to connect to remote servers because all data in transit between the client and the server is encrypted.

Layer 2 Forwarding

Cisco created the Layer 2 Forwarding (L2F) protocol as an early VPN tunneling protocol. L2F provides an authentication mechanism, but does not encrypt the traffic. It is rarely used today.

Point-to-Point Tunneling Protocol

Point-to-Point Tunneling Protocol (PPTP) was much more popular than L2F. RFC 2637 defines the standard generically, and Microsoft developed a Microsoft version. The RFC version doesn't provide native encryption of the data, but the Microsoft version uses Microsoft Point-to-Point Encryption (MPPE) to encrypt the traffic in the tunnel. PPTP uses Generic Routing Encapsulation (GRE) to encapsulate the packets.

A significant weakness of PPTP is that the initial authentication connection is not encrypted. Information in the initial session can include usernames and passwords that can be intercepted by attackers. Even though the passwords are sent in a hashed form instead of in cleartext, attackers can capture them with a sniffer and use the usernames and password information in a replay attack. PPTP uses TCP port 1723.



TIP In a replay attack, the attacker captures a valid transmission, manipulates it in some way, and then resends it. If the attacker is able to capture reusable credential information, the attacker is able to impersonate the original sender but with different data.

Layer 2 Tunneling Protocol

Layer 2 Tunneling Protocol (L2TP) combines the strengths of L2F and PPTP and is documented in RFC 2661. Although L2TP doesn't encrypt the traffic itself, it is commonly combined with IPsec for encryption and authentication and referenced as L2TP/IPsec. L2TP/IPsec provides a high level of security for VPNs and has been widely implemented.

A primary drawback with L2TP/IPsec is that IPsec packets cannot pass through a system using NAT. Recall that NAT translates private IP addresses to public addresses and public IP addresses back to private. Because of how the translation is done, it corrupts the IPsec packets. L2TP uses UDP port 1701.



TIP NAT Traversal (NAT-T) was created to overcome the problem of NAT corrupting IPsec traffic. When equipment supports NAT-T, it's possible to send IPsec through a device using NAT-T.

Internet Protocol Security

[Chapter 3](#) introduces Internet Protocol security (IPsec) as a secure protocol that provides mutual authentication of systems, integrity, and confidentiality. Recall that it uses the Authentication Header (AH) for authentication and integrity, and the Encapsulating Security Protocol (ESP) for encryption. It can be used with L2TP (as L2TP/IPsec) or by itself as a tunneling protocol. IPsec can work in two different modes:

- **Tunnel mode** In tunnel mode, the entire packet is encrypted and encapsulated into an IP packet. When tunnel mode is used for a VPN, this encapsulated packet is then sent over the Internet. When it reaches its destination, the encapsulated packet is decrypted and used on the internal network.
- **Transport mode** In transport mode, only the payload (the actual data) is encrypted. In other words, the IP header is not encrypted. In contrast, tunnel mode encrypts the entire packet, including the IP header. Transport mode is used within an internal network to encrypt traffic from one point to another.

The National Institute of Standards and Technology (NIST) published Special Publication (SP) 800-77, *Guide to IPsec VPNs*. It provides significantly more details on using IPsec for VPNs. You can view a list of SP 800 documents here: <http://csrc.nist.gov/publications/PubsSPs.html>.

Transport Layer Security

Many VPN solutions use TLS for the VPNs. As you might recall from [Chapter 3](#), TLS is commonly used to encrypt HTTP (as HTTPS) over TCP port 443. A TLS VPN allows clients to connect with a web browser and it also uses TCP port 443.



NOTE TLS VPNs originally started as Secure Sockets Layer (SSL) VPNs, and you'll often see them marketed as SSL VPNs even though they support both TLS and SSL. However, TLS is the designated replacement for SSL, and SSL is not recommended for use due to known vulnerabilities.

A TLS VPN is a good solution when the VPN is having trouble getting through a NAT device, such as when using L2TP/IPsec without NAT-T. It's also becoming popular because it doesn't require any additional configuration on the client. The client only needs a web browser and knowledge of the URL to connect to the organization's TLS VPN.

NIST published SP 800-113, *Guide to SSL VPNs*, in 2008. It mentions that an SSL VPN works the same as a TLS VPN, but at the time it was written, "SSL VPN" was a more popular term. Still, SP 800-113 provides a good in-depth description of TLS VPNs.



EXAM TIP An SSL VPN uses port 443 and allows clients to access internal resources.

Another VPN option is to use an application instead of a web browser to connect to the VPN server. As an example, OpenVPN is an open source software application that can be used for VPNs. Users connect to the VPN server with the OpenVPN application instead of a web browser. Many desktop operating systems and mobile operating systems support the use of OpenVPN. For example, the OpenVPN Connect app is available for iPhones and iPads.

Full Tunnel and Split Tunnel Implementations

Tunneling protocols can be implemented as either full tunnels or split tunnels. This refers to how much of a user's traffic is sent through an encrypted tunnel. As an example, imagine that Sally connects to a VPN through her employer's organization to access internal resources. This VPN traffic is always encrypted in the VPN tunnel. Later, she goes to Google.com to do some research. Is her research traffic encrypted? It depends.

In a full tunnel, all traffic from the user is sent through the VPN tunnel and encrypted. In the example, Sally's connection to Google.com goes from her computer, through the VPN tunnel to the organization's VPN, and then back out to Google. The benefit is that all traffic can go through the organization's protection technologies. For example, a next-generation firewall can scan the traffic and block malware. Similarly, a proxy server can filter traffic going to potentially malicious websites.

A split tunnel only encrypts the traffic directly related to the VPN. In the scenario, Sally's connection to Google is not rerouted through the VPN tunnel, but instead goes directly to Google.

Authentication

One of the primary methods of controlling who can access a remote access server is authentication. Users claim an identity (such as with a username) and prove their identity (such as with a password) during an authentication process. Based on a user's proven identity, access controls are checked to see whether the user is authorized to use the remote access server. If a

user cannot be authenticated, the remote access server denies access. The following sections identify some common remote access authentication protocols.

PAP

The Password Authentication Protocol (PAP) is an older authentication protocol. Usernames and passwords are passed over the network in cleartext, which makes them highly susceptible to sniffing attacks. PAP is rarely used by itself. However, it is sometimes used with other protocols that encrypt PAP before transmission.

CHAP

The Challenge Handshake Authentication Protocol (CHAP) is more secure than PAP because the actual password is never sent over the wire. Instead, a *nonce* (a number used once) is combined with a shared secret known only to the client and the remote access server. The result of combining the nonce and the shared secret is then hashed with Message Digest 5 (MD5).

Hashing was briefly introduced in [Chapter 1](#). As you'll recall, a hashing algorithm (such as MD5) creates a number referred to as a hash. More specifically, MD5 creates a 128-bit number, or a 128-bit hash.

Putting it all together, here's a big-picture view of how the process works. The client sends an authentication request to the remote access server. The server responds by sending back the nonce. The client then combines the shared secret with the nonce and hashes the result with MD5. This MD5 hash is then sent back to the server. Because the server knows the nonce (it created it) and it knows the shared secret, it can also calculate the hash on the server. If the MD5 hash sent by the client matches the MD5 hash calculated on the server, then the server authenticates the client.

MS-CHAPv1 and MS-CHAPv2

Microsoft improved CHAP with its versions of Microsoft CHAP (MS-CHAPv1) and MS-CHAPv2. Although MS-CHAPv1 is no longer used, MS-CHAPv2 is still used in many Microsoft implementations.

A primary improvement of MS-CHAPv2 is mutual authentication. In other words, the client authenticates to the server and the server also authenticates back to the client. This helps prevent the client from sending authentication data to an attacker impersonating the remote access server.

EAP

The Extensible Authentication Protocol (EAP) provides a more secure authentication process than many other remote authentication methods. EAP doesn't use a specific method, but instead security developers can extend existing authentication methods using EAP. The following list identifies some of the commonly used EAP methods:

- **Protected EAP (PEAP)** PEAP encapsulates and encrypts the EAP transmission in a TLS tunnel. PEAP requires a certificate on the server, but clients are not required to have a certificate. A major advantage of PEAP is support by Microsoft systems and MS-CHAPv2 often uses PEAP.



TIP [Chapter 14](#) covers certificates in more depth, but as an introduction, certificates support strong authentication and encryption services. A certificate authority (CA) issues certificates, and systems can query the CA to verify the validity of a certificate.

- **EAP-Tunneled TLS (EAP-TTLS)** EAP-TTLS also uses TLS to create a tunnel to encapsulate the authentication method. It exchanges authentication information as attribute-value pairs (AVPs), allowing it to secure many older authentication methods (such as PAP, CHAP, and MS-CHAP). Note that while PAP sends the authentication information in cleartext, the TLS tunnel ensures that the authentication information is encrypted while in transit. EAP-TTLS requires a certificate on the authentication server, but clients do not need a certificate.
- **EAP-TLS** EAP-TLS provides the best security of these methods. It requires the server and each client to have a certificate, providing mutual authentication.

RADIUS

The Remote Authentication Dial-in User Service (RADIUS) provides centralized authentication, authorization, and accounting (AAA) for remote clients. When a remote user attempts to connect, the remote access server passes the request and user credentials to a RADIUS server. The RADIUS server verifies the user's credentials and verifies the user is authorized to connect to the remote access server. After the user connects, the RADIUS server tracks the user activity in an accounting log.

Even though it has “dial-in” in the name, RADIUS is used for remote clients that connect through other sources, such as through a VPN.



EXAM TIP RADIUS is not only for dial-up. It also provides AAA services for VPN remote access connections. Other implementations include configuring RADIUS as an 802.1x authentication server to require authentication for wireless clients, commonly known as WPA2-Enterprise.

Although a RADIUS server can be used for a single VPN server, it is more commonly used with multiple remote access servers. For example, [Figure 4-11](#) shows several VPN servers hosted by

an organization in different cities. Remote access users can connect to any of the servers over the Internet and then have access to the network physically located in that city.

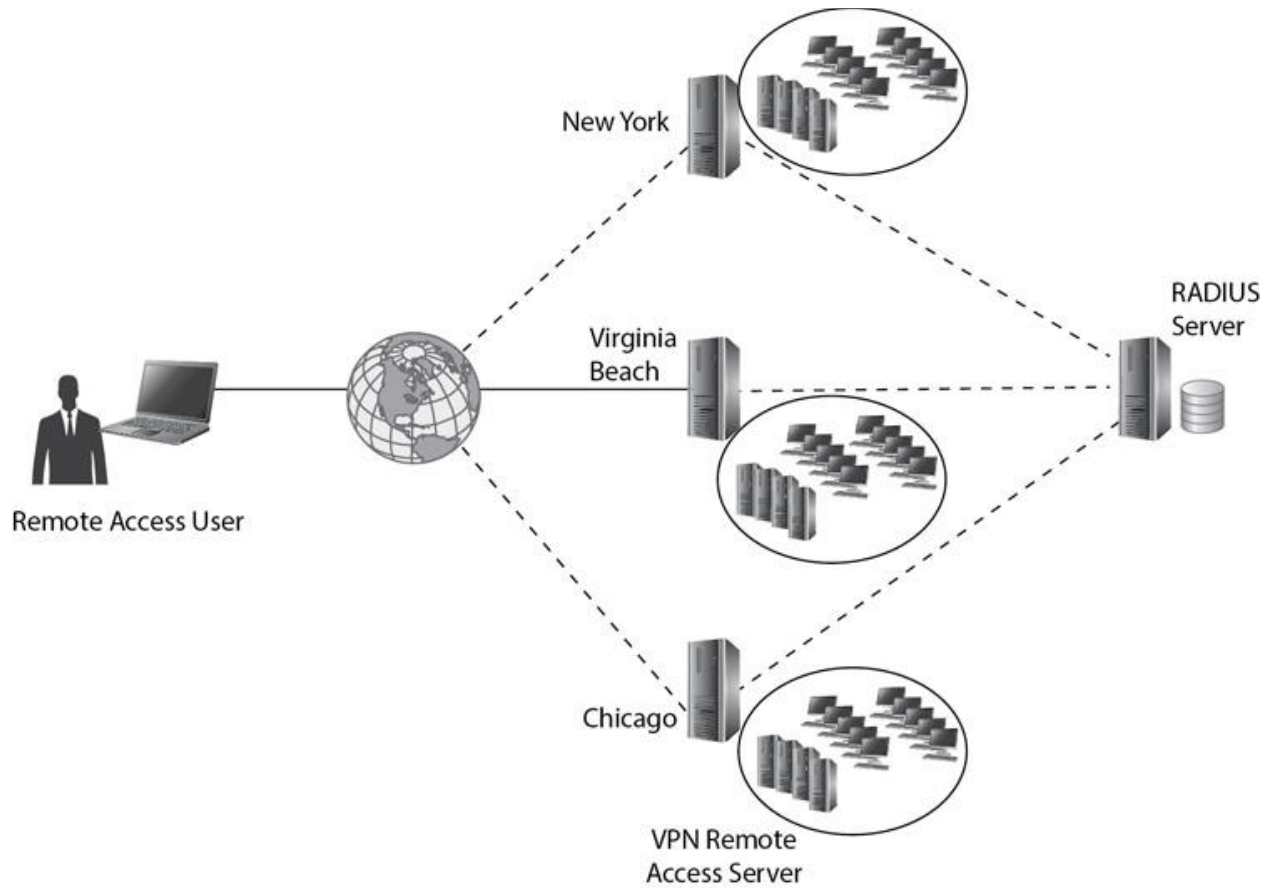


Figure 4-11 Using RADIUS for authentication

In the figure, a remote access user is connecting to the Virginia Beach VPN server over the Internet. The user passes authentication credentials to the VPN server and the VPN server then passes this authentication information to the RADIUS server.



TIP When a VPN server is using a RADIUS server for authentication, it is called a RADIUS client. In this context, the VPN server acts as the server for the VPN client but as a client for the RADIUS server.

The RADIUS server might have an authentication database with the credential information (such as a database of usernames and passwords) of all authorized users. It's also possible for the RADIUS server to check authentication with another server. For example, in a Microsoft network, the RADIUS server passes the authentication information to the domain controller. If

the credentials are verified, the information is passed back to the remote access server and the client is authenticated. If a user connects to one of the other VPN servers (such as in New York), the VPN server passes the authentication information to the RADIUS server and the rest of the process is the same.

A primary benefit is that the authentication is centralized and only a single authentication database needs to be maintained. Imagine for a moment that an organization uses VPNs to allow administrators to remote in and perform administrative work. Administrators can remote into any of ten different VPN servers located in offices around the country. When a new administrator is added or removed from the administrative team, only one entry is needed on the RADIUS server. In contrast, if RADIUS were not used for central authentication, each of the ten different VPN servers would have to be modified.

Historically, RADIUS has used UDP and it only encrypts the password during the authentication session. However, RFC 6613 provides specifications for using TCP and RFC 6614 provides specifications for using TLS to encrypt the entire authentication session. As I write this, both RFCs are still in the experimental stage.

Diameter

Diameter is another AAA protocol and it was created as an alternative to RADIUS. It isn't backward compatible with RADIUS, but it does provide several improvements. It uses TCP instead of UDP and supports both IPsec and TLS to encrypt sessions.

The name Diameter implies that it is twice as good as RADIUS. The radius of a circle is the distance from the center to an edge of the circle. The diameter of a circle is the distance from one edge of the circle to the other edge of the circle. In other words, the diameter of a circle is twice the distance of the radius.

TACACS+

Terminal Access Controller Access Control System+ (TACACS+) is another AAA protocol and is an alternative to RADIUS. Cisco originally created it as a proprietary protocol but later released it as an open standard. TACACS+ uses TCP so communication is more reliable than RADIUS transmissions. Also, TACACS+ encrypts the entire authentication process, while RADIUS only encrypts the password. TACACS+ uses TCP port 49.



EXAM TIP TACACS+ encrypts the entire authentication session. In contrast, RADIUS encrypts only the user's password.

Traffic Shaping and WAN Optimization

Traffic shaping refers to delaying certain types of traffic so that other traffic performs better. For example, audio and video streaming can consume a significant amount of bandwidth, so a traffic shaping policy would delay this type of traffic, giving preference to other traffic. Note that traffic shaping is not the same as packet filtering. Packet filtering drops undesirable traffic. Traffic shaping only delays it. The result is that other traffic is given priority.

Routers, firewalls, and dedicated hardware appliances are often used as traffic shaping devices. As an example of a dedicated appliance, it's possible to place a Linux server with two network interface cards (NICs) behind a border firewall. All traffic would go through the appliance and software on the server can perform bandwidth management for different types of traffic.

Wide area networks (WANs) often use VPNs to connect, and traffic shaping is useful to optimize the WAN traffic. As an example, imagine that a remote office connects to the main location for a business via a site-to-site VPN as shown in [Figure 4-12](#). Instead of connecting over the Internet, the locations are connecting via a leased T-1 at a speed of 1.544 Mbps. Compared to internal speeds of 1 Gbps or more, 1.544 Mbps is quite slow. Worse, the performance of this link can become even slower if users are regularly using it for streaming audio and video.

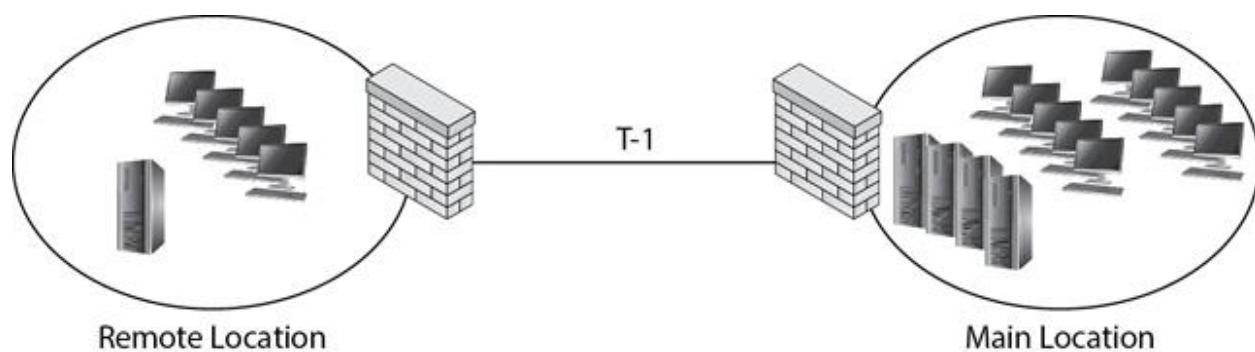


Figure 4-12 Site-to-site VPN

Traffic shaping will optimize the WAN link by delaying undesirable traffic. Note that this example focused on streaming audio and video, but an organization can choose to delay any type of traffic based on its needs.

Managing Network Access Control

Network access control (NAC) is a technology used to control which clients are granted access to a network. This is especially important when remote access solutions allow external clients into the private network. However, NAC can also control unhealthy clients within the network.

The overall goal of NAC is to inspect clients regularly to ensure they meet specific *health* requirements. Health is defined by criteria such as having up-to-date antivirus software installed and operational, having a host-based firewall enabled, and keeping the operating system up to date. These elements are defined in a health policy specified by administrators.

When NAC identifies an unhealthy client, it typically quarantines the client to prevent it from infecting other computers on the network. A quarantined network would also include resources that clients can use to improve their health, such as updates for the antivirus software or operating system.

When computers are completely controlled within an organization, administrators have several different tools and methods they can use to ensure that the systems are secure. They can automate the deployment of antivirus software and updates to the clients, automate the deployment of operating system updates, control what software is installed and running on the systems, and even control what the user can do on the system, such as what software the user can install.

However, remote access users may use a home computer over a VPN connection to access the private network. Because it's a home computer, the organization's administrators don't have direct control over the computer. A significant risk is related to malware. If the infected system connects to the internal network, it can easily infect other computers on the network.

[Figure 4-13](#) shows the overall process of how a NAC solution works. After a remote access user connects (via either dial-up or VPN connection), the remote access server coordinates with the NAC server to determine the health policy for the organization. The client is inspected to determine whether it meets the criteria of the health policy. If the client is healthy, it is granted access to the network. If the client is not healthy, the NAC server restricts access for the client to the quarantined network only.

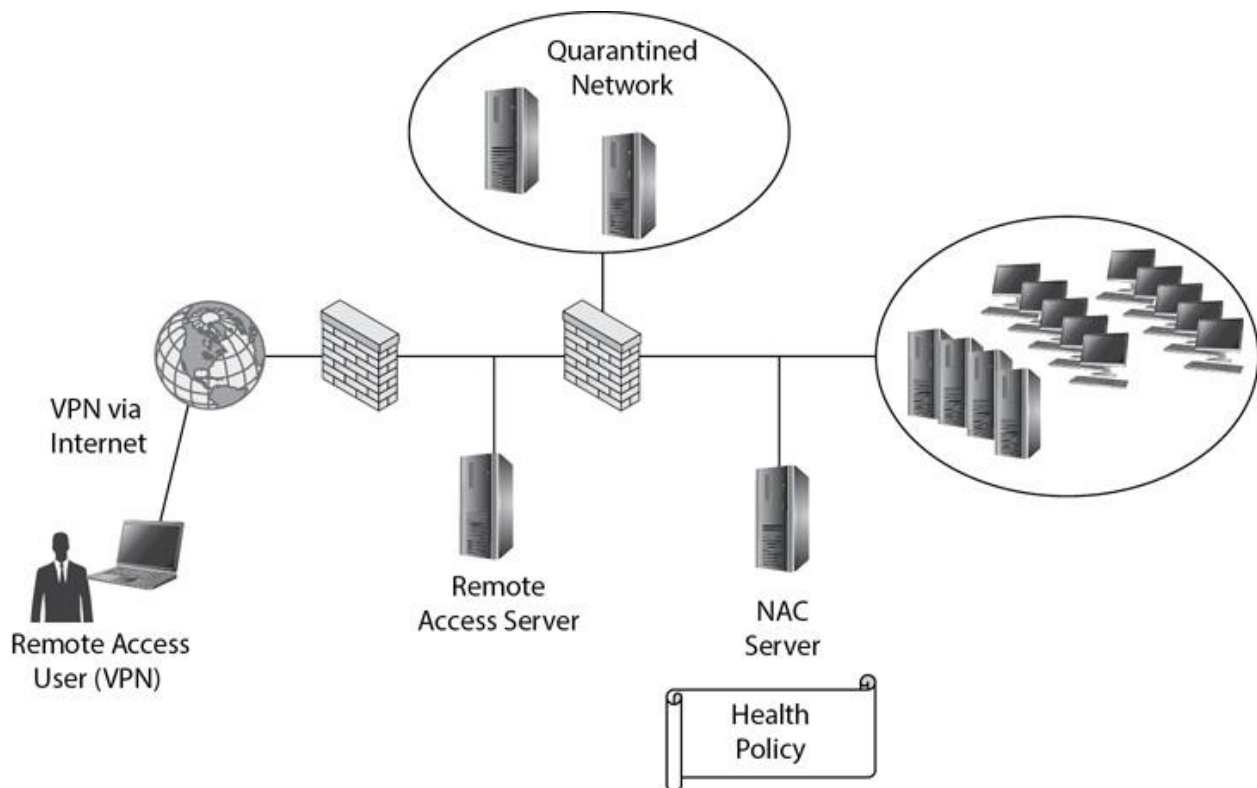


Figure 4-13 Using network access control

Quarantined clients have access to some resources that can help them remediate any issues that prevent them from accessing the network. As an example, the quarantined network can include patches to bring clients up-to-date or current antivirus signatures. After a client is remediated, the NAC system will grant it admission back into the network.

Although NAC is often used for remote clients, it can also be used for internal clients. This helps prevent unhealthy systems from staying operational on the internal network. As an example, many automated methods exist to deploy operating system updates to clients. However, there are multiple reasons that a critical update may not be installed on a system, such as if the system is powered off while a user is on an extended vacation, business trip, or leave of absence. When the user turns the computer back on and tries to access the network, NAC checks its health and quarantines it if it is not healthy.

Additionally, many mobile device management (MDM) systems can interact with an existing NAC system, or provide NAC services. An MDM system can identify approved mobile devices and inspect them to ensure they meet preset requirements such as being up to date. Unapproved devices can be redirected to a guest network. Approved devices that don't meet requirements can be redirected to the quarantined network. Healthy, approved devices can be given full access to the network.

Exploring Virtual Environments

Virtualization refers to replacing hardware with software. One of the first ways organizations did this was with servers. For example, imagine an organization has 100 servers within a data center. The organization can replace these servers with 20 powerful servers. Each of these physical servers can host five virtual machines (VMs) so the organization still hosts 100 servers.

While the cost of the 20 servers is high, an organization will quickly recoup its costs through other savings. For example, these 20 servers require less power; less heating, ventilation, and air conditioning (HVAC); and less space. Additionally, the organization reduces the hardware costs in future hardware refresh cycles. Instead of replacing 100 computers every three to five years, the organization replaces only 20 servers in the same timeframe.



NOTE Gartner reported in 2016 that more than 75 percent of x86-based servers were virtualized. Organizations primarily use VMware virtualization products, but Microsoft has been steadily gaining market share in this multibillion-dollar market. Some other virtualization providers include Citrix, Oracle, and Red Hat.

In addition to creating VMs, it's possible to create virtual networks within a host machine. These aren't the same as VLANs created on switches. Instead, the host machine allows you to create

and configure virtual switches and virtual routers running on the host machine. Administrators configure these virtual network devices to connect the host machines together and connect them to the live network.

Virtualization Terminology

Three important terms related to virtualization are host, guest, and hypervisor. The physical server is the host machine. Each VM running on the host machine is a guest and a host can run multiple guest VMs. The host runs hypervisor software such as Microsoft Hyper-V or VMware vSphere. The hypervisor is the software on the physical server that creates, runs, and manages the VMs. It also manages the resources for each of the VMs. These resources include the virtual data storage, memory, and networking components.

Because a host runs multiple guest VMs, it requires more resources than a typical server does. These physical resources include multiple processors, a significant amount of RAM, and fast drives. The host system shares the use of these resources with all the VMs.

Shared Storage

Guest VMs on host machines often share the same data storage, so it's important to ensure it has enough space and is fast. At a minimum, it should use a fast Redundant Array of Independent Disks (RAID) subsystem. Another option is to use solid-state drives (SSDs), which are much faster than traditional hard disk drives.

A security concern with using shared storage is confidentiality. You want to ensure that users within one VM are not able to access data in another VM. Strong access controls will typically provide enough protection to prevent unauthorized access.

[Chapter 11](#) covers data classifications in more depth. For example, the U.S. government classifies data as Top Secret, Secret, Confidential, and Unclassified. With this in mind, it's best to put guest machines with similar data classifications on the same host system. In other words, you wouldn't put a guest machine hosting Top Secret data on the same system hosting only Unclassified data.

Virtual Appliances

A virtual appliance is a preconfigured VM with an operating system and an application. The VM is ready to be added to a host machine as a guest. The benefit of a virtual appliance is that it is just a set of files that administrators can copy or import to a host machine. Of course, these virtual appliances are created for specific hypervisors. For example, a VM created for a Microsoft Hyper-V hypervisor will not run on a host machine using a VMware vSphere hypervisor.



NOTE Typically, appliance refers to a hardware device used for specialized tasks. For example, a UTM appliance is a hardware device that includes several security features on a single device. In contrast, a virtual appliance is software running as a guest VM on a host server.

As an example, Openfiler has created virtual appliances for file-based Network Attached Storage and block-based Storage Area Networking (NAS/SAN). The NAS features include support for the Common Internet File System (CIFS) and Network File System (NFS). The SAN features include support for Internet Small Computer System Interface (iSCSI) and Fibre Channel (FC). Openfiler has created virtual appliance images for a free version, and a full-featured Commercial Edition.

While Openfiler is one example, there is almost no limit to what vendors can create as a virtual appliance. This has caused many organizations to look at servers and virtualization differently. Instead of evaluating how many servers they need, they evaluate how many applications they want to host as VMs.

Virtual Desktop Infrastructure

A virtual desktop infrastructure (VDI) provides a virtual desktop for users. A user can log on to a server hosting VDI images and access a full-featured desktop including an operating system and all required applications. Some VDI-based systems support mobile clients. This allows clients to log on with a tablet or smartphone and access a specialized desktop for the mobile device.

When implementing VDIs, administrators need to decide if the desktops will be nonpersistent or persistent:

- **Nonpersistent** A nonpersistent desktop is the same for all users. If a user logs on and makes changes to the desktop, such as modifying the wallpaper, the changes aren't saved. Additionally, data saved by a user in a user profile is not saved, so users should save their data elsewhere, such as in a network share. When a user logs out, the desktop is reverted to a known state. The benefit is that less disk space is required on the server. Nonpersistent VDIs are also easier to maintain because they are all the same.
- **Persistent** A persistent desktop retains all changes made by a user. Each time a user logs on to a VDI system and makes changes, the changes are saved and retained. A drawback is that persistent VDIs require more disk space for every user that logs on to the server.

Continuity and Resilience

Virtualization also contributes to continuity and resilience, increasing overall availability. Many cloud providers include elasticity features with virtual servers that they rent. For example,

Amazon Elastic Compute Cloud (Amazon EC2) makes heavy use of virtualization. When customers rent cloud services through Amazon EC2, they rent access to virtual systems.

One of the benefits of elasticity is that capacity can easily shrink or expand based on the demand. In this way, users pay only for what they're using. For example, if a customer's website suddenly has a burst of activity and needs additional processing power or memory, it's relatively easy for EC2 to expand the available resources for the site. When the additional resources are no longer needed, the virtual system can shrink back to the original size.

Many hypervisors allow administrators to manage the guest machines using a point-and-click interface. This allows them to easily create backups of an entire VM with just a few clicks. Additionally, some hypervisor management tools allow administrators to move VMs between host systems via a central server. As an example, imagine one host system is running five VMs but increased usage is overloading the host system. Administrators can use a point-and-click interface to move one or more of the VMs to another host, freeing up resources on the original host system.

It's also much easier to restore a VM if it crashes, when compared to a physical server, assuming administrators keep up-to-date backups of the VMs. If a VM crashes, administrators can restore it by restoring the files. In contrast, if a physical server crashes, it takes much more to rebuild it from scratch.

Similarly, most hypervisors allow you to create snapshots of a system. These are very useful, especially before performing a risky operation such as an application upgrade. If the upgrade causes problems, you can revert the system to the state it was in when you created the snapshot.

Separation of Data Plane and Control Plane

Traditional hardware routers implement both the data plane and the control plane. A current trend is to replace hardware routers with software-defined networking (SDN), which separates these two planes. Before discussing SDN concepts, it's worth defining the data plane and the control plane:

- **Data plane** The data plane (sometimes called the forwarding plane) uses rules within an ACL to identify whether the router will pass traffic or block traffic. These are the same type of rules described in the "Packet-Filtering Firewall" section earlier in this chapter. The data plane is responsible for deciding if traffic should be forwarded or not. The data plane is proprietary. In other words, rules used by one router often have different syntax when compared to rules created by another router manufacturer.
- **Control plane** The control plane is responsible for identifying paths to other networks. Routers typically do this with routing protocols such as OSPF and BGP. These routing protocols allow routers to share information with each other to create a virtual map of the network.

Software-defined Networking

Software-defined networks separate the data and control planes using virtualization technologies. This allows an organization to replace many of the hardware routers and switches with software. This helps organizations move away from proprietary hardware and also gives them much more flexibility.

As an example, [Chapter 2](#) discusses Attribute-based Access Control (ABAC) models, which SDNs typically support. An ABAC model allows administrators to create policies that are much more versatile than rules within an ACL.

While SDN and network function virtualization (NFV) are relatively new technologies, they are growing fast. The 2015 Edition of the SDxCentral *SDN and NFV Market Size Report* includes several predictions of explosive growth. For example, it indicates sales of SDN products replacing traditional layer 2 and layer 3 hardware (switches and routers) will exceed \$10 billion in 2015 and grow to more than \$105 billion in 2020. For comparison, SDxCentral reported the SDN market size in 2013 was about \$1 billion.

The section “Physical and Logical Segmentation” earlier in this chapter discussed how routers and switches provide physical segmentation of networks. Additionally, it mentioned how switches can be used to provide logical segmentation within a network. SDNs also provide logical segmentation by separating the data and control planes.



EXAM TIP Routers and switches provide physical segmentation of networks. VLANs (created with switches) provide logical segmentation of networks. SDNs separate the data and control planes to provide logical segmentation of networks.

Protecting Virtualization Systems

As additional software running on the server, the hypervisor represents another potential risk. If an attacker can access the host machine from within a VM, the attacker potentially can access all the guest machines. The following list identifies some of the specific risks associated with virtualization, along with steps to protect against these risks:

- **Improperly configured hypervisor** The security of all the guest systems is directly related to the security of the hypervisor on the host. If administrators don't adequately restrict access, it becomes easier for unauthorized personnel to gain access to all the systems. The key here is to ensure that administrators have appropriate training to manage and secure the hypervisors.
- **Infected virtual appliances** Free virtual appliances sound very appealing. However, it's possible for attackers to infect these with Trojans or other malware, and once you add the virtual appliance, attackers have a foothold in your environment. This is similar to the many new

computers built in foreign countries that have been shipped to users with malware preinstalled. Just as a single infected physical machine in your network can spread malware to other systems, a single infected VM can also spread malware. A key here is to be suspicious of any files downloaded from the Internet and to only use trusted files from a trusted source.

- **Data leakage** Because VMs are simply files, it's very easy to copy an entire VM just by exporting and copying the files. If unauthorized personnel have access to these files, they might be able to import them to their own system running a hypervisor and access all the data within the VM. Additionally, attackers might be able to scan the data within the files and extract meaningful data. These VM files should be protected with the same level of security as the highest classification of data stored on the VM.

- **Improperly configured virtual network** Most network devices are configured by administrators with specialized knowledge about routers and switches. However, server administrators typically configure virtual network devices (such as virtual routers and virtual switches) on the host. They might not have the same level of knowledge and expertise and might configure them improperly. An improperly configured virtual network might incorrectly direct sensitive traffic to guest VMs on the host system.

- **VM escape** VM escape is an exploit that allows a guest to interact with the host system. As an example, contestants in the 2017 Pwn2Own hacking competition in Vancouver, Canada won \$105,000 by successfully performing a VM escape attack. They escaped from a Windows 10 guest operating system running as a VM in VMware Workstation. They exploited a bug in the Microsoft Edge web browser, a bug in the Windows 10 kernel, and a buffer vulnerability in VMware Workstation. They achieved a complete VM escape gaining access to the host system. Microsoft and VMware have patched these bugs, preventing the same type of exploit again. Keeping systems patched is one of the best ways to prevent successful VM escape attacks.



EXAM TIP VM escape is a known attack against virtual systems. If successful, an attacker can access the host system and all virtual systems within the host.

Understanding Cloud Computing

Cloud computing refers to any type of computing services provided over the Internet. It has become quite popular in recent years. Cloud computing is very cheap and sometimes even free, but it includes some security concerns.

Shared Responsibility Models

Three common services provided through the cloud are Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS). The lines between these cloud services are sometimes blurred, partly due to how cloud vendors market them. However, one

way to categorize them is to look at the different maintenance responsibilities, which correlate with security responsibilities. The *Cloud Computing Security Requirements Guide*, published by the U.S. Department of Defense (DoD), includes a figure similar to [Figure 4-14](#). [Figure 4-14](#) shows that the vendor has the most maintenance and security responsibility for SaaS, a little less for PaaS, and the least for IaaS.

| | | |
|------|--------|----------|
| SaaS | Vendor | Customer |
| PaaS | Vendor | Customer |
| IaaS | Vendor | Customer |

Figure 4-14 Maintenance and security responsibility



EXAM TIP Cloud computing includes Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS), and it’s important to know the difference between these three services. Customers have the most maintenance and security responsibility for IaaS and the least responsibility for SaaS.

SaaS

Software-as-a-Service (SaaS) is also known as on-demand software. It provides users with access to software or applications over the Internet. For example, Google provides several applications through the cloud, including Gmail for e-mail and Google Docs for other types of files, such as spreadsheets, presentations, and word processing documents. Users access the applications using a web browser. This allows the cloud-based applications to support different operating systems.

Domain Name System (DNS) also fits into the SaaS service model. [Chapter 3](#) discusses DNS and how it resolves host names to IP addresses. For example, if you want to reach Google, you type google.com as the address in your web browser, and your system queries DNS to return an IP address. Organizations that maintain resources on the Internet need to ensure that they register and maintain records for these resources. They could create and maintain their own DNS server, but it’s often easier to outsource this and use a cloud provider to manage the records on the cloud provider’s DNS server. A huge benefit is that the cloud provider manages the security of the DNS server and often has more knowledge and expertise to thwart DNS attacks.

The “Exploring Remote Access Solutions” section earlier in this chapter discussed the use of VPNs for remote access. However, VPNs can also be used in a SaaS model. Users connect to the VPN service provider and surf the Internet through it. The VPN encrypts all traffic and prevents attackers from capturing the traffic. This is especially useful for wireless users accessing Internet resources through public hotspots.

Similarly, a proxy server is typically used within an organization for caching and web filtering. However, you can also subscribe to proxy services from some cloud providers. Administrators configure the proxy settings within the web browsers, which causes all Internet access to go through the cloud-based proxy server. The proxy server performs web filtering just as an in-house proxy server does.

PaaS

Platform-as-a-Service (PaaS) provides users with a computing platform. This includes hardware, an operation system, and applications. As an example, many web-hosting companies rent access to servers that are preconfigured with an operating system and web hosting applications, such as WordPress. Customers may install the applications from a list of available options, or the vendor may install the applications. The vendor performs all the hardware maintenance on the server, and typically performs most of the software maintenance. For example, the vendor would keep the server up to date with current patches. The customer is responsible for keeping applications up to date, but the vendor may help.

IaaS

Infrastructure-as-a-Service (IaaS) is sometimes called Hardware-as-a-Service, and this is a good way to think of it. Customers rent access to hardware, such as servers and networking infrastructure, and the vendor maintains the hardware. The customer is responsible for installing operating systems and applications on the hardware, and is responsible for maintenance of the operating systems and applications.

Cloud Operation Models

When considering security with cloud-based computing, it's also worth examining the different operation models, or deployment models. NIST SP 80-144, *Guidelines on Security and Privacy in Public Cloud Computing*, defines the following four models:

- **Public cloud** Public cloud-based services are provided by third-party vendors and are available to anyone. As an example, Apple provides cloud-based storage via its iCloud service. Google hosts the Google Cloud Platform, which is a suite of cloud computing services developers use to create and serve cloud applications. Similarly, Amazon hosts Amazon Web Services (AWS), which provides a wide variety of cloud computing platforms.
- **Private cloud** A private cloud is available only to users within an organization. As an example, a company can set up cloud-based storage for its employees. The company would ensure the server is accessible via the Internet, but employees would be the only personnel authorized to access the server.
- **Community cloud** A community cloud is a private cloud that is shared by two or more organizations. For example, imagine that Globex set up a private cloud for Globex employees. Later, Globex and Acme sign a partnership deal and need an easy method to share files. Globex

can configure its private cloud so that Acme employees can access it. It is now a community cloud.

- **Hybrid cloud** A hybrid cloud is a combination of any two or more clouds. It can be a combination of different clouds or similar clouds. For example, imagine that both Globex and Acme host their own private clouds. Later, they sign a partnership deal with each other, and Acme shares part of its private cloud with Globex, creating a community cloud. The Globex employees will be able to access a single cloud-based service and have access to both the private Globex cloud and the Acme community cloud.

A private cloud provides the best security for sensitive data. However, it's worth remembering that a private cloud is still accessible via the Internet. If users don't use secure passwords, or if they reuse their credentials on other sites, attackers may be able to access the data in the cloud. While a community cloud is similar to a private cloud in that it is private between the companies sharing the data, it's important to ensure that only shared data is stored in the community cloud.

Storage

Storage in the cloud is almost limitless. An organization can easily rent the storage space it needs, and that space is often very cheap or even free. For example, Google Drive provides every user with 15GB of storage space for free. If you need more, you can rent 100GB for \$1.99 a month. Microsoft OneDrive offers 5GB of storage for free and 5TB of storage for Office 365 subscribers. These are moving targets though, and the pricing might be different by the time you read this.

However, free storage isn't necessarily safe storage. Organizations need to be careful about what they store in the cloud. A data breach with the cloud provider can result in a data breach for an organization.

When picking a cloud storage location, it's sometimes useful to consider the legal jurisdiction of the storage location. As an extreme example, imagine the cloud provider accesses an organization's data and sells it to a competitor. If the cloud provider is located in the same country, the organization can pursue a legal remedy. However, if the organization is in the United States and the cloud provider is in Russia, the organization may not have any legal options. Similarly, if a U.S. cloud provider outsources storage locations to other countries, sensitive data might wind up being stored in unfriendly locations.

Electronic discovery (e-discovery) refers to any process that searches data with the goal of using it as evidence. It also includes locating and securing the data so that it can be used in a civil or criminal case. Using a cloud provider from a local jurisdiction ensures that e-discovery processes can be used when necessary.

Privacy

Privacy of data posted on the Internet has always been a concern. A good guideline to adopt is that if you post data on the Internet, you have lost control of that data and might have also lost privacy of that data. This includes data that you store in the cloud.

Cloud computing attempts to provide users with privacy of their own data, but any cloud provider is vulnerable to attacks, potentially resulting in data loss or data leakage. When an organization hosts data on the Internet through a cloud provider, protection of that data is now partially dependent on the provider using adequate protection. If the provider has a failure, it can result in a data breach. Additionally, users often don't know about the severity of a data breach right away.

For example, Dropbox, a cloud-based storage provider, suffered a data breach in mid-2012. It was unclear at the time how many user accounts were compromised. Four years later (2016), attackers posted details of over 65 million accounts from the 2012 data breach. Admittedly, users should have changed their passwords over the four-year timeframe, but all users aren't security conscious. About a week before these details were made public, Dropbox forced all users to reset their passwords in 2016, which prevented anyone from using the hacked data after the reset.

Cloud customers also have a responsibility for security, and poor security practices by customers can lead to data breaches. In 2014, attackers accessed data on Dropbox again, but this was due to data breaches from other organizations and weak security practices of customers. Specifically, customers were using the same username and password on Dropbox that they used with other companies. When attackers gained access to user credentials, they methodically tried them on different services such as Dropbox. Similarly, after the 65 million Dropbox accounts were published in 2016, attackers could use these same credentials to access accounts on other cloud providers if customers used the same credentials on other sites.

Encryption is an excellent method of providing confidentiality. If you are using encryption to protect your data in the cloud, the data should be encrypted on the client end and not in the cloud. In other words, if you want to store sensitive data in the cloud, encrypt it before posting it to the cloud. If you use encryption tools made available by the cloud provider, it's very possible that the provider holds the decryption keys and employees of the provider can access the data or may inadvertently decrypt it, making it available to other users.

Data Control and Third-party Outsourcing

Organizations that store data in the cloud are effectively outsourcing some of the security requirements. With this in mind, an organization should consider implementing some security controls to increase the control of the data.

A service level agreement (SLA) is an agreement between an organization and a vendor. Organizations use SLAs when contracting services from service providers, including cloud service providers. An SLA typically stipulates performance expectations, such as minimum

uptime, and can include expectations for the availability and security of data. The SLA can mandate certain requirements such as the following:

- **Data portability** Data portability allows users to reuse data across interoperable applications. Cloud providers can implement different application programming interfaces (APIs) to support data portability.
- **Data destruction** When the data is no longer stored with the cloud provider, it's important that it is destroyed. [Chapter 11](#) discusses various methods of adequately destroying data, and it's important to ensure that the cloud provider uses similar methods to eliminate any data remnants of sensitive data.
- **Auditing** An organization owns the data, and this ownership doesn't change when the organization chooses to store the data with a cloud provider. Auditing processes can help ensure that employees of the cloud provider are not accessing the data inappropriately.
- **Data resilience** The cloud provider needs to implement appropriate security controls to ensure the resilience of the data. This includes appropriate backup and recovery processes to overcome any potential system failures.

Compliance

Organizations that must comply with laws and regulations need to be careful when using the cloud, especially for data storage. The security provided via the cloud may not be adequate to protect some data, but this doesn't relieve the organization of its legal responsibilities. For example, the U.S. Health Insurance Portability and Accountability Act (HIPAA) of 1996 mandates the protection of health-related data. This includes any data that directly relates to patients of medical facilities and any data collected about an individual's health. Many non-health-related organizations must comply with HIPAA because they collect information on users for insurance and health plans.

Organizations that process transactions from major credit cards need to comply with the Payment Card Industry Data Security Standard (PCI DSS). E-commerce transactions traverse the Internet, and data is often stored in cloud-based storage. Organizations that use cloud-based storage for credit card data need to ensure that the cloud-based provider meets the minimum standards defined by PCI DSS.

Some cloud providers are meticulous about complying with laws and standards, but other providers might not be so careful. Organizations need to be aware that they are ultimately responsible for protecting their data. If attackers gain access to sensitive data that an organization stores in cloud-based storage, the organization is still responsible for the data breach.

Chapter Review

Switches connect devices together in a network, and routers connect networks together within a local area network (LAN) and provide physical segmentation between systems and networks.

Virtual local area networks (VLANs) increase security by logically segmenting traffic with a switch, similar to how a router physically segments traffic between networks. Physical security methods protect routers and switches. Voice over Internet Protocol (VoIP) is used to transfer multimedia and voice communications over IP networks, including the Internet. Secure Real-time Transport Protocol (SRTP) provides confidentiality, authentication, and replay protection for VoIP signals.

Proxy servers improve the speed of Internet access for users and can also filter traffic. Retrieved web pages are cached on the proxy server, and the proxy server serves these cached web pages instead of using Internet bandwidth to retrieve the web page from the Internet again. Website filtering allows a proxy to block a user's access to specific website locations. Most proxy servers also have NAT installed to translate private and public IP addresses.

Firewalls are an important security element in a network and for individual computers. Packet-filtering firewalls can filter traffic based on IP addresses, ports, and some protocols by using the protocol ID. Stateful inspection firewalls identify active TCP and UDP sessions and can open and close firewall ports dynamically based on the needs of active connections. An application firewall includes different elements to examine specific commands used by different protocols, such as Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), and Simple Mail Transport Protocol (SMTP). Next-generation firewalls and unified threat management (UTM) devices add more filtering capabilities such as blocking malware, filtering spam, and URL filtering.

A DMZ is commonly created with two firewalls from separate vendors to achieve defense diversity. This is in line with an overall defense in depth strategy. Even if a vulnerability appears in one firewall, it's unlikely a vulnerability will appear in both firewalls at the same time. Additionally, to succeed in penetrating an internal network that employs defense diversity, an attacker must have in-depth knowledge about firewalls from two vendors.

Firewalls can be network-based or host-based. A network-based firewall protects traffic going into or out of an overall network, while a host-based firewall protects traffic for an individual system. Network-based firewalls are typically hardware-based, while host-based firewalls typically run as an additional software component on a server or desktop system. Using both network-based and host-based systems is part of an overall defense in depth strategy.

Many organizations provide access to their internal network through a VPN remote access solution. A VPN provides access to the private network over a public network such as the Internet. VPNs use a tunneling protocol to encapsulate the protocols needed on the internal network. Tunneling protocols include Secure Shell (SSH), Layer 2 Forwarding (L2F), Point-to-Point Tunneling Protocol (PPTP), Layer 2 Tunneling Protocol (L2TP), Internet Protocol security (IPsec), and Transport Layer Security (TLS). When using TLS, users often only need a web browser to connect to the VPN server.

Authentication protocols used for remote access include Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), Microsoft CHAP (MS-CHAPv1), MS-CHAPv2, Extensible Authentication Protocol (EAP), Remote Authentication

Dial-in User Service (RADIUS), Diameter, and Terminal Access Controller Access Control System+ (TACACS+). RADIUS, Diameter, and TACACS+ also provide authentication, authorization, and accounting.

Traffic shaping devices can delay certain types of traffic, such as audio and video streaming traffic, so that it doesn't consume as much bandwidth. This gives other traffic a higher priority. WAN optimization can use traffic shaping to improve VPN connections.

Network access control (NAC) helps protect internal networks by differentiating between healthy and unhealthy systems and restricting access of unhealthy systems to quarantined networks. NAC is very useful with remote access because remote clients aren't directly controlled by the organization and can have varying levels of security applied. NAC can also be used within the organization to ensure that unhealthy clients are isolated from healthy clients.

Virtualization refers to replacing hardware with software. Host systems include a hypervisor that creates, runs, and manages virtual machines (VMs), also known as guest systems. Virtual desktop infrastructures (VDIs) provide users with a virtual desktop that can be persistent or nonpersistent. Software-defined networking (SDN) separates the data plane and control plane, providing logical segmentation within a network. An important security step with virtualization is to keep all hosts and guests up to date with current patches.

Cloud computing includes Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS). A public cloud can be accessed by anyone. A private cloud is used by only a single organization. A hybrid cloud is a combination of any two or more clouds. Although cloud computing provides reduced costs for many users, it also includes increased risk. Data hosted in the cloud can be compromised due to errors or problems with the cloud provider. Encrypting data on the client before storing it in the cloud can help provide confidentiality of data, but encrypting data using a cloud provider's encryption services often does not provide adequate protection.

Questions

- 1.** Which of the following best describes the mapping of data held within a switch's table?
 - A.** IP address to port
 - B.** MAC address to port
 - C.** IP address to MAC address
 - D.** Physical port to logical port

- 2.** Which of the following provides physical segmentation for hosts and networks?
 - A.** SDNs

B. VLANs

C. Routers and switches

D. Proxy servers

3. Which of the following is the best choice to logically segment traffic on a network?

A. VLAN

B. EAP

C. SSL

D. TLS

4. Which of the following can provide security for VoIP?

A. RADIUS

B. TACACS+

C. Diameter

D. SRTP

5. Administrators recently installed a proxy server and configured all clients to use it when accessing websites. How does the proxy server filter traffic?

A. URLs

B. ACLs

C. MACs

D. VLANs

6. A packet-filtering firewall can block ICMP traffic, such as ping requests. How does a packet-filtering firewall identify ICMP traffic?

A. Based on the protocol number having a value of 1

B. Based on the port having a value of 1

C. Based on the state of the ICMP packet

D. Based on the protocol number of 1701

7. Which of the following choices provides the best protection against potentially malicious FTP commands?

A. Defense diversity

B. Packet-filtering firewall

C. Stateful inspection firewall

D. Application firewall

8. How can you provide defense diversity with a DMZ?

A. Use a single firewall.

B. Use two firewalls from the same vendor.

C. Use two firewalls from different vendors.

D. Ensure that only trusted partners are allowed access.

9. It's common to enable or install a firewall on a server to protect the server. What type of firewall is this?

A. Network-based

B. Hardware-based

C. Packet-filtering

D. Host-based

10. Your organization is planning to implement a VPN. Which of the following will provide the best security for the VPN?

A. SSL

B. PPTP

C. L2TP

D. L2F

11. Of the following choices, what is used for VPNs? (Select all that apply.)

- A. L2TP
- B. PPTP
- C. SSLTP
- D. TLS

12. What port does a TLS VPN typically use?

- A. 1701
- B. 1723
- C. 143
- D. 443

13. How can users access a TLS VPN? (Select all that apply.)

- A. With a web browser
- B. With a dedicated application
- C. With broadband access but never DSL access
- D. With an IMAP application

14. Of the following choices, what indicates the primary improvement that MS-CHAPv2 included over previous protocols?

- A. Support for biometrics
- B. Use of certificates
- C. Mutual authentication
- D. Use of a nonce

15. Which of the following identifies the correct representation of RADIUS?

- A. Remote Access Dial-in User System
- B. Remote Authentication Dial-in User Service
- C. Roaming Access Dial-in User Service

D. Remote Authentication Dialing User System

16. What port does TACACS+ typically use?

A. 25

B. 49

C. 53

D. 443

17. What can be used to examine the health of a client prior to allowing network access and restricting access of unhealthy clients to a quarantined network?

A. RADIUS

B. TACACS+

C. NAC

D. SRTP

18. Which of the following represents the greatest risk to virtual systems?

A. Confidentiality

B. VM escape

C. Increased costs for power and cooling

D. Loss of control of data in the cloud

19. Your organization is evaluating various cloud service models. Management wants to pick a model that places most of the security and maintenance responsibility on the cloud vendor. Which of the following is the BEST choice?

A. IaaS

B. PaaS

C. SaaS

D. RADIUS

20. An organization is sharing resources with another organization using cloud-based computing. Which of the following cloud operation models does this describe?

- A. Community
- B. Hybrid
- C. Private
- D. Public

Answers

1. B. A table in a switch maps the media access control (MAC) address to the physical port. Routers map network IP addresses to the physical port with the corresponding gateway IP addresses. The Address Resolution Protocol (ARP) resolves IP addresses to MAC addresses, but this data is not held in a switch. Physical ports and logical ports are not mapped together.

2. C. Routers and switches provide physical segmentation for hosts and networks. Software-defined networks (SDNs) and virtual local area networks (VLANs) provide logical segmentation for hosts and networks. Proxy servers do not segment networks.

3. A. A virtual local area network (VLAN) segments traffic on a network using a switch. Extensible Authentication Protocol (EAP) is used for authentication, not segmentation. Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are transport encryption protocols and do not segment traffic. Note that this is an example of how the SSCP exam may just be testing your knowledge of relevant acronyms. If you know what the acronyms represent, the answer to the question becomes obvious because EAP, TLS, and SSL have nothing to do with segmenting traffic.

4. D. The Secure Real-time Transport Protocol (SRTP) provides confidentiality, authentication, and replay protection for Voice over IP (VoIP) transmissions. Remote Authentication Dial-in User Service (RADIUS), Terminal Access Controller Access Control System+ (TACACS+), and Diameter are used to provide authentication, authorization, and accounting (AAA) for remote access.

5. A. A proxy server filters traffic using URLs and can block users from accessing websites. While the question didn't ask, a proxy server also caches web pages to improve overall performance of web traffic. None of the other answers are directly related to proxy servers. Routers and firewalls filter traffic using access control lists (ACLs). Switches can filter traffic using media access control (MAC) address filtering. A virtual local area network (VLAN) logically segments traffic and it is unrelated to a proxy server.

6. A. Packet-filtering firewalls can filter traffic based on IP addresses, ports, and protocol numbers, and a protocol number of 1 identifies Internet Control Message Protocol (ICMP) traffic. ICMP does not use port 1. A packet-filtering firewall is a stateless firewall, so it doesn't

evaluate the state of the ICMP packet. Protocol number 1701 is not used, but Layer 2 Tunneling Protocol (L2TP) uses port number 1701.

7. D. An application firewall (also called an application proxy or an application gateway firewall) can inspect commands used by individual protocols such as File Transfer Protocol (FTP) and block potentially malicious commands. Defense diversity refers to using firewalls from two different vendors in a demilitarized zone (DMZ). A packet-filtering firewall can only inspect individual packets for IP addresses, ports, and protocol IDs. A stateful inspection firewall can track the activity within TCP and UDP sessions, but can't interpret commands.

8. C. You can provide defense diversity with a DMZ by using two firewalls from different vendors. If a vulnerability appears in one, it's unlikely that a vulnerability will exist in the second firewall at the same time (unless the second is from the same vendor). A single firewall doesn't provide any diversity. An extranet (not a DMZ) would allow access only to trusted partners. Note that this question could also be presented as a "Select all that apply" question by asking which of the following are used for VPNs.

9. D. A host-based firewall is installed or enabled on individual hosts, such as desktop computers or servers, and provides protection for the host. Network-based firewalls protect the network rather than individual systems. Packet filtering identifies the method used by the firewall, and both network-based and host-based firewalls can filter packets.

10. C. Layer 2 Tunneling Protocol (L2TP) is the best choice of those given. It can use Internet Protocol security (IPsec) for encryption. Secure Sockets Layer (SSL) has been deprecated and should not be used. Transport Layer Security (TLS) is the designated replacement for SSL and can be used for a VPN. Although the terms are often used interchangeably, TLS is not the same as SSL. Layer 2 Forwarding (L2F) is a legacy VPN tunneling protocol that is rarely used because it does not encrypt the traffic.

11. A, B, D. Layer 2 Tunneling Protocol (L2TP), Point-to-Point Tunneling Protocol (PPTP), and Transport Layer Security (TLS) are used for virtual private networks (VPNs). There is no such thing as SSLTP in the context of VPNs.

12. D. A Transport Layer Security (TLS) virtual private network (VPN) typically uses TCP port 443, the same port as Hypertext Transfer Protocol Secure (HTTPS). Layer 2 Tunneling Protocol (L2TP) uses port 1701. Point-to-Point Tunneling Protocol (PPTP) uses port 1723. Internet Message Access Protocol version 4 (IMAP4) uses port 143.

13. A, B. Users typically access a Transport Layer Security (TLS) virtual private network (VPN) using a web browser or a dedicated application (such as OpenVPN Connect). A TLS VPN is not dependent on a specific type of Internet connection (such as broadband or DSL). Internet Message Access Protocol (IMAP) is used with e-mail, not VPNs.

14. C. MS-CHAPv2 uses mutual authentication, where the client authenticates to the server and the server authenticates to the client. MS-CHAPv2 does not directly support biometrics. MS-CHAPv2 can be used with the Extensible Authentication Protocol (EAP) to support certificates,

but it cannot do so on its own. CHAP uses a nonce (a number used once), so this isn't an improvement.

15. B. RADIUS is an acronym for Remote Authentication Dial-in User Service. The other choices are not valid.

16. B. Terminal Access Controller Access Control System+ uses TCP port 49. Simple Mail Transport Protocol (SMTP) uses TCP port 25, Domain Name System (DNS) uses TCP port 53 and UDP port 53, and Hypertext Transfer Protocol Secure (HTTPS) uses TCP port 443.

17. C. A network access control (NAC) system can check a system's health based on a predefined health policy and restrict the access of unhealthy clients to a quarantined network. Remote Authentication Dial-in User Service (RADIUS) and Terminal Access Controller Access Control System+ (TACACS+) are used to provide authentication, authorization, and accounting (AAA) for remote access. Secure Real-time Transport Protocol (SRTP) provides confidentiality, authentication, and replay protection for Voice over IP (VoIP) transmissions.

18. B. VM escape is a known attack against virtual systems. If the attack is successful, an attacker can access the host system and all virtual systems within the host. Loss of confidentiality (not confidentiality) is a risk that can be reduced with encryption. Virtualization reduces costs for power and cooling. Loss of control of data stored in the cloud is a risk associated with cloud computing, but organizations can use virtual systems internally to keep control of their data.

19. D. Cloud vendors have the most responsibility with the Software-as-a-Service (SaaS) model and customers have the least amount of responsibility. Cloud vendors have the least responsibility with the Infrastructure-as-a-Service (IaaS) model and customers have the least amount of responsibility. Customers and vendors typically share responsibility in a Platform-as-a-Service (PaaS) model, with the vendor typically managing a little more than the customer. Remote Authentication Dial-in User Service (RADIUS) is not a cloud service model.

20. A. A community cloud is a private cloud that is shared by two or more organizations. A hybrid cloud is a combination of any two or more clouds. A private cloud is only available to users within an organization. Public cloud-based services are provided by third-party vendors and are available to anyone.

CHAPTER 5

Attacks

In this chapter, you will learn about

- Comparisons of various types of attackers
- Differences between white hats, black hats, and gray hats
- Emerging issues related to advanced persistent threats
- Comparing various losses from insider threats
- Basic countermeasures used to thwart attacks
- Common denial of service (DoS) and distributed DoS (DDoS) attacks
- Botnets and zombies
- Reconnaissance and fingerprinting attacks
- Software security countermeasures
- Buffer overflow attacks, cross-site scripting, and cross-site request forgery
- Various password attacks and common protections
- Phishing, spear phishing, and whaling
- Wireless attacks and countermeasures
- Social engineering tactics and countermeasures
- The importance of training as a countermeasure

Comparing Attackers

Attackers break into computers and networks with malicious intent and often for personal gain. Sometimes, they try to take down a system or put it out of service. Other times, they try to get into the system to gain information. For example, an attacker may break into a system to steal credit card data from a business's customers. The attacker can then use the credit cards fraudulently. Alternatively, they may attempt to bring a system down as an act of revenge against an organization.

Make no mistake, though. These attackers are criminals. Just as someone who breaks into your home is a criminal, attackers who break into networks and individual systems are also criminals. They are breaking laws. Many attack individuals and networks for monetary gain without any apparent care to how the loss affects their victims.

In general, criminals must have motive, means, and opportunity to complete a crime. The motive for an attacker is often money, but can also be revenge, espionage, or other reasons. The means is often with some type of malicious code (discussed in [Chapter 6](#)), and the opportunity is available for anyone with Internet or internal access. Indeed, for a greedy criminal willing to risk jail time, there are often enough programs available on the Internet that they don't even have to write the malicious code, but only need to locate and download it. The following sections identify the different types of attackers and some names used to identify them.

Hackers and Crackers

Two terms you'll come across when studying IT security are hackers and crackers. The following two bullets provide traditional definitions (even though many people do not use these traditional definitions):

- **Hacker** Someone who is proficient with computers (often with programming skills) and uses these skills to gain and share knowledge with others. By this definition, a hacker does not break into systems with malicious intent or for personal gain.
- **Cracker** Someone who is proficient with computers and uses these skills to attack systems. A cracker *does* attack systems with malicious intent or for personal gain.

The English language has a long history of changing the use of words and their meanings, often with some people kicking and screaming along the way. For example, you probably don't use *thee* and *thou* in your daily speech. However, there was a time when everyone used these words regularly.

Similarly, the word *hacker* is evolving. Currently, the media widely uses the term *hacker* to identify anyone launching attacks on computers or networks. Even though some IT people are clinging to the traditional definition, it has lost the traditional meaning outside of these circles.



NOTE In this book, I'm using the term *attacker* to refer to anyone who attacks a network.

White Hats, Black Hats, and Gray Hats

Using the traditional definitions, hackers are good and crackers are evil. However, since the media began confusing these terms, the concept of white hats, black hats, and gray hats has

emerged. These designations are reminiscent of the older cowboy movies where the good guys wore white hats and the bad guys wore black hats.

White hats fall into the traditional definition of hackers and are sometimes called ethical hackers. Organizations often hire them as security consultants to perform vulnerability assessments or provide other security services.

In contrast, black hats fall into the traditional definition of a cracker. Black hats are criminals who attack systems with malicious intent or for personal gain.

Of course, good and evil don't always have definitive lines with good on one side and evil on the other. There are gray areas. Similarly, gray hats are individuals who have exceptional computer and networking skills, but they don't use them for personal gain or with malicious intentions. However, their activities may cross ethical boundaries.

As a comparison between the three, imagine that someone discovers a weakness or vulnerability in a system. A white hat would inform the owners of the system but not take any further action. A black hat would try to exploit the vulnerability without telling the owner. A gray hat may first try to let the owners know of the vulnerability, but then decide to publish the vulnerability if the gray hat doesn't think the owner is acting responsibly. By publishing the vulnerability, the gray hat makes the vulnerability known to the black hats, who can then use the information to exploit it.

Advanced Persistent Threats

An advanced persistent threat (APT) refers to a highly sophisticated group of people who target their attacks. They have both the capability and the intent to carry out successful attacks and are typically sponsored by a government. Once they identify a target or goal, they persistently attack until they succeed. Many countries reportedly have teams of IT personnel regularly probing the IT systems of organizations and other government networks.

In 2013 security firm Mandiant, now a division of FireEye, released a report titled *APT1: Exposing One of China's Cyber Espionage Units* (available for download by searching the Internet with the title). The report documents the activities of a Chinese group that Mandiant named "APT1" and is one of the first comprehensive reports on APTs. Mandiant is very clear in its conclusions and the report includes this statement: "The details we have analyzed during hundreds of investigations convince us that the groups conducting these activities are based primarily in China and that the Chinese Government is aware of them." The report documents the following activities by APT1:

- Compromised at least 141 companies in 20 different industries in many different countries, including the United States, Canada, and the United Kingdom
- Maintained access in compromised systems for an average of 365 days and as long as 58 months

- Stole a large quantity of intellectual property, including 6.5 terabytes of data from a single organization
- Developed over 40 different families of malicious software (malware)

Mandiant reported that the most commonly observed method APT1 used to gain access to a network was a spear phishing attack (described later in this chapter). This type of attack establishes a foothold by infecting a single system, and then the attackers quickly infect other systems in the network. For example, they often install backdoors on internal systems, providing the attackers with remote access to the systems.

An important point to remember here is that APT1 is only one APT operating out of only one country, China. There are more. As another example, the U.S. Department of Homeland Security (DHS) and the Federal Bureau of Investigation (FBI) documented Russian hacking of networks and endpoints related to U.S. elections, as well as networks of government, political, and private sector organizations. The DHS and FBI published their findings in late December 2016 in the Joint Analysis Report JAR-16-20296A, “GRIZZLY STEPPE – Russian Malicious Cyber Activity” (available for download by searching the Internet with the JAR number or the nickname GRIZZLY STEPPE).

The report identified two APTs operating out of Russia as APT 28 (nicknamed Fancy Bear) and APT 29 (nicknamed Cozy Bear) and documented how APT 29 infiltrated at least one political organization in the summer of 2015, and how APT 28 infiltrated a political organization in spring 2016. While the details of these infiltrations are quite complex, the initial infiltration apparently started with highly targeted spear phishing campaigns, similar to how APT1 gained a foothold in many organizations.

Similarly, other countries very likely have their own APTs. APTs don’t attempt to get all the information immediately. Instead, they are patient and willing to get a little bit of information at a time. However, they are also persistent. When one attack fails, they don’t give up. Instead, they look for other ways to attack and keep pursuing the target until they achieve some success.

After many successful attacks against computers and networks in the United States, the U.S. government created the United States Cyber Command (USCYBERCOM), which became operational in 2010. Its mission statement is as follows:

USCYBERCOM plans, coordinates, integrates, synchronizes, and conducts activities to: direct the operations and defense of specified Department of Defense information networks and; prepare to, and when directed, conduct full-spectrum military cyberspace operations in order to enable actions in all domains, ensure US/Allied freedom of action in cyberspace and deny the same to our adversaries.

Many other countries have created similar organizations to combat the threats of cyberwarfare.

Insider Threats

An insider is someone who works within an organization, such as an employee, consultant, or vendor, who has authorized access to the organization's assets. Insiders have the means and opportunity to cause a significant amount of damage and can cause the loss of confidentiality, integrity, and availability of an organization's resources.

Insiders can make mistakes or ignore security practices resulting losses. They can engage in malicious activity as revenge. And they can engage in data theft such as stealing trade secrets, customer information, or other sensitive information from the organization, which they then can attempt to sell for personal gain.

Accidental Losses from Insider Threats

Not all insider threats are deliberate theft or destruction of data. Instead, they are typically users not following security protocols or practices, resulting in losses. As an example, imagine an organization regularly tells employees not to click on links included in e-mail from unknown sources. A user that clicks on one of these links is an insider threat.

Consider the threats posed by APT 28 and APT 29. These APTs launched sophisticated spear phishing attacks. At least one user in each instance (yes, it only takes one user) responded inappropriately to these attacks. The APT 29 spear phishing attack included a link that activated malicious software (malware) if the user clicked it. The APT 28 spear phishing attack tricked users into changing their passwords through a fake website hosted by the APT 28 group. This allowed the APT 28 group to use these credentials to log on to the users' accounts.

As another example, a U.S. military data-mining program scoured the Internet looking at social media websites and other sites. The data was openly collected as part of the government's Outpost program and analyzed for different purposes. This data was apparently backed up in Amazon Web Services (AWS) Simple Storage Service (S3), which is a cloud-based storage platform. However, it wasn't protected with either encryption or permissions. A security analyst did a routine scan of Amazon-hosted data silos in late 2017, and found dozens of terabytes of data hosted there. While this data has since been protected, it was openly available for anyone to locate and download for some time.

Some other examples include the following:



NOTE Accidental threats aren't attacks. Employees who accidentally cause damage aren't doing so with malicious intent or for personal gain. However, some incidents caused by accidental threats can cause significant damage, so it's worthwhile knowing about them.

- **Responding to phishing attempts** Phishing is explained in more depth later in this chapter, but in short, it's an e-mail that encourages the user to respond with personal information or to click a link. The user can give up valuable company information or inadvertently install malicious software on their system. This provides a method for an outside attacker to attack the organization.
- **Forwarding malicious software (malware) or bringing malware from home** An uneducated user can forward viruses or other forms of malware via e-mail or bring it on a USB drive from a home computer. The malware could cause damage to internal systems or use other methods to gather information for attacks against the organization.
- **Unauthorized data access** If adequate access controls aren't in place, a user may be able to access data. The user may then accidentally modify or delete it, or even disclose the contents of the data to someone else without realizing its true value. Something as simple as an unauthorized employee gaining access to personal records of other employees can cause a great deal of havoc to an organization.
- **Losing hardware** Many organizations issue hardware such as laptops or USB drives to users. When users don't recognize the value of the hardware and the data it holds, they may not provide adequate protection, resulting in its loss or theft. As an example, many users have left laptops in conference rooms during a lunch break, only to discover them missing after lunch.

Revenge Attacks from Insider Threats

Some insiders are out for revenge. For example, Rajendrasinh Makwana was terminated from a contractor job at Fannie Mae in 2008. However, he retained access to the systems for at least a short period afterward. He later embedded a time bomb script in the Fannie Mae system. It was set to run on January 31, 2009, and would have deleted passwords on about 4,000 servers, erased all the data and backup data for the servers, and then powered them down. He also programmed the script to disable the ability to turn the servers back on remotely. Had the script run, the result would have been catastrophic.

Luckily, another engineer discovered the damaging script before it ran. Interestingly, Makwana's employer terminated him for allegedly making unauthorized changes to other systems. However, the employer didn't revoke his access right away, which gave him enough time to install the malicious script. A jury convicted Makwana and he was ultimately sentenced to over three years in prison.

Some organizations are using sophisticated monitoring systems to detect and prevent insider threats before they can successfully attack and cause damage. For example, in 2009 the Transportation Security Administration's (TSA) Colorado Springs Operations Center internal surveillance system identified the actions of Douglas Duchak as he was planting malware in a critical system about a week after he was terminated with a two-week notice. He received a two-year prison sentence.

Data Theft from Insider Threats

Data theft by insiders is also a concern. Insiders sometimes steal data with the hope of selling it, or using it in a new job with a rival company. In some cases, employees are just snooping, looking at files and folders to see what they can see. While these events are still occurring, they aren't reported as often, but instead handled within the organization. Effective access controls (including strong authentication and authorization mechanisms) combined with strong auditing techniques help to prevent successful attacks from insiders.

Script Kiddies

A script kiddie is a relatively inexperienced attacker who can run a script or application to launch an attack. The script kiddie rarely has the knowledge to program the script or application and often doesn't understand the details of what the attack is doing. You can think of a script kiddie as a bored teenager who downloaded some hacking tools and used them to launch attacks.

It's worth stressing that although script kiddies exist, many attackers today are highly proficient at programming techniques. They have detailed knowledge of different attack methods and are constantly modifying existing attacks to prevent detection.

Insider Threats vs. APTs

Years ago, many security professionals considered insider threats to be the primary threat to computer systems. However, this has changed as advanced persistent threats (APTs) and criminal groups have become more active.

Today, organizations experience much greater losses from APTs and criminal groups than they do from insider threats. This includes loss of data and money. As an example, the WannaCry ransomware attack affected more than 300,000 organizations worldwide in 2017. Ransomware attacks lock a user's computer or data, and demands a payment as ransom to release the computer or data. Symantec identified many similarities between this attack and other attacks attributed to the Lazarus Group. This includes the 2014 attack on Sony Pictures Entertainment and a 2016 attack on a Bangladesh central bank account.

Symantec indicated the Lazarus Group may not be an attack from a state-sponsored APT, but instead a cybercrime campaign. Many other cyber firms have connected the Lazarus Group with North Korea, indicating the government supports the group. In December 2017, the U.S. government attributed this attack to the Lazarus Group and blamed North Korea for it.

It's worth noting that many of these attacks are the result of insiders taking inappropriate actions, such as clicking links in e-mails or opening malicious attachments sent in e-mails. In other words, insiders are involved, but insiders are not launching the attacks.

This isn't to say that an organization shouldn't protect itself against malicious insiders. Many common security practices (such as implementing the principle of least privilege and separation of duties policies) limit the potential losses against malicious insiders. You can also think of this

as the privileges you give your users are the privileges that you give their attackers. By implementing the principle of least privilege for employees, it limits the potential losses from both malicious insiders and successful attackers.

Exploring Attack Types and Countermeasures

There are numerous attack types and countless variations. As attackers have launched new types of attacks, IT security professionals have identified new methods of protection. When one attack method no longer works, attackers develop others. It's a constant battle between the attackers and security professionals. As long as attackers are able to achieve their objectives through attacks, expect them to continue.

This section outlines many of the common types of attacks and some countermeasures. This is by no means a complete list, but it does provide an overview.

Basic Countermeasures

Countermeasures are security controls that reduce risks. They typically reduce vulnerabilities or reduce the impact of a threat. [Chapter 9](#) covers countermeasures and controls in more depth. The following sections introduce some basic countermeasures that help prevent many attacks.

Patching Systems

Many attacks are possible because of vulnerabilities discovered in an operating system. However, once vendors learn about these vulnerabilities, they develop, test, and release patches and updates to plug the security holes. As long as administrators keep systems up to date with current patches and updates, the systems are no longer vulnerable to that specific attack. However, unpatched systems remain vulnerable to the attacks.

This seems obvious to most security professionals. However, many attacks succeed because organizations don't follow basic security practices such as keeping systems up to date with current patches. As an example, the 2017 Equifax data breach exposed sensitive data for as many as 143 million people. Attackers exploited a vulnerability in an Apache web application. However, Apache Software Foundation had released a patch to remove the vulnerability over two months prior to the data breach. If Equifax had applied the patch, it would have prevented the data breach.

System Hardening

Administrators harden systems by increasing the security provided with the default configuration. Hardening includes removing or disabling unneeded protocols and services, ensuring the system is deployed with all current patches and updates, and enabling software firewalls.

System Isolation

Another basic countermeasure is system isolation or segmentation. [Chapter 4](#) covers the use of routers to provide physical segmentation for network hosts and using virtual local area networks (VLANs) to provide logical segmentation of network hosts. It's also possible to use virtualization technologies to isolate systems.

An extreme example of system isolation is to place systems in isolated networks without any Internet access. This is most commonly done with industrial control systems (ICSs) or classified networks.

Increasing User Awareness

User awareness programs help users understand risks, recognize risky behaviors, and understand their responsibilities related to security. A goal is to reaffirm safe computing practices and encourage users to change their behavior if they are engaging in unsafe computing practices.

Security professionals use a variety of methods to increase user awareness. This includes routine training, annual updates, periodic e-mails, and more. The value of increasing user awareness by any methods possible cannot be overstated.

One method some administrators use is sending out e-mails with catchy subject lines and telling about recent compromises caused by users. As an example, imagine a user in a different organization clicked a link that took the user to a malicious site, downloaded malware onto the user's system, and then infected other computers in the network. Sending out an e-mail to personnel with the subject line "Are you the weakest link?" can grab their attention. The e-mail can give short details about the incident, with a link to an article giving more details.

Intrusion Detection and Protection Systems

Network-based intrusion detection systems (NIDSs) can often detect attacks in progress and send alerts to administrators. Some NIDSs can take action to change the environment and block the attack, but only after it has started. A network-based intrusion prevention system (NIPS) is similar to a NIDS with one distinct difference—a NIPS is placed inline with the traffic, allowing the NIPS to inspect all traffic and identify and block potential attacks. The lines between an IDS and an IPS are sometimes blurred, so you may see them referred to as an IDPS, including in this chapter.

IDPSs are relatively common in networks today. In some networks, an IDPS may be a dedicated device. In other networks, an IDPS may be part of an appliance such as a unified threat management (UTM) device that combines multiple technologies together into a single appliance. Many next-generation firewalls also include IDPS capabilities. [Chapter 8](#) covers IDSs and IPSs in more depth.

Spoofting

Spoofting refers to impersonating or masquerading as someone or something else, and many attacks use spoofting techniques. [Chapter 3](#) presented the concept of spoofting in the context of IP addresses and media access control (MAC) addresses. Attackers can change the source IP address or source MAC address so that traffic appears to be coming from a different computer. Similarly, attackers can change the From field in an e-mail to make it appear as though the e-mail is coming from someone else.

Data Theft

Data theft refers to any attack that allows the attacker to exfiltrate data from an internal network. Attackers have stolen a wide assortment of data from just about every business and private sector. This includes data from banking and financial institutions, credit reporting agencies, small and large businesses, educational institutions, government and military organizations, and healthcare organizations.

Many of these attacks expose user personally identifiable information (PII). The Identity Theft Resource Center (ITRC), a U.S. nonprofit organization, publishes annual reports on data breaches in the United States. In 2017, the ITRC reported a record high of 1,579 data breaches exposing more than 178 million records. Over 50 percent of these breaches exposed user's Social Security numbers (SSNs) and 19 percent of the breaches exposed users' credit card or debit card data.

Almost 60 percent of these reported data breaches were the direct result of attacks, such as phishing, ransomware, and other malware attacks. Note that many of these successful attacks were caused by insiders not following security-related best practices, such as not opening attachments in e-mails received from unknown entities, or clicking links in e-mails. Other breaches occurred from unauthorized access, insider theft, accidental exposure, and physical theft.

APTs are responsible for a significant amount of data theft worldwide, including PII data, but they have stolen much more. As an example, APT1 reportedly stole 6.5 terabytes of compressed data containing intellectual property from a single organization. FireEye, in its report titled "APT28: A Window into Russia's Cyber Espionage Operations?" concluded that APT28 is engaged in espionage against political and military targets in Eastern European governments and militaries, and European security organizations.

FireEye has also tracked an APT they refer to as APT33. They identified malware used by APT33 tied to an Iranian persona and reported "We assess APT33 works at the behest of the Iranian government." APT33 targets many organizations in the aviation sector, aerospace sector, and energy sector. While it's unclear what data APT33 may have obtained, FireEye suspects they may be targeting these sectors for two primary reasons:

- **Military capabilities** FireEye indicates APT33 is collecting data to assess military aviation capabilities of Iran's regional rivals, such as Saudi Arabia. Additionally, they can use collected data to enhance Iran's domestic aviation capabilities.
- **Petrochemical production** Iran has expressed a desire to grow its petrochemical industry. By collecting up-to-date data on competitors in the region, APT33 can help Iran expand its production and improve its competitiveness in the region.

DoS

A denial of service (DoS) attack attempts to prevent a system from answering legitimate requests from users. The attack often attempts to overwhelm the system by consuming system resources. While the attacked system responds to the attack, it doesn't have enough resources to respond to regular requests. The goal is to deprive the organization or its customers from the use of a service.



EXAM TIP A denial of service (DoS) attack attempts to disrupt a service. It can do so by taking a system down or by consuming system resources with meaningless requests that overload its ability to respond to legitimate requests. The effect is that valid authorized users are unable to access a system.

As an example, a common DoS attack is a SYN flood attack that initiates the TCP handshake process but doesn't complete it. [Chapter 3](#) describes the three-way handshake used by TCP. Recall that a normal three-way handshake includes the following three packets:

1. A TCP client sends a packet with the SYN flag set to a system such as a server.
2. The server responds with a packet with the SYN and ACK flags set.
3. The client then completes the handshake process by sending back a packet with the ACK flag set.

In a SYN flood attack, the attacker withholds the third packet, leaving the half-open session on the server. For example, consider [Figure 5-1](#), which shows the process of a SYN flood attack. The first two packets are sent just like a regular three-way handshake, but the third packet is not sent. Instead, the attacker then sends out another SYN packet.

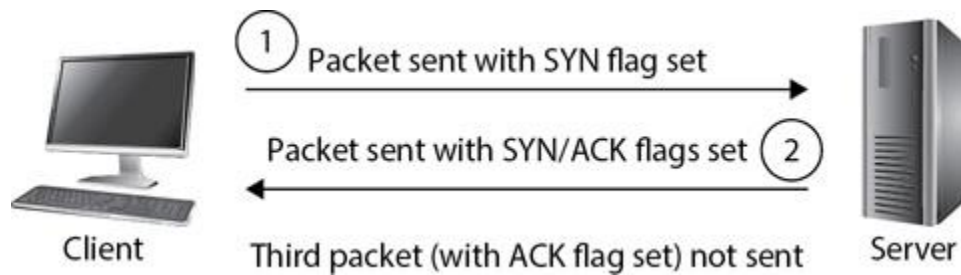


Figure 5-1 SYN flood attack



EXAM TIP In a SYN flood attack (also called a TCP SYN, TCP flood, and TCP half-open attack), the attacker floods a system with SYN packets but withholds the third packet in the TCP handshake process.

This is similar to a practical joker extending his hand to shake yours, you extending your hand in response, and the joker pulling away his hand at the last second. He sticks his hand out again, you extend your hand again, and he pulls his hand away again. You probably won't fall for this repeatedly. However, TCP will keep responding to SYN packets with SYN/ACK packets and continue waiting for the third packet.



TIP Most IDPSs can detect this type of attack. If the attack is coming from a single IP address, an IDPS can change the access control list (ACL) on the firewall or a boundary router to block all traffic from the offending IP address.

If this happens once or twice, it's no big deal. However, an attacker can initiate hundreds of half-open sessions in a short period. If the attack isn't detected, the server can keep each of these sessions open as it waits for hundreds of ACK packets. Each one of these sessions consumes resources, such as memory, processing power, and network bandwidth. The attack can slow down the server, and in extreme cases, it can result in resource exhaustion and actually crash the server.

The following are some other well-known DoS attacks:

- **Ping of death** An Internet Control Message Protocol (ICMP) ping packet is normally 32 bytes of data, although different operating systems may vary the size. A ping of death attack changes the size of the ping packet to more than 64KB. These large ping packets have crashed some systems, although most systems recognize the attack today and can avoid it.

- **LAND** A local area network denial (LAND) attack tricks a system into sending out packets to itself in an endless loop. In other words, the source and destination addresses are set to the attacked system. A LAND attack can shut down a system.



TIP Although many of these attacks will fail on protected systems, they are valuable to understand. Many current attacks are simply variations on attacks that have been successful in the past.

DDoS

A distributed denial of service (DDoS) attack is similar to a DoS attack except that it comes from multiple attackers at the same time. For example, an attacker may launch a TCP flood attack against a single system from hundreds or even thousands of systems. A key to the success of DDoS attacks is for the attacker to take control of multiple systems and then use these systems to launch the attacks. Attackers often use botnets (described in the next section) to launch DDoS attacks.



EXAM TIP The first step in a DDoS attack is often to compromise as many systems as possible. The attacker then directs these systems to launch the DDoS attack.

Botnets and Zombies

A botnet is a group of computers (often called *zombies*) controlled by an attacker. The term *botnet* is a combination of *robot* and *network*. The attacker (sometimes called a bot herder) manages a command and control center and directs the computers in the botnet to do the bidding of the attacker.

Computer users are usually unaware that their computer has become a zombie. Instead, malware infects computers without the user's knowledge. This allows attackers to control the computer from remote locations.

One way that a system is infected and joined to a botnet is through a drive-by download, as shown in [Figure 5-2](#). When a user visits an infected or malicious website, the website downloads the malware and attempts to install it on the user's system. Users sometimes unknowingly visit these sites during Internet surfing sessions. Other times, users click links included within malicious e-mails. As an example, the GameOver Zeus malware used drive-by downloads to join computers to botnets. Additionally, the attackers used malicious e-mails to encourage users to visit the malicious websites.

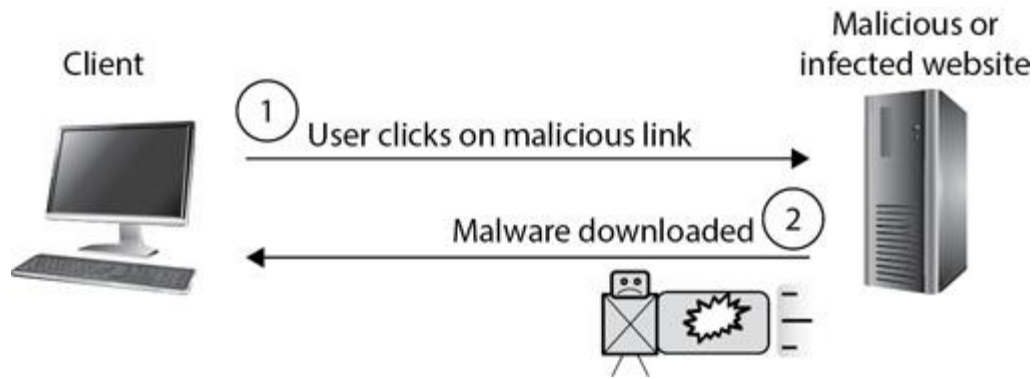


Figure 5-2 Single system infected in drive-by download



TIP Drive-by downloads are any downloads that occur without the user's knowledge after the user visits a website. They can include many different types of malware, including software that joins the computer to a botnet.

While a drive-by download is one way that computer systems are joined to botnets, it isn't the only way. Multiple types of malware can include code to join computers to botnets. [Chapter 6](#) covers malware in more depth.

[Figure 5-3](#) shows a botnet with multiple zombies in it. These zombies regularly check in with the attacker's command and control center, which is hosted on one or more servers controlled by the bot herder. Smaller botnets may have a single server hosting the command and control software. Larger botnets may have dozens of servers located around the world.

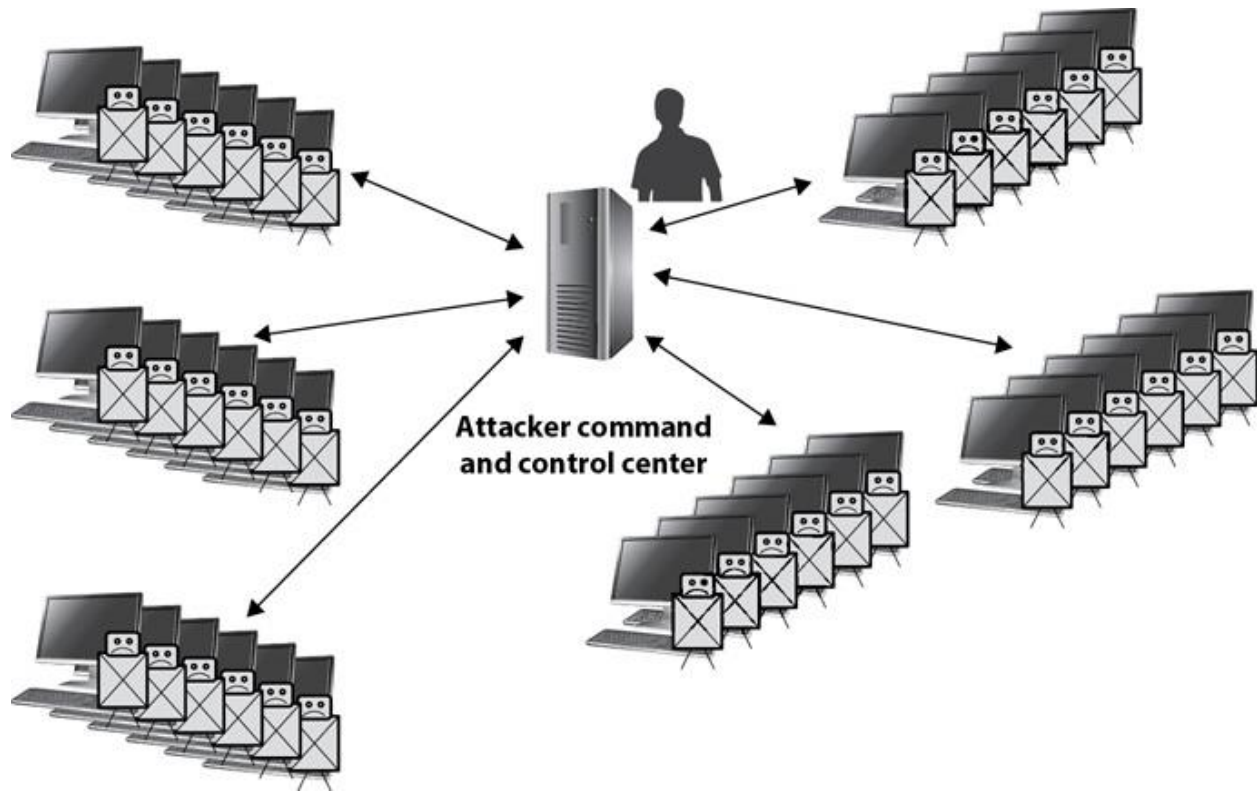


Figure 5-3 Zombies regularly check in with the command and control center.

The commands issued to the zombies in the botnet vary. Bot herders can direct zombies to bombard a server with a DDoS attack, as shown in [Figure 5-4](#). Attackers can also direct zombies to take other malicious actions, such as send spam and phishing e-mails, spread malware, or launch other types of attacks.

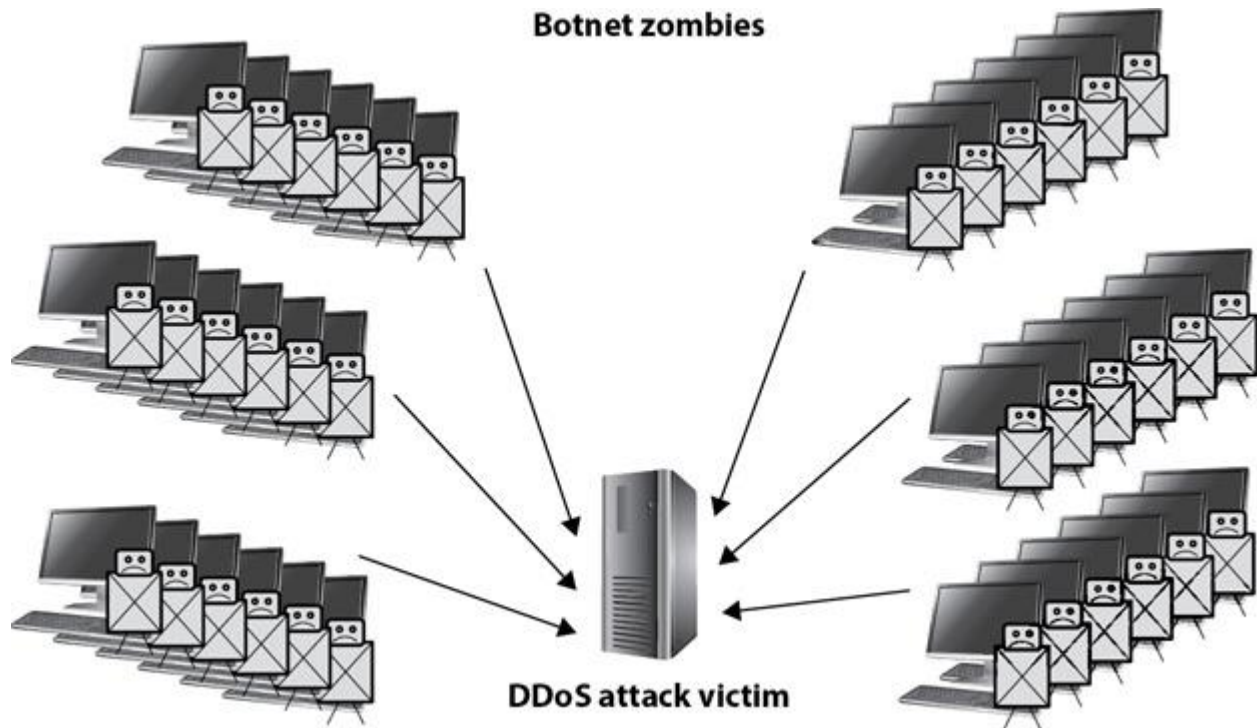


Figure 5-4 Zombies can launch DDoS attacks.

Criminals often rent out their botnets to other criminals, and if the botnet is large, the owners can generate a significant amount of money from these rentals. Many botnets control tens of thousands of computers, and some control more than a million. At its height, the Rustok botnet had more than a million zombies and sent out billions (yes, billions with a *b*) of spam e-mails daily. The Rustok botnet was taken down in March 2011, and IT watchers reported that spam levels immediately plummeted.



NOTE Microsoft launched the Microsoft Digital Crimes Unit (DCU) in 2010 and the Microsoft Cybercrime Center in 2013. They both work with other organizations to fight cybercrime, and have helped take down multiple botnets around the world.

Many current botnets are smaller, in part to evade detection. However, some large botnets still appear every now and then. In 2016, attackers gained control of about 100,000 Internet of Things (IoT) devices such as digital cameras, webcams, and digital video recorders. Attackers used Mirai malware to take control of the devices, and the botnet was commonly referred to as the Mirai botnet. This botnet then attacked Domain Name System (DNS) servers on the Internet, effectively crippling many popular websites such as Reddit, Tumblr, Amazon, Twitter, and more.

Sniffing Attack

A protocol analyzer (also called a sniffer or packet sniffer) captures packets transmitted over a network. After capturing packets, you can save them in a file (commonly called a capture) and browse them at your leisure. Automated tools can also scan captures to look for trends, possible attacks, and other items of interest. IT professionals use sniffers to troubleshoot and learn details about a network. Attackers use sniffers to capture data and exploit it for personal gain.

A popular sniffer is Wireshark. It's available as a free download and runs on most platforms, such as Windows, UNIX, and Linux.



NOTE You can read more about Wireshark and download a copy from here: www.wireshark.org. Wireshark was previously known as Ethereal.

If data is sent over the network in cleartext, an attacker can easily read it in a sniffer. However, if the data is encrypted, it can't be read. For example, [Chapter 3](#) mentioned the risks of unencrypted protocols such as File Transfer Protocol (FTP) and Trivial File Transfer Protocol (TFTP). Both FTP and TFTP send data over the network in cleartext by default. An attacker can capture these transmissions and read credential information, such as usernames and passwords, and the actual data. In contrast, FTP Secure (FTPS) uses Secure Shell (SSH) to encrypt the transmissions, preventing attackers from reading data.

[Figure 5-5](#) shows a screenshot of Wireshark with a capture. During this capture, I connected to a network drive and opened a file named passwords.txt. After stopping the capture, I browsed through it to locate packets that held data. In the figure, you can see arrow 1 pointing to packet 1470, which is an SMB Read AndX Response command. The center pane holds the details on this packet. Arrow 2 shows where the packet includes the name of the file (passwords.txt). The bottom pane shows the data in both a hexadecimal and text format, and arrow 3 shows the username and password in cleartext (Darril Gibson, IPa\$\$edSSCP).

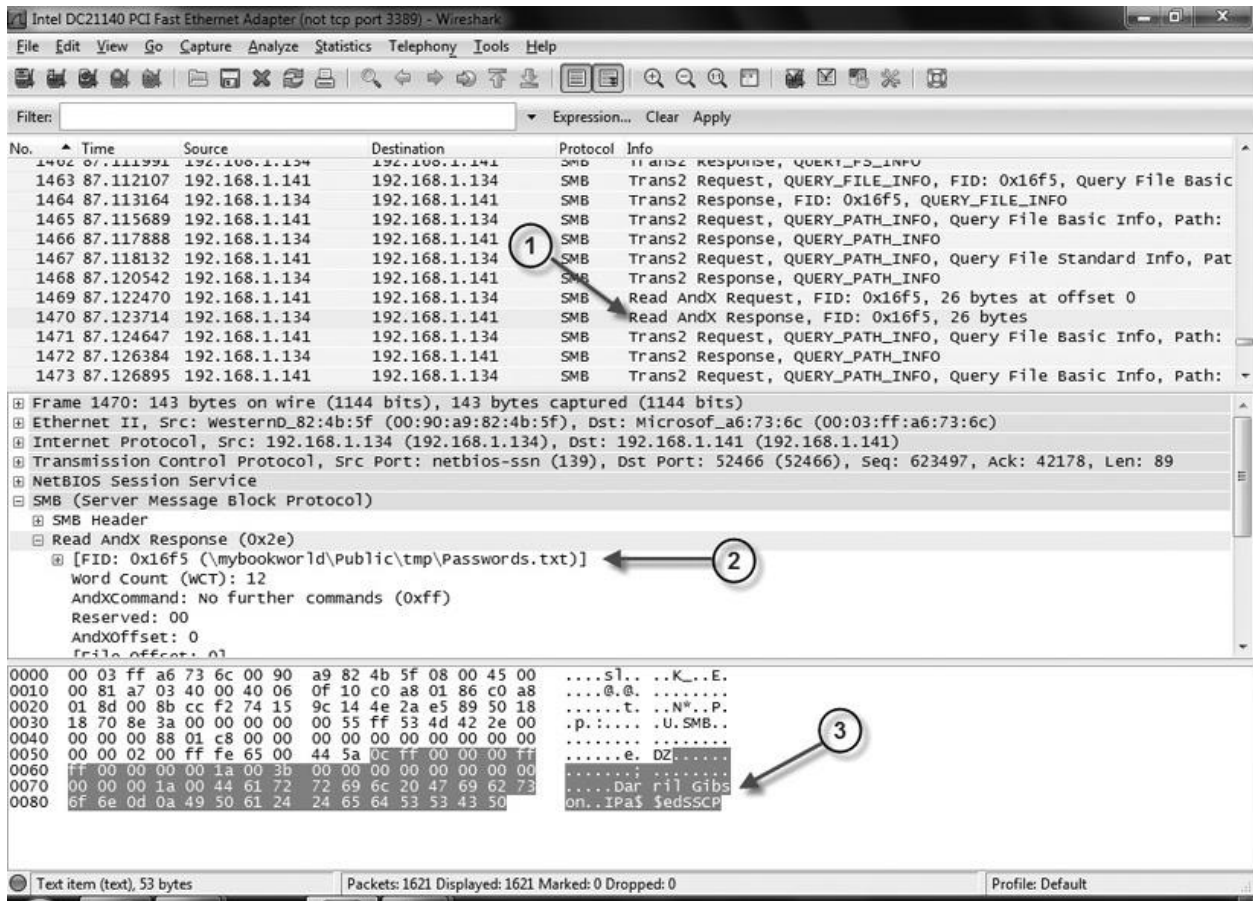


Figure 5-5 Viewing text data in Wireshark

Sniffers can work in either promiscuous mode or nonpromiscuous mode. In nonpromiscuous mode, the sniffer will capture only data sent directly to or from the computer's IP address. It will not capture any packets that do not have the computer's IP address as the source or destination address. In contrast, when the sniffer works in promiscuous mode, it will capture all data that reaches the computer regardless of the IP address. Technically, the sniffer configures the network interface card (NIC) to work in promiscuous mode for the application.



EXAM TIP Promiscuous mode allows a sniffer to capture all data that reaches the sniffer. When the sniffer runs in nonpromiscuous mode, it will only capture data sent directly to or from the system running the sniffer.

One of the ways an organization reduces vulnerabilities from sniffer attacks is to use switches instead of hubs. For example, consider [Figure 5-6](#), which shows both a hub and a switch in action. In the figure, PC1 and PC3 are sending data to, and receiving data from, each other. The

hub sends the data to all ports, so an attacker with a sniffer running on PC2 is able to capture the transmission with a sniffer (as long as the sniffer is running in promiscuous mode).

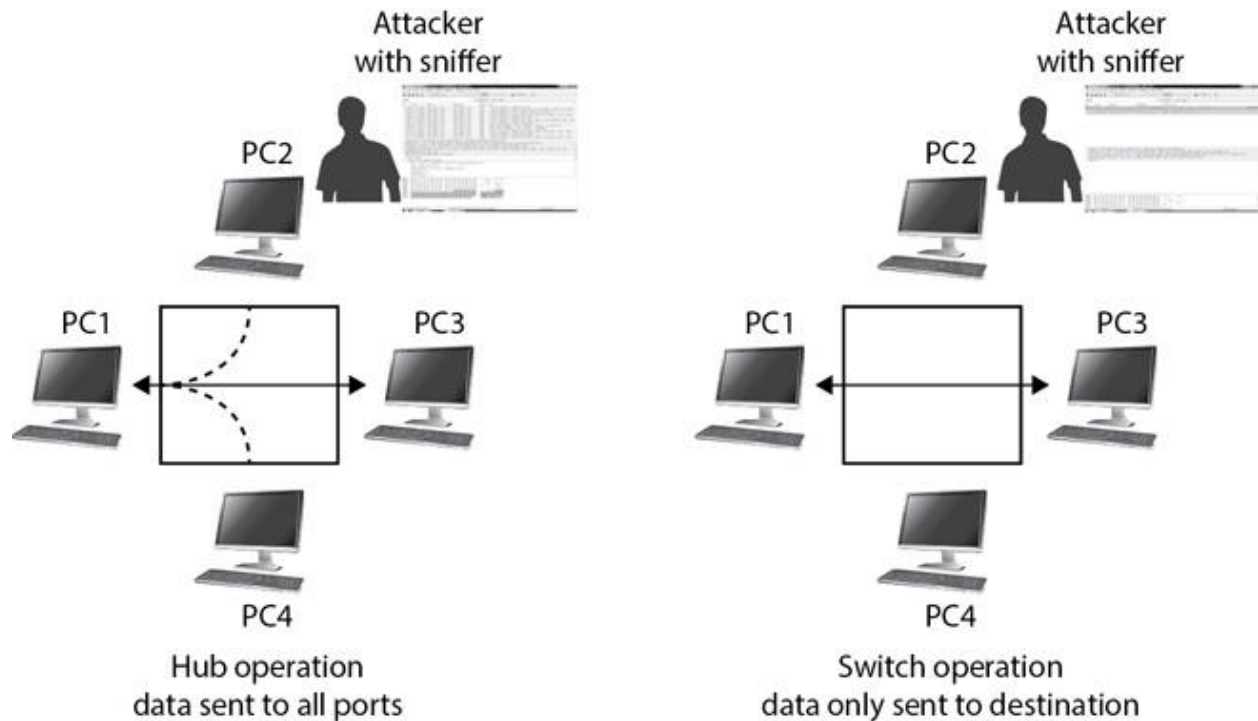


Figure 5-6 Comparing data sent via a hub and via a switch

However, when using a switch, the same data transmissions reach only PC1 and PC3. The attacker running a sniffer from PC2 is not able to capture the data because the traffic doesn't ever reach the switch port for PC2.

It's worthwhile noting that administrators can configure a mirrored port on most routers and switches. The device sends a copy of all data sent through the device to the mirrored port. If attackers can connect directly to the mirrored port, they will be able to capture all the traffic going through the device. This underscores the importance of physical security. Locating network devices in secure server rooms or locked wiring closets provides a necessary layer of security.

Reconnaissance and Fingerprinting Attacks

Many attackers gain a foothold into a network by first infecting a single system, often after a user responds inappropriately to a phishing e-mail. The attacker then uses various reconnaissance methods to identify other computers on the network by identifying active IP addresses. Next, the attacker attempts to learn details about each active IP address using fingerprinting methods such as port and service scans, and/or OS detection techniques.

Ping Sweep and Other IP Identification Methods

Attackers often use a ping sweep (sometimes called a ping scan) as part of a reconnaissance attack to identify active IP addresses. A ping sweep uses ICMP to identify what systems are operational within a range of IP addresses. This is similar to sending a ping to a system and waiting for a response. However, instead of sending a single ping to a single system, the ping sweep sends pings to all the IP addresses within a range, such as between 192.168.1.1 and 192.168.1.254. Tools can typically be configured with delays between each ping sent out in an effort to evade detection. However, many IDPSs can still detect this activity.

A ping sweep can be used as part of an overall attack strategy. The ping sweep performs reconnaissance to identify the system. Other tools (such as port scanners or vulnerability scanners) identify details on individual systems in a fingerprinting attack.

An Address Resolution Protocol (ARP) scan uses ARP instead of ICMP to identify active systems in a network. It sends out an ARP request with an IP address, and if a system on the network has the IP address, it responds with its MAC address. ARP uses broadcast and routers do not pass broadcasts, so an ARP scan is only effective at locating computers within a single network, or subnet.

It's also possible to use a SYN stealth scan to identify active IP addresses. If you look back at [Figure 5-1](#) you can see the three packets sent in a typical TCP handshake. In a SYN stealth scan, the attacker sends the first packet (SYN) to an IP address. If a computer is actively using the IP address, it responds with a SYN/ACK packet. Instead of sending an ACK packet to establish the TCP session, the attacker then sends an RST packet to reset the connection. However, the attacker now knows a system is actively using the IP address.

Port and Service Scans

The port scan attempts to detect what ports are open on a system. Attackers often use it as part of an overall fingerprinting attack to gather information on a system, such as what ports are open. [Chapter 3](#) covered well-known ports. As a reminder, well-known ports are in the range of 0 to 1023 and are commonly used by specific protocols. For example, SSH uses port 22, Simple Mail Transfer Protocol (SMTP) uses port 25, and Hypertext Transfer Protocol (HTTP) uses port 80.



EXAM TIP A fingerprinting attack attempts to gather information on specific computers, such as what protocols are running on the system or what specific functions the server performs. In contrast, a reconnaissance attack is a broad-based attack that attempts to identify systems on a network. For example, a ping sweep can identify computers based on their response to ping requests. A port scan on these computers provides more information on them.

If a port scan determines that port 80 is open, it indicates that HTTP is enabled on a system. Web servers use HTTP, so this also indicates the system is probably a web server. It could be an Apache web server running on a Linux system, an Internet Information Services (IIS) web server running on a Microsoft system, or something else. Attackers often use additional steps to identify the server's operating system and the web server software.

Port scanning tools can run sequentially through ports. You can detect them by checking firewall logs or by using automated IDPSs. For example, if the same source IP address checks for responses on ports 20, then 21, then 22, and so on, it's a clear indication of a port scan. More sophisticated port scanners can randomize the order of the port numbers that they check, such as 22, then 80, then 25, and so on. However, most IDPSs are aware of both tactics and can detect both sequential and randomized port scans.

Service scans go a step further to validate that a service is running. For example, if port 80 is open, it indicates that HTTP is running. The service scan then uses HTTP commands such as Get to retrieve an HTML document including header information.

Operating System Detection

Operating system (OS) detection techniques attempt to discover details about a system beyond open ports and active services. They typically analyze packets sent by the system to identify the OS and its version.

TCP/IP stack fingerprinting collects packets sent across a network and analyzes them to determine details about systems. RFCs don't require any specific values for many TCP and IP fields. Because of this, different OSs use different values. Two TCP/IP fields that can identify many systems are the initial Time To Live (TTL) field and the TCP window size field.

Some tools perform active fingerprinting by querying systems with probes. They then analyze the response to gain additional information. Many IDPSs can detect active fingerprinting techniques, and a variety of methods can thwart these attempts.



TIP Many automated scanning tools combine reconnaissance and fingerprinting capabilities. Both security professionals and attackers use these tools to scan networks and learn details about systems on networks.

Salami Attack

In a salami attack, an attacker performs a large number of minor actions that likely won't be noticed or reported but collectively can add up to big gains. Its name comes from a metaphor of slicing small slivers of salami sausage to create enough pieces to create a sandwich.

In one of the classic examples of a salami attack, an attacker shaves a penny off each of many financial transactions. Any individual who loses a penny is unlikely to notice or even complain about the loss. However, when attackers use the same tactic thousands or millions of times, they can pilfer a substantial amount of money. For example, programmers in the classic movie *Office Space* realized that they could write a program to shave fractions of pennies from transactions where they work. Unfortunately, they misplaced a decimal and ended up taking too much money, resulting in some comical results.

Salami attacks can also be used to gain information from databases. Instead of getting the entire database at a time, attackers can get small pieces of data and then combine them to reproduce the full picture. Espionage works this way when spies try to learn as much as they can using multiple sources. Similarly, social engineers (described later in this chapter) often try to gather small pieces of data from multiple sources. Both spies and social engineers can later compile all the information to gain actionable knowledge about their targets.



TIP [Chapter 11](#) delves into databases a little deeper, including the concept of inference. In an inference attack, an attacker collects multiple pieces of small, seemingly inconsequential information, but is able to put this together to infer or deduce sensitive information.

Man-in-the-Middle

Man-in-the-middle (MITM) attacks perform electronic eavesdropping by placing a system between two other parties. This third system captures the data sent between the two computers. For example, if one computer is exchanging data with a server, an MITM attack can capture all the data between the two systems. The attack is transparent to the two systems.



EXAM TIP A packet sniffer such as Wireshark is used in an MITM attack. The attacker captures packets and then analyzes the data to learn information. If systems transmit data in cleartext, the attacker can capture and read it.

An active MITM attack can be quite sophisticated, and [Figure 5-7](#) shows an example. A user is connecting to a web server over the Internet and the user's web browser indicates that a Transport Layer Security (TLS) session is established with Hypertext Transfer Protocol Secure (HTTPS). However, instead of having a secure TLS session with the web browser, the user has a secure TLS session with the MITM system.

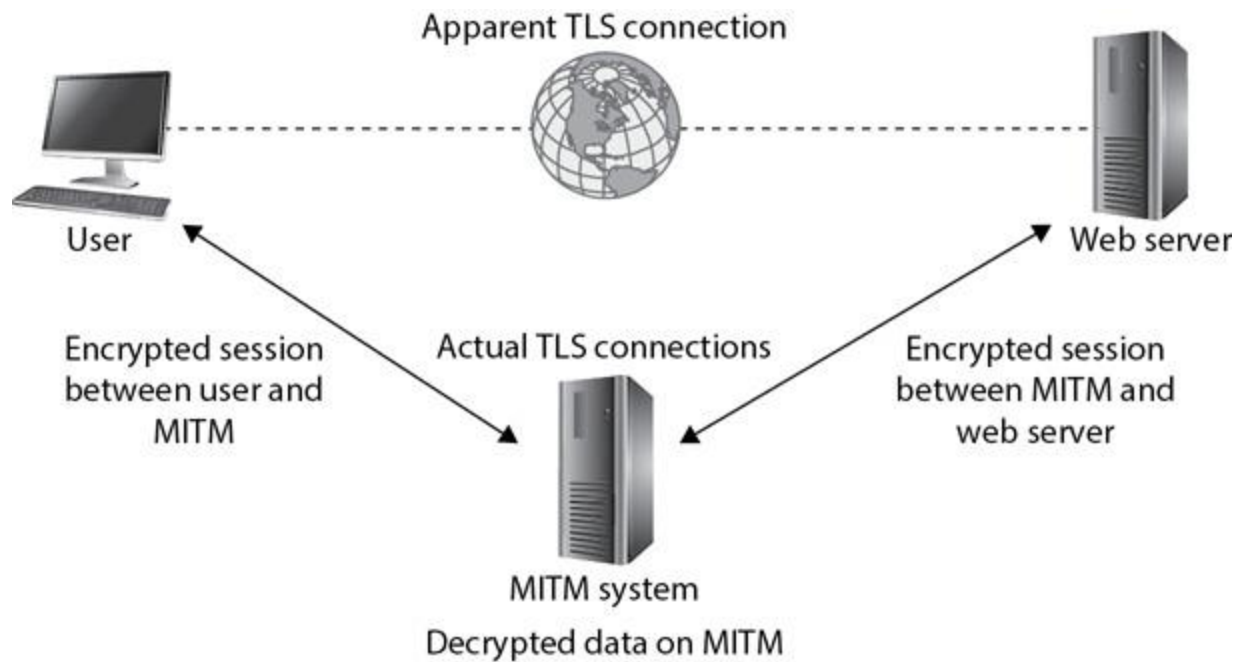


Figure 5-7 MITM attack using TLS

The MITM system establishes two TLS sessions. One session is secure between the user and the MITM system and the other session is secure between the MITM system and the web server. The MITM system is able to decrypt all of the data, allowing attackers to read it.

It's difficult to set up such an attack without control of either the user's internal network or the web server's internal network. Some organizations have set up an MITM system internally to track activity of users within the organization. For example, if the MITM system is established internally, it allows personnel to track all of the internal users' online activity. Before doing this, an organization would typically let employees know via an acceptable use policy that their online activity is tracked and that certain activity is prohibited.

If the web server is hosted at a disreputable hosting company, the hosting company may be able to implement an MITM system there. It can then capture user data such as credit card data and then run up charges on these credit cards.



TIP MITM attacks can be active or passive. In an active MITM attack, the attacker tries to capture data from specific systems. In a passive attack, the attacker attempts to capture data from any system.

MITM and TLS Decryptors

Many organizations use TLS decryptors (sometimes called SSL decryptors). These work as an MITM system, but as a monitoring method instead of an attack. Zombies in a botnet often transmit data to and from command and control servers in an encrypted format. Similarly, many attackers encrypt malware before sending it to clients. If an IDPS cannot decrypt this traffic, it would be undetected.

All Internet traffic goes through the decryptor and it can detect any HTTPS traffic. When it detects packets establishing a TLS handshake, it can use this information to establish two TLS sessions. One session is between the client and the TLS decryptor. The second session is between the client and the Internet-based web server. While data transmitted to and from the TLS decryptor is securely encrypted, it is in an unencrypted format on the TLS decryptor.

The TLS decryptor typically sends the unencrypted data to an IDPS for analysis. An IDPS can also host a TLS decryptor, but due to the processing power required, this may not be feasible. Also, most TLS decryptors include the ability to use whitelisting rules. These rules include a list of known safe sites. When internal employees create HTTPS sessions with known safe sites on the whitelist, the TLS decryptor doesn't expend processing power to decrypt the traffic.

Replay

A replay attack starts as an eavesdropping attack where the attacker uses a sniffer to capture the data. However, the replay attack goes a step further and uses the information to impersonate one of the systems in the session.

For example, two computers may establish a session and exchange authentication information. An attacker will capture this information and use it later to impersonate one of the systems.

Several methods help prevent the success of this type of attack. For example, Kerberos (a network authentication protocol) uses time stamps and requires systems to be within five minutes of each other. These time stamps restrict the time that an attacker can replay captured data. Additionally, many systems use unique session tokens within secure sessions to thwart replay attacks. For example, Challenge-Handshake Authentication Protocol (CHAP) uses a nonce (a number used once) to create session tokens. Because CHAP doesn't reuse these session tokens, attackers cannot use them in a replay attack.

Session Hijacking

Session hijacking is another sophisticated attack where an attacker captures information from an ongoing TCP/IP session and attempts to take over the session by impersonating one of the parties. Such an attack is possible with HTTP sessions where cookies are used. If successful, the attacker may gain access to either of the systems in the original session or launch a DoS attack on one of the systems to take it out of service.

The session hijacking attack often starts with a sniffing attack where the attacker gathers information about an ongoing session. The gathered information includes cookies passed back and forth between the web client and the web server.



TIP A cookie is a text file stored on a user's system. When a user visits a website, the website is able to read the cookie and use it to identify the user. The cookie also includes a session ID used to identify the current session to the web server.

As an example, Firesheep is a Firefox extension that allows attackers to listen for and exploit session cookies transmitted over a wireless network. An attacker can sit in a public wireless hotspot, capture session cookies, and use them to take over sessions such as Facebook and Twitter sessions. Faceniff is a similar application that runs on Android systems.

DNS Cache Poisoning

[Chapter 3](#) discusses the operation of Domain Name System (DNS) to resolve host names to IP addresses. DNS cache poisoning (sometimes called simply DNS poisoning) is an attack that attempts to redirect traffic away from legitimate servers.

As an example, when you type in [google.com](#) as the URL in your browser, your system queries a DNS server for the IP address. When it receives the reply, it stores the result (216.58.193.206) in the DNS cache. You can use the command **ipconfig/displaydns** to see the contents of the DNS cache on Windows systems. The Time To Live field indicates how long (in seconds) a result remains in cache. A successful DNS cache poisoning attack changes the correct cached IP address associated with a host to the IP address of another system. For example, it can change the IP address to 204.79.197.200, which is an IP address for [bing.com](#). Now, when a user types in [google.com](#), they will be taken to [bing.com](#) instead. Redirecting a user to [bing.com](#) instead of [google.com](#) might be fun as a practical joke, but it isn't really an attack.

DNS servers also have a cache. When one DNS server queries another DNS server to resolve a host name to an IP address, the DNS server receiving the response stores it in its cache. If the cache on the DNS server has been poisoned, the DNS server will send the incorrect IP address to any system that queries it.

Attackers sometimes try to redirect traffic to malicious sites. For example, they might redirect a user to a site that attempts a drive-by download or tricks a user into downloading and installing malware.

China uses traffic redirection as one of many techniques that block access to certain websites outside of China. For example, DNS servers within China redirect traffic of some popular U.S. websites to Chinese versions of the same sites. Some examples are: [twitter.com](#) is redirected to [weibo.com](#) and [facebook.com](#) is redirected to [wechat.com](#). In 2010, an Internet service provider (ISP) outside of China incorrectly configured DNS servers to receive information from Chinese DNS servers. This essentially caused a massive DNS cache poisoning attack redirecting traffic to Chinese servers.

One of the best protections against DNS cache poisoning attacks (beyond configuring them correctly) is the use of Domain Name System Security Extensions (DNSSEC). DNSSEC is a suite of extensions that validates DNS responses with a Resource Record Signature (RRSIG). When fully implemented, each DNS response includes the RRSIG, which is similar to a digital signature. The RRSIG provides assurances that the DNS record is valid and has not been poisoned.

Smurf and Fraggle Attacks

A smurf attack broadcasts ICMP ping packets to multiple computers on a network but spoofs the source address using the IP address of the attacked system. An ICMP packet normally includes the IP address of the sender in the source IP address field. However, by replacing the IP address with the victim's IP address, the ICMP packet appears to come from the victim's computer. By broadcasting the ping, all systems on the subnet receive the echo and respond by flooding the attacked system with echo replies.

Attackers often use an amplifying network in a smurf attack by sending a directed broadcast ping to the amplifying network. A directed broadcast ping goes through a router to the target network and then broadcasts the ping to all the computers on the target network. Each computer on this network then sends pings to the victim's computer. However, most routers block directed broadcasts, protecting a network from being used as an amplifying network.

A fraggle attack is similar to the smurf attack, but instead of using ICMP packets, it uses User Datagram Protocol (UDP) packets for the attack. It sends packets to UDP port 7 or UDP port 19. Port 7 is the echo port, which works similarly to a ping. Port 19 is the Character Generator Protocol (chargen) port. When chargen is enabled on a system, it will respond with a random character each time it receives any traffic on the port. However, chargen is rarely enabled on current systems.

Software Security as a Countermeasure

Many organizations host their own websites, and developers can take steps to protect these websites from various attacks. Successful attacks can infect the websites and infect user systems that visit the websites. Software security includes steps that developers within an organization can take to ensure that their software applications (including web applications) are secure.

Years ago, it wasn't common for an organization to develop its own software applications. However, now it is common for an organization to create simple web applications on its own instead of hiring an outsider. Many organizations let internal employees create them, and this practice is often successful. However, if the website developer isn't aware of basic security measures, including input validation, application reviews, code signing, and sandboxing, security problems can arise.

Using Input Validation as a Countermeasure

Input validation checks data before using it within the application. The goal is to verify that the data is valid and ensure that an application does not use invalid data. This prevents potential problems within the application and helps prevent attacks.

As an example of input validation, imagine an application needs a user to enter a number between 1 and 100 that the application will use in a mathematical equation. Input validation does a check to verify that the user actually entered a number in this range. If the user enters 0, a number over 100, or any non-numerical characters, it fails the input validation check, and the application won't use the entered data.

Applications should use input validation for any type of data. If users need to enter their first name, the application can check to ensure that the entered name is less than 35 characters and includes only valid characters for a name. Similarly, if users need to enter a five-digit ZIP code, input validation verifies that the input includes five digits.

Input validation doesn't verify the accuracy of the data. Instead, it only determines if it is valid. As an example, if you can enter 23456 as your ZIP code, a web application using input validation will accept it because it is a valid ZIP code. However, the application doesn't know if this is an accurate ZIP code for you. More advanced accuracy checks can verify that the ZIP code you entered matches the city and state you entered, but these are accuracy checks instead of validity checks.

If you've ever filled out a form on a web page, you've probably seen input validation in use. After you enter the data and submit it to the website, the web page checks your input. If something is wrong, the site redisplay the web page, often using something like a red asterisk (*) to show what items need to be addressed.

Input validation is valuable in preventing different types of attacks, such as buffer overflow, SQL injection, and cross-site scripting attacks. These attacks are discussed later in this chapter.



EXAM TIP Applications use input validation to check user input before using the data. Input validation often checks for specific characters such as < and > to prevent cross-site scripting attacks and can help prevent buffer overflow attacks.

Client Side vs. Server Side Input Validation

Input validation can be done on the client side and/or on the server side. With client-side input validation, the web page includes embedded code that validates the user input. If the user enters incorrect data, the code within the web page validates it and modifies the display to let the user know the data isn't valid.

Client-side validation saves a round trip to the server over the Internet. In other words, the web page doesn't send the input over the Internet to the web server, requiring the web server to send a page back to the user to correct the input. However, a malicious attacker can modify the web page on the client side and bypass the client-side validation. Because of this, client-side validation by itself doesn't provide reliable protection.

Server-side input validation validates the data when the server receives it. If the input is invalid, the server rebuilds the web page and sends it back to the client, typically with a message indicating what the user needs to correct. Web pages often use both client-side input validation and server-side input validation. The client-side input validation saves bandwidth and processing power on the server. The server-side input validation provides a final security check before using the inputted data.

Application Review as a Countermeasure

An application review attempts to discover vulnerabilities with an application before an organization releases it. One of the basic elements of a review is testing. A tester verifies that the application works as it should, and tests for any bugs or vulnerabilities.

Application reviews also include a review of the source code. This is a line-by-line examination of the source code of an application. The primary goal is to discover potential issues that developers can correct before the organization releases the application. Some organizations also use application reviews as a peer review to help developers improve their skills.

Code Signing as a Countermeasure

Code signing, or digitally signing software, provides security for the organization hosting a website by associating a certificate with the software. It also provides users with assurances of who wrote the software and that the software has not lost integrity. If malware infects digitally signed software, it will fail a check and alert the user. In some cases, the software is within a web page and the browser will raise an alert. In other cases, the operating systems will detect the failed check during the installation and raise the alert. Note that code signing verifies both the identity of the author and file integrity.

Vendors that release software purchase a certificate from a public certificate authority (CA). This certificate is associated with the software file and the vendor that released it. It includes the name of the vendor and a hash of the file.

When users download the software file, they also download the certificate. When they install the software, the system verifies the certificate matches the file and that the certificate is valid. [Chapter 14](#) discusses how a system validates a certificate, but in short, it queries the CA to ensure the CA issued the certificate, and that the CA has not revoked it. The system then displays a dialog box that displays information about the certificate. If the code doesn't have a certificate, the dialog box indicates the vendor is unknown. If malware infects the file, it changes the hash, and the certificate will no longer recognize the file. Instead, the system reports an error.

Although many legitimate software vendors are willing to purchase certificates to provide assurances to their customers, malware authors are not. If criminals purchased certificates for malware they wrote, law enforcement would be able to identify the criminals relatively easily. However, criminals don't want to lead law enforcement personnel to their door, so they don't purchase certificates and historically haven't digitally signed their malware.

However, in the past few years, criminals have been signing malware with code signing certificates that they stole from others. Symantec introduced Extended Validation (EV) Code Signing certificates, and the CA/Browser Forum released standards for these in 2012. EV Code Signing certificates provide a more thorough background check with the intent of providing users with a higher level of confidence for code signing certificates.

With this in mind, it's a safe bet that digitally signed software using EV Code Signing certificates is relatively safe. However, other code may not be safe. This doesn't mean that all unsigned software is malicious, but unsigned software should be viewed with some suspicion. Additionally, software using standard (not EV) code signing methods should be viewed with some suspicion.

Both Java applets and ActiveX controls (commonly used in websites) support code signing. Most browsers have settings that can restrict the use of unsigned software, or at least provide users with a warning if a website tries to run unsigned software.



TIP Configuring web browsers to block unsigned software greatly increases security. Unfortunately, some legitimate websites that use Java applets and ActiveX controls don't go through the expense of purchasing certificates for them. In other words, restricting a web browser to only digitally signed software increases security but often reduces usability.

Malicious ActiveX controls can gain a significant amount of control over a Windows operating system. ActiveX controls are largely limited to Microsoft Windows systems, but some ActiveX controls can run in Internet Explorer running on other operating systems. Java applets are theoretically portable to any operating system. However, due to differences in operating systems, the portability isn't 100 percent.

Sandboxing as a Countermeasure

Sandboxing runs computer programs in isolated areas of memory as a security control. Antivirus programs use sandboxing to check new or untrusted applications. If antivirus software determines that a program is malicious, it blocks the application from running normally. This prevents a malicious program from infecting a computer.

Security professionals often use virtualization to test applications. For example, they create a virtual machine, and then isolate it from the physical computer and other virtual machines. They

can then run the application within this isolated virtual machine. If the application causes problems, the problems are limited to the virtual machine.

Buffer Overflow Attacks

A buffer overflow occurs when an application receives more data than it expects and is unable to handle it gracefully. Applications have areas of memory called buffers that they use for temporary data storage. These buffers are specifically defined areas of memory, and applications are normally unable to access memory spaces beyond the buffers. However, there are times when an application gets overwhelmed by the amount of data sent to it, allowing it to access memory beyond the defined buffer.

A buffer overflow error exposes normally protected memory areas. If attackers can exploit a buffer overflow error, they are able to write additional data into the overflow area. Attackers often include a large volume of no operation (NOOP) commands in an attack to position the memory pointer. Attackers can then add malicious code after the NOOPs, allowing them to insert the malicious code into memory.



EXAM TIP A buffer overflow problem starts as a programming error. However, if attackers detect it, they can launch a buffer overflow attack. Successful buffer overflow attacks allow attackers to gain elevated privileges and insert malicious code into normally protected memory areas.

As an example, an application may expect a user to enter some specific data, such as a number between 1 and 10. An attacker may try to send different data, such as an escape sequence and multiple NOOP commands, to see what happens. If the different data is not expected, it can cause an error. If the application doesn't handle the error gracefully, the attacker can try to repeat the error but add more data. This additional data can include a malicious script that installs itself on the server.

Proper error handling within an application ensures that it can handle errors gracefully. Error handling routines provide user-friendly errors to users and log details giving administrators insight into what happened. They also ensure the application doesn't crash.

Input validation helps prevent buffer overflow attacks. Additionally, keeping systems updated protects against buffer overflow attacks. When a buffer overflow vulnerability occurs, vendors release patches that fix the problem. Patched systems are protected against known buffer overflow attacks. Of course, if patches are not applied, the system remains vulnerable.



EXAM TIP A key to keeping a system secure is to keep it updated with current patches and fixes. This protects against buffer overflow attacks and many other types of attacks.

Injection Attacks

Injection attacks include SQL injection and command injection attacks. In either type of attack, the attacker attempts to inject code into an application. If successful, the code runs a command outside of the application. Input validation is a primary prevention against these types of attacks.

SQL Injection

SQL (Structured Query Language) is the language used to communicate with databases. SQL statements provide the ability to read, write, modify, and delete data within a database. In many cases, programmers create richer websites by interacting with a database, and they use SQL statements to do so.

Normally, this works fine. However, many malicious attackers know about SQL and they can sometimes use these same SQL statements to interact with a database through an application or an online website. As an example, when entering data into a website form, an attacker can enter a specially crafted string of code that includes a valid SQL statement. This allows attackers to access the backend database. To get a better idea of how a SQL injection attack can access data in a database, it's useful to understand how dynamic SQL works.

Dynamic SQL Imagine a website that expects a user to enter a username and password into two text boxes on a website. If the credentials match, the website can retrieve the user's name from a table named Authentication. The two text boxes on the web page are named txtName and txtPassword, and the following dynamic SQL statement will retrieve the user's first name:

```
SELECT FirstName FROM Authentication  
WHERE login = 'txtName.Text' and password = 'txtPassword.Text';
```

The statement will take the data entered in both text boxes and dynamically add it to the SQL statement. For example, if Darril enters **Darril** as the username and **P@ssw0rd** as the password, the SQL statement is interpreted like this:

```
SELECT FirstName FROM Authentication  
WHERE username = 'Darril' and password = 'P@ssw0rd';
```



TIP This example shows how dynamic SQL can be used. Although such code will work, and developers who don't know any better sometimes use it, it is not recommended. Attackers can exploit sites using dynamic SQL statements with a SQL injection attack.

SQL Injection Attack An attacker with a little knowledge of SQL can enter specific strings of characters to access the database, and possibly cause damage. Two pieces of knowledge about SQL syntax are important here:

- A semicolon (;) indicates the end of a SQL statement.
- Two hyphens (--) indicate a comment that is ignored.

With this knowledge, the attacker can enter different data in the txtName and txtPassword text boxes discussed in the previous section. Suppose the attacker enters **whoever' OR 5=5; --** in the txtName text box and **password** in the txtPassword text box as part of a SQL injection attack. Dynamic SQL uses these entries to create the following statement:

```
SELECT FirstName FROM Authentication
WHERE username = 'whoever' OR 5=5; -- and password = 'password';
```

The WHERE clause checks for a login name of *whoever* in the Authentication table, but also includes an OR statement checking to see whether 5 equals 5. Because 5 always equals 5, the WHERE clause is always true. The statement retrieves the first name of all the users in the Authentication table. The semicolon ends the statement, and the two hyphens mark the rest of the text as a comment, which SQL ignores.

Although retrieving first names from the database isn't necessarily a problem, a SQL injection attack allows an attacker to inject any SQL statements. Attackers can retrieve all the data from the database, modify data, and even delete data.

Input validation helps prevent this type of attack. For example, a username should not include double hyphens or semicolons, and an input validation check can reject input using these characters.

Database developers also use stored procedures instead of dynamic SQL statements. A *stored procedure* is a group of SQL statements that executes as a small program. Instead of plugging the user input directly into the SQL statement, the data is sent to the stored procedure. The stored procedure handles the data differently and prevents a SQL injection attack.



EXAM TIP Input validation and stored procedures help prevent SQL injection attacks.

Some SQL injection attacks have infected systems with malware. For example, the LizaMoon SQL injection attack infected websites with scareware. When users visited the infected websites, they saw a pop-up window that indicated their systems were infected with malware. It encouraged users to install rogue antivirus software named Windows Stability Center (which is not a Microsoft product). Users who installed the bogus antivirus software were then encouraged to purchase the full version to remove reported malware.

Command Injection

In a command injection attack, an attacker attempts to insert operating system commands into an application or web page form. If successful, the command executes on the system hosting the application or web page.

A simple operating system command is **dir**, which lists the files in an operating system directory or folder. An attacker may be able to insert a specially crafted string of characters into a web page form, followed by the **dir** command. The string of characters provides access to the operating system command shell and then executes the **dir** command.

Of course, the attack can be much more malicious than just retrieving a list of files. If successful, the attacker is able to execute any command supported by the operating system. As with other injection attacks, input validation is the primary prevention.

Cross-Site Scripting

Cross-site scripting (XSS) is a relatively common attack that attempts to inject HTML or JavaScript code into a web page. The attacker modifies the code on the website, and when the user goes to the website, the code executes on the user's system.



TIP XSS is used instead of CSS as an acronym for cross-site scripting because CSS is the acronym for Cascading Style Sheets. Web pages use Cascading Style Sheets to control the look and formatting of HTML data.

In most cases, the website is a legitimate one, but after the attack, it can become a malicious website for visiting users. The attack doesn't cause direct problems for the website, but it can cause problems for the users and indirect problems for the website.

XSS attacks attempt to steal sensitive information from users' systems, such as cookies or other information. In some cases, the cookies can include session information, and an attacker can use this to launch a session hijacking attempt. If the website developer used the cookie to store user information such as passwords or credit card data, the attacker now has access to this information.

As with other types of injection attacks, a primary protection is server-side input validation. In this case, the website application can filter out any attempts to include HTML or JavaScript tags using the < and > characters. However, developers need to be much more proactive when developing web applications to prevent XSS attacks. The Open Web Application Security Project (OWASP) has published a list of rules developers can follow to prevent XSS attacks in their XSS Prevention Cheat Sheet. You can find it by going to <https://www.owasp.org> and searching on **XSS Prevention Cheat Sheet**.

Cross-Site Request Forgery

Cross-site request forgery (CSRF, pronounced as *sea surf*, and sometimes identified as XSRF) includes “cross-site” in the name, but it is different from an XSS attack. An XSS attack attempts to download malicious code from the website to the user’s system. A CSRF attack attempts to send a malicious command from the user’s system to the website. Both occur without the user’s knowledge.

In a successful CSRF attack, the attacker tricks the user into clicking a malicious link that includes a command. The website treats the command as if the user sent it and executes it.

As an example, consider an HTML query. HTML queries can include additional data. For example, if you launched Google, you could enter the following text into Google to do a search: **Darril “Security+”**. You can also do the same search with the following query embedded in a hyperlink:

```
https://www.google.com/search?as_q=Darril&as_epq=%22Security%2B%22
```

This query includes several elements worth pointing out:

- **?as_q** The question mark character indicates that the appended HTML link includes a query, and as_q indicates the query is using Google’s advanced search function.
- **=Darril** This equal sign identifies a search term, which is Darril in this example.
- **&as_epq=%22Security%2B%22** The ampersand indicates another search term, and as_epq indicates it is using Google’s advanced search function with quotes. It includes %22, which is the hexadecimal code for a quote symbol, and %2B, which is the hexadecimal code for the plus symbol. It looks like “Security+” in the query.

Websites often use the same type of HTML queries to send data entered in one page to a second page. For example, if mcgraw-hill.com allowed users to provide information such as their name and e-mail addresses, it could also have a web page that allows users to change their e-mail address. Users could log on, enter their new e-mail address in the form, and click a button to submit the change. When users click the button, it would send a query similar to this:

```
http://mcgraw-hill.com/change?action=set&key=e-mail&value=newe-mail@gmail.com
```



NOTE This is only an example and will not work on the McGraw-Hill website.

In some cases, the user doesn’t even need to be logged on to take this action. Many websites use cookies. When the user revisits a website, it reads the cookie and logs the user back on. If a user

clicks a link or enters the string into his or her web browser, the website can identify the user by reading the cookie and taking the requested action.

An attacker can use this information to craft a different HTTP link to change the user's e-mail address to the address of the attacker. Assuming the previous McGraw-Hill query actually worked, the attacker's query would look something like this:

```
http://mcgraw-hill.com/change?action=set&key=e-mail&value=attacker@attacker.com
```

Once attackers craft the HTTP link, they only need to encourage the user to click it. The link can be embedded in a phishing e-mail or in any type of HTML object, such as an image that includes a link. When the user clicks the link, it sends the command to the website.

Changing an e-mail address may sound trivial. However, if the site doesn't use security to validate the user before making a change, attackers may be able to change other types of information, such as the user's password. In some cases, CSRF attacks may enable attackers to make purchases on behalf of the user.

Websites protect against these attacks by verifying the user before making such a change. Most professional web developers know better than to take some actions without validating the user again. In other words, they use a separate step to log a user on instead of trusting the cookie actions. As an example, if you've ever bought anything from Amazon.com using your computer, your computer has an Amazon cookie that allows Amazon to identify you when you visit. However, if you try to make a purchase or modify any of your account details, Amazon challenges you to enter your credentials again. This prevents a successful CSRF attack.

One of the best protections users can take is to avoid clicking unknown links. This is especially true for any links received via e-mail from an unknown source.



EXAM TIP An XSS attack exploits the user's trust of a website and downloads code onto the user's system. A CSRF attack exploits the trust that a site has in the user's browser and attempts to send unauthorized commands from the user's system to the website.

Password Attacks

Attackers often try to learn passwords through different methods. If the attacker is able to learn the username and password of an account, the attacker can then impersonate the user by logging on with the user's credentials. A password cracker is a program that tries to discover, or crack, passwords.



EXAM TIP Strong, complex passwords can thwart most password attacks. A strong, complex password is sufficiently long and includes multiple character types such as uppercase letters, lowercase letters, numbers, and special characters. [Chapter 2](#) covers passwords in more depth.

Hashed Passwords

One of the ways that passwords are protected is by hashing them. A hash is simply a string created by executing a hashing algorithm against the password, and the hash will always be the same if the password is the same. For example, imagine a password is P@ssw0rd and the hash of the password is 1A2B3C. No matter how many times you execute the hashing algorithm against this password, it will always create a hash of 1A2B3C. (The actual hash would be longer, but I've shortened it here for brevity.) Instead of storing the password in cleartext, the hash of the password is stored in an authentication database.

When a user logs on, the system hashes the password entered by the user and sends the hash to an authentication server holding a database of hashed passwords. Credentials are typically encrypted when sent over a network to prevent a sniffing attack from capturing the hash. The authentication server then decrypts the credentials and compares the hash of the password provided by the user with the hash stored in the authentication database. If they match, the user is authenticated.

Older hashing functions such as MD5 are not recommended for hashing passwords. Instead, a strong hashing algorithm such as Secure Hash Algorithm 3 (SHA-3) should be used. SHA-3 can create hashes up to 512 bits long.

Common Password Attacks

The following list describes some common password attack methods:

- **Brute-force** In a brute-force attack, the attacker tries all possibilities until getting a match. This can be quite time consuming when performed manually, but password-cracking tools automate the attempts. Brute-force efforts generally take the most time and effort when compared with other password attacks. This attack can be thwarted by using strong passwords of sufficient length and with multiple character types. The longer and more complex a password is, the harder it is to crack with a brute-force attempt.
- **Dictionary** A dictionary attack attempts to use all the words in a dictionary or a database of different words. Password dictionaries usually include words from multiple languages in addition to commonly used passwords such as 123456, iloveyou, qwerty (the first five letters next to the TAB key), 7777777, letmein (let me in), trustno1 (trust no one), and qazwsx (the first six letters from top to bottom on the left side of the keyboard). Strong, complex passwords thwart dictionary attacks.

- **Social engineering** The simplest way to get a password is to ask for it. Although most users won't give up their password easily, social engineers often trick users. For example, phishing attacks (described later in this chapter) often ask users for their username and password, and some users respond.



TIP Dictionaries are available in multiple languages, so using words from different languages doesn't make a password more secure.

Rainbow Table Attacks

A rainbow table attack is a more sophisticated attack used to learn passwords. It uses comparative analysis techniques to identify passwords. To make this attack easier to understand, it's worthwhile examining each of the steps in a comparative analysis.

Imagine that an attacker has captured a hashed password either by infiltrating a network and exfiltrating a database containing password hashes or by capturing packets containing a hash with a sniffer. In either case, the attacker now knows that the hash of a password is 1A2B3C. The attacker can then use the following steps to discover the password from the hash:

1. The attacker guesses a password.
2. The attacker calculates the hash of the guessed password.
3. The attacker compares the calculated hash with the captured hash.
 - a. If the hashes are the same, the attacker knows the guessed password is the actual password.
 - b. If the hashes are not the same, the attacker repeats the steps.

The most time-consuming part of this process is guessing another password and then calculating the hash. However, a rainbow table shortens this time considerably. A rainbow table includes two columns. One column is a list of guessed passwords and the matching column includes the hashes of the guessed passwords. A program compares the captured hash with each of the hashes in the rainbow table. When it finds a match, it retrieves the actual password from the second column.

Attackers create rainbow tables using programs such as RainbowCrack and character set definition files that define the keyspace for the password. A typical character set definition file includes the following character set definition lines:

- **numeric:**

[0123456789]

- **alpha:**

[ABCDEFGHIJKLMNOPQRSTUVWXYZ]

- **alpha-numeric:**

[ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789]

- **loweralpha:**

[abcdefghijklmnopqrstuvwxyz]

- **loweralpha-numeric:**

[abcdefghijklmnopqrstuvwxyz0123456789]

- **mixalpha:**

[abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ]

- **mixalpha-numeric:**

[abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789]

- **keyboardascii-32-95:**

[!"#%&'()*+,-./0123456789:;<=>?@ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{|}~]

- **ascii-32-65-123-4:**

[!"#%&'()*+,-./0123456789:;<=>?@ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`{|}~]

- **alpha-numeric-symbol32-space:**

[ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789!@#\$%^&*()-_+=~`[]{}|:;'"<>.,?/]

The rainbow table generation program then creates passwords using all the characters in each character set. For example, the attacker can specify a password length up to 14 characters. The program will then create all possible combinations using between 1 and 14 characters, and using only the characters in the numeric data set (0123456789). As it creates the password, it also creates the hash and stores both in the rainbow table. When finished, the program does the same with each of the remaining character sets.

It can take several hours to create a rainbow table. The actual time depends on how many character sets are used, the length of the passwords, the hashing algorithm, and the processing power of the computer creating them. Once the rainbow tables have been created, attackers can use them to crack passwords using a simple rainbow table lookup process.

Several rainbow tables are available using different hashing algorithms, different character sets, and targeted against specific operating systems. Attackers can download the rainbow tables from sites that have already created them, or create the rainbow tables on their own. Although rainbow tables can be very large (such as 864GB for nine-character passwords hashed with the SHA-1 hashing function), they significantly reduce the time required to crack a password.

A primary protection against rainbow table attacks is key stretching. Key-stretching techniques salt the password before hashing it. A salt adds a random number of bits to the hashed password. In other words, an unsalted hash of P@ssw0rd may be 1A2B3C, but when it's salted, it could be 3602FD or just about any other hexadecimal number. [Chapter 14](#) discusses key-stretching techniques such as Password-Based Key Derivation Function 2 (PBKDF2) and bcrypt.



EXAM TIP Some password crackers use rainbow tables to speed up the process of cracking a password. Salting hashed passwords thwarts rainbow table attacks.

Spam

Spam is unsolicited commercial e-mail (UCE) and other unwanted electronic messages sent in bulk. It is most often associated with e-mail, but attackers send spam via other means, such as instant messaging. Companies have used spam to send unwanted advertisements. However, the bigger risk is that attackers now regularly use spam to launch different types of attacks.

The cost of sending spam is low for spammers. A single e-mail can include multiple e-mail addresses, making it relatively easy to send an e-mail to thousands of recipients at a time. Additionally, cyber criminals can rent access to botnets to send spam instead of creating their own botnet.

Criminals harvest e-mail addresses using a variety of methods. For example, some programs called *web spiders* crawl through the Web looking for the @ character to identify e-mail addresses. If you post messages on a website with your e-mail address, these spiders will likely find it and add your e-mail address to a spam list. These criminals then sell or rent the lists to others wishing to send spam.

Phishing Attacks

A phishing attack is spam with malicious intent. It encourages a user to respond with personal information or to click a link. Attackers send massive numbers of e-mails with the goal of fishing

for victims. Phishing e-mails often claim to be from a legitimate company and encourage the user to either click a link or provide a response.

Here's an example of a phishing e-mail that attempts to get users to respond with their logon credentials for a bank account.

Warning! This is a security alert from YourBank. We have noticed suspicious activity on your account. To protect you, we are putting a hold on your account until we can verify certain activities. Please click the following link and validate your account. If you do not validate your account, funds in your account will be frozen and remain inaccessible.

If you click the link, it takes you to a bogus website that looks similar to YourBank's website, but is actually a malicious website. It includes logon text boxes where you can enter your username and password. Of course, if you enter your credentials, the attacker will use them to remove your money.

Attackers quickly funnel money drained from accounts this way to offshore accounts. If the activity isn't discovered within 24 hours, the chances of recovering the money are extremely low. Thankfully, banks typically provide fraud insurance for personal bank accounts, so the account holders receive their money back if it is stolen. However, business accounts don't have the same protection. If an employee accidentally gives up credentials for the organization's bank accounts, attackers can drain the accounts. This can cause a financial catastrophe if the organization doesn't have fraud insurance.



TIP Phishing attempts have been circulating for years and most IT professionals, and even a lot of end users, know about them. Yet, they continue. Why? The short answer is that they continue to work. Many end users simply don't understand the risks and continue to respond to these e-mails. When users are educated and stop responding to the phishing e-mails, attackers will stop trying these methods (but they're sure to try different ones). In the meantime, it's important to educate users about the risks.

You've probably received some phishing e-mails that look like they are from a bank or a company where you don't have an account. Attackers send out mass quantities of these e-mails knowing that not every recipient has an account with the bank, but also knowing that many of them do. Imagine that an attacker sends out an e-mail to 10,000 recipients but only 1,000 of the recipients (10 percent) actually have an account with that bank. If an attack reaches 1,000 potential targets, it has some great possibilities of success.

Even though phishing has been going on for a long time, it still fools many people. Google released the "Handcrafted Fraud and Extortion: Manual Account Hijacking in the Wild" study in late 2014. The study found that phishing attacks succeed between 3 percent and 45 percent of the

time, with sophisticated attacks enjoying the most success. Additionally, the study found that attackers used the stolen credentials within 30 minutes after users provided them.

A phishing e-mail usually has several basic components:

- **Impersonation** A phishing e-mail impersonates a legitimate company. It may include actual graphics from the company and use similar fonts. For example, an e-mail that claims it is from PayPal may include PayPal graphics.
- **Identification of a problem** The e-mail usually indicates some type of problem that is of a concern to the reader. The phrase “suspicious activity” or something similar is often included in these e-mails.
- **A sense of urgency** Recipients are encouraged to take immediate action. Although the phishing e-mail doesn’t state it this way, the goal is to get the user to click without any more thought.
- **Dire consequences** The phishing e-mail warns of some dire consequences, such as closing an account, locking out an account, or freezing funds, if the recipient doesn’t respond. For individuals doing business through a company such as eBay or PayPal, the threat of their account being closed, or locked, sounds dire. This often propels them into taking immediate action.

Phishing and Drive-by Downloads

Some phishing e-mails just want users to click a link. If they click, the link takes them to a malicious website that attempts to start a drive-by download. Notice that in this case, it doesn’t matter whether a user enters any personal information or not because the attacker only wants the user to click the link. Drive-by downloads attempt to install malware and, when successful, often join the computer to a botnet.

Some phishing e-mails mention a failed delivery attempt from a service such as FedEx, and encourage the user to click the link for more information. Other e-mails have impersonated funeral homes. They mention a pending funeral and encourage the user to click for more details. These are just two examples, but attackers have used many more methods.

Spear Phishing and Whaling

Spear phishing and whaling are variants on the phishing theme. Phishing attacks send out massive amounts of spam e-mail without knowing the targets. In contrast, spear phishing and whaling attacks attempt to identify the targets.

A spear phishing attack targets a specific organization or group. For example, a spear phishing attack could target a military organization by sending e-mails to recipients with the appearance that Sergeant Jones sent it from within the organization. Because it looks like it’s from someone within the organization, recipients are more likely to open it. The contents could state something about a missed deadline and ask the user to provide some information in an attached document.

However, the attachment includes malware and attempts to install itself on the user's computer if the user opens it.

The From address provides a level of trust because it looks like it's coming from a fellow employee. Because of this, some users are more willing to open the e-mail and the attachment. In contrast, these users might not respond to a typical phishing attempt.



TIP Spear phishing can target any organization. If attackers are interested in money, they may try to target a bank. If they are interested in government espionage, they may try to target government employees. The key is that the spear phishing e-mails are targeted to a specific group.

In a whaling attack, the attacker attempts to identify an executive such as a chief executive officer (CEO), president, vice president, or manager. The attacker may be able to get the name and e-mail address information from public sources or through a social engineering attack with another employee. The subject line of a whaling attack may include words such as "Past Due" or "Lawsuit" that may get immediate attention from the executive.

If an attacker can get a response from a high-profile figure (a whale, if you will), the potential payoff is huge. For example, if an attacker is able to convince an executive to install keyloggers on the organization's systems, the attacker can steal highly valuable information and perhaps even account information to empty an organization's bank accounts.

Vishing

Vishing is a slight variation of a phishing attack. Instead of using e-mail, the attacker uses the telephone system or possibly a Voice over Internet Protocol (VoIP) system. The goal is similar to phishing. It attempts to trick the user into providing information that the attacker can use for personal gain.

A vishing attack starts by making a call. Some use a recording to alert people of some type of fraudulent activity with their bank or credit card, and encourage them to call a phone number immediately. If individuals return the call, an automated system often prompts them to enter information in the phone's keypad, such as a credit card number, expiration date, and more.

I once received a voicemail message mentioning suspicious activity on a credit card. I called back, and the woman who answered started asking personal questions that raised alarms for me. When I expressed distrust, the woman said, "You called us," as if that was supposed to tell me everything was OK. It wasn't and I ended the call. There weren't any problems with my credit card then, but if I kept answering questions, there would have been.

In contrast, after a business trip, one of my credit cards had some actual suspicious activity. My bank called and left a message asking me to call back. When I called back, the person provided information that made it clear she was an employee of my bank before asking me to provide information about my account. In other words, we provided mutual authentication.

Because vishing uses VoIP, it often takes advantage of additional features. For example, attackers can trick caller ID so that phone calls appear to be coming from a bank or someone else other than the attacker's phone number.

Smishing

Smishing is another variation of a phishing attack, but it uses text messages such as Short Message Service (SMS) messages commonly used with many smartphones. Attackers send a text message to users, indicating that there has been suspicious activity on an account. It then prompts the user to call a phone number to resolve the problem.

When the user calls the number, he or she often hears a voice recording that prompts the user to enter information such as an account number and personal identification number. It may prompt the user to give additional information, such as their birthdate, Social Security information, and more.

Zero Day Exploits

Zero day exploits are attacks that take advantage of vulnerabilities that are unknown to and/or unpublished by the vendor. However, even though the vendor might not know about the vulnerabilities, some attackers know about them. The two common scenarios for zero day vulnerabilities are

- **Unpublished and unknown by the vendor, but known by attackers** For example, an attacker may discover a buffer overflow vulnerability and launch an attack using this vulnerability. The attacker may share the knowledge with other criminals, and each of these attackers can launch zero day exploit attacks.
- **Known by the vendor, but not published** At some point, the vendor learns about the vulnerability. However, it takes time to develop and release a patch to fix it. For example, an antivirus company could discover an issue with a Microsoft operating system and inform Microsoft. Microsoft will investigate and write a patch to fix it, but its personnel also need to test the patch before releasing it. Attacks that attempt to exploit the vulnerability before Microsoft publishes the patch are zero day exploits.

Unpatched Systems and Zero Day Exploits

In some cases, people refer to a zero day exploit as an attack on a system shortly after the vendor releases the patch but before IT personnel apply the patch.

For example, Microsoft releases patches on “Patch Tuesday,” the second Tuesday, and sometimes the fourth Tuesday, of every month. However, IT personnel don’t always apply the patches immediately because they need to test the patches. They may not apply the patches until Wednesday or Thursday, or even later.

Attackers often download these patches as soon as they are available and get as much information on them as they can find. They reverse engineer the patch to discover the vulnerability the patch is fixing. Once they discover the vulnerability, they develop an attack method to exploit it.

Exploits of the vulnerability between the time when the vendor releases the patch and when IT personnel patch the system is sometimes called a zero day exploit. In this case, attacks on Tuesday or Wednesday after the patch is released may be called a zero day exploit.

Not everyone agrees with this definition though. Many security professionals say that if a patch is released, exploits against the vulnerability are simply attacks against unpatched systems. Still, this helps emphasize the importance of keeping systems patched and up to date. Many successful attacks take advantage of unpatched systems, and IT personnel can prevent these attacks by keeping systems up to date.

The danger from zero day exploits also helps to emphasize the importance of a defense in depth strategy. If the vulnerability is unknown to the public and the vendor hasn’t released a patch, IT personnel won’t know about it. However, other security practices can help prevent a successful attack.

For example, one of the basic elements of hardening a system, or making it more secure from the default configuration, is to disable unneeded protocols and services. Imagine that an attacker discovers a zero day vulnerability with FTP and launches an attack on a server. If the server has FTP enabled by default, even though it does not need FTP, it is vulnerable to this attack. On the other hand, if the server was hardened and FTP is disabled, it is not vulnerable to the attack.

Covert Channel

A *covert* action is a secret one that is not open or acknowledged, such as a clandestine spy operation, and a *channel* is a communications path between two computers. Combined, a *covert channel* is a mechanism used to transmit and receive information surreptitiously between two computers. Covert channels often use nontraditional communication paths, which helps them bypass traditional controls such as firewalls and IDPSs.



TIP Covert channels can be used as a secret method to pass information between two entities for malicious purposes.

For example, TCP uses a TCP handshake process to establish a TCP session. The primary pieces of information used in these packets, besides the source and destination data, are the SYN and ACK flags. However, there is room elsewhere in the TCP packet, such as in the protocol ID field, to add more data. For example, one person could add ASCII codes for individual characters in the protocol ID field before sending the TCP packet to another person. If both people know the packets include this data, they can use it as a covert channel to share information.

The ability to transfer small pieces of data in the protocol ID field is a well-known tactic. Many firewalls and IDPSs can check packets for this additional data and identify the covert channel. However, attackers can use other covert channels that aren't known to others.

Wireless Attacks and Countermeasures

Wireless networks have become very popular due to their convenience. However, convenience often has a cost, and with wireless networks, the cost is increased risk. Attackers have multiple methods they can use to attack a wireless network, but countermeasures are effective at blocking these attacks or thwarting their success.



NOTE [Chapter 3](#) covers wireless technologies commonly used with wireless networks.

WPA2 as a Countermeasure

Wi-Fi Protected Access 2 (WPA2) provides strong protection against attacks. WPA2-Enterprise uses an 802.1x authentication server and is more secure than WPA2-Personal. WPA2-Enterprise requires each user to log in with an account before they can access the wireless network. WPA2-Personal only requires users to know the SSID and a password or preshared key.

Wired Equivalent Privacy (WEP) has several significant weaknesses and should not be used. Wi-Fi Protected Access (WPA) is stronger than WEP, but WPA2 is stronger. WPA2 uses Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP). CCMP is based on the Advanced Encryption Standard (AES) and provides significantly stronger security than protocols used with WEP or WPA. WPA3 will be the strongest when it is fully implemented.

The Wi-Fi Alliance announced three security enhancements to WPA2 in January 2018. Vendors that create wireless devices are required to implement these enhancements in their devices to maintain the Wi-Fi CERTIFIED classification. Most users are unlikely to notice these changes as they are technical changes that aren't readily apparent. The Wi-Fi Alliance's announcement didn't include a timeframe of when these changes will be required by vendors. However, it will take time for vendors to fully implement these changes. The security enhancements are as follows:

- **Protected Management Frames (PMF)** PMF helps ensure the integrity of wireless network management traffic. Many devices are already using PMF.
- **Best practice checks** Device vendors are required to conduct several checks on their devices to ensure they are using best practices. The goal is to reduce common vulnerabilities caused by misconfigurations.
- **128-bit level cryptography** Devices are required to implement 128-bit level cryptographic suite configurations. This will help standardize network security.

WPA3 as a Countermeasure

The Wi-Fi Alliance announced WPA3 at the January 2018 Consumer Electronics Show (CES). While they didn't provide a lot of details, they did mention that it is the designated replacement for WPA2 and it will provide a significant upgrade to wireless security. They outlined four new security capabilities that WPA3 will provide:

- **Protection against brute-force dictionary attacks** Devices will send an alert to the network when it detects someone guessing a password. This is intended to prevent attackers from guessing a password offline and to protect users who create weak passwords.
- **Simplify the configuration process** Many wireless devices, such as IoT devices, have limited display interfaces, making it difficult to implement them securely. Users will be able to configure WPA3 devices on a network. They will initiate the process by scanning a QR code or possibly tapping a smartphone against the device.
- **Improved privacy on open networks** Users commonly connect to open networks, such as in a coffee shop, without the need for a password. However, by default, this traffic is sent over the wireless network in an unencrypted format. WPA3 devices will connect to open networks with individualized data encryption.
- **192-bit level cryptography** Devices are required to support 192-bit-level cryptographic suite configurations. These are formally defined as the Commercial National Security Algorithms (CNSA) suite.

WIDS and WIPS as Countermeasures

Wireless intrusion detection systems (WIDSs) and wireless intrusion prevention systems (WIPSs) monitor wireless frequency channels and can discover many problems. They monitor the 2.4-GHz and 5-GHz wireless frequency spectrums and can detect unauthorized access points and suspicious activity. This includes rogue access points, evil twins, MAC spoofing, and traffic patterns from known attacks. When they detect an unauthorized or suspicious activity, they can send alerts to notify administrators.

Rogue Access Points

A rogue access point (AP) is an unauthorized wireless AP set up within an organization's premises. Attackers connect the rogue AP to a wired network and then use it as a sniffer to capture traffic. After installing the rogue AP, attackers can collect traffic from a distance, such as in the organization's parking lot. The attacker simply uses a laptop to connect to the rogue AP and collects the traffic.

Physical security is a primary countermeasure against rogue APs connecting to a network. Protecting network devices (such as routers or switches) in secure server rooms or locked wiring closets prevents an attacker from connecting the rogue AP to the wired network.

Attackers sometimes set up rogue APs in public places such as coffee shops. They leave them open so that any user can connect and access the Internet via the AP. Often they'll name the AP something similar to the location (such as CoffeeAndFreeWiFi). However, the attackers can capture all data sent to and from the user.

Wireless spectrum analyzers can detect rogue APs. For example, Fluke's AirMagnet WiFi Analyzer Pro allows administrators to monitor all the channels in the 2.4-GHz and 5-GHz ranges to detect activity. It will identify rogue APs and monitor the overall health of the wireless network. Two other tools used to detect rogue APs are WIPSs and WIDSs, as described in the preceding section. Fluke's analyzer is an example of a WIDS.

Evil Twins

An evil twin is a type of rogue access point that appears to be legitimate. Attackers use the evil twin to trick users into connecting, and when users connect, attackers collect the data. Attackers use evil twins in one of two ways:

- **Within an organization** When evil twins are used within an organization, attackers name the APs with the same or similar service set identifier (SSID) names. Users are more likely to connect to APs with similar SSIDs. Organizations use WIDSs and WIPSs to detect these APs.
- **In a public place** Attackers set up APs in public places such as airports, hotels, coffee shops, parks, and restaurants. Users connect to get free or cheap wireless access, but the attacker uses the AP to collect their data.

The best protection against evil twins is to avoid connecting to public wireless networks. However, some basic countermeasures provide protection. They include

- Confirm that the wireless network is legitimate before connecting to it.
- Do not enter any login credentials that an attacker can capture and use.
- Do not install software from the public wireless network.

A more advanced countermeasure is to use a commercial virtual private network (VPN) service. For example, CyberGhost offers both free and paid versions of its CyberGhost VPN service. Users create secure connections to the VPN service through the free public wireless network. This prevents attackers running evil twins from collecting user data.



NOTE Another way of avoiding evil twins is to skip the free wireless networks and use your own smartphone for wireless access. For example, you can configure an iPhone as a personal hotspot and then connect your tablet or laptop to the iPhone. Many smartphones have similar capabilities, allowing users to tether another device to the smartphone. The trade-off here is that this consumes bandwidth from the user's smartphone data plan, which isn't free.

Wardriving

Wardriving is the practice of looking for wireless networks, typically from a car. Attackers can drive around with a laptop and a wireless connection. When they find a wireless network, they attempt to access it.

The best countermeasure against war driving attacks is to use secure wireless protocols. This includes WPA2. WPA2-Enterprise uses an 802.1x authentication server and is more secure than WPA2-Personal, which only requires a preshared key or password.

WPA Cracking Attacks

Attackers use WPA cracking attacks to learn the password for a wireless network. They use a wireless sniffer to capture wireless packets and look for a WPA four-way handshake (indicating a user has logged on). The handshake includes an encrypted version of the AP's password. Once attackers capture the encrypted password, they use an offline brute-force attack to discover the password.

It might take some time before a user authenticates and initiates the WPA handshake. To speed up the attack, attackers send data onto the network, causing the user's system to lose its association with the AP. This requires users to log on again, initiating the four-way handshake.

The best countermeasure against WPA cracking attacks is to use WPA2. If WPA2 is used, a strong wireless network password helps thwart WPA cracking attacks. Passwords can be up to 63 characters long.

WPS Attacks

Wireless access points typically include Wi-Fi Protected Setup (WPS). WPS allows users to configure wireless devices by clicking a button on the access point and clicking a button on the wireless device. These are typically physical buttons, but in some cases they are logical buttons

accessed via an interface, such as a web page. Other times, users configure the devices by entering a personal identification number (PIN). While using the buttons is an easier method for users, WPS always supports the use of a PIN.

A WPS attack uses brute force to discover the PIN. Because WPS doesn't include a lockout feature, attackers can keep guessing until they discover the correct PIN. When they guess the correct PIN, the access point responds with the SSID and password used for the access point. Reaver is an open source tool that automates the process. It can discover an eight-digit PIN in 10 hours or less.

The best countermeasure is to disable WPS when it is not in use. You can usually disable WPS within the AP interface. Unfortunately, this isn't always reliable. Some access points still respond to Reaver queries even with WPS disabled. If your AP is still responding to Reaver WPS queries with WPS disabled, check with the AP's vendor to see if it has released an update for the device.

Understanding Social Engineering

Social engineering is the practice of convincing people to give up sensitive data or to perform actions they wouldn't normally perform. A social engineer uses deception and fraud to trick or manipulate unsuspecting users.



EXAM TIP Although social engineering is a low-tech attack, it can still be a very effective method of gaining unauthorized access to an organization and the organization's IT infrastructure.

Sometimes a phone call to an employee can elicit a good deal of information. Attackers can learn the names of executives through public sources such as web pages, newsletters, and company brochures. Attackers can say something like, "I'm gathering some information for so-and-so," using the name of a high-level executive. This might be enough to get the employee to give out information.

With some basic information from a first phone call, attackers can build on that with other phone calls to other departments. Repeated a few times, attackers can build their knowledge, such as identifying how the organization formats e-mail addresses (perhaps as *firstname.lastname@company.com*) and learning the format of user accounts. Attackers may make a phone call to the IT department and ask something like, "I can't remember my password. Can you help me log in?" Such an inquiry may convince an uneducated help desk technician to change the password and give it to attackers. It's also possible for an unsuspecting employee to give up his or her own username and password, allowing attackers to log on with that employee's account.



TIP Pretexting is the practice of using gathered information to create another scenario and collect additional information. The pretext, or the invented scenario, increases the possibility that the person will give up additional information to the social engineer.

Social engineers appeal to someone's inherent desire to be helpful and liked. They exploit values such as courtesy and appeal to other people's vanity. Often they take the time to build familiarity and build trust. In some cases, they can invoke the name of an authority figure, such as an executive, to encourage a person to respond. If necessary, they can use intimidation tactics and threaten undesirable action to get what they want. They often create a sense of urgency to prevent people from analyzing the activity.

One of the primary methods to counteract social engineering is to educate users in the tactics of social engineers. Once employees recognize the value of information and common tactics used by social engineers, they are better able to identify and thwart a social engineer in action. The following sections identify some common social engineering tactics.

Tailgating

Tailgating (or piggybacking) occurs when someone passes through a controlled entry without providing credentials by following closely behind someone who has provided credentials. As an example, some organizations issue proximity badges to users. Employees swipe their badge in front of a proximity badge reader and it unlocks a door. All employees should use their own badge to open the door and gain entry. However, it's possible for one employee to open the door with a badge and other people to follow that person in without using their own badges. These other people are tailgating.

Imagine the beginning of the workday as the majority of the employees are coming into the building. It's unlikely that each employee will shut the door in the face of fellow employees. If the organization is large, employees may not know who is a fellow employee and who isn't. A social engineer could easily slip into the crowd and walk into the normally secure building without raising any suspicion.

Most people are courteous and consider slamming the door on someone else to be rude. When people have the choice of being either courteous or secure, they often choose the courteous path. Attackers know this. Often just a friendly smile is all that's needed if someone gives the social engineer a questioning look. Similarly, an attacker can load her arms with books or boxes and approach the door. A friendly employee may actually hold the door open for her to enter.

Educating users is a great start against tailgating, but it isn't always enough. Another way to stop tailgating is with a mantrap. The formal definition of a *mantrap* is a large enclosure that allows entry through one side and exit through the other, and can actually lock the person inside. Some are large steel bars formed as a revolving door that looks very similar to a cage.

However, instead of using cages to control entries, many organizations simplify this to turnstiles similar to what you see in a subway or bus station. Employees swipe their badge and the turnstile releases to allow one person through. When using a mantrap, even if it's a simple one like a turnstile, employees cannot keep the entrance open for a potential attacker, and it thwarts tailgating.



EXAM TIP Tailgating is a social engineering tactic where one person uses credentials to gain entry, while a second person follows closely behind without providing credentials. A mantrap protects against tailgating.

Impersonation

Another social engineering tactic is impersonation, where the social engineer impersonates someone. Impersonation can be over the phone, such as by invoking the name of someone in authority. It can also be in person, such as impersonating a technician.

For example, it's relatively easy to get a uniform that looks like it's from another company, such as the phone company. The social engineer can don the uniform, enter the building, and let a receptionist know he or she is there to fix some fictional problem. An uneducated employee might lead the attacker directly to a wiring closet, where the attacker can install a rogue access point.



EXAM TIP Impersonation is also known as masquerading or spoofing. In social engineering, the social engineer is attempting to masquerade as someone else.

Skimming

Skimming is the process of capturing information from credit cards at point-of-sale (POS) readers, gas pumps, and automated teller machines (ATMs). The attacker typically places a thin filament-based plastic, called a skimming film, into a card reader. When the victim inserts their credit card or debit card into the reader, the skimmer captures the data on the credit card or sometimes withdraws money from the debit account.

Attackers have gotten more sophisticated in these attacks, making it more difficult to detect the skimming film. Many gas pumps now have tamper-proof stickers that will be broken if an unauthorized individual opens the card reader and installs a skimming film. Experts recommend using ATMs and POS readers that are in the open and enjoy a lot of traffic. Attackers are less likely to insert skimming films in a reader that has a lot of regular traffic.

Dumpster Diving

Many attackers gain information by sifting through someone else's trash, a practice commonly known as dumpster diving. Dumpster diving can provide significant returns depending on how much information an organization or an individual throw away.

For example, if an organization regularly throws away papers that include information such as customer names, addresses, and credit card data, an attacker can retrieve this from the trash and use it. Similarly, if an individual throws away credit card offers, blank checks from a credit card company, or other personal data, an attacker can retrieve it and use it.

As a best practice, any information that contains any type of personal information should not be thrown away. Instead, these papers should be shredded using a crosscut shredder or burned to ensure they cannot be read.

Shoulder Surfing

Shoulder surfing is the practice of looking over someone's shoulder to gain information. For example, an attacker may try to watch someone enter a username and password to learn the user's credentials, or watch someone enter the numbers for a cipher lock on a door, a PIN for a badge, or even a PIN for a debit card.

Social Networking Attacks

Attackers are taking advantage of the popularity of social networks such as Facebook, Twitter, and YouTube to launch social engineering attacks. Attackers craft e-mails that look exactly like they came from the social network site and send them out to users.

For example, Facebook often sends notifications to users indicating that they have messages or pending invites and includes a link to the Facebook site. The attacker uses the same template letting the user know a message is waiting, but includes a link to a malicious website. If the user clicks the link, the website may attempt a drive-by download or encourage the user to install a Trojan horse. In other cases, the site is a fake one claiming to sell cheap drugs or casino gambling access or some other site intended to get the user to enter credit card data.

User Awareness as a Countermeasure

A primary method of combating social engineering attacks is by raising user awareness through training and other methods. When users recognize the tricks that social engineers use, they are less likely to fall for them.



EXAM TIP A primary purpose of user awareness campaigns and training is to change user behaviors from unsafe actions to safe actions. For example, if users recognize the malicious nature of phishing e-mails, they are less likely to click the link or respond.

User awareness training can't just be a one-time event, either. Instead, users need occasional retraining. Attackers are constantly modifying their attacks and attack methods. Likewise, users need updates on these different tactics, in addition to reminders of risks.

The following are some of the common methods used to educate and reeducate users:

- **Acceptable use policy** Most organizations include an acceptable use policy that informs users of the organization's expectations regarding the employee's use of computing resources. It lets users know what their responsibilities are and can include basic information on the risks of social engineering attacks. Employees normally review and acknowledge an acceptable use policy when they are hired and then periodically afterward, such as annually.
- **Initial training when hired** Many organizations include initial security training as part of a new employee's indoctrination. This could be a short presentation, a short video, an online presentation, or even a written document for the new employee to read and acknowledge. As risks and threats change, the organization modifies the initial training.
- **Annual refresher training** Refresher training reminds personnel about basic risks and trains them about new and current attack methods. Organizations often require employees to read and acknowledge an acceptable use policy at the same time.
- **Newsletters** If an organization sends out regular newsletters, it's relatively easy to add a short article or a sidebar about security. It reminds users about security while providing information on current threats.
- **Periodic e-mails** Just a simple e-mail reminding user of certain risks, with an example of a recent way that a victim was tricked into giving up information, can be useful.



NOTE Many training companies specialize in providing content for IT security. Organizations can purchase access to online training, computer-based training delivered to the desktop, videos, e-mail-based messages that an organization can tailor and send out, and even articles that an organization can include in a newsletter. By purchasing outsourced materials, the company is able to get access to current, relevant content.

Chapter Review

Traditional definitions indicate that hackers aren't malicious, but crackers are. However, the media generally uses the term *hacker* to identify anyone who launches attacks on computers or

computer systems with malicious intent. Some security professionals separate the good guys from the bad guys with terms such as *white hats* and *black hats* and refer to anyone who attacks with malicious intent as an attacker. An advanced persistent threat (APT) is a state-sponsored group that has both the capability and intent to carry out successful attacks. Insider threats often cause accidental losses by not following security policies, and allowing outside attackers into an organization.

Attackers are constantly modifying their attack strategies just as IT security professionals are constantly modifying their countermeasure strategies. Even though attacks have changed over the years, many attacks share common characteristics. By understanding past attack methods, it is easier to respond to new threats and understand the need for certain countermeasures. Basic countermeasures include patching systems, hardening systems, increasing user awareness, and implementing intrusion detection and protection systems (IDPSs).

Attackers often use spoofing techniques, such as modifying the source IP address, the source MAC address, or the From field in an e-mail. Data theft refers to any attack that allows an attacker to exfiltrate data. A denial of service (DoS) attack launches an attack on a single system from a single system. A distributed DoS (DDoS) attack launches an attack on a single system from multiple systems. A botnet is a group of computers (called zombies) controlled by an attacker through a command and control center. A botnet can control tens of thousands of computers, and some have controlled over a million computers. Criminals rent out the botnets to others for money, and these botnets are used to launch attacks and send out spam.

Some attacks attempt to gain information about a network and/or traffic sent over a network. Sniffing attacks use protocol analyzers to capture data sent over a network. If the data is sent in cleartext, the attacker can read the data with the protocol analyzer. Reconnaissance attacks (such as ping sweeps) attempt to discover active IP addresses on a network. Fingerprinting attacks (such as port and service scans and TCP/IP analysis) attempt to identify details of individual systems.

Man-in-the-middle (MITM) attacks capture traffic sent between two systems, typically with a sniffer or protocol analyzer. Advanced MITM attacks can create two TLS sessions, giving the user the impression, they have a secure HTTPS session, but all data is still viewable on the attacking system. A replay attack starts as an MITM attack. It then uses the captured data to impersonate one of the systems. A session hijacking attack captures information from a TCP/IP session and attempts to take over the session by impersonating one of the parties.

Domain Name System (DNS) cache poisoning attacks attempt to redirect users to alternate sites. Domain Name System Security Extensions (DNSSEC) helps prevent DNS cache poisoning.

Organizations that host their own websites need to implement software security. Countermeasures include input validation, application review, code signing, and sandboxing. Input validation is the most important countermeasure. It checks all input data before using it to ensure it is valid. It helps prevent many attacks, including buffer overflow attacks, SQL injection attacks, cross-site scripting (XSS) attacks, and cross-site request forgery (CSRF) attacks. A buffer overflow attack sends unexpected data to an application and exposes normally protected

memory areas. SQL injection attacks send embedded SQL statements as input to access backend databases. An XSS attack exploits the user's trust of a website and downloads code onto the user's system. A CSRF attack exploits the trust that a site has in the user's browser and attempts to send unauthorized commands from the user's system to the website.

Password attacks attempt to discover a user's password. Using strong, complex passwords helps thwart many password attacks. Passwords are typically stored in a hashed format. Older hashing functions such as MD5 should not be used to hash passwords. Newer hashing functions such as SHA-3 provide more security. Salting passwords with additional bits helps thwart rainbow table attacks.

Phishing is a form of social engineering that uses e-mail. The attacker sends out massive amounts of spam (unsolicited e-mail) hoping that someone responds by either clicking a link or responding to the e-mail with sensitive information such as a username and password. Simply by clicking the link, the responder can install malware on their system through a drive-by download. Variations on phishing attacks include spear phishing (targeted at an organization), whaling (targeted at a specific person, such as an executive), vishing (using telephones or VoIP), and smishing (using text or SMS messages).

Social engineering is the practice of using primarily nontechnical means to get people to give up sensitive data or perform actions they wouldn't normally perform. A social engineer uses deception and fraud to trick or manipulate people into giving up information they wouldn't normally give up. Tailgating occurs when one person enters a controlled space without providing credentials by following closely behind someone else who has used his or her credentials. Social engineers sometimes impersonate others, either in person or over the phone. Dumpster diving is the practice of sorting through trash to get potentially useful information that may have been thrown away.

A primary method of combating social engineering attacks is through training. The goal is to raise user awareness and change user behaviors from unsafe actions to safe actions. Training can be completed when a person is initially hired, annually, and/or periodically throughout the year. Training can consist of live presentations, online presentations, periodic e-mails, articles in newsletters, or any other means that an organization finds effective.

Questions

1. What is an APT?

- A.** A group, often sponsored by a government, that has the capability and intent to launch persistent attacks against an organization.
- B.** Software that alerts a user that their system is infected with malware, but won't remove the malware unless the user pays a fee.
- C.** An attack that redirects users to a bogus website.

D. A scan to detect open ports.

2. What is the difference between a DoS attack and a DDoS attack?

A. There is no real difference.

B. A DoS attack uses technical methods, but a DDoS attack uses nontechnical methods.

C. A DDoS attack is an attack from a single system, but a DoS attack is an attack from multiple systems.

D. A DoS attack is an attack from a single system, but a DDoS attack is an attack from multiple systems.

3. Of the following choices, what is a common DoS attack?

A. TCP flood

B. Tailgating

C. Smishing

D. Whaling

4. Thousands of computers have been infected with malware and are periodically directed to send out spam to other computers. What does this describe?

A. Zombies

B. Spear phishing

C. A botnet

D. Phishing

5. Which of the following best identifies a computer controlled by a botnet?

A. DoS computer

B. DDoS computer

C. Attacker

D. Zombie

6. An attacker is using Wireshark to capture and analyze TCP sessions. What is the best term that identifies this action?

- A. Dumpster diving
- B. Shoulder surfing
- C. Sniffing
- D. Vishing

7. A system has a protocol analyzer installed. What mode must the system operate in to capture all packets that reach it, including those that are not directly addressed to or from the system?

- A. Promiscuous
- B. Nonpromiscuous
- C. DoS
- D. DDoS

8. An attacker has infected a computer within an organization and wants to identify all active hosts within the organization's network. Which of the following methods will best meet this need? (Select three.)

- A. Ping sweep
- B. Port scan
- C. ARP scan
- D. SYN stealth scan

9. A user attempted to access <http://mcgraw-hill.com/> but was redirected to a website that advertises pharmaceutical drugs for sale. What does this describe?

- A. Session hijacking
- B. Cross-site scripting
- C. Whaling
- D. DNS cache poisoning

10. Of the following choices, what provides the best protection against buffer overflow attacks?

- A. SQL injection
- B. Input validation
- C. Cross-site scripting
- D. Code signing

11. An application has received more input than it expected and the resulting error has exposed normally protected memory. What is the best explanation for what happened?

- A. Phishing attack
- B. Salami attack
- C. Buffer overflow attack
- D. Session hijacking attack

12. What type of attack can access data in a database used by a website?

- A. Cross-site scripting
- B. Cross-site request forgery
- C. Rootkit
- D. SQL injection

13. A website is preventing users from entering the < and > characters when they enter data. What is the website trying to prevent?

- A. SQL injection attack
- B. Cross-site scripting attack
- C. Input validation attack
- D. Trojan horse

14. A user receives an e-mail indicating that the bank has detected suspicious activity on the user's bank account. The message indicates the user should log on immediately to prevent loss of funds. What is the best term to describe this attack?

- A. Sniffing

B. Session hijacking

C. Phishing

D. Tailgating

15. An attacker sends an e-mail to many members of an organization and spoofs the From address so that the e-mail looks like it came from within the organization. The e-mail tries to trick recipients into following a link. What is the best definition of this action?

A. Phishing

B. Spear phishing

C. Whaling

D. Vishing

16. A system has been attacked by an exploit that isn't published. What type of attack is this?

A. Scareware

B. APT

C. Pharming

D. Zero day

17. A user connected to a free wireless network at a coffee shop to access Facebook. Later, someone else started making posts on the user's page. What is the most likely cause of this?

A. Zero day exploit

B. WPS cracking

C. Evil twin

D. WPA cracking

18. An attacker uses nontechnical means to learn the e-mail address of a manager within a company. Which of the following best describes this attack?

A. Social engineering

B. Shoulder surfing

C. Smishing

D. Covert cramming

19. Of the following choices, what is the best method to prevent tailgating?

A. Education

B. Mantrap

C. Antivirus software

D. Access controls on the phone system

20. What is a primary goal of security-related user awareness training?

A. Increase use of e-mail

B. Change behavior

C. Implement technical solutions

D. Show how to use applications

Answers

1. A. An advanced persistent threat (APT) is a group of people (often sponsored by a government) that has the capability and intent to launch persistent attacks against organizations. Scareware is software that alerts a user their system is infected with malware, but won't remove the malware unless the user pays. Pharming is an attack that redirects users to a bogus website. A port scan is a scan that detects open ports.

2. D. A DoS attack is an attack from a single system, and a DDoS attack is an attack from multiple systems. Both typically use technical methods.

3. A. A TCP flood attack (also known as a SYN flood, TCP SYN, or TCP half-open attack) is a common DoS attack that withholds the third packet of the TCP three-way handshake. The other answers are not DoS attacks. Tailgating is a social engineering tactic. Smishing is a form of phishing using SMS messages. Whaling is a form of phishing against a single person, such as an executive.

4. C. A botnet is a group of computers that an attacker has taken over and now controls from a command and control center. The individual computers are referred to as zombies, but together they are a botnet. They may be directed to send out phishing or spear phishing e-mails, but that is the attack, not the network.

5. D. Computers controlled within a botnet are commonly called zombies. They are not referred to as DoS or DDoS computers, or attackers, although they can be directed to take part in a DDoS attack.

6. C. Sniffing is the practice of capturing and analyzing packets with a sniffer (a protocol analyzer). Dumpster diving refers to going through the trash looking for information. Shoulder surfing is the practice of looking over someone's shoulder to gain information, such as the password that a user enters to log on. Vishing is a form of phishing using telephones or VoIP.

7. A. The network interface card of the system running the protocol analyzer (or sniffer) must be in promiscuous mode. If it is in nonpromiscuous mode, the sniffer will only capture packets addressed directly to or from the sniffer. DoS and DDoS are not modes for a sniffer.

8. A, C, D. Ping sweeps, ARP scans, and SYN stealth scans are three methods that can be used to discover active IP addresses within a network. Hosts on a network have IP addresses, so this method can identify all active computers. A port scan attempts to discover what ports are open on a single system. While port scans can be run on all computers, they won't identify all hosts on the network.

9. D. A DNS cache poisoning attack attempts to redirect traffic away from legitimate servers. A session hijacking attack captures information from an ongoing TCP/IP session and uses it to take over the session, impersonating one of the parties. A cross-site scripting attack attempts to inject malicious HTML or JavaScript code into a web page. Whaling is a phishing attack that targets executives such as CEOs.

10. B. Input validation techniques validate data before using it and can help prevent a wide variety of attacks, including buffer overflow attacks. SQL injection is an attack that attempts to inject SQL code into an application. Cross-site scripting is an attack that attempts to inject HTML or JavaScript code into a web page. Code signing uses a certificate to digitally sign an application, but will not protect against buffer overflow attacks.

11. C. A buffer overflow occurs when an application receives more input than it expected and it is not able to handle the error gracefully. Attackers exploit buffer overflows to insert malware into systems. The best protection against a buffer overflow is to keep systems up to date. A phishing attack is sent through e-mail. A salami attack uses multiple small, usually unnoticeable actions, such as shaving a penny off a transaction. Session hijacking attempts to take over a session.

12. D. A successful SQL injection attack can access data in a database. Cross-site scripting injects HTML or JavaScript into a web page and runs the code on a user's system. A cross-site request forgery attack performs actions on behalf of a user without the user's knowledge. A rootkit is malware that takes over a user's system.

13. B. Cross-site scripting (XSS) injects HTML or JavaScript into a web page, and input validation techniques help prevent XSS attacks. The users are prevented from entering HTML or JavaScript tags that start with < and end with >. A SQL injection attack uses SQL code, but SQL

code does not use < or > characters. Input validation is a prevention technique, not an attack. A Trojan horse is an application that looks like it's something useful but is actually something malicious.

14. C. A phishing attack sends an e-mail to multiple recipients impersonating an e-mail from a legitimate company, indicating a problem, urging the recipient to take action, and warning of dire consequences if the recipient doesn't respond. A sniffing attack uses a protocol analyzer such as Wireshark to capture and analyze traffic. Session hijacking attempts to take over sessions and doesn't use e-mail. Tailgating is the practice of one person following another into a secure area while only the first person provides credentials.

15. B. Spear phishing is a phishing tactic that targets a specific organization. Phishing doesn't target individual organizations, but instead casts a wide net, hoping to catch someone. Whaling targets a specific individual, such as an executive. Vishing uses voice methods such as the telephone or VoIP.

16. D. Zero day exploits are attacks that take advantage of vulnerabilities that are unpublished and often include attacks that are unknown by the vendor. The other answers are known methods. Scareware is malware that scares users into thinking their system is infected with a virus and encourages them to install malware on their system. An advanced persistent threat (APT) is a group of people who have the capability and intent to launch extended attacks against organizations. Pharming is an attack that redirects users to bogus websites.

17. C. The most likely cause is an evil twin, which is a type of rogue access point (AP). An attacker likely created a free wireless hotspot in the coffee shop (perhaps on the attacker's laptop). When the user connected to it, the attacker captured the user's data, including logon credentials. This is a known attack but a zero day exploit is not widely known. Wi-Fi Protected Setup (WPS) cracking discovers the PIN of an access point and uses it to discover the access point's password. Wi-Fi Protected Access (WPA) cracking discovers the password on the access point by intercepting the four-way handshake and performing an offline brute-force attack.

18. A. Social engineering uses nontechnical (or low-technical) means to gain information, such as the names of people, e-mail addresses, and user credentials. Shoulder surfing is just looking over someone's shoulder, and although it may allow an attacker to see an e-mail address of a manager, it isn't the best answer. Smishing is a variant of phishing using SMS messages. There's no such thing as covert cramming.

19. B. A mantrap is the best method to prevent tailgating, which is the practice of one person following another into a secure area while only the first person provides credentials. Although education of employees can go a long way, ingrained courtesy sometimes overcomes security practices, and a person may actually open the door for a social engineer. Antivirus software, access controls, and the phone system aren't related to the social engineering practice of tailgating.

20. B. A primary goal of security awareness training is to change user behavior from unsafe practices to safe practices. It isn't related to applications such as e-mail, and end users aren't expected to implement technical solutions.

NOTES



CHAPTER 6

Malicious Code and Activity

In this chapter, you will learn about

- Different types of malicious code
- Common ways malware infects systems
- How attackers use remote access tools
- Stages of complex malware infections
- Differences between signature-based and behavior-based detection
- The importance of keeping antivirus signature definitions up to date
- The use of spam filters and content-filtering appliances
- How the principle of least privilege helps minimize infections
- Software security practices such as input validation and code signing
- Educating users about safe computing practices to reduce infections
- The Common Vulnerabilities and Exposures list

Identifying Malicious Code

Malicious code, or malicious software (malware), is any code or software that can be harmful or destructive to computers, networks, or the computing environment as a whole. Some of the common malicious code types include viruses, worms, logic bombs, Trojan horses, and ransomware.

Malware versions that are actively spreading and infecting computers are “in the wild.” In contrast, some researchers have discovered vulnerabilities in systems that attackers could exploit with malware. Researchers sometimes develop proof-of-concept code to demonstrate the potential exploits, but this code remains isolated within lab environments.

Virus

A virus is an application or a piece of code that causes unexpected and usually negative events on computers. Viruses have the ability to infect computers by copying themselves onto the

computer using a wide variety of methods. Once the virus is on a computer, it replicates itself and at some point executes its payload, which is the unexpected or malicious event.



TIP Many people use the term virus when referring to any type of malware. In other words, someone may be talking about a Trojan horse, a worm, a logic bomb, or even spyware, but call it a virus. This is not technically accurate, because there are differences among these forms of malware, but the misuse of the term is still common.

One of the key characteristics of a virus is that the infected file must be executed for the virus to run. Typically, this means that a user must run a program to execute the virus. For example, imagine that Bob has an infected game file on his computer. If Bob starts the game, it executes the virus, but if Bob never starts the game, the virus never runs.

Users can execute a virus indirectly too. For example, many systems support an auto-play function, though it is typically disabled on most systems today. If a user inserts a CD, DVD, or USB drive into a system, the system automatically looks for an application and runs it. If this application is infected, the virus runs and can infect the system.



EXAM TIP A virus can only run with some type of user interaction. Typically, this entails a user running a program. However, in some cases, users can indirectly run a program by inserting a USB or optical drive into a system.

Viruses typically have two primary components:

- Replication component
- Activation component

A virus replicates itself by copying its code onto other applications. When a user runs the infected application, the virus runs and looks for other applications to infect. For example, a virus may wait for a user to insert a USB drive into a system. When the virus detects the USB drive, it copies itself onto a file on the USB drive. Similarly, when a user inserts the USB drive into another system, the virus tries to copy itself onto files on the new system.

At some point, the virus activates and delivers its payload, which is the malicious part of the program. The virus could delete or corrupt data or operating system files. If the virus is part of a botnet, it could check in with a botnet server for instructions. These instructions may be to

download and install an additional virus, to send out spam, or even to launch an attack on another system. [Chapter 5](#) discusses botnets in more depth.



NOTE Historically, viruses have been destructive to the installed system. However, with the explosion of botnets and criminalization of malware in recent years, this isn't often the case today. With the exception of ransomware (described later in this chapter), a virus will rarely destroy data on the system. Instead, it tries to stay hidden on the same system as long as possible. This allows the malware to continue to collect information from the user, or continue to use the computer as a zombie within a botnet.

Viruses come in many different forms based on what they attack or how they work. The following list isn't all-inclusive, but it does describe many of the methods and forms of viruses used over the years. It's worthwhile mentioning that many viruses combine different techniques to make them more difficult to detect.

- **Stealth** Many viruses try to hide themselves from antivirus (AV) software by providing false or misleading information about the virus to the AV software. For example, a file stealth virus can misreport its actual size so the AV scan thinks the infected file is the same size as the original. Additional methods used by stealth viruses include using armor, polymorphism, or metamorphism. You can compare these techniques to how a stealth aircraft hides itself from traditional radar and infrared detection devices, and remains hidden from enemies.
- **Armored** An armored virus uses code to make it difficult for AV researchers to reverse-engineer the code. When AV experts detect a virus, they attempt to decompile the code to determine what it does and how it operates. An armored virus uses different techniques to thwart the researchers' attempts. For example, an armored virus can use a compiler in combination with encryption to make it difficult to reverse-engineer the code. Criminal virus developers often use complex and confusing code designed to hide what the virus is doing, along with encryption as a method of armor. Additionally, virus authors are constantly improving and inventing new armor techniques. An armored virus can also use stealth techniques to make it difficult for AV software to detect it.
- **Polymorphic** Some viruses can morph or mutate each time they replicate to another machine or even each time they run. Even though the file changes, the code used to replicate and deliver the payload remains the same. For example, by encrypting the file in a slightly different way each time the virus infects another file, it effectively changes the file. However, the virus replication and activation components will still behave in the same way. When the virus mutates, it's more difficult for AV software to detect it without additional virus signatures that recognize the mutated version. Virus authors use mutation engines to modify the viruses, and some viruses can mutate into thousands of different viruses.

- **Metamorphic** This is similar to a polymorphic virus, but with a subtle difference. A metamorphic virus actually mutates the code used to replicate and deliver a payload. Although the function is the same, the steps used to do so are different, making the virus more difficult to detect. In contrast, a polymorphic virus will always perform the same steps to achieve the same function, while only the file changes.
- **Boot sector** The virus is stored on the hard drive's boot sector and executes when the system is booted. This includes both cold starts and reboots.
- **Multipartite** A multipartite virus has multiple components. For example, it could combine a boot sector virus with a virus that infects one or more files.
- **Macro** A macro virus resides in a document such as Microsoft Word, Microsoft Excel, or Adobe Reader. When a user opens the document, it runs the macro launching the virus. Most current versions of applications that support macros restrict their use to prevent the execution of malicious macros.



EXAM TIP You should be familiar with the different types of viruses, such as armored, polymorphic, and metamorphic, when preparing for the SSCP exam.

Worm

A worm travels over the network, looking for systems to infect. Worms are similar to viruses, but they do not require any type of interaction to execute. Instead, they slither and worm through the network on their own. If a worm infects a network, it has the potential to add a significant amount of traffic on the network and can effectively slow legitimate traffic to a crawl.



EXAM TIP A distinguishing difference between a worm and a virus is that a virus must be executed, while a worm can travel over a network automatically.

For example, a graduate student at Cornell University launched the Morris worm (the first known computer worm) in 1988 with the goal of gauging the size of the Internet. It searched for computers and installed itself on each computer it found and then looked for other computers. It often installed itself multiple times on the same computer, consuming resources, and eventually making the computer unusable. It resulted in estimated damages between \$10 million and \$100 million and prompted the creation of the first Computer Emergency Response Team (CERT) at Carnegie Mellon University.

Worms use protocols to travel over a network. For example, a worm can use TCP/IP to discover computers on the network, access network shares, and then use file transfer protocols such as FTP and TFTP to copy files and infect the target computers. The ability of worms to travel over a network provides additional justification for a defense in depth strategy that includes the following components:

- Host-based firewalls on all systems
- Up-to-date AV software on all computers
- Reduced attack surface by removing unneeded protocols



EXAM TIP Networks include firewalls and other security devices at the boundary to protect the internal network from Internet-based threats. However, malware might still make its way onto the network. Host-based firewalls, up-to-date AV software, and hardening techniques such as removing unneeded protocols help protect systems from these threats.

Trojan Horse

A Trojan horse is an application that looks like it's one thing but is actually something different. It presents itself as a useful or legitimate program, but includes a malicious component.

Its name comes from the Greek story of the nine-year Trojan War. Greek soldiers burned their camp and sailed away after leaving a large wooden horse in front of the gates of Troy. The residents of Troy thought the war was over and thought the horse (a sacred animal to them) was a gift. They ultimately rolled the wooden horse into their otherwise secure city and celebrated their good fortune. However, soldiers hid inside the wooden horse and the rest of the Greek army was poised to return. The soldiers emerged in the middle of the night, killed the guards, and opened the gates to the city, letting the rest of their army in. The Greeks sacked the city, ending the Trojan War. It's a simple message I learned from a mentor who taught me how to play chess many years ago: "beware of gifts."



EXAM TIP A Trojan horse appears to be one thing but is something different. If a user installs new software and then notices suspicious activity, the user very likely installed a Trojan horse.

As an example of a Trojan horse, one attack sent tens of thousands of e-mails to users that looked like it came from CNN and listed several titles of recent news stories. Users who clicked any of the titles opened a website page that popped up a small window. The window indicated that the user's system didn't have the current version of Adobe's Flash Player, so the page

couldn't display the news story. It then prompted the user to install the current version. However, if the user clicked Yes to install the new version of Flash, it actually installed malware instead.

Variants of this attack continue. Users receive an e-mail with one or more links. If the user clicks a link, they are taken to a website and prompted to install an updated version of software, or an application under a false pretense. If the user agrees, the site downloads the malware to the user's system and encourages the user to install it.



TIP Adobe gives simple advice to protect yourself. Only update Adobe products by going directly to the Adobe website and downloading updates directly from there. The same advice is useful for any software or updates. Don't trust a pop-up window, but instead go directly to the source website to download updates for any applications.

RATs

A remote access Trojan or remote access tool (RAT) provides an attacker with control over a target computer via the Internet. Many Trojans install RATs, and then attackers can use the RAT to do anything on the target computer. This includes using the RAT to install malware, install keyloggers, and join the compromised computer to a botnet.

Many advanced persistent threats (APTs) use RATs to gain a foothold into a network. After infecting one system, they use various techniques to gather data about other systems on the network. [Chapter 5](#) discusses APTs in more depth.

Scareware

Scareware (sometimes called rogue-ware) is malware that describes itself as antivirus software. A user may visit a website and then see a pop-up similar to that shown in [Figure 6-1](#), which indicates that a virus has infected the user's system. However, every user who visits the website will see the exact same message. The website hasn't truly detected a virus. Instead, the pop-up attempts to trick the user into installing malicious software on their system.



Figure 6-1 Scareware pop-up trying to trick a user into installing malware



NOTE Although I created the pop-up window shown in [Figure 6-1](#), it is similar to some pop-ups that users see when visiting some malicious websites. There isn't a Trojan virus named fakemeout.takemymoney. However, scareware does try to fake you out with the intention of taking your money.

If the user clicks the Fix Now button, it prompts the user to download and install the bogus antivirus application. After installing and running it, the bogus software indicates it is checking the user's system, but it's all a show. Unlike actual antivirus software that truly scans a system to look for malware, scareware does not look for malware. However, it always reports that it has discovered multiple instances of malware on the user's system. It then informs the user that the free trial edition only scans their system. The user must pay to purchase the full version to remove the malware.

Although the extortionary nature of scareware is malicious enough by itself, most scareware applications do more. For example, the downloaded application could join the system to a botnet controlled by a criminal. It could download and install a remote access tool that allows a criminal to take control of the user's computer, or install a keylogger to capture all of the user's keystrokes and send the data off to an attacker to use later. The possibilities are limited only by the imagination and technical abilities of the scareware authors. History shows us that they have a lot of imagination when it comes to stealing money from victims, and their technical abilities continue to increase.

What should a user do after seeing a pop-up window like this? Close the pop-up window and leave the site. Users who don't have antivirus software installed should go directly to a reputable site to get it, and not try to install it from one of these warnings. Many brands of antivirus software are available, and some reputable companies offer free versions.

Ransomware

Ransomware is a type of Trojan that takes control of the user's computer or data and then demands a ransom from the user to get control back. There are many variants, but there are two primary themes. One type locks you out of your computer, preventing you from logging on. Another type encrypts your data files, allowing you to log on but preventing you from accessing any of your data.

The "Mandiant U.S.A. Cyber Security Ransomware" is an example of ransomware that locks you out of your computer. When users turn on an infected system, it displays a bogus but official-looking screen that indicates it's a joint effort between Mandiant Cyber Security, the U.S. FBI, the U.S. Department of Defense, and the U.S. Cyber Crime Center. Variants (known as the Police Virus, Trojan Reveton, and MoneyPak Virus) replace the names of U.S. agencies with those of agencies in other countries, such as the International Criminal Police Organization (Interpol). This type of malware isn't a joint effort with any official agencies. Instead, it's an attempt by criminals to extort money from users.



NOTE Information security company Mandiant has done a lot of research on attackers. It released a report in 2013 documenting activities of advanced persistent threats (APTs) operating in China. Mandiant was later purchased by FireEye, another security company.

The ransomware message accuses the user of participating in illegal activity such as distribution of pornography, copyrighted files, or malware. It then demands that the user pay a fine of \$300 (or some other amount) within 48 hours or they will face legal prosecution. The software directs users to purchase a MoneyPak or MoneyGram voucher at a store such as Kmart or Walgreens. They enter the voucher code into the ransomware screen to pay the attackers, which unlocks their computer.

CryptoLocker is an example of ransomware that encrypts user files. When it infects the user's system, it locates and encrypts data files on the user's system, attached USB drives, external hard drives, shared network drives, and sometimes cloud storage drives. It then demands payment of \$300 (or more) to get the decryption key. The message indicates that if the user doesn't pay within 72 hours, the attackers will delete the decryption key, causing the user to lose the files forever. Attackers sometimes offer to retrieve the decryption key after the time limit has expired, but at a much steeper price as high as \$2,300. Victims paid approximately \$30 million between September and December 2013.

Police forces disrupted the network running the original CryptoLocker ransomware in June 2014. However, CryptoLocker and many variants began appearing soon afterward. TorrentLocker is similar and often uses the CryptoLocker name, but the underlying code is different. The FBI is offering a \$3 million reward for information leading to the arrest or capture of a Russian man (Evgeniy Mikhailovich Bogachev) who they believe is the ringleader of the gang that ran the original CryptoLocker network.

Law enforcement agencies discourage victims from paying ransom. First, there's no guarantee that the criminals will release the user's computer or data, and there are many examples where they didn't, or asked for more money. Second, as long as criminals receive money from these attacks, there's no reason for them to stop.

Attackers released the WannaCry ransomware attack in 2017, infecting more than 200,000 computers in 150 countries in just four days, and more than 300,000 systems in total. It demanded a payment of \$300 within three days or \$600 within seven days. Government agencies in the United States, United Kingdom, and Australia concluded that North Korea was behind the attack. If true, it was likely an APT working in and for North Korea.

Notice that each of these ransomware attacks includes a sense of urgency. The Mandiant virus states that users will be prosecuted if they don't pay the fine within 48 hours. CryptoLocker states that attackers will delete the decryption key if users don't pay the ransom within 72 hours. WannaCry doubled the ransom if the user didn't pay within three days. Ransomware typically includes a timer that counts down as a constant reminder to the user to take action.

Installing Ransomware

Criminals have used a variety of methods to trick users into installing ransomware. Many methods include tricking users with phishing e-mails. Many phishing e-mails have mimicked UPS and FedEx tracking notices indicating the delivery service attempted to deliver a product but failed. These notices encourage the user to either click a link to a (malicious) website or open an (infected) attachment. Visiting the website or opening the attachment launches the attempt to install the malware.

As another example, attackers have sent users phishing e-mails that looked like they were from an official state agency. The e-mails said that a traffic camera caught the recipients speeding and that they now had to pay a fine to avoid the suspension of their driver's license and/or vehicle registration. The e-mails included two buttons, one for an invoice and one to view the camera images. Clicking either button took the recipients to a web page that prompted them to enter a CAPTCHA code. When they did, it downloaded a file. If they opened the file, it installed the ransom-ware and encrypted all the data on their system.

Keylogger

A keylogger (also known as a keystroke logger) is a tool that can capture all the keystrokes that users enter on a computer. Keyloggers can be hardware based or software based. A hardware-based keylogger plugs into the computer between the keyboard and the computer and captures

the data. For example, if the keyboard plugs into a USB slot, you would plug a USB-keylogger into the USB slot, and then plug the keyboard into the USB-keylogger. It could transmit the information wirelessly or store the data on the hardware.

Software-based keyloggers are more common. They simply run as an additional piece of software when the computer starts. For example, some parents use keyloggers to monitor children's computer activity. Keyloggers can monitor online chats, instant message content, e-mails, and data entered into Internet website forms. Many keyloggers also include the ability to capture screenshots. These same tools are available for employers to monitor employee activities.

Of course, criminals use the same type of keylogger software to capture information entered by users. Users don't knowingly install the keylogger. Instead, another infection such as a Trojan horse installs the keylogger.

Logic Bomb

A logic bomb is malicious code that executes in response to an event such as a point in time or in response to specific actions. For example, [Chapter 5](#) mentioned the attempted attack on Fannie Mae by a terminated administrator. After his employer terminated him from the job, he added a script that was set to run on January 31 that would have deleted passwords, data, and backups for about 4,000 servers. This is a logic bomb because it was set to run on a specific date.

Some keyloggers have a logic bomb component. They send data to an attacker after a user accesses specific websites, such as banking sites. For example, when the keylogger detects that the user accessed mybank.com and entered a username and password, the keylogger e-mails the URL and credentials to an attacker.

Rootkits

A rootkit is a set of programs that can run on a system largely undetected. It has root-level access to the system, similar to how a root-level administrator has full and complete control over a system. Attackers can often control the computer remotely after installing a rootkit, and the user typically doesn't notice any symptoms.

Rootkits have low-level control on systems, which allows them to identify when antivirus software is running. The rootkit can detect a system scan and hide itself from the scan or restrict the data returned to the scan. This has made rootkit detection difficult for traditional AV software. However, it's worth noting that attackers typically install rootkits via a virus or a Trojan. In other words, if you can detect and prevent attackers from installing viruses and Trojans on your system, you have a good chance at preventing them from installing rootkits.

Several companies have released rootkit detectors. For example, Sophos has released and updated a free dedicated tool designed to detect and remove rootkits from systems. If interested, you can read about it here: www.sophos.com/en-us/products/free-tools/sophos-anti-rootkit.aspx. Additionally, most popular AV programs advertise their ability to detect and remove rootkits.

This typically requires rebooting the system from a CD or bootable USB with antivirus software to detect the rootkit and clean the system.



TIP Microsoft 64-bit systems have added more protections against rootkits. They use the Kernel Patch Protection (also known as Kernel PatchGuard), which prevents rootkits from modifying the underlying operating system components (the kernel). Kernel PatchGuard requires digitally signed drivers. This ensures the integrity of the files. If malware modifies the files, the digital signature will detect the modification.

Mobile Code

Mobile code is any code that can transfer itself between systems without the user taking explicit action to install it. Attackers use malicious mobile code to install malware from remote servers via the Internet. Malicious mobile code is typically transmitted via interactive web applications. Some examples of mobile code include the following:

- **Scripts such as JavaScript or VBScript** Attackers embed the script into a web page, or within any file, that supports the script. When users click a link to retrieve a web page, the web server sends the page along with the embedded script. As the web browser displays the web page, it also runs the script. Some Hypertext Markup Language (HTML)-based e-mail applications support the use of scripts, so opening an e-mail can launch the script. Similarly, files such as Portable Document Format (PDF) documents and Microsoft Office documents support scripts. Attackers embed the code into these documents, which executes when the user opens the document. While keeping systems and applications patched with security updates protects against most of these threats, many users don't keep their systems patched.
- **Java applets or ActiveX controls** Java applets and ActiveX controls are mini programs that developers embed into web pages. Similar to scripts, they run when the web browser displays the web page. A primary method of protection is to block unsigned applets and controls. Using code signing (as discussed in [Chapter 5](#)) is one method that developers use to show their code is valid and provides assurances to users that the code is valid.
- **Documents that include macros** Many document types (such as Microsoft Office documents) support macros. While macros can be very helpful, an attacker can write malicious mobile code into a macro. When a user opens the file, it runs the macro. To reduce this risk, it's common to disable macros in applications that support them.

Backdoors and Trapdoors

A trapdoor or backdoor is code that is embedded in an application and provides access to the application, the application's code, or data via a covert method. Some backdoors provide access to the operating system.

Application developers often add backdoor code to an application for debugging purposes. For example, by entering specific keystrokes, they can display the value of variables within a running application. Although this is useful during the debugging process, a backdoor can be dangerous if developers leave it in the released version of the application. Instead, as a best practice, most developers remove all backdoor code prior to releasing the application.

Attackers often attempt to install a backdoor after infecting a system with a virus or Trojan. When successful, attackers can use backdoors to read data, install additional malware, or even take control of the user's computer from a remote location. Within the context of malicious code, both backdoors and trapdoors mean the same thing, although the term *backdoor* is more common.

Some attackers have successfully added backdoor code to released applications. For example, the Mydoom worm installed a backdoor that affected versions of Microsoft Windows and allowed attackers to control the infected computers remotely over the Internet. As with most other malware, after security experts discovered the vulnerability, the vendor (Microsoft in this example) wrote and released security updates to protect systems from the threat.



NOTE The Mydoom worm (released in 2004) was the fastest-spreading e-mail worm up to that date. It infected computers and then used them to send spam. The worm created a backdoor on port 3127, which allowed remote control of the infected PC. Variants appeared in 2005 and then again in 2009.

Another example is the Win32/Zbot (also called GameOver Zeus or simply Zeus) family of password-stealing Trojan horses. This is also a good example of a multipartite virus with backdoor functionality. Win32/Zbot tricked users into installing the Trojan onto their systems, and then joined their computers to a botnet. Interestingly, infected systems don't join a single huge botnet, but instead join one of many independently operated botnets. The attacker controlling the command and control center for each of these botnets then controls these infected systems through the botnet's backdoor.

Win32/Zbot successfully used a variety of techniques to steal financial-related information such as usernames and passwords. For example, it interfaced directly with web browsers to monitor traffic and information. The Trojan deleted cookies for websites, forcing users to reenter their credentials. It injected additional HTML code into banking website pages, requesting additional information from users such as Social Security numbers or birthdays. Win32/Zbot then encrypted this information and sent it to attackers controlling the botnet. Criminals logged on to the victim accounts with the stolen credentials and transferred funds to other bank accounts controlled by the criminals.



NOTE The FBI believes that the ringleader of the network that released the original CryptoLocker ransomware also created and distributed the Zeus banking Trojan. Security professionals estimate that attackers have successfully used Zeus to steal hundreds of millions of dollars. There are also indications that Zeus malware installed CryptoLocker onto infected systems.

Spyware

Spyware is software that can install itself on a user's system without the user's knowledge or consent. Although there was a time when spyware was relatively benign, it is much more dangerous today. For example, there was a time when spyware attempted simply to understand a user's surfing or buying habits so that it could send targeted advertisements to the user. However, spyware today is much more malicious and attempts to learn details to steal financial data, logon information, or a person's identity.

Most security experts consider spyware a form of malware because it can cause so much damage to an individual. Many spyware applications can identify websites that a user visits, usernames and passwords that a user enters, user's keystrokes even when not visiting a website, and even search a user's computer for information such as files including usernames and passwords. Spyware can also change settings on a user's system, such as settings to enable or disable a firewall, settings to enable or disable automatic updates for the operating system, or even settings for individual applications.

Due to the malicious nature of spyware today, most AV software includes the ability to search for and detect it. Years ago, many users used one application for spyware and another application for viruses and other malware. However, if a user gets an AV application that includes a spyware component, it isn't necessary to install a separate spyware application.



NOTE An advanced type of spyware is canvas fingerprinting, which uses a variety of techniques to track users through an entire surfing session. A large-scale study titled "The Web never forgets: Persistent tracking mechanisms in the wild" documents three advanced web tracking mechanisms. You can read about it here:

<https://securehomes.esat.kuleuven.be/~gacar/persistent/index.html>.

Malware Hoaxes

In its simplest terms, a malware hoax is a false message about a malware risk that doesn't exist. Attackers and uneducated users typically spread hoaxes through e-mail while urging other users to forward it to everyone they know. Many times these e-mails encourage users to take action that can harm their system.

Hoaxes commonly include a sense of urgency and often overuse capital letters and exclamation marks, with subject lines featuring something like "URGENT!!!," "WARNING!!!," or "VIRUS ALERT!!!" In contrast, legitimate warnings (such as those from a cyber security alert service or from an organization's administrator) simply provide the detailed information without stimulating emotion or fear.

In general, hoaxes just waste people's time. When users forward it to everyone they know, they perpetuate the hoax, along with consuming more bandwidth and wasting other people's time. Many organizations provide training to users so that they don't forward e-mail warnings to everyone. Instead, users are encouraged to forward these types of e-mails to system administrators if they think it represents a real threat. Although this adds to the administrator's workload, it's better than spamming the rest of the organization.

Analyzing the Stages of an Attack

While the previous descriptions of malicious code describe individual components of malware, they really don't paint the full picture of malware threats. Symantec released a white paper titled "Regin: Top-Tier Espionage Tool Enables Stealthy Surveillance" that describes the complexity behind some attacks that start with a serious malware threat. While the white paper (and the summary provided in this section) describes Regin, other attacks use similar complex methods to capture data after infecting a system. The ultimate goal of Regin is to collect data, and it does so through several different methods, including using Trojan RAT features, capturing screenshots of user computers, taking control of the mouse, stealing passwords, monitoring network traffic, retrieving deleted files, and collecting e-mail.



TIP This section describes the known stages of Regin. However, it's important to remember that this is just one example. Many families of similar malware include the same goals with the same level of complexity.

Regin 1.0 has been around since at least 2008, and Regin 2.0 has been used since at least 2013. This speaks to the technical sophistication of the attackers. Developers aren't teenage script kiddies doing random attacks. Rather, the developers appear to be a cohesive team with specific goals of monitoring targets and stealing data. There are six distinct stages to Regin, as shown in [Figure 6-2](#) and described in the following list:

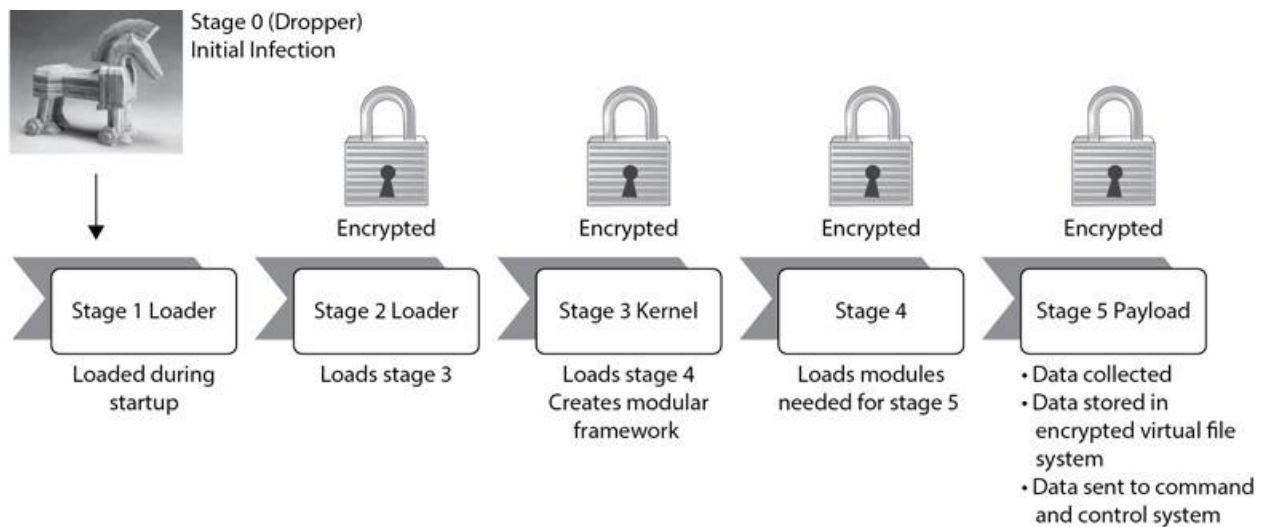


Figure 6-2 Regin stages

- **Stage 0** A dropper is a type of Trojan horse that attempts to drop malware onto a computer. Within Regin, stage 0 is the dropper stage and represents the initial infection. It could be from the user opening an infected file in an e-mail, clicking a malicious link in an e-mail, or any other way users inadvertently install Trojans. Stage 0 installs and executes stage 1 files. In many cases, stage 0 code is only stored in memory and disappears after stage 1 starts.
- **Stage 1** This stage loads operating system kernel drivers when the computer starts. These drivers decrypt, load, and execute stage 2. Two known names for these files are subclass.sys and adpu160.sys for version 1 and version 2, respectively. Stage 1 is the only stage that includes plainly visible code on the attacked computer. Regin encrypts stages 2, 3, 4, and 5 and stores them in files, or within nontraditional areas such as the registry, or within raw sectors of the disk drive. As mentioned previously, encryption is a method of armoring the virus to prevent AV researchers from decompiling and analyzing the code.
- **Stage 2** This stage loads another operating system kernel that decrypts, installs, and runs stage 3. Regin stores encrypted data for this stage within an NTFS extended attribute, or in some cases within one or more registry subkeys. This stage often hides running instances of stage 1, removing any plainly visible elements of Regin.



NOTE NTFS extended attributes are name–value pairs used for a variety of purposes. For example, Windows systems use them to mark files downloaded from the Internet, and provide a warning to users when they open these downloaded files. Regin uses the %Windir%, %Windir%\fonts, and %Windir%\cursors extended attributes to store encrypted blobs. You might know these as environment variables, which they are. However, Windows systems store environment variables as extended attributes.

- **Stage 3** This stage decrypts, installs, and runs stage 4 files and provides a modular framework for stages 3, 4, 5, and 6, consisting of hundreds of programmatic methods. These modules provide a wide range of functions such as compression and decompression routines, encryption and decryption routines, an encrypted virtual file system, and networking routines to send data to the attacker. Stage 3 is about six times larger than stage 2.
- **Stage 4** Stage 4 loads the relevant modules needed for stage 5 depending on the targeted payloads. Because Regin includes so many different modules, attackers can modify it to collect custom payloads. For example, one implementation might collect e-mail, another implementation might install a RAT, and so on. Note that this isn't a different version of Regin, but instead, the attacker just enables different features of Regin by implementing different modules.
- **Stage 5** Stage 5 collects data from a variety of sources and stores that data within the encrypted virtual file system. It can collect network data, passwords, computer files, deleted files, and screenshots. Stage 5 sends data to the attackers via one of several standard protocols, using special flags for identification.

While these six stages occur automatically after the initial infection, Regin also supports command and control operations similar to a traditional botnet. Attackers can contact the computer periodically and instruct it to send data. This often occurs using two stages. For example, the attacker might send an Internet Control Message Protocol (ICMP) datagram with instructions to send collected data to a specific server. The infected computer then responds with the payload using User Datagram Protocol (UDP) datagrams. Regin encrypts all data before sending it to the command and control computer.

Understanding Malware Delivery Methods

Attackers deliver malware through multiple methods. While most malware is delivered via the Internet, it can also be delivered from hardware devices such as USB drives. When end users understand how malware spreads, they are less likely to be tricked into installing malware on their systems.

Delivering Malware via Drive-by Downloads

In a *drive-by download*, web servers include malicious code that attempts to download and install itself on user computers when the users visit. These drive-by downloads can include many different types of malware, such as installing a RAT or joining a computer to a botnet.

Criminals rarely host these computers themselves because that would make it easier for law enforcement to find them. Instead, they look for vulnerable computers that they can modify to include the drive-by downloads. As an example, attackers might look for an unpatched server with a known vulnerability. If the vulnerability allows the attackers remote access to the server, they can use this access to install malicious code.



TIP Attackers modify vulnerable servers to install other malware too. For example, scareware uses a pop-up to indicate the site detected malware on the user’s computer and then attempts to trick users into installing fake antivirus software. It’s easy to think that the criminals own the site, but more often attackers just install their own pop-up on an innocent business owner’s site.

Here are the typical steps involved in a drive-by download, depicted in [Figure 6-3](#):

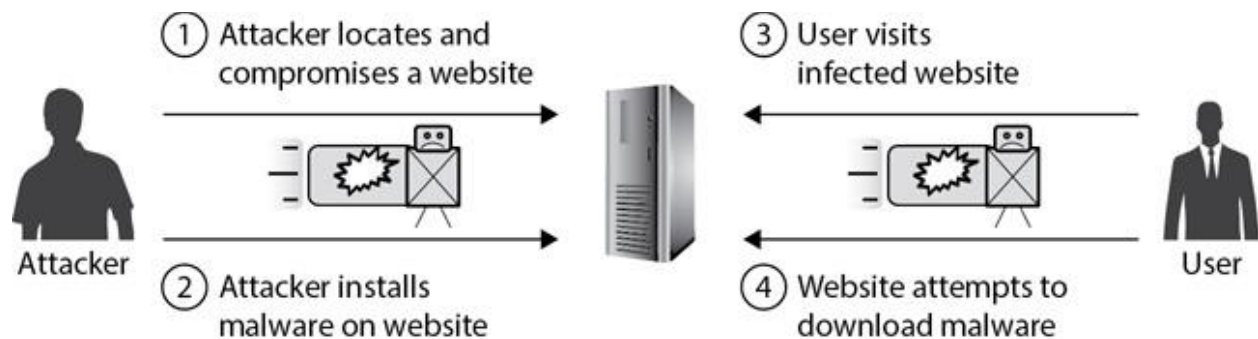


Figure 6-3 Drive-by download installing malware

1. Attackers compromise a website to gain control of it.
2. Attackers install a Trojan or other malware embedded in the website’s code.
3. Attackers attempt to trick users into visiting the site. Sometimes, they simply send the link to thousands of users via e-mail, hoping that some of them click the link. Other times the link is from another website or a malicious banner ad.
4. When users visit, the website attempts to download the Trojan or malware onto the user’s system.

Users can prevent successful drive-by downloads by keeping their systems up to date with operating system and application updates and with up-to-date antivirus software. This will prevent most infections, but it won’t prevent zero day exploits, so users still need to be vigilant. A *zero day exploit* is an attack that takes advantage of unpublished vulnerabilities that vendors either haven’t heard about or haven’t published updates for yet.

Many drive-by downloads prompt the user to complete the installation, and they won’t complete without the user’s permission. As an example, the “Trojan Horse” section earlier in this chapter mentioned how one attack prompted users to install an update to the Adobe Flash Player. If they agreed, it installed malware. Users should be encouraged to only install updates to applications from the site publishing the application. In other words, when updating Adobe Flash Player, users should get the update from a verified Adobe site.

Delivering Malware via Malvertising

Attackers sometimes deliver malware via malicious banner ads, commonly called malvertising. These look like regular ads, but they contain malicious code. Many of them include a link taking users to a server hosting a drive-by download. Others are Flash applets with malicious code embedded in them.

Malvertising has appeared on mainstream websites such as those belonging to The New York Times and Yahoo!. As an example, ads hosted on Yahoo! in 2013 and 2014 took users to sites hosting scareware. They included fake pop-ups as described in the “Scareware” section earlier in this chapter. In mid-2014, ads on Yahoo! took users to Eastern European sites hosting CryptoWall ransomware, a clone of CryptoLocker.

It’s worth noting that legitimate sites are not knowingly hosting these malvertisements. Instead, they sell ad space to advertising companies and these advertising companies unknowingly send the malicious ads to the legitimate sites. In some cases, criminals trick legitimate advertising companies into hosting the ads. For example, social engineers tricked people at the Gawker Media network to sell fake Suzuki advertisements. In other cases, attackers compromise websites and post their own malicious ads on the websites.

Attackers often use exploit kits, such as Magnitude EK, to create malicious advertisements. In 2017, attackers began using Magnitude EK to inject scripts into advertisements on vulnerable servers. When a user clicks a malicious advertisement, the script redirects them to a malicious server hosting malware such as ransomware. In many cases, users are quickly redirected to the site a user would expect after clicking the ad. This happens while the server is downloading the malware.

Delivering Malware via E-mail

One of the most popular methods of delivering malware is via e-mail. [Chapter 5](#) discusses spam, botnets, and various types of phishing attacks. As a reminder, criminals control hundreds of thousands of compromised systems in botnets and they rent these botnets out to other criminals who use them to send phishing e-mails.

While some phishing e-mails attempt to trick users into giving up personal information such as their credentials for a banking site, other e-mails try to install malware on the user system.

One method includes infected attachments. If users open the attachment, it can infect the user’s system. Attachments arrive as ZIP files, PDF documents, web pages (as HTML files), apparent graphics such as JPG files, and more. One way they disguise these documents is to simulate the extension of a safe file. For example, a JPG graphic file has a .jpg extension, such as picture.jpg. An attacker can rename an executable file to picture.jpg.exe. Because the operating system normally hides the extension, the file looks like picture.jpg, indicating it’s a relatively safe JPG file.

It's worth noting that attackers often spoof or modify the From e-mail address to make it look like it's from someone the recipient knows. For example, attackers monitor social media sites and identify your friends. They then send an e-mail to you but change the From field so that it looks like it came from your friend. This tricks some users into trusting the e-mail and makes them more willing to download an attachment or click a link.

Delivering Malware via USB Drives

As mentioned earlier in this chapter, some viruses replicate via USB drives. An infected USB drive will try to infect a computer system when a user inserts it. An infected system waits for a user to insert a USB drive and attempts to infect the drive.

When an organization allows users to transport data back and forth from home and work using USB drives, it becomes very easy for a user to transport viruses from home and work. Although an organization can implement multiple methods to protect its environment from malware, the organization usually has no control over a user's home network. If users are authorized to use USB drives, it's imperative that each system has AV software installed and kept up to date. Some organizations use additional methods to prevent systems from reading data from a USB drive, writing to a USB drive (to prevent data theft), or both.



EXAM TIP If users are authorized to transport data with USB drives, each system should be protected with AV software. This is the best method to protect a system without affecting the user. It's also possible to prevent copying of data with USB drives to provide stronger protection.

Implementing Malicious Code Countermeasures

Because malware is so prevalent, it's important to protect systems from malware. The primary method of protection is antivirus software and an aggressive defense in depth strategy. As an overview, the following methods are effective countermeasures against malware:

- Install antivirus software on all systems.
- Install antivirus and antispam software on e-mail servers.
- Install antivirus and content-filtering software on firewalls.
- Keep all antivirus and antispam software up to date.
- Perform regular antivirus and vulnerability scans.
- Keep all systems up to date.

- Educate users.

Antivirus Software

The primary method used to detect and prevent infections is the use of AV software. In addition to being able to detect and prevent infections, AV software can remove the malware, restore infected files, or quarantine infected files if the AV software can't remove the malware. The goal is to ensure that the malware cannot run and cause any more problems on a system. Researchers can examine and analyze quarantined files. However, the computer will not run quarantined files.

[Figure 6-4](#) shows an example of antivirus software. It includes the ability to perform on-demand scanning (shown as Scan Your PC in [Figure 6-4](#)), real-time scanning, and scheduled scanning:

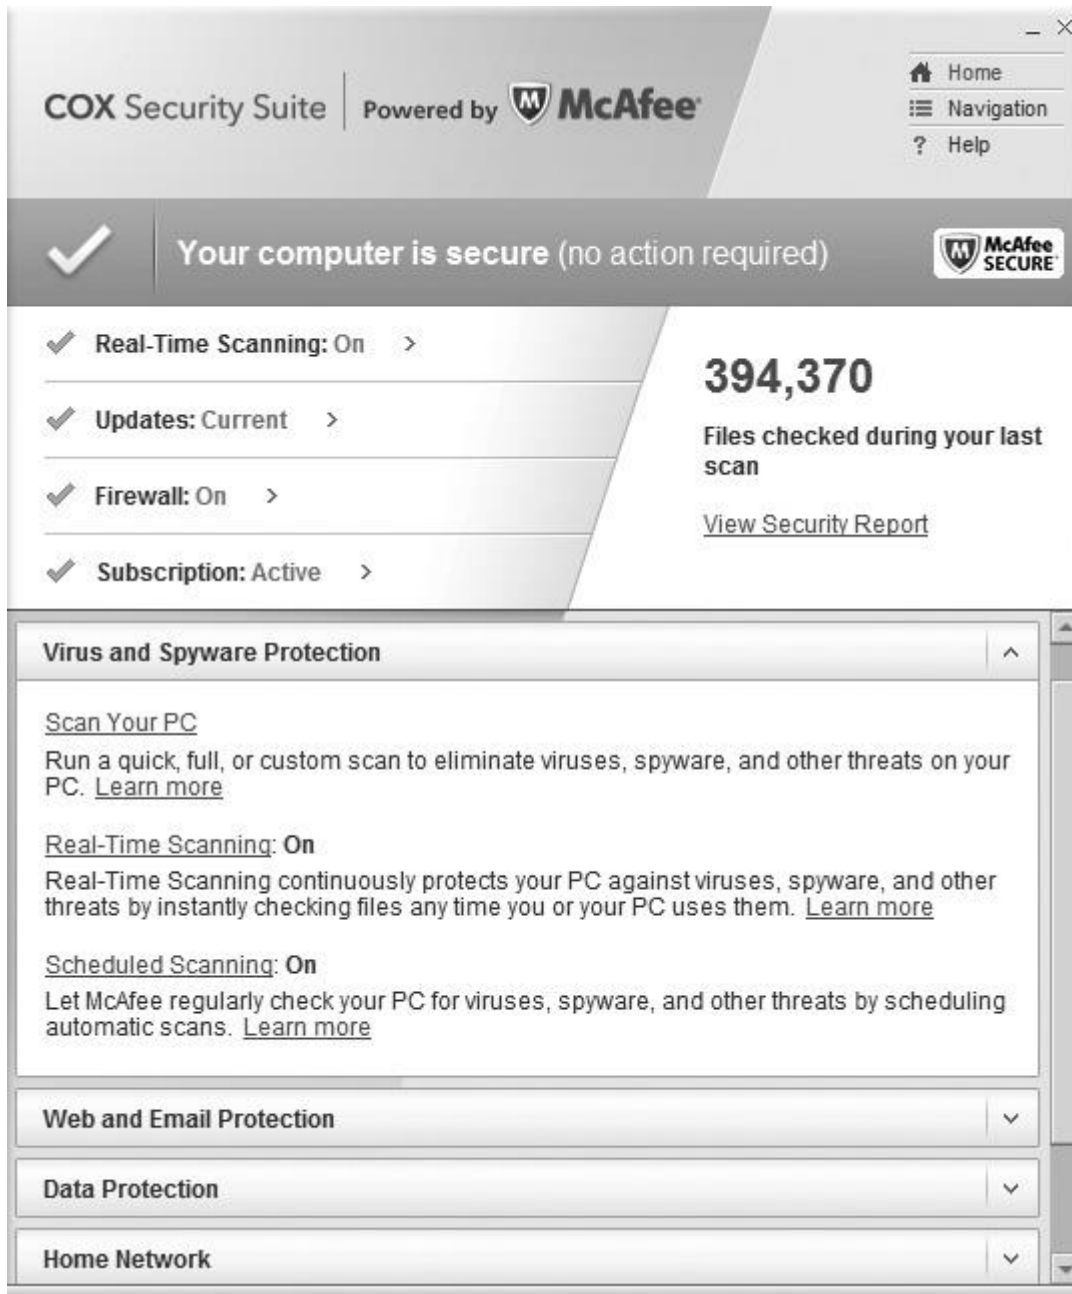


Figure 6-4 Antivirus software



NOTE Many Internet service providers (ISPs) provide free copies of antivirus software for their customers. For example, at the top of [Figure 6-4](#), you can see the COX Security Suite branding. Cox Communications provides this software to its Internet customers.

- **Real-time scanning** This provides protection as the user opens and uses files. For example, if a user opens an e-mail attachment, the AV software checks it to ensure that it doesn't have any malicious content. Real-time scanning checks any file that is opened or launched by either a user or the operating system.

- **On-demand scanning** If you suspect a computer is infected, you can perform an on-demand scan to check it. Most AV software allows you to run a scan to check memory, specific files, folders, or drives, or complete a full scan of the entire system. [Figure 6-5](#) shows an on-demand scan in progress on a system.

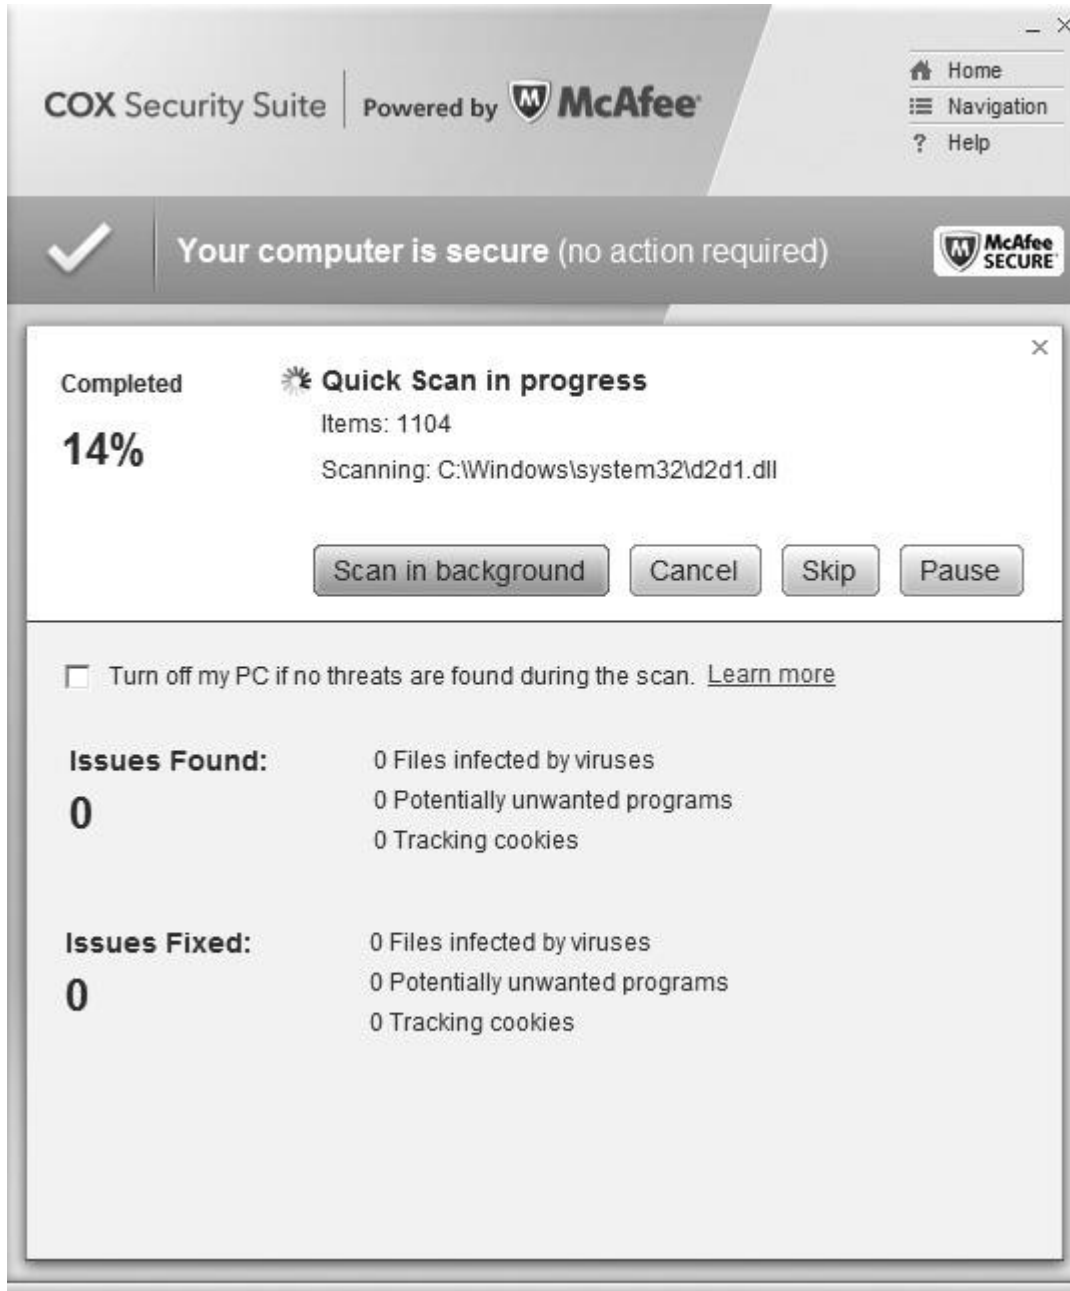


Figure 6-5 Performing an on-demand scan

- **Scheduled scanning** A scheduled scan occurs on a regular basis, such as once a week. Once scheduled, the AV software launches the scan based on the schedule.

Antivirus software uses a software engine to detect viruses using virus signatures. Additionally, some antivirus software uses behavior-based (sometimes called heuristic) scanning. The following sections describe both methods.

Signature-based Detection

Viruses have specific characteristics that AV software uses to identify them. The signature can be a unique characteristic such as a specific byte pattern within the virus. This is similar to how a person's fingerprint can uniquely identify an individual. Viruses in files, memory, and a drive's boot sector all have specific signatures that AV software uses to detect them.

When AV researchers discover malicious software, they identify its signature. AV companies add this signature definition to the AV signature database, which the AV software uses to detect the malware. When the AV software finds a file that matches a signature in the database, it alerts the user and often prompts the user to remove, repair, or quarantine the file.

You may remember from earlier in the chapter that polymorphic and metamorphic viruses change. A primary goal of attackers using these techniques is to prevent detection of the virus through signatures. Additionally, attackers have developed application tools that allow them to create new iterations of viruses rather easily. For example, the attacker can write the replication and activation component to create the core virus. The attacker can then run the virus through another application to create hundreds of different versions of the same virus. Some applications can generate as many as 500 different variants of a virus in a single month. AV software might have a signature of the original virus, but this signature is unlikely to detect all 500 variants of the virus.

Many AV developers have moved to fuzzy hashing techniques to identify malware variants. Fuzzy hashing techniques compare two different items and determine a level of similarity between the two. While fuzzy hashing might not detect all variants of a known virus, it does detect many of them.

Of course, as viruses morph to a state that prevents detection, AV researchers create signatures to detect the new variants. If users ensure their AV signature files are updated regularly, the AV software will discover these new variants. However, if users don't update the virus signatures regularly, the new variants may run undetected.

Staying Aware of Malware Threats

Malware is very dynamic. As an example, several international law enforcement agencies disrupted the crime ring that was running the GameOver Zeus botnet and the CryptoLocker ransomware in mid-2014. Many experts said this was the end of this criminal ring and the end of

CryptoLocker. However, CryptoLocker and some variants reappeared a few months later, and are now as strong as ever. One way to keep up with these changes is to subscribe to various newsletters or bulletins. Here are a few that you might find useful:

- **Virus Bulletin** This organization focuses on preventing, detecting, and removing computer threats, including malware and spam. They previously published a monthly bulletin, but now publish them at least monthly, and more often when appropriate. You can view recent or archived articles on their website (<https://www.virusbulletin.com/>) and it's free.
- **SANS Internet Storm Center** This site is part of the SANS Technology Institute and it provides information on current threats. You can sign up to receive daily podcasts at <https://isc.sans.edu/dailypodcast.xml> and view past diaries at <https://isc.sans.edu/diaryarchive.html>.
- **Cyber Awareness Alerts** These alerts provide timely information about current security issues, vulnerabilities, and exploits. They come out at least once a month and provide technical information for system administrators and experienced users. You can view archives here: <https://www.us-cert.gov/ncas/alerts>.
- **Cyber Awareness Bulletins** These provide weekly summaries of new vulnerabilities. They come out once a week and include information about security issues and vulnerabilities from the previous week. You can view archives here: <https://www.us-cert.gov/ncas/bulletins>.
- **Cyber Awareness Tips** These tips provide information and advice on common security topics. They are written for regular users and come out every two weeks. You can view archives here: <https://www.us-cert.gov/ncas/tips/>.
- **Current Activity Updates** These provide timely updates about high-impact security activity. This list can generate alerts from several times a day to several times a week. You can view archives here: <https://www.us-cert.gov/ncas/current-activity/>.

You can sign up for any (or all) of the Cyber Awareness notifications here: <https://public.govdelivery.com/accounts/USDHSUSCERT/subscriber/new>.

Behavior-based Detection

Attackers release new viruses into the Internet regularly, and older viruses are regularly mutating into variants. Because of this, it's entirely possible that at any time, a system will come across a previously unknown virus. This means that the AV software doesn't have a signature for the new virus. However, behavior-based detection can often detect these previously unknown viruses.

Behavior-based detection (sometimes called heuristic-based detection) attempts to detect a virus based on its behavior. In other words, if a file is behaving in a way that is not typical for a file, the heuristic engine can detect this behavior, identify the file as malware, and block the malware from infecting the system. For example, legitimate applications rarely modify, copy, or delete

system files. If an application does attempt to take one of these actions, the heuristic engine can detect and block it.

Many behavior-based scanners use a sandboxing technique, which relies on an isolated or confined area of memory on the system. The AV software ensures that new or modified applications run within an isolated sandbox for analysis before allowing them to run regularly. The idea is that the application can play within its own sandbox, but it can't interfere with applications or processes in other sandboxes or other areas of memory. The scanner observes the behavior of the application as it runs within the sandbox. If it doesn't display any suspicious activity, the AV software allows it to run.



EXAM TIP Virus signatures detect known viruses. It's important to update AV software signatures regularly to ensure the AV software can detect newly released malware. Behavior-based detection attempts to detect previously unknown viruses.

Keeping AV Signatures Up to Date

An important element of using AV software is to keep the signature definitions up to date. Because viruses can morph so quickly, most AV software includes the ability to check automatically for new signatures, and they do so once a day and sometimes as often as once every hour. Users only need to ensure that the AV software is set to download and install the updates automatically. If the AV software is not up to date with current signature definitions, systems quickly become more vulnerable to new attacks.



TIP AV software downloads signature files incrementally so that they include only differences. In other words, if you have all the current signatures except for two, the download only includes the two new signatures, not the entire signature definition file.

Users often don't understand the risks and the importance of updating signature definitions, so many organizations typically take control of this step. For example, some AV vendors sell server products where a central server can deploy both AV software and updates to all the clients within an organization. Administrators ensure the central server downloads the updates from the vendor site and then deploys the updates to clients within the organization. Many times these updates include more than new signature definitions. They can also upgrade and improve the antimalware application.

Spam Filters

Because attackers often deliver malware and malicious links via spam, a spam filter can provide an important layer of protection. A spam filter attempts to detect unsolicited e-mail and block it. Many e-mail servers include antispam software to filter spam. Additionally, many e-mail applications include antispam software to filter spam.

A danger of a spam filter is blocking legitimate e-mail. For example, a company that sells products and communicates to customers via e-mail doesn't want to block legitimate queries from customers. Because of this, most spam filters have low thresholds for spam. Given the choice of accepting some spam but never blocking legitimate e-mail, or blocking more spam but occasionally blocking legitimate e-mail, most organizations choose the former. They'll allow some spam to get through but do not want to block any legitimate e-mail.

Three technologies that have helped reduce spam are Sender Policy Framework (SPF), Domain Keys Identified Mail (DKIM), and Domain-based Message Authentication, Reporting, and Conformance (DMARC).

Sender Policy Framework

Sender Policy Framework (SPF) records provide a method to reduce spam by identifying spoofed e-mail. SPF records identify e-mail servers authorized to send out e-mail for a domain. For example, McGraw-Hill Education can publish an SPF record to DNS servers on the Internet that identifies the mail.mheducation.com server as the only authorized server to send out e-mail for the mheducation.com domain. Spammers sometimes spoof the source address, but the e-mail header still shows the server that sent the e-mail.

E-mail servers that receive e-mail can use SPF records to validate e-mail. When an e-mail server receives e-mail that appears to come from McGraw-Hill Education, it can check what server actually sent it and compare it to the public SPF record. If a different server sent the e-mail (not one identified in the SPF record), the receiving server marks the e-mail as spoofed. The server can then block all spoofed e-mail from this server to reduce spam levels. RFC 7208 formally defines SPF and it is currently listed as a proposed standard.

Domain Keys Identified Mail (DKIM)

Domain Keys Identified Mail (DKIM) uses digital signatures to authenticate e-mails. [Chapter 14](#) covers digital signatures in more depth, but in short, DKIM uses a certificate to provide authentication, integrity, and nonrepudiation. The sending e-mail server attaches a digital signature to outgoing e-mail. The receiving e-mail server uses the digital signature to verify the e-mail has been sent by the sending e-mail server.

Domain-based Message Authentication, Reporting, and Conformance

Domain-based Message Authentication, Reporting, and Conformance (DMARC) is an extension of both SPF and DKIM. It allows owners of a domain to publish a policy that identifies if they

are using SPF, DKIM, or both. Administrators of receiving e-mail servers can use this policy to validate incoming e-mail. DMARC is documented in RFC 7489.

Content-filtering Appliances

A content-filtering appliance is a device that filters traffic in and out of a network. It actively monitors data streams by inspecting datagrams, searching for malicious code or malicious behavior. [Figure 6-6](#) shows the overall concept. You install the device so that all traffic going to or from the Internet goes through the device. Basic content-filtering appliances filter malware and spam. More advanced appliances can act as firewalls, proxy servers, and intrusion detection systems. Notice that both malware and spam enter the appliance, but the device can block malware and spam, preventing it from getting into the internal network.

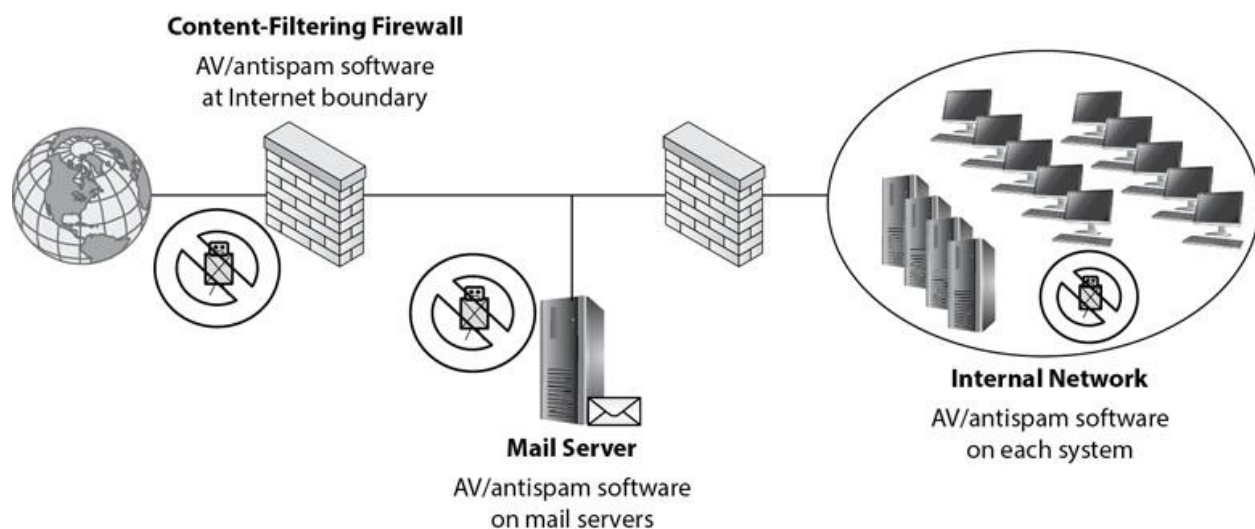


Figure 6-6 Content-filtering appliances

Similar to kitchen appliances, content-filtering appliances are rather easy to plug in and use. Think of a microwave kitchen appliance. It certainly has a lot of complexity within it, but all you have to do is plug it in, put your food in, punch some buttons, and in almost no time, you'll be enjoying a hot meal. You don't need to understand how the microwave works to enjoy its benefits. Similarly, an administrator doesn't have to understand the complexity of the content-filtering appliance to use it and enjoy its benefits.

Most content-filtering appliances include additional features and capabilities that administrators can adjust. For example, if the organization uses it as a proxy server, it checks the validity of website certificates and blocks sites that fail these checks. It can also filter URLs, controlling what websites users can visit.



NOTE [Chapter 4](#) presents proxy servers and how they can be used to control what websites users can access. Proxy servers can also reduce Internet bandwidth usage.

[Figure 6-6](#) also shows a defense in depth strategy to protect against malware by installing AV/antispam software at the boundary, on a mail server, and on all internal clients. Installing AV software in so many places might seem like overkill. After all, if most malware reaches systems via the Internet and e-mail, the content-filtering firewall and antivirus software on the mail server should catch it. Ideally, this is true. However, users can also inadvertently install malware via USB drives just by inserting the USB drive into their system. Additionally, zero day exploits might get past the content-filtering firewall and/or the mail server. AV software on the systems might be able to detect this malware using behavior-based detection methods. Together, these methods provide a defense in depth solution to protect against malware.



EXAM TIP Attackers commonly deliver malware over the Internet. Monitoring and scanning all Internet traffic is an important part of a defense in depth strategy to protect systems from malware.

Keeping Operating Systems Up to Date

It's also important to keep operating systems up to date. Malware often exploits software bugs or vulnerabilities in a system, and as long as the bug remains, the system is vulnerable. On the other hand, updated systems are not vulnerable to these bugs.

Vendors regularly write, test, and release patches for known vulnerabilities. For example, Microsoft releases patches on the second Tuesday of every month. The goal is to allow administrators time to plan for testing and installation of these patches on a regular basis. Although Microsoft does occasionally release patches for emerging threats outside of this schedule (out-of-band releases), these releases are relatively rare.

[Chapter 9](#) talks about hardening systems, or making them more secure from their default configuration. Keeping systems up to date is a key part of hardening a system.

Scanners

[Chapter 8](#) covers vulnerability scanners in more depth, but in short, vulnerability scanners can check systems for vulnerabilities and report their findings. Some common vulnerability scanners used by security professionals include Nessus, Nmap, and SAINT (Security Administrator's Integrated Network Tool).

Vulnerability scanners can check systems to verify that the operating system is up to date, has up-to-date antivirus software, and has an enabled firewall. A vulnerability scanner can also

verify that the systems are not running certain services and protocols, depending on how the administrator configures the scanner.

Beware of Shortened Links

It used to be that you could easily spot malicious links in e-mails. For example, if you saw a link claiming to be from mheducation.com, you could hover over it to see the actual link was www.hackers.ru. In this case, clicking the link would not take you to mheducation.com but instead would take you to a Russian website (because the *ru* top-level domain is in Russia).

As an example, [Figure 6-7](#) shows an HTML-based web mail with the user hovering over the second link. Although the link appears to be mheducation.com, if you look in the status bar at the bottom of the browser, you can see the actual link is <http://www.hackers.ru>. Many e-mail programs allow the user to hover over the e-mail in the same way and see the actual link as a tool tip.

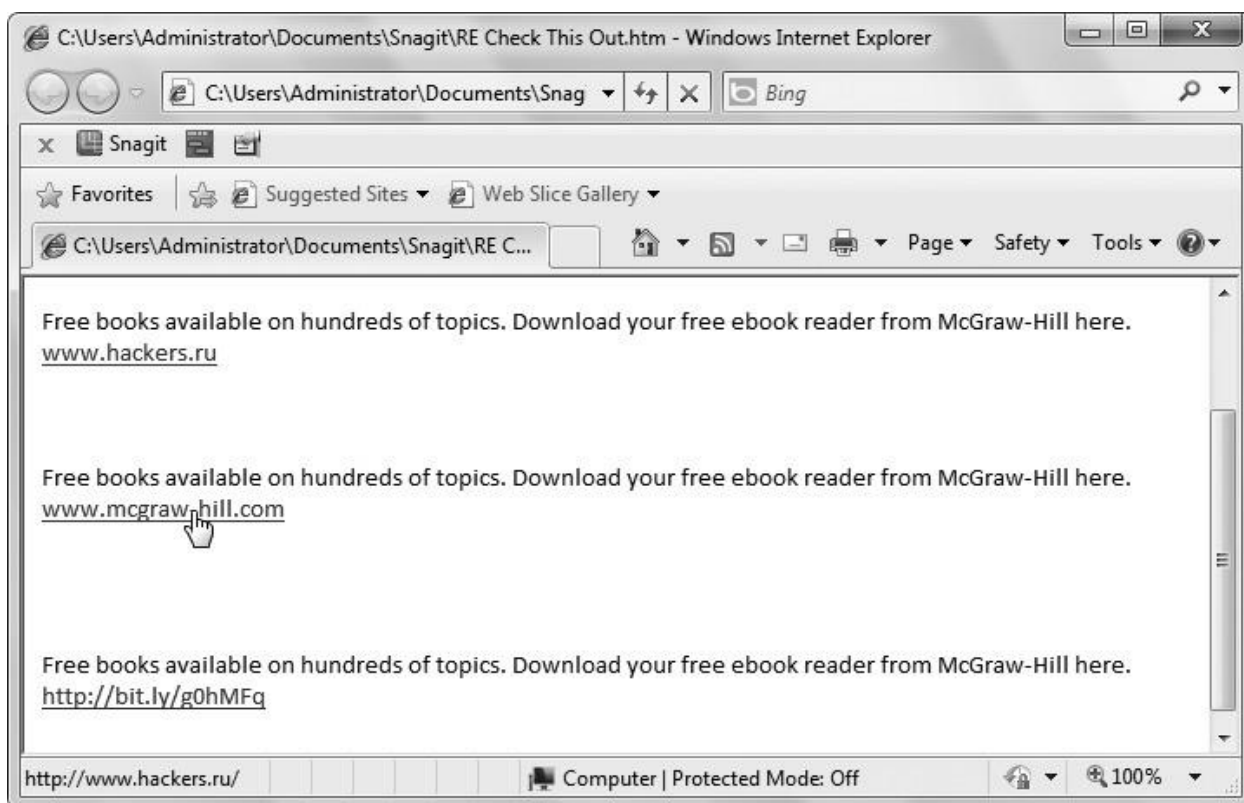


Figure 6-7 Embedding links in an e-mail

To prevent users from seeing the link, attackers sometimes use shortened URLs to mask the actual destination, as shown in the third link in [Figure 6-7](#). Legitimate users also shorten very long URLs with tools such as Bitly (<https://bitly.com>) or TinyURL (<https://tinyurl.com>). As an example, you can go to the Bitly website, plug in a long URL, and it gives you a short one.

Now, when a user clicks the bit.ly link, it takes the user to the actual long link. Although this is useful (especially for 140- or 280-character tweets), attackers have used shortened links for attacks. End users don't know where this link takes them until they click the link. By that point, it may be too late.

Some tools are available to expand URLs. For example, URL Expander (<http://urlex.org>) allows you to plug in a shortened URL and it expands it to a long URL. Users can copy the shortened URL from their e-mail or from a tweet and paste it into the text box on the <http://urlex.org> page.

As an example, you can use Bitly to shorten the URL of McGraw-Hill Education (<https://www.mheducation.com/>). The result is <https://bit.ly/2Fwm4Lz> (though it might be different if you do this at a different time). You can then go to <http://urlex.org> and plug in the shortened URL. Click Expand and it shows you that the expanded URL is <https://www.mheducation.com/>.

Sandboxing

As mentioned previously, behavior-based antivirus software sometimes uses a sandboxing technique. It runs new or untested code within an isolated environment. If the code is malicious, the code is unable to infect other applications or other areas of the system. Real-time scanning used by AV software uses sandboxing techniques.

Application sandboxing limits the environment where code can execute. One method that is often used is virtualization. It's possible to create virtual machines that are completely isolated from the host or the host's network. Additionally, AV researchers often use isolated virtual machines to analyze malware.

Least Privilege

Most operating systems require a user to have elevated permissions to install applications. Because of this, you can reduce the ability of malware to run on your system by logging on with an account that does not have elevated permissions. This follows a best practice of giving users only the rights and privileges they need to perform their job—the principle of least privilege. If users don't install applications as a part of their job requirements, they should not have the privileges to do so, and this reduces the ability of the user to accidentally install malware.

A best practice that many organizations follow is to require administrators to have two accounts. One account is a regular user account with little or no privileges, and the second account has elevated administrator permissions. The administrator uses the regular user account for normal, day-to-day activities and uses the elevated user account only when performing administrative work. These steps reduce the opportunities for malware to be installed.



TIP [Chapter 1](#) introduces the principle of least privilege as one of several fundamental security practices.

Software Security

[Chapter 5](#) discusses several software security methods that can also help protect against malware. This includes using input validation techniques and application reviews for web applications that an organization hosts. It also includes verifying that risky code has been signed with a digital certificate. Code signing helps reduce malware infections because criminals won't use code signing techniques to validate malware. It also helps identify modified code because a digital certificate will identify code that has been modified and display an alert.

Application Whitelisting and Blacklisting

Another malware protection method is application whitelisting or blacklisting. Application whitelisting identifies specific applications that can run on a system and blocks all other applications. This can be effective in an organization that has a limited number of applications on systems, but it does make it more difficult to add new applications. Application blacklisting identifies specific applications that cannot run on a system. This is more difficult because it requires you to know all potential malicious applications. [Chapter 9](#) covers application whitelisting and blacklisting as a method to secure endpoint devices.

Participating in Security Awareness and Training

Many users don't understand the risks associated with malware or how criminals try to trick them into installing malware. Because of this, it's important to educate users on safe computing habits. Security policies often mandate security awareness and training efforts to prevent security incidents before they occur. In other words, training is a preventive control. In contrast, detective controls attempt to detect incidents after they occur, and corrective controls attempt to reverse the effects of incidents after they occur.



TIP [Chapter 5](#) discusses the importance of educating users about social engineering. You can combine social engineering training with basic safe computing training.

The following list identifies some common items to repeat to users:

- Use spam filters and AV software.
- Don't open any attachments from unsolicited e-mails.
- Don't click any links in unsolicited e-mails.

- Be wary of all attachments and links in e-mails. Attackers often try to make e-mails look like they are from someone you know.
- Be wary of any warnings from a website, such as a warning that your system is infected with malware.
- Download files directly from the source website (not from a link provided through a third-party website).

Some organizations have security policies that require additional steps. For example, if a user notices a system is acting suspiciously as though it's infected with malware, the security policy may require that the user remove the computer from the network by disconnecting the network cable.

Common Vulnerabilities and Exposures

The Common Vulnerabilities and Exposures (CVE) list is a standardized list of known security vulnerabilities, exploits, and malware. Moreover, the CVE provides a standard way of naming information security vulnerabilities. The MITRE Corporation maintains the list, and the National Cyber Security Division of the U.S. Department of Homeland Security sponsors the CVE. You can access the CVE here: <https://cve.mitre.org/cve/>.



NOTE MITRE looks like an acronym, but it's just the name of the company and is presented in all uppercase. Although some of the early employees previously worked at the Massachusetts Institute of Technology (MIT) providing research and engineering services, MITRE is not a part of MIT.

The CVE has definitions for vulnerability and exposure that are worth knowing. A *vulnerability* is “a mistake in software that can be directly used by a hacker to gain access to a system or network.” An *exposure* is “a system configuration issue or a mistake in software that allows access to information or capabilities that can be used by a hacker as a stepping-stone into a system or network.”

Chapter Review

A significant security risk to computers and IT networks is in the form of malicious code, commonly called malware. Malware includes viruses, worms, Trojan horses, RATs, scareware, ransomware, keyloggers, logic bombs, rootkits, mobile code, backdoors, trapdoors, spyware, and malware hoaxes.

Viruses have a replication component to spread themselves and an activation component to deliver the payload of the virus. They use a variety of methods to thwart AV software and researchers. Stealth viruses attempt to hide themselves to escape detection. Armored viruses use techniques such as complex code and encryption to prevent a researcher from reverse engineering the code and discovering what the virus is doing. Polymorphic and metamorphic viruses change the virus file to escape detection through signature files.

Although viruses must execute, a worm can travel through a network without any human interaction. A Trojan horse appears to be one thing, such as a useful program, but is actually something different and malicious. Remote access Trojans (RATs) allow attackers to take control of a system from a remote location. Scareware is a type of Trojan. Scareware appears to be a useful antivirus program, but instead attempts to extort money from users. Ransomware locks a user's computer or data and demands ransom to return the control back to the user.

A keylogger captures keystrokes from users and can send the captured keystrokes to an attacker. Logic bombs execute in response to an event, such as a day or time. Rootkits take control of the operating system (at the kernel or root level) and have the ability to take complete control over what the user sees and does. A backdoor (or trapdoor) is code embedded into an application or operating system that allows an attacker to access data or even gain remote control access to the infected computer. Malware often installs RATs as backdoors into systems. Spyware is any software that is installed without the user's knowledge or consent and can be malicious. Most malware is delivered through the Internet via drive-by downloads, malvertising, and malicious e-mails.

The primary method of protecting systems from malware is the use of up-to-date AV software. AV software should be installed on every computer. It's also highly recommended to install AV software at the boundary between the Internet and the internal network and on e-mail servers. AV software uses signature definitions and behavior-based detection methods (sometimes called heuristics) to identify malware. Signature definitions identify known viruses based on a specific pattern within the virus. Behavior-based detection methods attempt to identify malicious behavior of previously unknown viruses. Because viruses can mutate into different variations, it's important to keep virus definitions up to date. Additionally, it's important to keep an operating system up to date.

Spam filters and content-filtering appliances help filter out malware before it reaches the internal network. Content-filtering appliances can also filter other traffic, such as by using URL filters to control which websites users can access and performing firewall and intrusion detection services.

Other methods of protecting against malware include using vulnerability scanners to identify potential vulnerabilities in the network, using sandboxing techniques to isolate potentially malicious code, and implementing the principle of least privilege for all accounts. Many software security methods such as input validation, application reviews, and code signing help protect against malware infections. Application whitelisting identifies applications that are allowed and blocks all other applications. Application blacklisting identifies applications that are not allowed and allows all other applications.

By educating users on the basics of safe computing habits, you can help them avoid risky computing behaviors that cause malware infections. This includes not opening attachments or clicking links from unsolicited e-mails, and being wary of all attachments and links in e-mails.

The MITRE Corporation maintains the CVE list, which provides a standardized method of describing malware. The National Cyber Security Division of the U.S. Department of Homeland Security sponsors the CVE.

Questions

1. Which of the following malware types alters its own code to avoid detection by antivirus software?

- A. Armored virus
- B. Metamorphic virus
- C. Polymorphic virus
- D. Ransomware

2. What type of malware can spread without any user intervention?

- A. Virus
- B. Trojan horse
- C. Worm
- D. Spyware

3. Joe, an employee in the accounting department, received an e-mail with a subject line of “Urgent” and an attachment labeled invoices.zip. The text says “Process these invoices as soon as possible.” When Joe opens the ZIP file, his system reboots, and he can no longer locate the e-mail with the ZIP file. Which of the following provides the best explanation of this activity?

- A. The ZIP file is ransomware.
- B. The ZIP file is a polymorphic virus.
- C. The ZIP file is a Trojan horse.
- D. The ZIP file is scareware.

4. An attacker has successfully infected a user's system. After doing so, the attacker begins scanning the network from the infected computer. Which of the following BEST describes this activity?

- A.** RAT
- B.** APT
- C.** Rootkit
- D.** Ransomware

5. After visiting a website, a user sees a pop-up indicating a virus has infected his system and offering free antivirus software. He downloads the free antivirus software, but finds that it won't clean the virus unless he purchases the full version. What does this describe?

- A.** Shareware
- B.** Rootkit
- C.** Freeware
- D.** Scareware

6. When Sally turns her computer on, she sees a screen indicating software has encrypted all of her data files. A message indicates she must pay \$300 within 48 hours to access the decryption key. What does this describe?

- A.** Logic bomb
- B.** Ransomware
- C.** Worm
- D.** Spyware

7. An employee configured malicious code to execute at midnight on February 2. What does this describe?

- A.** Logic bomb
- B.** Groundhog Day virus
- C.** Worm
- D.** Ransomware

8. What type of malware takes control of the operating system at the kernel level?

- A. Trojan horse
- B. Worm
- C. Keylogger
- D. Rootkit

9. A website developer wants to provide assurances to users that ActiveX controls used on the site are not malicious. What can provide this assurance?

- A. Input validation
- B. Code signing
- C. Code review
- D. Enabling cross-site scripting

10. What does antivirus software use to detect previously unknown viruses?

- A. Signatures
- B. Polymorphism
- C. Heuristics
- D. Armor

11. A virus is detected on a system based on the virus's behavior. What detected the virus?

- A. Heuristics
- B. A virus fingerprint
- C. A virus filter
- D. A signature

12. What should users do to ensure that antivirus software can detect recently released viruses?

- A. Update signatures
- B. Update the operating system

C. Update the AV software

D. Regularly purchase new AV software

13. Of the following choices, how is malware most often delivered today?

A. Over the Internet

B. Via an intranet

C. Via USB drives

D. Through company policies

14. A company authorizes users to transport data from work to home using USB drives. What's the best method of protecting systems from malware without affecting the user?

A. Install AV software on the network firewall

B. Install AV software on the e-mail server

C. Install AV software on each user's work computer

D. Prevent users from using USB drives

15. Of the following choices, what is the best technique you can implement on an e-mail server to reduce infection through e-mail?

A. Block all e-mail

B. Add a spam filter

C. Add a polymorphic filter

D. Remove all attachments

16. Of the following choices, which are valid methods to reduce malware infections? (Select all that apply.)

A. Don't open attachments from unsolicited e-mails.

B. Don't click links in unsolicited e-mails.

C. Don't send encrypted personal information via e-mail.

D. Don't follow shortened links from unknown sources.

17. Of the following choices, which one is a principle that prevents users from accidentally installing malicious software on their systems?

- A.** Nonrepudiation
- B.** Least privilege
- C.** Separation of duties
- D.** Accountability

18. Of the following choices, what network device can filter e-mail, spam, and malware?

- A.** Packet-filtering firewall
- B.** Proxy server
- C.** An intrusion detection system
- D.** Content-filtering appliance

19. Your organization mandates security training for users within its security policy to educate users about malware and methods to prevent malware infections. What is the best description of this effort?

- A.** A detective control
- B.** A corrective control
- C.** A preventive control
- D.** A technical control

20. What provides a standardized method of describing malware?

- A.** The Consortium of Antivirus Vendors (CAV)
- B.** The Consortium of Virus Authors (CVA)
- C.** The National Institute of Standards and Technology (NIST)
- D.** The Common Vulnerabilities and Exposures (CVE) list

Answers

- 1. B.** A metamorphic virus changes or mutates its code as it replicates itself to prevent detection. An armored virus uses techniques such as encryption to make it more difficult for AV researchers to decompile the virus. A polymorphic virus changes the file, but not the code. Ransomware takes over a user's computer and demands a monetary ransom to return control back to the user.
- 2. C.** Worms spread through a network without any user intervention. Viruses, Trojan horses, and spyware all require some level of interaction.
- 3. C.** The ZIP file is most likely a Trojan horse, appearing to be a ZIP file of invoices, but instead has a malicious component. While the ZIP file could have installed ransomware, the scenario doesn't indicate any actions other than rebooting the system and hiding the original e-mail and attachment. Similarly, the virus might have polymorphic characteristics causing itself to mutate, but simply hiding the e-mail doesn't indicate the virus mutated. Scareware uses a pop-up to indicate a system is infected with malware and attempts to trick a user into installing malware.
- 4. A.** A remote access Trojan or remote access tool (RAT) allows an attacker to use an infected computer to scan a network. While advanced persistent threats (APTs) often use RATs, any criminal can use a RAT to scan a network from an infected computer. A rootkit has root-level access on a system, but the scenario doesn't indicate root-level access. Ransomware takes control of a user's system or data and demands payment from the user.
- 5. D.** Scareware is malware that scares users into thinking a virus has infected their system and encourages them to install a free download. The free download appears as antivirus software that doesn't remove viruses unless users pay, but it often includes malware itself. Shareware is software that users are free to try and pay for if they like it and continue to use it. A rootkit takes over the system with root-level privileges. Freeware is free software.
- 6. B.** Ransomware takes control of a user's computer or data and demands a ransom to return control to the user. This scenario describes CryptoLocker. A logic bomb is malware that executes in response to an event such as a specific date and time. Worms infect computers over a network, and while worms deliver malware, not all worms include ransomware. Spyware is software installed on a user's system without the user's knowledge with the goal of spying on the user, not extorting money from the user.
- 7. A.** A logic bomb is malware that executes in response to an event such as a specific date and time. While February 2 is Groundhog Day, the scenario doesn't describe a Groundhog Day virus. Worms infect computers over a network, not on a specific day. Ransomware takes control of a user's computer or data and demands a ransom from the user.
- 8. D.** A rootkit is a set of programs that runs on a system, largely undetected, because it runs at the kernel level or root level of the operating system. A Trojan horse is malware that looks like one thing but is something else. A worm is a type of malware that spreads through a network without any user intervention. A keylogger captures keystrokes from users.

9. B. Code signing digitally signs ActiveX controls and provides assurances to users who created the control and that it hasn't been modified. Input validation helps prevent injection attacks, but it's used to protect the website, not provide assurance to users. Code review is a valuable tool to detect problems with applications before an organization releases them. Cross-site scripting is an attack and would not be enabled.

10. C. Antivirus software uses heuristics to detect previously unknown viruses. Signatures detect known viruses. Polymorphism and armor are techniques used by virus authors to prevent the detection of a virus.

11. A. Heuristics can detect malware based on the behavior of the malware and are designed to detect previously unknown viruses. There's no such thing as a virus filter or a virus fingerprint, although a virus signature does uniquely identify known malware similar to how a fingerprint can identify a person.

12. A. Antivirus software uses signature definition files to detect viruses, and these signatures must be regularly updated. It's not necessary to update the operating system, update the AV software, or purchase new AV software to detect recently released viruses.

13. A. The common way attackers deliver malware is over the Internet. While some attacks can come from internal intranet sources, they do not compete with the volume of attacks from the Internet. Unsuspecting users transmit viruses with USB drives, but this isn't as common as virus delivery over the Internet. Company policies would not deliver viruses.

14. C. Installing AV software on each user's work computer provides the best protection against a user inadvertently transporting malware from home to work. Installing software on the network firewall and on an e-mail server is a good practice, but it won't help if the virus is transported via a USB drive. Preventing the users from using USB drives will affect the users.

15. B. The majority of malware comes through spam, so a spam filter can reduce infections through e-mail. An e-mail server isn't very useful if it blocks all e-mail or removes all attachments. E-mail servers don't have polymorphic filters.

16. A, B, D. Common methods that help reduce malware infections are: not opening attachments from unsolicited e-mails, not clicking links in unsolicited e-mails, and not clicking shorted links from unknown sources. Sending e-mail in an encrypted format helps protect confidentiality, but it doesn't necessarily reduce malware infections.

17. B. The principle of least privilege specifies that users are given rights and privileges to do their job but no more. If a user doesn't need to install applications, the user is not given permission to do so, which reduces the possibility of the user accidentally installing malware. Nonrepudiation prevents a person from denying an action. Separation of duties divides tasks so that no single person or entity controls an entire process. Accountability ensures that user actions can be tracked and monitored.

18. D. Content-filtering appliances can filter e-mail to remove spam and malicious attachments. They can also act as a proxy server or an intrusion detection system (IDS). A packet-filtering firewall can only examine packets. A proxy server can filter websites but not e-mail. An IDS detects attacks, but not e-mail and spam.

19. C. Training is a preventive control because it attempts to prevent incidents from occurring. Detective controls attempt to detect incidents, and corrective controls attempt to reverse the effects of an incident. Technical controls use technology to implement the control, but training doesn't require technology.

20. D. The CVE is maintained by the MITRE Corporation and provides a standardized method of describing security vulnerabilities, exploits, and malware. There is no such thing as the Consortium of Antivirus Vendors (CAV) or Consortium of Virus Authors (CVA). NIST is a U.S. government entity that regularly publishes standard publications related to IT security, standards, and practices, but it does not maintain the CVE.

NOTES

CHAPTER 7

Risk, Response, and Recovery

In this chapter, you will learn about

- The definition of risk
- Threat sources, threat events, and vulnerabilities
- Impact assessments and sharing threat intelligence
- Different methods of managing risk
- The definition of residual risk
- Risk registers and the Common Vulnerability Scoring System
- Common risk management frameworks
- Risk assessments and threat modeling
- Differences between quantitative and qualitative risk analysis
- Common steps in risk assessments
- Common phases in the incident lifecycle

Defining Risk

The simple definition of *risk* is the probability or likelihood that a threat will exploit a vulnerability and cause a loss. A *threat* is any activity that can be a possible danger. A *vulnerability* is a weakness, and a *loss* represents a negative event for an organization. [Figure 7-1](#) shows these components and their relationships with each other.

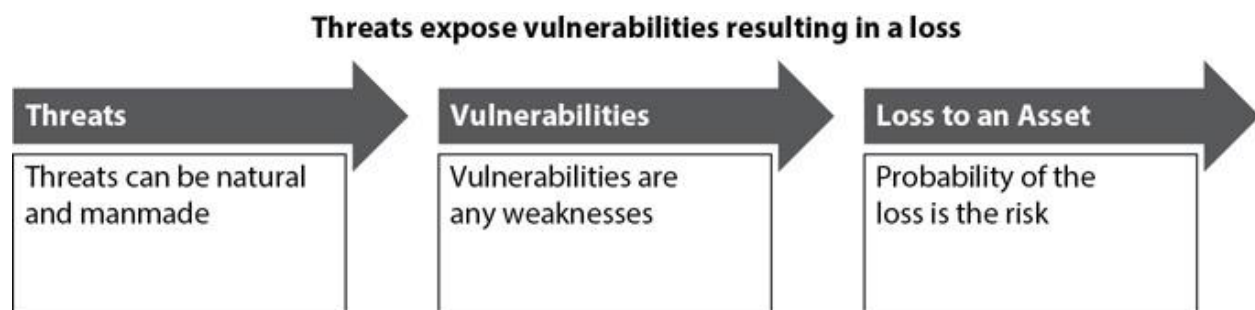


Figure 7-1 Components of risk

A common formula used to represent risk is

$$\text{Risk} = \text{Threat} \times \text{Vulnerability}$$

This formula isn't a mathematical algorithm. However, it does show the relationship between threats and vulnerabilities and risk. Small threats and insignificant vulnerabilities result in a low level of risk. More severe threats and serious vulnerabilities result in a higher level of risk.

When the threat and the vulnerability are combined (a threat exploits a vulnerability), it results in a loss. More significant threats or vulnerabilities can easily result in higher losses from the risks. In other words, some risks have higher impacts. However, if you can reduce the vulnerabilities, you reduce risk, resulting in lower potential losses.



TIP It's rarely possible to reduce threats. You can't stop an attacker from attacking, but by reducing vulnerabilities, you reduce the attacker's chances of success. In other cases, you can reduce the impact of the risk to limit losses.

For example, malware represents a threat to any computer. If a computer does not have antivirus (AV) protection, it is vulnerable to malware threats. If a virus infects an unprotected system, it can modify, destroy, or steal data, representing a loss to the owner of the computer. An organization can mitigate, or reduce, the vulnerability by using AV software, keeping it up to date with current antivirus signatures, and training users on how to avoid threats.



EXAM TIP Risk is the probability or likelihood that a threat will exploit a vulnerability and cause a loss. Risk mitigation is the practice of reducing risk. The primary method of reducing risk is to reduce or eliminate vulnerabilities by implementing safeguards or security controls. It's also possible to reduce risk by reducing the impact of the risk.

Total risk is the combined risk to all of the organization's assets, including all the threats and vulnerabilities. Total risk also includes the value of these assets.

NIST

The U.S. National Institute of Standards and Technology (NIST) includes a Computer Security Division and an Information Technology Laboratory (ITL) within that division. The ITL performs research to advance the development and productive use of information technology and regularly publishes documents known as Special Publications (SPs). The SP 800 series of reports

includes the results of this research performed in collaboration with industry, government, and academic organizations.

You can review a full list of the NIST ITL SP 800 series documents here:

<https://csrc.nist.gov/publications/sp>. You can also download any of these documents as PDF files.

NIST SP 800-30, *Guide for Conducting Risk Assessments*, provides a more comprehensive definition of risk. It defines risk as “a function of the likelihood of a threat event’s occurrence and potential adverse impact should the event occur.” It also mentions that a threat source initiates a threat event, which exploits a vulnerability, causing an adverse impact. [Figure 7-2](#) shows the risk model outlined in NIST SP 800-30, and the following list explains these elements:



Figure 7-2 Risk model

- **Threat source** A threat source is anyone or anything that can initiate a threat event. As an example, an attacker is an adversarial threat source.
- **Threat event** A threat source initiates a threat event. For example, an attacker can launch a SQL injection attack against a web server. The attacker is the adversarial threat source, and the SQL injection attack is the threat event.
- **Vulnerability** The extent of the threat event is related to the extent of the vulnerability. For example, if a website doesn’t use input validation techniques, it is more vulnerable to a SQL injection attack than a website that does use input validation techniques. Security controls, such as input validation techniques, reduce vulnerabilities.
- **Adverse impact** An adverse impact affects the confidentiality, integrity, and/or availability of IT systems or data, and in some cases can affect the image and reputation of the organization.



TIP [Chapter 1](#) presented the concepts of confidentiality, integrity, and availability (CIA). To recap, maintaining confidentiality means that unauthorized individuals don’t have access to data or systems. Maintaining integrity means that unauthorized individuals can’t modify data and systems. Maintaining availability means that data and systems are up and operational when needed.

While the SP 800-30 definitions and explanations are more technical, they do repeat the common themes of threats, vulnerabilities, impacts, and losses.

Identifying Threat Sources

While it's common to think of an attacker as a threat source, there are several other threat source types. NIST SP 800-30 categorizes them as adversarial, accidental, structural, and environmental:

- **Adversarial threat sources** Adversarial threat sources include individuals (internal or external) or groups that purposely attempt to exploit weaknesses. They are often referred to as threat actors. Internal threats are from malicious employees, contractors, or any other personnel granted access to an organization's resources and have the capability of launching an attack from within. External threats are from entities that don't have direct access to the organization's resources, such as attackers in a different city or country. Groups include both loose-knit groups of attackers and highly organized attackers (such as state-sponsored advanced persistent threats). Groups can also include competitors, suppliers, partners, and criminals.
- **Accidental threat sources** Accidental threat sources include erroneous actions taken by individuals in the course of executing their everyday responsibilities. For example, a user can ruin a laptop simply by spilling coffee onto the keyboard. Similarly, if someone accidentally kills the power or cuts communications cables to key systems, it results in a loss of availability. Users can accidentally cause damage to data, and the level of damage they cause is directly related to their level of access. Regular users have limited access to resources, so they can only cause a limited amount of damage (such as accidentally deleting some of their files). Privileged users such as administrators have much more access, so they can accidentally cause much more damage (such as deleting an entire customer database). While IT security topics often stress intentional threats, it's important to realize steps need to be taken to protect assets from accidental threats also.
- **Structural threat sources** Structural threat sources include failures of IT equipment, environmental controls, and software. As an example, if the environmental controls fail in a data center resulting in a loss of air conditioning, it can result in damaged equipment. Without air conditioning, servers get hot, and they can overheat and fail. Because of this, many organizations will power down servers when air conditioning fails. This protects the equipment, but also results in an immediate loss of availability.
- **Environmental threat sources** Environmental threat sources include natural and manmade disasters. Natural causes include weather events such as lightning storms, hurricanes, tornadoes, and floods, and can include nonweather events such as earthquakes or tsunamis. One example of a manmade disaster is a fire. Environmental threat sources also include failure of external services such as a telecommunications link, or power outages.



TIP Appendix D in SP 800-30 includes additional information on threat sources, including several tables listing threat sources and sample assessment scales used to categorize them.

Identifying Threat Events

Threat events include any type of activity or event that can result in a loss of confidentiality, integrity, or availability to a system. SP 800-30 matches threat events with threat sources to identify threat scenarios. A threat scenario is one or more threat events, caused by one or more threat sources, that cause an adverse impact.

Many times, you'll see the threat source and threat event combined as simply a threat. However, when performing risk assessments, security professionals often separate the two so that they can more accurately identify who the threat is (the threat source) and what they might do (the threat event).

NIST SP 800-30 includes an extensive list of potential adversarial threat events, grouped into several categories. These categories focus on known attack methods and activities, and the following list includes some of them:

- **Perform reconnaissance and gather information** Adversaries use a variety of techniques to learn more about their target. These techniques include gathering information from public sources and using tools such as sniffers and scanners to gather technical information. Attackers often try to trick users into installing malware on their computers. When successful, it allows the attacker to remotely control the computer and perform reconnaissance on the infected computer's internal network. These attacks from within the organization become more difficult to track because most detection tools focus on attacks coming from outside the organization.
- **Craft or create attack tools** Attackers create a wide variety of different attack tools. For example, they often use phishing attacks (including spear phishing and whaling attacks) to collect information and trick users into installing malware on their systems. They have been quite creative in their methods, including spoofing e-mails to appear to originate from funeral homes, delivery services, and more. Their goal is to get the user's interest. To complement these methods, attackers develop malware, which they include as attachments.
- **Deliver/insert/install malicious capabilities** Once attackers create the malware, they find ways to get the malware installed on user systems. Attackers often take the time to analyze methods that have worked before and create similar methods. For example, many users don't recognize the risk of opening a ZIP file, so attackers often embed malware within ZIP files, which they send as attachments with phishing e-mails. They have also attacked websites and installed malware on them. The malware then attempts to download itself automatically when users visit.
- **Exploit and compromise** Exploits include both technical and nontechnical methods. For example, social engineers circumvent physical security methods by tailgating closely behind authorized individuals to access secure locations. Technical methods exploit known vulnerabilities in unpatched systems. Additionally, attackers are always on the alert for zero day vulnerabilities that don't have patches. Many exploits allow attackers to assume control of compromised systems from remote locations.

- **Achieve results (such as causing adverse impacts and obtaining information)** Once attackers successfully exploit systems, they go after their target. This is often to obtain sensitive information from the victims, cause damage to the systems, and/or infiltrate a network to steal data. Adverse impacts include any results that affect the confidentiality, integrity, or availability of systems and data.

- **Maintain a presence or set of capabilities** Once attackers are within a system or network, they take steps to obscure their actions. This includes using methods to bypass or circumvent intrusion detection systems and auditing capabilities. Security firm Mandiant has documented several cases where an advanced persistent threat (APT) maintained a presence within networks for several months and sometimes for several years.

It's worth stressing that the previous list does not include all threat events. Attackers have a wide variety of methods they use to attack organizations. As security professionals find ways to thwart these attacks, attackers modify their methods using new and modified attacks.

Understanding Vulnerabilities

Vulnerabilities are any weaknesses in a system, network, infrastructure, or organization. A vulnerability can also be the absence of a security control or countermeasure. For example, a firewall protects systems and networks from Internet threats. A network without a firewall indicates a vulnerability. NIST SP 800-30 defines a vulnerability as follows:

A vulnerability is a weakness in an information system, system security procedures, internal controls, or implementation that could be exploited by a threat source.

If a threat can exploit these weaknesses, it results in a loss. Some examples of vulnerabilities include the following:

- **Poor or nonexistent AV protection** Computer systems need antivirus software with up-to-date virus signatures. If systems don't have AV protection, malware can infect them, resulting in costly losses.
- **Terminated employees** If employee accounts are not disabled or deleted after an employee is terminated, the account can be used to attack a system. The employee may be able to access the system and cause damage, or another employee may use the account to perform actions. If another employee accesses the account, logs will inaccurately attribute the actions to the terminated employee.
- **Weak access controls** Access controls include identification, authentication, and authorization techniques. Weak access controls allow unauthorized individuals to access, modify, and/or destroy data.
- **Poor or nonexistent change-management practices** An unintentional side effect of making a change to a system is that the change causes one or more problems. Strong change-

management practices examine prospective changes to predict the side effects and can mitigate problems caused by changes.

- **Poor hardening practices** Systems should be hardened, or made more secure from their default configuration. This includes changing defaults such as default passwords, implementing firewalls, keeping operating systems and applications up to date with patches, and more.
- **Lack of redundancies for key systems** Any single point of failure can result in an outage of an entire system if it fails. Administrators can add redundant technologies for disk drives, servers, communication lines, power, air conditioning, and even entire sites or locations. Adding redundancies allows systems to tolerate faults. For example, if a single drive in a RAID-6 array fails, the disk subsystem can tolerate the fault and continue to operate.
- **Uneducated users** Social engineers can trick uneducated users into giving up key information. On the other hand, social engineers aren't as successful when they try to trick educated users. Similarly, educated users are less likely to be tricked into installing malware onto their system by opening attached files or clicking links in malicious e-mail.



TIP The preceding list is by no means complete, but it does provide some common examples of vulnerabilities. An important point to remember about any vulnerability is that you can reduce weaknesses by implementing controls or safeguards. By reducing the vulnerabilities, you reduce the possibility of the risk resulting in a loss.

It's worth mentioning that an organization doesn't necessarily have to eliminate all vulnerabilities. For example, some legacy systems might be critical to the business, but include known vulnerabilities that cannot be mitigated. Similarly, some patches adversely affect applications. While testing a patch, administrators might detect the problem and report it to management, and management might decide to accept the risk of leaving the system unpatched.

Threats often exploit the weakest link in your organization, and once a threat exploits a vulnerability, the damage is done. A key method used to minimize losses is to identify and reduce vulnerabilities.

Understanding Impact Assessments

The probability that a threat can exploit a vulnerability provides the likelihood of occurrence. However, it's also important to examine the potential impact of the risk occurring as part of an impact assessment. The *impact* refers to the magnitude of harm that could result if a threat exploits a vulnerability. You can identify the magnitude of harm by estimating the potential loss of confidentiality, integrity, or availability. In some cases, the impact is negligible, making the risk a low priority. In other cases, the impact is quite high, making the risk a very high priority.

Consider the risk of a fire. Overall, if personnel within an organization use safe practices, the risk of a fire is low. Fires do occur, so organizations often install fire detection and fire extinguishing equipment to automatically detect and extinguish a fire. Still, if these methods fail, a fire can destroy most of the contents of a building in a very short time. The impact can be catastrophic. In this case, the likelihood of occurrence for a fire is low, but the potential impact is high, so organizations purchase insurance to reduce the overall risk by reducing the impact.

In contrast, consider the likelihood that students in a classroom might steal a \$10 mouse. The mice might be easy to remove and hide in a book bag, so the likelihood might be high. However, the impact of the loss is very low, so the overall risk is low.

Risk assessments consider the impact of the risk, although they use different methods such as quantitative and qualitative risk analysis methods. A quantitative risk assessment uses costs and monetary values to identify the impact. A qualitative risk assessment uses subjective opinions of experts to identify the impact. Both are covered later in this chapter.

Sharing Threat Intelligence

Cyber threat intelligence includes information that organizations can use to identify, assess, monitor, and respond to potential threats. It typically includes tactics, techniques, and procedures (TTPs) used by threat sources along with methods to detect known attacks. When organizations share this threat intelligence appropriately, both inside and outside the organization, they have a better chance of gaining knowledge from each other and thwarting more attacks.

In contrast, imagine that one organization is successfully attacked. Security professionals within the organization eventually discover the attack and take steps to stop it along with methods to detect, contain, and prevent a recurrence of the attack. However, management decides to keep this intelligence private within the organization. There are two drawbacks to this decision. First, criminals can use the same TTPs to successfully attack other organizations. Second, criminals are likely to slightly modify their TTPs in future attacks and may be doing so already. The organization that doesn't share its intelligence is unlikely to hear about intelligence shared by other organizations about these modified attacks.

NIST SP 800-150, *Guide to Cyber Threat Information Sharing*, provides some basic guidelines and recommendations related to sharing information on cyber threats. The following list summarizes some of these as paraphrased recommendations:

- **Establish goals and objectives related to information sharing** These goals and objectives should support businesses processes and security policies.
- **Identify internal sources that have cyber threat information** These include the tools and sensors used to collect information, and the tools that collect, collate, and analyze the data, converting into actionable intelligence.

- **Specify the scope of information sharing** It is rarely useful for an organization to share all internal information. Instead, an organization should identify the types of information it is willing to share.
- **Establish rules related to information sharing** Rules or guidelines identify when this information should be shared and who can authorize the release outside of an organization.
- **Participate in information sharing efforts** It is useful to identify public and private communities, government repositories, and commercial sites that routinely provide information on cyber threats, and other sources that provide information on cyber threats.
- **Provide feedback to add context, corrections, and/or improvements** Organizations that actively seek to enrich the data on TTPs help mature and improve the overall worth of the data to all.
- **Protect the security and privacy of sensitive data** When sharing threat intelligence, organizations should still maintain the confidentiality of any sensitive data.

While the previous list obviously doesn't cover all 43 pages of SP 800-150, it does give an overview of some of the important concepts. Some organizations will have more data to share than others.

An obvious question is "Who should an organization share their threat intelligence with?" Interestingly, SP 800-150 doesn't spell this out clearly, but that does add credibility to the document. It's published by a U.S. government entity, so it could easily recommend sharing your data with the Federal Bureau of Investigation (FBI), the Central Intelligence Agency (CIA), or the Department of Homeland Security (DHS). It doesn't.

However, SP 800-150 does mention that many public and private sharing communities exist. A community might be "based on geographic region, political boundary, industrial sector, business interest, or threat space" (such as one focused on phishing attack threats). Sharing partners could be vendors, service providers, industry sector peers, supply chain partners, business partners, law enforcement agencies, and, yes, government agencies. Appendix A includes seven different scenarios related to cyber threat information sharing.

The cyber security community has developed several technical specifications designed to enable automated sharing of threat intelligence data. The following list describes three standards that organizations might use:

- **Trusted Automated eXchange of Indicator Information (TAXII)** TAXII provides definitions for services and message exchanges. It doesn't provide standards on what data to provide, but instead identifies methods to share data. Interested parties can poll TAXII servers to download data or use subscriptions to automatically receive data. TAXII servers also allow entities to upload data to various intelligence categories.

- **Structured Threat Information eXpression (STIX)** STIX provides a standardized, structured language to represent cyber threat intelligence. The language allows personnel to report on active attacks, targets, threat indicators, threat, actors, and TTPs. STIX exchanges this data using TAXII services and messages.

Managing Risk

Risk management attempts to identify, assess, and mitigate risk to an acceptable level for the organization. As previously discussed, two key components of risk are threats and vulnerabilities, so these are key components to address when attempting to manage risk. It's important to realize that risk management is an ongoing process. Threats constantly change and evolve, and risk management practices need to change and evolve with them.

It isn't possible to reduce the overwhelming majority of threats. However, you can reduce vulnerabilities and reduce the impact of threats. For example, you can't stop a hurricane or an earthquake. However, you can reduce the impact of these natural threats by taking additional precautions, such as building stronger buildings or simply building them outside of a threat zone. Similarly, you can't stop an attacker from trying to launch attacks against your organization, but you can take steps to reduce the attacker's possibilities of success by strengthening security practices and reducing vulnerabilities.

Risk Treatment

Overall, an organization can select from several choices in response to different risks. [Figure 7-3](#) shows these choices, and the following list describes them:

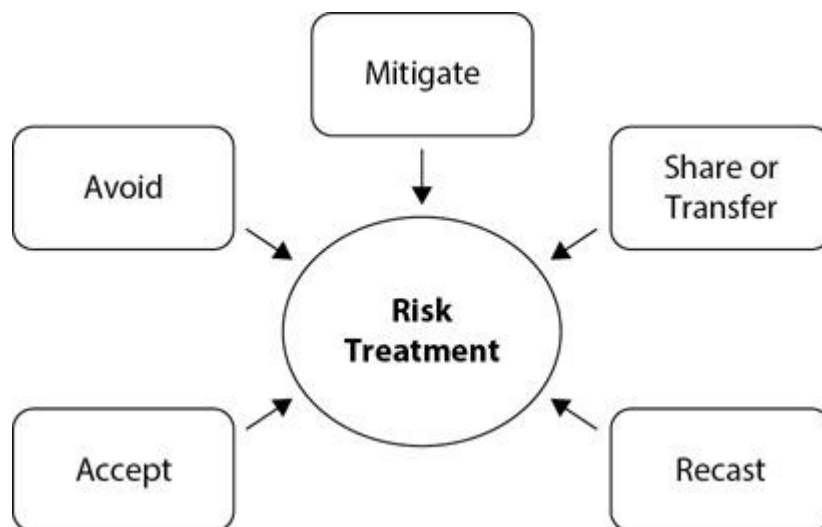


Figure 7-3 Risk treatment

- **Mitigate** You mitigate a risk by reducing vulnerabilities, and you reduce vulnerabilities by implementing controls or safeguards. For example, you can add AV software to reduce the risk

of malware infection. You can add intrusion detection systems to detect and prevent network attacks from penetrating your network.

- **Avoid** You can avoid a risk by simply avoiding the activity that introduces the risk. For example, if you want to avoid the risks of a hurricane, you can move your location to an inland location. If you want to avoid risks associated with hosting a website on the Internet, you can choose not to host a website.
- **Share or transfer** You can share a risk with another party, or transfer the risk to another party. The most common method of transferring a risk is with insurance. For example, you can purchase fire insurance, and if a fire destroys a building or a business, the insurance company covers the loss. It isn't always possible to transfer all responsibility of the risk to another party. For example, buying insurance might help an organization cover some losses, but the insurance doesn't necessarily cover all the losses.



TIP Previous documentation listed transfer as a risk management method but did not include share. However, that has been changing in recent years. Many security professionals recognize that it often isn't possible to transfer the entire risk, but it is possible to transfer some of it, or share the risk with an outside entity. For this reason, this option has changed from transfer to share or transfer.

- **Accept** You can also choose to accept the risk and its potential losses. This is appropriate when the asset value is low or if the cost to reduce the risk is higher than the value of the asset. For example, someone may steal a mouse, but the cost of securing mice to prevent theft is higher than the cost of the mice. Additionally, it's appropriate to accept the remaining risk after taking steps to reduce the risk to an acceptable level.
- **Recast** A newer risk treatment mentioned in the SSCP objectives is to recast a risk to a different severity. As an example, security professionals within an organization may have previously decided to accept a risk. However, based on current threats, they've decided to recast it to a more severe risk and implement controls to mitigate the risk. Another use of recasting is to downgrade the risk. For example, management may decide a risk isn't as great a threat as previously deemed and recast it to a lower level.

Residual Risk

Residual risk is the amount of risk that remains after mitigating, or reducing, risk. In other words, total risk – mitigated risk = residual risk, as shown in [Figure 7-4](#). It's important to ensure that residual risk is at an acceptable level to the organization. If the residual risk is too high, then the organization needs to find additional security controls to reduce the risk.



Figure 7-4 Residual risk

An old saying is that the only way to keep a computer 100 percent safe is to never plug it in and turn it on. Of course, the computer isn't very useful in this state. Instead, management decides what controls to implement to reduce the risk to an acceptable level. These decisions result in mitigating some risks and accepting other risks.

With this in mind, it's important to remember that management is responsible for any residual risk. Additionally, management is responsible for any losses associated with residual risk. Management chooses which mitigation methods to implement and which methods to reject. Their goal is to reduce risk to an acceptable level based on the needs of the organization.



EXAM TIP Residual risk is the risk that remains after reducing risk to an acceptable level. Management is responsible for approving (or rejecting) controls that can mitigate risk. Management is also responsible for any losses resulting from residual risk.

Of course, different organizations have different risk levels. A company that takes anglers out fishing won't intentionally go into a war zone. However, a naval warship will intentionally go into harm's way when necessary, even though the crew will take steps to reduce risk.

Management has an overall view of the entire organization, including understanding the organization's mission and vision. Security personnel won't necessarily know all the reasons why management might decide to implement one control to reduce a risk but reject another control and accept a different risk. Still, the key is that management is responsible for making these decisions, and is responsible for any losses resulting from residual risk.

Identifying Assets

An important step in risk management is identifying an organization's valuable assets. Assets with high value to the organization need to be protected by implementing controls or safeguards to reduce risks. While this seems obvious, it's not always apparent to security professionals what assets the organization considers to be of the highest value. For example, a web server that looks

just like any other server in the server room may be quietly generating thousands of dollars an hour in revenue. Technicians working in the server room might not recognize the server's value and instead prioritize other servers during a security incident.

Assets are valued using both tangible and intangible values. For example, consider a company that sells products via a website and generates an average of \$6,000 an hour. If the website goes down for an hour, the company loses \$6,000 in direct, tangible revenue. However, many other intangible losses aren't so apparent. Customers who visited the website while it was down may choose to shop at another company. These lost customers represent intangible losses. This includes the money expended to attract these customers and the loss of future sales to these customers. It also includes the money required to gain new customers. If a company suffers a sustained or highly public loss, it needs to expend advertising dollars to regain customer goodwill. Otherwise, even customers unaffected by the outage may choose to go elsewhere.

Of course, there's no need to waste time evaluating risk if an asset has very little value to an organization. It's relatively inexpensive to replace low-value assets if necessary, so they don't require much time or energy to protect.

Risk Visibility and Reporting

A key element of risk management is risk visibility and reporting. Ideally, an organization will be aware of the risks that threaten the organization's resources, but it does take time and effort to identify those risks. Processes such as risk assessments help identify these risks so that the risks are known and visible. Continuous monitoring systems such as antivirus software, data loss prevention (DLP) systems, firewalls, and intrusion detection systems monitor the resources and create reports or alerts when they detect problems.

Risk Register

A risk register (or risk log) is a central depository for known organizational risks and is a key method used to share threat intelligence. The risk register includes details about the known risks and can consolidate data from multiple sources. [Table 7-1](#) shows a sample table of a risk register for a web server hosting an e-commerce site. It shows sample threats along with estimated values for the probability of each of the threats occurring and the impact if it does occur. The risk score is the product of the probability and impact columns and helps security professionals prioritize the risk.



NOTE [Chapter 5](#) covers various attacks, including SQL injection and buffer overflow attacks. It also includes methods to mitigate these attacks, such as input validation and keeping systems up to date with current patches.

[Table 7-1](#) quickly shows that the greatest risk of those listed is a buffer overflow attack, with a score of 56. The table also shows what controls are in place. Because input validation is being used to prevent a SQL injection attack, the risk score is relatively low. If input validation methods weren't in place, the score would be much higher.

While [Table 7-1](#) is limited by the width of the book's page, actual risk registers don't have the same limitations. It's easy to add additional columns, such as the following:

| Name | Probability (P 1-10) | Impact (I 1-10) | Risk Score P x I | Controls | Remarks |
|------------------------|----------------------|-----------------|------------------|-------------------------|-----------------------------|
| SQL injection attack | 1 percent | 8 percent | 8 percent | Input validation | Control in place |
| Buffer overflow attack | 7 percent | 8 percent | 56 percent | Keep systems up to date | Systems not kept up to date |
| Data breach | 2 percent | 9 percent | 18 percent | Data on internal server | Control in place |

Table 7-1 Sample Risk Register

- **Description** Full descriptions of the risks so that personnel using the risk register understand them.
- **Planned controls** These are any controls that management has approved but have not been implemented yet.
- **Risk score with planned controls** Experts can predict the probability and impact of the risk with planned controls in place.
- **History of attacks** This column can provide links to an additional table that documents past attacks, along with explanations of how they were successful and what was implemented after the attacks.
- **Incident response** This provides details on the best response to the incident to mitigate losses.

An organization can create its own internal risk register, but there are also publicly available risk registers that are quite comprehensive. [Chapter 6](#) mentioned the Common Vulnerabilities and Exposures (CVE) list. This is a standardized list of known exploits, security vulnerabilities, and malware organized as a massive dictionary. It provides standardization for naming vulnerabilities and exposures. The National Vulnerability Database (NVD) includes CVE data plus much more, such as checklists, alerts, and vulnerability notes. Both the CVE and NVD are types of risk registers.

Common Vulnerability Scoring System

The Common Vulnerability Scoring System (CVSS) is an open standard that organizations can use to assess the severity of computer system security vulnerabilities. It outlines values for metrics that can then be plugged into a formula to determine the severity of any vulnerability. CVSS v3.0 was released in June 2015.

A CVSS assessment evaluates multiple areas and gives a numerical score between 0 and 10, with a score of 10 indicating the most severe vulnerability. These scores are based on several different metrics within several metric groups. The Base metric group includes both exploitability metrics and impact metrics.

- **Exploitability metrics** This group includes: an attack vector that evaluates how a vulnerability can be exploited (such as a local account, an adjacent network, or a network); an attack complexity that evaluates how easy or difficult (measured as Low or High) it would be to exploit the vulnerability; and privileges required.
- **Impact metrics** The impact is evaluated based on potential loss of confidentiality, integrity, or availability. Each area includes a value of None (no loss), Low (partial loss), or High (complete loss).

Each of the word values are preassigned numerical values as scores. For example, an attack complexity of Low is assigned a score of .777 and High is .44. Once the values are determined for a vulnerability, they are plugged into a formula to produce the relative score.

The CVSS also includes a temporal metric group and an environmental metric group. These are combined with the base metric group.

Temporal metrics evaluate the vulnerability at different times. For example, the risk is very low when an exploit is only a theory but very high when malware is in the wild and actively exploiting the vulnerability. Similarly, the risk is very high if there is no patch or other solution to mitigate the threat, but lower if a vendor has released a patch.

Environmental metrics assess the severity of the vulnerability within the environment. As a simple example, the base and impact metrics may be very high for a vulnerability on Linux systems running a specific application. However, if the only systems running the exploitable application are within isolated laboratory settings, this gives the vulnerability a very low score. In contrast, if the organization was running the application on dozens of servers accessible via the Internet, the environmental metric would be very high.

Risk Management Frameworks

A risk management framework provides a structured process that organizations can use to manage risk for IT systems. Two frameworks that are often used by organizations are NIST SP 800-37, *Guide for Applying the Risk Management Framework to Federal Information Systems*, and ISO 31000:2018, *Risk Management – Guidelines*, published by the International

Organization for Standardization. NIST SP documents are published by the U.S. government and freely available to anyone. Documents published by ISO can be purchased from the ISO website (<https://www.iso.org>).

Both frameworks refer to the process of identifying, evaluating, and selecting security controls. NIST SP 800-53, *Security and Privacy Controls for Federal Information Systems and Organizations*, provides a comprehensive listing of various security controls.



NOTE ISO may look like a typo for the International Organization for Standardization. However, it isn't an acronym but instead an abbreviation of the word isos, which means equal.

NIST SP 800-37

“Risk Management Framework (RMF)” (as identified in 800-37) provides a comprehensive six-step approach for managing risk within federal information systems. While it is focused on U.S. federal systems, any organization can apply the framework based on its needs.

[Figure 7-5](#) and the following list show the six steps discussed in SP 800-37. Each of these steps is applied to a single information system. Additionally, it is an ongoing process.

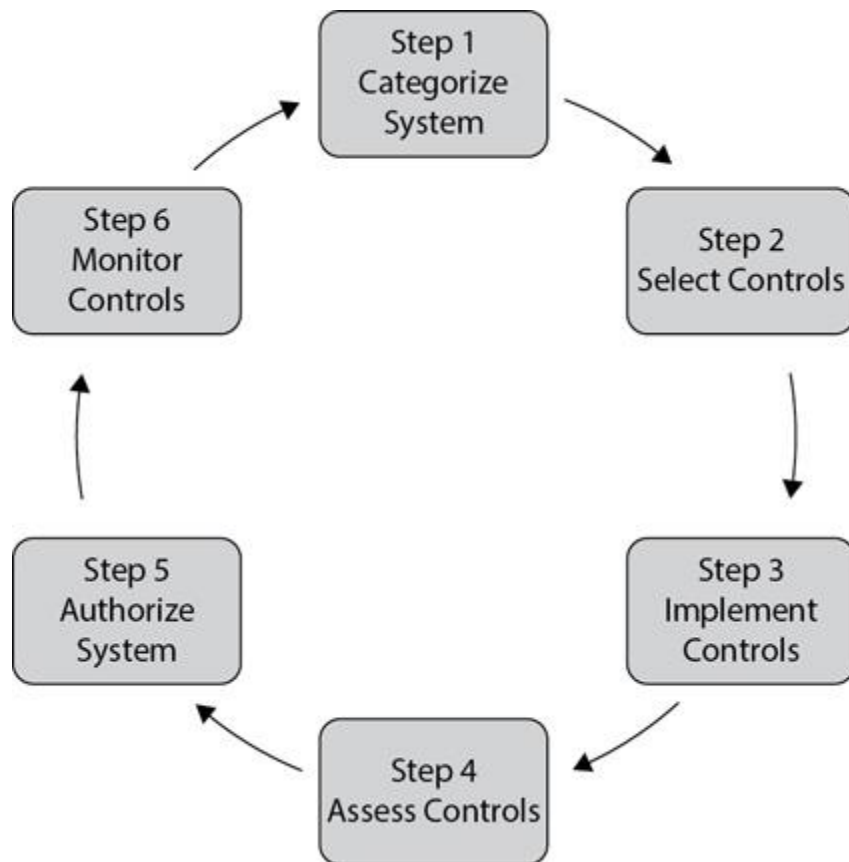


Figure 7-5 Risk Management Framework steps from SP 800-37

- **Step 1: Categorize the information system** Personnel within the organization evaluate the information system, along with the data that it processes, stores, and transmits. While an organization can use any internally defined categories, they can be as simple as Unclassified, Confidential, Secret, and Top Secret, based on the data that the system processes, stores, and transmits. The information system owner is responsible for this step.
- **Step 2: Select security controls** Next, personnel identify and evaluate common security controls that typically apply to the identified category. They then select appropriate security controls and develop a continuous monitoring strategy to assess the effectiveness of the controls. Last, they document and obtain approval of the security plan.
- **Step 3: Implement security controls** In this step, personnel implement the security controls. They also update the security plan by documenting the security control implementation. This typically includes a functional description of the control, along with planned inputs, expected behavior, and expected outputs.
- **Step 4: Assess security controls** In this step, personnel assess the effectiveness of the security controls based on the procedures identified in the security plan. Next, they create a security assessment report documenting their findings. If the security controls are not performing

as expected, personnel identify remedial actions and add recommendations to the assessment report.

- **Step 5: Authorize the information system** Next, personnel create a plan of action and milestones based on the security assessment report. They use this to submit a security authorization package to the system owner for approval. System owners then determine the risk that remains after the controls are implemented. If system owners determine that the risk is acceptable, they can choose to authorize the system and accept the remaining risk.

- **Step 6: Monitor security controls** In this step, personnel evaluate system and environment changes using standard change management processes. They perform ongoing security control assessments and make recommendations for remediation actions as appropriate. Any changes or updates are documented in the security plan, the security assessment report, and/or the plan of action and milestones, as appropriate.

This entire process can be repeated as necessary. For example, if a system will be processing different types of data, it may need to be recategorized. Similarly, if different security controls are implemented, they need to be evaluated, assessed, and monitored.

ISO 31000

ISO 31000:2018 includes a principles clause and a process clause that should be followed during the risk management process. The principles clause includes principles such as continual improvement, integrated, structured and comprehensive, and more. The process clause includes processes for risk assessments, risk treatment, and more.

The framework clause identifies the following five phases, similar to the six steps in the SP 800-37 framework. Just as with the SP 800-37 framework, these steps can be repeated as necessary.

- **Design** Security is built into the process from beginning to end, starting at the design phase. During this phase, security controls are identified and selected.
- **Implementation** Personnel implement the design of the system during this phase, including the selected security controls.
- **Evaluation** Security controls are evaluated during this phase to determine their effectiveness based on the goals set during the design phase.
- **Improvement** Based on findings during the evaluation phase, security personnel make recommendations for improvement.
- **Integration** Security personnel integrate approved recommendations into the system.

Performing Risk Assessments

A risk assessment is a point-in-time evaluation. It looks at the current situation and attempts to determine what risks exist and how to address them. As an example, you can perform a risk assessment to determine risks for e-mail servers in your organization. This assessment could examine current threats, system configuration, interconnectivity, the amount of e-mail processed, fault-tolerance capabilities, power and air conditioning requirements, and more. However, after the assessment is completed, any of these elements can change, affecting the results of the assessment.

One of the benefits of performing risk assessments is that it allows security personnel to identify and prioritize risks. This helps ensure the organization dedicates risk management resources to the most serious risks. A risk assessment will use either a quantitative analysis or a qualitative analysis.



EXAM TIP A risk assessment looks at risks at a specific time. In contrast, risk management is an ongoing practice. Organizations repeat risk assessments periodically to ensure that current risks are assessed.

Threat Modeling

Threat modeling is a type of risk assessment that attempts to predict an attack. It uses various tools to model aspects of an attack against a system, a network, an application, or data.

NIST SP 800-154, *Guide to Data-Centric System Threat Modeling*, provides a methodology that organizations can use as a guideline for threat modeling. It focuses on data-centric threats and protecting different types of data within systems. However, many of the methodologies discussed in the document can be applied to other types of threat modeling, such as focusing on threats against web servers.

Using SP 800-154 as a guide, threat modeling can be defined within four major steps.

- **Step 1: Identify and characterize the system and data of interest** A primary goal in this step is to characterize the system. This includes evaluating the data processed by the system, and how data will move between systems to other locations. Security objectives (such as confidentiality, integrity, and availability) are evaluated to determine if any one security objective is more important than another.
- **Step 2: Identify and select the attack vectors to be included in the model** Next, personnel identify and select potential attack vectors to be included in the model.



TIP An *attack vector* refers to a method an attack can use to exploit a vulnerability. SP 800-154 mentions that attack vectors typically include a source of malicious content, a potentially vulnerable processor of the malicious content, and the malicious content itself. As an example, imagine a user receives a malicious attachment sent via e-mail. The malicious attachment is the content in the attack vector, the e-mail client that allows the user to open it is the processor, and the e-mail server that sent the attachment is the source. In some situations, security professionals can evaluate all potential attack vectors against a system or, said another way, analyze the system's *attack surface*.

- **Step 3: Characterize the security controls for mitigating the attack vectors** In this step, personnel identify and document security controls that are currently in place to reduce the attack vectors and evaluate potential alterations or additions to security controls. For each recommendation, they attempt to identify the potential effectiveness of each alteration or addition. This can be simplified to assigning scores of low, medium, or high to each recommendation.

- **Step 4: Analyze the threat model** In the last step, security professionals analyze the threat model in an attempt to evaluate the effectiveness and efficiency of each security control. This can be a complex task and typically requires the use of quantitative analysis or qualitative analysis methods.

Quantitative Analysis

A quantitative analysis uses numerical-based data such as monetary figures to identify the actual cost associated with a risk. There are several variables associated with a quantitative analysis. If you can identify the values of these variables, prioritizing different risks becomes a simple matter of plugging the values into a formula. These values also help to justify costs for controls and countermeasures.

The following terms are important to understand when using a quantitative analysis:

- **Single loss expectancy (SLE)** This value identifies the actual loss of a single occurrence. For example, if a revenue-generating website is down for an hour, the company could determine that it results in a loss of \$10,000.

- **Exposure factor (EF)** The EF identifies the magnitude of loss as a percentage of an asset's value. As an example, imagine that a building has a value of \$500,000 and insurance covers 75 percent of the replacement cost. The EF is 25 percent (or \$125,000). This is often calculated automatically in the SLE, but you may run across it as a separate value that you need to calculate as: $SLE = SLE \text{ without EF} \times EF$.



TIP The SLE would include both tangible and intangible costs. Tangible costs would include the actual lost revenue while the website is down. Intangible costs include the costs associated with lost customers or advertising to rebuild the company image.

- **Annual rate of occurrence (ARO)** This identifies how often the SLE is expected to occur in a given year based on the current environment. You can use historical data to determine the ARO. For example, if the website experienced three one-hour outages each year in the past two years, it will probably experience three one-hour outages this year if something isn't changed.
- **Annual loss expectancy (ALE)** You calculate the ALE by multiplying $SLE \times ARO$. For example, if the SLE is \$10,000 and the ARO is three, the ALE is $\$10,000 \times 3$, or \$30,000. When the EF is mentioned but not included in the SLE, you would calculate the ALE with the following formula: $ALE = (SLE \times EF) \times ARO$.
- **Cost of control** A control will mitigate or reduce a risk. For example, if faulty hard drives caused an outage, you can replace them with a hardware RAID-6. If the RAID-6 costs \$5,000, then the cost of the control is \$5,000. Some controls have ongoing costs to consider. For example, administrators might need additional training to install and manage some systems, such as dedicated firewalls or unified threat management devices.

You can compare the ALE with the cost of the control to determine if it makes fiscal sense to purchase the control. In general, if the cost of the control is less than the ALE, the cost is justified. If the cost of the control exceeds the ALE, it may not be justified. If the cost of the control is about the same as the ALE, then the organization can perform a detailed analysis to determine the actual return on investment (ROI) for the control.

Consider the example of website outages caused by hard drive failures. Experts might predict that adding the hardware RAID-6 will eliminate these outages. You can then use the following information to see if it makes sense to purchase the hardware RAID-6:

- ALE before control: \$30,000
- ALE after control: \$0
- Cost of control: \$5,000
- Savings (ALE before control – ALE after control – cost of control): \$25,000

If the company does nothing, it can expect to lose \$30,000. If the company implements the control, it eliminates the outages related to hard drives and reduces the ALE to \$0. In other words, if the company spends \$5,000, then it can expect to save \$30,000, for a net savings of \$25,000. The cost of the control is justified.

It's also possible that by implementing a control, you don't expect to eliminate the outages, but instead you expect to reduce the ARO. Using the previous example, experts might predict that adding a hardware RAID-6 will reduce the ARO from three to one. Instead of failing three times, the website might fail once due to a hard drive failure. An ARO of one results in an ALE of \$10,000 ($1 \times \$10,000$). Is the cost of the RAID-6 still justified?

- ALE before control: \$30,000
- ALE after control: \$10,000
- Cost of control: \$5,000
- Savings (ALE before control – ALE after control – cost of control): \$15,000

You've reduced the ARO from three to one, so you've reduced the ALE from \$30,000 to \$10,000. This provides an initial savings of \$20,000. The cost of the control is \$5,000, so this provides a net savings of \$15,000. If you did nothing, you'd lose \$30,000, so this cost is still justified even if it doesn't eliminate the risk.

On the other hand, if another risk has an ALE of \$2,000 and the cost of the control is \$22,000, the control is not justified. In this example, the control will reduce the ALE to \$0:

- ALE before control: \$2,000
- ALE after control: \$0
- Cost of control: \$22,000
- Savings (ALE before control – ALE after control – cost of control): negative \$20,000

You're spending \$22,000 on a control to prevent a \$2,000 loss. In this situation, management would typically decide not to implement the control, but instead accept the risk.



EXAM TIP When preparing for the SSCP exam, it's important to have a solid understanding of SLE, ARO, and ALE. You should also be able to evaluate the cost of a control to determine whether it is justified.

An advantage of the quantitative analysis is that it provides a clear metric in the form of a monetary figure. You can use this to analyze the actual value of the control. However, the numbers used to create a quantitative analysis are not always easy to obtain. This is especially true when the SLE includes intangibles.

Qualitative Analysis

A qualitative analysis is subjective and often simply categorizes a risk using words such as low, medium, and high. For example, [Figure 7-6](#) shows the relationship between impact and likelihood ratings and how to use them to determine overall risk. Notice that there's a direct relationship between likelihood and impact. As the likelihood and impact increase, the risk also increases.

In [Figure 7-6](#), you can see that a low impact combined with a high likelihood results in a medium risk. It's important to realize this is just a guide and that the security professional doing the assessment will likely add in some subjective reasoning. As mentioned previously, a fire has a low likelihood but a high impact. However, it may be classified as an overall high risk because the result of a fire can be catastrophic.

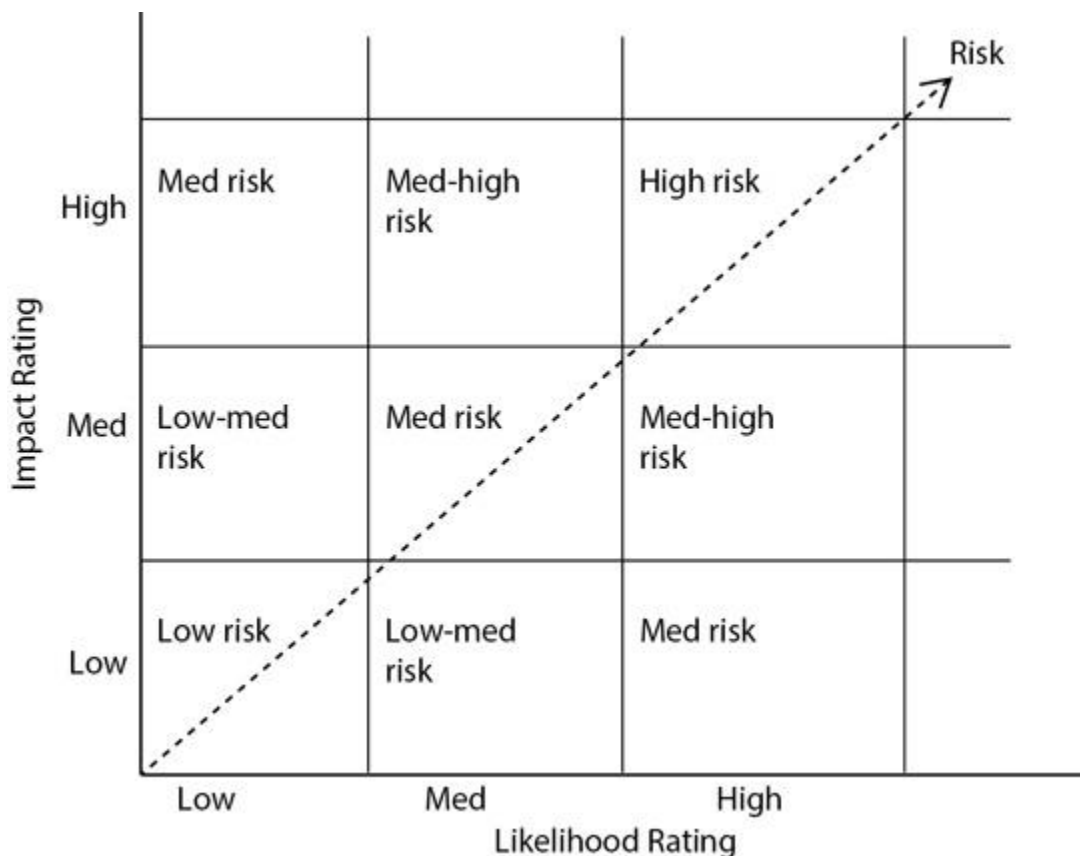


Figure 7-6 Qualitative analysis based on impact and likelihood ratings

It's also possible to assign values to the ratings. This allows you to prioritize the risks based on the result. For example, imagine that you've assigned a value of 1 for low, 5 for medium, and 10 for high. You can then use these values to determine a risk score by multiplying the likelihood rating by the impact rating. [Table 7-2](#) shows a risk level matrix that combines these ratings and shows the resulting risk scores. Notice that data such as this can be included in a risk register (as shown in [Table 7-1](#) earlier).

| | Low Impact (1) | Medium Impact (5) | High Impact (10) |
|-----------------------|--------------------|--------------------|----------------------|
| Low Likelihood (1) | $1 \times 1 = 1$ | $1 \times 5 = 5$ | $1 \times 10 = 10$ |
| Medium Likelihood (5) | $5 \times 1 = 5$ | $5 \times 5 = 25$ | $5 \times 10 = 50$ |
| High Likelihood (10) | $10 \times 1 = 10$ | $10 \times 5 = 50$ | $10 \times 10 = 100$ |

Table 7-2 Risk Matrix Using Qualitative Analysis



TIP An organization can use any words or values desired for the ratings. For example, instead of just using low, medium, and high, you can add in low-medium with a value of 2.5 and medium-high with a value of 7.5.

[Table 7-3](#) shows how these ratings can be used for specific threat events. It shows three adversarial threat events with risk scores of 5, 50, and 100. The threat event with a score of 100 (DoS attack on the Internet web server) represents the highest risk and needs to be addressed first.

| Threat Event | Likelihood | Impact | Risk Score |
|---------------------------------------|------------|------------|------------------|
| DoS attack on the intranet web server | Low (1) | Medium (5) | Low-medium (5) |
| Malware infection | Medium (5) | High (10) | Medium-high (50) |
| DoS attack on the Internet web server | High (10) | High (10) | High (100) |

Table 7-3 Risk Matrix Ratings Used for Adversarial Threat Events



NOTE [Chapter 5](#) covers various types of attacks, including denial of service (DoS) attacks. A DoS attack attempts to reduce a system's ability to provide a service. In [Table 7-2](#), DoS attacks on the web servers would hamper the web servers' ability to serve web pages.

An advantage of the qualitative analysis is that it provides a clear priority of the different risks. Simply by looking at the result of the analysis, it becomes a simple matter to determine what risks should be addressed first. A disadvantage is that it is difficult to use the results in a cost/benefit analysis when evaluating controls.



EXAM TIP A quantitative analysis uses numerical data such as monetary figures and is objective. A qualitative analysis uses scenario-based data and is subjective. Many organizations use both methods for different risk assessments.

Risk Assessment Steps

Risk assessments can be complex projects taking a lot of time and effort. However, following an outline of predefined steps can simplify them. This also helps ensure that the risk assessment is reproducible and repeatable, and that different security experts can create the same results from the same data. This also helps an organization identify trends between assessments.

With this in mind, it's important to look for previous risk assessments before starting a risk assessment. If they exist on the same system you're evaluating, you can use them as a template; if you're redoing a risk assessment, you only need to identify the changes since the last risk assessment.

An organization isn't required to follow any specific steps. However, NIST SP 800-30 provides four steps that any organization can follow to create reproducible and repeatable results. [Figure 7-7](#) shows the overall steps associated with a risk assessment as recommended in NIST SP 800-30, along with tasks within these steps. The following sections provide additional details on these steps.

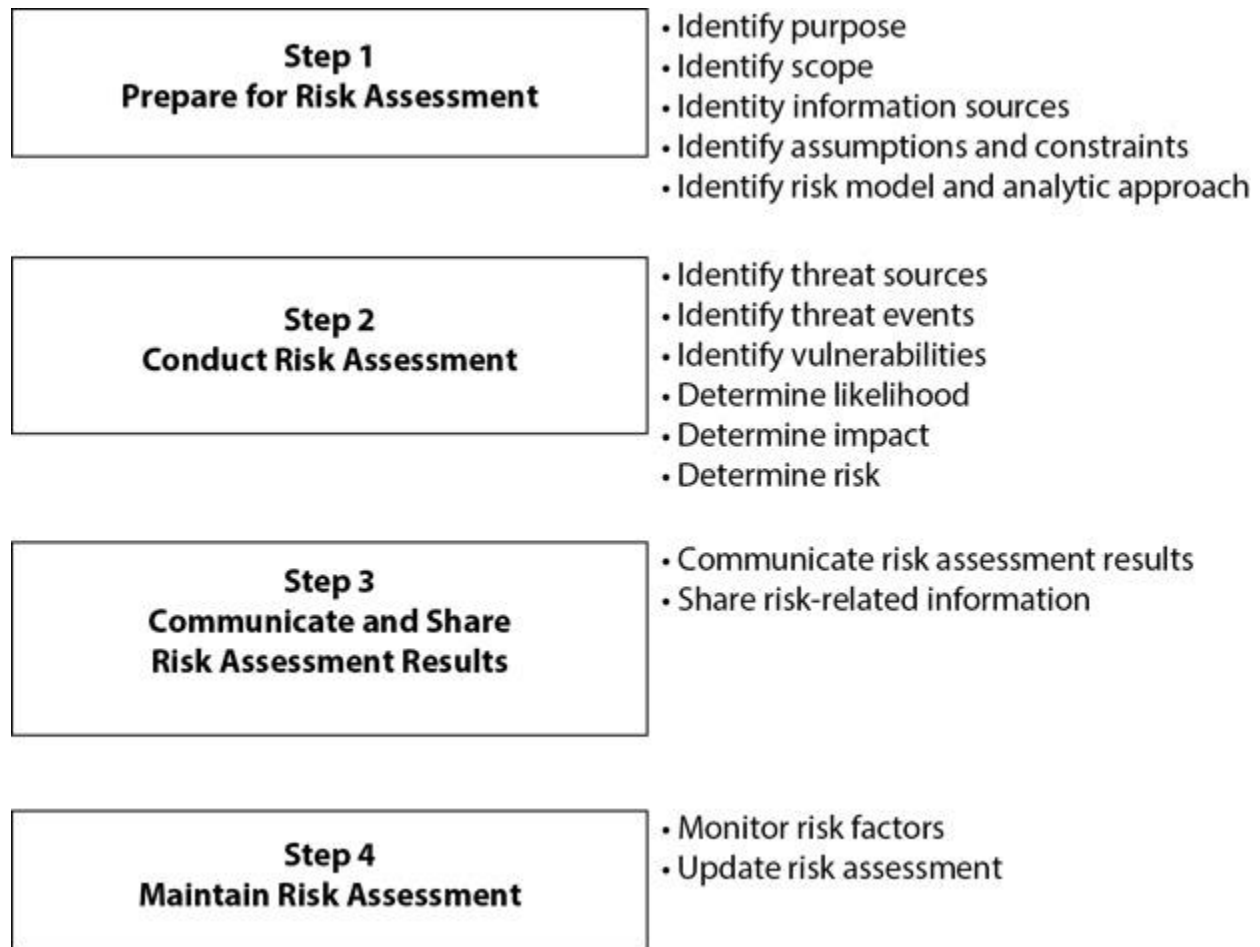


Figure 7-7 Risk assessment steps including tasks within each step

Step 1: Prepare for Risk Assessment

In the first step, the organization performs key tasks to prepare for the risk assessment by identifying several elements used during the assessment. This includes identifying the purpose and scope of the assessment and identifying sources of information for the assessment. Any large project is susceptible to scope creep, where the project slowly grows well beyond the original goals. As an example, if the goal of the assessment is to assess risks associated with e-mail, the study can easily extend into the network infrastructure, including routers and firewalls. It could include power and air conditioning supplied to the entire server room. It could also include malware protection on individual clients and much more. If the goal is to examine only the direct risks associated with the e-mail servers, the purpose and scope should identify this.

This step also includes identifying assumptions and constraints. An organization will typically have safeguards or controls in place designed to mitigate risks. Additionally, some controls may be in the purchasing pipeline but not implemented yet. A risk assessment will often include a list of both current and planned controls, along with clarifying that the risk assessment assumes all of these controls are in place.

Last, this step includes identifying the risk model and analytic approach for the risk assessment. As an example, some risk models refer only to threats, while other models refer to both threat sources and threat events. The two types of analysis are quantitative and qualitative, both of which were covered in the previous two sections.

Step 2: Conduct Risk Assessment

During the risk assessment, you identify threat sources, threat events, and vulnerabilities. You then perform an analysis to determine the likelihood of a risk occurring, the impact if it does occur, and the overall risk.

Threats can be internal or external, natural or manmade, and intentional or accidental. Some models suggest matching threats with vulnerabilities, but this can often be tedious and doesn't always provide any additional actionable data. Many current models (including the one outlined in NIST SP 800-30) do not require matching threat/vulnerability pairs.

After identifying likely threats and vulnerabilities, the assessment then attempts to identify the likelihood of a threat and the impact if it occurs. Security professionals rarely conduct risk assessment themselves. Instead, they often use surveys, questionnaires, and interviews with subject matter experts to gather information. For example, a survey may ask respondents to rank the likelihood of a threat event occurring and the impact if it does occur.



NOTE The specific numbers or words used to identify the likelihood and impact of a risk may be different from one organization to another. A specific scale isn't available that all organizations must use. However, as long as the assessment uses the scale consistently, assessors can choose any scale that works for them.

After identifying the likelihood and impact, it's possible to assign a risk score by combining the likelihood with the impact. A low likelihood combined with a low impact results in a low risk score, while a high likelihood combined with a high impact results in a high score.

The assessment can also assign numbers to the word values. For example, low, medium, and high could be 1, 5, and 10, respectively. It's then possible to calculate the risk score by multiplying the likelihood and the impact. For example, if the probability is 7 and the impact is 8, it provides a risk score of 56.

It's also possible to identify the risk score subjectively. For example, the possibility of a fire is low, while the impact would be high. Multiplying 1 (low) times 10 (high) gives a risk score of 10. However, the assessor can override this relatively low risk score because the impact could be catastrophic.

While a risk assessment typically does not recommend controls, it often evaluates prospective controls or safeguards to determine their value. As mentioned previously, a quantitative analysis makes it easier to evaluate the actual value of a control. However, you can also use a qualitative analysis to evaluate the control. The key is to calculate the risk score without the new control and with the new control.

Step 3: Communicate and Share Risk Assessment Results

After conducting the risk assessment, assessors provide the results to decision makers and other personnel within the organization. Risk assessment results include sensitive information, so the results aren't shared with everyone within the organization. However, if the results directly affect a department, it's common to share the results with department stakeholders.

While this step is typically done after completing the risk assessment, in some cases assessors will provide the results earlier. For example, if the assessment discovers a serious risk, assessors might report this to senior management right away.

The method of communicating and sharing the results will vary from one organization to another. One organization might use informal briefings before finalizing the risk assessment report. Another organization might decide to skip the briefings until the risk assessment report is finalized.

The contents of a risk assessment might also vary from one organization to another. However, you'll often see the following items:

- **Executive summary** This is typically about 5 to 10 percent the size of the full report and summarizes the risk assessment.
- **Risk assessment preparation** This section identifies the purpose, scope, and assumptions identified in step 1 of the assessment. It also documents the information sources, the risk model, and analysis methods used for the assessment.
- **Threat sources, threat events, and vulnerabilities** Depending on the needs of the organization, the assessment might focus on only adversarial threats or all possible threats.
- **Risk scores** The assessment includes the likelihood of the risks, the potential impact, and the risk scores. It's appropriate to document this information in tables or graphs so that it is easy to digest.
- **Controls or countermeasures** If the assessment evaluated controls or countermeasures, it should include details on them. Whenever possible, this section should document the cost of the controls and the predicted impact before and after implementing the controls.
- **Uncertainties** A risk assessment attempts to predict future risks, so it isn't always based on complete data. Instead, it relies on expert opinions, assumptions, and speculations. An uncertainty analysis documents uncertainties caused by missing data, subjective opinions, and

assumptions. This helps decision makers evaluate the usefulness of the risk assessment as events change. As an example, advanced persistent threats (APTs) from foreign countries have launched many attacks on companies within the United States. Even though there is some uncertainty about who sponsors specific APTs and their intentions, it is possible to predict future risks associated with APTs.

- **Appendixes** When appropriate, the assessment includes appendixes to document additional information such as details to support any of the conclusions.

Step 4: Maintain Risk Assessment

A risk assessment is a point-in-time report. In other words, it evaluates threats and vulnerabilities at a specific time. However, threats are constantly changing, so an organization must attempt to keep up with the changes with effective risk-management and risk-monitoring decisions.

Organizations periodically repeat or update risk assessments for key systems. The frequency is variable, but it is common to repeat a risk assessment on an annual or semiannual basis.

Maintaining a risk assessment also includes monitoring risk factors identified in the risk assessment. If risk factors change, the organization might choose to update the risk assessment outside of the normal schedule. It certainly isn't necessary to wait for an incident before reassessing risk.

Address Findings

An important step after completing a risk assessment is to address the findings. In many cases, the assessment will identify risks that need to be mitigated by implementing controls. Of course, management makes the decision about which risks to address. In general, a risk assessment does not recommend controls. However, it can evaluate recommended controls and assess their value.

One process that some organizations implement is to periodically review the findings of past assessments and audit the progress. Auditing the findings later (such as six months after the completion of an assessment) provides insight into the accuracy of the predictions. For example, an assessment might have predicted that implementing a specific security control could eliminate a risk. Six months later, you can determine if this was true or not. This step also helps ensure that the organization is implementing the controls approved by management.

Supporting the Incident Lifecycle

In the best of all worlds, your risk management practices will prevent any incidents. However, avoiding all incidents is highly unlikely, so organizations create plans to respond to incidents.

In the context of IT security, a security incident is any violation of policies or security practices that has the potential to result in an adverse event. NIST SP 800-61, *Computer Security Incident Handling Guide*, provides several definitions that are helpful in identifying incidents:

- *An event is any observable occurrence in a system or network.*
- *Adverse events are events with a negative consequence, such as system crashes, packet floods, unauthorized use of system privileges, unauthorized access to sensitive data, and execution of malware that destroys data.*
- *A computer security incident is a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices.*

Notice that not all events are security incidents. Administrators or security personnel investigate events to determine whether they are incidents. Additionally, not all computer security incidents necessarily result in adverse events. For example, imagine that the company security policy states that users should not use computers to do personal shopping on the Internet. If a user violates this policy, it is an incident, but it won't result in negative consequences (except perhaps for the employee).

The "Identifying Threat Sources" and "Identifying Threat Events" sections earlier in this chapter described many different types of incidents. Additionally, the following list describes some common types of incidents:

- **Denial of service (DoS)** Any type of attack that attempts to reduce a system's ability to provide a service is an incident. This includes both DoS and distributed DoS (DDoS) attacks. Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) attempt to detect and block DoS and DDoS attacks before they can cause any damage.
- **Malware** The discovery of any type of malicious software, such as a virus, worm, Trojan horse, and so on, is an incident. If AV software discovers the malware before it's installed on a system, it's a minor incident. However, if the malware is able to spread unabated through a network, it can quickly become a major incident.
- **Inappropriate usage** Most organizations have an acceptable use policy identifying what employees can do with IT systems. Any attempt to violate this policy is an incident.
- **Unauthorized access** If any individual or software program is able to access data or systems that the user or program is not authorized to access, it is an incident. Depending on the value of the data or system accessed, the incident may be a minor incident or a major incident.

As you study for various cyber security certifications, you'll find that the terminology used to define incident response varies. However, the underlying concepts are the same. As an example, [Table 7-4](#) compares the phases of the incident response lifecycle outlined in the SSCP objectives and the phases outlined in SP 800-61.

| SSCP Objectives | SP 800-61 |
|--|--|
| Preparation | Preparation |
| Detection, analysis, and escalation | Detection and analysis |
| Containment | Containment, eradication, and recovery |
| Eradication | |
| Recovery | |
| Lessons learned/implementation of new countermeasure | Post-incident activity |

Table 7-4 Incident Response Lifecycle Comparison

While the SSCP objectives list six phases and SP 800-61 lists four phases, there really isn't much difference between the two. The key is to understand the concepts within each of the phases. The following sections describe the phases listed in the SSCP objectives.



EXAM TIP The first phase in incident response is preparation, which includes creating an incident response plan. After detecting and verifying an incident, it's important to contain the incident as quickly as possible.

Preparation

Preparing for an incident is an extremely important first phase. With solid preparation, you have a much better chance of minimizing damage from an incident when it occurs. Preparation includes the steps to prevent an incident by ensuring that your systems are secure, but it also includes planning for incidents and planning the response.

Many organizations identify individuals in a computer incident response team (CIRT). These individuals are responsible for responding to an incident when it occurs. A CIRT typically includes members from different parts of the organization and with different skills, such as senior management, information technology, information security, legal, and personnel with public communications skills.



TIP Some organizations use different terms for their teams. CIRT is common, but you may see other terms such as computer security incident response team (CSIRT) or just incident response team (IRT). Each of the terms refers to the same group of people who respond to an incident.

The CIRT needs ready access to tools and resources when an incident occurs. Every minute they have to spend looking for these tools is a minute lost containing the incident, so it's common to locate the tools in an easily accessible location known by all the team members. The following list identifies some of these tools and resources:

- **Contact information** This includes contact information for all CIRT members and key personnel in the organization who may need to be notified of an ongoing incident.
- **Reporting forms** The organization may want to ensure that specific information is recorded about an incident. Precreated forms help the team members remember what information to document.
- **Forensic tools** Computer forensic tools help personnel collect and analyze evidence while observing relevant requirements to maintain the quality of the evidence. In case the incident results in legal proceedings, it's important for personnel to ensure they don't modify the data as they collect it and that they maintain proper custody of evidence after collecting it. [Chapter 13](#) covers forensic issues in more depth.
- **Documentation** This includes documentation on systems and the network infrastructure. It should include a ready reference of approved changes and the status of these changes. When team members identify something that looks suspicious, such as a system configuration, they need to be able to identify whether it is a valid change or a malicious change. For example, administrators typically disable Trivial File Transfer Protocol (TFTP) when hardening systems, unless it's needed. Some malware enables TFTP to forward files to other systems. If team members see TFTP enabled on a system, this documentation should allow them to verify if it is needed.
- **Software and hardware** The team may need access to hardware to perform simple tasks such as creating reports, doing research, or performing analysis. They also may need access to software used to rebuild systems.
- **War room** in the case of a serious incident, the team may need to use a central war room for coordination. This will likely be something like a general-purpose conference room used only as a war room during an incident. The incident response plan might designate a room to be used for this purpose if necessary.

This is by no means a complete list, and each organization's list will likely be different. However, the important thing to remember is that the CIRT will need ready access to tools and resources. Many of these items could be stored in a network location. However, an incident might take the network down. As an alternative, some organizations create a "crash cart" for incidents. It would include a dedicated computer with relevant forensic tools installed, additional hardware and software that may be needed, and copies of relevant documentation.

Detection, Analysis, and Escalation

The next phase of the incident response lifecycle is detection, analysis, and escalation. At this stage, the event is discovered, investigated, and analyzed. If analysis shows that the event is a security incident, personnel escalate it. Escalation typically involves reporting it to other personnel in the organization with expertise to contain it.

There are multiple methods of detecting an incident. It could be as simple as AV software detecting malware on a USB device and sending an alert to notify a user. Users could notice suspicious activity and report it. Systems may randomly lock up or crash. Administrators may see logs growing at an alarming rate, strange files appearing on servers, or simply an increase in network activity (and a slowdown in network response). After detecting these events, someone investigates them to determine whether an incident has occurred.

Intrusion detection systems (IDSs) provide automated detection of potential incidents and display alerts to IT personnel about them. It's important to realize that even though an IDS alerts personnel about activity, that activity isn't necessarily an incident. Consider a SYN flood attack. In day-to-day operations, network issues may prevent the third packet from reaching the server. If this happens once in an hour, it's highly unlikely that it's an attack. Then again, if a server is receiving 100 half-open connections a second, that is very likely an attack.



NOTE [Chapter 5](#) covers SYN flood attacks. As a reminder, they disrupt the TCP handshake process by withholding the third packet to complete the connection. The attacker sends the first packet (the SYN packet), the attacked system responds with a SYN/ACK packet, but the attacker never sends the final ACK packet. This consumes resources on the server and can prevent legitimate users from connecting to the server.

Security professionals who manage an IDS have to define the number of half-open connections that constitutes an attack. Clearly, if this happens once an hour, it isn't an attack. However, if it is happening 100 times a second, it is very likely an attack. They identify the number and set it as a threshold to indicate an event. If they set the threshold too high, they won't be notified when an attack is under way. If they set the threshold too low, they will receive false positives—notifications of possible attacks that are actually not attacks. Given the choice between not being notified of an actual attack or occasionally getting some false positives, most security professionals prefer the latter. They accept some false positives with the goal of ensuring the IDS does notify them of actual attacks.

With this in mind, personnel must investigate and analyze each potential incident. This analysis determines whether the event is an incident and, if so, attempts to prioritize and escalate the incident. Minor incidents don't affect critical systems or critical infrastructure, so they are escalated only within the IT department. Critical incidents affect mission-critical systems and can seriously degrade the organization's ability to perform its primary mission. Obviously, a

critical incident would mandate the use of all available resources to address the incident and escalation may result in recalling all members of an incident response team.

Containment

The next phase of the incident lifecycle is containment. Once personnel verify that an event is an incident, the incident must be contained as quickly as possible. The goal is to prevent the incident from spreading. For example, after malware infects a single system on a network, it typically tries to locate and infect other systems on the network. However, disconnecting the network cable on the infected system contains the malware to just that system. It cannot spread on the network. In contrast, if you didn't contain the malware, it could quickly spread to all the computers on the network, causing a major incident.



EXAM TIP Once an incident has been identified, it should be contained as soon as possible. This can't occur before detecting the incident, but you can contain an incident after you've verified that it is an actual incident and not a false positive.

The steps for containment might be different in different organizations. As an example, some malware can detect when the network cable has been removed or the network interface card (NIC) has been disabled. When detected, the malware can take steps to hide itself or even delete itself from the infected system. Some organizations may instead choose to disconnect the switch from any routers, effectively isolating the network. Security professionals can then monitor the network for any suspicious activity from the infected computer. Of course the risk here is that all other computers on the isolated network may be infected.

Eradication

Once you've contained the incident, you move to the eradication phase. For example, some multipartite viruses have multiple components, and you must remove each individual component. If you miss any component, the malware can reemerge and reinfect all the files you've cleaned.

Recovery

The next phase is recovery, and what it entails depends on the damage caused by the incident. For minor incidents that haven't caused any damage, recovery may not require any additional steps at all. As an example, if antivirus software detected malware and immediately quarantined it, no other action is required to recover a system.

However, a major incident may require the recovery of an entire server or even an entire location. For example, if a fire destroyed a building, recovery entails activating an alternative

location, moving critical systems and data to the alternative location, and bringing everything back online.



NOTE [Chapter 12](#) covers disaster recovery operations within the context of disaster recovery plans. It also covers alternative locations such as hot sites, cold sites, warm sites, and mobile sites. Each of these sites provides varying levels of readiness for disaster recovery.

Some malware can reconfigure systems, weaken security by enabling services or protocols, create new accounts used for backdoor access, or open previously closed ports on firewalls. All of the actions taken by the malware need to be reversed to complete the recovery.

Lessons Learned/Implementation of New Countermeasure

The last phase of incident response is to complete a lessons learned review and implement new countermeasures, if appropriate. This is often referred to as post-incident activity. In this phase, you examine the incident and the response. The goal is to determine whether the existing plans and procedures adequately addressed the incident or if there are response procedures that need improvement. Organizations that take the time to examine their response typically find ways to improve their existing processes. Occasionally, personnel recognize a need to implement a new countermeasure or security control.

One of the ways to perform a lessons learned activity is by holding a meeting with people involved in the incident and posing questions to generate discussions. Obviously, you wouldn't hold a meeting like this for a minor incident, but if an incident negatively affected the mission of the organization, a lessons learned meeting is warranted. The information gathered during this meeting is compiled into a post-incident report that can be used in future risk management activities such as risk assessments.

The following are some of the questions you can include in this meeting:

- Are there any lessons to be learned from this incident? What are they?
- Did we respond to the incident as quickly as possible? If not, what can we do to improve the response time for a subsequent incident?
- Was there anything that prevented the recovery from occurring in a timely manner?
- Was the documentation up to date? If not, what was not up to date? Is there a way that we can ensure it is up to date in the future?
- Did the CIRT have all the tools it needed? If not, what tools were not available? How can we ensure these tools are available if we have a similar incident?

- Are there any new countermeasures or security controls we can implement either to prevent a similar incident or to at least reduce the severity of the incident?

One of the challenges of performing a lessons learned meeting is to avoid finger pointing and “blame storming.” Instead of focusing on improving the process for a potential event in the future, some participants may focus on ensuring that they aren’t blamed for this event. If senior management stresses that the goal is to improve the process and not to place blame, participants may be more willing to provide constructive input. Of course, if management uses the post-incident report to hold individuals or departments accountable, it’s highly unlikely future lessons learned meetings will have any success.

Chapter Review

Risk is the probability or likelihood that a threat will exploit a vulnerability and cause a loss. It is sometimes represented in the formula of $\text{Risk} = \text{Threat} \times \text{Vulnerability}$. A threat is any activity that can be a possible danger and threat events include any type of activity or event that can result in a loss of confidentiality, integrity, or availability of an organization’s assets. A vulnerability is a weakness, and a loss represents a negative event for an organization.

Organizations attempt to identify potential threat sources and threat events. They then attempt to identify vulnerabilities that threats can exploit, and assess the potential impact of a threat exposing a vulnerability.

NIST publishes several documents related to IT security. Some documents related to risk mentioned in this chapter are SP 800-30 (*Guide for Conducting Risk Assessments*), SP 800-61 (*Computer Security Incident Handling Guide*), SP 800-150 (*Guide to Cyber Threat Information Sharing*), and SP 800-154 (*Guide to Data-Centric System Threat Modeling*).

Risk management (sometimes referred to as risk treatment) includes several elements used to reduce risk to a manageable level. This includes methods to mitigate (or reduce) risks, avoid risks, share or transfer risk, accept risk, and recast risk. Residual risk is the risk that remains after mitigating risks. Senior management is responsible for any losses resulting from risk they choose to accept, including residual risk.

One of the first steps in risk management is to identify assets. Once an organization identifies the assets that it considers valuable, it then takes steps to protect them. An organization addresses high risks by implementing security controls. Security controls typically attempt to reduce vulnerabilities or reduce the impact of a risk. However, if the valuable assets haven’t been identified, it’s possible that they won’t be protected and that resources will be expended to protect less valuable assets.

The Common Vulnerability Scoring System (CVSS) is an open standard that organizations can use to assess the severity of computer system security vulnerabilities. It has several precomputed formulas that security personnel can use to determine a risk score. Security personnel evaluate their system and assign metric values to the variables in the precomputed formulas.

A risk management framework provides a structured process that personnel can use to manage risk. NIST SP 800-37, *Guide for Applying the Risk Management Framework to Federal Information Systems*, includes the steps for one framework. ISO 31000, *Risk Management – Guidelines*, provides a slightly different framework.

A risk assessment is a point-in-time evaluation. It examines a system to identify threats and vulnerabilities, based on the current security controls. It then attempts to evaluate the likelihood of a threat exploiting a vulnerability and the impact, or magnitude of harm, from this event. Identifying the likelihood and the impact helps to determine the level of risk. Risk management processes address high risks. Threat modeling is a type of risk assessment that attempts to predict an attack.

Risk assessments commonly use either a quantitative analysis or a qualitative analysis. A quantitative analysis uses numerical figures to calculate risk. The single loss expectancy (SLE) identifies the loss from a single event. Annual rate of occurrence (ARO) indicates how many times you expect the event to occur. Annual loss expectancy (ALE) provides the total expected loss for a year and is calculated as $SLE \times ARO$. This is useful to evaluate the value of a control. If the cost of the control is less than the ALE, it is justified. If it is significantly higher than the ALE, it is not justified. If the cost of the control and the ALE are relatively close, they must be evaluated to determine a long-term return on investment (ROI).

Many organizations designate a computer incident response team (CIRT or IRT) to respond to incidents. It's important to know the concepts related to incident response. The SSCP exam outline lists the phases of the incident response lifecycle as (1) preparation; (2) detection, analysis, and escalation; (3) containment; (4) eradication; (5) recovery; and (6) lessons learned/implementation of new countermeasures.

Questions

- 1.** Which of the following choices best represents the definition of risk?
 - A.** The likelihood that a threat source can cause a threat event resulting in a vulnerability
 - B.** The likelihood that a vulnerability can exploit a threat and cause a loss
 - C.** The likelihood that a threat will exploit a vulnerability and cause a loss
 - D.** The likelihood that an incident can cause a vulnerability resulting in a loss

- 2.** You are involved in risk management activities within your organization. Of the following activities, which one is the best choice to reduce risk?
 - A.** Reducing threats
 - B.** Increasing vulnerabilities

C. Increasing impact

D. Mitigating risk

3. Which of the following choices identify valid threat sources? (Choose all that apply.)

A. Employee

B. Earthquake

C. State-sponsored attacker

D. Administrator

4. Which of the following can cause a negative impact on an organization's assets?

A. A threat

B. A risk

C. A weakness

D. A control

5. An organization has recently suffered an attack resulting in monetary losses. Security professionals within the organization identified the tactics, techniques, and procedures used by the attackers and successfully eradicated all elements of the attack. What should these security professionals do with the knowledge they gained while mitigating this attack?

A. Tell personnel within their department only.

B. Post it on public forums.

C. Keep it private.

D. Share it in accordance with the security plan.

6. What is the purpose of risk management?

A. Eliminate risks

B. Reduce risks to an acceptable level

C. Share or transfer risks

D. Identify risks

7. What should an organization do when the cost of a control exceeds the cost of a risk?

- A.** Implement the control
- B.** Accept the risk
- C.** Perform a risk assessment
- D.** Mitigate the risk

8. Which of the following helps ensure that an organization focuses risk management resources only on the most serious risks?

- A.** Risk assessment
- B.** Residual risk
- C.** Countermeasures
- D.** Qualitative analysis

9. You decide to manage risk by purchasing insurance to cover any losses. Which one of the following risk management treatments are you using?

- A.** Accept
- B.** Avoid
- C.** Mitigate
- D.** Transfer
- E.** Recast

10. Which of the following is an open standard that organizations can use to assess vulnerabilities, by evaluating key metrics of a system?

- A.** CVSS
- B.** RMF
- C.** ALE
- D.** SLE

11. An organization has implemented several controls to mitigate risks. However, some risk remains. What is the name of the remaining risk?

- A. Vulnerable risk
- B. Mitigated risk
- C. Alternate risk
- D. Residual risk

12. A risk assessment recommended several controls to mitigate risks, but only some of the controls were accepted and implemented. Who is responsible for any losses that occur from the remaining risk?

- A. The person completing the risk assessment
- B. Senior management
- C. IT personnel managing the systems
- D. Security personnel

13. Of the following choices, which one most accurately reflects differences between risk management and a risk assessment?

- A. A risk assessment is a point-in-time event, while risk management is an ongoing process.
- B. Risk management is a point-in-time event, while a risk assessment is an ongoing process.
- C. Risk assessments are broad in scope, while risk management is focused on a specific system.
- D. Risk management is one part of an overall risk assessment strategy for an organization.

14. Which of the following is a type of risk assessment that attempts to predict an attack?

- A. CVSS
- B. Quantitative analysis
- C. Qualitative analysis
- D. Threat modeling

15. You are completing a risk assessment using historical data. You've identified that a system has failed three times in the past year, and each of these outages resulted in approximately \$10,000 in losses. What type of analysis does this allow you to perform?

- A. Qualitative
- B. Quantitative
- C. Subjective
- D. Informative

16. Which of the following formulas will determine the annual loss expectancy (ALE)?

- A. $SLE - ARO$
- B. $SLE \times ARO$
- C. $ARO - SLE$
- D. SLE / ARO

17. You are completing a risk assessment and using historical data. You've identified that a system has failed five times in each of the past two years. Each outage resulted in losses of about \$5,000. What is the ARO?

- A. Five
- B. \$5,000
- C. \$25,000
- D. Impossible to determine with the information provided

18. You have completed a risk assessment and determined that you can purchase a control to mitigate a risk for only \$10,000. The SLE is \$2,000 and the ARO is 20. Is this cost justified?

- A. Yes. The control is less than the ALE.
- B. No. The control exceeds the ALE.
- C. Yes. The control exceeds the ARO.
- D. No. The control is less than the ARO.

19. What is the first phase in incident response?

- A. Analysis
- B. Containment, eradication, and recovery
- C. Detection
- D. Preparation

20. Of the following choices, what best represents the last phase in incident response?

- A. Containment
- B. Lessons learned
- C. Analysis
- D. Recovery

Answers

1. C. A risk is the likelihood that a threat will exploit a vulnerability and cause a loss. Threats do not create vulnerabilities, and vulnerabilities do not exploit threats. Similarly, incidents do not cause vulnerabilities.

2. D. Risk mitigation is the process of reducing risk. You can rarely reduce threats, but you can often reduce (not increase) vulnerabilities or reduce (not increase) the impact of a risk.

3. A, B, C, D. All of the answers are valid threat sources. Employees can be adversarial threat sources if they intentionally cause damage or accidental threat sources if they accidentally cause damage. An earthquake is an environmental threat source. A state-sponsored attacker is an adversarial threat source. An administrator is an employee and can intentionally or accidentally cause damage.

4. A. A threat source can cause a negative impact by exploiting a vulnerability. Risk is the likelihood that a threat will exploit a vulnerability and cause a loss; the risk doesn't cause the negative impact. A weakness is a vulnerability. Attackers can exploit a vulnerability, but the vulnerability doesn't cause the loss. Controls attempt to reduce risk by reducing vulnerabilities or reducing the impact of a risk.

5. D. They should share this threat intelligence with specific entities inside and outside the organization, as specified in a security plan. Telling personnel in only one department is unlikely to be effective at mitigating future incidents but sharing some elements of the attack may be a good method of educating users of tactics, techniques, and procedures (TTPs) used by attackers. Attackers monitor public forums, so posting TTPs might help them understand what they can use elsewhere, or how they can evade detection using similar TTPs. Security professionals recommend sharing threat intelligence judiciously, not keeping it private.

6. B. Risk management reduces risks to an acceptable level. It is not possible to eliminate risk. One method of managing risk is to share or transfer risk, but that is not the only method. Similarly, risk management processes identify risk, but risk management is much more than just identifying risk.

7. B. If the cost of a control exceeds the cost of a risk, the organization should accept the risk. The organization might implement the control if the cost of the control was less than the cost of the risk, indicating a cost savings. A risk assessment can analyze the value of the control, but you wouldn't need to do a risk assessment if you already know the cost of the control exceeds the costs of the risk. Mitigating the risk indicates you would implement the control, but based on the known costs, it's appropriate to accept the risk.

8. A. A risk assessment helps an organization prioritize risks so that it can focus risk management resources on the most serious risks. Residual risk is the risk that remains after implementing risk mitigation steps. Countermeasures are risk management resources that reduce risks. Risk assessments can use either a quantitative analysis or a qualitative analysis, but neither is superior to the other one in all circumstances.

9. D. Insurance is one of the ways that you can manage risk by transferring the risk to a third party. Risk acceptance doesn't take any further action to mitigate the risk. In risk avoidance, you avoid the activity that results in the risk. It's most common to try to reduce the risk using risk mitigation. Recasting a risk refers to changing it to a different severity.

10. A. The Common Vulnerability Scoring System (CVSS) is an open standard that organizations can use to assess the severity of computer system security vulnerabilities. It evaluates various metrics to identify a score. A risk management framework (RMF) is a generic term for a structured process to evaluate risks. Some RMFs are freely available to anyone, such as NIST SP 800-37, while others are only available by paying for them. Annual loss expectancy (ALE) and single loss expectancy (SLE) are terms associated with a quantitative risk analysis.

11. D. Residual risk is any risk that remains after implementing controls to mitigate the risk. It's often not cost effective to implement controls to eliminate all risks, so senior management must make decisions on what risk to mitigate and what risk to accept as residual risk. Vulnerable risk and alternate risk are not valid terms associated with risk management.

12. B. Senior management is responsible for making decisions on what risk to mitigate. The remaining risk is residual risk and senior management is responsible for any losses from this residual risk.

13. A. A risk assessment is a point-in-time event, while risk management is an ongoing process. Risk assessment is one element of a risk management strategy, and risk assessments are generally focused on specific systems with a limited scope, while risk management is much broader.

14. D. Threat modeling is a type of risk assessment that attempts to predict an attack. It uses various tools to model aspects of an attack against a system, a network, an application, or data.

The Common Vulnerability Scoring System (CVSS) is an open standard that organizations can use to assess the severity of computer system security vulnerabilities. Threat modeling typically uses either a quantitative analysis or qualitative analysis to analyze the threat model.

15. B. A quantitative analysis uses numerical figures to identify the actual costs associated with a risk. A qualitative analysis uses subjective terms such as low, medium, and high to analyze a risk. There is no such thing as an informative analysis.

16. B. ALE is the product calculated from the single loss expectancy (SLE) and the annual rate of occurrence (ARO), or $SLE \times ARO$. The ALE is not calculated by subtraction of the ARO and SLE or by dividing the SLE and ARO.

17. A. The annual rate of occurrence (ARO) is five because it happened five times each in the past two years. The single loss expectancy (SLE) is \$5,000 and the annual loss expectancy (ALE) is \$25,000.

18. A. Yes, because the cost of the control is less than the annual loss expectancy (ALE), the cost is justified. The cost of the control is \$10,000 and the ALE is \$40,000. The annual rate of occurrence (ARO) is how many times the loss occurred (20 in the example), but it is only useful when you multiply it with the single loss expectancy (SLE) to identify the ALE.

19. D. The first phase in incident response is preparation, which includes creating an incident response plan. The other answers are valid phases in incident response, but they aren't the first phase.

20. B. The last phase in incident response is lessons learned, sometimes called post-incident activity. Containment is important once an incident has been detected, analyzed, and escalated. Analysis occurs after an event is discovered and attempts to determine if it is an incident and should be escalated. Recovery is done after an incident has been contained and elements of the incident have been eradicated.

NOTES

CHAPTER 8

Monitoring and Analysis

In this chapter, you will learn about

- Events of interest
- Intrusion detection systems
- Security alerts and false positives
- Comparing network-based and host-based intrusion detection systems
- Comparing intrusion detection systems and intrusion prevention systems
- Detecting and preventing attacks
- Wireless intrusion detection and prevention systems
- Security information and event management tools
- Performing security testing and assessments
- Vulnerability assessments
- Penetration tests

Operating and Maintaining Monitoring Systems

Monitoring systems within an information technology (IT) network helps to prevent, detect, and correct potential security incidents. Intrusion detection systems provide continuous monitoring of a network and report events of interest that might be an incident. Intrusion prevention systems go a step further and prevent malicious traffic from entering a network, thereby preventing potential incidents. The following sections describe these systems in more depth.

Events of Interest

Monitoring systems including logging capabilities and they log a wide variety of events. *Events* are observable occurrences in a system or network and typically recorded in a log. *Events of interest* are known events that an organization wants to monitor. Some are potential security incidents that need to be investigated. Others are events that need to be monitored to ensure compliance with laws or regulations.

A protocol analyzer (sometimes called a packet sniffer or just a sniffer) can capture packets sent through a network. It's also possible to configure it to capture only specific packets based on

what traffic an organization wants to monitor. Packets are stored within a file commonly called a *packet capture* or a *packet dump*. Tools can then be used to analyze the packet capture to detect potential events of interest.

The following list describes some common events of interest:

- **Anomalies** An anomaly is anything that deviates from normal behavior. Anomaly-based detection methods (described later in this chapter) capture a baseline (sometimes called a security baseline) of activity and then compare current activity against the baseline to detect, log, and send an alert when a system detects an anomaly.
- **Intrusions** An intrusion typically indicates unauthorized access into a system or network. One of the most common types of intrusions occurs when an uneducated user clicks a malicious link or opens a malicious attachment within a phishing e-mail. When successful, attackers typically use the compromised system to gain more and more access within the computer's network.
- **Data exfiltration** Attackers often capture data within a network and send it out of the network so that the criminals can harvest it. This can include files, e-mail from e-mail servers, databases, and any other data stored within the network. Attackers typically encrypt this data before exfiltrating it, so a jump in encrypted data leaving the network is very likely an item of interest that needs to be investigated.



NOTE You can think of exfiltrate as the opposite of infiltrate. Attackers infiltrate a network (or break into it), collect data, and exfiltrate the data by sending it out of the network.

- **Unauthorized changes** Change management processes (covered in [Chapter 10](#)) help prevent outages caused by unauthorized changes. Even if an unauthorized change doesn't cause an outage, it should be investigated. Many monitoring systems can detect changes and log them.
- **Compliance monitoring** Many organizations need to monitor systems and networks to ensure they are in compliance with various laws and regulations. As an example, an organization that handles health-related data needs to comply with the Health Insurance Portability and Accountability Act (HIPAA). A continuous monitoring system can detect events of interest for this law (along with other laws and regulations) and log them.

Intrusion Detection Systems

Intrusion detection systems (IDSs) provide continuous monitoring of networks and hosts to help protect them from attacks. The goal is to detect an attack as it occurs and raise an alert so that administrators can take steps to stop the attack.

The IDS monitors and records events in a log. It provides alerts (such as via e-mail or a text message) when it detects potentially malicious events. It can also take action to modify the environment to stop the attack and reduce its effects. For example, it might modify an access control list (ACL) on a router to block traffic or close connections on a server.

IDSs are either network-based or host-based. A network-based IDS (NIDS) monitors network traffic, while a host-based IDS (HIDS) monitors traffic for an individual system such as a server or workstation.

[Figure 8-1](#) shows a network with both a NIDS and a HIDS. Notice that NIDS agents are located on different network segments, such as before the DMZ, inside the DMZ, and on the internal network. HIDS software is installed on each individual server and workstation.

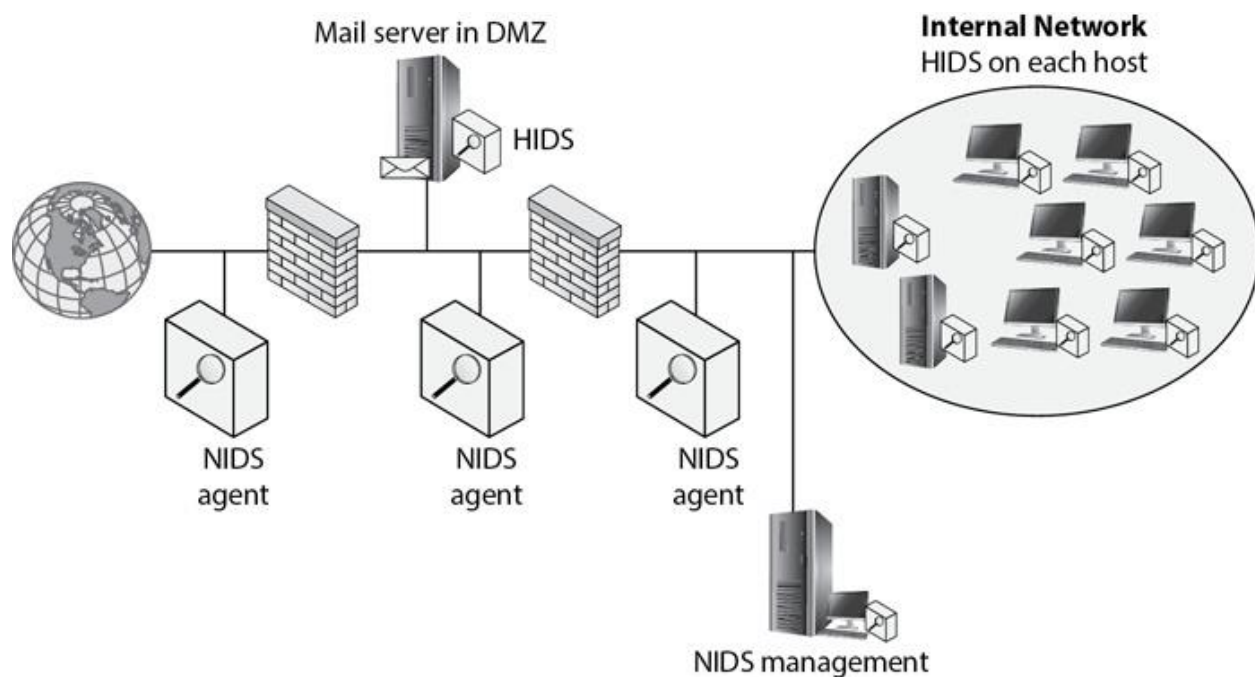


Figure 8-1 Intrusion detection systems protecting a network and systems

A NIDS collects information from the NIDS agents. Administrators install these agents on servers, routers, firewalls, or any other device that can monitor potentially malicious traffic coming into the network. In some cases, the NIDS agent forwards logs from these event source systems, and in other cases, the NIDS agent creates its own log and forwards it to the NIDS management system. A source system is any device that logs events. This includes devices such as firewalls, routers, proxy servers, and security appliances.

IDS Alerts

When an IDS detects an attack, it logs it, and typically reports it as an alert or an alarm. These reports come in many forms. Many IDSs include the ability to send an alert via an e-mail or a text message.

It's also possible to display the alert. As an example, a network operations center (NOC) typically has a wall of monitors in front of administrators and technicians. The administrators and technicians work on various tasks with desktop computers, but frequently glance at the monitors, and alerts displayed on these monitors will get their attention.

An alert or an alarm doesn't necessarily indicate an attack is under way. Personnel investigate the alerts to determine if they are valid. Unfortunately, it's common for an IDS to give false positives. In other words, an IDS might indicate a possible attack when an attack isn't underway at all.



EXAM TIP An alert provides notification of a potential adverse event. Personnel analyze the event to determine if it is an incident. An alert can be a false positive, which isn't an incident.

An IDS often triggers an alert when an event reaches a specific threshold. As an example, consider a port scan attack, where an attacker attempts to scan a system's ports to identify open ports. If the scan detects that port 80 is open, the attacker knows that the system is probably a web server running Hypertext Transport Protocol (HTTP) because port 80 is the well-known port for HTTP. A port scan will scan a list of ports and record what ports elicited a response (and were open) and what ports did not elicit a response.

If a remote system scans one port, that is probably not an attack. However, if an unknown external system scans all 1,024 well-known ports in a 60-minute period, it is very likely an attack.

Here's a trick question. If one port scan in an hour is not an attack, and 1,024 port scans in an hour is an attack, what is the lowest number between 1 and 1,024 that most likely indicates an attack? In other words, what should you set the threshold to so that the IDS detects a port scan attack?

There just isn't a good answer to that question. Port scanners allow attackers to randomize the order of the ports they scan and set delays between queries. If the attacker sets the delay to five minutes, the scanner does 12 port scans in an hour. If you set the threshold to 15, your system would be under attack but the IDS would not detect it.

In contrast, if you configure the threshold to 2 so that the IDS sends an alert if it detects two port scans in an hour, it will probably create many false positives. Your IDS will be known as the IDS that cries wolf and administrators will ignore it. Many administrators consider a threshold of two port scans in a minute to be too low because it will create so many false positives.

When choosing a number for this type of threshold, administrators recognize that the IDS will probably generate some false positives. They would rather see some false positives than configure the threshold so high that the IDS doesn't detect attacks in progress.



TIP Because attacker port scanning tools can randomize the order of ports they scan, and set delays between port queries, some IDSs use different techniques to detect a port scan. For example, an IDS might be configured to analyze the source and destination IP addresses for any isolated port queries. If the attacker is using the same computer (with the same source IP address) to query the same destination computer (with the same destination IP address), an IDS might be able to identify the trend within several hours, even if the attacking tool is only scanning three ports an hour.

Network-based Intrusion Detection Systems

A NIDS monitors the network traffic for any type of attack. It typically has several nodes or agents stationed around the network. These agents are installed on routers, firewalls, and network appliances such as unified threat management (UTM) devices. Each of these nodes monitors the traffic and reports its findings to a NIDS management server.

[Figure 8-2](#) shows one possible configuration. It has NIDS agents installed on two firewalls and one router, and each of these agents reports events to the NIDS management server. The server monitors the traffic, and if it detects an attack, it provides an alert to administrator or security personnel.

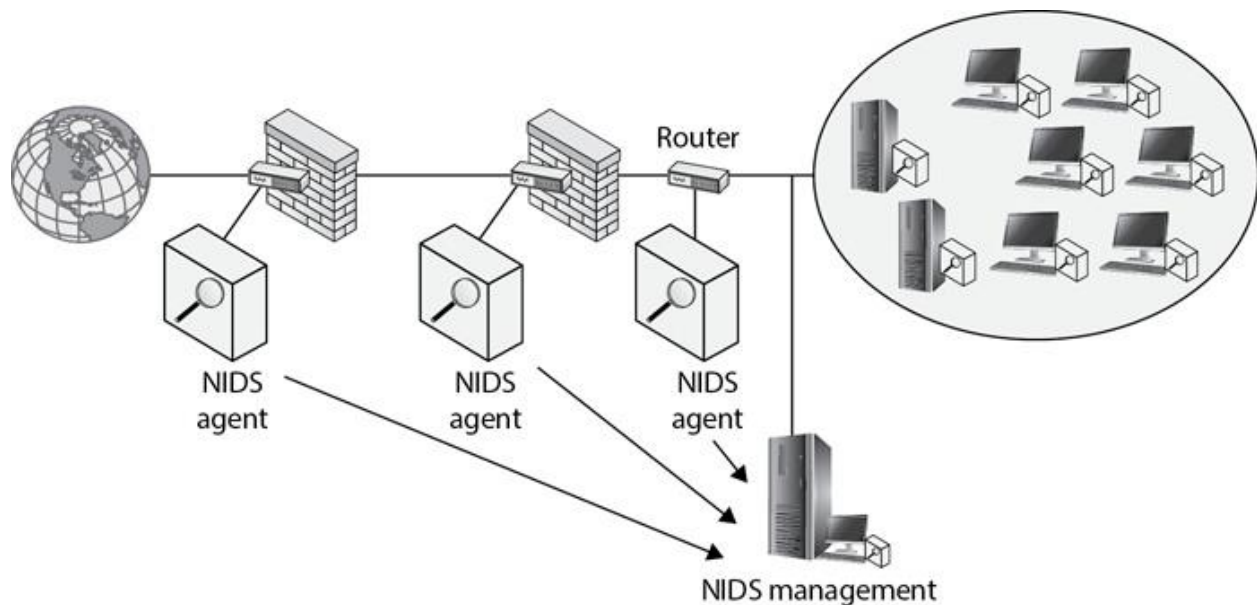


Figure 8-2 Network-based intrusion detection system with a NIDS management server

As discussed in [Chapter 4](#), many routers and switches support port mirroring. As a reminder, port mirroring sends a copy of all traffic passing through a device to a specific port, and this is one

way for agents to send traffic to a NIDS management server. Administrators configure the device to use port mirroring and connect the NIDS management server to the mirrored port.

One of the primary benefits of a NIDS is that the monitoring is transparent. Attackers don't have any clear indication that a NIDS is monitoring their activity. In contrast, it's easier for attackers to detect and disable a HIDS.



EXAM TIP A network-based IDS (NIDS) monitors traffic going through a network. It uses agents to monitor traffic on network devices, and the agents forward the traffic to a central management console.

Host-based Intrusion Detection Systems

A HIDS is installed on an individual system such as a server or workstation. It primarily monitors activity on the system such as system processes and running applications. It can only monitor activity on the host; it is not able to monitor overall network activity as a NIDS does. However, a HIDS can look at patterns on a host's disk and within its random access memory (RAM) in ways that a NIDS cannot.

Some organizations install a HIDS on every system in the network, while other organizations install them only on a few systems such as critical servers. The key is that a HIDS monitors systems only, while NIDS agents are installed on network components and monitor network activity. While antivirus (AV) software isn't technically a HIDS agent, many AV applications monitor the same resources (memory, disk usage, processes, and more) as a HIDS agent.

There are some drawbacks associated with a HIDS. First, the software can be very processor-intensive, slowing down an otherwise powerful computer. Second, installing a HIDS on each of your internal systems can be quite expensive. Last, many attackers can detect a HIDS, then disable it or erase the HIDS's logs.



EXAM TIP A host-based IDS is installed on a system such as a server or workstation. It monitors activity on the host, but cannot monitor network activity.

IDS and Defense-in-Depth

IDSs provide another layer of protection to systems and networks and are part of an overall defense-in-depth strategy. Firewalls monitor traffic going in and out of a network and can block much of the malicious traffic. Antivirus (AV) software detects malware and cleans, deletes, or

quarantines malicious files. The IDS focuses on intrusion attacks and can provide protection that is beyond the capabilities of the firewall and AV software.

[Chapters 4](#) and [6](#) discuss unified threat management (UTM) appliances and content-filtering appliances, respectively, and many of these appliances include intrusion detection components. They can also include spam filtering, malware protection, firewall services, proxy server components, and intrusion detection protection. A unified solution in a single appliance is typically easier to administer and manage.

Additionally, using HIDS with a NIDS solution is part of a defense-in-depth strategy. This is similar to installing AV software on each system even though AV software on a UTM appliance monitors all traffic coming into the network. The extra AV protection protects against a threat coming from another source, such as a USB drive inserted into a system. Similarly, installing a HIDS on each system protects against an attack launched internally, either from an internal attacker or from malware inadvertently released on an internal network.

Intrusion Prevention Systems

An intrusion prevention system (IPS) is an advanced implementation of an IDS. However, IPSs have been around long enough that they are often lumped together with IDSs. For example, National Institute of Standards and Technology (NIST) Special Publication (SP) 800-94 is titled *Guide to Intrusion Detection and Prevention Systems (IDPS)*.

However, there is a key difference between an IDS and an IPS. Specifically, a network-based IPS is placed inline with the traffic. For example, [Figure 8-3](#) shows an IPS placed at the boundary between the network and the Internet. It monitors all traffic going in and out of a network and actively blocks potentially malicious traffic.

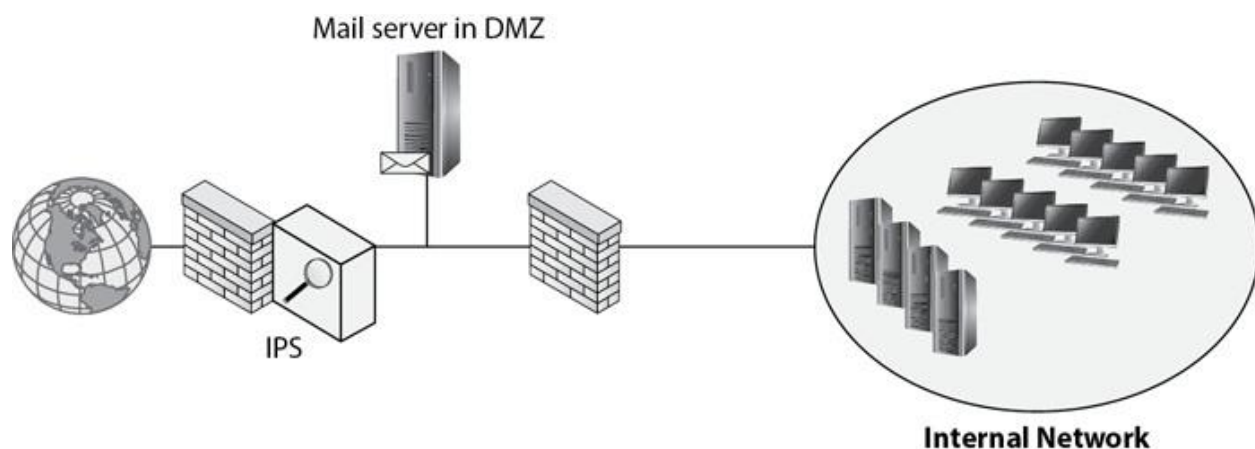


Figure 8-3 Intrusion prevention system placed inline



EXAM TIP A network-based IPS is placed inline with the traffic. It has the capability of modifying the environment to block an attack while it is in progress. In contrast, an IDS is not placed inline with the traffic, but instead monitors traffic with agents.

Although [Figure 8-3](#) doesn't show additional IDSs such as HIDSs on individual systems, they can be used with an IPS as part of a defense-in-depth strategy. An IPS can be host-based (as a HIPS) or network-based (as a NIPS), just as a traditional IDS can be host-based or network-based.

Notice that the IPS is placed inline with the firewall and that all traffic to and from the network goes through both the firewall and the IPS. The firewall protects a network from some malicious traffic, and the IPS helps protect the network from attacks that get through the firewall. Just as an IDS can modify the environment, an IPS can also modify the environment. For example, it might modify the ACL on a router if it detects a port scan from a specific IP address.



NOTE Many IDSs have the capability of providing a response that modifies the environment. Such systems are sometimes marketed as an IDS and sometimes marketed as an IDS/IPS or IDPS.

As another example, consider [Figure 8-4](#). It shows a server being attacked with a SYN flood attack. [Chapter 5](#) covers SYN flood attacks, but as a reminder, an attacker disrupts the Transmission Control Protocol (TCP) three-way handshake process, leaving half-open connections on a server. The attacker sends the first SYN packet, the server responds with the SYN/ACK packet, and then the server waits for the ACK packet. However, the attacker never replies with the ACK packet. The server can have hundreds of half-open connections waiting for the third packet in each.

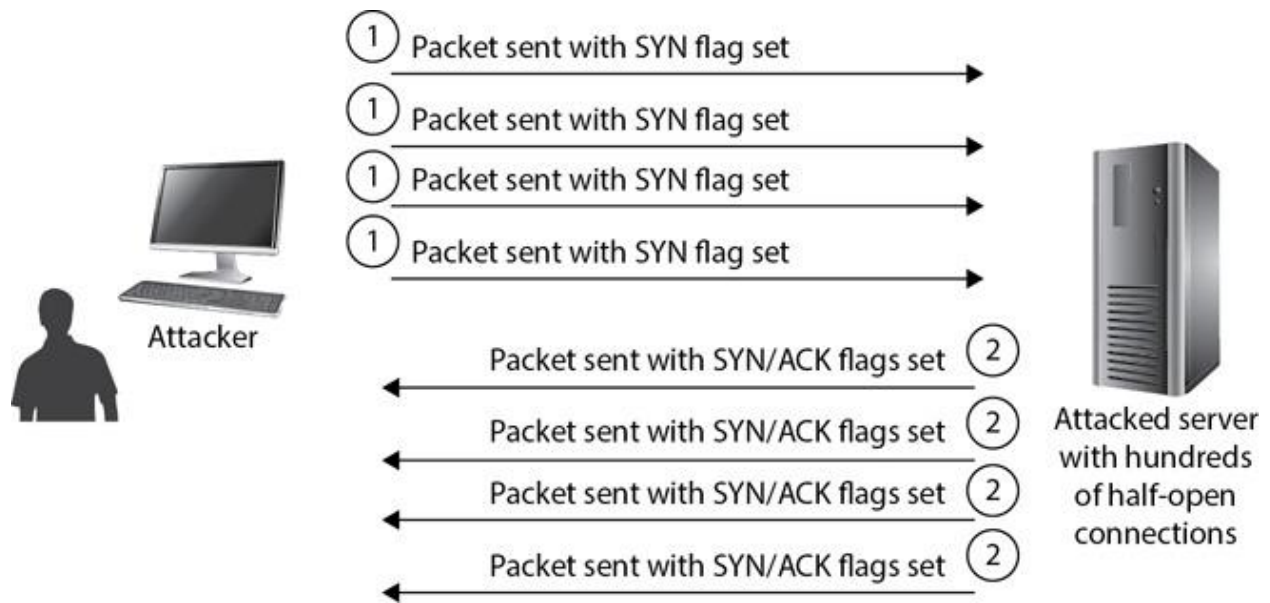


Figure 8-4 SYN flood attack

The IPS can detect this attack and block any more traffic from the attacker by modifying the ACL on the router to stop the attack. It can also close the half-open connections to free up the resources on the server that are being consumed by the attack, as shown in [Figure 8-5](#).

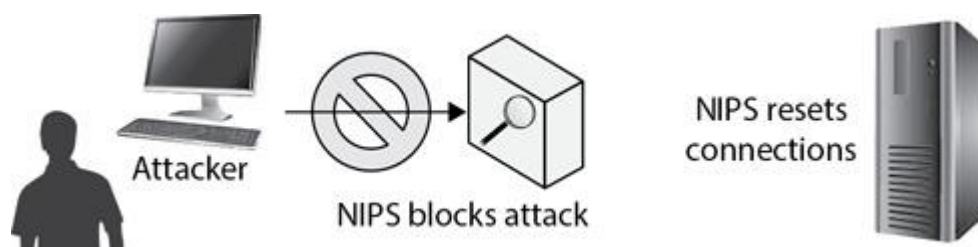


Figure 8-5 IPS blocking a SYN flood attack



NOTE Modifying the ACL on a router doesn't necessarily stop a SYN flood attack. Attackers often use botnets to launch distributed denial of service (DDoS) attacks. A protracted attack could be coming from tens of thousands of computers, none of which is the actual attacker.

It's worth emphasizing the differences between an IPS and an IDS. An IDS logs the event and provides notification. All IDSs and IPSs have this ability. The IPS is placed inline with the traffic and can take action to modify the environment and prevent the attack from entering the network.

Also, it's important to realize that neither an IDS nor an IPS launches attacks against an attacker. An IPS will change the environment to block the attack, but does not take any aggressive action against the attacker. Many organizations have security professionals with specialized skills who follow attackers through the Internet and launch attacks on them, but these skills are well beyond the scope of basic IT security professionals.

One of the basic risks of attacking an attacker is drawing attention to your organization. If an attacker attempts an attack and fails, that attacker may simply move on to another, easier target. However, if an attacker is attacked by an organization, that attacker may take it personally and launch protracted attacks against the company, holding a grudge for years. Additionally, there may be legal or liability issues with an organization launching an attack on a suspected attacker. Worse, because many attackers spoof their IP addresses or launch attacks through zombies controlled via a botnet, you may be attacking an unsuspecting victim.



EXAM TIP An intrusion detection system (IDS) is a detective security control. An intrusion prevention system (IPS) is a preventive security control.

Detection Methods

IDSs and IPSs use one of the following types of detection methods to detect attacks: signature-based, anomaly-based, or a hybrid of the two. Any type of IDS or IPS can use these methods.



EXAM TIP The two methods of intrusion detection are signature-based and anomaly-based detection. These are sometimes called knowledge-based (using signatures) and behavior-based (using a baseline to detect anomalies).

Signature-based

A signature-based (sometimes called knowledge-based) detection method is similar in concept to AV signatures. Recall that malware has unique characteristics, called signatures, that identify them. AV software has a database of signatures and compares potential malware against the database. If it gets a match, the software identifies the malware and attempts to clean or delete it.

Similarly, many attacks have unique characteristics documented in signature files. The IDS/IPS uses these signature files to identify and detect attacks. Both the vendor supplying the IDS and then the organization using the IDS must update these signature files regularly to detect newer attacks.

This shows a primary drawback with signature-based detection: It's good with known attack tools, but attacks may go unnoticed if an attacker uses modified attack methods.

Anomaly-based

Another method of detection is anomaly-based (sometimes called behavior-based). Administrators configure the device to monitor normal behavior to create a baseline. The device then observes activity and compares it to the baseline. If the current activity differs significantly from the baseline, the device issues an alert on the activity.

Anomaly-based systems go by many names, including behavior-based and heuristics-based. The important point is that they start by identifying normal activity and documenting this in a baseline, and then they detect abnormal activity that is outside the normal range of the baseline.

Artificial intelligence (AI) has advanced quite a bit in recent years and some companies are claiming that their anomaly-based systems are now using AI. As an example, IBM's Watson Explorer uses a cognitive miner and machine learning-based cognitive capabilities. Combined, these abilities allow Watson Explorer to analyze terabytes of information to detect patterns and build solutions. The goal is to provide administrators with actionable data as quick as possible.

A *false positive* is when an IDPS detects an attack but an attack is not occurring. In contrast, a *false negative* is when an attack is occurring but the IDPS does not detect it. All IDPSs are susceptible to false positives, and anomaly-based systems have a higher level of false positives than a signature-based system. For example, if an organization has a successful e-mail campaign resulting in a large customer response, this increased customer activity could be interpreted as an attack. Some IDPSs can be tuned to balance the acceptable level of false positives versus false negatives.

If the monitored system is modified, it's important to update the baseline. For example, if you are using a NIDS and you install software that increases network activity, the NIDS might generate many false positives if the baseline isn't updated.



EXAM TIP If the monitored environment is modified, it's critical to update the baseline in an anomaly-based system. Otherwise, the modification will be detected as an anomaly, resulting in a higher level of false positives.

Hybrid

A hybrid detection method simply uses a combination of both signature-based and anomaly-based detection methods. It's common for many IDSs and IPSs to use a combination of both detection methods to increase their overall performance.

Whitelisting

Although it isn't formally identified as a detection method for IDSs, the use of whitelisting is rising as an alternative to host-based IDSs. The whitelist detects abnormal activity by specifically identifying authorized activity. For example, an application whitelist identifies approved applications and processes. A system blocks applications and processes that aren't on the list.

Similarly, a switch can have a listing of system media access control (MAC) addresses authorized to communicate through a switch. Any systems not on the list that attempt to communicate through the path raise an alarm. This strategy isn't feasible for a large number of computers, but it can be useful for smaller networks. Still, it's important to realize that attackers can spoof MAC addresses. Whitelisting can be part of an overall strategy, but shouldn't be a primary method.

Wireless Intrusion Detection and Prevention Systems

A wireless intrusion detection system (WIDS) works similarly to an IDS, but for wireless networks, and a wireless intrusion prevention system (WIPS) works similarly to an IPS, but for wireless networks. They monitor wireless network traffic, looking for suspicious activity.

[Chapter 5](#) mentions WIDS and WIPS as countermeasures against wireless attacks. They are effective at detecting common attacks, such as from rogue access points, evil twins, and MAC spoofing. They use signature-based detection methods to detect traffic patterns from known attacks, and most include anomaly-based detection methods too. Additionally, they typically have the ability to inspect access points and identify configuration problems.

Analyze Monitoring Results

An important element of continuous monitoring is regular analysis of the monitored results. Administrators commonly use automated tools to assist in the analysis. The following list describes some common methods of analyzing the data:

- **Event data analysis** Individual events must be analyzed to determine if they are actual incidents. Administrators often use a variety of tools in the analysis. This includes logs on servers, routers, and firewalls in addition to logs collected by the IDS or IPS. In some cases, administrators need to collect additional data, such as by using a sniffer to capture and analyze the packets during an active attack, or the memory contents of an attacked system.
- **Reviewing security analytics and metrics** Administrators attempt to pull relevant data from the logs to create actionable information. For example, it's possible to measure the number of false positives for each week. These findings may show that the alert threshold is too high or too low and needs to be adjusted. Many systems include dashboards that make it easy to pull this data and view it in a timeline format.
- **Identifying trends** Metrics help identify a baseline and changes to that baseline. For example, an organization might be experiencing a steady increase in denial of service attacks.

Analysts can compare these attacks to performance issues with the attacked servers to determine if the security controls are providing adequate protection.

- **Creating graphics for visualization** You don't necessarily need pictures, but many tools allow you to create graphs and other graphics that are easily worth a thousand words. It is much easier to identify trends based on collected metrics with a graphic than it is to wade through text.
- **Document and communicating findings** After completing an analysis, it's important to document and communicate the findings. Organizations have different processes in place for this. Some organizations might have personnel use informal reports during the analysis, followed by a formal report after the analysis is complete. Based on the event, it might also be necessary to escalate the event.

Detection Systems and Logs

One of the challenges of detecting attacks is that attackers often try to erase their tracks. Logs will record what the attacker has done, but if the attacker is able to delete or modify the logs, the record is lost. With this in mind, it's important to realize that any logs on a local system that has been attacked are suspect. You need to investigate to determine whether the logs were modified during the attack.

On the other hand, logs stored on remote systems are the most valuable after an attack. For example, NIDS agents send a copy of their logs to a central NIDS management server. If a network firewall or resources within a DMZ are attacked, the logs on those resources may have been modified, but the logs on the NIDS server will still be intact and valid.



EXAM TIP After an attack, remote logs are the most valuable to re-create the events that occurred during and prior to the attack. Any logs on a local system should be treated with suspicion because the attacker may have modified them.

Detecting Unauthorized Changes

One of the ways that an IDS can detect an attack is by detecting unauthorized changes. In some cases, the changes can indicate a loss of integrity, and in other cases, they can indicate an attack in progress.

File Integrity Checkers

Many IDSs and antivirus applications use file integrity checkers to detect unauthorized changes to critical system files. A file integrity checker uses hashing techniques to detect these changes.

[Chapter 1](#) introduces hashing and [Chapter 14](#) covers it in more depth. You may remember that a hash is simply a number. A hashing algorithm creates the hash when it is executed against data, and no matter how many times the algorithm is executed against the same data, it will always create the same hash. However, if the data is modified, a new calculated hash will be different from the original hash. By comparing hashes created at different times, it's possible to determine whether the original file has lost integrity.

A file integrity checker calculates hashes on critical system files to establish a baseline. It then periodically recalculates the hashes of these files and compares the two hashes. If the hashes are different, it detects the change and sends an alert.

Of course, some files are modified as a part of normal system operation. For example, users modify their own documents regularly, so the file integrity checker doesn't monitor these files. However, it's rare for executable files to be modified except through authorized updates or malware infections. If the file integrity checker identifies unauthorized changes in executable files, it often indicates a problem such as a malware infection.

Unauthorized Connections

Many IDSs monitor the connections on systems and attempt to detect unauthorized connections. For example, if a server is hosted in an extranet and should be accessed only by a partner organization, an IDS can detect any unauthorized connections from unauthorized entities.

If a server is hosted within a DMZ and accessible to anyone on the Internet, detection of unauthorized connections isn't so easy. However, an IDS can identify connection attempts that are outside the scope of the server's purpose. For example, it's normal for a web server to have connections using HTTP and HTTP Secure (HTTPS). However, connection attempts using other protocols such as Telnet or Simple Mail Transport Protocol (SMTP) may not be authorized and indicate a potential incident.

Attackers often use scanning tools to learn information about a system, and these scanning tools often give telltale signs of their intentions. An IDS can detect these scans and send an alert.

Honeypot

A honeypot is a computer set up to entice would-be attackers. Administrators copy dummy data onto the server and configure it with weak security so that an attacker can easily hack into it. In essence, it's a trap designed to look appealing to attackers, just as a pot of honey looks appealing to Pooh Bear of *Winnie the Pooh* fame.

As an example, an organization can set up an FTP server with weak security that looks like it is hosting company data. Attackers may see the server, bypass the weak security controls, and attempt to download or upload data. While the attacker is attacking and probing the honeypot, that attacker is not in the actual network. Additionally, while the attacker is attacking, security professionals have an opportunity to observe the tactics used by the attacker. This observation

capability is especially useful when attackers use different methods or new exploits to attack a system.



EXAM TIP Attackers are constantly discovering new exploits, or new ways to exploit vulnerabilities. A honeypot gives administrators and security professionals an opportunity to observe new exploits.

A *honeynet* is an extension of a honeypot and includes two or more servers configured as honeypots. The goal of the honeynet is to simulate network activity similar to what an attacker would see in a live environment. In contrast, a single honeypot may not have any activity, enabling an intelligent attacker to determine that it is a trap.

Using Security Information and Event Management Tools

Security information and event management (SIEM) applications monitor an enterprise for security events. They are very useful in large enterprises that have massive amounts of data and activity to monitor, but smaller organizations also use them. Other names for these solutions include security information management (SIM), system event management, and security event management (SEM) applications.

Consider an organization with over 3,000 servers. When a situation arises on just one of those servers, administrators need to know as quickly as possible. The SIEM application provides constant monitoring and reporting capabilities to provide real-time reporting.



TIP Security information and event management applications are very useful for any enterprise. They have historically been very expensive, preventing smaller organizations from using them. However, open source solutions such as Security Onion are now available to smaller organizations.

A typical SIEM application can collect data from multiple sources. This includes firewalls, routers, proxy servers, IDSs, IPSs, proxy servers, and more. While this data is rarely logged in the same manner, the SIEM application uses aggregation and correlation techniques to combine the dissimilar data into a single database that can easily be searched and analyzed. In some cases, the SIEM application will include its own IDPS component, but it can collect data from an existing IDPS.

Much of the data that a SIEM application captures is called *machine data*. Machine data is the digital information created by a computing system's activity. On typical computers and servers,

machine data is all the logs created on the computers and servers. However, machine data refers to more than just the logs on typical computers and servers. It also includes data on mobile devices (such as smartphones and tablets), embedded systems, and Internet of Things (IoT) devices. An embedded system is any device that has a central processing unit (CPU), an operating system, and one or more applications that perform a dedicated function. IoT devices are embedded systems that have Internet access. They include digital cameras, smart televisions, and home automation systems.

Continuous Monitoring

A SIEM application provides continuous monitoring with a combination of tools. These tools within the SIEM application are combined into a single unified solution with several benefits, including the following:

- **Dashboard to monitor all activity** Administrators don't need to monitor multiple tools because the SIEM application provides a single interface (often called a *dashboard*). The dashboard can display the data in multiple ways, including graphics to provide a visualization of the data, timelines to show activity over time, and metrics that provide numbers to compare different data points. These various displays allow administrators to easily detect trends.
- **Database capabilities** Collected data is stored in a database that can easily be searched and analyzed for items of interest. Data includes logs collected from servers and network devices such as firewalls, routers, and switches. Log types include security logs, system logs, firewall logs, backup logs, and more. An IDPS typically includes a packet capture (sometimes called a packet dump) capability to capture all traffic that passes through the network, and the SIEM application can store these packet captures for later analysis.
- **Ability to define or fine-tune items of interest** SIEM applications come with predefined alerts, but they can be modified with different threshold levels or different items to search. For example, if management is interested in specific events occurring on a server, such as the use of privileges by anyone with an administrator account, they can have personnel configure the SIEM application to issue an alert each time these events occur.
- **Alerting capabilities** SIEM applications have the ability to issue an alert on any item of interest and provide different types of notifications. They can provide visual indications, such as pop-ups through the application's interface, and send notifications via e-mail or text messages.
- **Secure storage of data** Because a SIEM application is centrally located, the data can be protected from an attacker. In other words, the logs for individual monitored systems are stored remotely with the SIEM application. If any of the systems are attacked, the logs stored with the SIEM application remain intact.

When using a SIEM application, it is sometimes necessary to use time synchronization techniques to ensure the data remains meaningful. This isn't necessary if all the data comes from devices in a single location. However, if a SIEM application is collecting data from multiple locations and these locations are in different time zones, time synchronization techniques are

required. One method is to convert all time to Greenwich Mean Time (GMT), which is the time at the Royal Observatory in Greenwich, London.

Document and Communicate Findings

IT security administrators evaluate the output of monitoring systems looking for trends and events of interest. While many tools provide automated reports, administrators evaluate the reports and determine if they need to communicate their findings to others in the organization.

Imagine a SIEM application alerts on a potential DDoS attack and administrators verify the attack is real. The administrators would escalate the incident to appropriate personnel by communicating their findings.

An incident response plan typically identifies who to notify. As an example, a malware infection might only require the notification of a technician in the IT department to remove the malware using established procedures. A protracted DDoS attack might require notification of a chief technology officer (CTO) or chief information officer (CIO), again depending on the requirements dictated in the incident response plan.

Performing Security Tests and Assessments

Many organizations perform regular security assessments. Administrators use them to check the security posture of the organization and discover vulnerabilities. After performing a security assessment, they report their findings and recommendations, if appropriate, to management. Management can approve the recommendations to implement security controls to reduce the identified vulnerabilities, increasing overall security.

The most common type of security assessment is a vulnerability assessment. Organizations may take this a step further and perform a penetration test that attempts to exploit weaknesses discovered by a vulnerability assessment. Penetration tests are more intrusive and have the capability to cause damage if not approached cautiously.



TIP This section covers vulnerability assessments and penetration tests. [Chapter 10](#) covers security audits, which are also a form of a security assessment.

Vulnerability Assessments

Vulnerability assessments attempt to discover vulnerabilities and can take many forms. Some assessments attempt to discover whether an organization is susceptible to social engineering attacks. Other assessments attempt to discover whether the organization has technical weaknesses.

For example, phishing attacks are common, and an organization may want to test its employees' knowledge of phishing to see whether they are likely to respond. An online company, KnowBe4 (<https://www.knowbe4.com>), provides a free service that will send a bogus phishing e-mail to employees to determine how many employees are likely to respond to a phishing e-mail. Ideally, zero employees will respond, but if even one employee responds, it could compromise the organization's resources. Tests such as the one provided by KnowBe4 often show that many employees aren't educated about the risks of phishing e-mails and respond or click links.

Other security consultants use social engineering tactics to gain knowledge of an organization and then use this knowledge to gain more and more access. For example, an organization may hire a security consultant to see what information the consultant can learn about the organization just from calling employees or going through its trash.

Technical vulnerability assessments attempt to discover technical weaknesses. For example, vulnerability scanners can scan systems and networks to discover existing vulnerabilities.

APTs and Phishing

[Chapter 5](#) discusses advanced persistent threats (APTs) and spear phishing attacks. As a reminder, an APT is a highly sophisticated group who target their attacks and a spear phishing attack sends malicious e-mails to specific targets. APT 28 (nicknamed Fancy Bear) and APT 29 (nicknamed Cozy Bear) used spear phishing attacks very effectively to infiltrate organizations. The U.S. Department of Homeland Security (DHS) and the Federal Bureau of Investigation (FBI) released a Joint Analysis Report titled "GRIZZLY STEPPE – Russian Malicious Cyber Activity" (also called JAR-16-20296A) in late December 2016 providing details on these attacks. The spear phishing attacks included subject lines and text within the e-mail that are relevant to the recipient. These e-mails typically include a link to a malicious ZIP file.

If a user opens the ZIP file, it installs a backdoor on the user's system. This backdoor causes the user's system to initiate connections with the attackers' command and control server. Attackers configure the command and control server to respond by sending commands back to the user's system.

One method attackers use is to embed commands within a web page as HTML tags. These commands allow the attacker to perform internal reconnaissance on the infected user's network using standard commands. For example, **net view domain** will show the domain name and **net view /domain:domainname** (with the name of the domain in place of *domainname*) will list all of the computers in the domain. The attackers typically dump this list to a text file with a command like this: **net view /domain:domainname >> 1.text**. They encrypt the data and transmit the file back to the command and control server. Attackers then use this information to install multiple backdoors on other computers in the network.

Think about this for a moment. One employee in an organization opens a ZIP file and the result is a massive infiltration of backdoors installed on multiple systems. Once the attackers establish a foothold, they begin collecting data, compressing it into multiple files, and sending these files

to the attackers. A vulnerability assessment that checks to see if users are vulnerable to a similar spear phishing attack can be very valuable.

Test Types

Personnel performing a vulnerability assessment have different levels of knowledge, based on the type of test they are performing. There are three common test types:

- **White box or internal testing** Testers perform the assessment within the organization. They have full access to the internal network and know the network infrastructure, including what systems it hosts. Security professionals who are employees within the organization commonly do internal testing. Other names for this type of testing include clear box testing, glass box testing, and transparent box testing.
- **Black box or external testing** Testers don't have any knowledge of the internal network prior to starting the testing. Consultants hired to test an organization's security vulnerabilities often do black box testing.
- **Gray box testing** This is a cross between white box testing and black box testing. It can be done internally or externally, but the testers have at least some level of knowledge about the network.



TIP These test types are also called full knowledge (white box), zero knowledge (black box), and partial knowledge (gray box) tests.

[Figure 8-6](#) shows how a white box internal test is able to test multiple systems on the internal network, the DMZ, and even firewalls protecting the network. [Chapter 5](#) defines white hats, black hats, and gray hats. These terms are completely different from white box testing, black box testing, and gray box testing. For example, testers working for an organization with no malicious intent are white hat testers. White hat testers can perform white box testing, black box testing, or gray box testing.

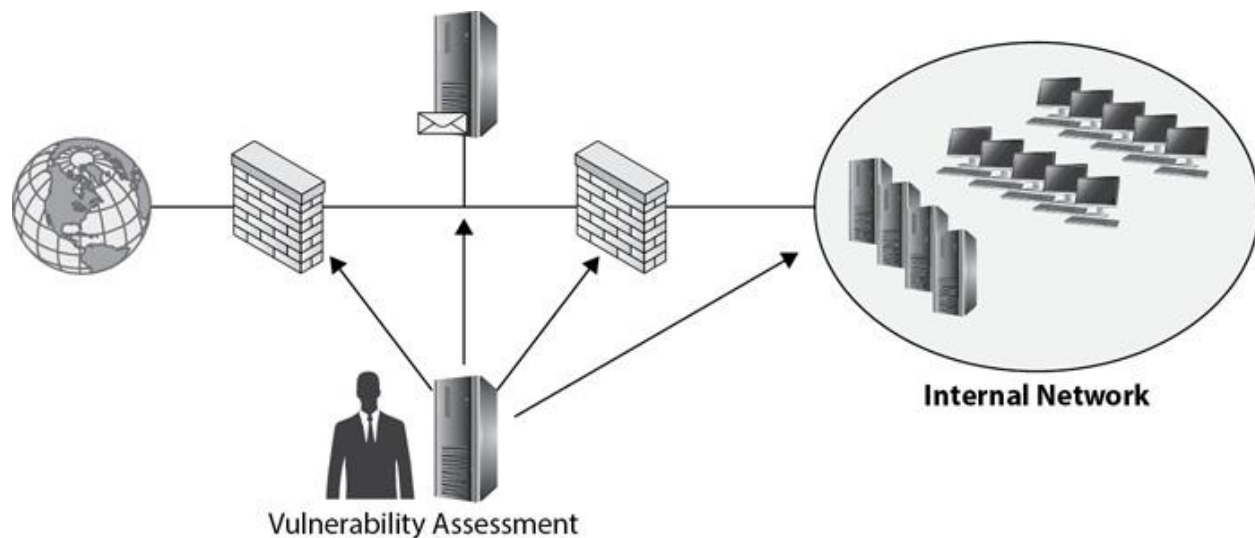


Figure 8-6 White box testing launched internally to check systems

It's worth stressing that many attackers use the same tools organizations use to assess an organization's vulnerabilities for potential attacks. In other words, if you can discover a vulnerability with one of these tools, the odds are very good that an attacker can also discover these vulnerabilities.

Vulnerability-Scanning Tools

Many vulnerability-scanning tools are available to automate scans of a network. Some tools are comprehensive and do multiple functions, while other tools perform a single function. These tools attempt to perform overall reconnaissance of a network. The result is a listing of the machines that are operating on the network along with the likely services running on each system.

Some of these tools start with an Internet Control Message Protocol (ICMP) sweep to identify computers based on their IP addresses. They follow this with a port scan on each of the computers to identify which ports are open on individual computers. Next, they try to identify more details about the system, typically by analyzing packets sent from the system.

As an example, [Table 8-1](#) shows the values of two observable items in a packet. The recommended default value for the initial Internet Protocol (IP) Time To Live (TTL) is 64, but this is only a recommendation and different operating systems have different values. It refers to the number of routers, or hops, a packet will traverse before it is discarded. Similarly, different operating systems use different TCP window sizes for the initial packets in a TCP session. In some cases, an operating system version may be using different values.

| Operating System | IP Initial TTL (hops) | TCP Window Size (bytes) |
|------------------|-----------------------|-------------------------|
| Windows | 128 | Varies (Auto-Tuned) |
| Linux | 64 | 5,840 |
| FreeBSD | 64 | 6,5535 |
| Cisco IOS | 255 | 4,128 |

Table 8-1 TTL and Window Sizes



NOTE Current Windows systems include a feature called Auto-Tuning. This allows the system to vary the TCP window size dynamically. However, the values used by Auto-Tuning are known, so by collecting several packets and comparing the different values, a vulnerability scanner can typically identify the operating system.

The initial TTL and TCP window size are only two items in a packet that identify details about a system. Vulnerability scanners use multiple comparisons to identify a system. This process is often referred to as *fingerprinting* the system.



TIP Scanning tools often use sniffers to capture and analyze packets. [Chapter 5](#) discusses protocol analyzers in more depth. Additionally, these tools often perform the same type of reconnaissance and fingerprinting techniques discussed in [Chapter 5](#).

Nmap (short for Network Mapper) can locate and document all the devices on a network as part of a reconnaissance scan. Nmap uses scanning methods first to discover all the devices based on their IP addresses, and then uses scans on each individual system to identify services running on each system. Nmap can often identify a system's operating system, services running on the system, the existence of a host-based firewall, and much more.

Nessus is a vulnerability-scanning tool used to discover vulnerabilities in individual systems. It can use Nmap to discover systems on the network during an initial reconnaissance scan. It then uses other methods for the fingerprinting scan. It can perform detailed analysis of individual systems, including auditing passwords for complexity, checking systems for compliance with specific security policies, and checking for the existence of personally identifiable information (PII) on a system.



EXAM TIP Many laws mandate the protection of PII. Tools such as Nessus can be used by organizations as part of a compliance check to ensure that they are in compliance with existing laws. Organizations keep the results of these scans to prove that the organization is actively involved in compliance monitoring.

There are many other vulnerability-scanning tools available. Security professionals testing networks for vulnerabilities can use these tools, but remember that criminals attempting to break into networks can also use them.

No matter which tool you use, you can expect the security professionals to provide detailed reports documenting the findings. Of course, an extremely important step in the vulnerability assessment process is to analyze these results and address any vulnerabilities that can be mitigated. Data in these reports is sensitive and should be protected. If attackers have access to the reports, they can use the information to launch attacks on documented vulnerabilities.

Reviewing Infrastructure Security Configurations

Infrastructure security refers to securing the hardware elements of a network. It's important to review these security elements periodically to verify they are up to date with current security practices.

Administrators performing daily tasks in a network often get used to the configuration and can overlook some things. Because of this, organizations often bring in an outside consultant to examine the infrastructure and provide recommendations. A security consultant will typically perform the following tasks and provide recommendations to improve security:

- **Assess network vulnerabilities** This can be done by running various vulnerability-scanning tools such as Nessus.
- **Verify appropriate host hardening** This check ensures that servers are running only the protocols and services they need, and all other protocols and services are disabled. It also ensures that host-based firewalls are enabled, and systems are up to date with current patches.
- **Check border security** This includes checking all paths into the network from external sources such as the Internet. The consultant would check network devices such as routers, firewalls, UTM appliances, and the DMZ configuration.

Instead of hiring outside consultants to perform these reviews, some organizations have dedicated IT security staff. If these personnel are given the time and resources to focus on security, they can perform these assessments on a routine basis.

Order of Steps

Vulnerability assessments typically use the following order of steps:

- 1. Gain permission from management.** The importance of this step cannot be overstated. Senior management from any organization typically interprets unapproved vulnerability scans as attacks. This often results in the loss of a job or even legal action. Most security professionals require written approval before starting an assessment. This is so important that if someone is willing to perform an assessment without written approval, they should be viewed with suspicion.
- 2. Discovery.** In this step, personnel perform vulnerability scans on the network to identify weaknesses. It often includes reconnaissance steps to document the environment, followed by fingerprinting scans to identify specifics about the systems, such as the operating system or running services. The detailed scans on individual systems can detect specific vulnerabilities for each of the scanned systems. Testers often do both credentialed scans and noncredentialed scans. A credentialed scan checks for vulnerabilities with a privileged account, and a noncredentialed scan checks for vulnerabilities without being logged on.
- 3. Analyze results.** Vulnerability scanners provide reports that personnel must analyze to determine their validity. It's possible for a scanner to misidentify a system or report a false positive. Security professionals who regularly run scans recognize the weaknesses and the strengths of the tools. For example, a scan may regularly report that a firewall is not enabled on a host simply because it doesn't recognize a third-party firewall. However, when new vulnerabilities appear on reports, they deserve extra scrutiny to determine their validity.



NOTE Some scanners can be configured to eliminate false positives for specific events. For example, it may be possible to eliminate a false positive if the scanner doesn't recognize a third-party firewall. The challenge is ensuring that the configuration change doesn't create any false negatives. In this example, you want to ensure the scanner still reports when a system doesn't have any firewall enabled.

- 4. Document vulnerabilities.** Next, personnel document the vulnerabilities in a report. Some organizations want this report turned in as soon as possible so that management has an idea of the vulnerabilities. For example, in one large network operations center where I worked, a security group performed detailed scans once a month. As soon as the group was able to compile the results, it published the report, making it available to key management personnel in other departments. In this situation, the security group simply performed the scans and reported the results, and other departments were responsible for identifying and recommending methods to remediate the vulnerabilities.

5. *Identify and recommend methods to reduce the vulnerabilities.* Next, personnel identify methods that can mitigate the vulnerabilities. These could be as simple as recommending a configuration change or as complex as recommending a detailed patch management process to ensure all systems are kept up to date.

6. *Present recommendations to management.* Management has the responsibility to examine recommendations and either approve or reject the recommendations. They have the overall view of the organization and the responsibility for any losses associated with residual risks for any rejected recommendations.

7. *Remediate.* Last, personnel implement any recommendations approved by management. It's worth stating the obvious here: If a fix isn't implemented, the vulnerability remains and attackers can exploit it. If management has approved a recommendation, it should be implemented as soon as possible.



EXAM TIP The last step (and arguably the most important step) in a vulnerability scan is to implement the fixes. Many organizations often repeat the scan after implementing the fix to verify that they have actually removed the vulnerability as expected.

Remediation Validation

Remediation validation is the process of verifying that approved fixes are implemented. Some approved changes can be implemented very quickly. For example, if a vulnerability assessment found that Telnet was enabled on many servers in the network, the remediation is to disable Telnet on all the servers. Administrators might be able to use Group Policy to remove Telnet from all the servers, assuming they are all Microsoft servers within an Active Directory domain.

In contrast, some changes might take much longer. As an example, an assessment using a different assessment tool might discover many vulnerabilities that weren't previously known. Security personnel might take the time to prioritize the vulnerabilities and create a plan of action to address each based on the severity of the vulnerabilities. IT administrators would then begin implementing the changes (typically using a change control process).

The remediation validation timeline depends on the timeline of fixes. However, the key is to take action to validate or verify that the fixes are implemented.

Audit Finding Remediation

Organizations often perform audits to verify that an IT security system is complying with the organization's security policy, guidelines and requirements from partners, and relevant laws. These audits can be performed by personnel within the organization or by personnel outside the organization.

As an example, the Payment Card Industry Data Security Standard (PCI DSS) is a standard used by organizations that handle many major credit cards. Its purpose is to help ensure these organizations comply with minimum IT security standards. Many organizations can certify compliance by answering questions within a Self-Assessment Questionnaire (SAQ). However, other organizations must bring in an external auditor to verify compliance. These external auditors are referred to as *assessors* and are certified by PCI Security Standards Council to perform the audits.

The auditor creates a report documenting the results of the audit, commonly called the audit findings. Audit findings verify an organization is in compliance with the requirements and document any instances in which an organization is not in compliance. If an organization is not in compliance, it creates a plan to resolve the issues.

Just as an organization should verify that approved changes from an assessment have been completed, it's also important to verify that audit findings have been addressed.

False Positives and the Importance of Results Analysis

In some cases, a vulnerability assessment may indicate that a vulnerability exists even though it doesn't. This is a false positive. For example, a scanning tool may indicate that a server is missing a patch, but investigation verifies that the patch is installed. This is often a weakness in the vulnerability-scanning tool, and some tools are more prone to false positives than other tools are.

However, the important thing to recognize is that false positives can occur. Just because the tool indicates a vulnerability exists, that doesn't mean it does. Administrators or security analysts examine the results to determine their validity.



TIP False positives from vulnerability tools are often unique to each tool. One tool may be able to effectively determine if a patch is installed, while another tool may report erroneously that a patch is missing.

Get Out of Jail Free Card

If an organization hires you to perform a vulnerability assessment or penetration test, it's important that you obtain written permission from management before starting. If not, you could face legal action after doing the test.

As an example, several security professionals were granted access to an organization's campus to secretly perform vulnerability testing. Some administrators detected their activity and did some digging. With the help of internal security personnel, they located the pseudo-attackers in an area of a building that was under renovation.

These pseudo-attackers were dressed as construction workers and working behind some tarps. With all the computer equipment they had running, it was clear they weren't construction workers. However, as soon as they were discovered, they quickly showed their written authorization, calling it their "get out of jail free" card from Monopoly fame. A phone call confirmed that the written authorization was valid and they were not arrested.

In contrast, if they didn't have written authorization, the result would have been much different. They would have been taken to the campus security office and perhaps arrested. They probably wouldn't have spent much time in custody, but what if the individuals who verbally OK'd the test weren't available? Worse, what if the managers who approved the test recalled a different version of what they authorized? The team might be accused of exceeding their authority.

Written documentation avoids all these problems.

Vulnerability tests can sometimes give false negatives also. As an example, imagine management decides not to implement an operating system patch on a server because it breaks an application. A vulnerability scan should indicate the system is vulnerable because the patch isn't applied, but if it doesn't, it's a false negative. This can give a false sense of security because a vulnerability scan may indicate that a network is secure, yet actual vulnerabilities exist.

Penetration Tests

Penetration testing takes a vulnerability assessment one step further. Instead of stopping after discovering vulnerabilities, a penetration test attempts to exploit the vulnerability. While a penetration test provides verification that an attacker can (or can't) actually exploit a system, it comes with significant risks.

It's important for testers to recognize the potential impact of security testing. It's possible for a penetration test to cause damage resulting in the loss of confidentiality, integrity, or availability of a system or data. This is never a goal, but due to the nature of penetration testing, it is possible, so testers must be cautious when testing.

Both internal and external penetration tests are possible, and they are classified as white box (full knowledge), black box (zero knowledge), and gray box (partial knowledge) tests, just as vulnerability tests are classified. The amount of information given to the testers is dependent on the goals of the test.

In many scenarios, penetration testers are given specific goals. As an example, they may be tasked with locating user databases within a network, exfiltrating them, and then cracking the password hashes for various users. If successful, they can provide management with a listing of usernames and passwords. As another example, testers may be tasked with gaining access to an employee's computer, using privilege escalation techniques to gain additional access, and executing scripts to map the network.

[Figure 8-7](#) shows the overall steps of a penetration test, and the following sections cover them in more depth.

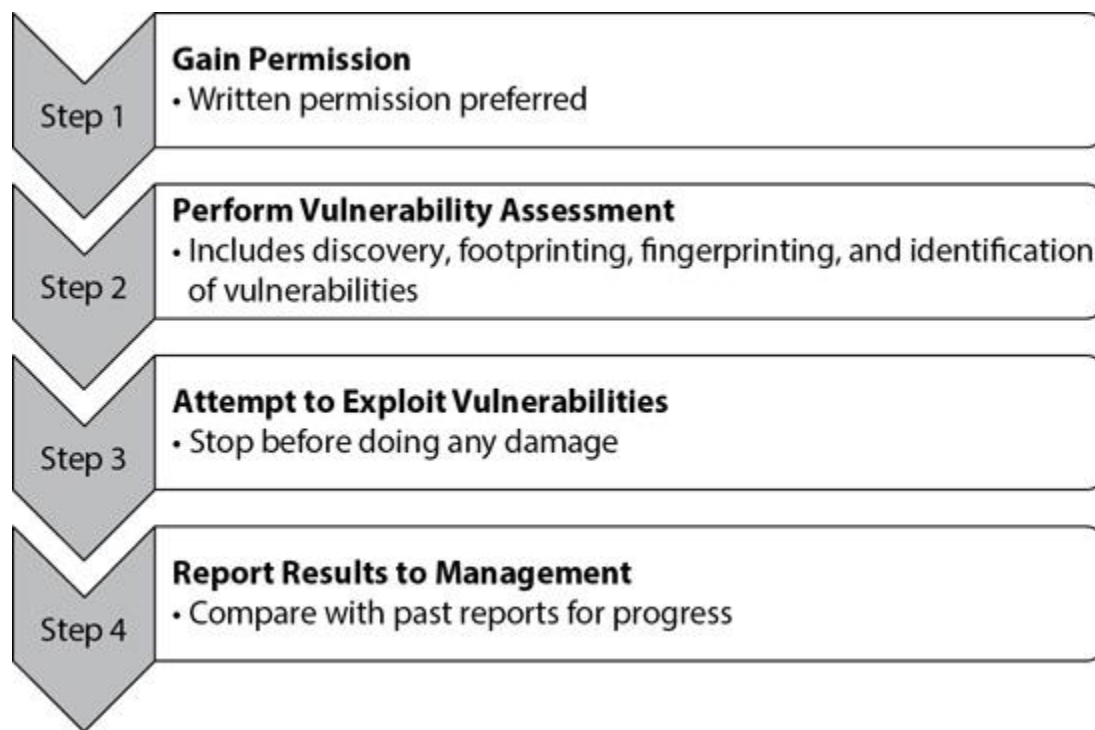


Figure 8-7 Penetration test steps



EXAM TIP A penetration test is an intrusive test. However, all penetration tests should stop before causing damage to a live system or network. The goal is to prove that exploiting a vulnerability can affect the mission of the organization. The test should stop before it actually affects the mission of the organization.

Gain Permission

Gaining permission is an even more important step in a penetration test than it is with a vulnerability test because it's possible that a penetration test can actually cause a loss in confidentiality, integrity, and/or availability of a system. The permission needs to be in writing and should identify the goals of the test.

Perform Vulnerability Assessment

In the vulnerability assessment stage, you go through similar steps as you would in a typical vulnerability assessment. While the actual steps of the vulnerability assessment can vary depending on the scope, it's very likely it will include the following components:

- **Discovery or reconnaissance** In this stage, you collect information about the targets using any available means. This could use tools such as Nmap to complete a reconnaissance scan, or

use existing documentation if the test is carried out as a white box test. Black box testers often need to use a variety of social engineering tactics to gain information about the organization and the network.

- **Fingerprinting** Once the network is mapped, individual systems are identified and probed to discover details about them. This includes the operating system, service packs, patches, open ports, running services, and running applications.
- **Identify vulnerabilities** The tester then analyzes the collected data to determine what vulnerabilities exist and can be exploited.

If this were a vulnerability test, the data would be put into a report and the test would be complete. However, a penetration test adds the next step: an attempt to exploit one or more vulnerabilities.

Attempt to Exploit Vulnerabilities

In this step, the tester attempts to use methods similar to what an attacker would use to exploit the vulnerabilities. At this stage, it's very important to understand the scope and purpose of the test. Most organizations want to know what can be exploited, but not at the expense of an outage. In other words, the tester should take the test as far as possible without causing any damage.

Of course, the danger with this strategy is that you don't always know how far is too far until you get there. Testers should be acutely aware of the potential damage that an exploit can cause and be very cautious when the exploit can be damaging.

Some exploits result in the disclosure of data. In this case, testers can gather as much data as possible without causing harm. The testers can show the data they've collected to demonstrate that an attacker can also access it. Other exploits can modify or corrupt data, and these types of tests should not be performed on live production servers. However, it is possible to create a test server configured identically to the production server. Testers then attempt the penetration test on the test server. If the test on the test server succeeds, it provides proof that the exploit is possible without affecting actual resources.

Report Results to Management

The results of the penetration test are compiled into a report, along with recommendations for controls or safeguards to mitigate any of the vulnerabilities. Just as with a vulnerability report, management decides on what risks to mitigate and what risks to accept.

Past reports are valuable to compare. For example, a past report may have documented a vulnerability, and management may have approved a control to mitigate the risk. If the vulnerability still exists, it deserves some investigation. Was the control implemented? Was it completely ineffective, or did it actually reduce the vulnerability? Has the threat technique been modified to circumvent the original control? Answers to these questions help determine the next steps to take.

Chapter Review

Events of interest are known events that an organization wants to detect and record. Some events of interest are potential security incidents that need to be investigated. Others are events logged for compliance monitoring.

An intrusion detection system (IDS) provides continuous monitoring and sends alerts to report suspicious activity. Administrators investigate alerts to determine whether they are actual attacks or false positives. NIDS agents are installed on network devices and send data back to a central network-based IDS (NIDS) management server. This server analyzes the activity throughout the network and provides notification when an alert occurs. A host-based IDS (HIDS) monitors activity on an individual system, such as a server or workstation, and provides notifications as alerts. Notifications are typically recorded in logs but can also be sent as e-mails or text messages. An intrusion prevention system (IPS) is placed inline with the traffic and can modify the environment to block attacks in progress.

IDSs and IPSs use two primary detection methods. A signature-based (or knowledge-based) system has a database of known attacks and compares live activity with this data. This is similar to how AV software uses signature files to detect malware. An anomaly-based (or behavior-based) system starts with a baseline of normal activity. It then monitors ongoing activity to determine whether there is a significant change or anomaly. Anomalies are reported as alerts and need to be investigated. Many systems use a hybrid model of both signature-based and anomaly-based detection methods.

One of the benefits of NIDSs over HIDSs is that the log files are stored remotely on the central server. A HIDS stores the logs on the host, and an attacker can modify or delete the logs, erasing the record of the attack.

Many times an organization wants to detect unauthorized changes. Tools such as file integrity checkers monitor files and can detect unauthorized modifications.

Honeypots are systems with fake data designed to attract attackers, diverting them from a live network. They also give security professionals an opportunity to observe attacks in action. Honeynets are groups of two or more honeypots to simulate a network.

Security information and event management (SIEM) applications are unified continuous monitoring systems. They collect data from multiple sources and typically include tools to aggregate, correlate, and analyze the data. IT administrators analyze the output of a SIEM application and escalate incidents when appropriate by communicating their findings to other personnel.

Most organizations perform security assessments to detect potential security issues. A vulnerability test attempts to discover vulnerabilities and provide recommendations to mitigate the associated risks. Testers can use nontechnical means, such as social engineering tactics, and use technical means with tools such as Nmap and Nessus.

Penetration tests go a step further than a vulnerability test. They attempt to exploit a vulnerability. Penetration tests are more intrusive and have the potential to cause damage, so they should be performed with utmost caution and stopped before inflicting any actual harm. Testers should obtain written permission prior to beginning either a vulnerability assessment or a penetration test. Both vulnerability assessments and penetration tests provide a report to management as the last step.

Both vulnerability tests and penetration tests can be white box (full knowledge), black box (zero knowledge), or gray box (partial knowledge) tests.

Questions

1. You want to monitor the network for possible intrusions or attacks and report on any activity. What would you use?

- A.** HIPS
- B.** HIDS
- C.** NIDS
- D.** AV software

2. You want to monitor a server for potential attacks. Of the following choices, what is the best choice?

- A.** HIDS
- B.** NIDS
- C.** Anomaly-based IDS
- D.** Signature-based IDS

3. Which of the following choices identifies a major drawback associated with a host-based IDS (HIDS)?

- A.** It is very processor intensive and can affect the computer's performance.
- B.** The signatures must be updated frequently.
- C.** It does not support anomaly-based detection.
- D.** It stores the logs on remote systems.

4. Of the following choices, what best describes an IPS?

- A. An active antivirus program that can detect malware
 - B. An inline monitoring system that can perform penetration testing
 - C. An inline monitoring system that can perform vulnerability assessments
 - D. An inline monitoring system that can modify the environment to block an attack
5. Which of the following identifies a system that requires a database to detect attacks?
- A. Anomaly-based IDS
 - B. Signature-based IDS
 - C. HIPS
 - D. NIPS
6. You have recently modified the network infrastructure within your network. What should be re-created to ensure that the anomaly-based NIDS continues to work properly?
- A. Signature database file
 - B. Baseline
 - C. Router gateways
 - D. Firewalls
7. How does an anomaly-based IDS detect attacks?
- A. It compares current activity against a baseline.
 - B. It compares current activity against a database of known attack methods.
 - C. It compares current activity with antivirus signatures.
 - D. It monitors activity on firewalls.
8. What logs are most valuable after an attack?
- A. Logs on a remote system
 - B. Logs on local systems that have been attacked
 - C. Logs for local firewalls

D. Logs for antivirus events

9. Which of the following can detect if a system file has been modified?

A. Encryption algorithm

B. Anomaly-based detection

C. Signature-based detection

D. File integrity checker

10. You are an IT administrator working in a 24-hour network operations center. One of your tasks is to evaluate alerts from various tools and determine if any are events of interest. Which of the following events would you most likely flag as events of interest? (Select all that apply.)

A. An apparent network intrusion caused by a user clicking a malicious link within an e-mail.

B. An unauthorized change updating an application on a web server.

C. An alert on a DoS attack that was proven to be a false positive.

D. Privacy data being sent out of the network by a user via e-mail.

11. Administrators in your organization recently installed a new SIEM application within your network. Which of the following best describes the capabilities of a typical SIEM application?

A. Collect, aggregate, and correlate data from multiple devices in a network

B. Protect a web server from attacks

C. Monitor traffic going through DMZ firewalls and alert on potential attacks

D. Perform automated vulnerability assessments

12. Your organization handles and transmits a significant amount of privacy data. As a result, it needs to comply with several laws and regulations. Which of the following tools would you most likely find within the network to provide continuous monitoring of all network traffic?

A. Vulnerability scanner

B. Penetration test

C. File integrity checker

D. SIEM

13. An external organization is performing a vulnerability test for a company. Officials from the company give this group some information on the company's network prior to the test. What type of test is this?

- A. White box test
- B. Gray box test
- C. Black box test
- D. Internal test

14. Your organization has contracted with a security organization to test your network's vulnerability. The security organization is not given access to any internal information from the company. What type of test will the organization perform?

- A. White box testing
- B. Gray box testing
- C. Black box testing
- D. Partial knowledge testing

15. An organization has hired you to perform a vulnerability assessment. Which of the following steps would you perform first?

- A. Document vulnerabilities
- B. Fingerprint systems
- C. Perform reconnaissance
- D. Gain approval

16. An outside security consultant recently performed a vulnerability assessment on your organization's network. She recommended several security controls to mitigate some risks. Which of the following steps would be the last step to complete the assessment?

- A. Discovery
- B. Analysis
- C. Remediation
- D. Document vulnerabilities

17. A security expert in your organization regularly scans your network to detect potential vulnerabilities using a vulnerability scanner. How would this vulnerability scanner most likely fingerprint a system?

- A.** With a biometric scanner
- B.** Using an ICMP sweep
- C.** By identifying its IP address
- D.** By analyzing packets

18. A vulnerability assessment reports that a patch is not installed on a system, but you've verified that the patch is installed. What is this called?

- A.** Anomaly-based vulnerability
- B.** Signature-based vulnerability
- C.** False negative
- D.** False positive

19. What's the primary difference between a penetration test and a vulnerability assessment?

- A.** A vulnerability assessment includes a penetration test, but a penetration test does not include a vulnerability assessment.
- B.** A penetration test is intrusive and can cause damage, while a vulnerability assessment would not cause damage.
- C.** A vulnerability assessment is intrusive and can cause damage, while a penetration test is passive.
- D.** They are basically the same, but with different names.

20. When should a penetration test stop?

- A.** After discovering the vulnerabilities
- B.** After discovering the threats
- C.** Before causing any damage
- D.** Before discovering the exploits

Answers

- 1. C.** A network-based IDS (NIDS) monitors the network in real time and sends alerts to report suspicious activity. A host-based intrusion detection or prevention system (HIDS or HIPS) cannot monitor network activity. AV software detects malware, but not necessarily network attacks.
- 2. A.** A host-based IDS (HIDS) can monitor a single computer (such as a server) for possible attacks and intrusions. A network-based IDS (NIDS) monitors network activity. Both HIDS and NIDS can use either anomaly-based detection or signature-based detection.
- 3. A.** A HIDS is very processor intensive and can negatively affect the computer's performance. All IDSs using signature-based detection need to have their signatures updated frequently. A HIDS will typically support both signature-based detection and anomaly-based detection. Another drawback of a HIDS is that it stores the logs on the local computer, and an attacker might be able to delete the logs.
- 4. D.** An intrusion prevention system is an inline monitoring and detection system that can modify the environment (such as by changing ACLs or closing half-open connections) to block an attack. Although it may be able to detect some malware such as worms, this isn't the best definition. It is not used for vulnerability assessments or penetration tests.
- 5. B.** A signature-based IDS compares activity against a signature file (or database of signatures) to identify attacks. An anomaly-based IDS requires a baseline. Both a host-based IPS (HIPS) and a network-based IPS (NIPS) can use either anomaly-based or signature-based detection methods.
- 6. B.** An anomaly-based NIDS compares current activity against a baseline to determine abnormal behavior, and this baseline should be updated when the network is modified. This provides the NIDS with an accurate baseline. A signature-based NIDS uses a signature database file. Router gateways and firewalls do not need to be re-created for the NIDS.
- 7. A.** An anomaly-based IDS detects attacks by comparing current activity against a baseline. A signature-based IDS detects attacks by comparing network activity with a database of known attack methods. An IDS does not use antivirus signatures. While an IDS monitors activity on firewalls, this doesn't identify how it detects attacks.
- 8. A.** Logs held on remote systems are the most valuable because an attacker is less likely to have modified them. Attackers can modify logs on a local system (including local firewall logs) during an attack. Although AV logs may be useful in attacks involving malware, they aren't very useful if attackers didn't use malware during the attack.
- 9. D.** File integrity checkers use hashing techniques to detect whether modifications of system or other important files have occurred. File integrity checkers use hashing algorithms, not encryption algorithms. Intrusion detection and prevention systems use signature-based detection to detect known attacks and use anomaly-based detection to detect previously unknown attacks, but these are not related to detecting modifications to system files.

10. A, B, D. Events of interest include intrusions (including those caused by a user clicking a malicious link), unauthorized changes, data exfiltration (including from an internal user), and more. A false positive is not an event of interest.

11. A. A security information and event management (SIEM) application collects data from multiple sources, then aggregates and correlates the data. It combines the dissimilar data into a single database and analyzes the data to detect potential events of interest. While it might detect attacks on a web server, it does much more. Similarly, it would probably monitor traffic going through demilitarized zone (DMZ) firewalls, but not just the DMZ firewall traffic. A SIEM application wouldn't perform vulnerability assessments.

12. D. A security information and event management (SIEM) application collects data from multiple sources and provides continuous monitoring. When configured to do so, a SIEM application can monitor all network traffic. None of the other answers provide continuous monitoring. A vulnerability scanner will scan a network at a point in time. A penetration test is performed during a specific period, but it ends at some point. A file integrity checker can detect modifications in files.

13. B. A gray box test (also called a partial knowledge test) is performed with some internal knowledge and can be performed either internally or externally. A white box test (also called a full knowledge test) is performed with full access to internal documentation. A black box test (also called a zero knowledge test) is performed without any inside knowledge of the organization. It is not an internal test because an external organization is performing it.

14. C. A black box test (also called a zero knowledge test) is performed without any inside knowledge of the organization. A white box test (also called a full knowledge test) is performed with full access to internal documentation. A gray box test (also called a partial knowledge test) is performed with some internal knowledge.

15. D. An important first step in a vulnerability assessment is gaining written approval from management. Discovery includes a reconnaissance scan and fingerprinting techniques, but should not be done without approval. Documenting the vulnerabilities comes after discovery and analysis.

16. C. The last step in a vulnerability assessment is remediation of vulnerabilities using controls approved by management. Discovery, analysis, and documentation all occur after gaining approval from management, but before remediation. In this scenario, the outside consultant would perform the assessment and provide a report to management. Management approves some or all of the security controls, and personnel then implement the controls to remediate the vulnerabilities.

17. D. A vulnerability scanner fingerprints a system by analyzing packets sent out by the system. In this context, fingerprinting is a metaphor and it doesn't use a biometric scanner. An Internet Control Message Protocol (ICMP) sweep identifies IP addresses on a network, but the IP addresses do not fingerprint a system.

18. D. A false positive occurs when a vulnerability assessment tool indicates that a vulnerability exists when it actually does not exist. A false negative occurs when a vulnerability assessment tool indicates that a vulnerability doesn't exist when it actually does exist. Anomaly-based and signature-based vulnerabilities are detection methods of IDSs and are not associated with vulnerability assessment tools.

19. B. A penetration test is intrusive; it includes a vulnerability assessment, attempts to exploit discovered vulnerabilities, and can cause damage. Vulnerability assessments are not intrusive and do not cause damage.

20. C. Most penetration steps stop before executing an exploit that can cause damage. Penetration steps include the basic vulnerability assessment components of identifying vulnerabilities, threats, and exploits, and then follow these steps with an attempt to exploit a discovered vulnerability.

NOTES

CHAPTER 9

Controls and Countermeasures

In this chapter, you will learn about

- Security controls, safeguards, and countermeasures
- Differences between deterrent, preventive, detective, corrective, and compensating controls
- Differences between administrative, technical, and physical controls
- Common steps to harden systems
- Differences between policies, standards, procedures, and guidelines
- Basic controls such as change control, configuration management, and patching systems
- Mobile device management methods and other endpoint device security concepts
- Use of RAID subsystems to provide fault tolerance
- Use of failover of clusters to protect against the failure of a server
- Load balancing methods such as round robin and source address affinity
- Different types of backups, such as full, differential, and incremental

Using Security Controls, Safeguards, and Countermeasures

[Chapter 7](#) provided information on risk and explained that risk is the probability or likelihood that a threat will exploit a vulnerability and cause a loss. [Chapter 8](#) presented information on performing vulnerability assessments to detect vulnerabilities. Both chapters mentioned that an organization mitigates risk by implementing security controls. This chapter digs deeper into security controls.

Security professionals often use the terms *controls*, *safeguards*, and *countermeasures* interchangeably. They are the means, methods, actions, techniques, processes, procedures, or devices that reduce the vulnerability of a system or reduce the probability that a threat will exploit a vulnerability and cause a loss. In this chapter, I've primarily used the term *controls*. However, if you see the terms *safeguards* or *countermeasures*, it's important to recognize they mean the same thing.

Security controls help prevent, detect, and correct problems related to security incidents. In general, a security incident is any event that results in a negative impact on the confidentiality, integrity, or availability of an organization's assets. This includes losses from data breaches,

denial of service (DoS) attacks, malicious software (malware) infections, power loss, and failure of environmental systems such as heating, ventilation, and air conditioning (HVAC) systems.

[Figure 9-1](#) shows the two ways that controls are used. They can either reduce vulnerabilities to reduce losses from a risk, or attempt to neutralize a threat to reduce losses from a risk. For example, antivirus (AV) software protects against malware. When administrators install AV software and keep it up to date, the system is less vulnerable to malware attacks. The AV software also helps to neutralize the threat of malware.

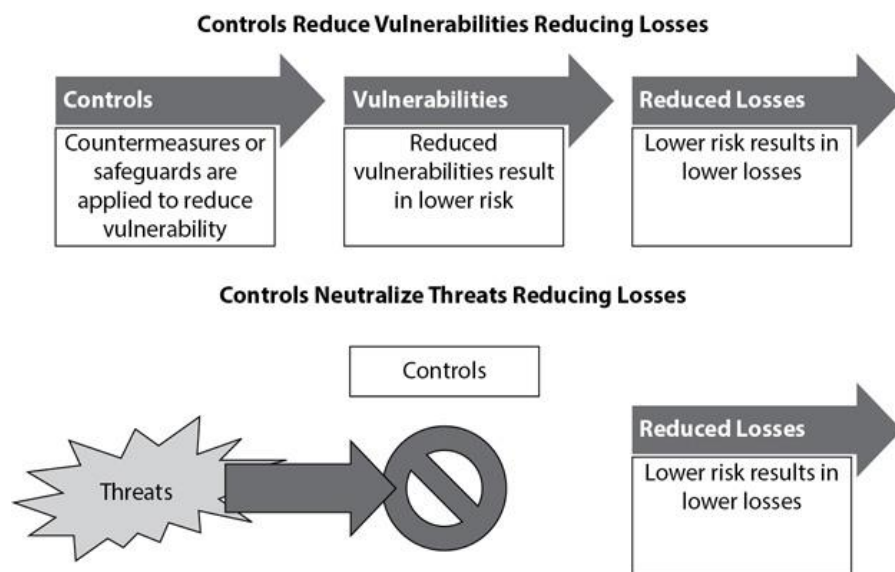


Figure 9-1 Controls reduce vulnerabilities, resulting in reduced losses.

Controls can be either technical or nontechnical. Technical controls use technical means to reduce risk and are merged with hardware, software, and firmware to either reduce vulnerabilities or reduce the impact of threats. Nontechnical controls are physical and administrative controls and include items such as written security policies, physical security controls, operational procedures, and personnel training.

For example, an intrusion prevention system (IPS), covered in [Chapter 8](#), is an example of a technical control that attempts to detect and block attacks. User training is an example of a nontechnical control that attempts to educate users and encourage them to avoid risky behavior such as clicking a link within a phishing e-mail.

Performing a Cost-Benefit Analysis

Organizations have a finite amount of money, so they must evaluate controls and purchase controls they perceive as having the best value. A cost-benefit analysis (CBA) attempts to determine the value of a control to the organization. A CBA compares the cost of a control with the potential benefits of the control to help determine if the cost of the control is justified. The CBA also compares the cost of implementing a control with the cost associated with losses if the control is not implemented.

Personnel often perform a CBA for each recommend control. This makes it easier for management to make an educated decision about purchasing the control. [Chapter 7](#) presents both quantitative analysis and qualitative analysis methods. A quantitative analysis uses monetary amounts that work well with a CBA.

When evaluating any control, it's important to evaluate both its expected effectiveness in mitigating a risk and its cost. Further, when evaluating the cost of the control, you must consider more than just the initial cost. Controls have several costs that management must consider before approving the purchase. These additional costs include the following:

- **Initial cost** The cost of the product. It could be a one-time purchase price or could entail monthly or annual recurring costs. For example, AV software has both an initial cost and a subscription cost to update definitions regularly. Of course, not using any type of AV software results in recurring costs too.
- **Implementation costs** Costs required to implement the control. Technical controls often require additional costs to implement them. Personnel may need to test the control for compatibility, resulting in labor costs. A control may require a vendor to install it initially, or may require training for administrators before the administrators can install and implement the control.
- **Compatibility costs** The costs associated with the usability of a system after implementing a control. Although security often requires a balance with usability, it's easy for security professionals to get overzealous and implement controls without considering the impact on usability. If the control negatively affects productivity or revenue, it can affect the mission of the organization.



EXAM TIP The cost of a control is justified if the cost of the control is significantly lower than the annual loss expectancy (ALE) without the control. The cost is usually not justified if the cost of the control is significantly higher than the ALE without the control. When the costs and savings are about the same, management may request a return on investment (ROI) analysis to determine whether the cost is justified. The ROI looks at the cost over a longer period of time.

Security Controls Lifecycle

Security controls go through a natural lifecycle. This includes selecting, implementing, assessing, and monitoring security controls. This is an ongoing process, and when monitoring indicates that a control is not effective, personnel must evaluate and select additional, or different, controls.

[Figure 9-2](#) shows the overall steps of selecting, implementing, assessing, and monitoring security controls. These steps are derived from National Institute of Standards and Technology (NIST)

Special Publication (SP) 800-53 Revision 4, *Security and Privacy Controls for Federal Information Systems and Organizations*. The steps are as follows:

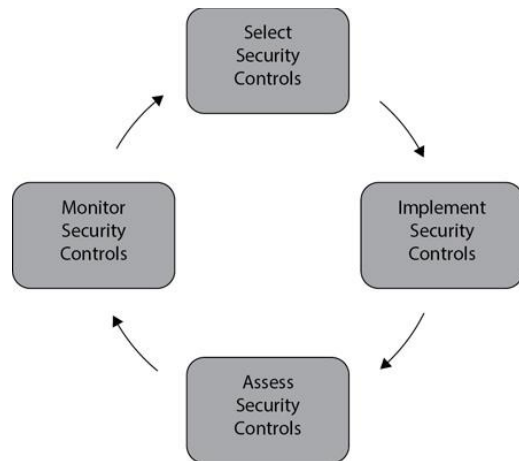


Figure 9-2 Security controls lifecycle

- **Step 1: Select security controls** After identifying the information system or process that an organization wants to protect, security personnel select and evaluate one or more security controls. The evaluation often includes a risk assessment, including either a quantitative analysis or a qualitative analysis, as described in [Chapter 7](#), and a cost-benefit analysis to determine the benefit of the control. While security and administrative personnel recommend security controls, management ultimately selects the security controls.
- **Step 2: Implement security controls** Next, personnel implement the security controls. This can be relatively easy to do or extremely complex, depending on the control. For example, creating a failover cluster to protect availability of a server requires research to purchase the correct hardware, training for personnel to install and maintain the hardware, and a phased implementation plan to put the failover cluster into service.
- **Step 3: Assess security controls** After implementing the security controls, personnel evaluate them to verify they are implemented correctly, operating as expected, and meeting the intended security requirements.
- **Step 4: Monitor security controls** Personnel monitor the security controls on a continuing basis to determine their effectiveness. They also monitor additional factors such as changes to the protected system and changes to external laws, regulations, policies, and guidelines.

It's worth stressing that the lifecycle of security controls is an ongoing process. Because threats are constantly changing, security personnel need to constantly monitor the environment for new threats and recommend security controls to mitigate associated risks.

Understanding Control Goals

Controls and countermeasures are primarily identified as one of the following three types: preventive, detective, or corrective. A preventive control attempts to prevent losses before they occur, a detective control detects violations, and a corrective control attempts to reverse the impact from a security incident. Other controls include compensating and deterrent controls, discussed later in this section.

Losses to confidentiality, integrity, or availability (CIA) can affect the organization's mission. [Figure 9-3](#) emphasizes that controls attempt to prevent, detect, and correct losses to any of these elements of the security triad.

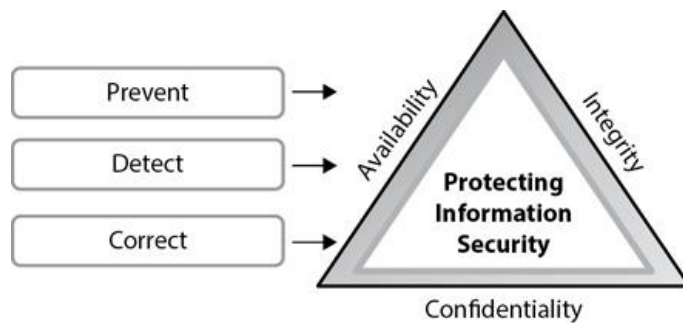


Figure 9-3 Controls prevent, detect, and/or correct losses to CIA.

Some controls combine all three objectives. For example, AV software running on an e-mail server can strip off malicious attachments to prevent them from reaching the user. AV software running on a user's system can detect malware, such as when a user inserts an infected Universal Serial Bus (USB) flash drive into a system. This AV software can also correct the problem by removing the malware from the infected USB flash drive.



EXAM TIP Controls have three primary goals: prevention, detection, and correction. Some controls have only a single goal (prevention, detection, or correction), while other controls attempt to achieve multiple goals.

Preventive

Preventive controls focus on preventing losses due to risks. The most basic type of preventive controls consists of written security policies, standards, and procedures. For example, many of the basic controls, such as separation of duties, least privilege, and strong authentication, are preventive controls identified in written security policies and implemented with different types of controls. Some security controls, such as fences, security cameras, and guards, can deter threats, so can be classified as both preventive and deterrent controls.



NOTE Most documentation uses the term preventive. However, the SSCP objectives (and the CISSP objectives) use the term preventative. Both terms mean the same thing in this context.

The following are some examples of preventive controls:

- Intrusion prevention systems
- Written policies and procedures
- Background checks on prospective employees
- Implementation of separation of duties and least privilege policies
- Strong access control processes (starting with strong authentication)
- Technical password policies that force users to create strong passwords and change them periodically
- Termination processes that ensure personnel disable accounts for terminated employees
- Classification of data (such as public, private, and proprietary) and implementation of varying levels of protection based on the classification
- Encryption of data (both at rest and in transit)
- Security cameras, fences, and guards



TIP Separation of duties can also be detective in nature if it includes a reconciliation or audit function. For example, if accounts payable is separated from accounts receivable, and a third party periodically reconciles (or audits) the accounts payable and accounts receivable data, the reconciliation process is a detective control that can detect anomalies.

Detective

Detective controls identify the event either as it is occurring or after it has occurred. Such controls provide evidence that preventive controls are working (or not working). Information can be used as evidence that a loss has occurred and can be used to provide an immediate response

by combining the detective control with a corrective control. The goal is to detect incidents as soon as possible so that either a technical control or a person can respond to resolve the problem.

The following are some examples of detective controls:

- Forensics analysis
- Physical inventories
- Intrusion detection systems
- Scripts to automate discovery of events of interest
- AV software that can detect malware installed on a system
- Audit logs (including logs on servers and network resources such as firewalls)
- Reconciliation (comparing different sets of data with each other, such as comparing a database inventory with a physical inventory)



NOTE Many of the controls mentioned throughout this chapter are covered in greater depth elsewhere in this book. I've pointed out the chapters for some of the controls in Notes or Tips. However, you can also use the Table of Contents and the Index to locate where I've covered any of these controls in more depth.

Corrective

A corrective control takes action to reverse the effects of an event and/or restore a disabled or failed control. Many corrective controls often work in conjunction with detective controls. After detecting the event, the corrective control goes into action. For example, after updating signatures on AV software, the AV software might detect an active virus and quarantine it. This reverses the infection. Similarly, a host-based intrusion detection system (HIDS) might detect that the host firewall is disabled and automatically enable it to correct the problem. The following are some examples of corrective controls:

- Intrusion detection systems (that can actively make changes)
- AV software that can remove or isolate malware
- Procedures to back up data and restore backups

- Disaster recovery and business continuity plans
- Scripts to automate corrective processes



EXAM TIP Many controls combine preventive, detective, and corrective measures. For example, AV software is detective when it identifies malware and corrective when it removes it. Similarly, a backup policy is a preventive control because it attempts to prevent the loss of data. The procedures to restore data from backups are corrective controls because they provide the means to restore data after a data loss.

Other Controls

Beyond preventive, detective, and corrective controls, you should also know about some additional controls, including compensating, deterrent, directive, and recovery controls. The SSCP objectives added compensating and deterrent controls in the previous edition and they are in this edition too. It could be argued that each of the following controls could be classified as preventive, detective, or corrective. However, classifying a control as compensating, deterrent, directive, or recovery provides a more focused classification:

- **Compensating controls** These controls are in place in case a primary control fails or is unavailable. For example, imagine a small organization employs the principle of separation of duties to ensure that no one employee controls an entire process in the accounting department. However, occasional illnesses or vacations require one person to do all the work. To compensate for this the organization can implement routine reconciliation audits as a compensating control. The audit will review and validate activity and can detect potential anomalies. You can also think of these as alternative controls that an organization implements when it isn't feasible or desirable to use the primary control.
- **Deterrent controls** These controls attempt to deter would-be attackers from attempting an attack or even deter users from trying to circumvent policies. For example, just the presence of a guard can deter an attacker from attempting to tailgate into an organization. Similarly, when a proxy server blocks access to a site from within an organization, it could display a web page reminding users of the acceptable use policy (AUP) and informing them that the proxy server records all their online activity. This message can deter users from purposely accessing unauthorized pages. Note that a deterrent control is similar to a preventive control. The primary difference is that a deterrent control encourages someone to *decide* not to take a specific action. In contrast, a preventive control doesn't rely on someone doing the right thing, but instead implements methods to prevent them from taking specific actions.
- **Directive controls** These are controls mandated by a higher authority. For example, the Health Insurance Portability and Accountability Act (HIPAA) directs organizations to implement

specific safeguards to protect private health information (PHI). The controls that an organization implements to protect PHI are directive controls.

- **Recovery controls** These controls are focused on recovering a system or mission after an outage, such as after a disaster. For example, reliable backups, along with detailed procedures on how to restore backups, are recovery controls. These are often considered corrective controls.



EXAM TIP When preparing for the SSCP exam, focus primarily on the preventive, detective, corrective, deterrent, and compensating controls. You may see directive and recovery controls too, especially if you pursue the CISSP certification.

Thousands of Controls

NIST SP 800-53 (*Security and Privacy Controls for Federal Information Systems and Organizations*) provides extensive coverage of controls. The first 45 pages of this 462-page document provide a solid foundation of the fundamentals related to security controls and the process of selecting security controls. The ten appendixes provide additional details. Appendix F, the “Security Control Catalog,” includes over 200 pages describing specific controls in several different categories.

If you’re interested in pursuing a career in IT security, I strongly recommend you download and peruse this document. You can get it from here: <https://csrc.nist.gov/publications/sp>.

One of the most important messages that you can glean from NIST SP 800-53 is that IT security has a lot of depth. There’s a lot to know and understand. You aren’t expected to know all of the material in SP 800-53 for the SSCP exam, but as you move forward in your IT security career, you’ll be much more valuable to employers if you understand these concepts in greater depth.

Comparing Security Control Implementation Methods

In addition to understanding security control goals (such as preventive or detective), you should also understand how controls are implemented. The three primary methods listed in the SSCP objectives are administrative, technical, and physical. [Figure 9-4](#) provides an overview, and the following sections explain these methods in more depth.

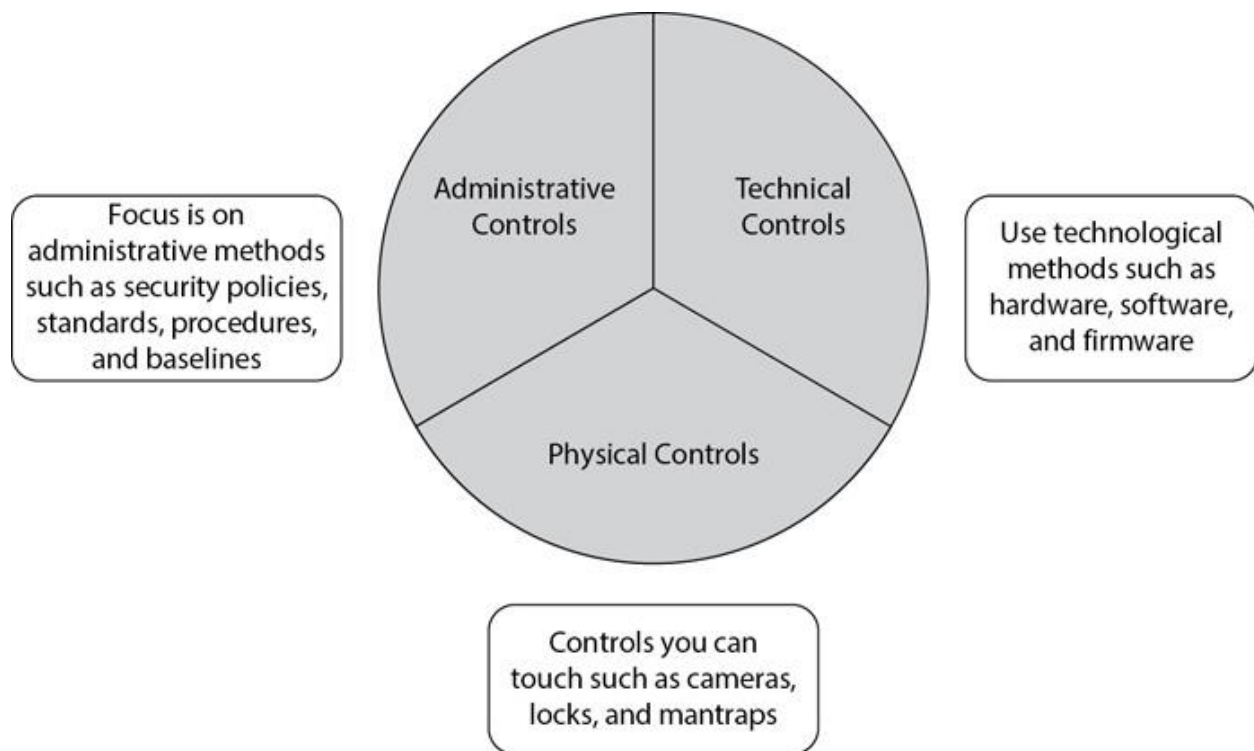


Figure 9-4 Control classes: administrative, technical, and physical

NIST SP 800-53 organizes controls into 18 families. In past versions, NIST attempted to match each of these families with specific control methods or classifications (such as management, technical, and operational). However, many controls apply to more than one control classification, so the current version has dropped that practice. [Table 9-1](#) lists the 18 families. SP 800-53 lists and describes multiple controls within each of these families.

| Family | Family |
|---------------------------------------|---------------------------------------|
| Access Control | Media Protection |
| Awareness and Training | Physical and Environmental Protection |
| Audit and Accountability | Planning |
| Security Assessment and Authorization | Personnel Security |
| Configuration Management | Risk Assessment |
| Contingency Planning | System and Services Acquisition |
| Identification and Authentication | System and Communications Protection |
| Incident Response | System and Information Integrity |
| Maintenance | Program Management |

Table 9-1 Control Families

Administrative Security Controls

Administrative controls (sometimes called management or operational controls) focus on the management of risk and the management of IT security. They include security policies, standards, procedures, guidelines, baselines, and other administrative practices that an organization uses to implement security. The following list shows some examples:

- Written documents such as security policies and standards that provide direction to employees
- Step-by-step procedures for routine operations such as performing backups and verifying proper operation of fire prevention and detection equipment
- Baseline templates to provide a starting point for security controls
- Clear processes for configuration, change, and patch management
- Tests and assessments such as risk assessments and vulnerability assessments
- Plans (such as incident response plans) that help an organization respond to incidents
- User awareness and training programs to help users take an active role in security



NOTE [Chapter 7](#) covers risk assessments and [Chapter 8](#) covers vulnerability assessments. Other concepts mentioned in the previous list are discussed later in this chapter.

Technical Security Controls

Technical (sometimes called logical) controls are those that are implemented through technical methods such as hardware, software, or firmware components. The following list shows some examples:

- Identification and authentication controls that allow users to prove their identity
- Technical password policies that enforce the written security policies, such as password aging settings that force users to change their password periodically
- Access controls that ensure only authorized entities have access to systems and data
- Auditing and accountability controls that track activity and detect unauthorized access

- Network access controls that restrict access to networks to clients that do not meet minimum security requirements
- Encryption protocols that scramble data and provide confidentiality
- Session timeout settings that lock a user's system or close a web browser session after a period of inactivity

Physical Security Controls

Physical security controls refer to the controls that you can touch. Many documents group them together with operational controls, and that is accurate. However, within the context of security, physical security controls are so important that it's worth listing them separately. NIST SP 800-53 identifies a family of Physical and Environmental Protection controls that falls into this class of controls. These include any controls that control, restrict, and/or monitor physical access and provide environmental protection (such as temperature and humidity controls). The following list shows some examples:

- Perimeter barriers such as fences and bollards outside the building and locked doors inside the building
- Mantraps that control how many people can enter the restricted area at a time and help prevent tailgating
- Cameras that can record all activity for any area of interest
- Lighting to illuminate potential areas where intruders may try to break into a building
- Access badges that personnel wear while inside secure areas
- Primary and backup power sources, such as generators and uninterruptible power systems
- Heating, ventilation, and air conditioning systems that control temperature and humidity

When planning physical security controls, it's important to consider personnel safety and always prioritize personnel safety over IT security. As an example, electronic door locks require power to open. If a building loses power, exit doors using electronic door locks should fail in such a way that the doors unlock. This allows personnel within the building to exit. In contrast, if the doors failed in such a way that they remained locked, a fire that affected power could trap personnel inside the room.

Combining Control Goals and Classes

When taking security tests such as the SSCP exam, you're often required to know both the control goal (such as preventive, detective, or corrective) and the control class (administrative, technical, and physical). For example, a locked door to a room is a preventive, physical security

control. It is a preventive control because it prevents individuals from going into the room, and it is a physical security control because you can touch the locked door.

Think about an audit log that security professionals use to identify security violations. How would you classify it? It is a detective, technical security control. It is a detective control because security professionals use it to detect security violations after they've occurred and been recorded in the log. It is a technical control because it uses technical means to record events.

It's impossible to list all the possible controls along with their combinations in this chapter. However, if you understand the control goals (such as preventive, detective, and corrective) and you understand the control, you'll find that it is relatively easy to match the two.

As an example, what is the goal of bollards placed by the entrance to a building? Bollards are vertical metal or cement poles. They act as barricades to prevent someone from driving a vehicle through the entrance, so they are preventive controls.

Similarly, if you understand the control classes (administrative, technical, and physical) and you understand the control, you'll find that it is relatively easy to match the two. Bollards are a physical security control because they use physical methods (something you can touch) for security. In other words, bollards are preventive, physical security controls.

Exploring Some Basic Controls

There are thousands of different security controls, and most (if not all) are documented in NIST SP 800-53. Obviously, this chapter cannot include all the controls covered in NIST SP 800-53. However, the following sections introduce some common controls that you may see on the SSCP exam.

Hardening Systems

Hardening a system is the practice of making it more secure from its default configuration. Administrators harden systems by following basic security practices such as those in the following list:

- **Remove or disable unused protocols** If a protocol is not being used, it should not be installed. Many operating systems include default protocols. Administrators should verify whether a system needs all installed and enabled protocols, and then either remove or disable any unused protocols. Every installed protocol increases the attack surface of the system. In contrast, by removing unused protocols, administrators decrease the attack surface. Attackers cannot attack a protocol that is not installed and running on a system. However, in some cases attackers can enable a protocol, so just disabling it doesn't provide the desired level of protection.
- **Remove or disable unneeded services** This is similar to removing or disabling unused protocols. Many systems include and enable default services. However, not all systems need these default services. Administrators should verify the services that are needed and disable unneeded services. Attackers cannot attack a disabled service. In some cases, administrators

disable protocols by disabling the related service. However, if malware can get administrative access to a system, it can enable the protocol by starting the service. As an example, some malware has used Trivial File Transfer Protocol (TFTP) to infect other systems on a network. If the related TFTP service was disabled, the malware simply enabled the service, which enabled the protocol.



EXAM TIP You can reduce the attack surface of a system by removing or disabling unused protocols and services. With these protocols and services removed or disabled, an attacker has fewer opportunities to attack a system.

- **Change defaults** If the system starts with any default configurations that represent a risk, change them before putting the system into service. The classic example is default passwords. For example, wireless routers included a default password for years that, if not changed, was easily used by attackers to access the device and modify key settings.
- **Keep systems up to date** Both operating system software and application software develop bugs. Attackers can exploit many of these bugs, so it's important to keep systems up to date. Vendors regularly release updates, patches, and fixes to resolve these bugs. These updates should be tested and applied as soon as possible.
- **Enable firewalls** Many systems have built-in firewalls, and many security suites provided by AV software companies include third-party firewalls. Enabling host-based firewalls reduces the possibility of successful attacks.
- **Include AV software** Threats from malware are constant today, and unprotected systems will quickly become infected. AV software protects systems from a wide variety of malware threats. It's also important to keep the AV software signatures up to date.

Equifax Data Breach

Equifax suffered a significant data breach in mid-2017 that resulted in the exposure of sensitive data for as many as 143 million people. Equifax is a consumer credit reporting agency. It collects and stores information on more than 800 million consumers, including data related to their credit history, and it provides this information to businesses. The data breach exposed personal data, Social Security numbers, birth dates, and addresses. Some of the records included driver license numbers and more than 200,000 records included consumers' credit card credentials.

However, the attack was preventable if Equifax had simply updated servers with a patch provided at least two months before the attack. The Apache Software Foundation released a patch for the Apache Struts web application in March. Criminals successfully attacked the server two months later in May. As many cybersecurity experts reported, this attack would not have been possible if administrators managing these Equifax servers had timely patched them.

Policies, Standards, Procedures, and Guidelines

Policies are high-level documents used to provide guidance to members of an organization. A policy provides direction to employees and is authoritative in nature. Standards document criteria such as a proven norm or method. They are typically external to an organization but can influence the organization's policies, guidelines, and procedures. Guidelines provide recommendations to members of an organization but aren't mandatory or authoritative in nature. Procedures provide individuals with specific action steps to accomplish tasks. [Figure 9-5](#) shows the relationship among policies, standards, procedures, and guidelines.

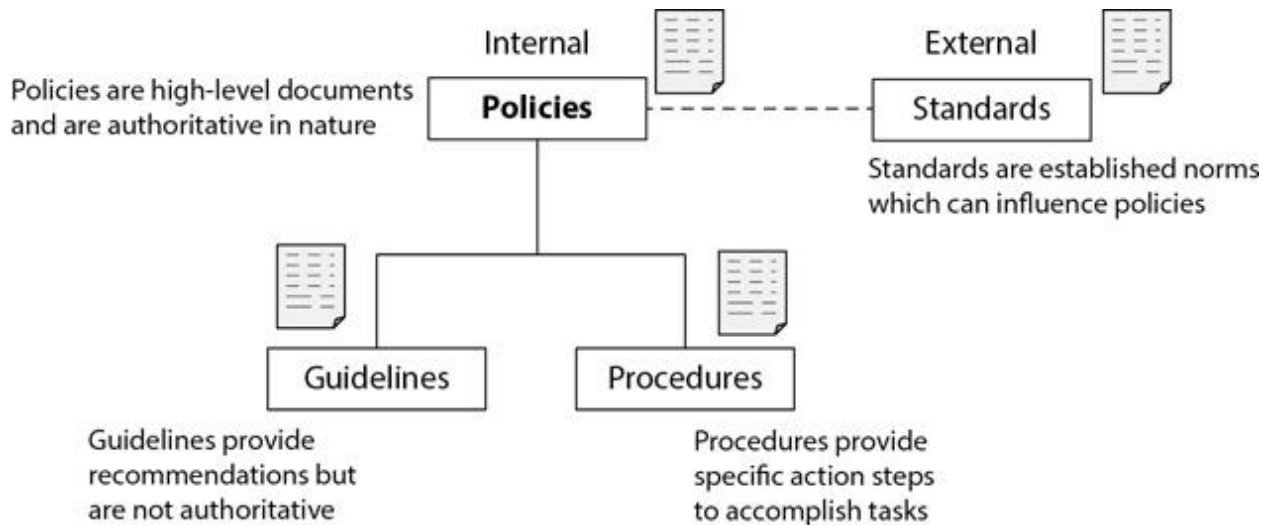


Figure 9-5 Relationship among policies, standards, procedures, and guidelines



EXAM TIP Policies are authoritative in nature and provide overall direction. Standards are established norms, which can influence policies. Guidelines and procedures are derived from policies. Guidelines provide recommendations but are not authoritative.

The following are some examples of policies:

- **Security policy** Many organizations have security policies that define the overall goals of security. A security policy defines what security means to an organization and outlines basic security requirements. Personnel within an organization create standards and procedures to meet the goals and requirements specified within a security policy.



NOTE [Chapter 12](#) covers security policies in more depth.

- **Acceptable use policy** This lets users know what they can and cannot do with IT systems owned and controlled by the organization. For example, some organizations allow users to send and receive personal e-mails with their organizational account, while other organizations restrict this.
- **Backup policy** Many organizations use a backup policy to define the data that should be backed up and how long backups should be retained. A backup policy would also include a statement mentioning that a copy of backups must be kept in an offsite location to protect against a potential disaster.

An organization may use configuration control procedures to dictate the configuration of systems. For example, the organization may identify specific baseline configurations for servers and baseline configurations for desktop systems. At any time, the organization can compare the current configuration of these systems against the baseline to ensure that the system is still meeting the requirements of the standard.

An organization can often choose what standards to follow, so standards aren't necessarily authoritative in nature. For example, the International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 20000 (previously known as BS 15000) is a standard based on the Information Technology Infrastructure Library (ITIL). An organization can use this as an external standard when creating its security policy. Management would ensure that the security policy addresses relevant elements of ISO/IEC 20000. Similarly, organizations can use NIST special publications in the SP 800 series as standards when creating or updating their security policies.

Guidelines are more generic in nature. For example, a guideline related to backups may recommend backing up all databases. Notice this is different from having an authoritative backup policy that states that all customer databases *must* be backed up.

Policies are implemented with procedures. For example, written procedures identify how users acknowledge the acceptable use policy. It may specify that employees need to review and sign the policy when hired and annually thereafter. Backup procedures would itemize the steps required to back up different types of data, such as a database or individual files, and the steps required to restore backups.

Response Plans

Many organizations prepare for disasters by creating different types of plans. A business continuity plan (BCP) helps an organization prepare for emergencies that can interrupt the mission of the business. The BCP uses a business impact analysis (BIA) to identify critical functions for an organization and then includes plans to keep these critical functions operating. A disaster recovery plan (DRP) identifies steps to recover critical systems after a disaster.

Security incidents can cause a significant amount of damage, so organizations often create incident response plans to prepare for them. An incident response plan provides detailed guidance for personnel responding to security incidents.



NOTE [Chapter 12](#) covers disaster preparation in more depth, including the differences between a BIA, BCP, and DRP. [Chapter 13](#) covers incidents and incident response in more depth.

Change Control and Configuration Management

Change control helps prevent unintended outages from occurring as a result of changes. There are countless stories of well-meaning technicians and administrators modifying systems to resolve one problem, but inadvertently creating another problem. In the worst-case scenarios, they end up disabling an entire system or service.

Instead of making a change on impulse when a change is required or desired, administrators submit a change request. A group of experts on a change review board reviews the request and tries to predict potential problems from the change. They approve the change if they predict it will not cause any problems and reject the change if they determine it will be detrimental to operations.

Configuration management ensures that information about a system's configuration is available and helps ensure that similar systems are configured similarly. Many organizations deploy systems using baseline configurations with images. In addition to ensuring that the system is deployed securely, configuration management ensures the system can later be checked and compared to the baseline to ensure that it hasn't been reconfigured.

Administrators develop and maintain system and security control documentation as a part of their daily tasks. The documentation can be contained within a change control system, a configuration management system, or a combination of both. If systems are compromised in an attack or from a disaster, this documentation is invaluable when returning the system to the state it was in prior to the incident.

Some organizations combine change control and configuration management into the same process or control, while larger organizations sometimes separate it into separate processes.



NOTE [Chapter 10](#) covers change management and configuration control in more depth, including the use of imaging to implement baselines.

Testing and Implementing Patches, Fixes, and Updates

As mentioned previously, vendors regularly release patches, fixes, and updates. Basic hardening practices ensure that systems have all relevant patches installed, including patches for operating systems and applications.

The systems development lifecycle (SDLC) is often used when planning, creating, testing, and deploying computing systems and computer applications. The SDLC includes five phases: design, implementation, maintenance, planning, and analysis. During the maintenance phase, personnel apply patches, fixes, and/or updates.

Some small organizations configure systems to apply all patches as soon as the vendor releases them. However, midsize and large organizations implement a patch management process. When a vendor releases a patch, administrators in these larger organizations often take the following steps, described in the following sections:

1. Evaluate the patch to determine if it is needed.
2. Test the patch to verify that it doesn't cause problems.
3. Apply the patch to systems that need it.
4. Audit systems to verify that the patch is applied.
5. Document the patch in a change management system.

Evaluating Patches

It's important to note that it isn't necessary to apply all patches to all systems. For example, a patch for a Microsoft Server product that implements a fix for Internet Information Services (IIS) is relevant for a server running the IIS web server application. However, it is not necessary or desired for a system not running IIS. Administrators first evaluate patches, fixes, and updates to determine what systems need them.

Testing Patches

Unfortunately, some fixes cause problems elsewhere. In the worst-case scenario, systems no longer boot after administrators apply a patch. If this happens to one or two desktop computers, it can be inconvenient, but if it happens to a critical server or hundreds of desktop computers, it can be catastrophic. Because of this, organizations test the patches before applying them to critical systems or to large numbers of computers.

Ideally, the patch won't cause any problems. However, many patches have caused problems in the past. In one case, a hardware vendor used a specially configured image on a huge quantity of computers that it sold. Using images is common, but due to how the vendor created this image, an update caused these computers to go into an endless reboot cycle. Organizations that owned computers with this image had to recover each system individually after applying the patch. However, organizations that tested the patch on a single computer discovered the problem and did not apply the patch to other computers.

Some problems created by patches aren't as extreme. For example, one patch prevented audio from working on a specific brand of computers. A smaller organization that I know of configured its computers to receive patches automatically (without testing). One of the patches disabled all audio on about 50 of its 80 computers. However, not everyone complained right away. Instead, the IT team (one full-time IT professional and one part-time intern) started getting random complaints. It took the team a couple of weeks to realize how many computers the patch affected and identify the source of the problem.

Some organizations use IT administrator computers for testing. In other words, after determining that systems need a patch, they deploy it to computers used by IT administrators. If the patch causes any major problems, administrators can troubleshoot and recover from the problem more easily than a typical end user could.

Applying Patches

If administrators determine that the patch is needed and doesn't cause any unacceptable problems, administrators apply the patch. Many organizations use server applications to automate the deployment of the patches. Microsoft provides Windows Server Update Services (WSUS) as a free download to manage Microsoft updates. System Center Configuration Manager (SCCM or ConfigMgr) is a server application that organizations can purchase from Microsoft. ConfigMgr includes the features of WSUS plus many more, such as the ability to inventory systems and push out applications. Additionally, third-party applications are available that organizations can use to manage updates.

After applying the patch, these automated tools verify that the systems have accepted it. In some cases, a system will block the installation of the patch. For example, a firewall can block patches if it isn't configured with an exception to accept them. It's also possible for an existing application to cause a conflict and block the patch.

Auditing Systems

Automated deployment applications such as WSUS provide administrators with a listing of all computers that have accepted the patch and a listing of computers that remain unpatched. Additionally, they periodically audit the systems to ensure that they all have approved patches. If a system is rebuilt or a patch is removed, these applications discover the problem. Administrators can configure WSUS and similar third-party applications to redeploy the update when they discover an unpatched system.

Unpatched systems are often vulnerable to attacks, so many vulnerability scanners check systems for patches. A vulnerability scanner can't deploy a patch, but it can provide an alert to security administrators that a patch is missing.

Documenting Patches

As mentioned previously, most organizations use a change management process to document any type of changes. This includes changes from patches, fixes, and updates. When administrators determine that a system needs a patch, they follow the change management procedures to apply the patch.

The change management system helps prevent unintended outages due to unapproved changes. It also provides documentation on approved changes. This documentation can be extremely useful if administrators need to rebuild a system from scratch.



EXAM TIP Whenever possible, it's best to test patches before applying them. As a follow-up step, administrators should audit systems to verify that patches are applied. Additionally, personnel should document applied patches in a change management system.

Endpoint Device Security

Many organizations implement controls to protect against risks associated with endpoint devices. In this context, endpoint devices refer to devices used by end users. They include mobile devices (such as smartphones and tablets), USB devices, laptops, and desktop computers.

[Chapter 3](#) includes a section on protecting mobile devices, such as smartphones and tablets, and some of the risks that can occur if they are lost. Common controls with these devices are password protection, encryption of data, remote wipe, and GPS capabilities. If devices are lost or stolen, these controls help reduce risks.

Other times, the risks are due to built-in cameras or phones transmitting audio without the user's knowledge. Because of these risks, many organizations prohibit the use of mobile devices in certain areas of an organization. Some organizations have mini-lockers close to the entrances to secure areas and require employees to lock their mobile device in the locker before entering, and retrieve it when they leave.

Bring Your Own Device

A trend in recent years is that employees want to bring their personal smartphones and tablets to work and connect into the work network. Bring your own device (BYOD) policies allow users to connect their personally owned devices to the organization.

Some of the benefits to the organization include reduced costs because they don't need to provide these devices to employees, and improved productivity for workers because they remain connected to the work network even when they aren't at work. For example, employees can configure their devices with their work e-mail and be available to answer e-mails during nonwork hours. Of course, this availability brings with it the challenge of maintaining a work-life balance.

Unfortunately, BYOD policies present so many challenges that some administrators have nicknamed it as bring your own disaster. Keeping these devices up to date is one challenge. Organizations have processes in place to keep organization-owned devices up to date with current patches and updates. However, pushing updates to employee-owned devices presents other problems. First, because updates might change how the device operates or how applications function, employees might not want to update their devices to a new operating system or with a new patch. Second, there are many different types of mobile operating systems (such as Apple's iOS and the Android operating system).

Android is the most popular operating system used worldwide on smartphones and tablets. However, it isn't run on a single brand of device. The Android operating system is open source, and multiple companies use it on their devices. For example, Google uses Android on many mobile devices that it sells. Samsung is one of many other companies that uses the Android operating system on multiple smartphones and tablets. However, because Android is open source, many mobile device vendors modify it for their devices. If the IT department is tasked with keeping BYOD devices up to date, they need to be able to manage all the different versions of Android. This includes all the Google devices, Samsung devices, and the different devices manufactured by other companies using Android. This requires knowing which devices users have and then keeping up with the updates for all of them.

Another challenge is managing the data. The user owns the device. Who owns the data? Also, how is the organization's data protected when it is on the user's device? Organizations often implement BYOD policies to address these challenges and include elements of the BYOD policies in acceptable use policies.

Choose Your Own Device

Choose your own device (CYOD) is an alternative that many organizations have adopted. Management creates a list of approved mobile devices that employees can connect to the network. Employees that have these approved devices (or purchase one of these devices) can connect them to the network.

This is similar to a BYOD policy, but it limits the number of device types that the IT department needs to manage. However, because the devices are still owned by the employees, the organization needs to create an acceptable use policy and to address issues such as who owns the data and how it is protected.

Using Corporately Owned Devices

Another alternative to BYOD and CYOD is corporate-owned, personally enabled (COPE) devices. These are smartphones and tablets that the organization purchases and issues to employees for work purposes. A huge benefit of COPE devices is that an organization can control what devices it purchases. For example, instead of managing different devices owned by employees, an organization can purchase only one or two models and issue these to employees.



EXAM TIP A bring your own device (BYOD) policy allows employees to connect any personally owned device to the organization’s network. Choose your own device (CYOD) is similar to BYOD, but employees can only connect devices on a preapproved list. Organizations purchase the devices and issue them to employees when using a corporate-owned, personally enabled (COPE) policy.

Mobile Device Management

Organizations typically have strong patch management processes in place, as discussed earlier in this chapter. Similarly, they have automated methods to configure systems, push out applications to systems, and update these applications when necessary. Historically, technical methods only accessed internal computers owned by the organization. However, with the explosion of BYOD, CYOD, and COPE policies, organizations now have a need to implement similar methods for smartphones and tablets.

Mobile device management (MDM) solutions provide these services. As an example, Microsoft’s ConfigMgr (mentioned in the “Applying Patches” section earlier in this chapter) includes management services for some mobile operating systems, including iOS and Android. Administrators use ConfigMgr to deploy updates, applications, and configuration settings.

MDM solutions typically have an enrollment process to register the mobile devices with the MDM solution. When using a COPE policy, administrators may manually enroll the devices. However, many MDM applications allow administrators to configure settings so that users can enroll their own device without administrator intervention. These processes do require users to authenticate while enrolling their devices.

Operating system developers (such as Microsoft, Apple, and various Android version developers) create application programming interfaces (APIs) that MDM developers can integrate within an MDM solution. These APIs simplify the way that mobile devices can be managed.

It’s worth repeating that hardware vendors, such as Samsung and Motorola, typically modify the open source Android operating system to meet their own needs. In other words, ConfigMgr (or

any MDM system) cannot necessarily modify any device running Android, but instead needs APIs specifically designed for the devices.



TIP You can visit this site for a list of Android devices: <https://www.bluestacks.com/blog/app-reviews/archive/list-of-android-devices.html>. Scroll down to the “Phones” or “Tablets” sections and you’ll see that devices are running many different operating system versions. Additionally, these vendors can update the operating system at any time. This becomes extremely challenging for administrators to manage without MDM software.

USB Devices

USB devices, such as USB flash drives and USB hard drives, are especially troublesome to many organizations. It’s easy to copy a significant amount of data to these devices, and it’s also easy to lose them. Lost USB devices can result in a significant amount of data loss or data leakage. Some organizations require the encryption of all data stored on any type of device, including USB devices, to protect against the loss of confidentiality.

Of course, another risk related to USB devices is the risk of transferring malware from system to system. If a USB device is infected with a virus, that virus can easily infect a system as soon as a user plugs the USB device into a computer. Similarly, if a system is infected with a virus, the virus can detect when a user plugs the USB device into the computer and infect the USB device.



EXAM TIP Two significant risks related to USB devices are data leakage and infection with malware. Losses due to data leakage can be mitigated by encrypting data at rest. Malware infections can be mitigated with antivirus software.

Virtualization

Virtualization was mentioned in [Chapter 4](#) with cloud computing. It is becoming more and more common today, and not just with cloud computing. Many organizations are realizing that they can get some significant benefits from virtualizing servers.

For example, an organization may have 500 servers hosted in a data center. The organization could instead purchase 50 very powerful servers, with each server hosting 10 virtual servers. Although the host servers are more expensive initially, they cost less in the long term. It’s cheaper to provide power and cooling to 50 servers than to 500. Virtualizing a data center also requires less space and less physical security requirements.

One of the biggest risks with virtual machines (VMs) comes from VM escape attacks. VM escape attacks give an attacker access to the physical system from within a VM. In other words, the attacker “escapes” from the VM. When successful, the attacker can gain full control over all the VMs on the host in a successful VM escape attack. The primary protection against VM escape attacks is to keep the host system and all VMs up to date with current patches.

Thin Clients

The price of computers is low enough that it’s common for each user to have his or her own computer. However, in some cases, it’s better to use a thin client to reduce some risks. A thin client has very little hardware and only minimal software installed on it. It connects over a network to a server for all its data and applications. If the thin client is stolen, no data is lost. In contrast, a computer with an operating system and applications is a thick client.

Virtual desktop infrastructure (VDI) is one method used with thin clients. Users connect to the server and access a VDI that includes an operating system and all the applications needed by the user. Because the VDI is stored on the server, the clients do not need much processing power. Some VDIs can serve desktops to users running mobile devices.

The primary consideration when using a VDI is deciding on persistent or nonpersistent desktops. A persistent desktop will save all user changes and require more resources to host different desktops for each user. A nonpersistent desktop will not save any user changes. Each user will see the same desktop each time they log on.

Application Whitelisting and Blacklisting

Application whitelisting includes a list of approved applications. When implemented, it prevents users from installing or running any applications that are not on the list. Blacklisting includes a list of nonauthorized applications. It prevents users from installing or running any application on the list. Some MDM applications support both whitelisting and blacklisting. Administrators create a list and configure the MDM application to use it. The MDM application can then scan devices to ensure they are in compliance.

Apple maintains the App Store, which is effectively a massive whitelist. Apple employees review all apps before publishing them in the App Store. Unless users jailbreak their Apple devices to bypass security, they can only install apps from the App Store. An exception to this is that an organization can create apps for its employees and configure Apple devices to accept these apps.

Endpoint Encryption

Many devices support whole disk or full device encryption. This helps prevent attackers from accessing data on stolen or lost devices. As an example, the iOS operating system encrypts all data on the device. This includes text messages, photos, contacts, notes, and more. The data is encrypted automatically as long as the device includes a screen lock, which the iOS enables

automatically. Similarly, Android devices support encryption and automatically encrypt user data.

Jailbreaking and Malware

Jailbreaking bypasses security and grants the user root-level access to the iOS operating system. It is similar to rooting a device running an Android operating system. It voids the warranty, so it isn't commonly done in most countries. Interestingly, it is common in China.

Free jailbreak tools also install apps for Chinese app stores such as Taig and Cydia. These app stores host apps similar to Apple's App Store. However, the owners of these alternative app stores receive the revenue instead of Apple.

Alternative (non-Apple) sites have hosted pirated software and software that included malware so they do not provide the same reliable whitelisting services that you'll find on Apple's App Store. As an example, WireLurker is malware that targets jailbroken Apple devices. It collects sensitive information (such as user's call logs, phone book contacts, and more) from the user's device. Users have downloaded apps infected with WireLurker more than 350,000 times from non-Apple sites such as Taig and Cydia.

WireLurker also infected some non-jailbroken devices via infected Apple Mac computers. The Macs were infected from apps downloaded from another alternative site, the Maiyadi app store. When users connect their iOS devices to the infected Mac using a USB cable, the malware attempts to install itself on the iOS device.



TIP Apple previously had a backdoor to encrypted data, which represented a potential risk. Apple personnel could unlock any device with a key they controlled. However, since the release of iOS 8, Apple no longer includes this backdoor key. This ensures iOS 8 isn't accessible to potential attackers too. Similarly, Android devices create the encryption key on the device, and this key isn't accessible by the manufacturer.

Trusted Platform Module

Many laptop computers support disk encryption using a Trusted Platform Module (TPM), which is a hardware chip embedded in the computer's motherboard. If the laptop includes the TPM, users can enable it via the operating system. However, if the laptop didn't ship with a TPM, you'd need to replace the motherboard to add one. The TPM stores cryptographic keys, which the operating system can use to encrypt and decrypt data.

The TPM includes a unique RSA asymmetric key (called an endorsement key), which the manufacturer embeds into the TPM before shipping it. When users enable the TPM, it creates a storage root key, which encrypts application keys. Applications (such as Microsoft BitLocker)

create application keys to encrypt an entire disk. Note that it doesn't encrypt the entire disk drive. Instead, it partitions the disk drive into separate disks. The BitLocker partition remains unencrypted, but BitLocker encrypts the other partition where the data resides.

Sandboxing for Secure Browsing

Many web browsers (such as Google Chrome and Apple Safari) support sandboxing, which ensures that web pages and web applications have limited access to your system. Think of a sandbox in a children's playground. Parents can place their children in the sandbox and let them play. They know their children will remain safe as long as their children stay in the sandbox, which is isolated from any risks outside of the sandbox.

Similarly, web browsers run within isolated areas of a computer, which function as a virtual sandbox. The sandbox prevents web applications from accessing your personal files or hardware such as your webcam. Additionally, web browsers have limited permissions within the sandbox. These limited permissions are sometimes called protected mode or low integrity mode. Even if a user visits a malicious website that downloads malicious code, the malicious code isn't able to do any damage without first breaking out of the sandbox.

User Awareness and Training Programs

User awareness programs attempt to help users within the organization understand security procedures and practices. These programs help users understand their responsibilities related to security, recognize risky behaviors, and change their behavior if they are engaging in unsafe computing practices.

For example, without security training and education programs, many users don't understand the tactics used by social engineers and may readily give up important information to them. Similarly, users may not understand the risks of jailbreaking a device and installing apps from other app stores. Training topics are sprinkled through this book to stress the importance of awareness and training programs.

Understanding Fault Tolerance

A single point of failure is any component within a system that may cause a significant system outage if it fails. Fault-tolerant controls help eliminate outages from single points of failure by adding redundancies. If a system develops a fault, the system can tolerate the fault and continue to operate. Organizations often add fault tolerance for disks, servers, connections, and sites.



NOTE This section covers fault tolerance for disks, servers, and connections. [Chapter 12](#) covers alternative sites and how they can provide fault tolerance for entire locations.

Fault Tolerance for Disks

A Redundant Array of Independent Disks (RAID) uses multiple disks to improve the performance of the disk subsystem, and most RAID systems provide fault tolerance for the disks. You can implement RAID with software or hardware. For example, some operating systems have built-in support for software RAID. Some motherboards have built-in support for hardware RAID, and you can purchase external hardware RAID disk subsystems.

Whenever possible, use hardware RAID to get better performance. In software RAID, the operating system provides the services to support the RAID and that can slow down overall system performance. Some of the common RAID configurations are RAID-0, RAID-1, RAID-5, RAID-6, and RAID-10.

RAID-0

RAID-0 uses two or more disks in an array to improve the performance of both read and write operations. However, RAID-0 does not provide any fault tolerance. If one drive in a RAID-0 disk subsystem fails, the entire array fails.

RAID-1

RAID-1 uses two disks in a mirror configuration, as shown in [Figure 9-6](#). Each file is written completely to each drive in the array. If one of the drives ever fails, the other drive still includes complete copies of all of the files.

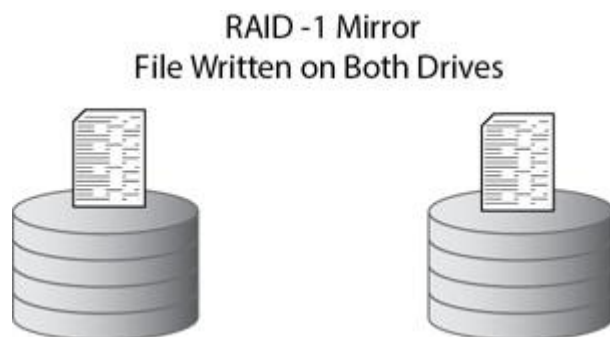


Figure 9-6 RAID-1 mirror

Most RAID-1 disk subsystems also include the ability to read from both drives at the same time. This improves read performance because the disk controller can read half the file from one drive while simultaneously reading the other half of the file from the other drive.

Hardware RAID-1 disk subsystems can typically recover automatically when a drive fails. For example, if one of the drives fails, the system recognizes the failure and stops using the failed drive. However, because all the data is on the drive that didn't fail, the system continues to operate without interruption. Of course, if one of the drives fails, the array is no longer providing fault tolerance. Administrators should replace the failed hard drive as soon as possible.



EXAM TIP A RAID-1 mirror provides fault tolerance by mirroring data on one drive to a second drive.

RAID-5

RAID-5 uses three or more drives in an array and uses striping with parity. It writes data onto the drives in stripes, and it uses the equivalent of one drive for parity data. If any single drive fails, the array can use the parity information to reconstruct the data and the array continues to operate.

RAID-5 uses either odd or even parity to give the number of bits in a stripe either an odd or even value. For example, if odd parity is used and the number of 1 bits in the stripe is even (such as two 1 bits or four 1 bits), then the parity bit will be a 1 to make the total number of 1 bits odd (such as three or five). If odd parity is used and the number of 1 bits in the stripe is odd (such as one 1 bit or three 1 bits), then the parity bit is a 0 to keep the total number of 1 bits odd.

Consider [Table 9-2](#). It shows the decimal values of 0 through 3 and their binary equivalent represented as data bits. In the second data row (considered a stripe in RAID-5), a decimal value of one (1) has the data bits of 0 and 1, giving an odd number of 1 bits. The parity bit is 0, so the number of 1 bits remains odd. In the fourth stripe with a decimal value of three (3), the data bits are 1 and 1, which is an even number of 1 bits, so the parity bit is a 1.

| Decimal Value | Data Bit (2^1) | Data Bit (2^0) | Odd Parity Bit |
|---------------|--------------------|--------------------|----------------|
| 0 | 0 | 0 | 1 |
| 1 | 0 | 1 | 0 |
| 2 | 1 | 0 | 0 |
| 3 | 1 | 1 | 1 |

Table 9-2 Odd Parity

Now imagine that each of these columns represents a disk drive and each of the rows represents a stripe of data, as shown in [Table 9-3](#). The equivalent of two drives carries the data, while the equivalent of one drive carries the parity bits.

| | Data Bit (2^1) | Data Bit (2^0) | Odd Parity Bit |
|--------|--------------------|--------------------|----------------|
| Stripe | 0 | 0 | 1 |
| Stripe | 0 | 1 | 0 |
| Stripe | 1 | 0 | 0 |
| Stripe | 1 | 1 | 1 |

Table 9-3 RAID-5 with Odd Parity

If any of these drives fails, the system can still calculate the contents of the drive simply by calculating the remaining data. If the remaining data in a stripe has an odd number of 1 bits, the missing data in the stripe is a 0. Similarly, if the remaining data in a stripe has an even number of 1 bits, the missing data in the stripe is a 1.

[Table 9-4](#) shows the same RAID-5 array with a failed drive. Even if you don't know what the data was originally but you do know that the RAID is using odd parity, you can easily calculate the data in each of the stripes. This allows a RAID-5 array to suffer the loss of an entire drive and continue to operate.

| | Data Bit (2¹) | Failed Drive | Odd Parity Bit |
|--------|---------------------------------|---------------------|-----------------------|
| Stripe | 0 | Unknown | 1 |
| Stripe | 0 | Unknown | 0 |
| Stripe | 1 | Unknown | 0 |
| Stripe | 1 | Unknown | 1 |

Table 9-4 RAID-5 with Odd Parity and a Failed Drive



TIP RAID-5 stripes hold more than a single bit in each stripe, but for illustration purposes, it's much easier to show the values as single bits. RAID-5 stripes are typically 64KB.

It's important to stress that while a RAID-5 has the equivalent of one drive devoted to parity, the parity data is not written on a single drive. [Figure 9-7](#) shows a three-drive RAID-5 array with multiple stripes. Notice that each stripe has two drives holding the data, with one drive holding the parity, but the parity data for different stripes is contained on different drives. Early implementations of striping with parity put the parity bit on a single drive, but engineers discovered that the drives performed better when the engineers spread the parity bit across multiple drives.

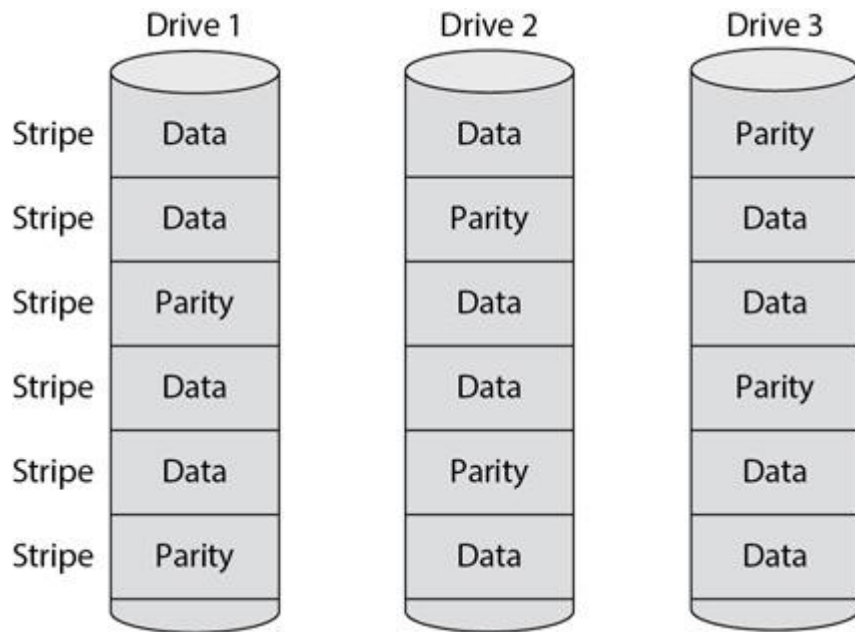


Figure 9-7 RAID-5 striping with parity

Both software-based and hardware-based RAID-5 disk subsystems can automatically detect when a drive fails and automatically continue to operate. However, hardware-based RAID-5 disk subsystems often have extra spare drives in the array. If a drive fails, the system detects the failure, logically removes the failed drive, logically adds a spare drive into the array, and rebuilds it. This all occurs automatically without any human intervention. The array also provides some type of notification to administrators of the failure so that they can replace the failed drive.

RAID-6

RAID-6 is an extension of RAID-5. The formal definition of RAID-6 is a RAID that uses two parity blocks, though the parity blocks can be implemented differently. As a simple example, the second parity block can be a copy of the first parity block but stored on a different disk. The benefit is that a RAID-6 can survive the failure of two disk drives. RAID-6 requires a minimum of four disks, with the equivalent of two disks dedicated to parity.

Many storage vendors recommend the use of RAID-6 over RAID-5 today. This is partly due to the unrecoverable read error (URE) rate of standard SATA drives and the size of current drives. As an example, think of a seven-disk RAID-5 with each disk 2TB in size. While UREs are relatively low, one URE can prevent the reconstruction of a RAID-5 disk subsystem. Further, the chance of a single URE from the 12TB of storage on the remaining six drives after a hard drive failure is as high as 62 percent.

However, with RAID-6, the disk subsystem will survive the failure of a single hard drive and a URE. If administrators quickly replace the single failed hard drive, they will be able to re-create the RAID.

Interestingly, most vendors gloss over the URE problem and instead stress that the RAID-6 can survive the failure of two hard disks. Technically, it can continue to operate with two failed disks. However, the disk subsystem is likely to experience a URE with two failed hard disks, and administrators will not be able to rebuild the array.



TIP More advanced RAID implementations are also available. For example, RAID-10 combines the benefits of a RAID-1 mirror with a RAID-0 striped array and uses at least four drives. This provides both increased performance and increased fault tolerance.

Failover Clusters

A failover cluster provides fault tolerance for one or more servers. A simple two-node failover cluster ensures a service continues to operate even if a server fails. For example, if an organization is hosting a database and wants to ensure that the database is always available, even if a server fails, the organization could use a failover cluster.



EXAM TIP Failover clusters allow a service to continue to operate even if a server fails. A failover cluster provides fault tolerance at the server level.

[Figure 9-8](#) shows a simple two-node cluster with Node 1 active and Node 2 inactive. In this configuration, Node 1 is running the database management software and accesses the database on an external drive. When users query the database, they contact an apparent server for the cluster, and the clustering software routes the traffic to the active node.

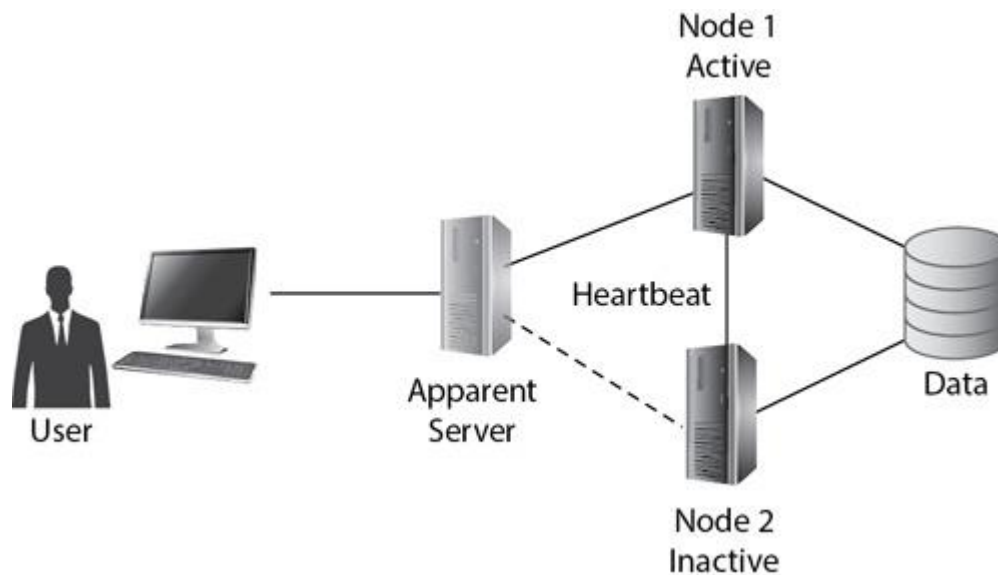


Figure 9-8 Two-node failover cluster



NOTE Node 1 and Node 2 are both physical servers that you can touch. However, the apparent server is only an outward representation of the failover cluster. The clustering software presents this outward representation to users and internally routes all traffic to the active node.

Node 2 in this cluster is currently inactive. It is running the same database management software and has access to the same database on the external drive. More importantly, it is constantly monitoring the health of Node 1 using a heartbeat signal. If Node 1 ever fails, Node 2 will know about it immediately. Depending on how administrators configure the failover cluster, Node 2 may try to restart services on Node 1 or take over services as the active node for the failover cluster.

If Node 2 takes over the services, the cluster software reroutes all requests to the failover cluster to Node 2. At this point, Node 1 is inactive and Node 2 is active. Users may notice a momentary glitch, but overall there is no loss of data.

Although a two-node cluster is easier to illustrate and conceptualize, clusters can be much more complex. For example, a two-node cluster can support two services instead of just one, with each node being active for one service and inactive for the other service. If either of the services fails, the other node can take over for both services. As another example, an eight-node cluster can support multiple services, with each node hosting one or more services and every other node capable of supporting the additional load if any other server fails.

Load Balancing

Load-balancing clusters are another type of cluster, but their primary purpose isn't to provide fault tolerance. Instead, a load-balancing cluster spreads the operations load among multiple servers, increasing their ability to handle larger loads. This allows you to add more and more servers to compensate for increased demand.

As an example, an e-commerce site may be extremely busy during the holiday season. A load-balancing cluster allows the organization to add more servers to handle the additional website load. When the holiday season is over, the organization can remove the additional servers. Load balancing can be managed by a hardware load balancer or by software installed on each of the servers in the load-balancing cluster.

[Figure 9-9](#) shows a sample load-balancing cluster. Each of the web servers hosts the same web applications and websites. When a user accesses the website, the connection is sent to the load balancer. The load balancer then directs the connection to one of the web servers. There are two primary methods that load balancers use to distribute the load: round robin and source address affinity.

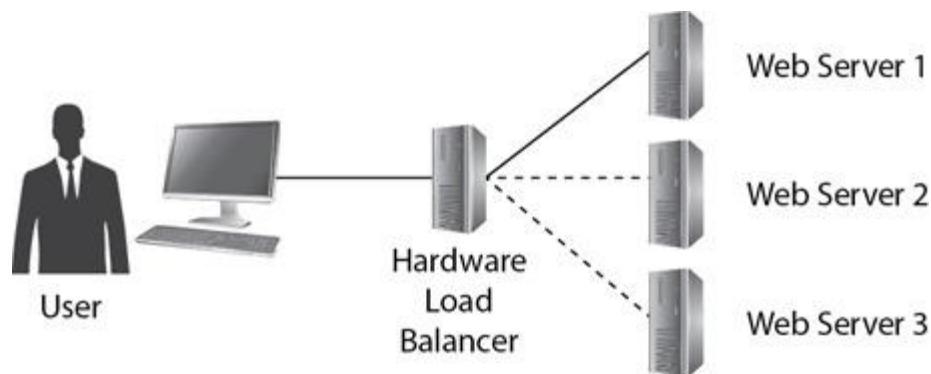


Figure 9-9 Load balancing

In a round-robin configuration, the load balancer directs the first connection to Web Server 1, the next one to Web Server 2, and the next one to Web Server 3. The load balancer continues to redirect the connections in the same rotating pattern. Some load balancers can detect the load on servers and redirect new connections to the least used server.

Source address affinity redirects traffic based on the user's source IP address. Some web applications need users to contact the same server throughout the session. As an example, one application may store data on the server when a user adds items to a shopping cart during a session. In this scenario, the load balancer uses the customer's IP address and source IP address affinity to ensure that the connection is redirected to the same server throughout the session.

Redundant Connections

If an organization must stay connected to the Internet, it can implement two or more connections to ensure connectivity. Similarly, an organization can implement redundant connections to ensure that a main location remains connected to a remote office.

Organizations that use two or more connections often use one high-speed connection for primary use and a second, lower-speed connection as a redundant connection in case the first connection fails. For example, one connection might be a broadband connection, and the second connection might be a slower fractional T1 connection (which is only a fraction of a full 1.54-Mbps T1 connection).

Understanding Backups

The importance of backing up data cannot be overstated. Data will be lost at some point if an organization is using computers. Hardware failures, software bugs, accidental deletions, natural disasters, and malicious attacks can all destroy or corrupt data. However, the difference between a minor inconvenience and a major catastrophe is the existence of a backup. If a reliable backup exists, personnel can restore the lost data. If a backup isn't available, the lost data can result in significant losses.



EXAM TIP Backups are useful only if they can be restored. Organizations typically include test restore procedures in their backup plan. A test restore goes through the steps to restore data from a backup and verifies the usefulness of the backup.

A backup policy identifies the data that is important to an organization, which identifies the data personnel need to back up. Backup policies include retention policies that identify how long the organization should keep backed-up data. These policies are authoritative in nature and can help justify the purchase of backup materials. Tape drives, tapes, backup software, and dedicated backup servers can all add up to significant costs. However, if a backup policy dictates the need for the backups, the policy provides justification for the costs.

In contrast, without a backup policy, administrators will make their own decisions on what to back up based on their own experience. These decisions are often driven by recent data losses. For example, if a high-ranking user recently lost data, then administrators might decide that backing up user data is most important. However, if no user data is permanently lost over a year's timeframe, administrators might decide that backing up user data isn't as important. Also, without a backup policy, administrators may never test these backups until someone loses data. If the backups aren't reliable, this is too late. The data is lost. Worse, when administrators ask for funds to replace worn-out tapes, management might decide not to approve the purchase.

A backup policy also stresses the importance of keeping a copy of backups in a separate geographical location. If backup tapes are kept with the server and a fire in the server room destroys the server, it will also destroy the backups. Even if the backups are stored elsewhere in the same building, a tornado or flood or other disaster can destroy all the backups. However, if the backups are stored offsite, they will be available even if a catastrophe destroys the site.

Traditional backups use full backups combined with either differential or incremental backups. These backups often store the backups on tapes, but the use of disks for backups is becoming more common. Backing up to tape is much slower, but it provides more data storage at an overall lower cost. As an alternative, many organizations use a disk-to-disk-to-tape strategy. This is useful when backing up online databases because the backup can slow down database performance. Administrators back up the data to a disk, taking advantage of the fast disk speed. They then back up the data from the second disk to tape.

Full Backups

A full backup simply backs up the entire contents of the target data. This is useful if there isn't much data to back up or the data doesn't change very often. If the backups are small, you can back up the data every day.

However, most backups are rather large and it takes a long time to back up the data and a long time to restore it. Alternatives to full backups include full/incremental backup strategies and full/differential backup strategies. If you want to minimize the amount of time that it takes to back up the data, you can use a full/incremental backup strategy. If you want to minimize the amount of time that it takes to restore the data, you can use a full/differential backup strategy. Both strategies start with a full backup.

Full/Incremental Backup Strategy

In a full/incremental backup strategy, you back up all the data to start the backup plan. Then on a regular basis you back up only the data that has changed since the last backup. For example, consider [Table 9-5](#) as you read the following list, which describes the backups performed on each day:

| Sunday Morning | Monday Morning | Tuesday Morning | Wednesday Morning | ... |
|-----------------------|---|---|--|------------|
| Full backup | Incremental Backs up changes since full on Sunday | Incremental Backs up changes since incremental on Monday | Incremental Backs up changes since incremental on Tuesday | ... |

Table 9-5 Full/Incremental Backup Strategy

- Full backups are completed early on Sunday morning at 1 A.M. using an automated system.

- At 1 A.M. on Monday morning, the incremental backup backs up only the changes made since the last backup (the full backup on Sunday, in this case).
- At 1 A.M. on Tuesday morning, the incremental backup backs up only the changes since the last backup (the incremental backup on Monday, in this case).
- At 1 A.M. on Wednesday morning, the incremental backup backs up only the changes since the last backup (the incremental backup on Tuesday, in this case).
- Throughout the week, the incremental backups would be about the same size.



NOTE [Table 9-5](#) shows backups until Wednesday for brevity, but administrators would continue with the backups through the rest of the week (on Thursday, Friday, and Saturday). On Sunday the process starts again with a full backup.

If a data loss occurs, you can use the backups from Sunday and then use each of the backups from individual days to restore the data. For example, imagine that a hard drive failure occurs on Wednesday afternoon. You would restore the full backup from Sunday, then the incremental backup from Monday morning, then the incremental backup from Tuesday morning, and then the incremental backup from Wednesday morning.

Some of the recovered data would only be included in the full backup, while some of the recovered data would only be included in the backup from Tuesday, and some would only be included in the backup from Wednesday. Here a few examples:

- One person might create a monthly report but not modify it on Monday, Tuesday, or Wednesday. It would only be in the full backup.
- Another person might create a weekly report every Monday but not modify it on Tuesday, Wednesday, or any other day of the week. It would only be included in the Tuesday incremental backup.
- Help desk personnel might modify files seven days a week. Changes would be included in each incremental backup file (Monday morning, Tuesday morning, and Wednesday morning).



TIP A full/incremental backup strategy minimizes the time needed for backups during the week. However, the recovery can take longer because multiple backups typically need to be restored when recovering the data.

Full/Differential Backup Strategy

In a full/differential backup strategy, you back up all the data to start the plan, and then on a regular basis you back up all of the data that has changed since the last full backup. For example, consider [Table 9-6](#) as you read the following list, which describes the backups performed on each day:

| Sunday Morning | Monday Morning | Tuesday Morning | Wednesday Morning | ... |
|----------------|---|---|--|-----|
| Full backup | Differential Backs up changes since full on Sunday | Differential Backs up changes since full on Sunday | Differential Backs up changes since full on Sunday | ... |

Table 9-6 Full/Differential Backup Strategy

- Full backups are completed early on Sunday morning at 1 A.M using an automated system.
- At 1 A.M. on Monday morning, the differential backup backs up all the changes since the last full backup on Sunday.
- At 1 A.M. on Tuesday morning, the differential backup backs up all the changes since the last full backup on Sunday, including the changes from Monday.
- At 1 A.M. on Wednesday morning, the differential backup backs up all the changes since the last full backup on Sunday, including the changes from Monday and Tuesday.

Throughout the week, the differential backups would steadily increase in size.

If a data loss occurs, you can use the backups from Sunday and then the most recent differential backup. For example, imagine data was lost due to a hard drive failure on Wednesday afternoon. You would need to restore the full backup from Sunday and then restore the differential backup from Wednesday morning.



TIP A full/differential backup strategy takes longer to back up during the week than a full/incremental strategy. However, the recovery time is reduced because a maximum of two backups are needed to recover the data.

Image-based Backups

Image-based backups create a backup of the operating system and all data. They are commonly used to back up virtual machines (VMs) but can also be used on regular computers or servers. The backup creates a single file, called an image, of the system at a point in time.

Administrators often use image-based backups before performing a risky operation such as an upgrade to the operating system or software. If the upgrade causes problems, administrators can re-create the server using the image backup.

Some image-based backups run in differential mode. After the image is captured, the backup system tracks all changes to the image until the image is removed. These differential-based images can significantly slow down the system if they're allowed to track changes indefinitely.

Chapter Review

Security controls include both countermeasures and safeguards, and security professionals use these terms interchangeably. They are the means, methods, actions, techniques, processes, procedures, or devices that reduce the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. The primary goals of security controls are to prevent, detect, and correct loss of confidentiality, loss of integrity, and loss of availability of a system.

Preventive controls attempt to prevent losses before they occur, detective controls attempt to detect violations either as they are occurring or after they've occurred, and corrective controls attempt to reverse the impact of the event. Compensating controls provide an alternative if circumstances prevent an organization from implementing a primary control. Deterrent controls encourage people to decide not to take certain undesirable actions that could result in an incident.

Control methods (or classifications) include administrative, technical, and physical. Administrative controls focus on the management of risk and the management of IT security. Technical controls are implemented using technical methods, such as hardware, software, or firmware components. Physical controls refer to the controls you can touch.

System hardening is the practice of making a system more secure from the default settings and includes removing or disabling unused services and protocols, changing defaults, keeping systems up to date, enabling firewalls, and using AV software.

Policies are high-level documents used to provide overall guidance to an organization. Standards document criteria such as a proven norm or method and are typically external to an organization, but they can influence an organization's security policies. Guidelines and procedures are derived directly from the organization's security policy.

Business continuity plans and disaster recovery plans are examples of response plans that help an organization respond to a disaster. Change control procedures help prevent outages due to unauthorized changes. Configuration management procedures help ensure that systems are configured similarly using baselines. Patch management procedures evaluate patches, test

relevant patches, and apply relevant patches that don't cause unacceptable problems. Automated tools audit systems to ensure patches are applied and change management systems document approved patches.

Endpoint device security typically focuses on mobile devices such as smartphones and tablets. If an organization allows employees to connect their personally owned devices to the organization's network, it will typically create a bring your own device (BYOD) or a choose your own device (CYOD) policy and add elements of the policy to an acceptable use policy. Alternatively, an organization may opt for corporate-owned, personally enabled (COPE) devices to maintain more control of mobile devices. Mobile device management helps ensure that BYOD, CYOD, and COPE devices are kept up to date.

User awareness and training programs raise the overall security posture of an organization by helping users understand their responsibilities associated with security and reduce risky behaviors.

Fault tolerance helps eliminate outages from single points of failure by adding redundancies. Redundant disk systems (such as RAID-1, RAID-5, RAID-6, and RAID-10) allow a system to continue to function even if a disk fails. Failover cluster configurations allow a server to continue to operate even if a server fails. Load-balancing solutions spread the load among multiple servers. Redundant connections allow an organization to stay connected to the Internet or to other locations, even if a connection fails.

Backups ensure that data is not lost even if the original data is lost due to a hardware failure, software bug, accidental deletion, natural disaster, or malicious attack. Backup policies identify what data to back up, how long to retain backups, and the importance of storing backups in separate geographical locations.

Backup types include full backups, incremental backups, and differential backups. Each strategy includes a full backup, which backs up all the important data. A full/incremental backup strategy completes a full backup and then backs up only the data that has changed since the last full or incremental backup. Incremental backups are about the same size and take the least amount of time to back up. However, they can take longer to restore because multiple incremental backups may need to be included in the restoration.

A full/differential backup strategy starts with a full backup and then backs up all the changes since the last full backup without regard to what was backed up during previous differential backups. Differential backups get steadily larger because they are backing up more and more data, requiring more time to complete as the week progresses. A full/differential backup strategy is the quickest to restore because only the full and the last differential backups are needed. Image-based backups create an image of the entire operating system and all applications.

Questions

- 1.** Which of the following provides the best definition of a security control?

A. The means, methods, actions, techniques, processes, procedures, or devices used to prevent attackers from launching attacks on systems

B. A detective method that identifies threats

C. A corrective method that reverses the impact of an incident

D. The means, methods, actions, techniques, processes, procedures, or devices used to reduce the vulnerability of a system or the possibility of a threat exploiting a vulnerability

2. Your organization is considering adding a security control. The CBA indicates that the ALE without the control is approximately 25 percent the cost of the control. Which of the following is the best choice given these results?

A. Implement the control.

B. Do not implement the control.

C. Start a configuration management request.

D. Start a change management request.

3. What are the primary objectives of security controls?

A. Prevent, detect, and correct

B. Prevent, detect, and block

C. Detect, correct, and block

D. Detect, correct, and remove

4. Which of the following security controls attempts to avoid security incidents?

A. Preventive

B. Compensating

C. Corrective

D. Detective

5. What type of security control is an audit log?

A. Technical

B. Corrective

C. Detective

D. Preventive

6. Which of the following security controls can restore a failed or disabled control?

A. Preventive

B. Corrective

C. Detective

D. Deterrent

7. What type of control are procedures to back up and restore data?

A. Administrative

B. Corrective

C. Detective

D. Deterrent

8. Which of the following choices best describes an administrative security control?

A. A control that you can physically touch

B. A control implemented using hardware, software, or firmware

C. A control that focuses on the management of risk and the management of IT security

D. A control that focuses on preventing losses due to risks

9. A computer system records events into a security log, and administrators periodically review the log for security incidents. Which of the following best describes this security control?

A. A preventive, technical security control

B. A detective, physical security control

C. A preventive, physical security control

D. A detective, technical security control

10. Which of the following tasks would administrators complete when hardening a server?
(Select all that apply.)

- A. Removing unneeded protocols
- B. Keeping a system up to date
- C. Disabling firewalls
- D. Adding AV software
- E. Disabling unused services

11. Of the following choices, which provides high-level guidance to employees?

- A. Procedure
- B. Policy
- C. Action steps
- D. Disaster recovery plan

12. What is the overall goal of a change management process?

- A. To slow down changes
- B. To ensure that systems are configured similarly
- C. To enable stakeholders to deny unwanted changes
- D. To reduce unintended outages from unauthorized changes

13. What should administrators do after learning that a vendor has released a patch that is relevant for servers they manage?

- A. Apply the patch.
- B. Test the patch.
- C. Audit systems to see whether the patch is applied.
- D. Document the systems where the patch is applied.

14. Which of the following helps ensure that mobile devices have all relevant patches?

- A. BYOD
- B. CYOD
- C. MDM
- D. COPE

15. Your organization has recently implemented a CYOD policy. Management wants to ensure that devices are kept up to date with current patches and fixes. Which of the following should be implemented?

- A. Use a change management policy.
- B. Create a list of preapproved devices.
- C. Implement an MDM solution.
- D. Implement a COPE policy.

16. What is an important benefit to organizations that use virtual servers? (Choose all that apply.)

- A. VM escape capabilities
- B. Better control of data with cloud computing
- C. Reduction of costs associated with power and cooling
- D. Reduction of costs for physical security

17. You have two disk drives and you want to provide fault tolerance by mirroring the two drives. What should you use?

- A. RAID-0
- B. RAID-1
- C. RAID-5
- D. RAID-6

18. You need to ensure that a service continues to run even if a server fails. What should you implement?

- A. RAID-1

- B. RAID-6**
- C. Failover cluster**
- D. Warm site**

19. Administrators are configuring a group of e-commerce-based web servers in a load-balancing configuration. Customers need to return to the same server each time they visit within a session. Which of the following configurations will support this goal?

- A. Source address affinity**
- B. Round-robin**
- C. Failover cluster**
- D. Incremental**

20. You don't have enough maintenance time during the week to perform full backups, so you decide to implement a backup strategy that takes less time to do backups during the week. Of the following choices, what strategy will minimize the amount of time needed to restore a backup after a failure?

- A. Full**
- B. Incremental**
- C. Full/incremental**
- D. Full/differential**

Answers

1. D. A security control provides the means, methods, actions, techniques, processes, procedures, or devices that reduce the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. You can't prevent attackers from launching an attack, but you can reduce their possibilities of success by either reducing vulnerabilities or reducing the impact of the threat. Security controls can be preventive, detective, and/or corrective, but it isn't accurate to limit a control to only one of these types.

2. B. Based on the results of the cost-benefit analysis (CBA), the most likely choice is to not implement the control. A CBA identifies the cost. If the cost of the control is significantly lower than the annual loss expectancy (ALE) without the control, the cost is justified, but the scenario indicates the ALE is 25 percent the cost of the control (indicating the control cost is four times the ALE). Configuration management refers to the starting point of a system and doesn't use

configuration management requests. You would start a change management request if you planned to implement the control.

3. A. Security controls attempt to prevent, detect, and/or correct losses to confidentiality, availability, or integrity. Some preventive controls attempt to block threats, and some corrective controls attempt to remove vulnerabilities, but neither blocking nor removing is referred to as a primary objective of security controls.

4. A. A preventive security control attempts to avoid security incidents by preventing them. A compensating control provides an alternative when a primary control fails or is unavailable. A corrective control attempts to reverse the effects of an incident and can restore a failed or disabled control. A detective control attempts to detect incidents either as they are occurring or after they've occurred.

5. C. An audit log is a detective security control because it identifies events either as they are occurring or after they've occurred. Technical is a method or class of control. A corrective control acts to reverse the effects or impact of an incident. A preventive control prevents the event from occurring.

6. B. A corrective control can restore a failed or disabled control. A preventive control attempts to prevent incidents and a detective control attempts to detect incidents. A deterrent control attempts to dissuade or deter personnel from trying to circumvent security policies or otherwise cause an incident.

7. B. Procedures to back up and restore data are corrective controls because they take action to reverse the effects of data loss. Administrative is a method or class of control. Deterrent controls attempt to deter personnel from causing a security incident. A detective control identifies events either as they are occurring or after they've occurred.

8. C. An administrative security control focuses on the management of risk and the management of IT security. A physical control is one that you can touch. A technical control is implemented with hardware, software, or firmware. A preventive control focuses on preventing losses due to risks.

9. D. This describes an audit log, which is a detective, technical security control. It is detective because it identifies incidents after they've occurred. It is technical because it is implemented with technology. It doesn't prevent the events, and the log is a digital file, not a physical item that you can touch.

10. A, B, D, E. Some common steps when hardening a server include removing unneeded protocols, keeping a system up to date, adding antivirus (AV) software, and disabling unused services. Disabling firewalls is not a method used to harden a server.

11. B. Policies provide high-level guidance to employees and are authoritative in nature. Procedures include action steps to accomplish tasks. A disaster recovery plan can include one or more procedures to use in response to a disaster.

12. D. The overall goal of a change management process is to reduce unintended outages from unauthorized changes. The intent isn't to slow down changes, but that is often a side effect. A change management process gives stakeholders the ability to propose changes, but several entities within the organization grant or deny changes. A configuration management process ensures that systems are configured similarly.

13. B. Administrators should test patches before applying them to detect potential problems. If the patch is applied as soon as it is released, it can result in unforeseen problems. Systems are audited after the patch is applied to ensure that the patch was successfully applied. Change management and configuration control procedures are used to document patches.

14. C. Mobile device management (MDM) solutions help ensure that mobile devices (such as smartphones and tablets) have all relevant patches. A bring your own device (BYOD) policy allows employees to bring their own devices to work and connect them to an organization's network. A choose your own device (CYOD) policy is similar to a BYOD policy, but employees can only connect devices on a preapproved list. Corporate-owned, personally enabled (COPE) devices refer to mobile devices that an organization purchases and issues to employees.

15. C. A mobile device management (MDM) solution can manage mobile devices that connect to an organization's network. This includes ensuring they are kept up to date with current patches and fixes. None of the other answers ensures that devices are kept up to date. A change management policy would be used to approve changes before they are deployed. A choose your own device (CYOD) policy includes a list of preapproved devices. A corporate-owned, personally enabled (COPE) policy is different from a CYOD policy.

16. C, D. Organizations often use virtualization to reduce costs associated with power and cooling and associated with physical security. Virtualization requires fewer physical servers to power, cool, and physically secure. Virtual machine (VM) escape is one of the biggest risks associated with virtual servers. Virtual servers can be used with cloud computing, but cloud computing reduces control of data.

17. B. RAID-1 provides fault tolerance by mirroring two drives. RAID-0 does not provide fault tolerance. RAID-5 uses three or more drives. RAID-6 is an alternative to RAID-5 and uses four or more drives.

18. C. A failover cluster includes two or more servers (called nodes), and if one server fails, the other server can handle the load for that server. RAID-1 and RAID-6 provide fault tolerance for disk drives. A warm site is a type of alternative site used to provide fault tolerance for an entire location.

19. A. A source address affinity configuration uses the customer's IP address to ensure the customer returns to the same web server during a session. A round-robin configuration directs connections to the servers, one after another. A failover cluster provides fault tolerance, not load balancing. Incremental refers to a backup strategy.

20. D. A full/differential backup strategy takes the least amount of time to restore because it will require only two backups to restore: the full backup and the last differential backup. The scenario says that there isn't enough time to do a full backup during the week, and you can't do an incremental backup by itself. A full/incremental would require you to restore the full backup and each of the incremental backups since the last full backup, which usually requires restoring more than just two backups.

NOTES

CHAPTER 10

Auditing and Management Processes

In this chapter, you will learn about

- The value of auditing to enforce accountability
- Using clipping levels to define acceptable thresholds
- Using audit trails to re-create events leading to security incidents
- Types of audit logs used to create audit trails
- Using security audits to verify compliance
- Using baselines for configuration control
- Preventing outages with change management

Understanding Auditing and Accountability

[Chapter 1](#) introduced the AAAs of security: authentication, authorization, and accounting. When used together, they ensure that only authorized entities (such as users and applications) can access systems or data and that a record of such activities exists. [Chapter 2](#) presented access controls, including the three primary authentication factors, and different access control models. Auditing provides the accounting component, which tracks and records individual actions.

There are two primary methods used to perform auditing. They are different processes, but both methods are part of an overall auditing strategy:

- **Auditing activity through logs** Audit logs record activity and can be inspected by auditors at any time to reconstruct events. Logs may be on servers, firewalls, routers, or any other type of device.
- **Auditing activity through an inspection process** Periodic security audits inspect an organization to determine whether the organization is following required policies and procedures. For example, a security policy may state that inactive accounts should be disabled immediately and deleted within 60 days after an employee leaves. An inspection can detect if this policy is being followed.



NOTE The “Exploring Audit Logs” section later in this chapter covers audit logs and the “Performing Security Audits” section covers the second type of auditing using an inspection process.

Holding Users Accountable with Audit Logs

If the organization uses strong authentication practices, administrators, security personnel, and auditors can use the audit logs for accountability. As an example, [Figure 10-1](#) shows a Security log entry from Windows 10. The entry shows that a user named Darril Gibson read a file named PayrollData.rtf in the C:\Payroll folder using the WINWORD application.



Event 4656, Microsoft Windows security auditing.

General Details

A handle to an object was requested.

Subject:
Security ID: DarrilLaptop\Darril Gibson
Account Name: Darril Gibson
Account Domain: DarrilLaptop
Logon ID: 0x5ACFC586

Object:
Object Server: Security
Object Type: File
Object Name: C:\Payroll\PayrollData.rtf
Handle ID: 0xee8
Resource Attributes: -

Process Information:
Process ID: 0x2594
Process Name: C:\Program Files (x86)\Microsoft Office\root\Office16\WINWORD.EXE

Access Request Information:
Transaction ID: {00000000-0000-0000-0000-000000000000}
Accesses: SYNCHRONIZE
ReadAttributes

Access Reasons:
SYNCHRONIZE: Granted by D:(A;ID;0x1200a9;;;BU)
ReadAttributes: Granted by D:(A;ID;0x1200a9;;;BU)

Log Name: Security
Source: Microsoft Windows security Logged: 3/19/2018 12:00:27 PM
Event ID: 4656 Task Category: File System
Level: Information Keywords: Audit Success
User: N/A Computer: DarrilLaptop
OpCode: Info

Figure 10-1 Windows 10 Security log entry

It could be that Darril is a hiring manager and needs access to that file on a regular basis. It could also be that Darril is an administrator completely uninvolved in the accounting or payroll department and abused his administrative privileges by accessing the file.

It's worth repeating that strong authentication is an important element here. If weak passwords are used, Darril could claim that he never looked at the file and suggest that someone guessed his password and impersonated him. On the other hand, if the organization implemented multifactor authentication and Darril was required to log on with a smart card and personal identification number (PIN), he won't be able to deny that he accessed the file. In this case, the audit logs provide nonrepudiation.

Other access control methods, such as video cameras and badge reader logs that record when someone enters or exits a secure area, provide added accountability. If these other methods documented that Darril was in the general location, they provide additional evidence for accountability.



EXAM TIP Audit logs combined with strong authentication and authorization practices provide nonrepudiation. When logs record specific user actions, users are unable to deny they took an action. At least, they don't have any credibility if they do try to deny taking the action.

Sometimes, a higher authority mandates specific accounting practices. For example, the Federal Information Processing Standards Publication 200 (FIPS Pub 200) published by the National Institute of Standards and Technology (NIST) specifies that federal agencies must meet at least the following minimum requirements:

- (i) create, protect, and retain information system audit records to the extent needed to enable the monitoring, analysis, investigation, and reporting of unlawful, unauthorized, or inappropriate information system activity; and*
- (ii) ensure that the actions of individual information system users can be uniquely traced to those users so they can be held accountable for their actions.*

Although this squeezes many words and phrases into a single long sentence, it does provide a good summary of the requirements of audit logs for federal agencies. It's also a good summary of the requirements any organization can use when implementing auditing and accountability.

Auditing with Logs

Audit logs record details such as who, what, when, and where for any events of interest. For example, [Figure 10-1](#) (shown earlier) allows you to identify who took an action, what action the user took, when, and where:

- **Who** The local user account Darril Gibson on the computer named DarrilLaptop (DarrilLaptop\Darril Gibson).
- **What** Darril read the ParyrollData.rtf file using the WINWORD application.
- **When** The entry was logged on March 19, 2018 at 12:00 P.M.
- **Where** The file was in the C:\Payroll folder on the Windows computer.

This example shows Darril read the file. If Darril modified or deleted the file, the audit log could also show those actions, depending on how administrators configured logging.



TIP In Microsoft and most other computing systems, administrators must specify what to audit. If a system automatically audited all types of access to every single file, folder, and other system resource, the logs would become astronomically large. They would consume drive space and contain so much data that retrieving useful information would be difficult for administrators.

The primary purpose of auditing is to be able to reconstruct events leading up to and occurring during an incident. A secondary purpose is to serve as a deterrent. In other words, if users are aware that logs record their activities, they may be less likely to engage in unauthorized activities.

As an example, imagine an organization has an acceptable use policy (AUP) that states that users should only use computers and networking resources for work purposes. The AUP specifically prohibits accessing specific categories of websites, such as gambling sites. Non-work-related sites could easily be malicious and attempt to infect user systems with malware. However, if users know they shouldn't be visiting these sites and that logs are recording their activity, they may be less likely to visit these sites. The net effect is a reduction of incidents.

Proxy servers can block user access to unauthorized sites. Many organizations configure the proxy server to display a warning page when users attempt to visit unauthorized sites. These warnings often repeat the relevant information from the AUP. They also show the user information that the proxy server has logged, such as the username, the URL of the blocked site, and the date and time when the user attempted to access the site.



EXAM TIP Audit logging can reduce incidents simply by deterring users from engaging in unauthorized behaviors.

Clipping Levels

Many automated accounting systems generate alerts after the system detects a preset number of events. The system accepts or ignores these events until they reach a predetermined threshold, or *clipping level*. A classic example is a report based on failed logon attempts. Instead of reporting every single failed logon attempt, a system uses a clipping level to generate a report only if an account has exceeded a set number of failed logon attempts in a given period.

For example, administrators can configure an accounting system to generate a report if a log records three failed logon attempts in a 30-minute period for a single user account. If a user enters the wrong password once or twice, the system ignores those attempts as acceptable, normal user errors, and doesn't generate a report. However, if a user enters the wrong password three times within 30 minutes, it exceeds the clipping level, and the accounting system generates an alert.

Some automated auditing tools such as security information and event management (SIEM) tools, discussed in [Chapter 8](#), perform audit-reduction tasks. In other words, they can examine a log and retrieve relevant data from the log while ignoring the irrelevant data. These audit-reduction tools use clipping levels to identify events of interest based on predetermined thresholds.



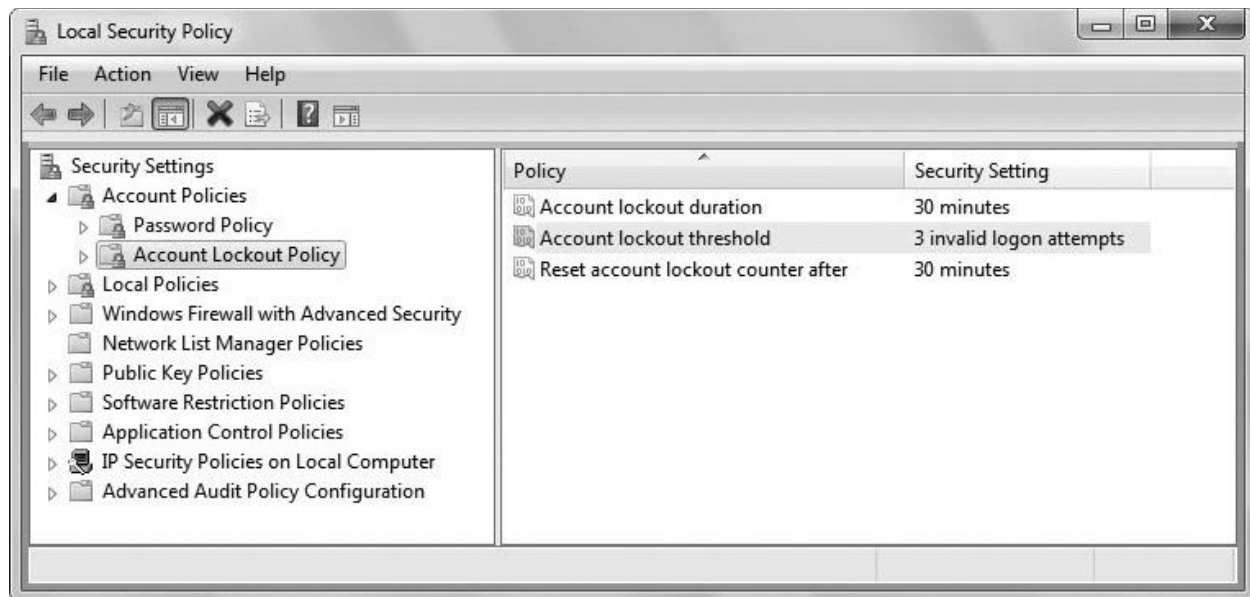
EXAM TIP A clipping level ignores events such as normal user errors until the auditing system reaches a predetermined threshold. However, the accounting system generates an alert after it detects the number of events has reached the threshold. Clipping levels such as this ignore individual logon failures and only act after reaching the threshold.

Account Lockout

Many systems use an account lockout policy that locks out an account after a predetermined number of failed logons. These policies use clipping levels and only lock out the user after the number of failed logons exceeds a threshold.

For example, the following illustration shows an account lockout policy within the local security policy for a Windows system. It has the Account Lockout Threshold set to “3 invalid logon attempts,” so if a user enters the wrong password once or twice, the system doesn't lock out the account. The Reset Account Lockout Counter After setting is set to 30 minutes. If users don't

enter the wrong password three times within a 30-minute period, the system doesn't lock the account. After 30 minutes, the system resets the incorrect password counter to zero.



You can also see that the Account Lockout Duration setting is 30 minutes. If the system locks out an account, it automatically unlocks the account after 30 minutes without any administrative intervention. If this setting is 0, the system locks the account until an administrator unlocks it.

Understanding Audit Trails

An audit trail includes one or more logs that track events occurring on a system or network. When auditors have access to all the logs, they can re-create the events that occurred leading up to the incident and determine what actually occurred during the incident. In other words, the audit logs allow auditors and administrators to re-create all the activity, or the trail of activity, created by an individual or process.

For example, if an attacker launched an attack on a web server in a demilitarized zone (DMZ), you could use firewall logs on the external firewall to identify the path the attacker used to get into the DMZ. You could then use the logs on the web server itself to track what the attacker did once he reached the web server. Logs on the web server could include host-based firewall logs, system logs, and security logs.

In some cases, the "who" identifies a specific account, such as a user account. In other situations, the "who" identifies the source of activity, such as an IP address. It's worth noting that many attackers launch attacks through other locations or spoof the source IP address. In other words, the logged IP address doesn't always provide the actual IP address of the attacker.

An audit trail is a technical detective control. However, when configuring or evaluating an audit trail, you should consider the following items to ensure it is effective as a detective control:

- **Does the audit trail provide a full accounting of user activity?** For example, does the audit trail record who, what, when, and where?
- **Is access to the logs controlled?** If anyone can access the logs, it makes it easy for attackers to erase the logs after an attack.
- **Are copies of logs stored on remote servers?** This ensures that even if an attacker erases or corrupts a log file after an attack, the remote server still holds a copy of the original logs.

Exploring Audit Logs

There are many different types of audit logs on computing and networking systems. Operating systems include regular logs that record activity on the computer; applications frequently record information in application logs based on the needs of the application; and network devices (such as routers and firewalls) regularly record different types of data to document network activity.

This section describes many of these logs. It also describes how administrators and auditors use these logs to track activity. Together, these logs can help an organization create a comprehensive audit trail documenting activity on individual systems and networks.

Operating System Logs

Operating systems include logs that track activity on individual computers and servers. Operating systems also provide the means to review these logs, such as the Event Viewer in Windows systems. Older Microsoft operating systems included at least the System, Application, and Security logs, but newer Windows operating systems include additional logs that record information on specific applications and services installed or running on the system.

Security logs can record activity from any auditable activity and can record both success and failure events. In other words, if a user tries to open a file and successfully opens it, it is a success event. However, if the user doesn't have appropriate permissions, the system blocks the user's access and can record it as a failure event. Security logs don't record all activity by default, but instead require administrators to configure logging based on the organization's needs. For example, an administrator can configure object access logging for objects such as folders, files, or printers, and then when a user accesses one of these objects, the security log records the event. Other auditable security events include deleting or modifying files or folders, modifying system settings such as the system time, modifying permissions, and more.



TIP It's not useful to audit all object access. Instead, administrators choose to audit access to sensitive or proprietary data. For example, an organization may choose to record all access to classified research and development data, but not record access to publicly available data.

Some of the logs used in Windows systems include

- **Security log** Records security-related events such as when an account logs on or off or any auditable event. Administrators often configure specific items to audit based on the needs of the organization.
- **System log** Records events from the operating system, such as when the system boots or is shut down, and when a driver or service stops or starts.
- **Application log** Records events from applications. These events can be from end-user applications, such as a web browser, or server-based applications, which include database server applications and networking services, such as a server running the Domain Name System (DNS) service. Applications can log data in the generic Application log or create an application-specific log and log events in the application-specific log.
- **Setup log** Records events related to the setup of certain applications.

Microsoft has included additional features with logs in Event Viewer. For example, [Figure 10-2](#) shows the Event Viewer in Windows 8 displaying some of the information from a successful logon event. Event Viewer allows you to view details from any of the captured events. You can also see Forwarded Events and Subscriptions in the left pane of this figure. Administrators create subscriptions on a central system that collects forwarded events from other systems. This allows administrators to monitor events from multiple systems in the network from a centralized computer.

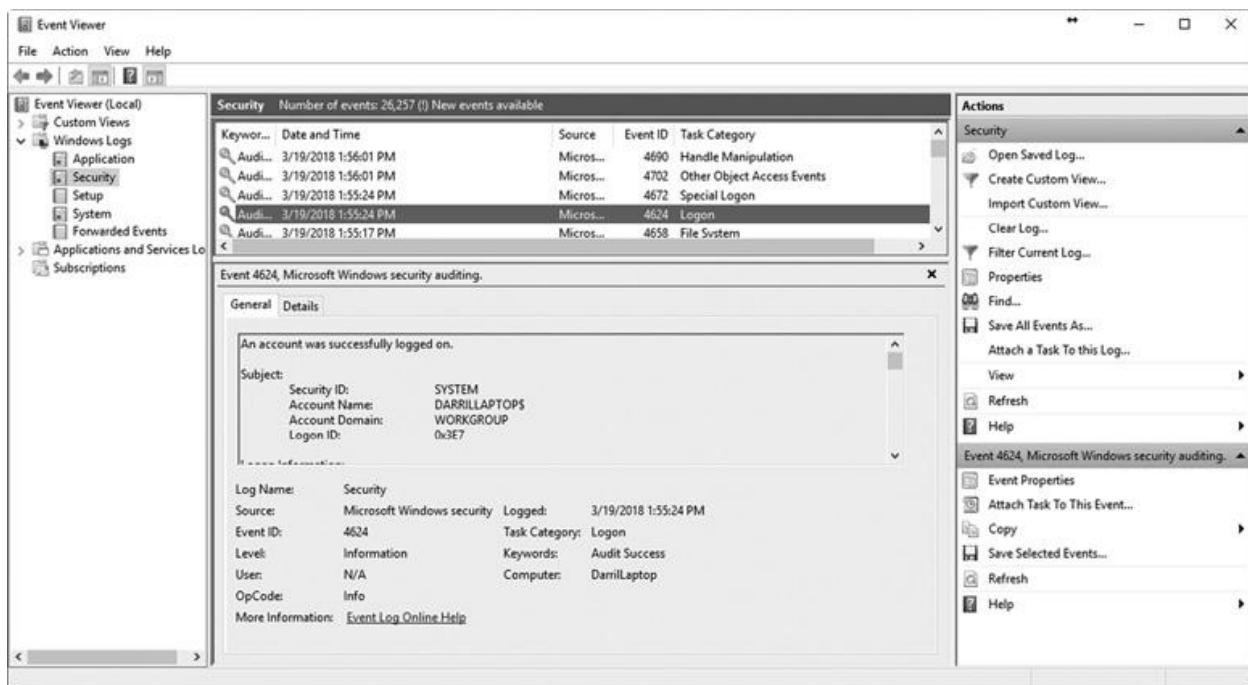


Figure 10-2 Microsoft Event Viewer

If you manage 10 servers hosting Microsoft SQL Server databases, you can forward relevant events to an 11th system to monitor all of these servers. The original servers will still hold the original logs. However, the monitoring server hosts a copy of the logs in a remote location. Even if an attacker modifies the logs on one of the servers after an attack, the remote server still holds all the log entries within the Forwarded Events log.

Storing Logs on Remote Systems

[Chapter 8](#) mentioned that attackers could modify logs stored on a local system (rather than a remote system). Administrators and investigators should view these local logs with suspicion after an attack.

However, if you can store the logs on remote systems, it is much less likely that an attacker will be able to modify these remote logs. In other words, by storing logs on remote systems, you can help retain their integrity. A common way to store logs on remote systems is to forward them to a SIEM application.



EXAM TIP Many attackers attempt to erase the audit trail by erasing or modifying individual logs. Their goal is often to get in, attack the system, erase or modify the logs, and then get out. However, logs on remote systems are more difficult for attackers to erase or modify.

*Nix Logs

UNIX and UNIX derivatives (such as Linux) are often referred to as *nix systems, and these systems have their own types of logs. The log file locations sometimes vary from one version to another, but there are some similarities. For example, system and accounting log files are usually in one of the following directories: /var/log, /var/adm, and /usr/adm.

Syslogd is a service that typically receives all the log entries and then moves them to a log based on how it is configured. The syslog.conf file specifies the configuration. If the configuration file is not available or is blank, syslogd will move the log entries to the syslog (/var/log/syslog). If the file is available, syslogd can specify any log for log entries. As an example, the following line will ensure all mail messages are sent to one log:

```
mail.* /var/log/maillog
```

Some of the common *nix log files and their contents include the following:

- **syslog** Many security processes and other utilities write to syslog. It's often a good place to view overall system activity.

- **sudo** This log records all attempts to use the su command. The su command (sometimes called the super user command) grants the user elevated privileges. Attackers often attempt to use su in various attacks, and uneducated users can misuse it.
- **auth** This log records authentication events on a Secure Shell (SSH) server. You can use it to determine whether a system intrusion has occurred. For example, a brute-force attempt to log in to the system will show multiple failed SSH login attempts, usually from multiple usernames.
- **maillog** The system records inbound and outbound mail activity in this log.

Proxy Server Logs

Most proxy servers can log all activity through the proxy server. This allows an organization to track websites that users access with their web browsers, how much time they spend on the sites, and all attempts to visit unauthorized sites. [Chapter 4](#) presented proxy servers. As a reminder, an organization typically uses a proxy server for caching and URL filtering.

You may recall that an organization can use a proxy server to retrieve web pages from the Internet on behalf of internal clients. The proxy server uses Network Address Translation (NAT) to translate private addresses from internal clients to public addresses used by Internet-based websites. The proxy server caches web pages. If a user requests a cached web page, the proxy server sends the page to the user without retrieving it from the web server again. This reduces Internet bandwidth usage.

URL filtering prevents users from accessing unauthorized websites. If a user attempts to go to an unauthorized website, the proxy server blocks the request. As mentioned earlier in this chapter, proxy servers can also log these unauthorized attempts. When a user visits an unauthorized site, some organizations display a warning similar to that shown in [Figure 10-3](#). When users see that the proxy server is logging their activity, they may be less inclined to circumvent the policy.



Figure 10-3 Proxy server warning

Of course, it's possible that a user has a legitimate need to access unauthorized websites. Organizations typically include a process allowing users to ask for an exception to the policy.



EXAM TIP Some logs such as proxy server logs can serve as a deterrent to unacceptable or undesirable behavior. Although this isn't a good deterrent to stop criminals, it can help to deter many users from engaging in unsafe computing practices.

Firewall Logs

Most firewalls include the ability to log all activity sent in and out of a firewall. This includes traffic that the firewall blocks, traffic that it passes, or both. These logs can be very useful when trying to reconstruct the events from an attack. For example, after discovering an attack on an internal system, you can use the logs from firewalls at the edge of your network to determine whether the attack came from an external attacker on the Internet or from an internal source.

SIEM applications commonly use firewall logs to detect attacks such as port scans or TCP flood attacks. The SIEM application regularly monitors the firewall logs, and when it detects a potential attack, it can log and report the event. Some active intrusion detection and prevention systems (IDPSs) can also take action to block an attack. For example, an IDS can modify the configuration of the firewall to block an attack in progress.

Firewall logs can also provide details on botnet activity. If you suspect that one or more of your internal systems has become a zombie within a botnet, you can check the host-based firewall logs of the system and the network-based firewall logs for suspicious activity. For example, if 20 of your systems are regularly connecting with a system that has an IP address from another country and you don't recognize that address, it could be an indication of botnet-related activity. The foreign IP address might be the address of a command-and-control server used to provide instructions for zombies within the botnet, or another server that attackers hijacked to provide instructions to the zombies.

Reviewing Logs

One of the most important steps of auditing is reviewing audit logs. Reviewers should look for any types of anomalies in the logs. For example, a long string of failed attempts to log on to the administrator account indicates a potential brute-force attack. If the logs then show a successful logon event for the administrator account, it indicates an attacker probably guessed the password for the account and is now accessing the system with administrative privileges. In this situation, it would be important to examine all the activity by this administrator account.

Admittedly, reviewing logs for anomalies isn't easy to do. If you're reviewing thousands of log entries and the log includes one or two abnormal entries, it's easy to overlook the abnormal entries. However, many automated systems, such as SIEMs, are available that can look for and issue alerts for specific events. Most SIEMs include filters that allow administrators to filter out unwanted log entries.

In some organizations, administrators get overwhelmed with day-to-day activities and skip the reviewing step. However, at the very least, administrators should archive the logs so that they can review them later, if necessary. For example, [Figure 10-4](#) shows how you can right-click any log in Microsoft's Event Viewer to save the contents of a log. Administrators and auditors can open saved logs in Event Viewer by selecting Open Saved Log and browsing to the location of the saved file.

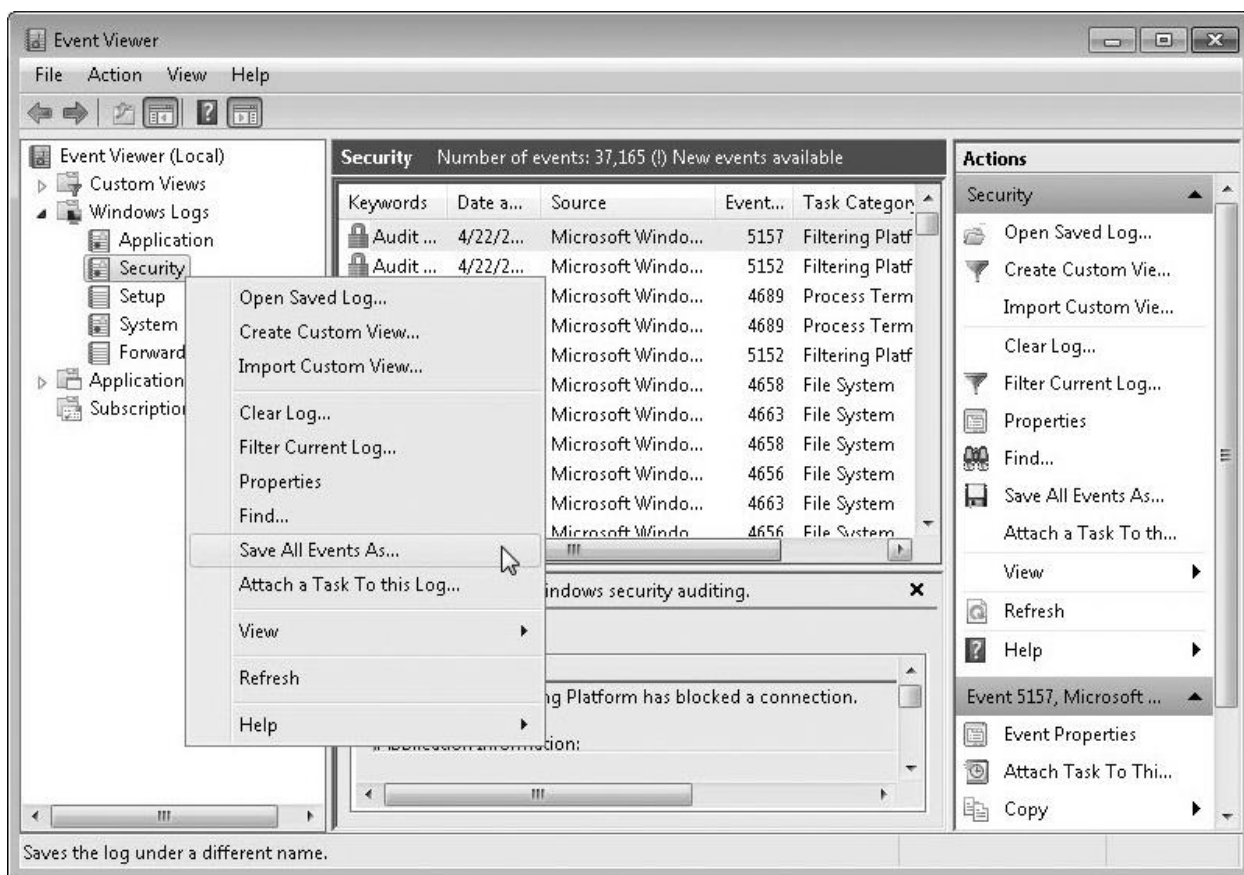


Figure 10-4 Archiving a log

When you archive logs like this, you can review them later. In other words, even if you don't notice an incident until much later, you can still re-create the events by reviewing the archived logs.

Managing Audit Logs

In addition to storing logs on remote systems to maintain their integrity, you can take several other steps to improve the security related to logs. A primary step is to protect the logs with permissions.

Access to logs should be restricted to personnel (such as administrators) who have a need to view or maintain the logs. If anyone can clear the logs, users might erase logs accidentally. Similarly, if anyone can clear the logs, attackers will be able to do so to erase their activity after an attack.



TIP Most systems include the ability to detect when someone has cleared a log. For example, in Windows systems, the first log entry in a new log records details on when someone cleared the previous log. If an organization normally archives and clears logs on Friday, but a log shows that someone cleared it on a different day, it indicates suspicious activity that deserves further investigation.

Many organizations mandate the retention of logs for a specific period such as three years or longer. This allows an investigator to go back in time to reconstruct events even if several months or years have passed. Log retention policies require someone to archive the logs periodically. In smaller organizations, administrators do this manually. Larger organizations use automated processes to archive logs. In some scenarios, organizations write the archived log data to a write once, read many (WORM) drive to prevent someone from rewriting the archived logs.

Performing Security Audits

A security audit examines an organization's practices and operations to determine whether they conform to the organization's policies or applicable laws. An organization can perform an internal audit to examine its practices, or external auditors can come in to examine the organization's practices. The audit documents the organization's policies, processes, controls, testing, and results. The following sections discuss various audit practices.

Periodic Audit and Review

Audits can be performed on a periodic basis, such as once a year, after an event such as a security incident, or when required by laws or guidelines. Audits are detective controls because they detect events after they've occurred.



EXAM TIP Security audits help an organization identify vulnerabilities in its processes and procedures. After a security audit, it's important to implement fixes to mitigate or reduce the discovered vulnerabilities.

As an example, an organization can perform monthly vulnerability scans to check for unused accounts. Many policies mandate disabling accounts when an employee leaves, or even if the employee takes an extended leave of absence. These scans can detect accounts that are enabled but unused. Similarly, an organization can perform an annual audit of processes within the organization to determine if these processes comply with the organization's security policy. Organizations also perform these types of audits after an incident to determine if personnel are complying with existing policies. For example, after a successful malware attack, an audit can help determine if the infection was due to personnel not following existing policies.

Organizations that must follow certain laws and guidelines often use compliance audits to ensure that personnel are complying with the relevant laws and guidelines. Internal compliance audits provide proof to outside entities that the organization is taking steps to ensure compliance. For example, if a U.S. organization handles any type of health- or medical-related data, it must comply with the Health Insurance Portability and Accountability Act (HIPAA). If an incident occurs showing a violation of the law, the organization can show documentation from past audits as proof that the organization is taking active steps to be compliant with the law. However, if no past audits exist, the organization may find itself incurring significantly higher fines and penalties after a security incident.

In other situations, organizations hire external auditors, but typically only when they are required to do so. For example, the Payment Card Industry Data Security Standard (PCI DSS) requires some organizations that handle credit card data to hire an external auditor after the organization has suffered a data breach.

Auditing Passwords

A password audit checks to see whether users are following the policies related to passwords. For example, a password policy may dictate that regular users use strong passwords with at least three of the four character types (uppercase, lowercase, numbers, and symbols) and be at least eight characters long. A password audit could use a password-cracking tool to determine whether any passwords fail to meet these requirements. The password audit could also check to see if users are regularly changing their passwords in compliance with a written password policy.

I know from experience that just asking users to use strong passwords and expecting that they will, or just asking them to change their passwords regularly and expecting that they will, works about as well as getting a cat to come to you when you call. In general, you can't expect most users to take secure actions on their own.

Instead, organizations use technical controls to implement written password policies. For example, Microsoft operating systems use a technical password policy within Group Policy to force users to create strong passwords and change them regularly. Third-party tools are also available that can enforce the use of strong passwords, but they do cost money. However, if a

password audit shows that users are using weak passwords and not following the written policy, the cost of a password enforcement tool is justified.



NOTE If an organization is using best practices with passwords, auditing passwords becomes much more difficult. Ideally, passwords are hashed with strong hashing algorithms and include a salt. This makes it more difficult for a password cracker to identify the password from the hash.

Auditing Security Policies

Many organizations require an annual review of a written security policy to ensure that the policy still meets the needs of the organization. Auditors can examine historical data from the past year to see whether users are following the security policy. They can also examine details of security incidents to help them determine if the policy is helping the organization prevent security incidents.

For example, some security policies look great on paper, but employees never hear about them. This often results in users repeating the same mistakes year after year. Social engineers may continue to trick users, or employees might not follow specific procedures simply because they are unaware of the policies. In this case, the audit may discover a lack of security awareness and recommend specific actions such as security training to improve employee awareness.



NOTE [Chapter 12](#) discusses security policies in more depth. A security policy is a written document that provides overall direction for an organization's security.

ISACA

ISACA is an international organization that provides standardization for IT professionals involved in auditing. It includes more than 110,000 members from more than 180 countries working in various IT-related positions, such as IS auditor, internal auditor, consultant, educator, IS security professional, and chief information officer (CIO).

Two of ISACA's top auditing certifications are the Certified Information Systems Auditor (CISA) and the Certified Information Security Manager (CISM) certifications.

- **CISA** The CISA certification identifies individuals who audit, control, monitor, and assess an organization's information technology and business systems.

- **CISM** The CISM certification is management-focused and identifies individuals who design, build, and manage enterprise information security programs.



NOTE ISACA's original name was the Information Systems Audit and Control Association. It dropped the long name in 2008 and now goes by just the acronym ISACA.

ISACA developed the Control Objectives for Information and Related Technology (COBIT) framework. COBIT is a robust framework that helps managers identify processes to manage IT systems and it is one of the popular management frameworks. COBIT 5 is based on five principles and seven enablers, which are generic enough that organizations of any size can apply them.

Exploring PCI DSS Requirements

The Payment Card Industry (PCI) Security Standards Council created and regularly updates the Payment Card Industry Data Security Standard (PCI DSS). The PCI DSS requires that organizations that process credit card payments comply with specific steps to protect customer account data, and it includes a strong focus on protecting credit card data. The goal of PCI DSS is to reduce overall credit card fraud, and toward that goal it specifies 6 control objectives with 12 supporting requirements related to security management, policies, procedures, network architecture, software design, and more.

[Table 10-1](#) provides an overview of the PCI DSS (Version 3.2) objectives along with the specific requirements associated with these objectives.

| Control Objectives | Requirements |
|---|---|
| Build and Maintain a Secure Network and Systems | <ol style="list-style-type: none"> 1. Install and maintain a firewall configuration to protect cardholder data. 2. Do not use vendor-supplied defaults for system passwords and other security parameters. |
| Protect Cardholder Data | <ol style="list-style-type: none"> 3. Protect stored cardholder data. 4. Encrypt transmission of cardholder data across open, public networks. |
| Maintain a Vulnerability Management Program | <ol style="list-style-type: none"> 5. Protect all systems against malware and regularly update antivirus software or programs. 6. Develop and maintain secure systems and applications. |
| Implement Strong Access Control Measures | <ol style="list-style-type: none"> 7. Restrict access to cardholder data by business need to know. 8. Identify and authenticate access to system components. 9. Restrict physical access to cardholder data. |
| Regularly Monitor and Test Networks | <ol style="list-style-type: none"> 10. Track and monitor all access to network resources and cardholder data. 11. Regularly test security systems and processes. |
| Maintain an Information Security Policy | <ol style="list-style-type: none"> 12. Maintain a policy that addresses information security for all personnel. |

Table 10-1 PCI DSS Control Objectives



TIP Although these objectives and requirements only apply to organizations that must comply with PCI DSS, you can slightly modify the steps to apply them to any organization. Replace “cardholder data” with “data” throughout and this list becomes a generic list for any organization. If you want to read more about the PCI DSS standards, you can access them from here: https://www.pcisecuritystandards.org/pci_security/.

Compliance Reports

Organizations that fall under the PCI DSS must submit regular compliance reports. These reports identify the business and its use of credit cards. They also document the business’s network, including its hardware, software, and connectivity, and how the business complies with each of the 12 requirements. However, organizations do not send the reports to the PCI Security Standards Council. Instead, organizations send the reports to the acquiring bank (the bank handling the credit card transactions) and the global payment brand (such as Visa, MasterCard, and American Express).

Compliance Audits

Businesses are also required to perform periodic audits to verify their compliance with the PCI DSS requirements. Businesses that don't have a high volume of credit card use can perform the compliance audits internally. These organizations can often just answer a Self-Assessment Questionnaire (SAQ) and submit it to the acquiring bank. However, businesses that have a high volume of credit card use must hire an independent expert to perform an audit. The PCI Security Standards Council certifies Quality Security Assessors (QSAs) and Approved Scanning Vendors (ASVs) to perform these audits. A QSA is an external data security firm that is qualified to perform an onsite audit. An ASV is an external data security firm using an authorized vulnerability scanner. An Internal Security Assessor (ISA) is an employee of the organization who has obtained the ISA certification after attending training at PCI DSS. An ISA can perform some onsite audits.

The PCI Security Standards Council can fine businesses that don't comply with the PCI DSS requirements, or revoke the businesses' authorization to process credit card payments. This provides a high level of motivation for businesses to ensure they comply with the PCI DSS requirements.



NOTE It's worthwhile stressing that the PCI DSS provides a minimum baseline for security. Simply being in compliance does not provide any guarantees against an attacker compromising customer and credit card data. Organizations often implement many other safeguards beyond what is required from the PCI DSS.

PCI DSS Merchant Levels

Organizations have different requirements based on their merchant level. It's easier for an organization identified as Merchant Level 4 to verify compliance with PCI DSS requirements than it is for an organization identified as Merchant Level 1. The details that classify what level a merchant falls under vary for different credit cards. The following list summarizes the general guidelines for PCI merchant levels:

- **Merchant Level: 1** The merchant processes over 6 million transactions annually. The merchant must submit an annual Report on Compliance (ROC) completed by a QSA, along with quarterly network scans completed by an ASV.
- **Merchant Level: 2** The merchant processes between 1 million and 6 million transactions annually. The merchant must submit an annual SAQ or submit an annual Report on Compliance (ROC) completed by a QSA (depending on specific credit card requirements). The merchant must also ensure quarterly network scans are completed by an ASV.

- **Merchant Level: 3** The merchant processes between 20,000 and 1 million transactions annually. The merchant must submit an annual SAQ and ensure quarterly network scans are completed by an ASV.
- **Merchant Level: 4** The merchant processes less than 20,000 transactions annually. The merchant must submit an annual SAQ and ensure quarterly network scans are completed by an ASV.

Auditing Physical Access Controls

If an organization has physical access controls, it can periodically perform an audit to determine whether these controls are effective. Physical access control audits can also verify if personnel are using the controls or attempting to bypass them.

As an example, many proximity card badge readers have the capability of recording the date and time when an individual enters or exits a secure area. If personnel consistently use their badges to enter and exit the area during the day, the logs would show this. However, if users frequently tailgate each other without using their badges, the logs would show inconsistencies. Each entry showing a user entered the area should have a matching entry for the user when they exit the area.

[Chapter 5](#) covers tailgating in more depth. As a reminder, tailgating occurs when someone passes through a controlled entry point without providing credentials. Instead, the person follows closely behind someone else who has provided credentials. In this scenario, the first person opens the door with a proximity badge and the second person tailgates the first person, bypassing the physical security control.

Physical access logs can also identify whether users are entering secure areas outside of normal work hours. For example, if an employee normally works from 9 to 5 during the week but a log shows the employee entering at 10 P.M. one evening, it might indicate a problem. However, if the logs don't exist, or there isn't a method of reviewing the logs, the organization may not detect this potentially malicious activity.



EXAM TIP A physical access audit can detect whether tailgating is occurring. If employees can tailgate each other, a social engineer can use the same technique to get into a secure area.

Understanding Configuration Management

Configuration management is the practice of ensuring that systems are configured with security in mind and that configuration is tightly controlled. Many organizations use configuration management to configure systems with a secure baseline and to ensure the system remains secure with the same configuration.

As an example, administrators ensure that servers start with similar system configurations that make them more secure from the default configuration. Administrators use various hardening techniques such as removing or disabling unnecessary services and protocols, and changing default passwords for all the servers. They then tweak the configuration of different groups of servers to meet the needs for these groups. For example, they configure web servers differently than they configure file servers.

The “Holding Users Accountable with Audit Logs” section earlier in this chapter mentioned FIPS Pub 200, which mandates minimum accounting requirements for IT systems. It also lists the following minimum requirements for configuration control:

(i) establish and maintain baseline configurations and inventories of organizational information systems (including hardware, software, firmware, and documentation) throughout the respective system development life cycles; and

(ii) establish and enforce security configuration settings for information technology products employed in organizational information systems.



TIP Based on the FIPS Pub 200 requirements to enforce security configuration settings, some federal agencies combine configuration management with change management. However, it’s worth stressing that these are separate processes.

In addition to taking steps to ensure that systems begin with specific configuration settings, an active configuration control plan will mandate that systems be checked periodically. This ensures that the systems maintain the secure settings from the baseline configuration. This process works hand in hand with a strong change management program that attempts to prevent unauthorized changes and the unintended outages caused by unauthorized changes.

Using Imaging for Configuration Management

A common configuration management practice is to use imaging technologies. For example, an organization can deploy each desktop system with a common image that includes the operating system, basic applications, and specific settings needed for usability and security. Images help ensure that every deployed system starts with a secure baseline.



EXAM TIP System baselines are an important element of configuration control and are often implemented with images. When administrators know the starting point (the baseline

configuration), they can easily check systems to determine if the configuration is different from the baseline.

[Figure 10-5](#) shows the overall process of imaging. You start with a clean computer called the reference computer and install an operating system on it. You can then install applications and configure the operating system based on the needs of the organization. For example, you can configure security settings such as auditing, a password policy, an account lockout policy, a password-protected screen saver, and much more. After testing the image, administrators capture it using imaging software. They then deploy the image to other systems, and these imaged systems start with the same configuration as the reference computer.

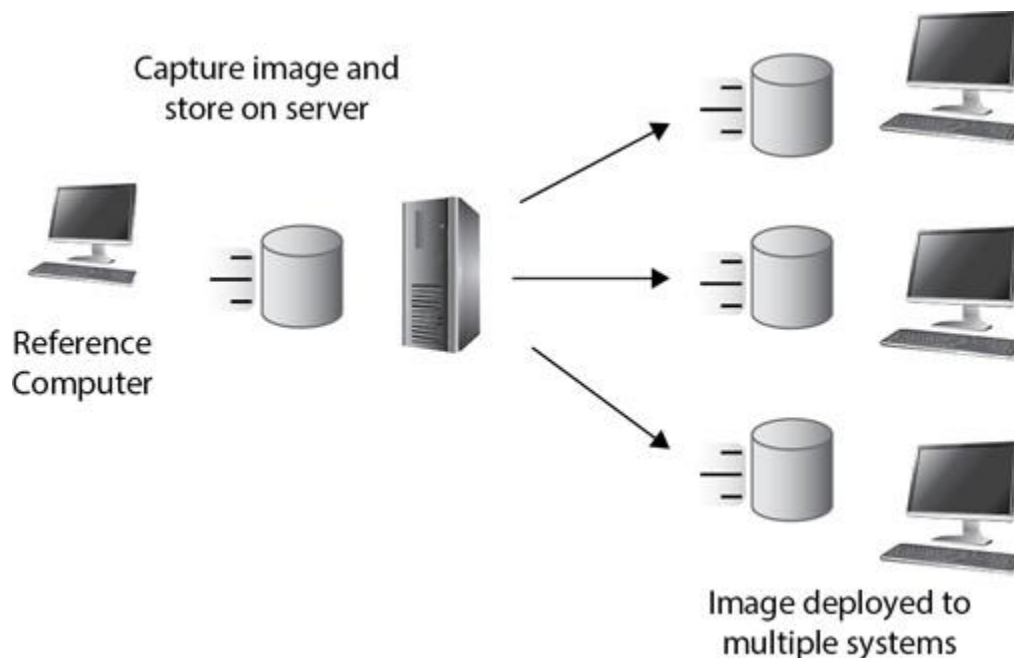


Figure 10-5 Using imaging for configuration control

Images also make systems easier to maintain. Help desk professionals only need to learn the configuration of one system. No matter what user they are helping, the configuration is the same, enabling help desk professionals to troubleshoot and restore the system quicker.

In contrast, when an organization has 100 computers and each of these computers has a different configuration, support personnel must first figure out how the system should work normally and then take steps to restore it to its original configuration. Effective configuration management programs reduce the overall total cost of ownership of systems by reducing the repair times.

Using Group Policy for Configuration Management

In Microsoft networks, administrators use Group Policy to configure systems. Group Policy allows centralized administration of systems. An administrator can configure one setting with a Group Policy object (such as opening a port on a firewall), and Group Policy applies the change to all the computers within the scope of the Group Policy object. Whether the organization has

50 computers or 50,000 computers, Group Policy can apply the change consistently to all relevant computers in the organization.

Group Policy also supports targeting specific systems, such as all the computers in a specific department or only computers running Windows 7. Administrators with a little knowledge of Group Policy can easily implement configuration control to all their Microsoft systems with very little administrative effort.

Using Other Tools for Configuration Management

Other tools are available for configuration management and these are especially valuable when automating management of non-Microsoft systems such as Linux and other Unix-like systems. A challenge with many of these tools is the learning curve for administrators to master the language of each tool. System administrators often find it challenging to cross the line from administration to writing scripts or developing apps in a programming language. The ones that do are typically the most valuable within an organization. The following list mentions some of these tools:

- **Puppet** Puppet uses Puppet's proprietary language or Ruby domain-specific language. Many commands must be entered from a command-line interface. It comes in two versions: Open Source Puppet (which is free) and Puppet Enterprise (which has a cost).
- **Chef** Chef uses Ruby to create system configuration "recipes." Administrators can group multiple recipes together as a "cookbook" to identify packages to install, services to run, and files that should be written.
- **Ansible** Ansible is written in Python, a popular language used to create websites. Administrators can script commands in YAML (YAML Ain't Markup Language). YAML uses a human-readable data serialization language. It uses SSH to encrypt data sent over the network.

Understanding Change Management

Many organizations have a change management or change control process in place. The primary purpose is to ensure that changes don't result in unintended outages. Networks are complex, and many of the systems within a network interact with each other. Making a change in one system can easily affect other systems and result in an unintended outage. However, if experts examine these changes before technicians or administrators implement them, the organization can prevent many of the unintended outages from changes.

For example, imagine that a server hosts an application that requires opening TCP port 4567 on a server. Administrators open the port when they configure the server application, and the application runs as expected. Later, a security professional performs a vulnerability assessment and notices TCP port 4567 is open. Thinking this port is a vulnerability, the security professional decides to close the port and reduce the server's attack surface. Of course, this action breaks the server application. Users begin to complain, and the organization calls in the application developers to troubleshoot the problem. However, it might take them hours, or even days, before they discover the closed port.

Even though the security professional's intentions are good, the action still caused an outage. You could fill volumes with similar examples of well-intentioned IT personnel making changes, often to correct a minor problem, and in the process creating other problems that cause outages.

In contrast, a change control process ensures that appropriate personnel examine all changes before approving them. It also provides a documentation trail. If the server ever suffers a catastrophic failure, administrators can restore it to its last state, including all the changes applied to the server.

Although many organizations have adopted some level of change management, there is a wide diversity of maturity among organizations. For example, some organizations have a policy for change management, but follow it only for major changes and still suffer outages from many minor changes.

Change Management Process

A change management process gives stakeholders a formal mechanism to propose changes to a system. Multiple entities then examine the change request and either approve or reject it. The following list describes the overall steps in a change control process, as shown in [Figure 10-6](#):

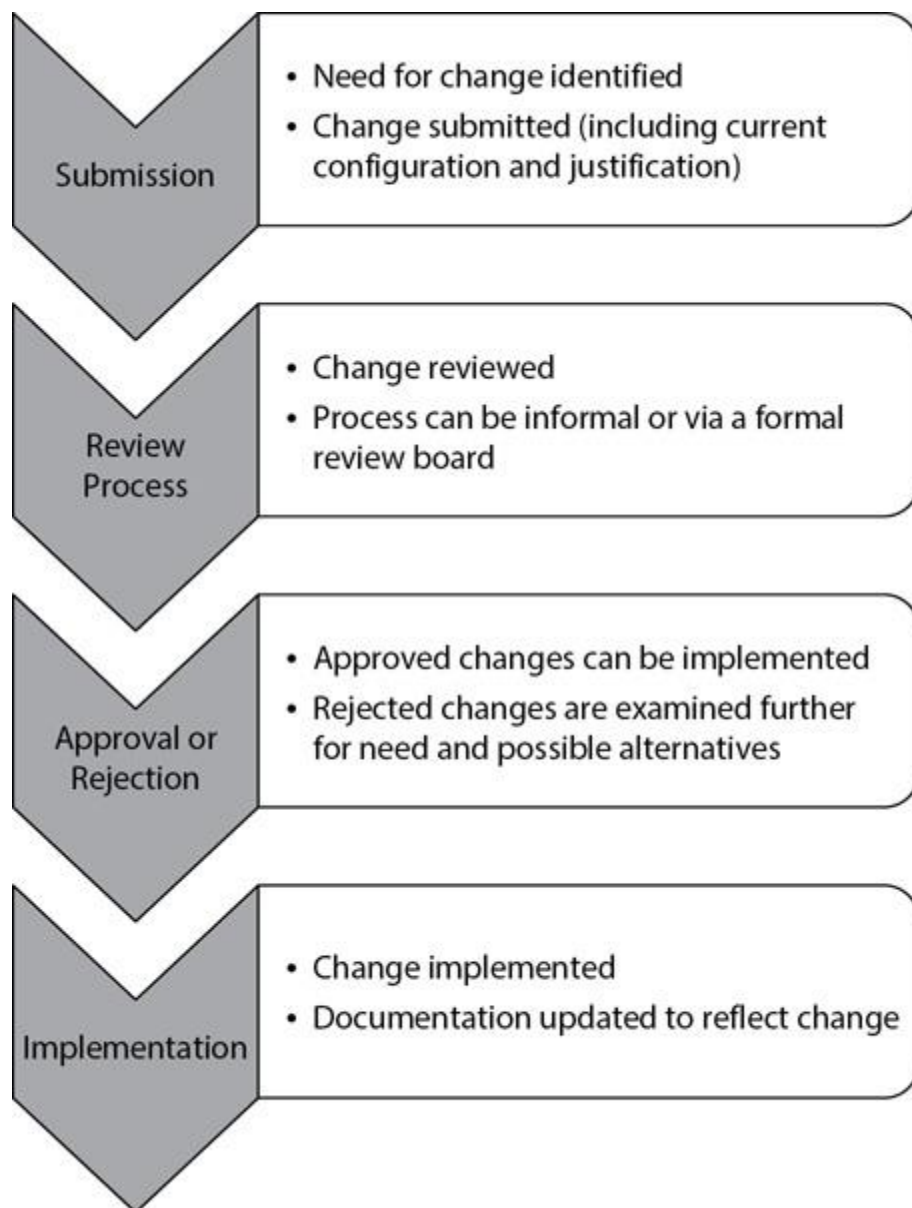


Figure 10-6 Change control process steps

1. Change submitted Someone identifies the need for the change and submits a change request. The change request identifies the current configuration and the desired change, along with justification for the change.

2. Change reviewed The review process allows experts within the organization an opportunity to review the proposed change. For simple changes, this can be a quick, informal process, but for major or potential problematic changes, this can require a thorough review by a formal change review board (CRB) or change authorization board (CAB).

3. Changed approved or rejected The CRB or CAB either approves or rejects the change request. They also analyze rejected changes to see whether a change is required, and if so, they look for suitable alternatives.

4. Change implemented Personnel implement and document the change. It is also important to ensure that appropriate personnel know about the approved change. If administrators don't know about changes, they might notice them and then take steps to undo them. If the change is complex, administrators may implement the change on a test system first to ensure it doesn't cause an unexpected negative impact.

As an example of a change management process, one organization where I worked used an intranet website to submit, review, and approve changes. Experts in different departments reviewed these changes and either approved or suspended them. If each expert approved the request, the CRB approved the request. However, if any expert suspended the request, the CRB reviewed the request at the next meeting.



EXAM TIP A change control process allows stakeholders to propose changes and gives the organization an opportunity to review the changes before implementation. The goal is to reduce unintended outages from unauthorized changes.

Identifying Security Impact

When reviewing change requests, it's important to identify potential security impacts. As an example, imagine someone wants to open incoming port 3389 on a border firewall to enable the use of an application. While it will enable the application, it also opens the port commonly used by the Remote Desktop Protocol (RDP). Many attackers also use RDP to access internal networks, so opening this port brings additional vulnerabilities.

Chapter Review

Auditing is an important component of enforcing accountability. When an organization implements strong authentication and authorization methods, administrators and security personnel can use auditing methods to track user activity and hold individuals accountable for their activity. Most auditing methods use clipping levels to generate alerts only when the system detects the number of events exceeds a predetermined threshold. The auditing system ignores events that don't meet the threshold. Together, logs create an audit trail that administrators and security professionals can use to reconstruct the events that occurred leading up to and during an incident.

Many different types of audit logs exist. Operating system logs record events on individual systems. These include security logs, system logs, and application logs. Security logs record security events, such as when a user accesses a file. System logs record system events, such as

when a service stops or starts. Application logs record application events based on the needs of the application. Proxy logs and firewall logs record activity through proxy servers and firewalls, respectively. An important step related to logs is to review them on a regular basis.

Security audits examine an organization's practices and operations to determine whether they conform to existing policies or applicable laws. Audits are often required to ensure that an organization is complying with specific external requirements (such as PCI DSS) or laws (such as HIPAA), and internal audits show an attempt to comply with these external requirements. In some situations, organizations are required to hire external auditors to perform audits.

Configuration control ensures that systems are configured using a baseline. This ensures that systems are configured in a secure manner and that configuration is similar (if not identical) for similar systems. Imaging is a common method used to deploy systems with similar needs and configuration requirements. Change management helps prevent unauthorized changes resulting in unintended outages, while also identifying potential security impacts.

Questions

1. Your organization uses strong authentication and authorization mechanisms and has robust logging capabilities. Combined, what do these three elements provide?

- A.** Guaranteed security
- B.** Prevention of unintended outages from unauthorized changes
- C.** Accountability
- D.** Configuration control

2. An accounting system ignores logon failures until an account has three logon failures within a 30-minute period. It then generates an alert. What is the accounting system using?

- A.** Account lockout
- B.** Password policy
- C.** Snipping level
- D.** Clipping level

3. Which of the following statements best describes a benefit of using clipping levels?

- A.** Clipping levels ignore baselines and generate alerts when they detect security violations.
- B.** Clipping levels ignore normal user errors but generate alerts when these errors exceed a predetermined threshold.

C. Audit trails use clipping levels to record all potential alerts for accountability.

D. Clipping levels ensure systems generate alerts when they detect any potential security violations.

4. A user entered an incorrect password three times. Now, the user is no longer able to log on. What caused this to occur?

A. Password policy

B. Account lockout policy

C. Clipping level

D. Audit trail

5. What do you call a group of one or more logs used to re-create events leading up to and occurring during an incident?

A. A configuration control program

B. A change management program

C. A security audit

D. An audit trail

6. What type of control is an audit trail?

A. Preventive control

B. Corrective control

C. Detective control

D. Physical access control

7. What type of log on a Microsoft system records auditable events, such as when a user deletes a file?

A. System

B. Security

C. Application

D. Forwarded Events

8. Of the following choices, what is an example of an auditable event logged in an operating system's security log?

- A. Access through a firewall
- B. Accessing a website through a proxy server
- C. Reading a file
- D. The date and time when a service starts

9. Of the following choices, what is the best example of a log used as a deterrent for internal employees?

- A. Proxy server log
- B. Network firewall log
- C. Security audit
- D. Change management log

10. You suspect that many internal systems may be part of a botnet. What log would you review to verify your suspicions?

- A. Network-based firewall logs
- B. Host-based firewall logs
- C. Operating system logs
- D. System security logs

11. Your organization is updating its security policy. Management indicated that they want to address best practices associated with audit logs. Of the following choices, which ones are recommended best practices with audit logs? (Select all that apply.)

- A. Review the logs regularly
- B. Archive logs for later review
- C. Periodically overwrite logs
- D. Store logs on remote servers

12. What is the purpose of reviewing logs?

- A. Detecting potential security events
- B. Preventing potential security events
- C. Correcting potential security events
- D. Deterring potential security events

13. Security professionals within your organization recently completed a security audit. Which of the following are valid steps to take after the audit is complete? (Select three.)

- A. Evaluate controls
- B. Approve changes
- C. Implement fixes
- D. Delete the security audit

14. An organization handles credit card data from customers on a regular basis. What provides the security objectives and requirements that the organization must follow?

- A. PCI DSS
- B. HIPAA
- C. FIPS Pub 200
- D. NIST SP 800-53

15. A badge reader records employee names, dates, and times when employees enter and exit a secure server room. An auditor reviewed the logs and noticed that they showed that many employees entered the room, but the logs do not show when all the employees exited the room. What does this indicate?

- A. The badge reader is operational
- B. Tailgating
- C. The mantrap is not being used
- D. Unauthorized entry

16. Who would measure the effectiveness of an organization's security controls?

- A. An administrator
- B. A manager
- C. An auditor
- D. A data owner

17. Of the following choices, what is a primary method used for configuration control?

- A. Baseline
- B. Change management requests
- C. Security logs
- D. Password audits

18. Which of the following are valid methods used for configuration control? (Select three.)

- A. Imaging
- B. Microsoft's Group Policy
- C. Change management
- D. Proxy server logs

19. Of the following choices, what helps prevent unintended outages caused by system modifications?

- A. Security audit
- B. Change management
- C. Configuration control
- D. Audit trail

20. Application developers are currently testing a new application that salespeople can use while traveling. The application works when it's tested in the internal network, but not when testers attempt to use it by connecting to the internal network via the Internet. Developers recommend a firewall modification to resolve this issue. Which of the following will evaluate potential security impacts from this modification?

- A. Configuration management processes

B. Vulnerability assessment

C. Firewall logs

D. Change management processes

Answers

1. C. Authentication, authorization, and accounting (AAA) provide accountability, and logging provides the accounting element. Although AAA increases security, it does not guarantee security. Change management prevents unintended outages from unauthorized changes. Configuration control ensures systems are deployed with a secure baseline and maintain approved configuration settings for system stability.

2. D. A clipping level uses a predetermined level as a threshold. A classic example is three or five logon failures in a short period, such as within 30 minutes. Although many operating systems use account lockout policies to lock the account after a predetermined level, the question doesn't ask what happens to the account, but instead asks what the accounting system is using to ignore the first two logon failures and only generate the alert after three logon failures. A password policy ensures that users have strong passwords and change them regularly. Snipping level isn't a valid term associated with accounting systems.

3. B. Clipping levels ignore normal user errors, but generate alerts when these errors exceed a predetermined threshold. Clipping levels are not associated with baselines. It is possible to configure an audit trail without clipping levels, but if clipping levels exist, the audit trail does not ignore them. Clipping levels do not generate alerts when they detect any potential errors or security violations, but instead only generate alerts when they detect the number of events has exceeded a predetermined threshold.

4. B. An account lockout policy locks out an account after a predetermined number of failed logins. The password policy ensures that users create strong passwords and change them often. The account lockout policy is using a clipping level by ignoring failed login attempts until it detects a preset threshold, but the clipping level doesn't lock the account. An audit trail is one or more logs used to reconstruct events leading up to and occurring during an incident.

5. D. An audit trail is one or more logs that can re-create events leading up to and occurring during an incident. Configuration control helps ensure that systems are configured in a secure manner. A change management program allows stakeholders to request changes and helps reduce unintended outages from unauthorized changes. A security audit examines an organization's policies and procedures to determine whether the organization follows these policies and procedures.

6. C. An audit trail is a technical detective control, because it uses technology and can detect incidents after they occur. A preventive control attempts to prevent incidents. A corrective control attempts to reverse the impact of an incident after it has occurred. A physical access control is an item that you can physically touch.

7. B. The Security log records auditable events, such as when a user accesses or deletes a file (if resource auditing is enabled). The System log records system events such as when a service stops or starts. The Application log records application events. The Forwarded Events log shows events forwarded from other systems as part of an event subscription.

8. C. A security log records auditable events related to resources, such as when a user reads, modifies, or deletes a file. Firewall and proxy server logs are not operating system logs. A system log would record events such as when a service stops or starts, but not security events.

9. A. A proxy server log can serve as a deterrent for internal employees. If employees know that the server is monitoring and logging their activity, they may be less likely to engage in activity that violates the security policy. A network firewall log can capture activity for traffic to and from the Internet but does not provide much of a deterrent for internal employees. A security audit is not a log. A change management log documents changes for a system.

10. A. Network-based firewall logs record traffic on the network, and because many systems are involved, network-based firewall logs is the best choice. Each of the other logs are local logs on individual systems. This would require checking logs on multiple systems, rather than checking logs on a single network-based firewall.

11. A, B, D. Reviewing the logs, archiving the logs, and storing logs on remote servers are recommended strategies to retain the integrity of audit logs. If you periodically overwrite the logs without first archiving them, it's possible to overwrite log entries, erasing them forever.

12. A. Security professionals and auditors can detect potential security events by reviewing logs after the event has occurred. Reviewing the logs doesn't prevent an incident that has already occurred and reviewing the logs does not enable security professionals and auditors to correct the effects of an incident. While logging some activity, such as proxy servers, can deter events, reviewing the logs doesn't deter the activity.

13. A, B, C. A security audit often recommends changes in the form of additional security controls. Personnel evaluate the recommended controls, management approves changes, and personnel then implement the fixes. It isn't appropriate to delete the audit after it has been completed.

14. A. The Payment Card Industry Data Security Standard (PCI DSS) provides 6 control objectives and 12 supporting requirements that organizations must follow if they process credit card payments from customers. The Health Insurance Portability and Accountability Act (HIPAA) covers organizations handling health- and medical-related data. Federal Information Processing Standard Publication 200 (FIPS Pub 200) identifies standards required by federal agencies. NIST SP 800-53r4 provides information on recommended security controls.

15. B. Logs that include entries showing employees entered a secure area but do not include entries showing they exited indicate tailgating is occurring. Some employees are using their credentials to exit (and the logs show them exiting), but other employees are following closely behind these employees without showing their credentials (and the log doesn't include entries for

these employees). While it is possible the badge reader has a problem, it is recording some employees exiting, so this isn't the most likely cause. Mantraps prevent tailgating, and if a mantrap is in use, employees would be forced to use it. There isn't any indication of unauthorized entry.

16. C. An auditor would measure the effectiveness of a security control. An internal auditor might have other roles, such as an administrator, a manager, or a data owner. However, when measuring the effectiveness of security controls, they are acting as an auditor.

17. A. A baseline is a primary method used for configuration control and it ensures that systems start in a known state. Automated or manual processes periodically examine the systems to verify the system still has the same configuration settings from the baseline. An organization doesn't approve and implement all change management requests, so examining the requests does not give an accurate representation of the server configuration. Security logs and password audits aren't typically used for configuration control.

18. A, B, C. Imaging, Microsoft's Group Policy, and change management processes are all used for configuration control. Proxy server logs are often used for auditing, but not for configuration control.

19. B. A change management program allows stakeholders to request changes and helps reduce unintended outages from unauthorized changes. A security audit examines an organization's policies and procedures to determine whether those who work in the organization follow these policies and procedures. Configuration control helps ensure that systems are configured in a secure manner and similarly to each other. An audit trail is one or more logs that can re-create events leading up to and occurring during an incident.

20. D. Change management processes evaluate recommended modifications and changes and can identify potential security impacts from a change. Configuration management processes help ensure systems are configured consistently. A vulnerability assessment evaluates a system or network for weaknesses. Firewall logs can show activity, but they won't evaluate potential changes.

NOTES

CHAPTER 11

Security Operations

In this chapter, you will learn about

- Data classifications such as Confidential, Sensitive, Private, and Public
- The importance of marking and labeling data with its classification
- The difference between data at rest and data in motion
- Different elements of data management policies
- Components of a database such as tuples, rows, primary keys, and foreign keys
- Risks related to data inference and data diddling
- Regulatory requirements related to PII and PHI
- Asset management components
- The differences between certification and accreditation
- The Common Criteria and its evaluation levels
- Using a risk management framework with a certification and accreditation process
- Different phases of a system development lifecycle

Handling Data

Data is one of the most valuable resources an organization has, so it's important to understand the value of data and some of the risks related to data. Different types of data have different values, so an organization protects data based on its value. Highly classified or sensitive data requires a higher level of protection. Similarly, when an organization no longer needs particular data, it uses one of several different methods to destroy the data depending on its value.

Classifying Data

One of the first steps in protecting data is identifying what data needs protection by classifying it. For example, an organization will often take additional steps to encrypt and protect highly classified data, such as research and development data that took years to develop. However, data published in company brochures or company websites does not deserve the same level of protection. By classifying data, it becomes easier for anyone within the organization to identify its worth to the organization and the level of protection it deserves.

For example, the U.S. government classifies data using terms such as Top Secret, Secret, Confidential, and Unclassified. The higher the classification, the more protection it requires. These aren't the only U.S. government classifications, but they provide a good example. The following list shows the definitions assigned to each of these classifications:

- **Top Secret** The public disclosure of Top Secret material could cause exceptionally grave damage to U.S. national security.
- **Secret** The public disclosure of Secret material could cause serious damage to U.S. national security.
- **Confidential** The public disclosure of Confidential material could cause damage to U.S. national security.
- **Unclassified** The public disclosure of Unclassified material will not cause damage to U.S. national security. People can obtain information in this category through the Freedom of Information Act.

To have access to any of the classified material, an individual must have the associated clearance. For example, someone with a Secret clearance may be able to view Secret material, but not Top-Secret material.



TIP The principle of need-to-know is widely used to protect data. For example, just because individuals have a Secret clearance doesn't mean they can access all Secret material. The clearance level only indicates they are cleared to view the material if their job requires it, but it does not automatically grant access to the material. If their job does require access to the material, an administrator grants them appropriate access.

Many private organizations also classify data, although they rarely use the same Top Secret, Secret, Confidential, and Unclassified labels used by the U.S. government. An important point to remember is that organizations can use any labels they desire. For example, an organization can use Class 0, Class 1, Class 2, and Class 3. However, end users might find them to be a little mysterious. What's the highest classification, Class 0 or Class 3? One organization might say Class 0 is highest, and another organization might say Class 3 is highest. Of course, both organizations are correct, because the organization defines the labels in a security policy or data policy.

Data owners are responsible for identifying data classifications. In civilian organizations, data owners are typically the chief executive officer (CEO), chief information officer (CIO), and/or department heads. A data owner ensures that a security policy or data policy includes the labels and accompanying descriptions.

Here are some examples of data classifications that organizations can implement, listed from highest classifications to lowest:

- Confidential, Private, Sensitive, Public
- Proprietary, Internal Use Only, Sensitive, Public
- Private Highly Restricted, Private Restricted, Private, Unrestricted
- Highly Confidential, Restricted, Private, Public
- Regulatory, Confidential, Internal Use Only, Public
- Confidential, Restricted, Internal Use Only, Public
- Highly Confidential, Restricted, Private, Unrestricted

[Figure 11-1](#) shows the Confidential, Private, Sensitive, and Public labels that any organization can use. It also includes possible definitions that an organization can use within its security policy or data policy. Notice that Confidential data is at the core and deserves the highest level of protection.

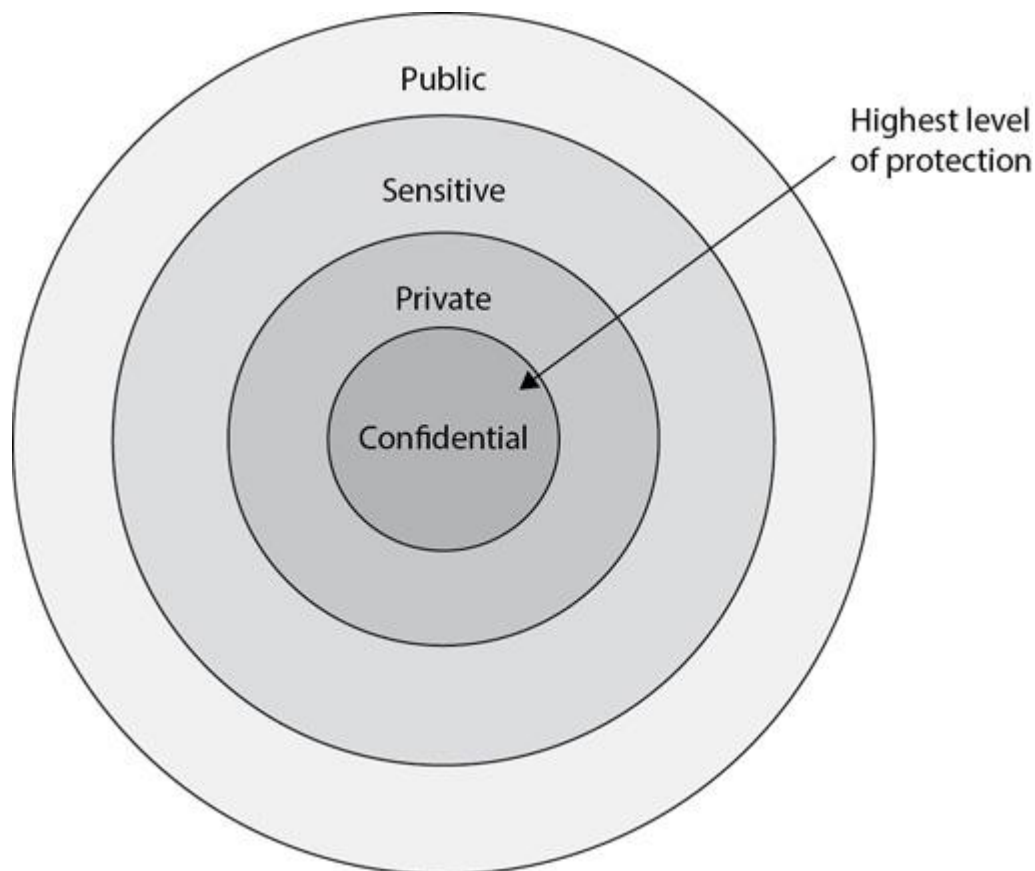


Figure 11-1 Data classifications used by some organizations

- **Confidential** This includes proprietary information central to the operation of a company. It could include research and development information for upcoming projects, trade secrets for products sold by the company, and so forth. If someone discloses an organization's confidential data, it could result in an exceptionally grave negative impact on the mission of the organization.
- **Private** This includes information used privately within the company, such as employee records and customer data. Private data is intended for internal use only, and if it is disclosed, it could result in a serious negative impact on the mission of the organization.
- **Sensitive** This includes information that requires special precautions for its protection, such as financial information about a company. If someone discloses sensitive data, it could result in a negative impact for the organization.
- **Public** This is data that is either publicly available or would not cause any harm if it was publicly available. If disclosed, this data would not have any negative impact on the organization.



TIP Organizations aren't required to use any specific data labels and there isn't a universally accepted standard, so it's entirely possible an organization could have different labels or define these a little differently. The important thing to grasp is that organizations do classify material and provide higher levels of protection for data that requires it.

Marking and Labeling Data

After identifying data classifications, personnel mark or label the data. In general, *marking* refers to using digital methods to identify the data classification. *Labeling* refers to using physical labels to identify the data classification. This ensures that users can immediately identify the classification files and media holding the files. For example, after printing classified data, users may be required to place it within a folder with a cover sheet indicating the classification of the data. Digital files can include watermarks, headers, or footers that show the classification so that it is easily identifiable after a user prints it.

It's also important to label hardware that holds or processes classified information. This includes portable media such as USB drives and tapes, cabinets where portable media is stored, and computer systems that process classified data. Organizations commonly use color-coded stickers with the classification name. As an example, they could use red stickers with the label Proprietary. This allows employees to identify the classification based on the name and the color of the label.

An alternative is to purchase a large quantity of drives (such as red USB drives) and train employees to use only these drives for classified data. It's also possible for administrators to configure technical policies to encrypt data on these drives automatically. These technical policies identify the drives using a universally unique identifier (UUID).

Roles and Responsibilities

Multiple people within an organization have different responsibilities for how data is handled, and those responsibilities can be identified by their role. Some of the common roles and responsibilities are

- **Management** Personnel in management and leadership positions define the data classifications and specify the requirements to protect data with different classifications. A data policy or security policy includes these definitions, and management personnel ensure the policy is available to everyone within the organization.
- **Data owner** Data owners have primary responsibility for protecting the data based on its classification and requirements stated within a security policy. When data owners create data, they identify the appropriate classification and mark or label the data according to requirements dictated in the data policy. They periodically review the data and modify the classification if necessary based on changes within the organization.
- **Custodian** In some cases, the data owner delegates various data-related tasks to a data custodian. For example, a data custodian could perform regular backups of the data based on requirements stated within a backup policy.
- **Administrator** Administrators use available access control mechanisms to provide access to the data to personnel who need access. For example, administrators assign permissions to personnel to grant access, typically by direction from the data owner.
- **User** Users access the data. Users are also responsible for protecting the data they handle.



TIP Data owners are ultimately responsible for protecting the data. This includes ensuring that it is marked or labeled correctly and that security controls match the requirements of the data classification.

Protecting Data from Cradle to Grave

Personnel protect classified data from the moment someone creates it until someone destroys it. The level of protection is based on the classification. If classified data is stored or processed on a computer system, then the system must be protected at least at the same level as the classified

data. Similarly, if classified data is transmitted over a network, all access points on the network must also be protected to at least the same level.

If a system or hard drive fails, it's important to remove all remnants of data before disposing of it. For example, just because a hard drive fails during normal operation, that doesn't mean that all the data on the drive is inaccessible. Experts working in a clean room can open failed disks and often retrieve data remnants. Although using a clean room to access data on a failed disk is expensive, it can be worthwhile to attackers if they can retrieve sensitive data.



TIP The “Removing Data Remnants” section later in this chapter covers various methods of removing residual data.

Data at Rest and Data in Motion

Data at rest is any data that is in computer storage, such as on system hard drives, portable USB drives, flash drives, storage area networks, backup tapes, and so on. Data in motion (sometimes called data in transit) is any data transmitted over a network. This includes data transmitted over an internal network using wired or wireless methods and data transmitted over public networks such as the Internet.

Encryption is often used to protect data at rest and data in motion. Encryption converts plain-text data into cipher text, making it unreadable until it is decrypted. Strong encryption algorithms ensure that unauthorized entities are unable to read encrypted data.



EXAM TIP Encryption protocols such as Advanced Encryption Standard (AES) encrypt data at rest. Protocols such as Transport Layer Security (TLS) and Internet Protocol security (IPsec) protect data in motion. [Chapter 14](#) covers these and other encryption protocols.

Additionally, access control methods such as strong authentication and authorization practices help to protect data. Access controls help ensure that only authorized people have access to a network where data is stored and transmitted. However, access controls aren't useful when transmitting data over a public network.

It's also important to protect data during active processing (data in use), and this is a function of the application processing the data. For example, imagine a web application retrieves customer data from a database, including credit card data, to process a purchase. Application developers often use specific buffers for sensitive data and purge these buffers when the data is no longer needed. In this example, the data goes through several stages:

- **Data at rest**• Sensitive data is stored in an encrypted format within the database. However, the database decrypts it before sending it to the web application.
- **Data in motion**• The web server and database server create a secure channel to transfer the data. For example, they can use TLS to encrypt the sensitive data while transmitting it over the network.
- **Data in use**• The web server decrypts the sensitive data and the web application processes it in an unencrypted format. When the web application is done with the data, it clears the data buffers holding this data.

Data Management Policies

Many organizations implement data management policies to help employees understand the value of different types of data. These policies also identify methods of protecting sensitive data. These policies help an organization protect data from unauthorized disclosure and misuse throughout the lifetime of the data. The following subsections identify some of the common elements of data management policies.

Storage Media

Storage media refers to how data at rest is stored. It can refer to data on hard drives, solid state drives (SSDs), portable USB drives, USB flash drives, optical disks, backup tapes, or any other media storing data. Data management policies may restrict the use of some media and require special handling of any media that holds sensitive data.

One of the biggest risks of most storage media is that it's highly portable. It's very easy for someone to copy mass amounts of data onto the media, and the organization can lose control of it. For example, a malicious insider can steal valuable data by copying it onto a portable USB drive. It's also possible for an honest employee to lose a USB flash drive after copying data onto it.

A common way to protect the confidentiality of sensitive data on storage media is with encryption. For example, many USB drives are available that automatically encrypt data when users copy data to them. Several technical security controls allow administrators to configure whole disk encryption for portable devices, and hard drives within computer systems.

In addition to encryption, the media should be labeled to indicate the highest classification of data stored on the media. When the media is not in use, employees should secure it in a storage cabinet or safe.

Transmission

Data transmitted over a network is data in motion. Attackers can use sniffers or protocol analyzers (covered in [Chapter 5](#)) to capture and read the data transmitted in cleartext. However,

encrypting sensitive data before transmitting it thwarts a sniffing attack. Data management policies dictate the minimum encryption methods of specific types of data.

[Chapter 3](#) introduced several encryption protocols that encrypt data before transmission. Secure Shell (SSH) is often used to encrypt traffic such as Telnet and File Transport Protocol (FTP) on internal networks. IPsec can be used to encrypt traffic in internal networks and traffic sent over the Internet, such as with virtual private networks (VPNs). TLS is commonly used to encrypt data transmitted over the Internet, such as with Hypertext Transport Protocol Secure (HTTPS).

Archiving and Retention Requirements

Archiving refers to making a backup of data and keeping it for long-term storage. [Chapter 9](#) covers backup strategies in more depth. *Retention* refers to how long archives and other data are kept. Data policies often identify both archiving and retention requirements.

Laws or regulations may require an organization to keep data for a minimum timeframe, and sometimes indefinitely. For example, laws in the United States require that e-mails sent to and from officials in several U.S. public offices be kept forever. These archives are kept for historical purposes. Similarly, many banking regulations require institutions to keep certain data for three, five, or seven years.

Retention requirements also identify how long to retain backups. For example, an organization can decide to keep copies of backups for the last three years or the last 30 days. Many factors go into these decisions, and an important one is cost. The number of backup tapes required to keep archive copies of backups for the last three years can be quite high.

Additionally, many companies consider legal consequences of long-term storage of data. For example, a court order might direct a company to turn over all e-mail sent or received by a certain individual over the last five years. However, if the company retention policy states that e-mail is only kept for a year, the company only needs to retrieve e-mails from the last year. This assumes that personnel are following the policy and the company has only a year's worth of e-mails. In contrast, if the organization retains e-mail for five years, IT personnel will have to sift through five years of e-mails to comply with the court order. This can be time consuming and labor intensive.

Removing Data Remnants

Data remnants refers to the residual data left on media. Data management policies usually include information on how to eliminate data remnants. This allows personnel to reuse the media or safely dispose of the media without losing confidentiality of the data. Methods of eliminating data remnants include sanitizing (or purging) the media and destroying it.



TIP Residual data is sometimes referred to as data remanence. However, remanence technically refers to residual magnetic flux on traditional hard drives. It does not apply to other media such as CDs, DVDs, and solid state drives.

- **Sanitizing (or purging)** Sanitization methods remove all residual elements of data, preventing someone from accessing it. After sanitizing media, personnel can reuse it or dispose of it.
- **Destruction** In some cases, sanitizing methods are not reliable enough, so personnel destroy the media instead. For example, if media holds highly classified data, an organization's data policy may require personnel to destroy the media instead of using methods to purge it of all data remnants.

If employees throw paper documents into the trash, dumpster divers can retrieve them. Instead, employees should destroy paper documents containing sensitive data either by passing them through a crosscut shredder or by burning them in an incinerator. It's important to use a crosscut shredder because it cuts the paper into smaller pieces that resemble confetti. In contrast, strip-cut shredders cut the paper into long strips, which a patient attacker can reconstruct more easily.

Most administrators understand that simply deleting a file or formatting a drive doesn't remove the data. Someone with the right software and a little bit of expertise can often retrieve the data. To combat this, organizations implement different procedures to remove data. In some cases, an organization approves a specific application to write a pattern of ones and zeros to remove data. These programs can "shred" a single file by overwriting it, or completely erase the entire disk by overwriting it.

Magnetic-based media such as traditional hard drives and backup tapes store data by magnetizing elements of the media. *Degaussing* uses a powerful magnet to remove data and is an effective method of sanitizing backup tapes. You can also degauss hard drives, and this usually destroys the drives. Optical media such as DVDs and CDs do not use magnetic methods, so degaussing does not remove data from this type of media. Similarly, degaussing won't remove data from an SSD because SSDs use integrated circuitry instead of the magnetic flux used on hard drives. Instead, optical media and SSDs must be physically destroyed.

SSDs present a serious risk to an organization when it is sanitizing systems. If you put a traditional hard drive next to an SSD, they look the same. Personnel might think that they can degauss both and then throw them away. However, the SSD will still hold all the data.

When destroying media, the goal is to ensure it is unusable, which is the most secure method of protecting confidentiality of data. Destruction methods include shredding it (for media such as optical media) and burning or incinerating it (for backup tapes). Policies often require personnel to open hard drives and destroy the platters. For example, personnel can remove the platters and send them through large shredders or incinerate them.



TIP Data policies identify how to destroy data and eliminate data remnants. It's common to degauss backup tapes to remove data. Programs that repeatedly write a series of ones and zeros onto a drive overwrite or shred files.

Cloud Computing

As an example of a data management policy, the U.S. Department of Defense (DoD) released its *Cloud Computing Security Requirements Guide* in January 2015. It identifies specific requirements for storing data in the cloud.

The guide includes six information impact levels. It defines data within each level and defines the security controls required when storing the data in the cloud. For example, Level 4 is Controlled Unclassified Information such as personally identifiable information (PII) or protected health information (PHI), and Level 6 includes any classified information up to Secret.

Requirements to store the data increase for higher levels. For example, the list of requirements for storing Secret Level 6 data in the cloud is quite extensive and includes several elements such as legal jurisdiction and location requirements. Secret Level 6 data must be located in a facility under the legal jurisdiction of the United States.

The guide also identifies specific requirements for encryption of data, destruction of data, and disposal of storage media. For example, personnel can reuse media holding Level 4 data after purging the media using an acceptable method. In contrast, media holding Level 6 data cannot be reused to hold anything other than Level 6 data.

Deduplication

It's common for multiple individuals to have a copy of the same file. If this file is stored in multiple locations, it can take up a significant amount of wasted storage space. *Deduplication* refers to ensuring that a file is stored only once on a system, even if multiple users have access to the same file.

For example, employees commonly have private shares on servers where they can store their data. Imagine if 100 employees have shares on the same server, and each employee has a copy of a 5MB company policy document. This single file takes up 500MB of space if there are 100 separate copies of it. In contrast, with deduplication the file will be stored only once and take up only 5MB of storage. Each employee will appear to have a copy of the document in their private share. However, the file that the employees see acts as a pointer to the single document.

Additionally, many deduplication systems support multiple versions of the same file without keeping multiple copies of the file. Consider the same 5MB file. If five users make slight modifications to their copy, a deduplication system can record their changes in separate files.

The system keeps the original 5MB file intact. If any of these five users opens the file, the system applies the changes that the specific user has made to the file. This is all automatic and often done without the user's knowledge.

There is a trade-off between deduplication and encryption. If the files are encrypted (and the deduplication software cannot decrypt them), then the deduplication software cannot identify any two files that are the same. While there are methods of weakening the encryption so that the deduplication software can identify identical files, this sacrifices security. Any method of weakening the encryption increases the possibility that an attacker can decrypt the data. Unless you weaken security in encrypted files (certainly not recommended), you'll rarely see gains from deduplication software.

Data Loss Prevention

Data loss prevention (DLP) systems attempt to monitor data usage and prevent the unauthorized use or transmission of sensitive data. Different types of DLP systems can monitor data in motion and data at rest.

Network-based DLP devices have the ability to scan all traffic that exits a network. This includes e-mail and attachments. These devices can look for specific data and block it before it leaves the network.

For example, a company may want to ensure that e-mail or attachments never include Social Security numbers. A Social Security number is in the format of *xxx-xx-xxxx*. The device can scan all outgoing traffic looking for this match. When it detects a match, it blocks the traffic from leaving the network and sends an alert to an administrator. Similarly, if a company uses classification labels such as Confidential, Private, and Sensitive, it can configure the device to look for these words along with any other words or patterns of interest.

Some DLP solutions can limit the use of hardware. For example, an organization may want to prevent users from copying certain data to external USB drives or printing this data to a printer. The DLP software scans the data before copying it or sending it to the printer. If it includes restricted data, such as Social Security numbers or specific classification labels, the DLP software blocks it.

Software DLP methods have the ability to scan data stored on a system to detect inappropriate data. As a simple example, an application can scan data stored on a shared server, looking for MP3 files or other data types that the company has stated should not be stored there. These applications can monitor all data stored on a system and periodically provide reports to administrators.

DLP systems can usually perform deep inspection of the data. For example, even if several files are embedded in a ZIP file or other archive file type, the DLP system can still extract the data to inspect it.



EXAM TIP Data loss prevention (DLP) systems help prevent data from leaving a network. A network-based DLP can scan all traffic looking for specific data and block transmissions when it detects restricted data transmissions.

Social Network Usage

Social networks also present risks related to the loss of data. These include sites such as Facebook, Twitter, and LinkedIn.

[Chapter 5](#) mentions some of the risks with social networking sites in the context of social engineering and phishing. Attackers frequently send phishing messages to users of social networking sites, attempting to trick them into following a link to a malicious website or downloading malicious software. However, there are additional risks associated with data from social networking sites.

For example, a chain letter was floating around asking people on Facebook to list ten things about themselves that others don't know. It included questions about the person's favorite things (like favorite color or favorite movie), the school(s) from which they graduated, and the name of their first pet. Users were encouraged to fill out the questionnaire, send it back, and forward it to ten of their friends.



TIP Many websites use cognitive passwords created with multiple personal questions during registration. When users want to retrieve a forgotten password, the system queries them with these questions. As long as users answer the questions correctly, the system allows them to reset their password.

People who answered were at risk of giving out their answers to secret questions at banks and other online sites. For example, if you forget your password or log on to a bank's website from a different computer, the site will often challenge you. You're required to give the name of your first pet, for example, and if you can answer correctly, the system grants you access. If attackers have this information, they can use it to hack into your accounts.

Employees can accidentally provide inside information about their company when they spend time on these social networking sites. Because of the potential damage, many organizations prohibit the use of these types of sites on company computers. The rules are spelled out in acceptable use policies, which employees read and acknowledge at least once a year.

Individuals are often tricked into giving up highly personal information on themselves. As an example, Cambridge Analytica reportedly published one or more personality quizzes that captured information on over 50 million Facebook users as early as 2015. Users responded thinking it would be fun to get a free analysis of their personalities. However, Cambridge Analytica allegedly used this information to target ads to users based on their responses. While Cambridge Analytica's practices made the headlines, many organizations are using similar practices to capture data about users.

Social media sites are free to use, and when they're well designed, they can become quite popular. Facebook reportedly has over 2 billion users worldwide. And while the site is free to users, it's worthwhile remembering the old phrase "if you're not paying for the product, you are the product." While the phrase doesn't apply 100 percent to all free sites, it is worth remembering any time users are asked to provide personal information on themselves.

Information Rights Management

Information rights management (IRM) refers to the different methods used to protect sensitive information from unauthorized access. IRM is a subset of digital rights management (DRM). DRM is much broader and typically refers to protecting media such as music and video. IRM focuses on documents. It allows people to specify access permissions to different documents that they share, even after the documents leave their control.

For example, a user can protect a document with IRM before e-mailing it to an associate at a different company. IRM permissions can be applied to the document before it is sent, restricting what the recipient can do. IRM can prevent recipients from forwarding, copying, modifying, printing, faxing, or pasting the content if that's what the sender wants to restrict. It's also possible to restrict users from copying the content using the Print Screen functionality and to set expiration dates on files. After the expiration date, users are unable to open the file.

Understanding Databases

Databases are stores of information that users can easily access and modify as needed. Large databases require sophisticated database management systems (DBMSs) to manage database access. Some common DBMSs in use today include products from Oracle, IBM, and Microsoft. These systems provide access to relational databases.

Understanding Database Tables

In a relational database, data is stored in tables, and this data is linked to other tables based on predefined relations. As an example, [Figure 11-2](#) shows a Parts table and a Suppliers table with different elements identified. The Parts table holds information on parts, such as parts used for manufacturing a bicycle. The Suppliers table holds information on the company that supplies these parts. The SupplierID field links the two tables.

| Parts Table | | | | | |
|--------------|-------------|-----------|---------------------|-------------|-----|
| | Primary Key | | Column or Attribute | Foreign Key | |
| | PartID | Part Name | QtyOnHand | SupplierID | ... |
| | 123 | Widget | 12 | 745 | |
| Row or Tuple | 124 | Sprocket | 23 | 745 | |
| | 125 | Wheel | 15 | 745 | |
| | 126 | Horn | 4 | 723 | |
| | ... | | | | |

Field

| Suppliers Table | | | |
|-----------------|-------------|---------------|--------------|
| | Primary Key | | |
| | SupplierID | Supplier Name | Phone |
| | 723 | We Make Noise | 555-765-4321 |
| | 745 | Acme Supplies | 555-123-4567 |
| | ... | | |

Figure 11-2 Tables in a database



TIP Actual databases will have many more than just two tables. Many production databases used within organizations have hundreds of tables. A data dictionary describes the purpose of each table and its relationship with other tables. The ellipsis (...) in [Figure 11-2](#) indicates that the table can have many more columns and many more rows, but for brevity, the figure only shows a few of the columns and rows. Some tables have dozens of columns, and they can have millions of rows.

Key elements of relational databases include the following:

- **Primary key** The primary key is used to uniquely identify each row or tuple. It ensures that each item has only a single row and that the item can easily be identified using the primary key. A primary key is normally an integer, but it can be any string of characters. In the Parts table, the

primary key for the sprocket is 124, which will always reference the sprocket part and nothing else.

- **Row or tuple** Each row, or tuple, contains a unique data element. For example, the row for a sprocket identifies details about a sprocket identified with a primary key of 124. The row also identifies how many sprockets are currently on hand in the QtyOnHand column.
- **Column or attribute** Columns, or attributes, provide additional details on the data. For example, the Parts table could have additional columns, such as the cost of the part, a description, the color, and more.
- **Foreign key** A foreign key provides a relationship to another table. Notice that three of the parts have the same SupplierID of 745. The Suppliers table has the details of the supplier with a SupplierID of 745 (Acme Supplies). The foreign key in the Parts table points to the primary key in the Suppliers table. This allows you to enter the details for the suppliers once in the Suppliers table, instead of multiple times (in each row of the Parts table).
- **Field** A field provides an individual piece of information within any row or tuple. For example, the field that matches the primary key of 123 and the Part Name is Widget.

Understanding Relations

Tables are related to each other using the primary and foreign keys. The primary key ensures that each row or tuple within a table is unique. The foreign key in one table points to a unique primary key in another table to create the relationship. For example, [Figure 11-3](#) shows the foreign key in the Parts table linked to, or related with, the primary key in the Suppliers table.

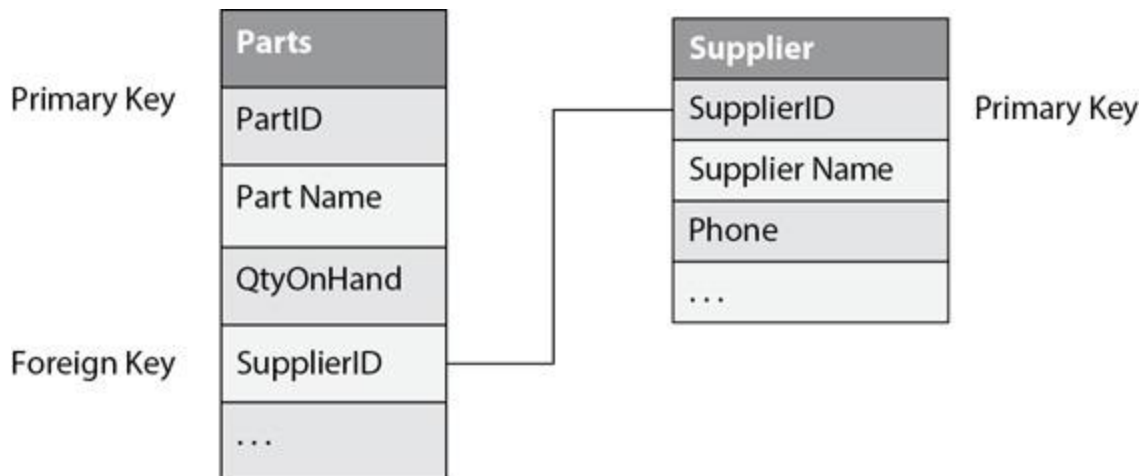


Figure 11-3 Table relationships

One of the benefits of the primary key and foreign key relationship is that data doesn't need to be repeated. For example, if you look back at [Figure 11-2](#), you can see that the supplier with a SupplierID of 745 is supplying the first three parts. The Suppliers table shows that SupplierID

745 is Acme Supplies. If you entered all of the information on Acme Supplies (such as name, phone number, address, sales contact, and so on) in each row of the Parts table, it would become a lot of needless work. You would have to repeat the same information in each of the rows, and each time that you reentered the data in a different row, the likelihood of making mistakes would increase.

However, by creating the relationship between the tables, you need to enter the data only once. If any of the data on Acme Supplies ever changes (such as a different sales contact or phone number), you only need to make the change once in the Suppliers table. If the supplier data were within the Parts table, you'd need to make the change in every relevant row or tuple of the Parts table.

Data normalization is the process of separating data into multiple related tables to reduce duplication of data. Duplicating the same data in more than one table creates dual entry (or multiple entry) problems, but data normalization eliminates these issues.



EXAM TIP Rows within a database are known as tuples. A primary key uniquely identifies each row in a table. A foreign key in one table points to a primary key in another table to link the tables together through a common relationship.

Using Views

One way that databases provide security is with views. A *view* is a virtual table that provides access to specific columns in one or more tables. It doesn't actually hold any data but presents the data in the underlying table or tables. A database administrator can grant access to a view without granting access to a table, thereby limiting what a user can see and manipulate.

For example, imagine that people in the company need to access phone numbers for employees and this information is contained within a table named Employee in a company database. However, the Employee table holds more information (such as salary data), and administrators don't want to grant users access to this additional information. Instead of granting access to the table, a database administrator can create a view and grant users access to the view, as shown in [Figure 11-4](#).

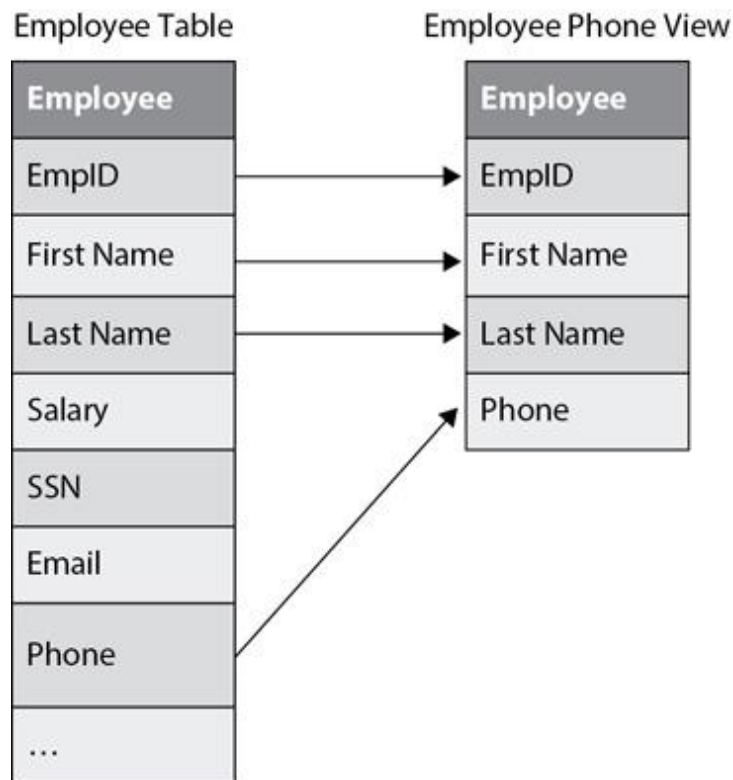


Figure 11-4 Database view

Notice that the Employee table has other information such as salary and Social Security number (SSN) data that should remain confidential. The view shows several columns from the table, including the employee's name and phone number, without showing the confidential data.

Additionally, database administrators can grant access based on the users' needs. For example, an administrator can grant access for users to read or modify the data in the view. A key point is that users will only be able to read and modify columns in the view. When users modify data in a view, they modify the data in the underlying table. However, the view still restricts their access to other columns in the underlying table.

Communicating with Databases

Relational databases have their own language used to communicate with them, known as Structured Query Language (SQL). Although different companies often include their own unique additions to enhance SQL, they share some commonalities. SQL statements are primarily divided into two categories:

- **Data Definition Language (DDL)** DDL statements are used to create the structure of the database, including tables, primary keys, and relationships among the tables.
- **Data Manipulation Language (DML)** DML statements are used to add, modify, retrieve, and delete data from the databases. For example, a SELECT statement retrieves data, an

INSERT statement adds data, an UPDATE statement modifies data, and a DELETE statement deletes data.

OLTP vs. OLAP

Databases are organized differently depending on how they're used. There are two primary methods:

- **Online transaction processing (OLTP)** E-commerce websites that access backend databases use OLTP. For example, each time a user makes a purchase, the web application records the transaction in the database as it occurs. OLTP databases need to be quick, and data normalization is one technique that ensures they provide the best performance. Additionally, OLTP databases record all transactions in logs before making the changes in the database. This allows the database management system to easily reverse failed transactions before writing them to the database.
- **Online analytical processing (OLAP)** Many organizations have very large databases known as data warehouses. Retrieving usable data from these data warehouses often requires additional steps. OLAP reorganizes data from data warehouses into multidimensional cubes for easier retrieval of useful information. These cubes are not normalized, but instead include redundant data. They are typically slower than OLTP databases, but speed isn't the most important consideration. However, business managers are able to access a significant amount of information and convert it into actionable data from OLAP databases. For example, a manager might want to compare how many blue widgets the company sold last December to how many green widgets the company sold last December. This helps the manager determine how many widgets of each color to preorder for this December. This process is also known as *data mining*.

Data Inference

One of the security issues with data is *data inference*, the ability of someone to gain knowledge by piecing together unclassified data to determine classified or secret information.

As an example of data inference, the U.S. Air Force once had a problem with snipers taking potshots at jets as they flew sorties out of a certain airfield in a foreign country. The Air Force sent patrols out when sorties weren't planned, but the snipers only appeared when the jets were flying sorties. Even though the times and dates of the sorties were highly classified, the enemy learned of these sorties by accessing unclassified information.

Because the times and dates of the sorties were not consistent from day to day or even week to week, pilots and crew were notified when to report to work through schedules posted on a bulletin board. A foreign worker was able to view these bulletin boards and passed the information on to other foreigners working as spies collecting information. The spies were able to deduce that the posted work times corresponded to the times of the sorties. Although this is a lesson that is sometimes forgotten, military organizations recognize that any piece of information can be valuable to an enemy and take steps to protect all information. The U.S. Navy had a

simple phrase to emphasize this during World War II, “loose lips sink ships,” as part of a larger campaign known as “careless talk costs lives.”



EXAM TIP Data inference attacks attempt to collect public or unclassified pieces of information to predict or guess an outcome. They can also attempt to gather large quantities of information and attempt to learn details of the information through deduction.

Data inference can sometimes work in the reverse also. In other words, if people have access to summary data (also called aggregate data), they can sometimes use this to learn details. For example, if personnel in accounting have access to total payroll data for any given day, they wouldn't necessarily know the actual pay for an individual. However, examining the total pay for the day before a person is hired with the total pay the day after a person is hired may allow someone to deduce the actual pay of the new hire.

Data Diddling

Data diddling is the unauthorized changing of data before entering it into a system or while entering it in a system. It is a form of fraud, and anyone involved with data entry has the opportunity to change it.

As an example, imagine Bob works in accounting and handles payroll. He is responsible for manually approving and entering overtime hours for employees. While entering the data, he adds more hours for his friend, which increases his friend's paycheck. Similarly, he and his friend could decide to do this together. Bob would increase his friend's overtime hours, and his friend would give him half of the extra pay.

Regulatory Requirements

Laws mandate the protection of some data. This includes personally identifiable information (PII), health- or medical-related data, and other types of data. This section covers some of the regulatory requirements related to different types of data.

Personally Identifiable Information

PII is information that personally identifies an individual. The National Institute of Standards and Technology (NIST) Special Publication (SP) 800-122, *Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)*, defines PII as

any information about an individual maintained by an agency, including:

(1) any information that can be used to distinguish or trace an individual's identity, such as name, Social Security number, date and place of birth, mother's maiden name, or biometric records; and

(2) any other information that is linked or linkable to an individual, such as medical, educational, financial, and employment information.

Although this definition is focused on a federal agency, it applies equally well to any organization that holds information on employees or customers that can identify an individual personally. This also extends to information such as credit card data, any type of account access data, and driver's license numbers.

Several laws mandate the protection of PII, and if an attack compromises PII, organizations have a responsibility to inform the affected individuals. Some businesses consider these laws unnecessary and a barrier to business growth because they add expenses. On the other side of the fence, some people suggest that without laws forcing companies to disclose this information, they may not readily do so.



EXAM TIP Several laws mandate the protection of PII. Companies face serious fines if they do not protect PII. Additionally, if a data breach occurs, companies are obligated to let customers know of the breach.

As an example, Equifax suffered a significant data breach in 2017 exposing private data for as many as 143 million people. In March 2018, Equifax reported that as many as 2.4 million additional Americans had partial driver's license information exposed. This partial driver's license information wasn't as sensitive as the private data exposed in the 2017 data breach. The following list provides a rough timeline of the incident:

- **May 14, 2017** Attackers breached Equifax servers.
- **July 29, 2017** Equifax detected the security breach.
- **July 30, 2017** Administrators patched the vulnerability.
- **September 7, 2017** Equifax officially announced the security breach.
- **September 11, 2017** U.S. senators asked Equifax to provide information on the attack.
- **September 13, 2017** The Equifax CEO testifies before Congress.
- **September 14, 2017** The Federal Trade Commission (FTC) states it is investigating the incident.

- **September 26, 2017** The Equifax CEO retires.
- **September 27, 2017** San Francisco is the first city to sue Equifax for failing to notify consumers in a timely manner.

The company ultimately faced over 70 class-action lawsuits for not protecting PII and not reporting the data breach in a timely manner. They also were also investigated by the Consumer Financial Protection Bureau, the House Financial Services Committee, the Senate Finance Committee, the New York Attorney General, and the New York Department of Financial Services.

This example raises a question that people have often debated after a data breach: Did Equifax reveal the information due to legislative requirements or out of genuine concern for its customers?

HIPAA

The Health Insurance Portability and Accountability Act (HIPAA) of 1996 mandates protection of health-related data within the United States. Prior to HIPAA becoming law, it was possible for just about anyone to call your doctor to ask for information on your health, and the doctor's office may have given it out freely. Today, any organization that gives out health-related data about a patient without authorization is subject to serious fines.

HIPAA identifies protected health information (PHI) as any information concerning the health status, provision of health care, or payment of health care for an individual. Although this primarily affects medical organizations such as hospitals, laboratories, and doctor's offices, it also extends to other, nonmedical organizations. For example, any organization that offers a health care plan and collects information about employees related to the plan also falls under the HIPAA rules.



EXAM TIP HIPAA covers any U.S. organization that processes health information. Failure to follow the requirements of HIPAA can result in significant fines and penalties.

SOX

The Sarbanes-Oxley (SOX) Act of 2002 mandates specific protections for data related to publicly held companies (companies registered with the U.S. Securities and Exchange Commission). SOX requires high-level officers (such as CEOs and CFOs) to verify personally the accuracy of financial data. Part of the goal is to prevent some of the fraud that was apparent in companies such as Enron and Tyco International.

For example, the Enron scandal resulted in a loss of \$11 billion to shareholders after the company's stock price plummeted from over \$90 per share to less than \$1 a share when it went bankrupt. Many experts attributed the failure to corrupt accounting practices. Later, several Enron executives were indicted and sentenced to prison.

Because the officers must personally verify the data, organizations are required to take extra precautions to ensure the accuracy and integrity of the data. Both internal and external audits are also required to ensure that these publicly held companies are in compliance.

GDPR

The General Data Protection Regulation (GDPR) is a European Union (EU) regulation. It supersedes the Data Protection Directive (also known as Directive 95/46/EC). Both mandate the protection of personal data of EU residents. The GDPR was adopted in April 2016 and became enforceable in May 2018. The transition period was intended to allow organizations to adopt relevant practices to comply with the law.

A primary purpose of the GDPR is to protect personal data of EU residents and it strictly restricts the transfer of data of EU residents outside of the EU. Organizations that collect or process personal data on EU citizens must comply with all elements of the GDPR, and this includes organizations that operate outside of the EU. As an example, Facebook and Google are U.S.-based companies that routinely collect data on users. Both companies must comply with the GDPR when collecting and exporting data on any EU citizens. If they violate the privacy rules in the GDPR, they can be fined up to 4 percent of their global revenue.

With 11 chapters and 99 articles, the GDPR is rather complex and filled with typical legislative legalese. Many organizations seek outside specialists to help them ensure they are complying with all elements of the regulation.

Training

Higher authorities sometimes mandate specific awareness and training requirements related to data. For example, Federal Information Processing Standards Publication 200 (FIPS Pub 200), *Minimum Security Requirements for Federal Information and Information Systems*, published by NIST, specifies that organizations must meet at least the following minimum requirements for awareness and training:

(i) ensure that managers and users of organizational information systems are made aware of the security risks associated with their activities and of the applicable laws, Executive Orders, directives, policies, standards, instructions, regulations, or procedures related to the security of organizational information systems; and

(ii) ensure that organizational personnel are adequately trained to carry out their assigned information security-related duties and responsibilities.

This compliance starts with the organization defining data classifications. The U.S. government has done this with data classifications such as Top Secret, Secret, Confidential, and Unclassified. Next, the organization ensures that users are aware of the different types of data classifications. Training helps users understand the classifications, how the organization protects the data, and the users' responsibilities and requirements for protecting the data. Audits (internal and external) often use documentation of training to help verify that an organization is complying with basic requirements.

Managing Assets Through the Lifecycle

Asset management refers to the steps used to manage important assets within an organization throughout their lifecycle. It helps the organization maintain control of its assets. In this context, assets include hardware, software, and data. The *lifecycle* refers to the time period from when the asset is first obtained or created to when it is discarded or destroyed.

On a larger scale, asset management is often associated with configuration and change management. In other words, an organization identifies all of its systems, but it also has documentation that identifies the current configuration of these systems. When a system is changed, the change management process ensures that the configuration documentation is kept up to date.



TIP Asset management systems help an organization track its assets. It's common for asset management systems to track hardware and software owned by an organization.

Hardware Inventory

Hardware management entails identifying relevant hardware that arrives at a company, periodically inventorying the hardware, and removing the hardware from the inventory when it is disposed of. This includes items such as desktop computers, laptops, servers, routers, and switches.

Not all hardware is added to an inventory. For example, a mouse is very inexpensive and the labor required to inventory all the mice in an organization is much more expensive than the cost of replacing a few stolen or lost mice. Internal policies drive the decision to add hardware to the inventory. For example, a company policy may state that all hardware valued at more than \$100 should be added to the inventory and anything less than \$100 should not be added.

There are many different ways to manage a hardware inventory. A basic method is to record the model and serial number of each piece of hardware, place that information into a spreadsheet, and periodically inventory the hardware. Small organizations may use this method. Another inventory method labels systems with bar codes when the hardware arrives. The organization then periodically scans all the systems to verify that they haven't disappeared.

Some organizations use radio-frequency identification (RFID) tags. The RFID tags transmit information on each inventory item, and RFID receivers can read data from the RFID tags. This is much easier than using bar codes. An individual has to scan each system manually with a bar code system. However, with RFID tags, an individual can just walk around the room with an RFID receiver to complete an inventory.

When an organization wants to combine the system with configuration management, it often has enterprise tools that it periodically uses to inventory the systems electronically. For example, an enterprise application will periodically check systems over the network to ensure that they all have relevant patches applied. This same application can be used as part of the hardware inventory.

Before disposing of hardware, it's important to ensure that it doesn't include any sensitive data. It's common to remove and destroy hard drives, or use methods to purge data before disposing of systems. The "Removing Data Remnants" section earlier in this chapter provides specifics of how to do so.

Software Inventory and Licenses

Some asset management systems track software. This helps an organization track operating systems and software applications throughout their lifecycle. It also helps administrators ensure that they are not violating licensing terms. For example, if an organization purchased 100 copies of an application, it needs to be certain that the number of installed copies of the application is 100 or less. Installing 200 copies of the application would be the same as software piracy.

Just as enterprise applications can check systems for hardware inventory and patch management, many can also inventory the systems and identify software installed on each system. For example, Microsoft systems support Windows Management Instrumentation (WMI). By sending WMI queries to devices over the network, administrators can retrieve a significant amount of information on the systems. This includes the hardware, operating system version, updates, and installed applications. Enterprise applications also use WMI queries to retrieve this information.

Data Storage

Adding data to an asset management system is more difficult because it's so easy to make a copy of files. However, it is possible to add media to asset management systems. This includes external drives and backup tapes.

It's also important to follow procedures mentioned earlier in this chapter related to handling data. This includes ensuring the data is properly classified, marked, and labeled, and ensuring the data is properly destroyed when it is no longer needed.

Certification and Accreditation

Certification and accreditation are two separate processes used to test, evaluate, and approve systems for specific purposes. The certification process includes several steps to evaluate,

describe, and test a system. This includes all the security controls that are used to mitigate risks to the system. Once a system is certified, an accrediting authority provides a formal declaration, which approves the system operation.



TIP Systems are often operating during the certification process. However, there may be specific restrictions on what the system can be used for until it is formally accredited.

The certification and accreditation process provides a level of checks and balances for systems. A single entity doesn't certify and accredit systems. Instead, one entity certifies them, and then another entity has the responsibility of verifying that the information is accurate and, if so, accrediting the system. A management authority normally completes the accreditation. When the authority accredits the system, it takes responsibility for the risks associated with the system and any losses associated with the system. Because the accrediting authority takes responsibility for the approval, this discourages simple "rubber stamp" approval. Instead, the accrediting authority verifies the authenticity of information provided by the certification process. Of course, because an organization expects the accrediting authority to examine the data closely, personnel take extra steps to ensure the data's accuracy before submitting the certification paperwork.



EXAM TIP Within the U.S. government, a Designated Approving Authority (DAA) provides official accreditation by approving a system for operation at a specific level of risk. It does not guarantee that a system is free of risk.

The U.S. government uses several publications from the NIST SP 800 series as standards for the certification process. For example, NIST SP 800-47, *Security Guide for Interconnecting Information Technology Systems*, provides information on the certification and accreditation of interconnected systems.

Certification, Accreditation, and Security Assessments

In addition to NIST SP 800 documents, FIPS Pub 200 mandates specific requirements related to certification, accreditation, and security assessments as follows:

(i) periodically assess the security controls in organizational information systems to determine if the controls are effective in their application;

(ii) develop and implement plans of action designed to correct deficiencies and reduce or eliminate vulnerabilities in organizational information systems;

(iii) authorize the operation of organizational information systems and any associated information system connections; and

(iv) monitor information system security controls on an ongoing basis to ensure the continued effectiveness of the controls.



TIP [Chapter 10](#) mentioned FIPS Pub 200 in the context of accounting and configuration management. Recall that it specifies minimum requirements that federal agencies must meet in several different IT security areas. Nongovernment organizations often adopt these same standards as best practices for IT security.

Common Criteria

The Common Criteria for Information Technology Security Evaluation (commonly called Common Criteria or simply CC) is a framework used to evaluate systems. It provides assurances that the specification, implementation, and evaluation of a system's security has gone through a rigorous and standardized process. The Common Criteria was created with the help of several government organizations, including those of the United States, Canada, United Kingdom, France, and Germany. It's recognized as an international standard and has superseded several other standards.

Evaluation Assurance Levels (EALs) provide a level of assurance for a product or a system. EAL has seven levels that indicate the level of quality assurance steps that personnel have taken to ensure the reliability of security features:

- **EAL1** Functionally tested. The threats to security are low, so a high level of security assurance is not required.
- **EAL2** Structurally tested.
- **EAL3** Methodically tested and checked.
- **EAL4** Methodically designed, tested, and reviewed. Many commercial operating systems achieve EAL4.
- **EAL5** Semiformally designed and tested.
- **EAL6** Semiformally verified design and tested.
- **EAL7** Formally verified design and tested. This provides the highest level of assurance. It's applicable in extremely high-risk situations, where the assets have exceptionally high value or when the impact of a threat can result in catastrophic losses.



EXAM TIP The Common Criteria uses EALs to determine the level of assurance of a system or product. EAL1 is the lowest level of assurance and EAL7 is the highest level. Many commercial operating systems achieve EAL4.

There's a trade-off in both cost and time to achieve higher levels of assurance. For example, the U.S. Government Accountability Office (GAO) indicates that it can take between 10 and 25 months to develop a system at EAL4, and the cost can range between \$150,000 and \$350,000.

TCSEC Superseded

You may run across documentation that talks about the Trusted Computer System Evaluation Criteria (TCSEC) and all the different components of it. U.S. government entities previously used TCSEC and it was widely known as the "orange book" because its cover was orange. The orange book included specific standards used to evaluate and classify systems in use to process and store classified information, or considered for use to process and store classified information.

However, a Department of Defense directive in 2002 canceled the use of the orange book. It has been superseded by the Common Criteria since 2005. Studying TCSEC does provide interesting information on how standards have been developed over the years, but it isn't relevant to certification and accreditation requirements currently used within the U.S. government.

Using a Risk Management Framework

NIST SP 800-37, *Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach*, provides guidelines on transforming the traditional certification and accreditation process into a six-step risk management framework. [Figure 11-5](#) shows the steps in the framework, and the following text provides an overview of each of these steps:

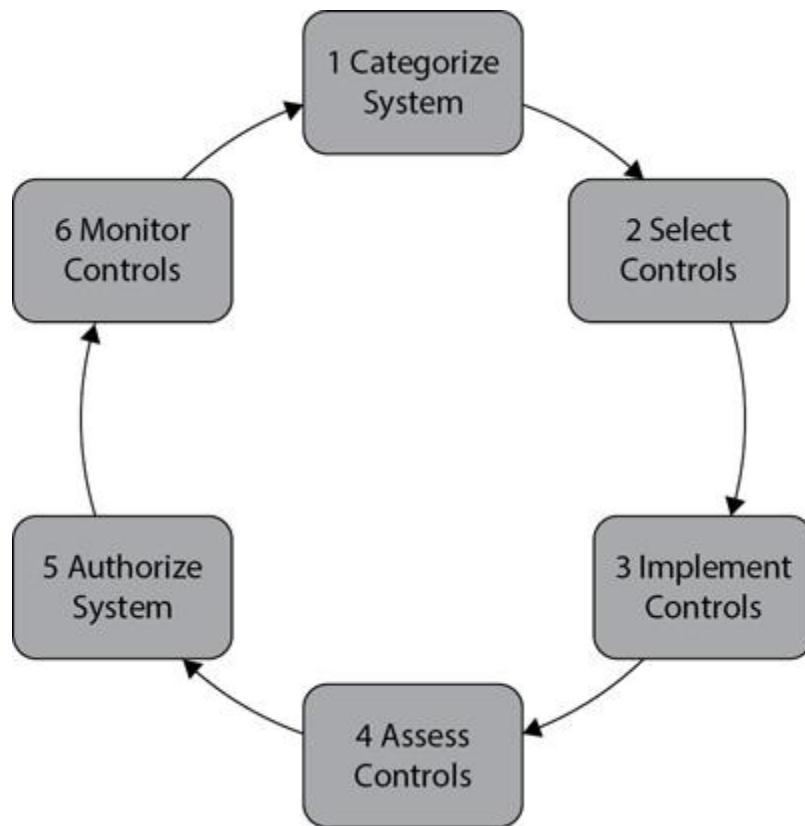


Figure 11-5 Risk management framework

1. Categorize system. Personnel examine the system, including the data it processes, stores, and transmits. They combine this information with an impact analysis to determine the security category of a system.

2. Select controls. Personnel select a set of baseline security controls based on the category of the system. For example, a system processing confidential data would start with a baseline of controls to protect confidential data. Personnel then tailor or supplement these controls based on the system and the needs of the organization.

3. Implement controls. In this step, personnel implement the selected security controls. They also document information on the controls to show how they are used within the system and what risks they mitigate.

4. Assess controls. Personnel periodically examine the security controls to verify they are implemented correctly, operating as intended, and working as expected.

5. Authorize system. If the level of risk is determined to be acceptable based on the assessment of the implemented controls, the system is authorized. Note that the system is often already operating before this formal authorization.

6. Monitor controls. Personnel monitor security controls on a continuing basis to determine their effectiveness against new threats and system changes. For example, any system change documented through a change control process may require a repeat of the six-step process.

The risk management framework recognizes that there are specific tasks required during different phases of the certification and accreditation process, but it also allows an ongoing reevaluation of the system on a regular basis. Risk doesn't stand still. Threats continue to morph, and mitigation controls implemented yesterday don't necessarily thwart the threats of today. Many security professionals believe that the risk management framework is better able to address these changes than the two-step certification and accreditation process.

Understanding Security Within the System Development Lifecycle

The system development lifecycle (SDLC) is a model used to track systems (and sometimes software projects) from cradle to grave. In other words, the model identifies each of the different phases from the initiation of the project through to the ultimate disposition of the project at the end of its lifetime.



TIP Different sources sometimes show different phases for the SDLC. This discussion centers on the phases documented in NIST SP 800-64, *Security Considerations in the System Development Life Cycle*, of initiation, development/acquisition, implementation/assessment, operations/maintenance, and disposal.

[Figure 11-6](#) shows the SDLC phases, and the following list provides a short overview:

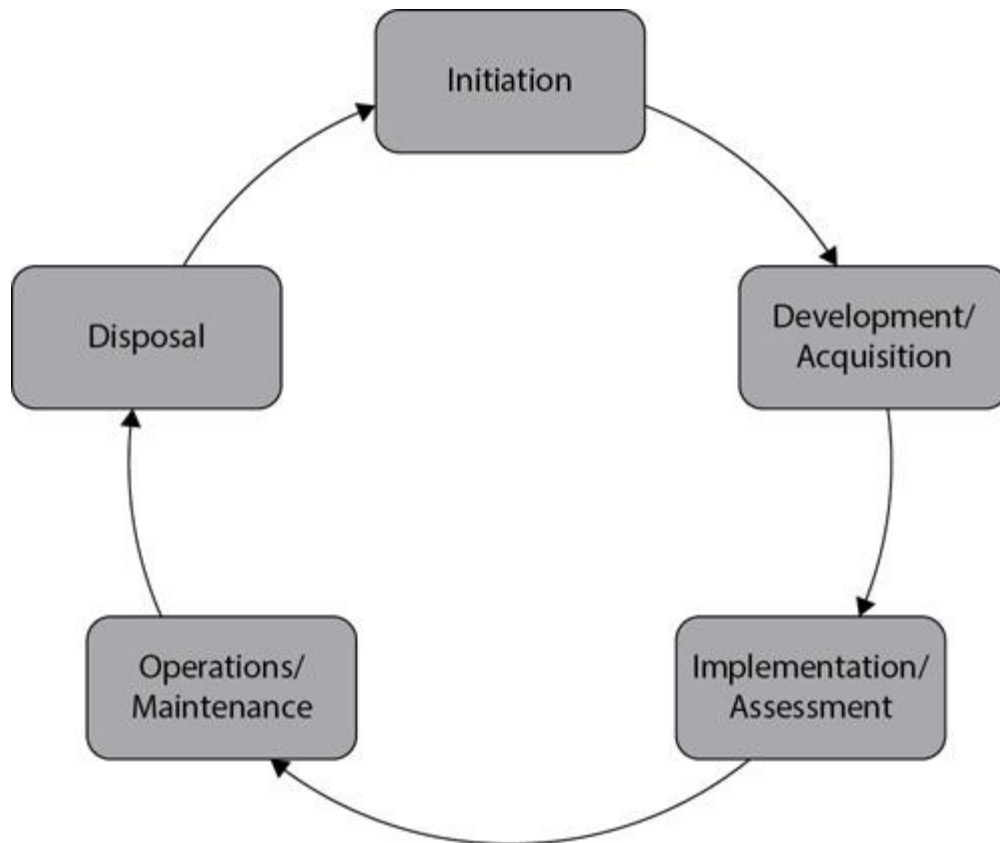


Figure 11-6 SDLC phases

- **Initiation** Personnel identify the need for the system in this phase. This includes documenting the purpose of the system and high-level requirements.
- **Development/Acquisition** Personnel design, purchase, program, develop, or create the system. In many models, this phase includes other processes such as an acquisition cycle or a software development cycle.
- **Implementation/Assessment** Personnel install, configure, and test the system in this phase. Testing and evaluation steps determine whether the system meets the originally identified need.
- **Operations/Maintenance** During this phase, the system performs its primary mission. Personnel perform regular maintenance on the system to ensure it continues to operate as desired.
- **Disposal** In this phase, personnel remove the system from service. A key component of this phase is sanitizing all media. If it is being replaced by another system, it should be removed only after the other system has completed the implementation/assessment phase. Personnel should also update inventory records to reflect the final disposition.

While it's informative to understand the tasks required at different phases, much of this is outside the scope of the SSCP exam. From the perspective of the SSCP exam, the most important

element is to understand the security issues at each of these stages. The following sections focus on the security elements.

Initiation Phase

Personnel focus on security concerns throughout the SDLC, and this includes the very beginning, the initiation phase. If security issues are not addressed early, then it's very possible that security fixes you try to add later will not be adequate or will cost too much to implement. It's important to look at the data that will be processed, transmitted, or stored on the system and understand the value of the data to the organization. As a reminder, the system should be protected at least at the level of the classification of the data. By identifying the classification of the data that the system processes, the organization then knows the minimum level of protection required by the system.



EXAM TIP Security issues should be addressed starting in the initiation phase and continuing all the way through to the disposal phase. Ignoring security early in any project makes it much more difficult to address problems adequately later.

Development/Acquisition Phase

During the development/acquisition phase, an organization conducts a risk assessment to identify baseline controls to protect the system and its data. The risk assessment allows the organization to identify risks, develop security requirements to address these risks, and evaluate security controls to ensure they mitigate risks to an acceptable level.

Implementation/Assessment Phase

In the implementation/assessment phase, the organization installs the system and begins testing it. In addition to testing the system to ensure that it meets the needs of the organization, personnel test the security controls to ensure that they meet the required security specifications. If the security controls don't meet the requirements, it's appropriate to add more security controls. However, these new controls should go through the same testing process.

Operations/Maintenance Phase

The operations/maintenance phase is normally the longest phase of a system. One of the biggest risks during this phase is that changes to the system will reduce the effectiveness of security controls. Organizations use change management and configuration management practices to ensure that changes to a system do not reduce the system's security. Configuration management often includes an auditing component to check a system periodically. These checks not only ensure the system retains the required configuration, but also can detect unauthorized changes.

Disposal Phase

When the system's lifetime ends, personnel sanitize it prior to disposal. Also, it's common to migrate data from the system to ensure that the data will be available after the system has been disposed of.

Personnel use various methods to sanitize media, as described in the earlier "Removing Data Remnants" section. The method they use is dependent on the data held and processed by the systems. Higher classifications require more stringent sanitization methods, such as destroying the media. Personnel may use programs to remove files with lower classifications. The key here is that they need to sanitize all media to ensure that there are not any data remnants on the media prior to disposal. This includes all disk drives, tapes, and nonvolatile memory. Systems should also be checked to ensure that optical media (such as CDs and DVDs) have been removed from the systems before disposing of them.



EXAM TIP If data remnants are not removed from systems, anyone who has access to the system after it has been thrown away will be able to read the data. Drives with data that has a higher level of classification may have to be destroyed, while drives with data of lower levels of classification may be scrubbed with programs that write different patterns of ones and zeros onto the media repeatedly.

During the disposal phase, it's also valuable to review the useful lifetime of data. Just because you can save data indefinitely doesn't mean you should. If data doesn't have any long-term use, it's often a good idea to dispose of it at some point. Of course, if requirements such as laws mandate keeping data for a specific period, it should not be disposed of until this time has passed.

Chapter Review

One of the first steps used to manage and protect data is to classify it based on its value to the organization. The classification drives the steps required to protect the data. Higher-classified data deserves greater protection, while lower-classified data does not warrant the extra costs associated with extra protection.

Some data classifications in the U.S. government are Top Secret, Secret, Confidential, and Unclassified. The private sector doesn't have standard terms for classifying data. However, some common classifications used in the private sector are Confidential (or Proprietary), Private, Sensitive, and Public, with Confidential representing the highest level. Organizations commonly implement data management policies to help protect data throughout its lifetime. Policies identify requirements when storing, transmitting, archiving, and retaining data. Access controls and encryption methods protect data at rest and data in motion.

Data is commonly stored in databases, and more specifically, in related tables within a database. Rows within a table are known as tuples. Tables have a primary key column to identify the row uniquely, and tables are related to each other by using a foreign key in one table to point to a primary key in another table. Database normalization separates data into multiple tables to prevent the duplication of data. SQL commands are used to communicate with databases. The two primary methods of organizing databases are online transaction processing (OLTP) and online analytical processing (OLAP). Web e-commerce applications typically use OLTP databases because of their high transaction processing speed and store transactions in transaction logs. Managers use OLAP databases to retrieve actionable data.

Data inference is a risk with data that allows an individual to gather individual unclassified pieces of data to predict or guess an outcome. Similarly, inference may allow someone to view summary or aggregate data and identify details through deduction.

Personally identifiable information (PII) is protected through several laws, and organizations are often required to disclose whether control of any PII is lost or a database has been breached. The Health Insurance Portability and Accountability Act (HIPAA) mandates the protection of certain health information, and the Sarbanes-Oxley (SOX) Act mandates that executives of publicly held companies vouch for the integrity of a company's financial data. The General Data Protection Regulation (GDPR) is an EU regulation that mandates the protection of personal data of individuals within the EU.

Asset management ensures that an organization knows what hardware, software, and data it owns and helps protect those assets. It's common to tie asset management systems to configuration management and/or change management systems.

Certification and accreditation are processes used to test, evaluate, and approve systems for specific purposes. Certification includes steps to evaluate, describe, and test a system to identify and mitigate risks. After certification, an accrediting authority provides a formal declaration approving the system for operation. In the U.S. government, a Designated Approving Authority (DAA) provides the official accreditation of systems.

Some organizations are applying lifecycle approaches to certification and accreditation. For example, NIST SP 800-37 provides a risk management framework used to identify specific tasks required during the certification and accreditation process. NIST SP 800-64 uses the SDLC model to track systems throughout their lifetime. When using any model, it's important to address security in every stage of the lifecycle, including the very first stage.

Questions

1. Of the following choices, what type of data requires the least amount of protection?

A. Confidential

B. Public

C. Private

D. Sensitive

2. Which of the following provides the best confidentiality protection for data at rest?

A. Marking it

B. Labeling it

C. Backing it up

D. Encrypting it

3. Who is responsible for classifying data?

A. Management

B. User

C. Administrator

D. Owner

4. An organization wants to reduce risks associated with proprietary data transmitted over the network. What can it do in its data management policy to achieve this objective?

A. Restrict how long data is retained

B. Specify how data is deleted from storage media

C. Require the encryption of data in motion

D. Require the encryption of data at rest

5. Which of the following methods will reliably remove all data from a backup tape?

A. Erasing

B. Degaussing

C. Diddling

D. Sanitizing

6. Which of the following is a secure method of sanitizing optical media?

- A. Degaussing
- B. Overwriting
- C. Shining
- D. Destroying

7. A company wants to reduce the amount of space used to store files used and shared by employees. What can it use to reduce the amount of storage space used?

- A. Data loss prevention (DLP) systems
- B. Deduplication
- C. Information rights management (IRM)
- D. Retention policies

8. Users within an organization have recently sent sensitive data outside the organization in e-mail attachments. Management believes this was an accident, but they want to prevent a recurrence. Which of the following is the best method to do so?

- A. Implement a network-based intrusion prevention system (IPS)
- B. Provide training to users
- C. Ensure the data is marked appropriately
- D. Implement a network-based data loss prevention (DLP) system

9. Of the following choices, what best describes a tuple?

- A. A column in a database
- B. A row in a database
- C. A primary key
- D. A foreign key

10. Which of the following is a virtual table and allows a user access to a limited amount of data within a table?

- A. View

- B. Tuple**
- C. Row**
- D. Foreign key**

11. An attacker has collected several pieces of unclassified information to deduce a conclusion. What is this called?

- A. Data mining**
- B. Database normalization**
- C. OLAP**
- D. Data inference**

12. An employee makes unauthorized changes to data as he is entering it. What is this?

- A. Data diddling**
- B. Data entry**
- C. Data inference**
- D. Data deduplication**

13. A database includes information on customers. This information can be used to distinguish or trace a customer's identity. Of the following choices, what best describes this data?

- A. PII**
- B. Tuple**
- C. Data inference**
- D. PHI**

14. A database includes health-related information on employees of a U.S. organization. Management wants to ensure it protects this data and complies with relevant laws. Of the following choices, which one identifies the regulation they should follow?

- A. SOX**
- B. HIPAA**

C. Equifax

D. PII

15. The CEO of a publicly held company in the United States is required to verify the accuracy of a company's financial data. What requires this activity?

A. HIPAA

B. SOX

C. NIST SP 800-64

D. NIST SP 800-37

16. A U.S. company collects information on customers and users around the world and stores much of it in databases in the United States. Management wants to ensure the company is complying with relevant laws in the EU. Which of the following is the most relevant law?

A. GDPR

B. Directive 95/46/EC

C. HIPAA

D. SOX

17. Within the U.S. government, who can formally approve a system for operation at a specific level of risk?

A. Certification authority

B. NIST

C. Senator

D. Designated Approving Authority (DAA)

18. Which of the following EALs indicates a system was methodically designed, tested, and reviewed, and is the level of assurance assigned to many commercial operating systems?

A. EAL0

B. EAL1

C. EAL4

D. EAL7

19. Which of the following is an international standard that provides a framework to evaluate the security of IT systems?

- A.** ITSEC
- B.** TCSEC
- C.** Common Criteria
- D.** Orange book

20. Of the following choices, what is a primary task to accomplish in the disposal phase of a system's lifecycle?

- A.** Migrate all data to other systems
- B.** Delete all data
- C.** Remove data remnants from systems before disposal
- D.** Back up all data to tape

Answers

1. B. Public data requires the least amount of protection. An organization would want to ensure that public data on a website is not modified, but the organization puts the data on the website to make it available to the public. The other choices indicate sensitive data that is important to an organization and deserves varying levels of protection.

2. D. Encryption provides the best confidentiality protection for data at rest. While it is appropriate to mark or label it, this isn't as strong as encrypting it. Backing it up provides protection for availability.

3. D. The data owner is responsible for classifying data. Management is responsible for defining data classifications, such as in a data policy. Users access the data, but they are not responsible for assigning or modifying data classifications. Administrators grant access to users based on the individual user's need and often at the direction from the data owner.

4. C. Encrypting data in motion can protect it against loss of confidentiality when it is transmitted over a network. Retention policies restrict how long data is retained, but do not affect data in motion. Destruction policies dictate how to delete data or destroy media but aren't related to data in motion. Encrypting data at rest doesn't ensure that it is encrypted when it is transmitted. Some data-at-rest encryption methods will decrypt the data before transmitting it over a network.

5. B. Degaussing will reliably remove all data from a backup tape. Degaussing uses a powerful magnet to erase the data. Other methods of erasing data from a tape don't necessarily erase all the data. Data diddling is the unauthorized changing of data before or while entering it into a system and is unrelated to removing data from a tape. Sanitizing the tape is the goal, but it is not a method.

6. D. Optical media must be destroyed to ensure it doesn't include any remaining data. Degaussing isn't effective because optical media doesn't use magnetic methods of storing data. Overwriting isn't effective because the media might still have data remaining after overwriting it. Shining is not a valid method of sanitizing media.

7. B. Deduplication ensures that a file is stored only once on a system, even if multiple users have access to the same file. DLP systems attempt to monitor data usage and prevent the unauthorized use or transmission of sensitive data. IRM refers to the different methods used to protect sensitive information from unauthorized access. Retention policies restrict how long data is retained.

8. D. The best solution is to implement a network-based DLP system. It can scan outgoing data to look for sensitive data and block any it finds. An IPS focuses on incoming traffic to block attacks and doesn't necessarily scan outgoing traffic. Thus, it wouldn't necessarily stop data sent as an e-mail attachment. Training is appropriate, but training doesn't necessarily prevent accidents. Also, training wouldn't stop a malicious insider from this action. The data should be marked with its appropriate classification, but there isn't any indication that it isn't. Also, even if the data is marked, it doesn't necessarily prevent accidents.

9. B. A tuple is a row in a database. Columns are also known as attributes. A primary key uniquely identifies a row in a table and is related to a foreign key in another table to create a relationship between two tables.

10. A. A view is a virtual table that provides access to specific columns in one or more tables, allowing a user to access a limited amount of data. A tuple is the same as a row and it contains a single record, but it is not a virtual table. Tables are linked together with foreign keys.

11. D. Data inference occurs when someone is able to piece together unclassified information to predict or guess an outcome. Data mining is the process of retrieving relevant data from a large database. Database normalization is the process of dividing the data into multiple tables in a database to reduce duplication of data. Online analytical processing (OLAP) uses techniques to make it easier to retrieve data.

12. A. Data diddling is the unauthorized changing of data before entering it into a system or while entering it into a system. While the employee is involved with data entry, data entry by itself doesn't indicate the employee is making unauthorized changes. Data inference occurs when someone is able to put together unclassified pieces of information to predict or guess an outcome. Data deduplication stores a file only once on a system, even if users attempt to store copies in multiple locations.

13. A. Personally identifiable information (PII) is any information that can be used to distinguish or trace an individual's identity and includes items such as their name, Social Security number, birthdate, birthplace, and more. A tuple is a row in a database table. Data inference occurs when someone is able to piece together unclassified information to predict or guess an outcome. Protected health information (PHI) is information referring to an individual's health and is protected by HIPAA.

14. B. The Health Insurance Portability and Accountability Act (HIPAA) mandates the protection of health-related information. The Sarbanes-Oxley (SOX) Act mandates specific requirements to guarantee the accuracy of data held by companies registered with the U.S. Securities and Exchange Commission. Equifax is a U.S.-based company that suffered a data breach and was required to report the loss to customers based on laws that protect personally identifiable information (PII).

15. B. The Sarbanes-Oxley (SOX) Act requires high-level officers (such as the CEO and CFO) to verify personally the accuracy of a company's financial data. The Health Insurance Portability and Accountability Act (HIPAA) mandates the protection of protected health information. NIST publishes documents that provide guidance to federal agencies and can be used by private companies, but these documents do not require any actions by private companies.

16. A. The General Data Protection Regulation (GDPR) is a European Union (EU) regulation that mandates the protection of EU residents personal data. The GDPR supersedes Directive 95/46/EC (also known as the Data Protection Directive). The Health Insurance Portability and Accountability Act (HIPAA) mandates the protection of health-related information. The Sarbanes-Oxley (SOX) Act requires high-level officers (such as the CEO and CFO) to verify personally the accuracy of a company's financial data.

17. D. A DAA provides official accreditation by approving a system for operation at a specific level of risk. The certification authority does not approve a system, but instead evaluates, describes, and tests a system. The National Institute of Standards and Technology (NIST) provides recommendations of standard best practices, but it does not certify or accredit systems. Senators do not certify or accredit systems.

18. C. Most commercial operating systems achieve Evaluation Assurance Level 4 (EAL4) when evaluated by the Common Criteria (CC). EAL4 indicates the operating system has been methodically designed, tested, and reviewed. EAL0 is not a valid level of CC. EAL1, the lowest level of assurance, indicates the system has been functionally tested. EAL7 is the highest level of assurance and indicates the system has a formally verified design and has been tested, but EAL7 ratings are not as common as EAL4, especially for operating systems.

19. C. The Common Criteria provides a framework for evaluating the security of IT systems, and several countries (including the United States, Canada, United Kingdom, and France) have adopted it. The Information Technology Security Evaluation Criteria (ITSEC) is a standard used in Europe, but it has been largely replaced by the Common Criteria. The Trusted Computer System Evaluation Criteria (TCSEC or orange book) was used by the U.S. government, but it has been superseded by the Common Criteria.

20. C. Data remnants should be removed before disposal to ensure that data does not fall into the wrong hands. Depending on the value of the data, it could be migrated, deleted, or backed up, but this does not remove the need to ensure that data does not remain on the system before disposing of it.

NOTES

CHAPTER 12

Security Administration and Planning

In this chapter, you will learn about

- Security policy contents and characteristics
- Raising the awareness of security policies
- Business continuity plans
- Business impact analysis
- Disaster recovery plans
- Difference between a BCP and DRP
- Alternative locations such as hot sites, cold sites, and warm sites
- Security organizations such as NIST and US-CERT

Understanding Security Policies

A security policy is a written document that provides an organization with a high-level view of its security goals. [Chapter 9](#) presented and contrasted the differences among policies, standards, guidelines, and procedures. To quickly review, policies are high-level, authoritative documents. Standards document criteria such as a proven norm or method and may specify requirements for a process or a technology. An organization can choose to adopt some standards, while other standards apply to an organization based on partnerships or relationships with other entities. Guidelines provide recommendations for members of an organization, but they are not mandatory. Procedures provide the action steps to accomplish tasks.

A security policy is an administrative control that focuses on the management of risk and IT security. Senior management creates or at least endorses the organization's security policy. The policy reflects the culture of the organization and provides the overall direction for security. Others within the organization use the security policy as an authoritative document to implement controls to enforce and follow the requirements of the security policy.



NOTE A security policy can be an overall security policy for an organization. Organizations can also create security policies for individual systems that define the system and security

protections related to the system. Within this chapter, the focus is on an overall security policy for an organization.

A security policy typically goes through several stages:

- **Initial stage** Personnel draft the security policy based on the needs of the organization. This might be a formal, nearly complete draft of the policy or an initial proposal to senior management identifying the needs and objectives of the policy.
- **Approval stage** Senior management approves the policy. It might take several iterations between the initial stage and the approval stage to create a document that senior management approves. Once approved, it provides direction for personnel to enforce the policy.
- **Publication stage** The policy is provided to relevant personnel so that they can follow and implement it.
- **Implementation stage** A security policy is an important first step to provide security within an organization, but it isn't the final step by any means. After senior management has approved a security policy, personnel must take steps to implement and enforce it.
- **Maintenance stage** Periodic reviews (such as once a year) ensure that the policy remains up to date, meets the needs of the organization, and addresses current threats.



EXAM TIP A security policy is an administrative control that provides the high-level authority for an organization. They are authoritative in nature and provide direction to personnel that implement them using other security controls.

While it's possible to have a single security policy covering all the elements of security for an organization, this can easily become very large and difficult to maintain. Instead, organizations often create several security policies, with a single security policy driving the others. As an example, [Figure 12-1](#) shows the relationship between an organization's security policy and other supporting policies. For instance, the security policy can mandate the creation of an acceptable use policy (AUP) and require users to review and acknowledge the AUP periodically. Personnel then create an AUP to support this requirement. Similarly, the security policy may state that the organization needs to prepare for potential disasters and ensure that critical functions continue to operate during disasters. Personnel then create a business continuity plan, identify critical functions with a business impact analysis, and create disaster recovery plans to respond to different types of disasters.

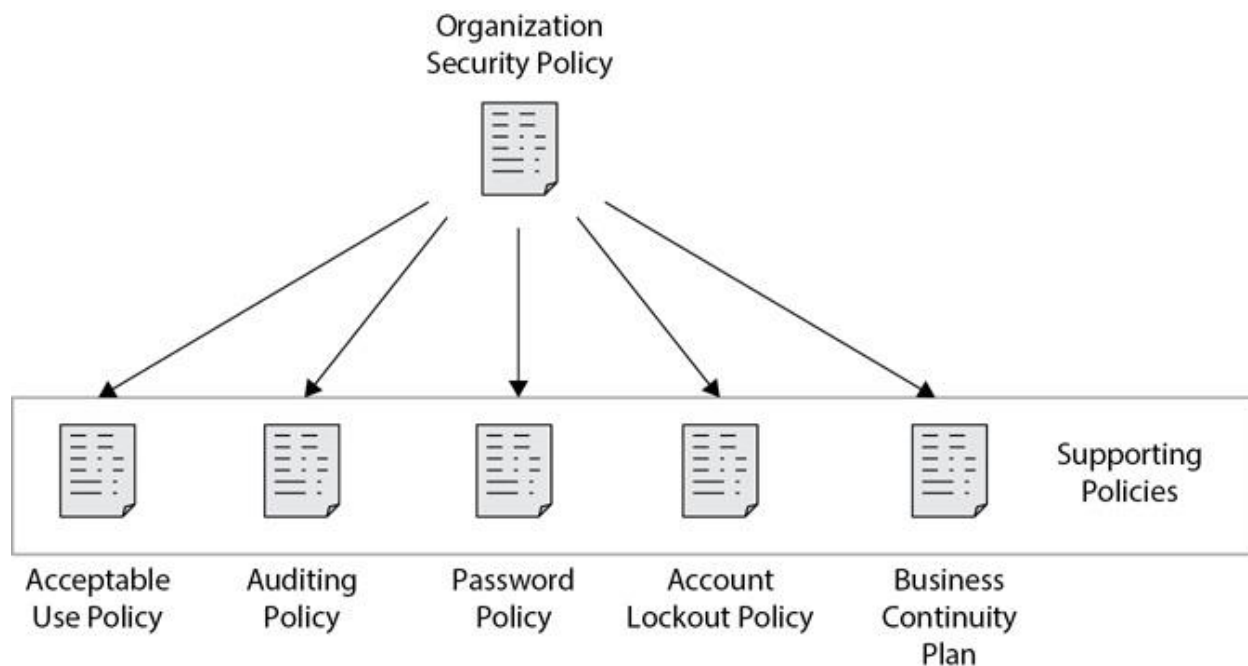


Figure 12-1 An organization security policy and supporting policies

Some organizations use a similar approach but include the supporting policies as appendixes to the security policy. The primary document is the organization’s security policy and it refers to the appendixes. Each appendix addresses separate security concerns, such as an acceptable use policy, an auditing policy, and so on.

Security Policy Characteristics

There isn’t a single standard for creating a security policy. Depending on the size and needs of the organization, a security policy can be just a few pages long or 50 pages or more. However, you’ll see some common characteristics in many effective security policies, including the following:



TIP Items in this list are not mandatory items to include within a security policy and the list is not all-inclusive. However, the list does provide a cross-section of topics that many organizations include in their security policies.

- **Organization mission statement** A basic mission statement helps members of the organization understand the overall vision of the organization and how security can be used to enhance the mission and vision. The goal is for security methods and procedures to support the mission of the organization and its overall business model.

- **Statement of accountability** A security policy defines the roles and responsibilities of users, management, and IT staff. This section provides a basis for ensuring compliance with the policies and procedures.
- **Data classifications** An organization often defines the classifications of data (such as Confidential, Sensitive, Private, and Public) in the security policy. This section explains how to determine data classifications and provides some guidance on protecting data based on the classification.



NOTE [Chapter 11](#) covers data classifications in greater depth. It includes a discussion of the common classifications used in the U.S. government and classifications used in private organizations.

- **Backup plans** Security policies often dictate the requirements for overall backups. This includes backup requirements, offsite storage of backup copies, and how long to retain backups.
- **Classification of resources** A security policy can identify specific resources that need higher levels of protection. As an example, consider an organization that hosts an e-commerce website with access to a customer database. The security policy can mandate the protection of the e-commerce website and require the database to be stored on a different server. Administrators choose the details of how to implement these requirements. They might choose to host the web server within a demilitarized zone (DMZ) protected with an intrusion detection system, and store the database on a separate server in the internal network.
- **Network access** This section defines the types of access authorized on the network. For example, the organization may restrict the use of any wireless devices on a network. If the organization allows wireless access, the policy may specify minimum encryption standards. Similarly, the policy might mandate the creation of a separate network with Internet access for wireless devices, but prohibit these devices from accessing internal network resources.
- **Risk management** In this section, the organization defines an overall view of risk management. This could include direction on risk assessments and vulnerability assessments, such as how often to complete these assessments. An organization may also provide direction on types of analysis that assessors should use. For example, the policy might direct the use of quantitative analysis or identify specific values to use in a qualitative analysis.
- **Auditing** An auditing policy defines what should be audited at a minimum and how often. For example, it can mandate extra auditing requirements for privileged accounts and require that nonadministrator personnel audit actions by these accounts at least every six months.

- **Business continuity** Most organizations include an overall statement regarding business continuity and disaster preparation. Separate documents provide the details, but the security policy provides the overall goal.
- **Incident response** An incident response section defines a security incident for an organization and provides information on what should be done in response to an incident. It could include overall direction on goals to prevent, detect, and respond to security incidents and require the creation of an incident response team.
- **Physical security** This section provides an overview of physical security requirements. It may mandate the creation of a zone document for the building. The zone document identifies the specific protections required when processing different types of classified data and specific zones within the building. For example, a server room may process confidential data and require protections such as a single video-monitored entrance/exit that can be accessed only with a smart card and PIN.
- **Acceptable use** Most organizations draft acceptable use policies to let users know what is acceptable use of computer equipment and networks. The security policy may mandate the creation of an acceptable use policy and identify how users acknowledge it, such as by signing an acknowledgment that they have read the policy. The policy will often specify when or how often users should acknowledge it, such as when they are initially hired, and annually thereafter.



EXAM TIP An acceptable use policy (sometimes called an acceptable usage policy or AUP) ensures that users understand their responsibilities when using an organization's IT resources. Organizations commonly require users to review and acknowledge these types of policies regularly, such as once a year.

- **Enforcement section** Many policies identify an overall vision of how a security policy is to be enforced. For example, the policy may state an overall goal of using electronic monitoring and technical controls to enforce policy elements whenever possible. Similarly, it might mandate periodic training to ensure personnel are aware of the security policies.
- **Passwords and other authentication requirements** A policy can generically state the overall authentication requirements. For example, it may state that a technical password policy must ensure that users create strong, complex passwords and that they change them regularly, or that access to certain systems requires the use of two-factor authentication. Many organizations require administrators to have separate administrator accounts and use longer and more complex passwords for the administrator accounts than they use for their regular user accounts. When services require separate service accounts, these can have different password requirements.
- **Account lockouts** The policy can require that accounts automatically lock after users enter the incorrect password too many times. Further, it can require technicians or administrators to unlock the accounts manually.

- **Hardware usage** The security policy can restrict or prohibit the use of some hardware. For example, USB flash drives have specific risks such as the ability to transfer malware from system to system, or allow an employee to steal massive quantities of data. Many organizations restrict the use of USB flash drives to reduce these risks. A security policy can mandate the use of technical methods to prevent the usage of USB drives altogether, or to prevent systems from reading data from a USB drive or writing data to a USB drive. The policy can also require technical controls to detect and report any user attempts to use a USB drive.

- **Ethics statement** An ethics statement identifies the minimal acceptable behavior by members of the organization. The security policy can include an ethics statement identifying the organization’s expectations related to ethics and directing an employee of what to do when faced with an ethical dilemma. In general, most organizations expect employees to resolve ethical dilemmas within the organization by going to their supervisor, or their supervisor’s supervisor, whenever possible.

Enforcing Security Policies

Of course, security policies are useful only when personnel know about them, implement them, follow them, and enforce them. This is one of the reasons that senior management must endorse and support security policies. If not, the policies become meaningless documents.

[Figure 12-2](#) shows some of the methods used to enforce security policies. You’ll find these controls covered throughout this book. For example, [Chapter 2](#) covers authentication and access controls; [Chapter 7](#) covers risk management topics; and [Chapter 9](#) covers several different types of administrative, technical, and physical security controls.

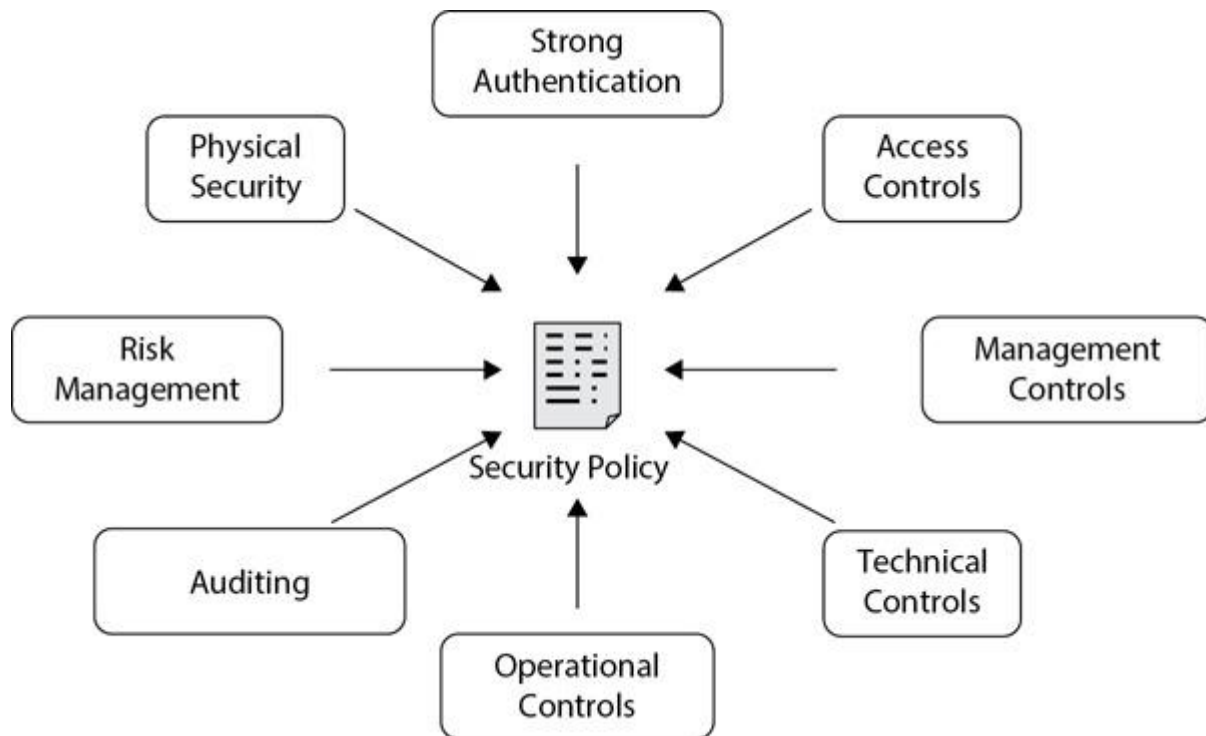


Figure 12-2 Enforcing a security policy



TIP The elements shown in [Figure 12-2](#) aren't meant to be a complete list. However, they do help illustrate that an organization uses multiple methods to support a security policy.

These elements work together to provide a defense in depth security solution by increasing the overall security posture of the organization. An important point worth stressing is that, ideally, the security policy comes first, and then personnel implement security controls in support of the security policy. The goal is for these controls to enforce the security policy and support the overall mission of the organization.

Value of a Security Policy

The obvious value of a well-crafted security policy is that it increases security and security awareness. If the security policy can prevent or at least reduce security incidents, it can protect an organization against direct financial losses, as well as indirect losses through damage to its reputation. However, there are also some less obvious benefits.

Direct Losses vs. Indirect Losses

Direct losses refer to the immediate losses resulting from an incident, such as the loss of revenue. Indirect losses refer to other factors such as the loss of customer goodwill. Accountants refer to direct costs as any cost that can be traced to a specific cost object.

As an example, the Equifax data breach of 2017 resulted in losses of \$164 million in 2017. They reportedly expect to lose an additional \$275 million in 2018, with total expected losses of \$439 million through the end of 2018. Many of these costs were related to technology and security upgrades, free identity theft services for consumers affected by the attack, and legal fees.

The company purchased an identity protection company for \$63 million after learning of the data breach, but before disclosing it to the public. They then used this to provide free identity protection services to consumers affected by the data breach. This is an example of a direct cost.

Shares of Equifax were trading close to \$143 before they announced the data breach. After reporting the data breach, shares fell to \$116. This reflects an indirect loss.

Even if a security incident occurs, the security policy may limit the organization's liability. As an example, consider a similar attack against two different fictitious organizations, ACME and EMCA. The attack exploits a zero day vulnerability and results in the loss of customer data. ACME doesn't have a security policy, and personnel have not documented their security practices. EMCA does have a security policy, and has a significant amount of documentation showing that it regularly follows security practices. ACME will be more vulnerable to litigation

claiming that the company has not taken steps to protect customer data. In contrast, EMCA may be able to demonstrate that it did everything it could to prevent losses, so its liability would typically be limited.

Similarly, when employees sign security policies (such as an acceptable use policy) to acknowledge that they have read the policies and are aware of their contents, they can be held accountable to follow the security policies. Most employees won't knowingly violate security policies. However, some employees may purposely try to circumvent them, which can cause additional harm to an organization. For example, if an employee surfs pornographic sites using a company computer (counter to the AUP), fellow employees may see the objectionable material and raise sexual harassment issues. In the absence of a policy, the company may be liable for the sexual harassment. With a policy in place, the company can hold the offending employee accountable.

Security Policies Becoming More Common

Organizations are increasingly required to have security policies, typically because a higher authority mandates them. For example, the Payment Card Industry Data Security Standard (PCI DSS) requires organizations processing credit card data to meet multiple security control objectives and requirements, as detailed in [Chapter 10](#). One of the requirements is to maintain a policy that addresses information security.

Some organizations aren't required by law or regulation to have a security policy and they consider it too expensive or too time-consuming to draft, implement, and enforce. This view often changes after one or two costly security incidents. The lack of a security policy can often cost much more in the long run.

Without a policy providing guidance, employees must use their own judgment to decide what they think is important to the organization. At the very least, these decisions will be inconsistent and can easily result in the loss of confidentiality, availability, and integrity of IT systems and data. Worse, the lack of direction often results in wasted efforts, rework, and higher costs.

Data is often one of the most valuable assets within an organization, but not all data has the same value. Without guidance, employees could easily overprotect public data and under-protect highly confidential data. For example, if IT personnel don't recognize the value of research and development data, they may not use strong access controls to protect it or may not back it up often enough. An attacker may be able to hack into the system and steal the data, or a common hardware failure may result in the loss of a significant amount of research data. Both scenarios can result in the organization's losing a significant amount of money.

It's also worth noting that if an organization doesn't have a security policy, it may be susceptible to other legal problems, including liability issues. For example, if an organization doesn't have an acceptable use policy and an employee uses a computer to hack into other systems or networks, the company can be held liable. However, an organization reduces its legal liability by having documentation (such as a signed AUP) proving it has informed employees of unacceptable practices.

Complying with Codes of Ethics

[Chapter 1](#) covered the Code of Ethics required by (ISC)². As you'll recall, you must subscribe to this Code of Ethics to earn and keep the SSCP certification. The Code of Ethics includes responsibilities such as protecting society, acting honorably, providing diligent and competent service, and advancing and protecting the profession.

However, the (ISC)² doesn't have a corner on ethics. Many organizations have their own ethics statement that provides guidance to personnel related to their conduct within the organization. When employees sense a conflict based on the organization's stated ethics, they have a responsibility to raise this internally within the organization rather than reporting it externally.



EXAM TIP Employees have an obligation to abide by the ethics of their employer. In other words, employees have an implied loyalty to their employer as long as they remain employed.

It's unlikely that any security policy will cover every possible issue that can present itself within IT security. This is where ethics come in. They provide a framework of norms and principles of correct conduct beyond the security policies. With this in mind, employees of an organization have a responsibility to ensure that secure practices are still followed even if there isn't a specific policy that covers a particular challenge.

One point worth stressing is that employees should not retaliate against attackers, even if an organization doesn't have a policy addressing counterattacks. The exception is when it is part of the cybersecurity professional's job. For example, if an attacker launches a SYN flood attack on one of your servers, you may be able to analyze the logs and determine the source of the attack based on the IP address. You may be tempted to launch an attack against the attacker, but this is not ethical. It can also result in unexpected results.

First, the attacker may be spoofing an IP address. Even though you think you know where the attack originated, you may be attacking an innocent company server that the attacker hijacked and used to launch the attack. Second, when you launch an attack, you are escalating the situation. If you block the attack, the attack stops and the incident is over. If you attack back, the attacker may take it personally and launch multiple malicious attacks against your organization for days, months, or even years afterward.

Policy Awareness

After an organization creates and approves a security policy, it's important to ensure that the appropriate personnel know its contents. Having a policy hidden away in a desk is almost the same as having no policy at all.

Different members of an organization require different levels of understanding of a security policy. In other words, regular employees need to understand some topics within the policy, but not necessarily all the topics that security and IT personnel need to understand. As an example, all employees must be aware of an acceptable use policy.

In contrast, imagine the policy states that two-factor authentication is required when accessing the network via a virtual private network (VPN). IT and security personnel would design and implement an appropriate two-factor authentication system (such as with smart cards and usernames and passwords). Personnel using the VPN would need smart cards to authenticate via the VPN. However, users who never use VPNs don't need to know about the policy requirement. Similarly, users who do use VPNs don't need to understand the technical details of multifactor authentication, as long they know they must use a smart card with their username and password.

The following sections cover some common methods that help ensure personnel know the relevant details within the security policy.

Easy-to-Read Language

The policy should be written in language that is simple to understand. If the policy is too complex for the target audience, they won't understand it and very likely will not follow it.

As an example, if the policy states that passwords must have at least 30 bits of entropy, most people won't understand what that means. However, personnel will understand the policy if it states that passwords must have at least eight characters, include at least one character from all four character types (uppercase letters, lowercase letters, numbers, and special characters), and must not include the user's name.

Warning Banners

One of the ways that you can let users know of relevant security policy contents is with warning banners. Administrators can create warning banners that appear each time a user logs on. Warning banners can define who is considered an authorized user and warn unauthorized users against accessing the system.

Some organizations go further and provide information to users on their expectation of privacy within the warning banner. For example, many organizations regularly monitor user activity. A warning banner can remind users that tools are monitoring their activity to detect improper or illicit use of computer systems and networks. The warning banner could also list penalties for noncompliance with the organization's acceptable use policy. These warning banners remind users of their responsibilities and help serve as a deterrent.

Training Sessions

Some elements of a security policy may require training to implement. Organizations typically target training to different groups based on the needs of individuals within the group. For

example, an incident response team may require forensic analysis training to ensure they are able to respond to incidents, identify what happened, and limit the potential damage.

Regular users might only need annual training on the acceptable use policy and periodic training on emerging threats. For example, attackers continue to create different types of attacks and change techniques in older attacks. Periodically educating employees about current attacks such as phishing attacks can help employees recognize malicious e-mails and avoid malware infections.

Many organizations provide this type of training by using online methods, such as via an internal website. However, even a simple e-mail can inform users of a current threat. For example, one phishing attack sent a ZIP file to users as an e-mail attachment, with the e-mail message indicating that FedEx was unable to deliver a product to them and instructing them to open the file for more information. However, the ZIP file included malware designed to infect their system when users unzipped the file. Sending an e-mail to users within the organization to inform them of the phishing attack not only helps them to avoid opening the file, but also helps them to recognize how phishing attacks morph over time.

Security Flyers and Posters

Some organizations use security flyers and posters to raise security awareness and remind people of the security policies. The goal is often to get the users to recognize that security is their responsibility too, not just something the company does.



EXAM TIP Some common methods used to increase user awareness of security policy elements that are relevant to them include training, warning banners, and posters. At the very least, relevant portions of the policy must be understandable and accessible to employees.

Updating Security Policies

Security policies are dynamic documents. An organization that creates a security policy but never updates it can't expect the policy to remain effective. Organizations change as time moves forward. Similarly, threats and vulnerabilities change. By periodically reviewing and updating a security policy, the organization keeps up with these changes.

It's common to schedule periodic reviews of a security policy as a maintenance stage. For example, many organizations schedule annual reviews. During this review, the organization can examine changes to determine whether the security policy warrants an update. Changes in the industry, business practices, laws, threats, and many other aspects may require changes in the security policy.



EXAM TIP Security policies should be reviewed regularly, such as once a year. These reviews will identify where a policy should be updated. If an organization's security policy is updated, supporting policies should also be reviewed to ensure that they still provide the necessary support.

Additionally, an organization will often reexamine a security policy after a security incident. For example, if a security breach results in the loss of valuable data to an organization, personnel should examine the security policy and supporting practices. They may find that the security policy is inadequate and needs to be updated. Or, they may find that the security policy itself is adequate but personnel didn't follow the policy. In that case, personnel might need additional training, or the policy might need to be more accessible to personnel who are responsible for implementing the policy.

Understanding BCP and DRP Activities

Business continuity planning includes all the processes and procedures to prevent the loss of mission-critical services for an unacceptable length of time. A business continuity plan (BCP) includes multiple elements, such as a business impact analysis (BIA) and a disaster recovery plan (DRP). A BIA identifies critical business functions and processes, and a DRP helps an organization restore critical systems to full functionality as quickly and efficiently as possible after a disaster.

Several events can result in service disruptions, including any of the following:

- Attacks
- Disasters
- Accidents
- Human error
- Equipment failure
- Software or application failure

It's worth stressing that business continuity planning isn't a small endeavor and it requires the support of senior management. This ensures that business continuity efforts are adequately funded and that all elements of the organization support the planning efforts. When organizations take business continuity planning seriously, it entails several distinct action steps. [Figure 12-3](#) shows the overall steps to develop a BCP.

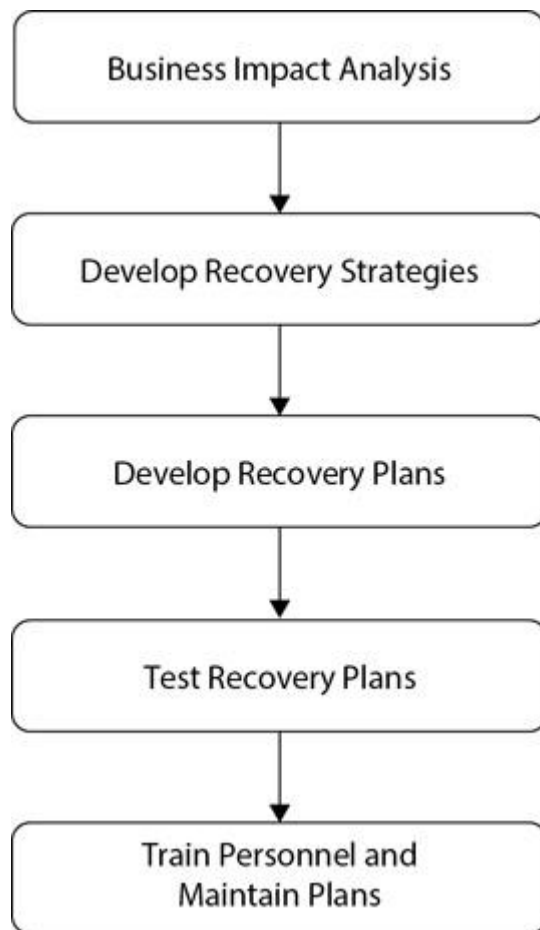


Figure 12-3 Overall steps for a BCP

The business impact analysis (discussed in more depth in the next section) identifies critical functions within the organization and provides an important starting point. The organization uses the results of the BIA to develop recovery strategies, such as the use of alternative sites and different backup strategies. Once personnel identify recovery strategies, they develop and test plans to support the strategies. These plans include detailed steps used in response to disasters or emergencies. Last, personnel responsible for implementing the plans get appropriate training and begin implementing the plans. Personnel periodically review the plans and update them as necessary.

[Figure 12-4](#) illustrates the relationship between a BCP and supporting policies. For example, a BCP starts with a BIA to identify critical functions. The BCP also includes backup policies outlining backup requirements and one or more DRPs that outline recovery strategies.

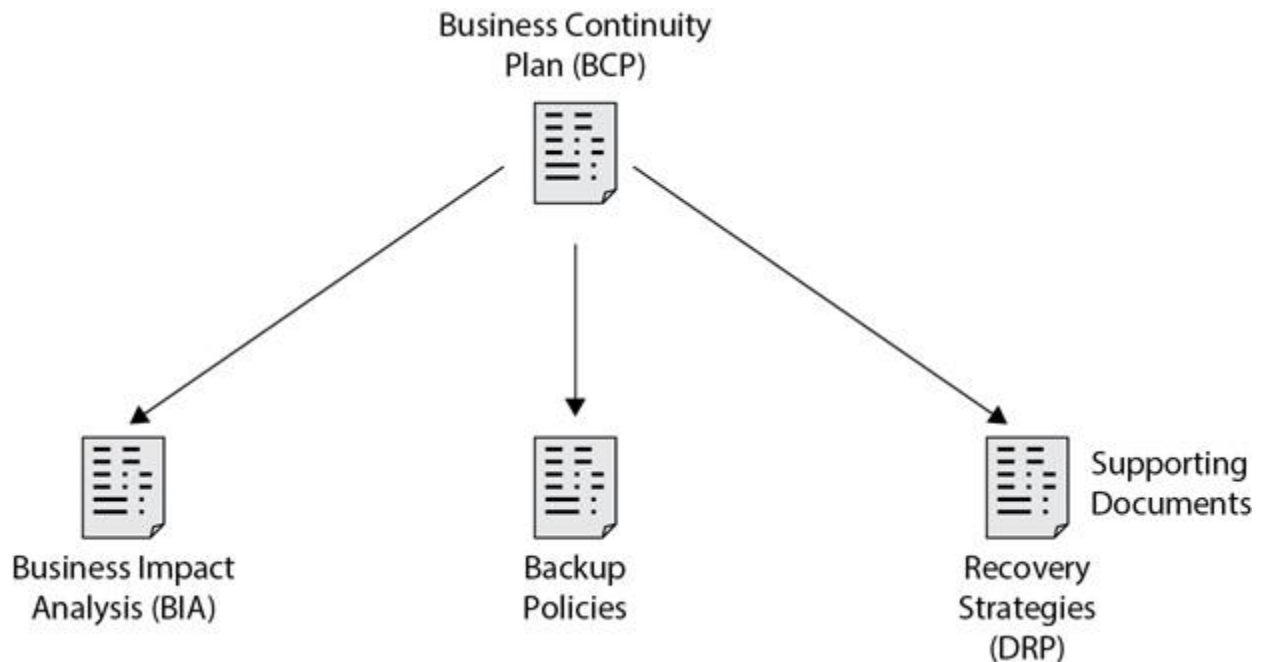


Figure 12-4 BCP and supporting elements



TIP The BCP is the primary document used for business continuity. It includes a BIA and other supporting documents such as backup policies and one or more DRPs.

Some organizations are in high-risk zones for specific types of events, so they often create supporting documents specifically for these events. For example, an organization that is in a high-risk zone for hurricanes would have a document specifically targeted at ensuring that critical operations continue during and after a hurricane.

Before the hurricane season begins, the organization dusts off last year's hurricane preparation plan and reviews it. It may decide that it wants to ensure that it is able to continue to operate critical functions for at least ten days after a hurricane hits. This drives other requirements, such as ensuring the location has at least ten days' worth of diesel fuel for generators and at least ten days' worth of food and water for personnel remaining onsite.

This plan would also identify steps to take as a hurricane is approaching, such as what to do when the hurricane is within 72, 48, 24, and 12 hours of striking. For example, when the hurricane is within 72 hours, personnel check to ensure they have enough supplies on hand and begin bringing everything inside to prevent potential damage. When the hurricane is within 48 hours of the facility, the organization tests the generators and ensures that everything can continue to operate if the storm hits. When the hurricane is within 24 hours, emergency personnel report to work and all other personnel go home.



TIP The actual steps that the organization takes are identified in the hurricane preparation document and may be different from the steps identified in this timeline. However, the point is that the members don't have to guess about what they should do. Instead, they simply follow the checklists. Even if this is the first time that some of the personnel are experiencing a hurricane, they have the benefit of others' experience documented in these checklists.

Business Impact Analysis

A BIA is an important part of a BCP. It identifies the impact to the organization if any business functions are lost due to any type of incident. This information helps an organization identify what business functions are critical to continued operation. It also helps identify which resources are required to support these critical business functions.

For example, will the business stop functioning if a weather event causes a 24-hour or 48-hour power outage in a building? The answer depends on what functions are performed in the building and whether any of these functions are critical. If this building houses primary servers used in an online e-commerce website, the loss of power could cause the entire business to cease functioning. In this case, the purchase of generators for long-term alternative power may be justified. However, if the building houses products and shipping facilities, a 24-hour or 48-hour power outage may result in only a minor impact on operations.

The BIA evaluates both direct costs and indirect costs. As mentioned previously, direct costs are associated with the immediate loss, such as an immediate loss to sales, while indirect costs are expenses related to recovering from the loss. For example, if an e-commerce website normally generates \$5,000 in revenue an hour, then the direct costs are \$5,000 for every hour the server is down. However, the organization may have to spend money on public relations after an outage to assure customers that their personal information is secure, and this extra money represents indirect costs. Additionally, some customers may leave forever, causing additional indirect costs.

Maximum Acceptable Outage

One of the most important findings of the BIA is the maximum acceptable outage (MAO) for different systems or services. This is sometimes called the maximum tolerable downtime (MTD) or the maximum tolerable outage (MTO). If an outage lasts longer than the MAO, it affects critical business functions supported by these systems or services.

As an example, consider an online website generating \$5,000 an hour during peak hours. After 60 minutes of downtime, it loses \$5,000 in direct revenue. Management may determine that it has an MAO of 60 minutes. This includes all the elements required to support the online website, including web servers, backend databases, Internet access, and the network infrastructure.

The same company might ship products from a separate shipping warehouse. Management may decide that warehouse outages up to 48 hours will not affect critical business functions. As long as they communicate the delays with customers, the customers will accept some shipping delays. This results in an MAO of 48 hours for the shipping functions.



TIP Maximum allowable outage (MAO) is sometimes called maximum tolerable outage (MTO) or maximum tolerable downtime (MTD). All three terms refer to the maximum outage time before the outage affects critical business functions.

Recovery Time Objective

The MAO helps an organization determine the recovery time objective (RTO), which is the maximum amount of time that personnel can take to restore a system or service after an outage. If personnel take longer than the RTO to restore the system, the outage will affect critical business functions. If the MAO is 60 minutes, the RTO is also 60 minutes.

Administrative and security personnel also use the RTO to identify fault tolerance and redundancy methods. This includes Redundant Array of Independent Disks (RAID) subsystems and redundant servers configured in a failover cluster or load-balancing cluster. Implementing fault-tolerant and redundancy methods helps prevent the outages from occurring at all. For example, if a hard drive fails in a RAID, the system continues to operate. Similarly, if a server fails in a failover cluster, another server picks up the load and the service continues to operate.

These options can be expensive. However, if a web server generates \$5,000 in revenue an hour and the MAO is 60 minutes, management may decide that the cost is justified.

Recovery Point Objective

A recovery term often associated with databases is the recovery point objective (RPO). It identifies the maximum amount of data that an organization is willing to lose. This is sometimes difficult to identify. If you ask managers how much data they are willing to lose, you'll often hear "none" as the answer. However, after you identify the cost of restoring data up to the moment of failure, you will often get a more realistic answer.



NOTE Although an RPO is most often associated with databases, the same concepts can be applied to any type of backup.

If a database is relatively static, with just a few manual updates during the week, then management may determine that the RPO is a single week. In other words, backing up the database once a week is enough to meet this requirement. If a disaster deletes or corrupts the database, administrators can restore the database and personnel can manually enter the updates since the last backup.

In contrast, consider a database recording online transactions for a web server generating \$5,000 in sales an hour. Weekly backups aren't enough. If personnel back up the database on Sunday and a catastrophic database failure occurs on Wednesday, it results in the loss of all the customer transactions since Sunday. In this scenario, the RPO may be up to the moment of failure.

Although it's certainly possible to create a backup strategy for databases to recover a failed database up to the moment of failure, such a strategy can be expensive. That's why it's important to complete the BIA to help determine the actual needs of the organization related to the RPO. If the BIA determines the organization has a need to restore the database up to the moment of failure, it justifies the extra expense to support this RPO.



TIP Database backup strategies can include full and differential backups and the backup of database transaction logs. The transaction logs record every database transaction since the last backup (full or differential) and can be used to restore the database up to the moment of failure. Transaction logs are typically transferred to a different server for storage to prevent their loss if the primary server suffers a catastrophic failure.

BIA Output

One of the major deliverables of the BIA is a document that identifies the actual monetary losses that can result from the outage of a critical business function. This is invaluable when you are evaluating controls that can mitigate these losses. You may remember from [Chapter 9](#) that the controls are valuable only if they cost less than the losses they're trying to protect. For example, you wouldn't spend \$1,000 to protect a \$10 keyboard. Similarly, you don't want to spend business continuity funds protecting noncritical business functions. Although an e-commerce web server may be critical to an organization, an intranet web server used to track long-term projects may not be as critical.

The BIA identifies how much money can be lost if the e-commerce web server fails. This provides management with an idea of the justified costs to keep the system operational even during a major incident. When evaluating multiple business functions, the BIA also attempts to prioritize them.

[Figure 12-5](#) shows one way that a BIA can prioritize different functions by analyzing the probability of an incident occurring and the impact to the organization if the incident occurs. Incidents with a high probability and a high impact have the highest priority, while incidents

with the lowest probability and lowest impact have the lowest priority. By graphing different types of disasters and identifying the systems that they can affect, an organization can identify the incidents that deserve the greatest attention.

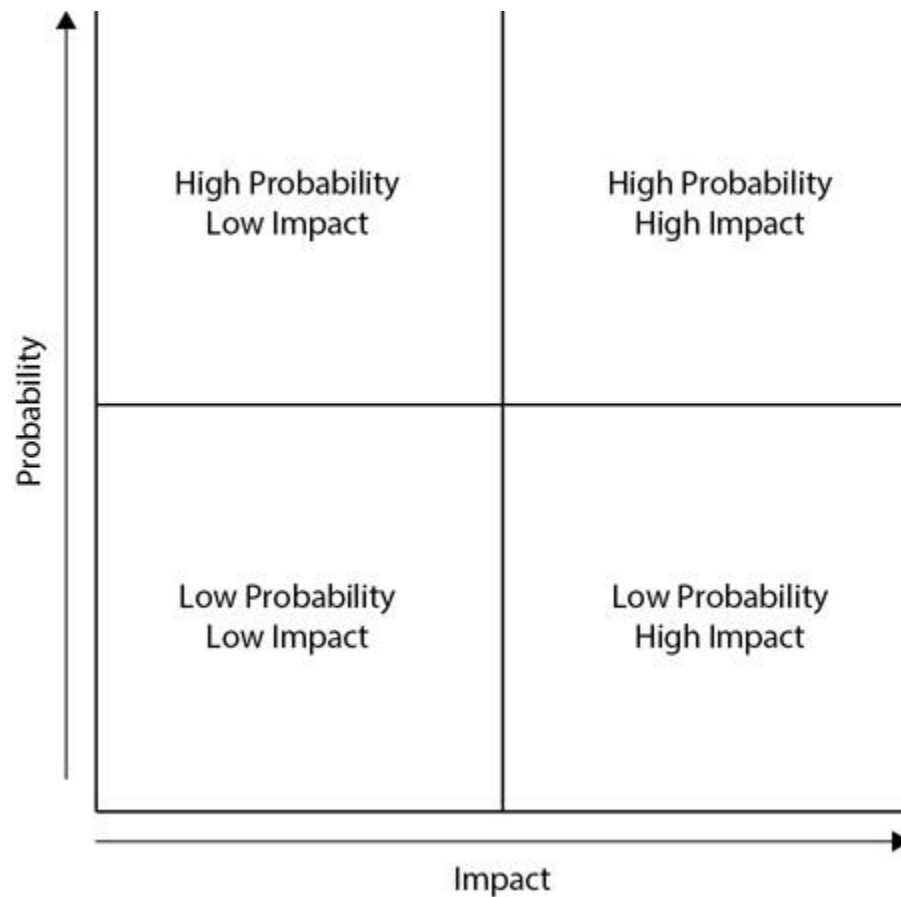


Figure 12-5 Probability and impact analysis

Historical data helps an organization predict the probability of specific disasters. For example, an organization in San Francisco may consider the threat of an earthquake to be significant, while the threat of a hurricane may be nonexistent. Similarly, an organization in Miami may consider the threat of a hurricane to be significant, while the threat of an earthquake may be nonexistent. These are extreme examples, of course, and many other examples aren't so obvious. However, by analyzing and charting the results, the organization can prioritize different risks.

It's important to remember that historical data won't necessarily cover all potential negative incidents. Still, it is possible to predict some of these negative events. As an example, an organization might never have suffered any losses from a cyber-attack in the past. However, with the level of ongoing cyber-attacks, it's very likely the organization will be attacked at some point in the future. Personnel can predict the potential impact of a data breach based on the value of the organization's data.

Disaster Recovery Plan

A DRP provides an organization with a plan to restore critical operations after a disaster. The overall goal is to provide employees who are recovering the systems with clear-cut steps to restore the systems.

Disasters are major events, and people don't always think clearly when they occur. If employees have steps to follow within a checklist, they can ensure that they restore systems in a logical and orderly manner. On the other hand, if employees don't have steps to follow, they may spend a lot of time trying to determine what to do, and then more time redoing their steps when things don't work as expected.



EXAM TIP Disaster recovery and fault tolerance are not the same thing. Disaster recovery helps an organization recover after a disaster. Fault tolerance helps ensure that a system or component continues to function after a failure. [Chapter 9](#) covered fault-tolerance strategies such as RAID and failover clusters.

Emergency Response Plans and Procedures

Disaster recovery plans include emergency response plans and procedures, and an organization will often have several different DRPs in place, ready to respond to different types of disasters. For example, organizations typically have plans for responding to disasters that can hit any organization, such as a fire or a crippling cyber-attack. Organizations will also have other plans that are based on local risks.

These additional plans address weather and nature-based events such as floods, tornadoes, hurricanes, cyclones, and earthquakes. While it's difficult to predict exactly when these events will occur, reviewing the history of an area can predict the likelihood that one of these events will occur at some point in the future. For example, no one should be surprised if a tornado touches down in Oklahoma during May. Similarly, no one should be surprised if a hurricane threatens land bordering the Gulf of Mexico during September or October.

Comparing a BCP and a DRP

It's important to recognize the differences between a BCP and a DRP. Many people combine them as though they are synonymous, but they aren't. In short, the BCP has a much wider scope and helps an organization continue to operate. In contrast, a DRP is a part of the BCP with a primary purpose of restoring a system to operation after an incident results in a major failure.

For example, if an organization is in the midst of a disaster and needs to move operations to an alternative location, the BCP would provide the details on the alternative location and what

business functions it will perform. After the disaster occurs, the DRP would provide the details on how to restore operations for specific business functions.



EXAM TIP The BCP provides the information to keep critical functions running during a disaster (such as which critical functions to move to an alternative location). A DRP has a narrower focus and identifies how to recover one or more individual systems after a failure.

Even though a BCP and a DRP are not the same thing, you may run across documentation that treats them as though they are the same. However, (ISC)² documentation distinguishes between the two.

Restoration Planning

A key part of business continuity planning is restoration planning. Putting the concepts together, they typically occur in the following order:

- **Identify critical business functions** The BIA provides this information. Additionally, the BIA identifies the systems that support these critical business functions. For example, an online website might include one or more web servers, a backend database, Internet access, and a supporting network infrastructure.
- **Identify restore targets** The BIA identifies the MAO, and personnel use this to identify the RTO (and RPO if relevant).
- **Create plans to restore systems** These plans identify how to restore systems supporting critical business functions within the required times.

These plans need to be extremely clear and detail the exact order of restoration. For example, consider the online website with a backend database hosted on a database server. If the web application expects to connect with the database, it might have unexpected problems if it cannot connect to it. If that's the case, restoration plans might require technicians to restore the database server first, and then restore the web server. If the database is hosted on a group of servers within a failover cluster, the failover cluster might require servers to be restored in a specific order to ensure the failover cluster remains operational. Similarly, if the website is hosted within a load-balancing cluster of multiple servers, the cluster might require technicians to restore these servers in a specific order.

All of this might be common knowledge to the technicians and administrators who maintain these systems. However, you never know if these technicians and administrators will be available to restore the systems during a disaster. When the restoration plans provide clear directions on what to restore, in what order to restore it, and how to restore it, it makes it easier for any technician or administrator to complete the restoration in the shortest amount of time.

Testing and Drills

The first draft of any plan is rarely perfect. However, testing helps identify the problems so that they can be resolved before an actual disaster. NIST SP 800-84, *Guide to Test, Training, and Exercise Programs for IT Plans and Capabilities*, provides some excellent guidance on testing methodologies. It discusses tabletop exercises and functional exercises in-depth.

- **Tabletop exercises** These are discussion-based exercises where personnel meet in a conference room or classroom setting. A facilitator presents a scenario such as a hurricane or tornado and asks the exercise participants questions related to their roles. Talking through the scenario gives participants a better chance to identify potential flaws in the written plans.
- **Functional exercise** Personnel actually perform the steps outlined in the plan as a method of validating the plan. Functional exercises can be focused on specific elements of a plan or be very broad in their scope. For example, you can give a technician a plan to restore a database server and determine if the technician is able to restore it on a test server using the plan. The actual database server remains online, so the test does not affect actual operations. The exercise can also encompass a full-scale exercise designed to address all the plan elements. While it is only a simulation, it allows personnel to perform their roles as they would in an actual emergency.

In some cases, plans require personnel to take specific actions. A classic example is a fire. If a fire occurs, personnel should exit the building using a predefined path, which helps the maximum number of people exit the building in the shortest amount of time. A simple way to test this type of plan is to hold a drill. The drill requires personnel to respond as if the emergency is occurring. In the case of a fire drill, people hear the fire alarm and exit the building. Fire departments are often happy to participate in these types of drills and provide feedback to the organization on how to improve their plans.

Alternative Locations

If an organization determines that it must be able to continue operations even if a location suffers a catastrophic failure, it designates an alternative location. For example, banks and investment companies need to provide access to customer data even if a weather event (such as a tornado, flood, or hurricane) causes so much damage to one location that they can no longer operate at that location.

Alternative locations include hot sites, cold sites, warm sites, and mobile sites. [Figure 12-6](#) provides an overall view of each of these alternative location types, and the following sections explain them in more depth.

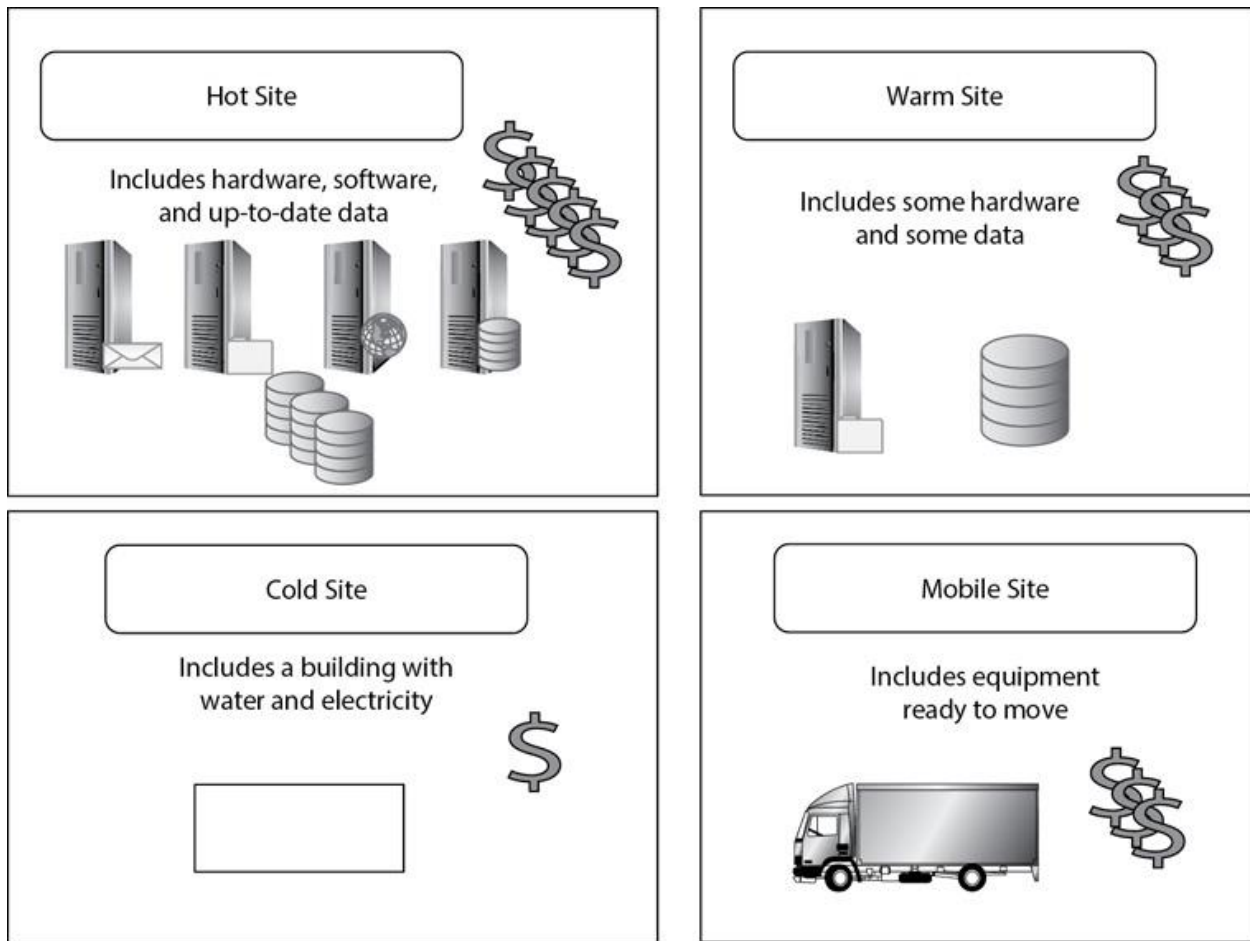


Figure 12-6 Comparison of alternative locations

Hot Site

A hot site includes all the resources necessary to take over operations of another location in a very short time, typically within an hour and sometimes sooner. A hot site includes hardware such as servers and the network infrastructure, up-to-date data, and personnel to manage the functions of the alternative location.



TIP An extreme version of a hot site is a mirrored site. A mirrored site is identical to the primary location and provides 100 percent availability. As data changes at the primary location, real-time data transfers send the updates to the mirrored site. This ensures that the mirrored site is always up and operational. In contrast, a hot site might take as long as an hour or so to become operational.

Hot sites are valuable if an organization needs to prepare for a catastrophic failure that can occur at any time. For example, many banks and investment companies have hot sites ready to take over in case of major incidents, such as the terrorist attacks of September 11, 2001.

Organizations that operate in earthquake or tornado zones often need to have hot sites available far away from the primary location. Hot sites are the most expensive of the alternative location types, but they are the easiest to test, because everything is in place.

It's rare for a hot site to be used only as an alternative location for an organization. Instead, the organization typically uses it for less critical operations the majority of the time, but keeps it in a ready state to take over critical operations at a moment's notice. For example, the organization could use the location to provide basic customer services, but also host the equipment and infrastructure as a hot site. If an emergency requires the organization to activate the hot site, personnel at the alternative location scale back customer services temporarily and dedicate more resources to take over for the primary location. As additional personnel arrive at the hot site, it will eventually assume all operations.

Similarly, regional locations often have the capability of taking over another region's services. For example, an organization can have major locations in the western and eastern parts of a country. If either region's location fails, the other location can take over the critical operations of the failed location.

Cold Site

A cold site is a building with a roof, running water, and electricity. It doesn't include the necessary hardware, software, or personnel. In the event of an emergency, personnel move all the resources to the cold site location, hook it up, and configure the site for operation.

It's difficult to test a cold site because nothing exists at the location. However, a cold site is significantly cheaper to maintain than other types of sites.

Warm Site

A warm site is a compromise between a cold site and a hot site. The organization makes compromises with costs and time. Instead of ensuring that the site can be activated within minutes, the organization may decide that 24 hours' notice is enough to bring it online.

For example, an organization may operate in a hurricane zone. Because hurricanes travel relatively slowly, communities often have more than a day's notice of a possible hurricane landfall. When a storm is approaching, the organization can use the time to activate a warm site. This includes deploying personnel to the warm site and updating all the data.

Mobile Site

When an organization doesn't want to use a dedicated location as an alternative location, it can create a mobile site. For example, it's possible to set up the inside of a storage container as an

alternative location. These are the same storage containers that 18-wheeler tractor-trailers use to haul goods, and they are as easy to move as any other loads hauled over the highways.

Organizations outfit these storage containers with all the equipment needed to support critical operations. If an emergency requires activating the mobile site, the organization moves it to the desired location and hooks it up to power. Depending on its needs, the organization might use an external heating and ventilation system, or bring one with the container and then mount it on the roof or outside the storage container. Similarly, the organization might connect directly to a broadband Internet service provider (ISP) for Internet access at the mobile location, or use a satellite or cellular connection.

Mobile sites like this will have all the necessary equipment, but the software and data are typically out of date. When the mobile site reaches its destination, personnel will ensure that the equipment is up to date with all appropriate patches and approved changes and that the data is up to date.

Mobile sites don't have to be as elaborate as a fully equipped storage container. For example, one small company in Virginia Beach has a simpler mobile site plan. When a hurricane is approaching, the owner packs the company's primary server into her car and drives west. When she settles somewhere, she plugs a USB wireless adapter into the server, configures her phone as a personal hotspot, and uses it to provide Internet access for the server.

The cost of a mobile site is similar to the cost of a warm site—somewhere in the middle of a cold site and a hot site. The actual cost varies depending on what an organization chooses to use for its mobile site. Clearly, the packed storage container will be much more expensive than the owner's car, smartphone, and a single server.

Identifying Security Organizations

Several security organizations exist that provide different levels of assistance and documentation for IT security personnel. It's worthwhile knowing about some of these organizations and their available documentation.

NIST

Several chapters within this book mentioned the National Institute of Standards and Technology (NIST). NIST includes a Computer Security Division and an Information Technology Laboratory (ITL). The organization regularly performs research in the area of information technology in direct support of federal agencies. However, NIST uses federal funds to perform this research, so it makes this research publicly available as Special Publications (SPs).

SP 800 Documents

The SP 800 series of documents includes over 100 well-researched documents on a wide range of IT security topics, including cloud computing, risk management, intrusion detection and prevention systems, protection of wireless transmissions, protection of personally identifiable

information (PII), and much more. You can view a full list of these documents (and download any of them) at <https://csrc.nist.gov/publications/sp800>.



TIP NIST also publishes Federal Information Processing Standards (FIPS) documents, which define specific minimum requirements for many federal agencies. These documents often provide guidance for specific federal agencies, but they also provide sound guidelines for private companies. You can access them here: <https://csrc.nist.gov/publications/fips>.

Some of the documents that directly relate to security policies and continuity planning include

- **SP 800-18** Guide for Developing Security Plans for Federal Information Systems
- **SP 800-30** Guide for Conducting Risk Assessments
- **SP 800-34** Contingency Planning Guide for Federal Information Systems
- **SP 800-37** Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy
- **SP 800-41** Guidelines on Firewalls and Firewall Policy
- **SP 800-53** Security and Privacy Controls for Information Systems and Organizations
- **SP 800-58** Security Considerations for Voice Over IP Systems
- **SP 800-61** Computer Security Incident Handling Guide
- **SP 800-64** Security Considerations in the System Development Life Cycle
- **SP 800-77** Guide to IPsec VPNs
- **SP 800-83** Guide to Malware Incident Prevention and Handling for Desktops and Laptops
- **SP 800-84** Guide to Test, Training, and Exercise Programs for IT Plans and Capabilities
- **SP 800-94** Guide to Intrusion Detection and Prevention Systems (IDPS)
- **SP 800-97** Establishing Wireless Robust Security Networks: A Guide to IEEE 802.11i
- **SP 800-100** Information Security Handbook: A Guide for Managers
- **SP 800-111** Guide to Storage Encryption Technologies for End User Devices

- **SP 800-113** Guide to SSL VPNs
- **SP 800-115** Technical Guide to Information Security Testing and Assessment
- **SP 800-122** Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)
- **SP 800-123** Guide to General Server Security
- **SP 800-124** Guidelines for Managing the Security of Mobile Devices in the Enterprise
- **SP 800-127** Guide to Securing WiMAX Wireless Communications
- **SP 800-183** Networks of ‘Things’

Cybersecurity Framework

NIST published Version 1.1 of the *Framework for Improving Critical Infrastructure Cybersecurity* in April 2018. Security personnel commonly refer to this as the Cybersecurity Framework (CSF). It focuses on using business drivers to guide cybersecurity activities, while considering cybersecurity risk as part of an overall risk management process. The CSF has three primary elements: Framework Core, Implementation Tiers, and Framework Profiles.



TIP The CSF has quickly gained a lot of popularity within the cybersecurity community. Version 1.1 is a relatively short read at only 55 pages and is worth your time if you are working in cybersecurity, or want to break into the cybersecurity field.

The Framework Core is a set of cybersecurity activities, desired outcomes, and references that can be applied to most organizations. It includes industry standards, guidelines, and best practices within five primary Framework Core Functions: Identify, Protect, Detect, Respond, and Recover. Each of these functions can be applied to different categories of a cybersecurity program.

Table 2 (a comprehensive 19-page table) lists between three and six categories within each of these functions. As an example, the Protect function includes a Data Security category. It includes the following eight subcategories, which represent best practices:

- Data-at-rest is protected.
- Data-in-transit is protected.
- Assets are formally managed throughout removal, transfers, and disposition.

- Adequate capacity to ensure availability is maintained.
- Protections against data leaks are implemented.
- Integrity checking mechanisms are used to verify software, firmware, and information integrity.
- The development and testing environment(s) are separate from the production environment.
- Integrity checking mechanisms are used to verify hardware integrity.

Each subcategory includes multiple references that personnel can use to get a better understanding of the topic. Personnel can identify if they are currently implementing this best practice, and if not, use the references to help them improve their current practices.

The Implementation Tiers can be used to identify how an organization views cybersecurity risk. The Tiers are

- **Tier 1: Partial** The lowest tier, Tier 1 indicates that risk management processes are not formalized and there is a limited awareness of cybersecurity risk.
- **Tier 2: Risk Informed** This tier indicates that risk management practices are in place but may not be followed. Personnel may be aware of risk, but an organization-wide approach to managing risk has not been established.
- **Tier 3: Repeatable** This tier indicates that risk management practices are formally approved and documented in policies. Personnel have the knowledge and skills necessary to perform their roles. Personnel at all levels of leadership regularly communicate to each other on cybersecurity risk.
- **Tier 4: Adaptive** The highest tier, Tier 4 indicates that the organization adapts its practices based on previous and current cybersecurity activities. It has processes in place for continuous improvement. Executive management evaluates cybersecurity risk in the same manner that they evaluate financial risk.

The Framework Profiles can help an organization ensure that the Functions within the Framework Core are aligned with the business requirements. As an example, cybersecurity personnel can evaluate all the Categories and Subcategories within the Framework Core and identify which tier they are in. This can provide the organization a baseline.

Personnel can then work with senior management to identify a target profile based on senior management's vision, risk tolerance level, and amount of resources management is willing to devote to cybersecurity.

Monthly Bulletins

NIST also publishes monthly bulletins available here: <https://csrc.nist.gov/publications/itl-bulletin>. Each monthly bulletin focuses on a single topic of interest related to computer security. For example, the March 2018 bulletin, titled “Safeguards for Securing Virtualized Servers,” summarizes technical guidelines related to securing hypervisor platforms.

ITL vs. ITIL

The Information Technology Laboratory (ITL) is not related to ITIL. ITIL (formally known as the Information Technology Infrastructure Library) is a group of books documenting best practices for IT management. The United Kingdom Office of Government Commerce created the original ITIL books, but ITIL has gone through several changes over the years. Axelos, Ltd. currently owns all of the ITIL intellectual property, and it licenses the use of ITIL materials to other organizations. Axelos also accredits organizations as licensed Examination Institutes for the ITIL-related certifications.

In contrast, a U.S. government entity creates the ITL documents, which are freely available to anyone. The SP 800 series of documents are public-domain documents rather than intellectual property. Additionally, ITL doesn't sponsor any certifications.

US-CERT

The United States Computer Emergency Readiness Team (US-CERT) provides response support and defense against cyber-attacks for several government entities in the United States and for some industry and international partners. It is operated as part of the National Cyber Security Division (NCSD) of the Department of Homeland Security as a public/private partnership.



TIP As many as 250 organizations use “CERT” in their name and deal with IT security incidents and responses. US-CERT is not affiliated with any of these organizations, although it does work with some of them. For example, US-CERT coordinates with the CERT Division of the Software Engineering Institute (SEI), a federally funded research and development center at Carnegie Mellon University.

US-CERT sponsors the National Cyber Awareness System, which includes several methods of providing information to users and technical experts. [Chapter 6](#) mentioned several of the newsletters that US-CERT publishes, and you can sign up for any of them here: <https://www.us-cert.gov/ncas>.

SANS Institute

While the SANS Institute is a private institution that sponsors other security certifications and sells training for these certifications, it also publishes some free resources. For example, the SANS Reading Room includes more than 2,780 free white papers on a wide variety of computer security topics. You can access the SANS Reading Room here: www.sans.org/reading-room/.



NOTE SANS looks like an acronym but it is the name of the company. However, the name is derived from Sysadmin, Audit, Networking, and Security.

SANS also sponsors the Internet Storm Center, mentioned in [Chapter 6](#). The Internet Storm Center monitors threats related to malware and provides free information via diaries (<https://isc.sans.edu/diaryarchive.html>) and podcasts (<https://isc.sans.edu/dailypodcast.xml>).

CERT Division

The first CERT Coordination Center was created at Carnegie Mellon University in response to the Morris worm that struck in 1998, and was later renamed as CERT Division. It is a federally funded program located in the Software Engineering Institute (<https://www.sei.cmu.edu/about/divisions/cert/index.cfm>). The CERT Division works closely with the government, industry, law enforcement, and academia personnel to improve security of computer systems and networks. You can read more about it here: <https://www.sei.cmu.edu/about/divisions/cert/index.cfm>.

Chapter Review

A security policy is a written document, authoritative in nature, that provides a high-level view of the security goals for an organization. External standards may influence the security policy, but the organization decides what standards to follow. Personnel develop procedures based on the guidance from the security policy.

Senior-level management should create, or at least endorse, the organization's security policy. Once the security policy is created, personnel within the organization implement and enforce it. Users must be made aware of the contents of the security policy through training, warning banners, posters, and other awareness methods, and they should be asked to read and sign an acceptable use policy. It's common to review a security policy periodically, such as once a year or after a security incident.

Business continuity plans (BCPs) include processes and procedures to prevent the loss of mission-critical services. A BCP starts with a business impact analysis (BIA) and includes one or

more disaster recovery plans (DRPs). It's important to recognize that a BCP is not the same as a DRP, but a DRP is often a component of a BCP.

The BIA identifies the maximum acceptable outage (MAO) time (also called maximum tolerable downtime, or MTD) for critical services and systems. If an outage for one of the critical services or systems exceeds this time, it impacts critical business functions. The MAO drives the recovery time objective (RTO), or the maximum amount of time that the organization can take to restore a critical service or system. The recovery point objective (RPO) refers to the point in time to which a database needs to be recovered if the system fails.

A DRP provides the steps required to restore a system after an outage. Organizations test both BCPs and DRPs using various testing methods. Tabletop exercises allow personnel to talk through the steps of a given scenario. Functional exercises simulate an event and allow personnel to walk through the steps of a given scenario.

Some BCPs require the use of alternative locations. A hot site includes all the hardware, software, and up-to-date data required to take over operations at a moment's notice. A cold site includes a roof, water, and electricity, but little else. A warm site is a compromise between a hot site and a cold site. Some organizations use a mobile site instead of a fixed location, which allows the organization to move its operations to another location during an emergency.

NIST includes a Computer Security Division and an Information Technology Laboratory. The ITL has published over 100 Special Publications in the SP 800 series on a wide range of IT security topics. All of these special publications are available for free download here: <https://csrc.nist.gov/publications/sp800>. Other organizations that provide security-related materials include US-CERT and the SANS Institute.

Questions

- 1.** Of the following choices, which provides the highest level authority within an organization?
 - A.** Standards
 - B.** Policies
 - C.** Guidelines
 - D.** Procedures

- 2.** Which of the following choices best describes an organization's security policy?
 - A.** An authoritative written document that identifies an organization's overall security goals
 - B.** A nonauthoritative written document that identifies an organization's overall security goals
 - C.** A technical control that mitigates risks

D. A baseline used to ensure that systems are secure when deployed

3. Which of the following best describes the purpose of a security policy?

A. Ensures personnel understand their responsibilities

B. Ensures personnel use strong authentication

C. Informs personnel of management priorities related to security

D. Provides guidance on management controls

4. An organization wants to ensure that users are aware of their responsibilities related to the use of IT systems. What should the organization create?

A. A video monitoring system

B. An audio monitoring system

C. An acceptable use policy

D. An account lockout policy

5. Sally notices that Homer appears to be stealing from the company. What should Sally do?

A. Confront Homer

B. Ignore the activity because it doesn't concern her

C. Call the police

D. Report the activity to a manager

6. Which of the following choices are effective methods of ensuring that employees know the relevant contents of an organization's security policy? (Choose all that apply.)

A. Providing training

B. Using warning banners

C. Using posters

D. Storing the policy in the company vault

7. An organization has a security policy in place. What can personnel within the organization do to ensure it remains relevant?

- A. Perform audits
- B. Perform training
- C. Review it
- D. Test it

8. Which of the following is the most important element of business continuity planning?

- A. Support from senior management
- B. Availability of a warm site
- C. The backup plan
- D. Cost

9. What is the purpose of a BIA?

- A. To identify recovery plans
- B. To drive the creation of the BCP
- C. To test recovery plans
- D. To identify critical business functions

10. Which one of the following is a valid step to perform during a business impact analysis?

- A. Identify alternative locations
- B. Create a plan to restore critical operations
- C. Identify resources needed by critical business functions
- D. Identify minimum outage times for key business services

11. What is MTO in relation to business continuity planning?

- A. Minimum time for an outage
- B. Maximum time for an outage
- C. Minimum tolerable outage

D. Maximum tolerable outage

12. Which of the following best describes maximum tolerable downtime?

A. The maximum amount of downtime before a business loses viability

B. The point in time in which a failed database should be restored

C. The maximum amount of time that can be taken to restore a system or process

D. The minimum amount of time that can be taken to restore a system or process

13. What is RPO in relation to business continuity planning?

A. Restoring potential outage

B. Recovery point objective

C. Restoration process option

D. Recovery process options

14. What is RTO in relation to business continuity planning?

A. Recovery terminal objective

B. Recovery time objective

C. Recovery tolerable outage

D. Recovery tolerable objective

15. An organization's location has been hit by a tornado and the organization is moving to an alternative location. What provides the direction for this action?

A. BIA

B. BCP

C. DRP

D. Hot site

16. An organization decides to designate an alternative location to be used in case of an emergency. The organization doesn't need anything other than an open building with water and electricity. What type of site best meets this need?

- A. Hot
- B. Warm
- C. Cold
- D. Distant

17. An organization decides to designate an alternative location for operations during a disaster. The site must be up and operational within minutes of an outage at the primary location. What type of site best meets this need?

- A. Mobile
- B. Cold
- C. Warm
- D. Hot

18. An organization is updating its business continuity plan (BCP) and wants to implement an alternative location that is the easiest to relocate. What type of site best meets this need?

- A. Cold
- B. Hot
- C. Mobile
- D. Warm

19. Of the following choices, what is a U.S. government entity that regularly publishes Special Publications (known as SP 800 series documents) related to IT security?

- A. ITIL
- B. NIST
- C. CERT Division
- D. US-CERT

20. Which of the following organizations provides regular cybersecurity alerts about current security issues, vulnerabilities, and exploits as part of the U.S. National Cyber Awareness System?

- A. ITL
- B. NIST
- C. CERT Division
- D. US-CERT

Answers

1. B. Policies are high-level documents and provide authoritative direction. Standards can influence a policy, but an organization chooses what standards to follow. Guidelines provide recommendations, but they are not mandatory. Personnel create procedures based on the policy.

2. A. A security policy is an authoritative (not nonauthoritative) written document that identifies an organization's overall security goals. It is a management (or administrative) control, not a technical control or a baseline; however, technical controls and baselines are created based on the direction from the security policy.

3. C. Of the choices, the best description is that the security policy informs personnel of management priorities related to security. The other answers provide some specific goals of a security policy but do not address the overall purpose. An acceptable use policy helps ensure personnel understand their responsibilities. Technical controls ensure personnel use strong authentication, but a security policy covers more than just technical controls. Similarly, a security policy provides guidance on more than just management controls.

4. C. An acceptable use policy (also called an acceptable usage policy) lets users know what is acceptable use of computer equipment and networks. Monitoring systems ensure users follow policies, but they don't ensure users know the policy. Account lockout policies lock out users after too many failed password attempts.

5. D. An employee's loyalty should be to the organization, so Sally should report this activity to a manager or supervisor. An organization's ethics policy often includes procedures for reporting these types of incidents. Confronting Homer may not solve the problem but instead may result in Homer causing problems for Sally, especially if Homer is stealing from the company. Because continued employment is based on the success of an organization, losses are the concern of every employee. The organization should make the decision of whether or not to call the police.

6. A, B, C. Providing training, using warning banners, and using posters are all effective methods of ensuring that employees know the relevant contents of a security policy. If the security policy is stored in the company vault, it won't be accessible to employees.

7. C. A security policy should be reviewed on a regular basis (such as once a year or after a security incident) to ensure that it is still relevant. Audits help to prove that the security policy is being used and enforced. Training ensures that people know the contents of the security policy. It's appropriate to test a BCP or a DRP, but not a security policy.

8. A. Of the available answers, the most important element is support from senior management. While an organization might decide it needs a warm site, not all BCPs require warm sites. A security policy may mandate the creation of backup plans, but this is separate from the BCP. The cost is a concern, but the requirements drive the cost, and without support from senior management, business continuity planning may not receive any funding.

9. D. The business impact analysis (BIA) identifies critical business functions and is a part of the BCP. Personnel create recovery plans later in the process, after creating recovery strategies. The BCP drives the creation of the BIA, not the other way around as suggested by answer B. You can only test the plans after personnel have created them.

10. C. A core goal of a BIA is to identify critical business functions and the resources needed by these critical business functions. Identifying alternative locations is part of a business continuity plan (BCP). A disaster recovery plan (DRP) is a plan to restore critical operations. A BIA does identify maximum acceptable outage times, but not minimum outage times.

11. D. The maximum tolerable outage (MTO), sometimes called maximum allowable outage (MAO) or maximum tolerable downtime (MTD), identifies the maximum amount of time that a system can be down before critical business functions are affected. The *T* does not represent time, and the *M* does not represent minimum.

12. A. The maximum allowable outage (MAO), sometimes called maximum tolerable downtime (MTD), indicates the maximum amount of downtime a business can tolerate and still maintain viability as a business. Recovery point objective (RPO) indicates the point in time to which a failed database should be restored. Recovery time objective (RTO) represents the maximum amount of time that can be taken to restore a system or process after an outage. MTD is not related to minimum timeframes.

13. B. RPO represents recovery point objective and indicates the point in time to which a failed database should be restored. The other answers are not valid terms for RPO within business continuity planning.

14. B. RTO is an acronym for recovery time objective and represents the maximum amount of time that can be taken to restore a system or process. The RTO is derived from the maximum allowable outage (MAO). The other answers are not valid terms for RTO within business continuity planning.

15. B. The business continuity plan (BCP) provides direction for moving to an alternative location after a disaster; the primary purpose is to continue to provide critical business functions. The business impact analysis (BIA) helps an organization identify what functions are critical. The disaster recovery plan (DRP) has a narrower focus and helps an organization recover one or more systems after the disaster has passed. A hot site is a possible type of alternative location, so it doesn't provide direction for the action.

16. C. A cold site is a building with water and electricity and nothing else (such as no computer equipment or data). A hot site includes everything and is ready to take over operations within a

short time after an outage. A warm site is a compromise between the two. A distant site isn't a valid term for an alternative location.

17. D. A hot site includes everything and is ready to take over operations within minutes after an outage. A mirrored site (not available as an answer) is a type of hot site that can take over operations almost immediately. A mobile site will take much longer to relocate and set up. A cold site is a building with water and electricity, and it would take much longer to become operational. A warm site is a compromise between a hot site and a cold site and takes longer to become operational than a hot site.

18. C. A mobile site is the easiest to relocate. Hot, cold, and warm sites use a designated location that the organization either purchases or leases.

19. B. The National Institute of Standards and Technology (NIST) publishes documents in the SP 800 series related to IT security. ITIL is a United Kingdom project. CERT Division is a federally funded program located in the Software Engineering Institute at Carnegie Mellon University, and the United States Computer Emergency Readiness Team (US-CERT) provides response support and defense against cyber-attacks. While CERT Division and US-CERT publish documents related to IT security, they don't publish SP 800 series documents.

20. D. The U.S. Computer Emergency Response Team (US-CERT) manages the National Cyber Awareness System, which provides cybersecurity alerts, bulletins, tips, and updates. The other organizations provide information, but not alerts through the National Cyber Awareness System.

Find answers on the fly, or master something new. Subscribe today. [See pricing options.](#)

CHAPTER 13

Legal Issues

In this chapter, you will learn about

- Computer forensics
- The incident lifecycle
- The three phases of a computer forensic investigation
- The importance of preserving and protecting evidence
- How to protect volatile RAM and data on disk drives
- The difference between computer abuse and computer crime
- How mandatory vacations and job rotation help prevent fraud
- Some of the laws and regulations designed to protect PII

Exploring Computer Forensics

Computer forensics is the science of examining and inspecting computer systems for evidence about an event or crime. The examination includes recovering data stored on or transmitted by any electronic device. A primary goal when identifying and recovering digital evidence is to preserve its integrity. This ensures that forensic experts can analyze it and that the evidence can be used in legal proceedings if necessary.

You've probably watched one of the CSI television shows where detectives use crime scene investigation (CSI) techniques to piece together the elements of a crime. Investigators are careful to ensure they don't modify the evidence as they collect it, and they protect the evidence after collecting it. Computer forensic experts have the same goals as they try to piece together evidence to investigate computer abuse or computer crime. The only difference is that computer forensic experts typically use different tools and techniques when they are investigating an incident. The following sections describe some of the important concepts to consider related to incident handling and computer forensics.

Supporting the Incident Lifecycle

As an IT or security professional, you are sure to observe security incidents, so you should be prepared to participate if asked. With that in mind, it's important to understand the overall incident lifecycle process. [Chapter 7](#) included a section on responding to incidents and covered methods documented in National Institute of Standards and Technology (NIST) Special

Publication (SP) 800-61, *Computer Security Incident Handling Guide*. This section covers the topics from a slightly different perspective using the terminology in the SSCP objectives.

As a reminder from [Chapter 7](#), a security incident is any violation of policies or security practices that has the potential to result in an adverse event. Negative consequences include system crashes, unauthorized access to data, and the execution of malicious code on a computing system. [Chapter 1](#) introduced the CIA triad (confidentiality, integrity, and availability). An adverse effect includes any event that results in the loss of confidentiality, integrity, or availability and is also considered a security incident.

Once an incident occurs, the primary goal of incident handling is to contain any potential damage and then repair the damage. The following sections discuss the individual phases in an incident lifecycle.



TIP Incident handling processes are generally the same. However, different documentation labels these processes with different terms. Similarly, different organizations may label the processes differently in their security policies. Within an organization, the correct terminology is what the managers of the organization have documented within these policies.

Preparation

The first step an organization takes in incident handling is preparation. Preparation includes implementing security controls (or countermeasures) to prevent, detect, and correct security incidents, as discussed in [Chapter 9](#). Organizations typically create an incident response policy that outlines its plans to prepare for and respond to an incident.

Detection, Analysis, and Escalation

IT and security personnel often don't know an event is an incident until they respond and begin investigating and analyzing the event. For example, a user may report erratic behavior on a computer. After investigating and analyzing the user's computer, IT administrators could discover malware infecting the user's system. Deeper analysis would identify the extent of the malware infection.

In this scenario, the user's report is the event. Once administrators detect and analyze the malware, they can promote the event to an identified incident. The administrators may find that they can resolve the incident by running antivirus software and either quarantining the malware or removing it completely.

Personnel escalate an incident if they realize they cannot resolve it. This is similar to a help desk technician escalating a help desk ticket when the technician cannot solve the customer's problem. A more experienced technician takes over and attempts to resolve the problem.

Similarly, if technicians or administrators discover a security incident but realize they cannot resolve it, they take steps to escalate it. Often, this simply means they report the incident to their supervisor or a higher-level administrator to get some help.

Personnel do not escalate all incidents. For example, if a technician discovers malware and uses antivirus software to remove it, there is no need to involve other personnel. However, the technician would still document it using the organization's procedures.

Containment

After discovering an incident, personnel respond in accordance with the organization's incident response plan, which might include actions to take after discovering specific incidents. For example, it's important to contain an incident to limit potential damage. If an administrator discovers a computer is infected with malware, the incident response plan might specifically state that the IT administrator should disconnect the network interface card (NIC) to isolate the computer from the network. If the malware hasn't spread yet, this action contains the malware to the infected computer.

If multiple computers within a network are suspected of being infected, containment may require administrators to isolate the network by manipulating routers or switches. As an example, a network can be connected together with a switch and connected to other networks via a router. By disconnecting the cable from the switch to the router, the network is isolated.

The incident response plan obviously cannot include specific actions for every possible incident. However, it will include generic guidance such as protecting evidence while investigating the incident, containing the incident, escalating the incident, and reporting the incident.



TIP Some malware can detect when a cable has been removed from a NIC and take different actions in response. However, if the switch is isolated from adjoining networks by removing its connectivity to a default gateway or other router, the malware would likely just view this as a network problem.

Eradication

Eradication is the complete removal of all elements of an incident. In the scenario of a simple malware infection, eradication means that the malware has been detected and quarantined or removed. However, it can include much more.

As an example, advanced persistent threats (APTs) typically start with a remote access Trojan (RAT) installed on a single computer when a user clicks a malicious link within a phishing e-mail. Attackers quickly exploit this RAT and look for other computers in the network that they can exploit. It's common for attackers to create multiple backdoor accounts that they can use to

attack other systems. Eradication starts by disabling the RAT, but also includes identifying all the backdoor accounts and removing them.

For APT attacks, complete eradication can be extremely difficult, especially if the RAT was first installed weeks or months before the incident was discovered. Organizations often seek outside help from dedicated cybersecurity organizations or government agencies to fully eradicate an attacker's presence.

Recovery

Recovery includes completing the steps required to return all systems to normal operation and confirming that the systems are operating normally. Depending on the extent of the incident, this can be very simple or very complex. For a simple malware infection, recovery may just be quarantining or removing the malware and running another antivirus scan to verify it remains clean.

For serious incidents, recovery can include multiple actions, such as completely rebuilding systems from images and installing all appropriate patches, restoring data from backups, changing passwords, and tightening network security. After attackers have successfully infiltrated a network, it's common for them to try again. Because of this, recovery may include increasing system logging or network monitoring to detect subsequent attacks.

Reporting

Reporting an incident refers to informing personnel within the organization and entities outside the organization about the incident. Internally, an organization's incident response plan might require personnel to report any security incident to a member of the incident response team. As an example, technicians who identify and remove malware from systems might still be required to report the incident to security personnel, even though they resolved the incident.

Many laws and regulations require organizations to report certain types of incidents. This is especially true for any data breaches that expose personally identifiable information (PII) or protected health information (PHI). These laws typically require the organization to report the data breach to affected individuals. This helps ensure that people are informed when a data breach has exposed their personal information. They can take appropriate countermeasures such as subscribing to a credit monitoring service.

Lessons Learned

Taking the time to identify lessons learned helps ensure that an organization examines incidents in an attempt to learn from them. This is often done by holding a lessons learned meeting after an incident. It's extremely important to avoid any "blame storming" during these meetings because it puts people on the defensive and makes them reluctant to participate. Instead, the goal should be to discover things like what happened, when it happened, how the staff responded, if the existing procedures were adequate, and what could have been done differently.

As a simple example, after discovering and resolving a malware infection, security personnel will often try to determine the steps leading up to the malware infection. They may find that the malware infection occurred for one of the following reasons:

- Antivirus software was not installed or was disabled, or signatures were not up to date.
- Up-to-date antivirus software did not block the malware.
- A user brought the malware in from home with a USB drive.
- A user opened an attachment in a phishing e-mail, which installed the malware.
- A user clicked a link in a phishing e-mail, taking the user to a malicious website.
- A user visited a malicious website, which downloaded malware in a drive-by download.
- A system didn't have up-to-date patches.

By understanding what allowed the malware infection, security personnel are better prepared to recommend countermeasures to prevent a recurrence. For example, the following list identifies possible countermeasures security personnel can recommend for the malware infection:

- Use technical security controls to install and update antivirus software
- Evaluate the effectiveness of the antivirus software and consider replacing it
- Implement antivirus software at other locations, such as on e-mail servers and at the network boundary (between the Internet and the internal network)
- Use technical security controls to block access to USB drives except for those purchased by the organization
- Purchase antivirus software for users to install on their home computers
- Provide training to users on the dangers of transferring data between home and work computers via USB drives
- Provide training to users on phishing e-mail tactics
- Provide training to users on drive-by downloads
- Terminate employment of users that ignore the training
- Implement or improve a patch management program

Implementing Countermeasures

Management has the responsibility of selecting which countermeasures to implement based on the organization's acceptable level of risk. Once management approves the countermeasures, other personnel begin implementing them. Of course, it's possible that the organization will experience a similar incident even after implementing new countermeasures. This repeats the process. The result might be that the organization implements new countermeasures or replaces existing countermeasures.

Handling Evidence

During any incident, personnel are likely to encounter evidence. Ensuring that the evidence is handled properly is important. This includes preserving the evidence by preserving the scene, protecting the evidence, and creating a chain of custody to document how the evidence has been protected.

In the medical field, the first responder is the first medically trained person to arrive on the scene of an incident that requires medical knowledge. It can include emergency medical services (EMS) personnel, emergency medical technicians (EMTs), police, firefighters, or other personnel. In the IT field, the first responder is the first forensics-trained person on the scene. This may be someone with advanced forensics training. It can also be an IT administrator who understands the basics of preserving forensic evidence and can manage the scene until a forensic expert arrives.

Medical first responders aren't expected to have the same level of knowledge as a doctor. Instead, they're expected to have enough knowledge to stabilize an injured person's health until a more qualified professional arrives or they transport the injured person. In situations with multiple injuries, medical first responders perform triage to prioritize patient treatment based on the severity of their medical conditions. Similarly, IT first responders aren't expected to have the same level of knowledge as full-time forensic experts. However, they should know the basics of incident response and perform triage during an incident. They prioritize the impact on different systems during an incident and prioritize their response based on the severity of the incident for different systems.



EXAM TIP First responders are the first forensics-trained personnel on the scene of a computer incident. One of their primary responsibilities is the preservation of evidence. NIST SP 800-86, Guide to Integrating Forensic Techniques into Incident Response, discusses first responders in more depth.

Preserving the evidence is extremely important in a computer forensic investigation. If the evidence is not preserved, other personnel cannot verify the analysis. Additionally, evidence that isn't preserved is considered tainted. Tainted evidence can be misleading, indicating that

innocent people are guilty and that guilty people are innocent. Tainted evidence loses its value, and a court of law won't allow its use in legal proceedings.

Along these lines, many organizations recommend disconnecting computers from a network as an initial step in a computer investigation. This isolates the system and prevents remote access. This ensures someone cannot clear logs or delete data on the system via a remote connection. It's also an important step to take if personnel suspect the system is infected with malware, such as a worm, that may infect other systems on the network.

As soon as evidence is collected, investigators create a chain of custody. A *chain of custody* form documents who obtained the evidence, where it was obtained, when it was obtained, and how it was controlled since it was first collected. Maintaining control of evidence ensures that the evidence is never out of someone's control and subject to tampering.

Three Phases of a Computer Forensics Investigation

A computer forensics investigation is typically described as occurring in three phases of acquiring, authenticating, and analyzing the evidence, as shown in [Figure 13-1](#). IT professionals need to be aware of the three phases and understand their role if an incident occurs. For example, if IT professionals sound the alarm to report an incident, they need to know what their responsibilities are related to protecting the evidence prior to security professionals arriving to secure the evidence.

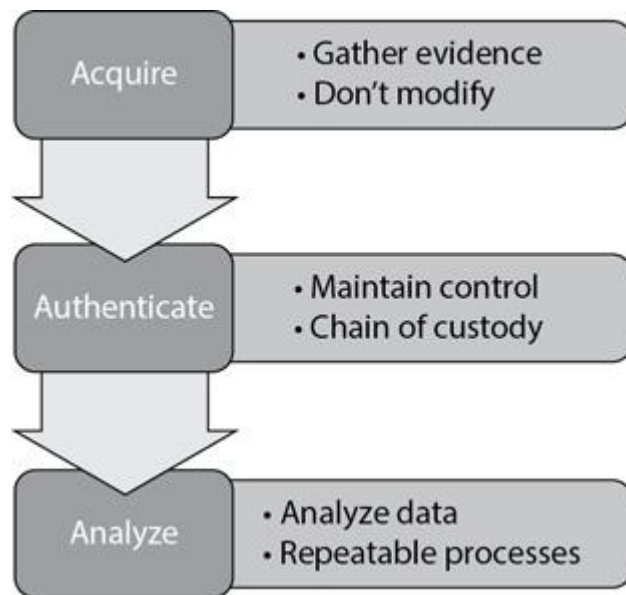


Figure 13-1 Three phases of a computer forensic investigation



EXAM TIP The goal of a computer forensics specialist is to acquire, authenticate, and analyze electronic evidence. It's extremely important that no one modifies the evidence during this process.

Acquire the Evidence

In the acquisition phase, the investigators collect the evidence. It's very easy to modify the evidence during this phase if personnel don't take steps to preserve the evidence. As a simple example, opening a file modifies it because the file now shows a new access time.

Forensic experts use bit-copy or drive-imaging tools to capture the contents of a drive without modifying the data. This copy is a forensic duplicate image. In addition to copying data, this process includes copying free and slack space on the drive or partition. The copy must be an exact bit-by-bit copy of the original in order for legal entities to consider it forensically sound. For example, [Figure 13-2](#) shows the overall process of copying a drive. Experts analyze the copy instead of the original drive. If they want to validate their analysis, they make another copy of the original drive and repeat their analysis. Additionally, the original drive remains available as evidence in a court of law if required.



Figure 13-2 Using bit copy-tools to capture drive contents

In addition to capturing the contents of a drive, it's often useful to capture the contents of a computer's memory. This often has information on recently run processes and applications and the information can be useful during the investigation. However, the computer's memory is volatile and will be lost when the computer is turned off, so it's important to keep the system on until forensic experts have an opportunity to capture the memory contents. Forensic experts have other tools available that can read and copy data from volatile memory.



EXAM TIP It's important that anyone who responds to an incident does not access any files on a system or power a system down. These steps can modify or destroy evidence that may be valuable for an investigation. Before examining hard drives, you should copy them using a bit-copy tool.

Forensic experts typically use one or more of the available forensic toolkits when collecting, authenticating, and analyzing evidence. These toolkits allow investigators to capture data more

efficiently, with less loss and less effort. [Table 13-1](#) shows some of the tools that are available. Note that this isn't intended to be a comprehensive list of all available forensic toolkits, but instead provides a list of some examples.

| Forensic Toolkit | Comments |
|---|---|
| Forensic Toolkit (FTK) | Includes a disk imaging program, scanning tools to examine disks, and multiple add-ons for additional capabilities. Developed by AccessData. |
| Computer Online Forensic Evidence Extractor (COFEE) | Captures and analyzes data from Microsoft systems. Developed by Microsoft and distributed internationally by the International Criminal Police Organization (INTERPOL). |
| The Coroner's Toolkit (TCT) | A suite of free tools used to capture and analyze data from UNIX-based systems. Developed by Dan Farmer and Wietse Venema. |
| The Sleuth Kit (TSK) | A Coroner's Toolkit replacement that captures and analyzes data from UNIX and Windows systems. Developed by Brian Carrier. |
| EnCase | A widely used suite of products sold by Guidance Software, which also provides EnCase training and certification. |
| SANS Investigative Forensic Toolkit (SIFT) | A VMware appliance that includes multiple other tools (such as the Sleuth Kit). Developed by the SANS Institute. |

Table 13-1 Common Forensic Toolkits

Comparing Volatile Data and Persistent Data

Volatile data is data that will be lost if a system is powered down. Data stored in random access memory (RAM) is an example of volatile data. RAM often holds valuable data related to recent activity on a system that can be used in an investigation. However, if the system is powered down before this data is collected, the data is lost. Several computer forensic tools are available that can capture volatile data before the system is shut down. However, there is very little possibility of retrieving information from volatile RAM after power has been removed.

In contrast, some data on a system is *persistent data*, which is data stored on a hard drive or other type of media such as an optical disc or USB flash drive. Persistent data remains even when a system is powered down.



EXAM TIP Systems that need to be investigated should not be powered down until an investigator has had an opportunity to gather data. Forensic tools can capture data in volatile RAM before the system is powered down, but the data is lost once the system is powered down.

NIST SP 800-86, *Guide to Integrating Forensic Techniques into Incident Response*, provides a wealth of details on how to collect data from different sources. This includes collecting data from data files, operating systems, network traffic, and applications.

Authenticate the Evidence

The chain-of-custody form (described in the “Handling Evidence” section) provides proof that the evidence has been controlled since it was collected and records every interaction with the evidence since it was collected. In the absence of a chain-of-custody form, courts may question the validity of the evidence and refuse to allow its use in legal proceedings.



EXAM TIP A chain-of-custody document provides proof that evidence has been controlled and protected. Without the chain-of-custody document, a court will likely question the validity of the evidence and disallow its use.

If an investigator confiscates a laptop as evidence but leaves it on his desk, anyone with access to the desk may be able to add, modify, or corrupt the data on the laptop. Even if no one tampered with the laptop while it was on the desk, the laptop may not be admissible as evidence in a court of law without a chain of custody. Even if the investigator locked the laptop in a safe, if he neglected to document his actions with a chain of custody, legal authorities are less likely to believe it was controlled at all times. Controlling the evidence and documenting the control with a chain-of-custody form provides proof that the evidence has been controlled and was not subject to tampering.

You can also authenticate files and drives by comparing a hash value of the original to a hash value of the copy. [Chapter 14](#) covers hashing in more depth, but as an introduction, a hashing algorithm creates a hash, which is simply a number. A hashing algorithm will always create the same hash when executed against the same data, including a forensically sound copy of the same data.

Imagine the hash of the original disk is 123. It should always be 123 no matter how many times you execute the hashing algorithm against the disk. If it is ever anything other than 123, it indicates the original disk has been modified. Similarly, the hash of a forensically sound copy of the disk should be 123. If it's different, it indicates that the copy is not a true bit-by-bit copy of the original disk, and courts may not accept the results from analyzing the copy as admissible evidence.

In reality, a hash is much longer than just three digits. For example, Message Digest 5 (MD5) creates a 128-bit hash expressed as a 32-digit hexadecimal number. However, for the explanation, it's much easier to conceptualize numbers such as 123 rather than numbers like this: dc647eb65e6711e155375218212b3964.

Analyze the Evidence

In the last phase of a computer investigation, the forensic experts analyze the evidence. One of the key requirements in this phase is to use repeatable processes by analyzing copies of the data whenever possible. This ensures the original data is not modified and can be copied again if necessary.

In some cases, it is necessary to analyze the original. For example, magnetic media hard drives often hold data remnants as data remanence. This data remanence is not copied in a bit-by-bit copy but might hold valuable information. The only way to analyze data remanence is by analyzing the original hard drive. This type of analysis should only be done by trained personnel, and only when necessary.

The science of computer forensics has advanced significantly in the last few years. There are many more college and university courses and degrees focused on this science, and there are also several specialized certifications available. Additionally, forensic tools continue to improve, giving forensic experts more and more capabilities.

Legal and Ethical Principles

When responding to incidents, it's important to apply legal and ethical principles. Applying legal principles simply means identifying relevant laws and abiding by them. These laws vary in different jurisdictions, such as in states and different countries, but organizations have an obligation to identify and follow them.

When an organization decides to add forensic capabilities, it's important to document some basic legal and ethical principles along with guidelines within a forensics security policy. These principles and guidelines can be anything the organization chooses to add, and many professional organizations have provided suggestions of items to include.

The Association of Chief Police Officers (ACPO), a British organization recently replaced by the National Police Chiefs' Council (NPCC), published the *ACPO Good Practice Guide for Digital Evidence*. Version 5.0 includes the following four legal and ethical principles:

- **Principle 1** No action taken by law enforcement agencies, persons employed within those agencies or their agents should change data which may subsequently be relied upon in court.
- **Principle 2** In circumstances where a person finds it necessary to access original data, that person must be competent to do so and be able to give evidence explaining the relevance and the implications of their actions.
- **Principle 3** An audit trail or other record of all processes applied to digital evidence should be created and preserved. An independent third party should be able to examine those processes and achieve the same result.

- **Principle 4** The person in charge of the investigation has overall responsibility for ensuring that the law and these principles are adhered to.

NIST SP 800-86, *Guide to Integrating Forensic Techniques into Incident Response*, stresses the importance of clearly defining the roles and responsibilities of personnel performing forensic activities. It mentions that forensic processes should be consistent with the organization's policies and all applicable laws. It also stresses the importance of maintaining a chain of custody to support the admissibility of evidence into legal proceedings.

NIST SP 800-72, *Guidelines on PDA Forensics*, suggests using the Daubert method, which is a set of standards that help ensure the reliability of evidence. When applied to digital evidence, it's important to consider the following reliability factors derived from the Daubert method and set forth in SP 800-72. Following these reliability factors will help ensure the evidence can be used in a court of law.

- **Testability** Has the scientific theory or technique been empirically tested? (Some tools, such as the EnCase Toolkit, have been used successfully in court cases, but some have not.)
- **Acceptance** Has the scientific theory or technique been subjected to peer review and publication? This ensures that flaws in the methodology would have been detected and that the technique is finding its way into use via the literature.
- **Error Rate** What is the known or potential error rate? Scientific measures generally have associated error rates, which can be estimated with a fair amount of precision.
- **Credibility** What is the expert's qualifications and stature in the scientific community? Does the technique rely upon the special skills and equipment of one expert, or can it be replicated by other experts elsewhere?
- **Clarity** Can the technique and its results be explained with sufficient clarity and simplicity so that the court and the jury can understand its plain meaning?



TIP It's rarely clear at the beginning of an investigation that evidence will be needed in a court of law. Because of this, personnel often overlook important evidence, handle it improperly, or accidentally destroy it early in an investigation. However, by always following proper forensic practices, it ensures the protection of evidence in case it is ever needed.

Forensics and Fraud

While most scientists and forensic experts seek the truth, some have strayed from the law. This includes falsifying evidence, falsifying lab results, and misrepresenting evidence in trials.

As an example, Fred Zain was a forensic laboratory technician who falsified DNA results. Before this was discovered, he was regularly sought out as an expert witness, though his credentials were sketchy at best. He reportedly failed organic chemistry, flunked an FBI course on forensic science, and fabricated his resume.

Investigations later indicated that he falsified evidence resulting in as many as 180 wrongful convictions. Senior Circuit Court Judge James Holiday reported that Zain engaged in misconduct and outright fraud so often, that any testimony by him should be considered invalid, unreliable, and inadmissible. Another investigation indicated that Zain often filed reports on tests that weren't completed, reported negative results as positive, and described inconclusive results as conclusive.

Even though there were many hints that Zain was engaged in fraud during his career, supervisors, and other experts chose to ignore them. For example, two workers reported seeing him record test results from blank test results, but their complaints weren't taken seriously. Additionally, an FBI laboratory director told state police that Zain failed basic forensic science courses, but no action was taken.

There aren't any infamous examples of cybersecurity forensic experts engaging in the same level of fraud. However, it is certainly conceivable that some are tempted to falsify evidence in a misguided sense of justice. Hopefully, when others see these actions, they will report them and the actions will be investigated.

[Chapter 1](#) includes the four Code of Ethics Canons that SSCP candidates should follow. As a reminder, they are

- Protect society, the common good, necessary public trust and confidence, and the infrastructure.
- Act honorably, honestly, justly, responsibly, and legally.
- Provide diligent and competent service to principals.
- Advance and protect the profession.

Comparing Computer Abuse and Computer Crime

Computer abuse and computer crime mean two separate things. In computer abuse, an individual is dishonest or unethical and using a computer in an unauthorized manner. In contrast, a criminal commits illegal acts with a computer. Both of these concepts are highlighted in [Figure 13-3](#).

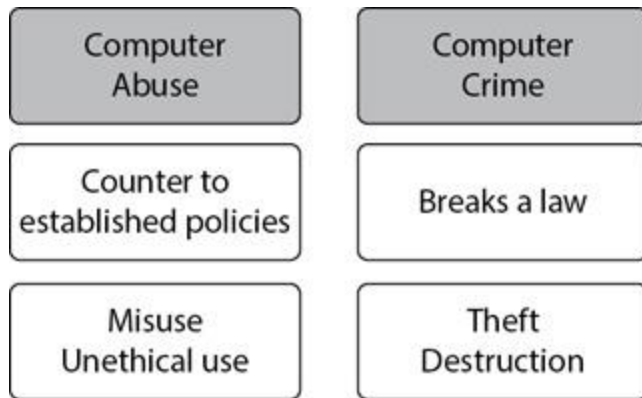


Figure 13-3 Computer abuse and computer crime

For example, in computer abuse, an individual may misuse a computer system in violation of an organization's security policies. The user may be using the computer to shop for personal items during company time despite knowing that this is against the established acceptable use policy.

Cybercrime (or computer crime) is any crime that involves a computer. The criminal may use a computer as a tool to commit the crime, or the computer may be the target of the crime. The purpose of cybercrime is often monetary gain for the criminal, but some criminal acts are simply destructive, with the purpose of negatively affecting a computer or organization.



EXAM TIP Computer abuse may result in the loss of a job or position, but does not actually involve breaking laws. Computer crime includes illegal activities and can result in prosecution, fines, and/or jail time.

The three commonly quoted elements needed to commit a crime are means, motive, and opportunity (sometimes referred to as means, opportunity, and motive, or MOM). By understanding these, you have a greater chance at reducing crime. More specifically, if you can remove the means, motive, or opportunity to commit a crime, you have a better chance of preventing the crime. The means is the ability of the criminal to commit the crime, the motive is the reason (such as for money or revenge), and the opportunity is simply the chance to commit the crime.

The following are some of the common terms related to computer abuse and computer crimes:

- **Cyberbullying** Cyberbullying has become common among some young people. The bully uses the Internet to attack a victim through harassment, coercion, or intimidation. For example, a cyberbully may spread lies about a person via social media such as Facebook or Twitter, or send harassing text messages to the victim or to others about the victim.

- **Cyberstalking** Cyberstalking includes cyberbullying actions, but it goes further. For example, the stalker may gather information on the victim, monitor the victim's activities, and steal the victim's identity. Adults are typically involved with cyberstalking, while children are more likely to be involved in cyberbullying. A number of laws are in place addressing cyberstalking making it a criminal act.

- **Cyberwarfare** Cyberwarfare is a politically motivated attack on entities in another country. The goal is typically sabotage to disrupt the capabilities of an entity in the target country or espionage to gather intelligence on the target. For example, the U.S. government concluded that North Korea attacked Sony in late 2014 and labeled this cyberwarfare. Similarly, the *Washington Post* reported that attackers accessed computer networks in the White House and that the *Post's* sources indicated the attackers were Russian state-backed groups. Attack groups sponsored by a state (or government) are typically referred to as advanced persistent threats (APTs).

Understanding Fraud and Embezzlement Crime

Fraud is the use of deception for unlawful gain or unjust advantage. It includes engaging in wrongful or illegal activities with the intention of gaining undue or illegitimate advantage over another person, group, or entity. Embezzlement is a specific type of financial fraud where an individual steals money or property from an employer, company, clients, customers, or even partners. It also includes misappropriating money or assets held in trust.

MCI WorldCom gained notoriety for a massive accounting scandal that resulted from fraudulent accounting methods. CEO Bernard Ebbers helped grow the business that ultimately became MCI WorldCom. Later, courts held him responsible for the company's fraudulent financial reporting that resulted in the loss of over \$100 billion to investors (including 401K investors), and he was ultimately convicted of fraud and conspiracy. He is serving a 25-year prison sentence.



TIP [Chapter 11](#) mentions the Sarbanes-Oxley (SOX) Act of 2002. It requires CEOs and other executives to personally verify the accuracy of financial data. Part of the goal of the legislation is to prevent fraudulent financial reporting similar to what occurred with MCI WorldCom.

Bernie Madoff has become a household name related to fraud. He managed a Ponzi scheme for years under the guise of an investment company. Many people regarded him as an outstanding investment banker who ran a successful wealth management company. In reality, he routinely skimmed new investments and used these investments to pay what looked like great dividends to current investors. Investigators calculated the amount of missing funds at about \$65 billion. He is currently serving a 150-year prison term. Interestingly, investigations such as one by the U.S. Securities and Exchange Commission failed to uncover the fraud, which went undetected until Madoff admitted the Ponzi scheme to his sons in 2008.

[Chapter 5](#) mentions salami attacks where an attacker takes multiple minor actions to get large gains. For example, prosecutors indicted Michael Largent for collecting more than \$50,000 from some online brokerage firms by collecting between a penny and \$2 at a time. He used scripts to open over 58,000 accounts under phony names with the brokerage firms. These firms verified his bank account by depositing small amounts of money into his account. However, when his bank questioned the thousands of small payments, he didn't have any good answers and the FBI later arrested him. Largent was later sentenced to 15 months in prison and ordered to pay over \$200,000 in restitution.

Many companies use job rotation and mandatory vacations to reduce the risks of internal fraud and embezzlement. These policies also reduce illegal data modifications and general misuse of information.

Mandatory Vacations

Many banking institutions implement mandatory vacation policies. These require employees to take a vacation outside of the workplace for a specified period, such as at least five or ten consecutive workdays. The goal is to require someone else to perform their job functions, which increases the possibility of exposing any suspicious activities.

For example, imagine that an employee identifies a loophole in the flow of money within an organization and discovers that he can divert a small percentage of money in each transaction. However, if he ever stopped diverting the money, the loss would appear as an anomaly. If the organization forces him to take time off work, the diversion would stop for a short time, and someone would discover the fraud. Knowing this is a possibility will often deter an employee from exploiting the loophole in the first place.

Similarly, imagine if an investment banker takes millions of dollars from specific investors but never actually invests the money. Instead, he uses the money to pay healthy dividends to other investors who previously gave him money. He gives these investors fake documents indicating that they have a substantial amount of wealth. In reality, this crooked investment banker has stolen all the money. As long as the investment banker is at work, he can answer the queries from the investors and no one will discover his actions. However, if the organization forces him to take regular vacations, the Ponzi scheme will collapse while he is away.

Admittedly, a mandatory vacation policy doesn't guarantee that fraud won't occur. Consider this same scenario with two or more people colluding to create and manage this Ponzi scheme. If that's the case, forced vacations won't necessarily discover the actions. The accomplices can cover for anyone who is on vacation. However, when two or more people are involved, it increases the chances that someone will make a mistake. Additionally, law-breaking criminals committing fraud aren't necessarily trustworthy, so any one of them might report the actions at some point and place all the blame on one of the other accomplices.

While mandatory vacations can be an effective security control, they aren't foolproof. However, when used as part of an overall defense-in-depth plan, mandatory vacations can help reduce the chance of fraud and embezzlement.



EXAM TIP A mandatory vacation policy requires employees to take vacations for a minimum period (such as five or ten consecutive days). The goal of the policy is to reduce the chance of fraud.

Job Rotation

Job rotation is the practice of moving employees periodically, such as once a year or so, between different jobs. The goal is for employees to do different jobs at different times, which also prevents any single employee from doing the same job for years at a time.

[Chapter 1](#) covered separation of duties and explained how it prevents a single person from having complete control over a process. For example, accounting departments frequently separate accounts payable tasks from accounts receivable tasks to reduce the risks. Even though a separation of duties policy separates these tasks, employees working in an accounting department typically have the skills to do both jobs. In other words, someone working in an accounts payable division also has the skills to perform accounts receivable tasks. A job rotation policy periodically moves personnel between the two jobs. Just in case an employee discovers a loophole in the process that he or she can use to defraud the company, the job rotation policy limits how long the fraud could continue. It also increases the chance of discovering the fraud, which acts as a deterrent.

Many organizations use job rotation for cross-training. For example, organizations often rotate management personnel through different jobs within an organization to gain a better overall understanding of the organization. The primary purpose of this practice is to help build the knowledge and skills of personnel rather than deter and detect fraud. Job rotation also helps in disaster recovery if key personnel are unavailable.



EXAM TIP Job rotation used as a security practice can reduce the risk of fraud within a company. It is commonly combined with the security practice of separation of duties.

Understanding Privacy Issues

Although the explosion of the use of computers has certainly helped people share information, it has also resulted in people having easy access to massive amounts of personal data about individuals. If organizations don't protect this data, it can negatively affect an individual's privacy. Even if individuals are very careful with their personal data and reveal it only when required to do so for business transactions, they are still affected if the business suffers a data breach. If these businesses don't protect the customer's data, the customer loses.

Because of this, many laws require organizations to protect personal data. Despite these laws, data breaches continue. [Chapter 5](#) discussed the Identity Theft Resource Center (ITRC). As a reminder, the ITRC reported a record high of 1,579 data breaches exposing more than 178 million records in 2017.

Legal experts agree that the key to avoiding liability for any data leakage is due diligence. In other words, if an organization is actively taking steps to protect data and prevent privacy compromises, but still suffers a loss, it is less likely to be held liable. The message seems to be that “incidents happen,” but if you’re doing your best to prevent the incidents, you’ll have less legal liability.

[Chapter 11](#) introduced personally identifiable information (PII) and the importance of protecting it. PII is information that can be used to personally identify an individual and includes components such as those shown in [Figure 13-4](#). [Chapter 11](#) also introduced the Health Insurance Portability and Accountability Act (HIPAA) of 1996. HIPAA mandates the protection of protected health information (PHI), which is any PII that includes information concerning the health status, provision of health care, or payment of health care for an individual.

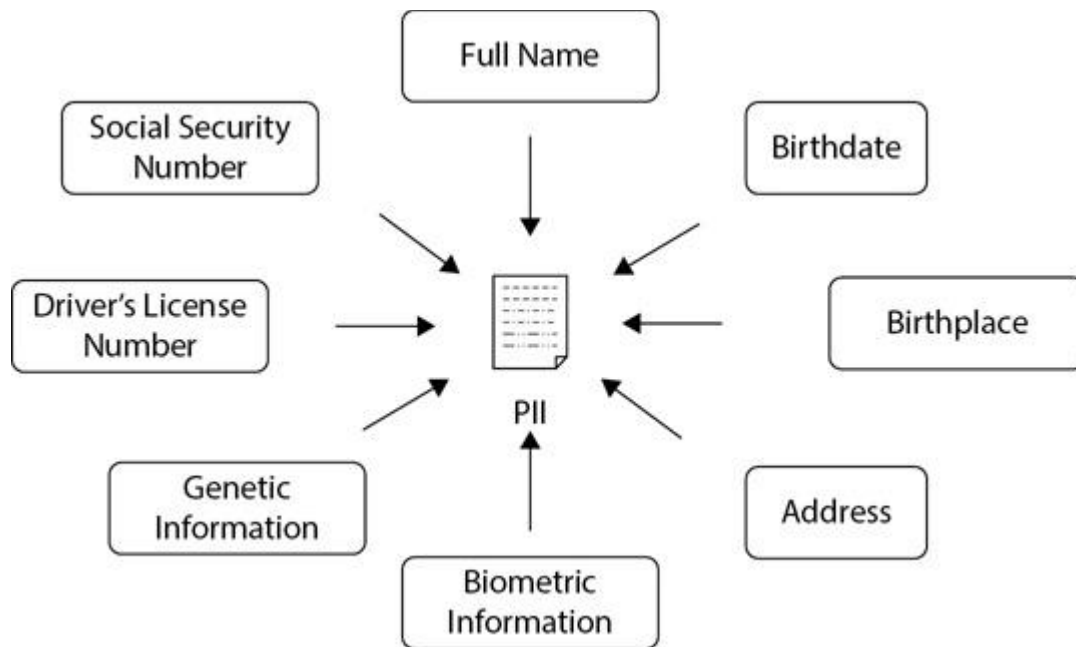


Figure 13-4 Components of PII

Defining a Data Breach

Generically, a data breach is any incident where unauthorized personnel gain access to sensitive, protected, or confidential data. It often occurs when criminals from outside an organization launch attacks allowing them to access data held inside an organization.

The Identity Theft Resource Center focuses on the exposure of personal data and defines a data breach as “an incident in which an individual name plus a Social Security number, driver’s

license number, medical record or financial record (credit/debit cards included) is potentially put at risk because of exposure. This exposure can occur either electronically or in paper format.” Outside of the United States, this definition of a data breach includes a national identification number instead of a Social Security number. Note that this matches the definition of PII. Because medical records are included, it also matches the definition for PHI.

Some data breaches only expose usernames, passwords, and e-mail addresses without exposing other information that identifies an individual. The ITRC doesn’t report the number of records exposing only this data, primarily because its focus is on identity theft. However, many other definitions include this data compromise as a data breach. Attackers can do a lot of damage if they gain access to valid user-names, passwords, and e-mail addresses.



TIP You don’t need to have all of the components shown in [Figure 13-4](#) to have PII. If any single piece of information combined with a name can uniquely identify an individual, it is PII.

The following sections cover some of the regulations that mandate the protection of PII.

General Data Protection Regulation

[Chapter 11](#) introduced the General Data Protection Regulation (GDPR). As a reminder, it is a European Union (EU) regulation that became enforceable in May 2018 and supersedes the Data Protection Directive (Directive 95/46/EC). The overall purpose is to protect the personal data of individuals within the EU.

Organizations that operate in the EU and/or collect data from any EU residents are required to comply with the GDPR. It requires organizations to obtain freely given, specific, informed, and unambiguous consent before collecting personal data. The European Commission defines personal data as “any information relating to an individual, whether it relates to his or her private, professional or public life. It can be anything from a name, a home address, a photo, an e-mail address, bank details, posts on social networking websites, medical information, or a computer’s IP address.”

Three relevant terms within the GDPR are data subject, data controller, and data processor.

- **Data subject** The GDPR defines a data subject as “an identified or identifiable natural person.” In short, a data subject is a person. However, the GDPR does not define *natural person*. The three likely possibilities are 1) a person who is a citizen of the EU, 2) someone who is living in the EU, or 3) someone who is a citizen of the EU and living in the EU. This is likely to play out in the courts in the coming years. The GDPR does specify that parental consent is required to collect data on any child below the age of 16 years.

- **Data controller** A data controller is an entity that collects data about EU residents. Data controllers must comply with certain rules related to how they collect and use personal information.

- **Data processor** A data processor is an entity that processes data on behalf of a data controller but doesn't have control over the data. As an example, an external payroll company that processes data for another organization's employees is a data processor. Cloud providers that store data are also data processors. Data processors have a responsibility to keep data secure by preventing unauthorized access, disclosure, or accidental loss.

It is possible for an entity to be both a data controller and a data processor.

The penalties for noncompliance are fines up to 4 percent of an organization's annual global revenue.

California Supreme Court Rules That ZIP Codes Are PII

The California Supreme Court determined in early 2011 that a person's Zone Improvement Plan (ZIP) code is PII. In this case, Williams-Sonoma stores collected a customer's ZIP code each time that the customer made a credit card purchase. It then combined the customer's name and ZIP code in a search to discover the customer's address (which the customer didn't provide) and used this information for marketing purposes.

The court determined that a ZIP code is not required for a credit card transaction, but is readily understood to be part of an address. Because a company can combine a ZIP code with the person's name to discover the person's address, the California Supreme Court classified it as PII. The decision restricts businesses in California from requesting a ZIP code during a credit card transaction. The California Supreme Court's justification in this case is that the business can use this personal information for business purposes. A similar lawsuit is playing out in Massachusetts.

While the California Supreme Court determined that a ZIP code is PII when combined with a customer name, it is one of the few jurisdictions to do so. If you buy gas with a credit card in a different jurisdiction, you are often required to enter the ZIP code.



EXAM TIP Data inference (mentioned in [Chapter 11](#)) is the ability to piece together knowledge with seemingly unconnected information. In this case, the company was using a data inference technique to get a customer's address from online databases.

Connecticut's Public Act No. 08-167

In 2008, Connecticut passed a law titled “An Act Concerning the Confidentiality of Social Security Numbers.” Although the title focuses on Social Security numbers, the contents of this law go beyond just these numbers.

For example, it states, “Any person in possession of personal information of another person shall safeguard the data, computer files, and documents containing the information from misuse by third parties....” It then defines personal information as “information capable of being associated with a particular individual through one or more identifiers, including, but not limited to, a Social Security number, a driver’s license number, a state identification card number, an account number, a credit or debit card number, a passport number, an alien registration number or a health insurance identification number....”

Intentional violations of this law can result in a fine of \$500 for each violation, with a maximum fine of \$500,000 per event. In other words, if a database is released that includes this information on 1,000 or more people, it can result in a fine of half a million dollars.

A notable distinction with this law is that unintentional violations are not subject to the law. In this case, if an organization discloses the information, it must prove that the disclosure wasn’t intentional. The law does not consider a data breach by an attacker an intentional violation as long as the organization took steps to protect the information.

Other laws, such as the Social Security Number Protection Act of 2010 and the Social Security Number Protection Act of 2011, provide additional protection of Social Security numbers. The 2010 act prohibits federal, state, and local agencies from printing any part of a Social Security number on a check; some organizations did so to ensure only the intended recipient could cash it. The 2011 act requires the Department of Health and Human Services to eliminate the unnecessary collection, use, and display of Social Security numbers on Medicare cards and other communications.

Children’s Online Privacy Protection Act

The Children’s Online Privacy Protection Act (COPPA) is a U.S. federal law that applies to the collection of information on children under the age of 13. The goal is to protect the privacy of these children by preventing websites or online services from collecting information on them without parental consent.

For example, UMG Recordings was fined \$400,000 after it was found to have knowingly collected information on children under the age of 13. The company created a website directed at children to promote a 13-year-old pop star known as Lil’ Romeo. UMG Recordings collected extensive information on the children, including their full name, e-mail address, home address, phone number, interests, and birthdate. Even when the user’s birthdate clearly indicated that the user was under the age of 13, UMG Recordings still collected and used the information. For example, it used the collected e-mail addresses to send e-mail marketing about other musicians and other websites.

Although it is possible to collect information on children under 13, the organization must gain parental consent, which is often challenging. Because of this, most organizations simply do not collect information on these children to remain in compliance of the law.

California Online Privacy Protection Act of 2003

The California Online Privacy Protection Act (OPPA) of 2003 became law in 2004. Notice the acronym is OPPA without an initial *C* to differentiate it from the Children's Online Privacy Protection Act (COPPA). It requires operators of commercial websites to conspicuously post and comply with a privacy policy if they collect PII on any visitors who reside in California. In this context, a commercial website is one created by or for an organization. In other words, OPPA doesn't apply to personal websites created by or for individuals.

Note that this law doesn't just apply to websites created by organizations in California or hosted on servers in California. Instead, it applies to any organization that makes its website available to any California residents. In other words, it applies to any website on the Internet, because California residents with Internet access can access pages hosted anywhere in the world.

Although it's doubtful that California law will extend to remote places like Tanzania, it will certainly affect websites operating within the United States. Because of this, most operators of commercial websites in the United States make certain that they include a privacy policy that complies with this law if they collect any information on the user, such as name, e-mail address, or other data, used to track the user.

The privacy policy must be easy to locate (for example, websites typically make it accessible from the home page) and include the following:

- What information is gathered by the website
- How the information is used and shared
- The effective date of the policy
- A description of changes since the policy was first created

Legal and Regulatory Concerns

The previous sections covered some specific laws related to privacy. [Chapter 4](#) covers cloud computing, including various cloud server models and cloud security, and [Chapter 8](#) covers various monitoring systems, including intrusion detection and prevention systems. It's worthwhile to connect some of these topics within the concept of legal and regulatory concerns.

Jurisdiction refers to the legal authority that a legal body has over an organization. If an organization is operating in Maryland, for example, it must comply both with the laws in Maryland and with U.S. laws. However, the jurisdiction can easily extend to other legal entities if the organization conducts interstate or international business.

However, if an organization chooses to store data in the cloud, the jurisdiction for this data might be different. Imagine that some of an organization's data is stored in a data center in Virginia and some of its data is backed up in a data center in Canada. The organization now must comply with laws in Virginia and Canada. Additionally, the GDPR provides another example of how jurisdiction can change. Simply by collecting data from an EU resident, an organization is obligated to comply with the GDPR regardless of the organization's physical location.

Cloud-based surveillance solutions are becoming quite popular, but they bring other concerns. As an example, organizations can install surveillance cameras that capture video and store the video in the cloud. This can change the jurisdiction and, depending on contracts with the surveillance company, can also raise questions about data ownership.

Electronic discovery (also known as *e-discovery*) refers to any process that attempts to locate and secure electronic data for a civil or a criminal legal case. Legal entities often request such data with a subpoena and organizations have a legal obligation to comply. However, collecting this data and turning it over to legal authorities takes time and money.

Organizations can save time and money in potential future events by limiting the amount of data they maintain. As an example, imagine an organization has a policy to delete all e-mail older than one year, and that policy is followed. Later, a lawsuit results in a subpoena that requires the organization to turn over all its e-mail related to a specific topic. Personnel would then be required to go through all the e-mail for the past 12 months to find relevant e-mails. In contrast, imagine the organization doesn't have a policy and the e-mail servers have e-mails from the last ten years. Personnel would then be required to comb through all ten years' worth of e-mails.

Monitoring systems collect data within an organization and must also comply with the same jurisdiction issues. However, they can also help an organization prove compliance with laws and regulations. Still, there are limitations on which data monitoring systems can collect. Organizations need to ensure they do not violate privacy laws within their jurisdiction when collecting data.

Chapter Review

A security incident is any violation of policies or security practices that has the potential to result in an adverse event. An incident lifecycle includes several different phases. The first phase is preparation, which includes implementing security controls to prevent, detect, and correct security incidents. When an event occurs, personnel investigate it in the detection, analysis, and escalation phase. If analysis verifies that an incident has occurred, personnel may escalate it depending on the severity of the incident. During the containment phase, personnel attempt to isolate the incident to limit potential damage. The next phase is to eradicate all elements of an incident. This phase can be combined with the recovery phase to return all systems to normal operations. After an incident has occurred, it's important to examine the responses in a lessons learned phase. In many situations, this results in a recommendation to implement additional countermeasures.

First responders are the first forensics-trained personnel on the scene, and one of their primary responsibilities is the preservation of the scene. The three phases of a typical computer forensic investigation are acquiring the evidence, authenticating the evidence, and analyzing the evidence.

When acquiring the evidence, one of the most important concerns is to protect the data and prevent data modifications. Forensic tools allow experts to capture data in volatile RAM before a system is powered down and use bit-copying procedures to copy the contents of hard drives. Several forensic tools are available such as FTK, COFEE, the Sleuth Kit, EnCase, and more.

In the authentication stage, a chain-of-custody document is used to provide proof that the evidence was controlled after it was collected. The chain-of-custody document ensures that the evidence is admissible in a court of law. The analysis of the evidence is the last stage. Some of the same forensic tool suites used to collect evidence include tools that can be used to analyze evidence.

Computer abuse indicates that an individual is using a computer or network in an unauthorized manner, such as counter to an organization's security policies. Computer crime indicates that the user is committing illegal acts with the computer. Computer crime is punishable by fines and/or jail time.

Computers can be used in fraud and embezzlement crimes resulting in substantial losses to an organization. Mandatory vacations and job rotation are two security policies that an organization can use to reduce the success of fraud and embezzlement crimes. These two policies are often used with other security controls as part of an overall defense-in-depth strategy.

PII includes any information that can identify an individual. Privacy issues related to the protection of PII have become more important as the amount of data that organizations collect on their customers has increased. If an organization doesn't take precautions to protect customer data, attackers may breach the organization's database and access its customer data. Many laws have been enacted on the international, national, and state levels requiring organizations to protect PII.

Legal and regulatory concerns apply to cloud computing and various monitoring systems. They typically identify requirements to maintain the privacy of data. Jurisdiction refers to the legal authority that a legal body has over an organization. It can change depending on the physical location of an organization, where the data is held within the cloud, and where the data was collected from, such as from an EU resident. Electronic discovery refers to any process that attempts to locate and secure electronic data for a civil or a criminal legal case.

Questions

1. What is the first phase in the incident lifecycle?

A. Event

B. Preparation

C. Detection

D. Escalation

2. Which of the following *best* describes a primary goal of incident handling?

A. Collecting evidence

B. Documenting evidence

C. Containing any potential damage

D. Improving security controls

3. Which one of the following is most likely to be performed during a lessons learned phase in the incident lifecycle?

A. Chain of custody

B. Hashing

C. Escalation

D. Countermeasure recommendations

4. Which of the following identifies a primary responsibility of a first responder after a computer security incident?

A. Capturing images of disks

B. Capturing data in RAM

C. Interviewing witnesses

D. Preserving the scene

5. What forensic evidence can be lost if a system is powered down before the evidence is collected?

A. Data on the disk drive

B. Data on a USB drive

C. Data in memory

D. Data in files

6. Managers suspect that an employee in your organization has committed fraud. You are told to secure his computer as part of an incident response. Which of the following are appropriate steps? (Select three.)

- A.** Disconnect the computer from the LAN
- B.** Power the system down
- C.** Prevent anyone from accessing the system
- D.** Take pictures of the system

7. Which of the following is a benefit of a chain-of-custody form?

- A.** It helps ensure that evidence is protected.
- B.** It helps ensure that evidence is controlled.
- C.** It helps ensure that evidence is not modified.
- D.** It helps ensure that evidence is admissible in court.

8. Of the following choices, what are valid phases in a computer forensic investigation? (Select three.)

- A.** Prosecution based on evidence
- B.** Authenticating evidence
- C.** Analyzing evidence
- D.** Acquiring evidence

9. Which of the following choices is the most important consideration when gathering evidence as part of a computer forensic investigation?

- A.** Ensuring that systems are turned off as soon as possible
- B.** Ensuring that a record of files on a system is recorded by accessing the system
- C.** Ensuring that users can log on to the system
- D.** Ensuring that evidence is not modified

10. A forensic expert wants to examine data on a hard drive of a confiscated computer. Which of the following actions should the expert complete first?

- A. Ensure that the computer has UPS protection
- B. Create a bit copy of the disk
- C. Disable the antivirus software on the computer
- D. Move the hard drive to another system and examine it on the other system

11. Which of the following helps to prove that collected evidence has been controlled since it was collected?

- A. COFEE application
- B. Chain of custody
- C. SIFT
- D. Audit logs

12. In general, what elements need to come together for a crime?

- A. Means, motive, and opportunity
- B. Criminal, software, and hardware
- C. Discovery, theft, and benefit
- D. Attacker, attackee, and method

13. Countries sometimes engage in espionage against other countries. What is this called?

- A. Cyberbullying
- B. Cyberstalking
- C. Cyberwarfare
- D. Cyberterrorism

14. What is the purpose of mandatory vacations in relation to security?

- A. To ensure that employees do not burn out
- B. To ensure that employees take time to relax
- C. To reduce the payroll of an organization

D. To reduce the chance of fraud

15. What security practice moves employees into different positions periodically to reduce the risk of fraud?

A. Separation of duties

B. Job rotation

C. Mandatory vacations

D. Risk mitigation

16. A business in Florida gathers customers' names and ZIP codes and uses them to identify the customers' addresses. What is occurring?

A. Violation of the GDPR

B. Data breach

C. Data inference

D. Violation of COPPA

17. An organization collects customer data such as their name, e-mail address, physical address, and phone number. What term best describes this information?

A. PHI

B. PII

C. COFEE

D. FTK

18. An organization regularly collects information on customers for marketing purposes. It uses this information to personally identify the customers. Who is responsible for protecting this data?

A. The organization.

B. It depends on whether the customers gave permission to collect the data.

C. It depends on whether a data breach occurred.

D. Customers.

19. What law requires an organization to get a parent's consent prior to collecting information on children under 13?

- A. COPPA
- B. OPPA
- C. GDPR
- D. E-Privacy Directive

20. What law requires organizations to post a privacy policy on their website?

- A. SOX
- B. PHI
- C. OPPA
- D. COPPA

Answers

1. B. The first phase in the incident lifecycle is preparation, which includes implementing security controls to prevent, detect, and correct security incidents. An event isn't necessarily an incident. Events are investigated to determine if they are incidents during the detection, analysis, and escalation phase.

2. C. Primary goals of incident handling include containing any potential damage and repairing the damage that has occurred. While incident handling often includes evidence collection and using chain-of-custody forms to document collected evidence, these are not primary goals. The lessons learned phase often improves security controls, but this is not a primary goal.

3. D. The lessons learned phase examines the incident and the response, and often results in countermeasure recommendations. The goal is to prevent a recurrence of the incident. A chain of custody validates that evidence has been controlled since it was collected. Hashing provides assurances that data (such as a bit-by-bit copy of a disk) is identical. Escalation refers to getting more people involved in an incident.

4. D. A primary responsibility of first responders is the preservation of the scene. Forensic experts acquire and analyze the evidence, which includes capturing images of disks and data in RAM. Investigators interview witnesses.

5. C. Data in memory (volatile RAM) is lost if a system is powered down. Power isn't required to retain data on disk drives or USB drives, so it is not lost when powered down. Data in files is retained on disk drives.

6. A, C, D. Organizations often have policies to disconnect a computer from the local area network (LAN) to isolate it and protect evidence on the system. No one should be allowed to access the system as a precaution to ensure that data (and potential evidence) is not modified. In some situations, taking pictures of the system can help preserve evidence. You should not power the system down, as doing so will delete any data in volatile RAM.

7. D. A chain-of-custody form provides proof that evidence has been protected and helps ensure that the evidence is admissible in court. Note that it only documents how the evidence has been protected and controlled, but it doesn't actually protect or control the evidence. It also doesn't ensure that the evidence is not modified.

8. B, C, D. A computer forensic investigation typically includes the three phases of acquiring, authenticating, and analyzing evidence. Prosecution is not a part of a computer forensic investigation but may occur based on the results of an investigation.

9. D. It's important that evidence is not modified when collecting evidence. If a system is turned off, data in volatile RAM is lost. If files on a system are accessed, it modifies the files by showing new access times. If a user logs on to a system, it can modify information about the last user that logged on.

10. B. Disk drives should be copied with an approved bit-copy tool prior to examining any of the data because the process of examining the data can result in the loss of data. An uninterruptible power supply (UPS) protects a system from a power loss, and antivirus (AV) software protects a system from malware, but neither has anything to do with protecting evidence. Moving the drive to another system and examining it on the other system will modify the original evidence.

11. B. A chain-of-custody document provides proof that collected evidence has been controlled since it was collected by identifying exactly who has control of it and where it was located since it was collected. Microsoft developed the Computer Online Forensic Evidence Extractor (COFEE) to help law enforcement forensic experts. SANS Investigative Forensic Toolkit (SIFT) is another set of applications that computer forensic experts use. Audit logs can be used in a forensic investigation but they don't validate evidence.

12. A. The three commonly quoted elements for a crime are means (the ability to commit the crime), motive (such as money or revenge), and opportunity (the chance to commit the crime). Crimes can be committed without software or hardware. Criminals can commit crimes (such as vandalism or destruction) without theft. Similarly, criminals can commit crimes (such as theft) without an attack.

13. C. Cyberwarfare is a politically motivated attack on entities in another country and is done for sabotage and/or espionage. Cyberbullying occurs when one person harasses, coerces, or intimidates another person using the Internet. Cyberstalking is more serious than cyberbullying and is a criminal act. Cyberterrorism is the use of the Internet to launch terrorist attacks.

14. D. A mandatory vacation policy can reduce the chance of fraud by requiring other employees to take over the tasks and responsibilities of a vacationing employee. While vacations are good to help employees relax and reduce the chance of burnout, these matters aren't as relevant to security issues as reducing the chance of fraud. While some companies do implement mandatory vacations to reduce payroll expenses during periods of low activity, this is not a security policy.

15. B. Job rotation is a security practice of rotating employees into different positions periodically to reduce the risk of fraud. It is often combined with a separation of duties policy, but separation of duties by itself does not require employees to rotate between jobs; only that a single employee doesn't control all elements of a process. A mandatory vacation security policy requires employees to take time away from work to increase the chances of suspicious activity being discovered. Risk mitigation is any practice that reduces risks, but it is not limited to only rotating employees to different positions or only reducing the risk of fraud.

16. C. This is an example of data inference, because the company is collecting small pieces of information to get other information that isn't provided directly. The General Data Protection Regulation (GDPR) doesn't restrict the collection of data, but instead mandates the protection of data that a company collects. The question doesn't indicate the company is not protecting the data adequately. A data breach occurs when unauthorized individuals access stored data, but there's no indication of unauthorized access in this scenario. The Children's Online Privacy Protection Act (COPPA) protects the privacy of children under 13.

17. B. This information is personally identifiable information (PII) because it can be used to identify the customers personally. Protected health information (PHI) is information concerning the health status, provision of health care, or payment of health care for an individual. COFEE is a forensic tool used by law enforcement agencies. The Forensic Toolkit (FTK) is another forensic tool used by cybersecurity professionals.

18. A. Organizations are responsible for protecting personally identifiable information (PII), which is information that can identify individuals. This is true even if the customers give permission to the organization to collect the data, and even if a data breach has not occurred. Customers are not required to protect the information that they give to an organization.

19. A. The Children's Online Privacy Protection Act (COPPA) requires organizations to get a parent's consent prior to collecting information on children under 13. The California Online Privacy Protection Act of 2003 (OPPA) requires operators of commercial websites to post a privacy policy on the website if the website collects personally identifiable information (PII). The General Data Protection Regulation (GDPR) requires the protection of European Union (EU) residents and requires parental consent if a child is under the age of 16. The GDPR supersedes the E-Privacy Directive.

20. C. The California Online Privacy Protection Act of 2003 (OPPA) requires operators of commercial websites to post a privacy policy on the website and comply with it if the website collects personally identifiable information (PII). Sarbanes-Oxley (SOX) Act is a U.S. law that requires executives of publicly held companies to validate the integrity of their financial data.

HIPAA is a law focused on the protection of protected health information (PHI) and requires organizations to post a privacy policy. However, PHI is not a law. The Children's Online Privacy Protection Act (COPPA) is a U.S. federal law that applies to the collection of information on children under the age of 13.

NOTES

CHAPTER 14

Cryptography

In this chapter, you will learn about

- Basic cryptography concepts and terms
- Hashing algorithms and integrity concepts
- Common symmetric encryption algorithms
- Differences between symmetric and asymmetric encryption
- Asymmetric encryption using public and private keys
- Asymmetric encryption and digital signatures
- Hiding data within data using steganography
- The components of a public key infrastructure
- Different types of cryptanalysis attacks

Understanding Basic Cryptography Concepts

[Chapter 1](#) introduces confidentiality and integrity (in the context of the CIA security triad of confidentiality, integrity, and availability). Cryptography includes confidentiality and integrity, but it's worth stressing that many cryptography processes provide authenticity too. Confidentiality from access controls starts with proper identification and authentication as described in [Chapter 2](#). Appropriate access controls then grant personnel access to data based on their proven identity. As an introduction for this chapter, the following list describes confidentiality, integrity, and authenticity within the context of cryptography:

- **Confidentiality** The goal of confidentiality is to prevent unauthorized access to data. Beyond access controls, the primary method of preventing loss of confidentiality is with encryption. This chapter focuses on encryption techniques to protect the confidentiality of data. Two common encryption techniques are symmetric encryption (with one key) and asymmetric encryption (with two keys).
- **Integrity** The goal of integrity is to prevent unauthorized modifications to data and systems. Beyond access controls, hashing is one of the primary methods of providing assurances that data has not lost integrity. A hash is calculated on data at different times, such as on Monday and then again on Tuesday. If the two hashes are the same, it proves that the data is the same and has maintained integrity. This chapter describes hashing and covers several hashing protocols.

- **Authenticity** Validation processes provide assurances that data is coming from a known source and is valid or reliable. Many encryption and hashing protocols provide authenticity. Additionally, some processes, such as digital signatures, have authenticity as a key goal. This chapter covers all of these concepts.



EXAM TIP Cryptography concepts have a lot of depth. However, the SSCP exam doesn't expect you to design cryptographic algorithms or be able to describe the inner workings of them. You do need to have some basic knowledge so that you can recognize differences between some basic concepts.

Cryptography Terminology

Cryptography has many terms that might not be familiar to you. The following list introduces some of these terms and topics, all of which are covered in greater depth in this chapter:

- **Plaintext** Data that is not in an encrypted format. Plaintext data (sometimes called cleartext data) is easily readable by anyone. You can think of it as ordinary text, such as what you're reading on this page.
- **Ciphertext** Data that is in an encrypted, or scrambled, format. Authorized personnel can decrypt ciphertext to return it to plaintext and read it. When strong encryption protocols are used, unauthorized personnel cannot read ciphertext.
- **Encryption** The process of converting plaintext data into ciphertext data to prevent loss of confidentiality. [Figure 14-1](#) shows the overall process of encrypting plaintext data into ciphertext data using an encryption algorithm. Decryption reverses the process by decrypting the ciphertext data with a decryption algorithm to re-create the original plaintext data.

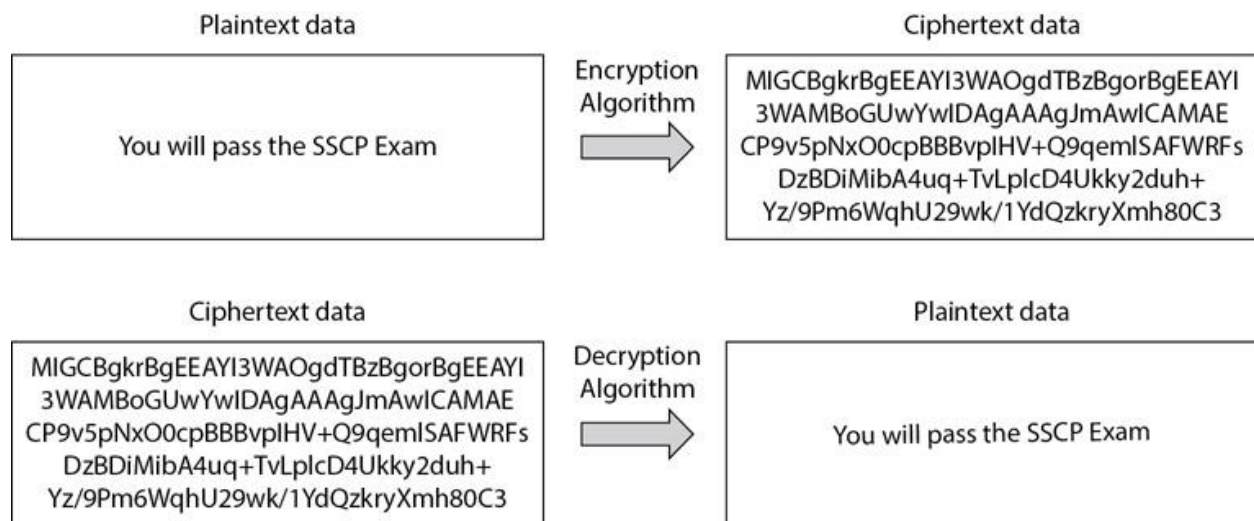


Figure 14-1 Encryption and decryption process

- **Encryption algorithm** A mathematical process that converts plaintext data to ciphertext data. Encryption algorithms include a process to decrypt the data. Encryption algorithms are often published so that anyone can see how they work. This allows experts to analyze and test the algorithms to verify their strength. For example, you can find the algorithm for Advanced Encryption Standard (AES) by doing an Internet search on “Advanced Encryption Standard algorithm” or “AES algorithm.”
- **Symmetric encryption** Data is encrypted and decrypted with the same key. For example, if an encryption algorithm uses a key of 101 to encrypt data, it will use the same key of 101 to decrypt the data. The algorithm would use a different key (such as 234) to encrypt and decrypt a different set of data. In practice, keys are much longer than three digits, such as 256, 512, 1024, or 2048 bits long.
- **Asymmetric encryption** Data is encrypted and decrypted with two keys. These two keys are a matched pair known as a public key and a private key. A private key can decrypt data that has been encrypted with the matching public key. Similarly, a public key can decrypt data that has been encrypted with the matching private key. The private key is always kept private and the public key is freely shared.



EXAM TIP Encryption helps prevent loss of confidentiality. The two primary categories of encryption are symmetric (using one key) and asymmetric (using two keys, known as a public key and a private key).

- **Cryptography** The science of scrambling, or encrypting, data. Cryptography includes the study of algorithms and other methods that help protect the confidentiality of data.
- **Cryptographic system** A system that includes the algorithms needed to support a particular method of encryption or decryption. Although most systems support standard algorithms, administrators sometimes need to add features to a computer to support a newer algorithm.
- **Cryptanalysis** The science of studying cryptographic methods. Security experts and attackers often use cryptanalysis methods to discover weaknesses in existing algorithms. Attackers attempt to use cryptanalysis to defeat a cryptosystem with the goal of compromising the confidentiality or integrity of protected data.
- **Hashing** The process of creating a hash from a file or a message with a hashing algorithm. A hash is simply a number. A hashing algorithm will always produce the same hash when executed against the same data. In other words, as long as a file or message never changes, a recalculated hash will always be the same no matter when it is calculated.

Data Sensitivity

The primary purpose of encrypting data is to preserve its confidentiality. However, this isn't required for all data. For example, [Chapter 11](#) discusses data classifications such as Top Secret, Secret, Confidential, and Unclassified. Data that is classified as Top Secret, Secret, or Confidential is sensitive, so it's appropriate to protect this data with encryption protocols. Similarly, organizations often use encryption methods to protect intellectual property, personally identifiable information (PII), and protected health information (PHI), all of which are also covered in [Chapter 11](#).

Additionally, the data sensitivity drives the selection of the encryption protocols. Higher data classifications (such as Top Secret) warrant the use of the strongest encryption protocols. In contrast, lower data classifications (such as Confidential) do not need the strongest encryption protocols.

You might be wondering, "Why not use the strongest encryption protocols for everything?" The reason is the performance cost. It takes processing power to encrypt and decrypt data, and if all data is encrypted and decrypted, it can slow the system down. Also, encrypted data typically consumes a lot more space than unencrypted data. As an example, [Figure 14-1](#) shows the plaintext data as a single sentence, but the ciphertext data is about five times as many characters. The size difference varies with different algorithms, but ciphertext data is larger than plaintext data. This takes more space when the encrypted data is saved on media such as a hard disk and consumes more bandwidth when transferred over a network.



NOTE While these examples are using data classifications of Top Secret, Secret, Confidential, and Unclassified, the same concepts apply no matter what data classification labels an organization uses. Typically, an organization would not encrypt any public data, but would often encrypt any other data that is not public.

Regulatory Requirements

Many jurisdictions have implemented regulatory requirements related to cryptography. [Chapter 13](#) covers some of the laws that mandate protection of privacy information. This includes the protection of personally identifiable information (PII) and protected health information (PHI), though some jurisdictions use different labels than PII and PHI. As a reminder, PII is any information that can personally identify an individual. PHI is similar but it also includes information concerning the health status, provision of health care, or payment of health care for an individual. One of the ways that organizations protect PII and PHI is by encrypting data at rest (stored on media such as a hard drive) and data in motion (transmitted over a network).

The Payment Card Industry Data Security Standard (PCI DSS) applies to any organizations that process credit card payments. [Chapter 10](#) provides an overview of PCI DSS in the context of security audits and lists the 6 control objectives and 12 supporting requirements. The overall goal is to protect credit card data. This includes data at rest and data in motion, and PCI DSS discusses the use of encryption as one of the methods to protect the data.

For data in motion, PCI DSS stresses the use of strong cryptography and security protocols such as Transport Layer Security (TLS), Internet Protocol security (IPsec), and Secure Shell (SSH). PCI DSS version 3.1 specifically discourages the use of SSL and early TLS versions because they are not strong cryptographic protocols.

[Chapter 11](#) discusses the Health Insurance Portability and Accountability Act (HIPAA) of 1996. It mandates the protection of protected health information (PHI), and one method to do so is encryption of PHI data at rest and in transit. Another method of protecting personal data is by using the Safe Harbor method. This uses anonymization and de-identification processes to eliminate 18 patient identifiers (such as names, telephone numbers, social security numbers, and more) to ensure that the data does not identify an individual. This is useful when collecting and analyzing data for research.

[Chapters 11](#) and [13](#) discuss the General Data Protection Regulation (GDPR). It strongly recommends the use of encryption to protect all personal data. If an organization chooses not to encrypt data, it needs to use alternative mechanisms to protect the data. However, if the organization suffers a data breach and it hasn't encrypted personal data, the organization may have difficulty proving that its alternative mechanisms were adequate.

Participating in Security Awareness and Training

As a security expert in your organization, you may be asked to participate in security awareness and training events to share your knowledge. For example, your organization may require users to encrypt certain data and digitally sign all of their e-mails. However, users may not understand the importance of these actions. By sharing your knowledge, you can help raise the overall security awareness within your organization.

User training isn't always formal. Many times, technicians and administrators provide informal training to end users while troubleshooting problems. Even though this is informal, it can often provide lasting meaning to the users.

Enforcing Integrity with Hashing

[Chapter 1](#) introduced hashing as one of the methods used to enforce integrity. Recall that a hash is simply a number created by executing an algorithm against a file or message. As long as the original data is not changed, the hash will always be the same. Knowing this, you can execute the hashing algorithm on the file or message at one point to create a hash, and then later execute the same hashing algorithm again. If the hashes are the same, you know the original data has not been modified. It has retained integrity.



EXAM TIP Hashing algorithms are used to enforce integrity. A hashing algorithm creates a hash. The hash is calculated at two different times, and if the hashes are the same, you know that the data has not lost integrity.

In contrast, if the hashes are different, you know that the file or message has lost integrity. You don't necessarily know why the file is different. A virus may have infected it, someone may have purposely modified it, or it may have simply lost a bit or two as it traveled over a network. However, you do know it is different and you should treat it with suspicion.

Consider [Figure 14-2](#). Lisa is sending a message (“You will pass the SSCP Exam”) to Homer. Before sending it, she calculates the hash and then sends both the message and the hash. Homer receives both the message and the hash. He can then recalculate the hash on the received message and compare the two hashes.

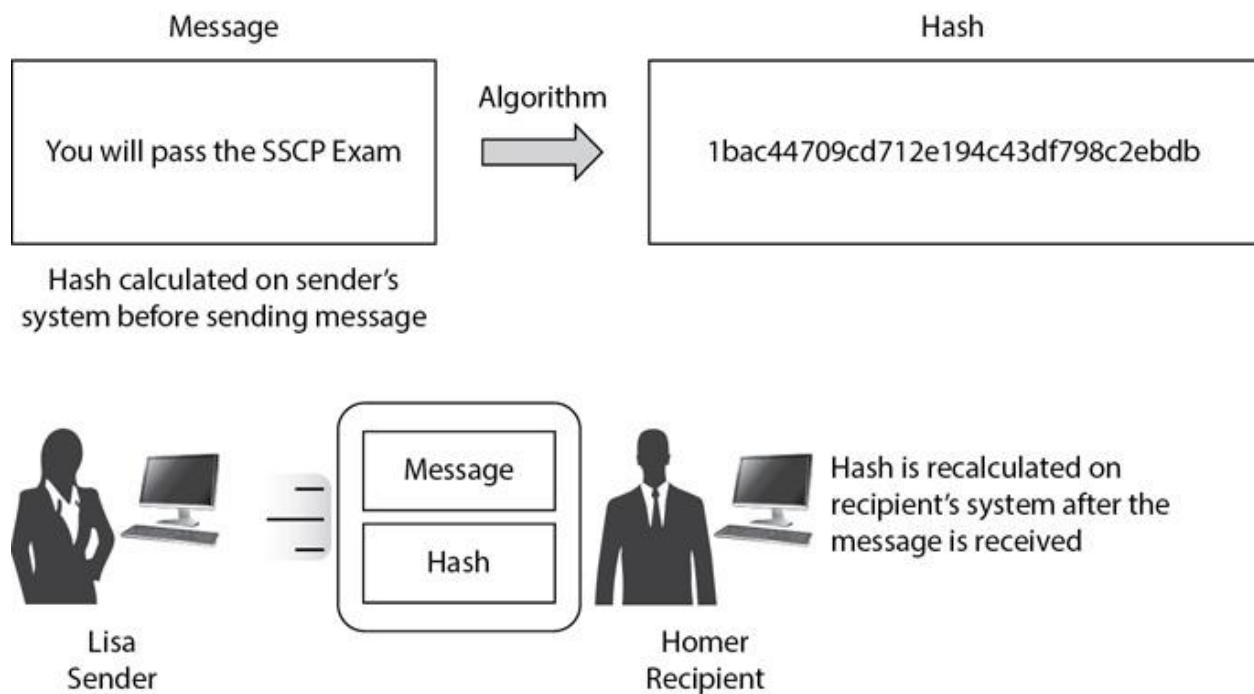


Figure 14-2 Using a hash to verify integrity

If the recalculated hash is the same as the sent hash, then the message has not lost integrity. If the recalculated hash is different, Homer knows that the message Lisa sent is not the same as the message that he received.



TIP If Lisa wanted to ensure the confidentiality of her message, she could have encrypted it to prevent unauthorized disclosure. However, because this discussion is focused on integrity, the message is not encrypted.

Hashing Algorithms Provide One-Way Encryption

An important point about hashing algorithms is that they provide one-way encryption of data. They create a hash by reducing the message or file to a message digest, which is a fixed-length number. However, in most cases, you can't reverse this process to discover the original data.

For example, the message "You will pass the SSCP Exam" results in a hash of 1bac44709cd712e194c43df798c2ebdb. However, you can't use this hash to determine the original message. The hash could have come from a 10KB message, a 5MB file, or data of any other size and content. As an example, consider this hash: df076ba5c4ac4d04cb286bde900c68f2. Just looking at the hash, there's no way that you could determine that it is a hash of a computer file on my computer.



EXAM TIP Hashing algorithms provide one-way encryption and are sometimes referred to as one-way functions. You can create a hash from a file or a message. However, you can't discover the original file or message if you only have the hash.

Hashing Algorithms

The following list describes some common hashing algorithms:

- **Message Digest 5 (MD5)** This creates a 128-bit hash, typically displayed as 32 hexadecimal numbers. It was widely used in applications, but security experts identified vulnerabilities with it in 2004. The U.S. government and security experts consider MD5 cryptographically broken, and it is not recommended for use in applications. The hash shown in [Figure 14-2](#) is the actual MD5 hash of the statement "You will pass the SSCP Exam" created from an online MD5 generator.
- **Secure Hashing Algorithm 1 (SHA-1)** This is a 160-bit hash. It was commonly used by U.S. government agencies such as the U.S. military until 2010. Due to vulnerabilities, it is no longer recommended for use.
- **Secure Hashing Algorithm 2 (SHA-2)** This is an improvement over SHA-1 and can create hashes of 224, 256, 384, or 512 bits. Although SHA-2 hasn't been broken, it functions similarly to SHA-1, and many experts believe it's just a matter of time before it is cracked. It's worth noting that SHA-256 is a version of SHA-2.
- **Secure Hashing Algorithm 3 (SHA-3)** The U.S. National Institute of Standards and Technology (NIST) sponsored the NIST hash function competition from 2007 to 2012. NIST

selected a subset of the Keccak algorithm as SHA-3. The algorithm is completely different from SHA-1 and SHA-2 and creates hashes of 224, 256, 384, or 512 bits.

- **Hash-based Message Authentication Code (HMAC)** The HMAC algorithm works with other hashing algorithms such as MD5 and SHA versions. It uses a shared secret key that adds some randomness to the hash. Only the sender and the receiver know the shared secret key. A primary benefit of HMAC over other hashing protocols is that it adds authenticity to the message. This also provides nonrepudiation because it verifies to the recipient that the message was sent by the other entity that knows the secret key.

While MD5 and the SHA algorithms provide integrity, HMAC provides both integrity and authenticity. [Figure 14-3](#) shows the extra step used with HMAC-MD5. In this example, both Lisa and Homer know the secret key is “You Can Pass.” Lisa’s system first calculates the MD5 hash on the message and then uses the secret key to create the HMAC-MD5 hash. She then sends both the message and the hash. Homer’s system then calculates the MD5 hash and uses the secret key to calculate the HMAC-MD5 hash. As long as both hashes are the same, Homer knows that the message has not lost integrity and that Lisa sent it.

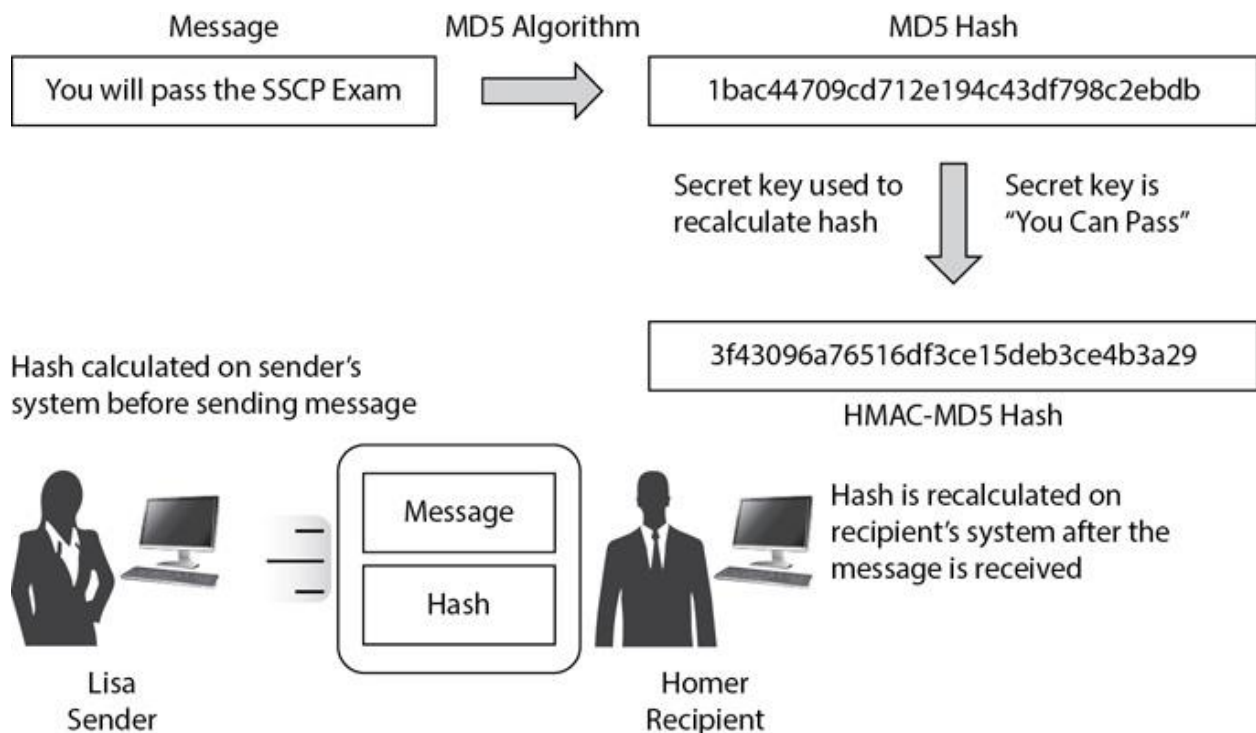


Figure 14-3 Using HMAC to provide integrity and authenticity

In contrast, imagine Bart wanted to impersonate Lisa and send a message to Homer. Because Bart doesn't know the secret key, he might guess one (such as "I won't write on chalkboards"). However, this secret key creates a completely different HMAC-MD5 hash. Homer's system would calculate the HMAC-MD5 hash using the secret key of "You Can Pass." After comparing both, his system would recognize that the keys are different and show an error. This indicates to Homer that Lisa did not send the message.



EXAM TIP MD5 creates a fixed-size 128-bit number (represented as 32 hexadecimal characters) from any message or file. SHA-1 creates 160-bit hashes. SHA-2 and SHA-3 create 224-, 256-, 384-, or 512-bit hashes. HMAC adds authenticity with a shared secret key.

Verifying Integrity with a Hash

Hashes can be verified automatically or manually. For example, applications that use digital signatures will automatically check the hash to verify the integrity of a message and report the results to the user. The section “Digital Signatures” later in this chapter shows the process used to check the hash automatically.

Other times, you might want to verify the hash of a file manually. For example, if you download a file from a website, you may want to verify that the file you downloaded is the same as the file that the website developers posted. Website developers sometimes calculate the hash of the file before posting it. They then post both the hash and the file. After you download the file, you can manually calculate the hash to see whether it’s the same as the posted hash.

As an example, imagine you download a file named `runthis.exe` from a website. On the website, you see that the SHA-1 hash is posted as `286bde900c68f2df076ba5c4ac4d04cb27a61ffb`. After you download the file, you can calculate the SHA-1 hash of the file and verify it is the same as the posted hash. If it is, you know the file hasn’t lost integrity. If it is different, you know the file you downloaded isn’t the same as it was when the web developer posted it.

There are many free hashing tools available on the Internet. I’ve downloaded the `sha1sum` tool, which you can access here: <http://gcapremium.com/data/sha1sum.zip>. Figure 14-4 shows a command prompt window with three commands executed. The first one is `dir`, which shows the files in this folder (or directory). I’ve copied the downloaded file (`runthis.exe`) and the two files used by the `sha1sum` tool to the SSCP folder. The second command (`sha1sum compute runthis.exe`) calculates the hash on the downloaded file. If I compare this to the file posted on the website, I can see that it is the same. This tells me the file has not lost integrity.

```
C:\windows\system32\cmd.exe
C:\SSCP>dir
Volume in drive C is Windows
Volume Serial Number is 260A-97DE

Directory of C:\SSCP
05/12/2015  09:23 AM    <DIR>          .
05/12/2015  09:23 AM    <DIR>          ..
02/22/2015  01:57 PM             11,280,580 runthis.exe
05/12/2011  07:47 AM              5,632 SHA1Sum.exe
05/12/2011  07:47 AM              4,608 SHA1Sum_library.dll
          3 File(s)             11,290,820 bytes
          2 Dir(s)          2,746,597,146,624 bytes free

C:\SSCP>sha1sum compute runthis.exe
286bde900c68f2df076ba5c4ac4d04cb27a61ffb  runthis.exe

C:\SSCP>sha1sum compute runthis.exe
286bde900c68f2df076ba5c4ac4d04cb27a61ffb  runthis.exe

C:\SSCP>sha1sum compute runthis.exe
286bde900c68f2df076ba5c4ac4d04cb27a61ffb  runthis.exe

C:\SSCP>
```

Figure 14-4 Calculating the hash with sha1sum.exe

As a reminder, no matter how many times I calculate the hash on the same data, I will always get the same hash. You can see this in the commands labeled 3 and 4 in [Figure 14-4](#).

If the hash of the downloaded file is different from the posted hash, you would know that the file you downloaded has lost integrity. You don't necessarily know what caused the problem, but you do know that you shouldn't trust this file. It may have just lost a bit or two during the download, but it also might be infected with malware.



TIP Md5sum.exe is another freeware tool that calculates hashes. While sha1sum.exe calculates a SHA-1 hash, md5sum.exe calculates the MD5 hash for files. Both are available as free downloads from multiple sites on the Internet.

Salting Passwords

[Chapter 5](#) discusses rainbow table attacks and how salting passwords helps prevent rainbow table attacks. Salting adds a random number of bits to a password. To get a better picture of this, it's valuable to understand how passwords are stored.

Passwords are rarely stored on a system in cleartext. A common way that systems and applications store passwords is in a hashed format. For example, each of the following bullets shows a password and the corresponding SHA-1 hash of the password. Instead of storing the passwords, systems store the hashes.

- **123456** 7c4a8d09ca3762af61e59520943dc26494f8941b

- **IW1IIP@33SSCP** f281f938734ccbb10aaf2a17f6ced2198be5582d
- **YouW1IIP@33SSCP** c33d3e769a77776bd2557029d0e9a1d6e0db4a28



TIP I mentioned earlier that you normally cannot reverse a hash to discover the original data. The exception is with weak passwords. Without salting, it is trivial for an attacker to discover the actual password from the SHA-1 hash of 7c4a8d09ca3762af61e59520943dc26494f8941b.

When a user logs on and enters a password, the system hashes the password and compares the hash of the entered password with the hash of the stored password. If they're the same, the system knows the user entered the correct password. This is also useful when transmitting the user's credentials over a network. Instead of sending the user's password, a system can send the hash of the password that the user sent.

While a hash of 7c4a8d09ca3762af61e59520943dc26494f8941b looks secure (certainly more secure than the weak password of 123456), attackers found that they could reverse-engineer the hash to discover the password. As mentioned in [Chapter 5](#), they developed rainbow tables, which are large tables of passwords (created from predefined character sets), and calculated hashes of the passwords. Attackers can then perform a simple rainbow table lookup process that compares the hash of a captured password with each of the hashes in the rainbow table. When they have a match, they have discovered the password.

A salt is a set of random bits added to a password before hashing it. These additional bits significantly change the password. For example, by adding the hexadecimal number 9 (1001 in binary), it changes the hash of the previous three passwords as shown in the following bullets. You can see that just a few bits changes the hash significantly. This makes it difficult for attackers to discover even the weak password of 123456.

- **123456** 99e2b89f3cda731add47da0c4a7698c1e4c80ba3
- **IW1IIP@33SSCP** 52d02acb0467a856f56c2974cb0d1d5d675042e7
- **YouW1IIP@33SSCP** 949245c0ab1c894eaa8ac42263c8747070707acb

Exploring Symmetric Encryption

Symmetric encryption uses a single key to encrypt and decrypt data. Symmetric encryption algorithms are constant, but the keys used to encrypt different sets of data are changed frequently to avoid compromising security. For example, the Advanced Encryption Standard (AES) always uses the same algorithm to encrypt and decrypt data, no matter how many times it's used. The key used to encrypt a set of data is also used to decrypt that data.

However, symmetric encryption algorithms use different keys to encrypt and decrypt other sets of data. For example, the algorithm might use a key of 123456 to encrypt and decrypt one block of data, but then use a key of 789321 to encrypt and decrypt another block of data. These keys are shortened for brevity but would be much longer. For example, AES has a default key length of 128 bits.



EXAM TIP Symmetric encryption uses a single key to encrypt and decrypt data. This key should be protected and changed regularly. In contrast, asymmetric encryption uses two keys (called public and private keys), where one key encrypts the information and the other key decrypts the information.

Many applications use symmetric encryption. For example, [Chapter 2](#) presents Kerberos as a network authentication protocol and discusses how it uses tickets for authentication. Kerberos uses symmetric encryption to protect the confidentiality of these tickets.

ROT13

The simplest way to demonstrate symmetric encryption is with ROT13, an encryption scheme similar to that used by ancient Romans. I've seen it used occasionally online by people attempting to make their messages harder to read. For example, if someone posts movie spoilers in online forums using ROT13, a reader can't easily read them.

ROT13 rotates letters 13 spaces. If you wanted to post a message saying, "The butler did it," with ROT13, you would rotate each letter forward 13 spaces. If you get to Z, then you continue with A. T rotated 13 spaces is G, h rotated 13 spaces is u, and so on. The ultimate message is "Gur ohgyre qvq vg." Anyone receiving this ciphertext can rotate the letters backward 13 spaces to read the plaintext. The ancient Romans used the same rotation method, but always rotated it 13 spaces.



TIP On sites such as <http://rot13.com>, when you enter a message, the site provides the ciphertext in ROT13 form. Similarly, you can enter the ROT13 form and decipher the message.

There are variations on ROT13. For example, you could use a rotation algorithm to encrypt and decrypt the data, but instead of always using the key of 13, you could use a different key at different times. On Monday, for instance, you might use a key of 3 to rotate the letters three spaces forward. To decrypt the data, you would use the same key of 3 but rotate three spaces backward. This allows you to modify the key regularly so that if other people discover the data, they won't necessarily know what the key is and may have trouble decoding it. Although this is a

very weak encryption algorithm compared to what's in use today, it does allow you to easily see the difference between the algorithm and the key.

This symmetric rotation algorithm could be defined as follows:

- **Encryption** Rotate X spaces forward to encrypt, where X is the key.
- **Decryption** Rotate X spaces backward to decrypt, where X is the key.
- **Key** The key is any number. Data encrypted with a key of X can be decrypted with the same key of X . In other words, if you use a key of 3 to encrypt the data, you would use the same key of 3 to decrypt the data. The challenge is to ensure that all authorized entities know what key to use.

Encryption algorithms are usually published and well known. For example, if you wanted to look up the encryption algorithm used for any of the symmetric encryptions discussed in this chapter, you could probably find them on the Internet. However, because data is encrypted with different keys, attackers will not be able to read ciphertext with only the algorithm. They also need the key.

Creating Strong Keys

Strong encryption algorithms create longer keys and change them often. This prevents attackers from cracking the encryption by guessing the key and ensures that the algorithm can use different keys. This is difficult to conceptualize with large keys, but consider keys of 2 bits, 3 bits, and 4 bits as shown in [Table 14-1](#).

| Bits | Calculation | Possible Key Combinations |
|------|-------------|---------------------------|
| 2 | 2^2 | 4 |
| 3 | 2^3 | 8 |
| 4 | 2^4 | 16 |
| 5 | 2^5 | 32 |

Table 14-1 Number of Bits Used in Keys

If an algorithm used only 2 bits, it could only create four keys, and attackers could easily crack the code in a brute-force attack using each of the four keys. However, with more bits, the algorithm can create more keys. Notice that adding a bit doubles the number of keys you can create. This is consistent with each bit you add.

Some older encryption algorithms used 56-bit keys, allowing them to create as many as 72,057,594,037,927,936 keys (over 72 quadrillion keys). While the number looks large, it is quite small by today's cryptographic standards. Attackers using basic computers can use a brute-

force attack to guess the correct key and crack a code in minutes, if not seconds. Of course, faster computers can go through the keys quicker.

When using the same algorithm (such as AES), longer keys provide stronger encryption and take longer to crack. For example, Advanced Encryption Standard (AES, covered in more depth later in this section) can use keys of 128 bits, 192 bits, or 256 bits. Imagine that an attacker could identify a 128-bit key using current computers within about 100 years. If you used AES-256 instead (with a 256-bit key), it would increase the time required to identify the key to a trillion years or so. The point is that using a longer key with the same encryption algorithm makes it theoretically impossible for an attacker to discover the key.



NOTE The length of time it takes to crack AES using 128 bits may be closer to 2,000 years, but this is theoretical and highly debatable. For example, the amount of processing power available to attackers even 100 years from now is difficult to predict.

In addition to using longer encryption keys, the encryption algorithm should create random keys and rotate between different keys. Ideally, the algorithm wouldn't ever use the same encryption key twice within the same set of data. If the algorithm did use the same key twice while encrypting a set of data, the code would be easier for an attacker to crack.



TIP Attackers can crack a code when weak keys are used or not changed often. Strong encryption algorithms start with a secure algorithm, use large keys, and change the keys often.

Comparing Block and Stream Ciphers

Two common types of ciphers used in symmetric encryption are block ciphers and stream ciphers. Both ciphers use a single key to encrypt and decrypt specific pieces of data, but they divide the data differently.

A block cipher encrypts fixed-length groups of bits, or blocks of data, individually. It starts by dividing the plaintext data into fixed-length sizes, such as 128-bit blocks of text. It then encrypts each block individually.

In contrast, a stream cipher encrypts the bits as a stream of data rather than as fixed-length blocks. Stream ciphers are susceptible to more vulnerabilities when not implemented properly. For example, an important rule related to stream ciphers is never to reuse a key. Wired Equivalent Privacy (WEP) is an early wireless security protocol that broke this rule, making it easy to crack. It's worth stressing that stream ciphers are secure when used properly.

Advanced Encryption Standard

Advanced Encryption Standard (AES) is one of the most popular symmetric encryption algorithms. It is a fast, efficient, 128-bit block cipher. NIST spent five years examining multiple encryption algorithms looking for a replacement for an older encryption algorithm. NIST ultimately selected the Rijndael algorithm, and both public and private sectors have adopted it for use. This is similar to how NIST selected a subset of the Keccak algorithm as SHA-3 to replace SHA-1 and SHA-2.

Many different applications use AES for encryption. For example, Wi-Fi Protected Access 2 (WPA2) uses AES to encrypt wireless transmissions. AES is used to encrypt entire hard drives with Trusted Platform Modules (TPMs) and can encrypt data on USB flash drives. These are just a few examples of AES's popularity as a symmetric encryption standard, and it is important to realize why—AES is strong and efficient.



EXAM TIP AES is a strong, efficient symmetric algorithm and is used in many applications. NIST selected it as a replacement for Data Encryption Standard (DES) over 15 other competing algorithms.

AES uses keys of varying sizes to add strength to the cipher:

- AES is sometimes listed as AES-128 and uses 128-bit keys.
- AES-192 uses a 192-bit key.
- AES-256 uses a 256-bit key.

Larger key sizes make the encryption significantly harder to break, but there is an associated resource cost. For example, AES-256 is stronger than AES-128, but systems can encrypt and decrypt AES using 128-bit keys quicker than they can encrypt and decrypt AES-256.

The U.S. National Security Agency (NSA) requires the use of AES-256 (256-bit keys) for data classified up to Top Secret.



TIP You can't accurately compare the strength of different encryption algorithms based simply on their key sizes alone. For example, imagine if I created a simplistic cipher called Gibson's Cipher and used 2048-bit keys. It wouldn't be accurate to say this cipher is stronger than AES-

256 simply because Gibson's Cipher uses bigger keys. However, it is accurate to say that AES-256 is stronger than AES-128 because it's the same algorithm with larger keys.

Other Symmetric Encryption Algorithms

The following list shows many of the other symmetric encryption algorithms that you may come across. However, none of these have the combined strength and efficiency of AES, so their usage is significantly less common today than it has been in years past. Additionally, attackers and security experts have discovered vulnerabilities in some protocols, so experts recommend discontinuing their use. This includes DES and RC4.

- **DES** DES is an older 64-bit block cipher that uses 56-bit keys. It has been broken and is rarely used anymore. It was a standard selected by the National Bureau of Standards (now known as NIST) but was replaced by AES.
- **3DES** Triple DES (or three DES) improves DES by encrypting data in three passes with three separate keys. It is also a 64-bit block cipher and it was one of the standards evaluated by NIST with AES but was not selected. It supports 56-bit, 112-bit, or 168-bit keys. It is a slower and processor-intensive block cipher but it is strong.
- **Blowfish** Blowfish is a 64-bit block cipher that is still used in many applications. It supports 32-bit to 448-bit keys and was another standard evaluated by NIST as a replacement for DES. It is a strong cipher and is still in use, although its use is being overshadowed by AES.
- **Bcrypt** Many Unix and Linux distributions use bcrypt to protect passwords stored in the shadow password file. It salts the password with additional bits before encrypting it with Blowfish. Bcrypt is one of many key-stretching techniques. Key stretching transforms a potentially weak password into a more secure password.
- **Password-Based Key Derivation Function 2 (PBKDF2)** PBKDF2 is another key-stretching technique. It adds a salt of at least 64 bits and then applies a cryptographic function (such as a cipher or hashing algorithm) multiple times. PBKDF2 can be used for passwords similar to bcrypt or to create stronger cryptographic keys used for symmetric encryption.



TIP Key-stretching techniques convert a potentially weak key or password into a more secure one. A common method of key stretching is to add additional bits to the password. Bcrypt and PBKDF2 are two techniques commonly used to create stronger passwords even if a user creates a relatively weak password. Key stretching can help thwart brute-force and rainbow table attacks.

- **IDEA** The International Data Encryption Algorithm (IDEA) is a 64-bit block cipher using 128-bit keys that was also submitted as a replacement for DES. Many implementations of Pretty

Good Privacy (PGP) used IDEA for encrypting e-mails, but IDEA is rarely used with current implementations of PGP.

- **RC4** Ron Rivest invented Rivest's Cipher (RC4), and it has been a widely used symmetric stream cipher. It has been used with Transport Layer Security (TLS) to encrypt Hypertext Transfer Protocol Secure (HTTPS) traffic. Security experts have identified multiple vulnerabilities in RC4, and it is no longer recommended for use.

[Table 14-2](#) summarizes key points about some popular symmetric encryption algorithms.

| Algorithm | Block or Stream Cipher | Keys |
|-----------|------------------------|-----------------------------------|
| AES | 128-bit block | 128-bit, 192-bit, or 256-bit keys |
| DES | 64-bit block | 56-bit key |
| 3DES | 64-bit block | 56-bit, 112-bit, or 168-bit keys |
| Blowfish | 64-bit block | 32-bit to 448-bit keys |
| IDEA | 64-bit block | 128-bit keys |
| RC4 | Stream cipher | 40-bit to 2048-bit keys |

Table 14-2 Symmetric Encryption Algorithms

Exploring Asymmetric Encryption

Asymmetric encryption (also known as public key cryptography) uses two matched keys identified as a public key and a private key. There are several important points to remember related to these keys:

- Anything encrypted with the public key can only be decrypted with the matching private key.
- Anything encrypted with the private key can only be decrypted with the matching public key.
- The private key is always kept private and never shared.
- The public key is freely shared and publicly available.

Asymmetric encryption requires a public key infrastructure (PKI, described in more detail later in this chapter) to create, manage, distribute, validate, and revoke certificates. A public key is embedded in a freely shared certificate. For example, [Figure 14-5](#) shows a certificate provided by www.gcgapremium.com after a user requests a secure session. This certificate is using a 2048-bit public key, and the matching private key is stored privately on the www.gcgapremium.com web server.

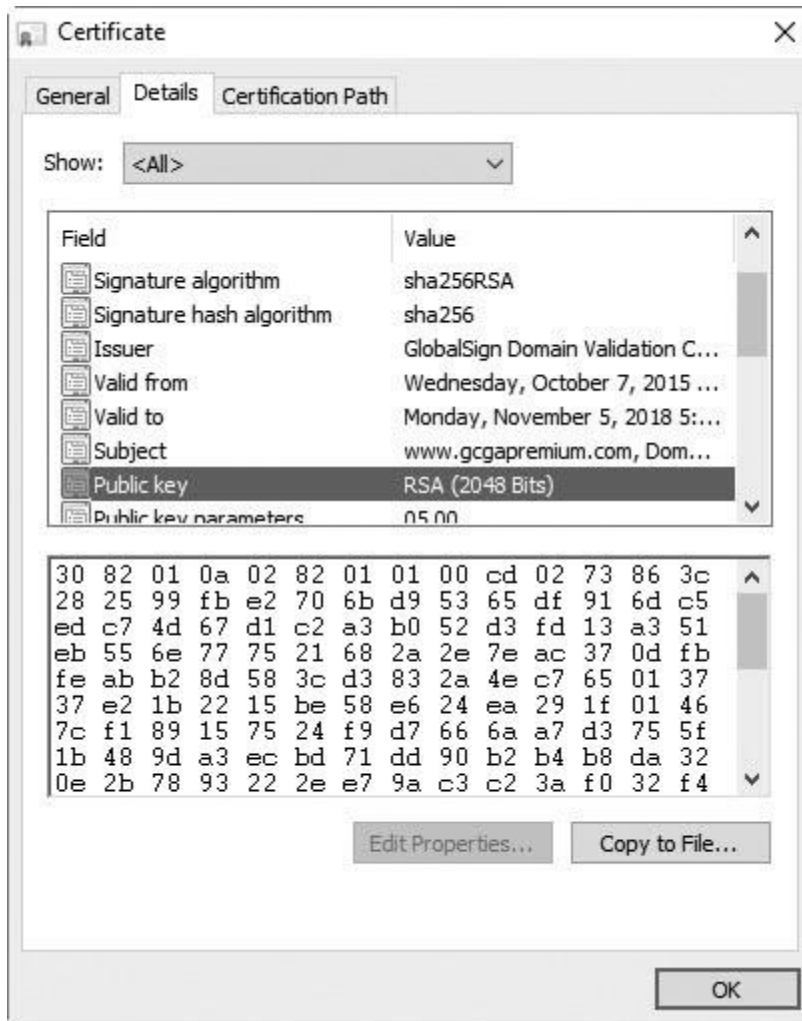


Figure 14-5 Certificate with public key selected

A common question people ask is, “Which key encrypts and which key decrypts?” The answer depends on the goal. The following are two common uses of these keys in asymmetric encryption:

- **Privately share a symmetric key between parties over a public network** In this case, the public key encrypts and the private key decrypts. The “Transport Layer Security” section later in this chapter describes how this works.
- **Digitally sign an e-mail** In this case, the private key encrypts and the public key decrypts. The “Digital Signatures” section describes how this works.



EXAM TIP Public keys are embedded in certificates and freely shared. The public key is matched with a private key, and the two matched keys work together to encrypt and decrypt data.

RSA

RSA is widely used in public key cryptography for privately sharing symmetric keys (key exchange) and digital signatures (described later in this chapter). Ron Rivest, Adi Shamir, and Leonard Adleman publicly described the RSA algorithm in 1978, and its name is derived from the first letters of their last names. In the field of cryptography, it is rather amazing that an algorithm created in 1978 is still in use today.

The reason why it has lasted so long is that it can adapt to using larger keys as computing power increases. Early implementations used smaller keys such as 512 bits. These were secure for a very long time but can be cracked relatively quickly today. NIST recommended using 1024-bit keys for any data that needed to be protected through the year 2015. Beyond 2015, NIST recommends using 2048-bit keys. Security experts predict that 2048-bit keys will likely be secure until 2030. Beyond 2030, 3072-bit keys are currently recommended. However, processing power will be much greater in 2030, so applications may use 4096-bit keys instead. These recommendations are based on computing power and the likelihood that cryptanalysis techniques will be able to crack 2048-bit keys in the future.

RSA uses large prime numbers to create secure, matching public and private keys. These large prime numbers are multiplied together to create a composite number that is then manipulated to create the public and private keys. Although it's easy to multiply the two large prime numbers together to create the composite number, it is not easy to factor the composite number to discover the original two prime numbers.

In other words, when the prime numbers are large enough, it's not possible to use the composite number to identify them, at least not in a reasonable amount of time. The original prime numbers remain secret, and more importantly, the private key derived from these numbers remains secret. Due to the increasing speed of computers, the goal of "large enough" prime numbers has been a moving target. Similarly, the size of public and private keys has been steadily increasing to stay ahead.



EXAM TIP RSA uses large prime numbers. Its strength lies in the fact that factoring the composite number created from these prime numbers is computationally infeasible as long as the prime numbers are sufficiently large.

Secure Sockets Layer

Secure Sockets Layer (SSL) is the original protocol used to encrypt HTTPS sessions. Version 1 was never released due to security flaws. Version 2 was released in 1995, but it also had several

security flaws. Netscape released version 3 in 1996, but there wasn't any formal method of updating it. The Internet Engineering Task Force (IETF) ultimately created Transport Layer Security (TLS) as the replacement and has been steadily updating and improving TLS over the years.



TIP SSL has not been defined as a formal RFC standard. However, the IETF released RFC 6101 as a Historic document that describes SSL in depth.

A Google security team disclosed a vulnerability with SSL named POODLE (Padding Oracle On Downgraded Legacy Encryption) in late 2014. While many organizations had already switched to TLS by this time, this announcement prompted more to do so. Soon afterward, many organizations disabled SSL in their applications, and many applications now specifically block its use.

Transport Layer Security

Transport Layer Security (TLS) version 1.0 is an upgraded version of SSLv3, and just as SSLv3 has vulnerabilities, TLSv1.0 has vulnerabilities and should not be used. TLSv1.1 was defined in April 2006 and upgraded as TLSv1.2 in August 2008. While TLSv1.1 and v1.2 initially included some backward compatibility with SSL, the IETF recommended in 2011 that this “be phased out with all due haste.” This helps eliminate potential downgrade attacks such as the POODLE attack mentioned previously.

As of mid-2018, TLSv1.3 is in draft stage. While some web browsers began implementing it as the default TLS version, some incompatibilities caused most of them to revert to TLSv1.2. This could change as TLSv1.3 goes through the draft stage and these issues are resolved.

TLS encrypts Internet transmissions as HTTPS over port 443. TLS is also used for other purposes, such as securing virtual private network (VPN) transmissions, but this section is focused on TLS's use within HTTPS. It describes how TLS uses asymmetric encryption to share a symmetric key privately.

It's important to realize that TLS uses both symmetric and asymmetric encryption, as follows:

- **TLS uses asymmetric encryption to share the session key privately** The public and private asymmetric keys are used to encrypt and decrypt the session key.
- **TLS uses symmetric encryption to encrypt session data** The session key is then used to encrypt and decrypt the data within the session.

[Figure 14-6](#) shows the overall process. In this example, Bart is establishing a secure session with a web-based server, such as a web-based e-mail site. Bart establishes an HTTPS session before

entering his credentials to access the e-mail, and the session ensures his credentials aren't sent over the Internet in cleartext. As you look at the figure, you can trace the following steps:

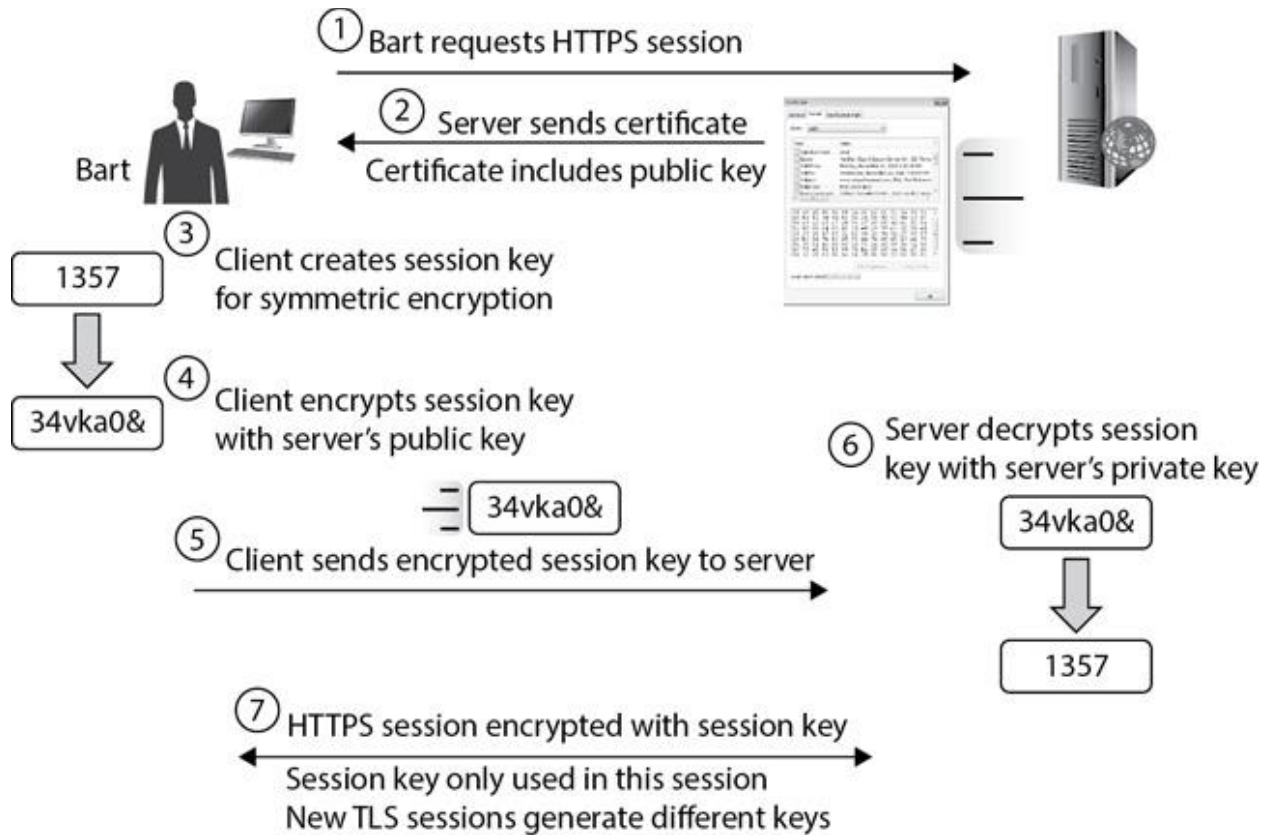


Figure 14-6 Creating a secure HTTPS session with TLS

1. Bart initiates the process by requesting an HTTPS session. For example, he could simply click a link on a website page, such as an account logon page.
2. The server responds by sending its certificate with the server's public key embedded within it. This public key is matched with a private key known only by the server.
3. After receiving the certificate, the client (Bart's computer) creates a session key (which will be used as a symmetric key) that will be used later to encrypt session data. At this point, only the client knows the session key. For brevity, this discussion is using a simplistic session key of 1357, but it would be much larger in practice. For example, TLS can use a 256-bit AES key.
4. The client encrypts the session key with the server's public key. Recall that only the matching private key (which only the server knows) can decrypt the encrypted session key.
5. The client sends the encrypted session key to the server over the Internet. Even if an attacker captures this encrypted session key, it won't be of any use because the attacker won't have the matching private key needed to decrypt the session key.

6. When the server receives the encrypted session key, it decrypts it with its private key. At this point, both the client and the server know the session key, but no one else knows it.

7. The session key encrypts data within the HTTPS session using symmetric encryption. If attackers intercept any of these pages, they won't be able to read the data because they don't have the session key. After the session ends, the client and server discard the session key. If the user initiates another HTTPS session, the two systems repeat the process. Session keys have the shortest lifespan when compared to other keys.



TIP Although [Figure 14-6](#) doesn't show it, the client validates the certificate before using it. This includes ensuring that the certificate hasn't expired and that the certificate authority hasn't revoked the certificate. The "Validating Certificates" section later in this chapter describes this process in a little more depth.

A logical question to ask here is why not use asymmetric encryption to encrypt the entire session. The reason is that symmetric encryption is much faster than asymmetric encryption. Depending on the algorithm used, symmetric encryption can be as much as 100 times faster than asymmetric encryption.



TIP Historically, HTTPS sessions were primarily used to encrypt sessions that transferred sensitive data such as credentials and credit card data. Other data was sent via HTTP in cleartext. However, more and more websites are moving to HTTPS by default. A report by Search Engine Journal (SEJ) in May 2017 stated that 31 percent of the top 100,000 domains use HTTPS by default, and that number continues to climb.

SSL Decryptors

An SSL decryptor (sometimes called a TLS decryptor) is placed between user systems and the Internet. Organizations often use them to decrypt TLS session traffic and examine the decrypted data. SSL decryptors can often help detect potential problems.



TIP An SSL decryptor is technically a TLS decryptor because it processes TLS sessions. However, they are often still called SSL decryptors.

Imagine a user goes to a website and establishes a secure session using HTTPS. By default, the session is created between the user's computer and the website and the data is encrypted along the entire path.

However, if you add an SSL decryptor between the user's computer and the Internet, it will establish two TLS sessions. One session is between the user's computer and the SSL decryptor. The second session is between the website and the SSL decryptor. These two sessions are encrypted but the data on the SSL decryptor is in an unencrypted format. The SSL decryptor can then inspect this unencrypted data for potential issues. In some scenarios, the SSL decryptor sends the unencrypted data to an intrusion detection system (IDS) for analysis.

Consider the following two issues that may not be detected without an SSL decryptor:

- **Botnet traffic** After a computer has been joined to a botnet, it will often exchange data and commands with a command and control server. This data is often encrypted in a TLS session and unreadable unless it is decrypted.
- **Drive-by downloads** When users click a malicious link, they are often taken to a server that attempts to download malware. When the session is encrypted, antimalware software won't be able to detect the malware until it is decrypted.

In contrast, with an SSL decryptor in place, the botnet traffic and malware are likely to be detected after they are decrypted. An SSL decryptor can only decrypt TLS sessions. It can't decrypt traffic that was encrypted on client computers outside of a TLS session.

As an example, advanced persistent threat (APT) attackers often collect data within a network, encrypt it, and then send it to an APT-controlled server. This data is typically encrypted on client computers before being sent. Even if the infected computer establishes a secure TLS session with the APT-controlled server, the SSL decryptor will see the data in an encrypted format because it was encrypted before the TLS session. However, this can still provide a warning. It isn't common to encrypt data before encrypting it again and sending it in a TLS session.

Diffie-Hellman

Diffie-Hellman is one of the first public key algorithms used to privately share symmetric keys between two systems. Whitfield Diffie and Martin Hellman published the method in 1976, and it has been adopted in many public key algorithms. For example, TLS supports the use of Diffie-Hellman to exchange session keys.

The focus of Diffie-Hellman is on secure key exchange. As a comparison, RSA provides secure key exchange and is also a digital signature algorithm used to create digital signatures. The "Digital Signatures" section later in this chapter describes digital signatures.

Elliptic Curve Cryptography

Elliptic curve cryptography (ECC) uses a mathematically defined curve to create public and private keys. A benefit of ECC is that it provides strong security with smaller key sizes when compared to other methods such as RSA. As an example, a 256-bit ECC key provides similar strength as a 3072-bit RSA key. ECC is often used in smaller mobile devices because the smaller keys require less processing power.

Several algorithms support ECC. For example, the Elliptic Curve Digital Signature Algorithm (ECDSA) uses ECC methods for digital signatures. Similarly, elliptic curve Diffie-Hellman (ECDH) is a variant of Diffie-Hellman that uses ECC for key exchange.



NOTE Reports in 2013 indicated that the NSA influenced NIST to include a flaw in a recommended elliptic curve. Later that year, RSA Security LLC issued an advisory recommending the discontinued use of ECC based on Dual Elliptic Curve Deterministic Random Bit Generation (Dual_EC_DRBG).

Advances in quantum computing have raised some concerns about the security of ECC. Experts predict that a sufficiently strong (hypothetical) quantum computer will be able to discover keys used in public key cryptography. Because ECC uses smaller keys, ECC is likely to be a target of attackers as quantum computing capabilities increase. However, as of mid-2018, this remains theoretical because quantum computers are not sufficiently strong to solve the complex mathematical problems associated with cracking public key cryptography.

Secure Shell

Secure Shell (SSH) uses public key cryptography to encrypt data in motion. [Chapter 3](#) discusses the use of SSH in more depth. It is a secure alternative to Telnet and similar protocols such as rlogin, rsh, and rexec, which all send credentials over a network in cleartext. Secure File Transfer Protocol (SFTP) and Secure Copy (SCP) both use SSH to encrypt the data transfers.

Protecting E-mail with S/MIME

Secure/Multipurpose Internet Mail Extensions (S/MIME) is the standard used to encrypt and digitally sign e-mail. Almost all e-mail applications that support encryption and digital signatures use S/MIME. The following sections describe how S/MIME uses digital signatures and encrypts e-mail.

Nonrepudiation

Nonrepudiation prevents individuals from believably denying they took an action. This is similar to how you cannot believably deny making a charge on your credit card if you provided your signature during the transaction. Of course, you can deny it, but your claim isn't believable.

Electronic signatures provide another example of nonrepudiation. Mortgage companies, banks, realtors, and many other organizations use electronic signatures when sharing documents over the Internet. Electronic signatures are valid and legally binding, just as a signature made with a pen on paper is legally binding.

DocuSign is one of the companies many organizations use for online electronic signatures. Here's how it works. A company representative uploads documents to a DocuSign site and identifies the customer that needs to sign it, using the customer's e-mail address. The representative then marks the document where the customer needs to initial or sign it. The customer receives an e-mail with a link for the document. They agree to use the electronic signature and follow the instructions to sign and initial the document. The customer cannot later say they never received the document, because the electronic signature provides legal proof that they did.

Digital signatures (described in the next section) are another method used for nonrepudiation. [Chapter 10](#) discussed audit logs and audit trails. When secure identification and authentication processes are in place, these audit logs and audit trails also provide nonrepudiation.

Digital Signatures

A digital signature is an encrypted hash of a message. It uses asymmetric encryption and is encrypted with the sender's private key. Digital signatures provide several security benefits, including authentication, integrity, and nonrepudiation.

- **Authentication** The digital signature provides proof to the recipient that the claimed sender is the actual person who sent the e-mail. In this context, the digital signature provides authenticity.
- **Integrity** The e-mail is hashed before it is sent, and the hash is included with the e-mail. The recipient hashes the received message and verifies that the hashes are the same. If the hashes are the same, that verifies that the message has not been modified. In other words, it verifies the message has not lost integrity.
- **Nonrepudiation** The sender creates the digital signature with the sender's private key. Because only the sender has the sender's private key, the sender cannot deny sending it.



TIP A Digital Signature Algorithm (DSA) creates digital signatures. RSA includes a DSA component, and many applications use RSA to create digital signatures. The Elliptic Curve Digital Signature Algorithm (ECDSA) is another DSA, but unlike RSA, ECDSA does not include a key-exchange component.

[Figure 14-7](#) shows the overall process of creating and using a digital signature. In this example, Maggie is sending an e-mail to Lisa and using the digital signature for authentication, integrity, and nonrepudiation. As you look at the figure, you can trace the following steps:

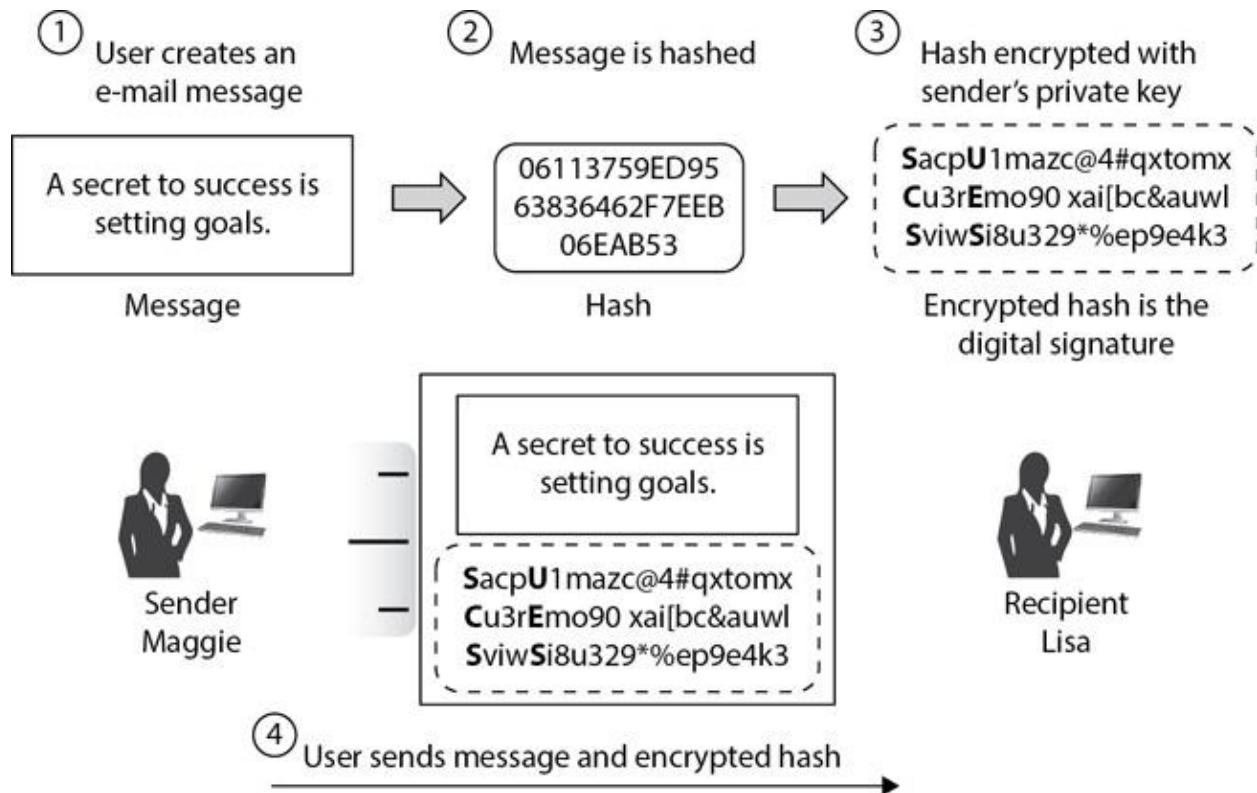


Figure 14-7 Creating a digital signature

1. Maggie, the sender, first creates the message. In this example the message is, “A secret to success is setting goals.”
2. A hashing algorithm creates a hash from the message.
3. The hash is encrypted with the sender’s private key, creating the digital signature. As mentioned in the “Enforcing Integrity with Hashing” section earlier in this chapter, this hash provides integrity.

There are two key points to remember at this point:

- Only the sender has the sender’s private key.

- Only the sender’s public key can decrypt the encrypted hash.

4. Maggie sends both the message and the encrypted hash to Lisa, the recipient.

Note that the message isn’t encrypted as part of the digital signature. In other words, the digital signature does not provide confidentiality. However, if it is important to maintain confidentiality of the message, you can encrypt an e-mail. The “Encrypting E-mail” section later in this chapter shows the encryption process. While it is possible to encrypt and digitally sign the same e-mail message, this section focuses only on the digital signature.

When the recipient receives a digitally signed e-mail, the receiving system verifies the digital signature. [Figure 14-8](#) shows the steps taken after a user receives a digitally signed e-mail. As you look at the figure, you can trace the following steps:

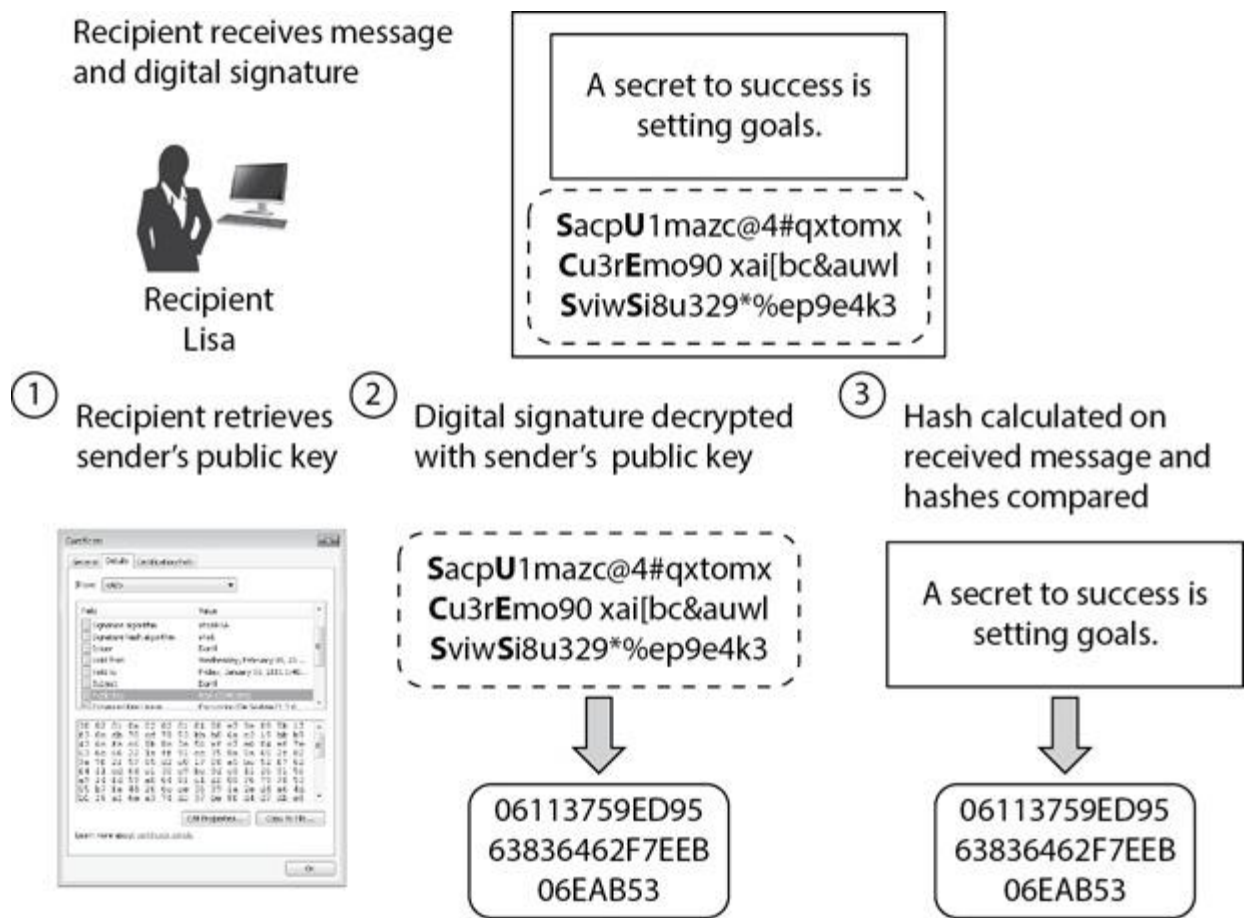


Figure 14-8 Verifying a digital signature

1. The recipient retrieves the sender’s public key by retrieving the sender’s certificate. Within an internal network (such as in a Microsoft domain), this process is automated without requiring any user interaction. In other situations, users may be required to install the sender’s certificate on their system manually. The certificate only needs to be retrieved once.

2. The digital signature is decrypted with the sender's public key. When successful, this step shows the original hash of the message created on the sender's system. This is an important step. It will succeed only if the sender's private key encrypted the hash.

- If the decryption of the digital signature is successful, it proves that it was encrypted with the sender's private key. Because only the sender has the sender's private key, the sender must have encrypted it. This provides both authentication and nonrepudiation. The recipient knows the sender sent the message (authentication), and the sender cannot later deny sending the message (nonrepudiation).

- In contrast, if someone tried to impersonate the sender and included a bogus digital signature, this step would fail. The impersonated sender's public key would not decrypt the digital signature, so the process would show an error.

3. The recipient's system calculates the hash on the received message. It then compares the calculated hash with the decrypted hash from the digital signature.

- If both the received hash and the calculated hash are the same, the message has not been modified (or has not lost integrity).

- If the hashes are different, the recipient knows that the received message is not the same as the sent message. The recipient wouldn't know how the message was modified. The differences could be due to tampering from an attacker, a problem with the network, hardware, or software, or something else. However, the recipient knows the content of the message should not be trusted.



EXAM TIP The primary purpose of a digital signature is to provide authentication. However, a digital signature also provides integrity and nonrepudiation.

Encrypting E-mail

You can also use asymmetric encryption to encrypt e-mail. In simplistic terms, here's what happens:

1. The sender retrieves the recipient's public key.
2. The sender encrypts the e-mail with the recipient's public key.
3. The recipient decrypts the e-mail with the recipient's private key.



EXAM TIP An important point to remember about e-mail encryption is that the sender requires the recipient's public key.

Although the preceding steps simplify the process, e-mail applications implement the encryption process with some additional steps. The following explanation provides the underlying details.

In the “Transport Layer Security” section, you saw how TLS uses public key cryptography to privately share a session key and how the session key is used to encrypt the session data. E-mail encryption typically uses a similar process. The primary difference is in how the sender retrieves the certificate holding the recipient's public key.

The certificate retrieval process is automated in many internal networks. For example, in a Microsoft domain, the user could use Microsoft Outlook to draft an e-mail similar to that shown in [Figure 14-9](#) and click the Encrypt button. When the user clicks Send, Microsoft Outlook attempts to retrieve the certificate of the recipient based on the e-mail address. In other situations, users may be required to install the recipient's certificate on their system manually.

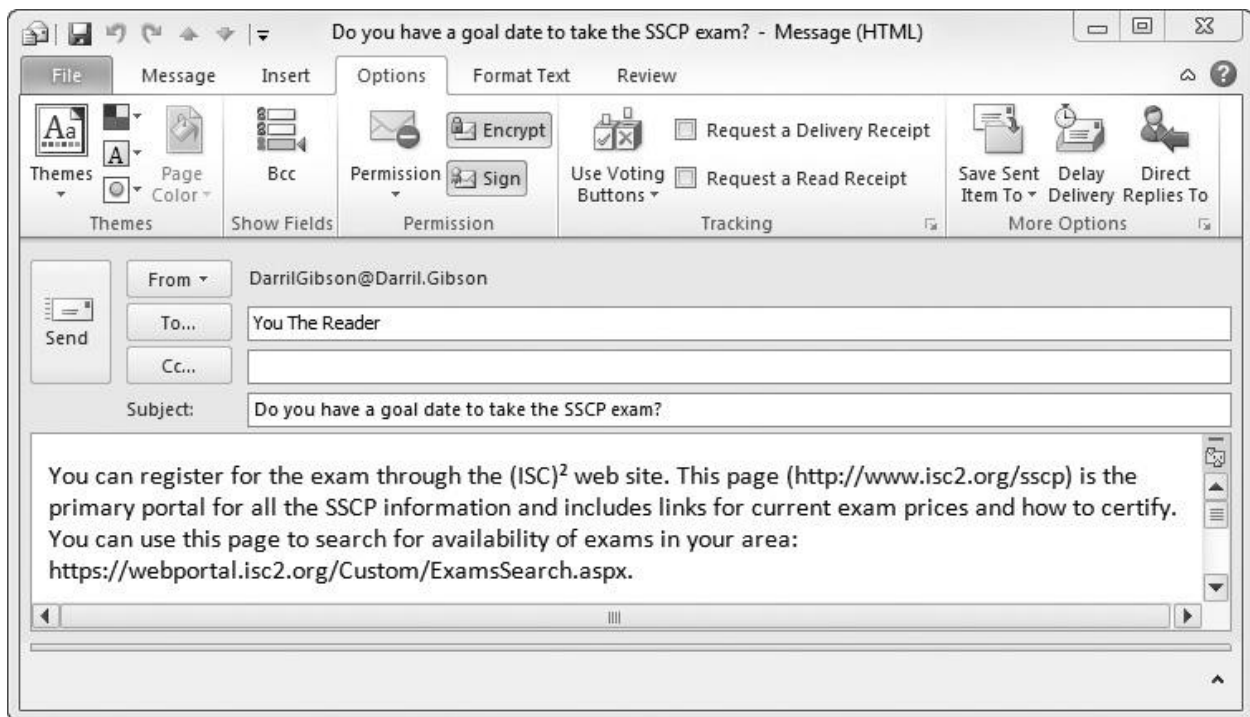


Figure 14-9 Encrypting an e-mail in Microsoft Outlook



NOTE [Figure 14-9](#) also shows the Sign button, which is used to create a digital signature. Although it's technically possible to encrypt an e-mail without digitally signing it, e-mail

applications rarely allow you to do so. Instead, if you try to encrypt an e-mail, the application will force you to add a digital signature.

Once the sender's system has the recipient's certificate, it uses the steps shown in [Figure 14-10](#) and described in the following list to encrypt the e-mail:



Figure 14-10 Encrypting an e-mail with symmetric and asymmetric encryption

1. The sender's system creates a session key. This is similar to how TLS creates a session key. For brevity, this discussion is using a simplistic session key of 2468, but the key would be much larger in practice.
2. The sender's system uses symmetric encryption to encrypt the e-mail with the session key.
3. The sender's system uses asymmetric encryption to encrypt the session key with the recipient's public key.
4. The sender's system sends the encrypted e-mail and the encrypted session key to the recipient.
5. The recipient's system retrieves the recipient's private key.
6. The recipient's system uses the private key to decrypt the encrypted session key.
7. The recipient's system decrypts the e-mail content by using the decrypted session key.

Like the encryption process, the decryption process is automated. For example, the recipient may only need to open the e-mail to view the decrypted content.



TIP Symmetric encryption is significantly faster than asymmetric encryption. You can improve encryption performance by using both symmetric and asymmetric encryption to encrypt the e-mail.

Protecting E-mail with DKIM

DomainKeys Identified Mail (DKIM) is an e-mail authentication mechanism that helps detect e-mail spoofing. It was derived as a combination of enhanced DomainKeys created by Yahoo and Identified Internet Mail created by Cisco.

Attackers often send e-mails with forged sender addresses to trick users into thinking that the e-mails are valid. DKIM allows other e-mail servers to easily detect that these e-mails are forged.



TIP [Chapter 6](#) discussed Sender Policy Framework (SPF) records as a method of identifying spoofed e-mail. The SPF record is published on a Domain Name System (DNS) server that is authoritative for the domain. DKIM records are also published on a DNS server that is authoritative for the domain.

If you get an e-mail that indicates it is from an e-mail address of orders@amazon.com, you will likely think it came from amazon.com. However, it's a trivial matter to change the From address of any e-mail. Any attacker can change the From address, use a subject line related to an order, and ask the recipient to click a link to verify something. The link from an attacker is very likely malicious. Someone who receives this e-mail but is not an Amazon customer may accuse Amazon of sending out malicious spam, even though Amazon was not involved.

However, with DKIM, this type of forged e-mail is thwarted. An organization would first publish a DKIM DNS record. Next, it would configure its e-mail servers to attach a DKIM signature to every e-mail it sends out. Any e-mail server that receives an e-mail can match the DKIM DNS record with the signature in the e-mail. If they match, the e-mail is known to be valid. If the signature verification fails, servers send messages indicating a failure.

DKIM isn't mandatory, so an e-mail without a DKIM signature doesn't clearly indicate it is forged. Domain-based Message Authentication, Reporting and Conformance (DMARC) helps with this. DMARC allows organizations to publish a policy indicating if they are using DKIM,

SPF, or both on outgoing e-mail. DMARC also tells a receiving server what to do if the DKIM and/or SPF methods fail, such as marking it as junk or rejecting the message.

While DKIM can use keys ranging from 512 bits to 2048 bits, RFC 6376 recommends using keys at least 1024 bits long. Mathematician Zach Harris found that many organizations use short 384-bit keys. In October 2012, he found that he could spoof these keys in as little as 24 hours on his laptop, and spoof 512-bit keys in about 72 hours using more sophisticated resources.

PGP and GPG

Pretty Good Privacy (PGP) can encrypt, decrypt, and digitally sign e-mails between individuals. PGP is an application, originally created by Philip Zimmermann, that provides confidentiality through encryption, as well as authentication, with the use of digital signatures.

For example, if you and I wanted to be able to send e-mails between each other and keep them confidential, we could use PGP. I could send you my certificate, including my public key, and you could add it to PGP running on your system. Similarly, you would send me your certificate.

PGP creates a session key, encrypts the e-mail with it, and then uses the public key to encrypt the session key. When my system receives the encrypted e-mail and the encrypted session key from you, PGP uses my matching private key to decrypt the session key and decrypts the e-mail with the decrypted session key. Of course, I could use the same process to send you an encrypted e-mail.

In addition, either of us could digitally sign our e-mails to prove that we sent them. For example, I could use PGP to create a digital signature with my private key. If you had a certificate with my public key, PGP would use it to decrypt the digital signature, providing you with assurances that it was created from my private key.

PGP started as open source software freely available to anyone. However, it has gone through several iterations of ownership and is currently proprietary software owned by Symantec.

GNU Privacy Guard (GnuPG or GPG) was later developed as a free, open source alternative to PGP. It complies with the specification for OpenPGP defined in RFC 4880.

A significant difference between S/MIME, PGP, and GPG is how certificates are validated. Remember that a certificate has a public key and is freely shared. This public key is matched to a private key held by the certificate owner. S/MIME uses a centralized certificate authority within a PKI. Certificate authorities and PKI systems are described in more detail in the “Understanding Public Key Infrastructure Systems” section.

PGP and GPG use a decentralized web of trust (WOT). In a WOT, several individuals vouch for the trustworthiness of certificates. When enough people indicate that the certificate is trustworthy, others trust it. A WOT is also discussed in more detail in the “Understanding Public Key Infrastructure Systems” section to provide a better comparison between a centralized PKI and a decentralized WOT.

Other Encryption Schemes

Many encryption schemes fit neatly into the symmetric or asymmetric categories, but a few fall outside the scope of these categories. For example, cryptographers can use steganography to manipulate files and embed detailed messages within them. However, steganography simply hides the data within the file. IPsec is another encryption scheme that doesn't fit neatly into the symmetric and asymmetric categories, but it is important in the context of cryptography.

Steganography

Steganography is the practice of hiding data within data or hiding data in plain sight. It doesn't actually encrypt the data, but it can provide a level of confidentiality. If others know that the message is embedded in the data, they can retrieve it.

As an example, some programs modify the least significant bit of individual bytes within a file. For example, consider the byte represented in [Table 14-3](#). The value of the byte is 255 (eight 1's), but if you modify the value of the least significant bit (2^0), the new value is 254. If this value is used to represent the tint of a color or one of the values in a sound file, the change between 255 and 254 is so slight that it is indistinguishable.

| 2^7 | 2^6 | 2^5 | 2^4 | 2^3 | 2^2 | 2^1 | 2^0 | Decimal |
|-------|-------|-------|-------|-------|-------|-------|-------|---------|
| 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 255 |
| 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 254 |

Table 14-3 Modifying the Least Significant Bit

If someone modified the last bit of each byte in a 1MB file, he or she could embed a message in these 1 million bits. As an example, the letter A is represented in hexadecimal as 41 and binary as 0100 0001. A person could embed the letter A by ensuring the least significant bit in 8 bytes in a row includes these bits as the least significant bit in each of the 8 bytes.

[Table 14-4](#) shows the least significant bit from 8 bytes in a graphics file. To embed the letter A in these bytes, a steganography tool ensures that the last bit in these 8 bytes is 0100 0001. Notice that the first byte ends with a 1, so it is changed to a 0. The second byte ends in a 1, so it doesn't need to be changed. The steganography tool would change the last bit in bytes 1, 4, and 7, and leave the other bits the same. This shows how you can embed the letter A in a file. It's possible to repeat this process for a full message.

| | Orig 2 ⁰ | Letter A | Modify? | 2 ⁰ with Message |
|--------------------|---------------------|----------|---------|-----------------------------|
| Last bit in byte 1 | 1 | 0 | Yes | 0 |
| Last bit in byte 2 | 1 | 1 | No | 1 |
| Last bit in byte 3 | 0 | 0 | No | 0 |
| Last bit in byte 4 | 1 | 0 | Yes | 0 |
| Last bit in byte 5 | 0 | 0 | No | 0 |
| Last bit in byte 6 | 0 | 0 | No | 0 |
| Last bit in byte 7 | 1 | 0 | Yes | 0 |
| Last bit in byte 8 | 1 | 1 | No | 1 |

Table 14-4 Embedding the Letter A

The changes are not distinguishable to most users. However, when people know what to look for, they will be able to find and read the message. For example, it's possible for someone in one country to post a picture on a website with an embedded message and enable someone in another country to retrieve the picture and read the message.

While most people will overlook the hidden message, security professionals use steganalysis techniques to detect the use of steganography. A common method is the use of hashing. If you have a hash of the original file, you can create a hash from the modified version. If the two hashes are different, it indicates that the file has been modified and may have a hidden message embedded within it. Additional tools allow experts to analyze the bytes within the file to detect the message.



EXAM TIP Steganography attempts to hide information by embedding it into a file such as a picture or streamed audio or video file. Steganalysis techniques attempt to detect the use of steganography.

IPsec

[Chapter 3](#) introduces Internet Protocol security (IPsec) and explains how IPsec provides security for data traveling over a network. The two protocols within IPsec are Authentication Header (AH) and Encapsulating Security Protocol (ESP). AH provides authentication and integrity, while ESP provides confidentiality, authentication, and integrity. You can use AH alone, ESP alone, or both protocols together.

AH provides data origin authentication by assuring both systems of the identity of the other system. It creates a hash derived from packet data and uses this hash to provide integrity in addition to authentication. AH does not encrypt the data or provide confidentiality.

ESP encrypts the data within the packets to provide confidentiality. Although it's possible to configure ESP without integrity, it is not recommended. In other words, when you use ESP for encryption, you will also include the same features provided by AH (authentication and integrity).

[Chapter 4](#) presents IPsec as a tunneling protocol for VPNs. IPsec can be used with Layer 2 Tunneling Protocol (L2TP/IPsec) or by itself to create a secure connection between a VPN client and a VPN server.



TIP Both IPv4 and IPv6 support IPsec. IPv6 mandates the support of IPsec for any IPv6 applications. However, IPsec is an extension to IPv4, so there are compatibility problems in some situations. For example, IPsec traffic can't go through a system using Network Address Translation (NAT).

Understanding Public Key Infrastructure Systems

A public key infrastructure (PKI) includes all the components necessary to create, manage, distribute, validate, and revoke certificates. The “Exploring Asymmetric Encryption” section earlier in this chapter explained the use of public and private keys to encrypt information and digitally sign e-mails. The PKI provides the necessary support for these processes, including distributing public keys in digital certificates.

Although you aren't expected to create and manage a PKI for the SSCP exam, you should understand the basic components. This section covers certificates, certificate authorities (CAs), and certificate validation.



TIP A PKI is created based on the X.509 protocol. The X.509 standard identifies many of the components and formats used by the PKI, CAs, and certificates.

Certificates

The “Exploring Asymmetric Encryption” section earlier in this chapter presented some information on certificates. For example, [Figure 14-5](#) showed a certificate with the public key selected. Although the distribution of the public key is an important purpose of certificates, they do have many other uses. [Figure 14-11](#) shows a certificate from VeriSign with multiple purposes.



Figure 14-11 Certificate with multiple purposes

The following are some of the common purposes of certificates:

- **Authentication** Certificates can prove the identity of users and computers. For example, smart cards have embedded certificates used to prove the identity of the user. Remote systems can prove their identity to your system, and your system proves its identity to remote systems with the use of certificates.
- **Encryption** Certificates can be used to assist with encryption. For example, the public and private keys encrypt a symmetric key so that it can be privately shared.
- **Protecting e-mail** This protection includes both encryption and digital signatures.
- **Code signing** Applications and scripts can be digitally signed. This provides authentication of the developer and integrity for the code. If the code is modified (such as with malware), the signature embedded in the certificate will no longer match.

- **Wildcard** A wildcard certificate can be used for multiple domains that share a common root domain and start with an asterisk (*). As an example, a single wildcard certificate of *.getcertifiedgetahead.com can be used for blogs.getcertifiedgetahead.com and training.getcertifiedgetahead.com.
- **Domain validation** A domain-validated certificate indicates that the certificate requestor has some control over the DNS domain associated with the certificate. This control can be validated by responding to an e-mail or publishing a DNS record.
- **Extended validation** An extended validation certificate provides additional validation beyond a domain-validated certificate. Before issuing the certificate, the CA takes additional steps to verify the requesting entity's identity. Web browsers display the company name next to the URL or instead of the URL.



TIP As mentioned in [Chapter 5](#), attackers have been signing malware with code signing certificates, which they stole from others. Extended Validation (EV) Code Signing certificates help mitigate this problem. The Certificate Authority/Browser Forum (CABF) released standards for EV Code Signing certificates in 2012.

When entities want a certificate, they provide the required information to a CA in the form of a request. The CA embeds this information in the certificate.

[Figures 14-12](#), [14-13](#), and [14-14](#) show the different tabs of a certificate, giving you an idea of the information that is available. [Figure 14-14](#) shows that GlobalSign issued this certificate to www.gcgapremium.com.

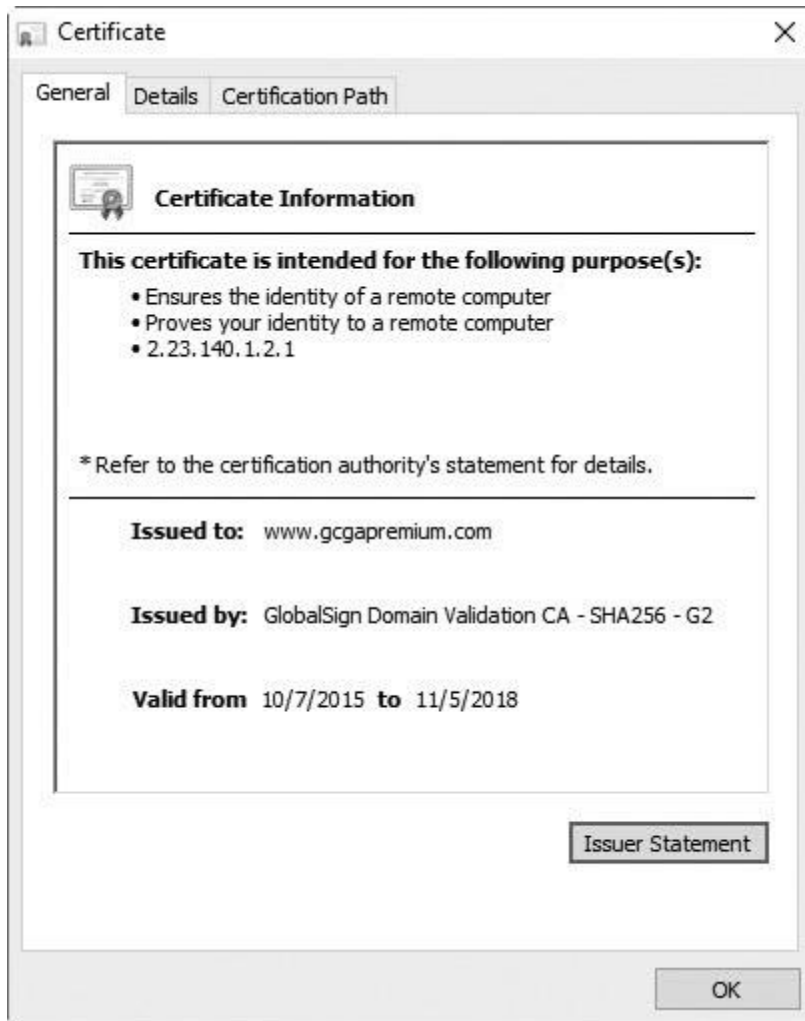


Figure 14-12 Viewing the General tab of a certificate

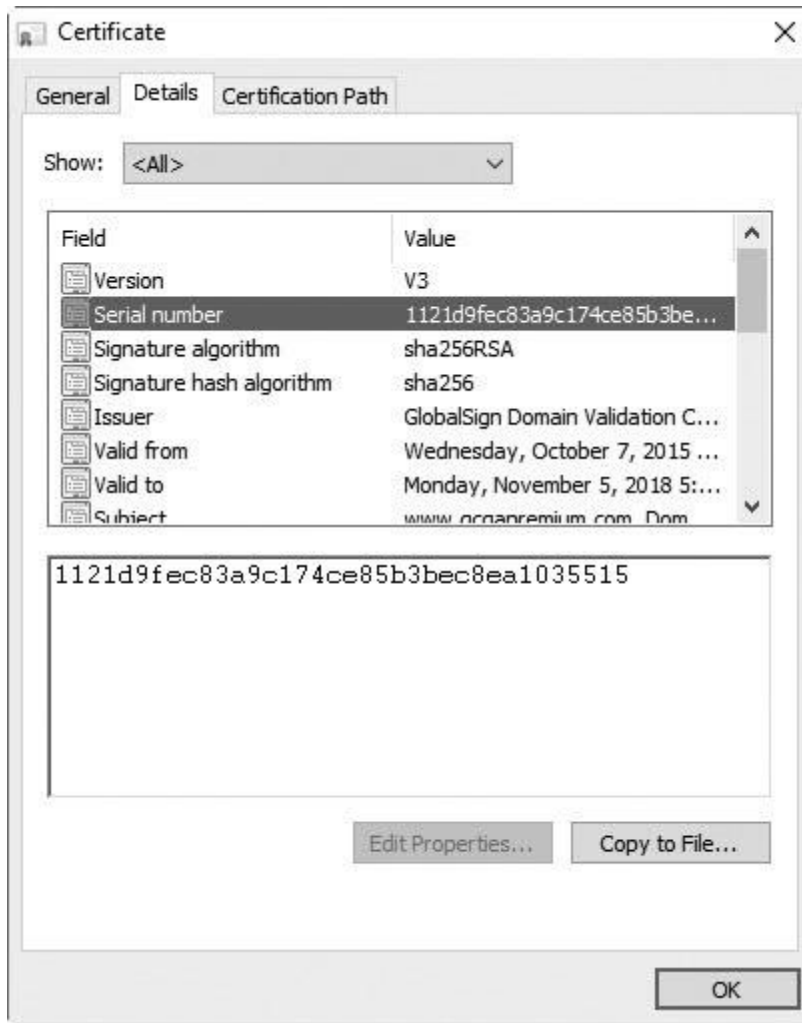


Figure 14-13 Viewing the Details tab of a certificate



Figure 14-14 Viewing the Certification Path tab of a certificate

Certificates include the following information:

- Who it was issued to (such as a website, a server, or a user)
- Who issued it (the CA)
- Validity dates (which allow the certificate to expire)
- Serial number (used to uniquely identify the certificate)
- Public key (matched to a private key that is not included in the certificate)
- A certification path showing the trust chain of the certificate to the root CA

Certificate Authority

A CA issues and manages certificates through their lifetimes. CAs can be public (such as VeriSign) or private within an organization. Public CAs sell and validate certificates and provide assurances to users that certificates are valid.



TIP A large organization can use a registration authority (RA). An RA accepts certificate requests, verifies the data, and passes the request on to the CA. The RA is optional and never issues certificates.

For example, if you wanted to do business on the Internet via your own website, you could purchase a certificate from a public CA to create a secure HTTPS session. Then when a customer wanted to make a purchase, your website would send the certificate to the customer's computer to initiate the secure session.



TIP The “Transport Layer Security” section earlier in this chapter describes the process of establishing an HTTPS session.

First, you would create a matching public key and private key. Many applications are available to create these matching keys. You would then send the public key and other required information to the CA. The CA would create the certificate and send it back to you, and you would then use it to configure your website.

It's also possible to create a private CA and issue certificates internally. The benefit is that the organization with a private CA doesn't need to purchase certificates from a public CA. However, the drawback is that certificates issued by a private CA aren't trusted by default. If external clients receive these certificates, they will receive an error indicating that the certificate is not trusted.

Understanding the Certificate Trust Chain

An important part of the PKI is the trust chain, also known as a hierarchy of trust. For example, if you look back to [Figure 14-14](#), you can see that the Certification Path tab shows two levels. GlobalSign is the top level and in this example is the root CA. The root CA issues certificates to other CAs in the trust chain, known as intermediate CAs.

In [Figure 14-14](#), GlobalSign Domain Validation CA SHA256 - G2 is an intermediate CA. The intermediate CA then issued a certificate to www.gcgapremium.com. Systems that trust the root

CA or the intermediate CA will automatically trust the certificate issued to www.gcgapremium.com.

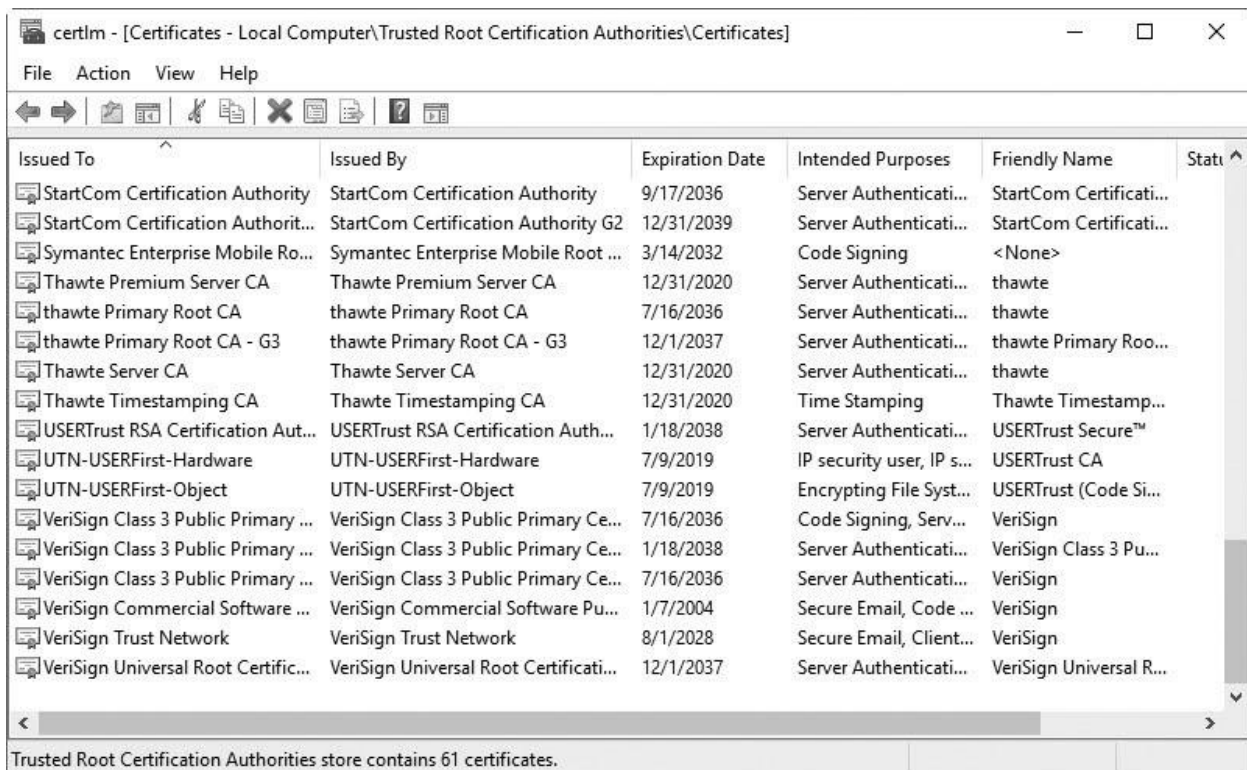


TIP It's possible to have multiple layers of intermediate CAs. In other words, a root CA can issue a certificate to an intermediate CA. The intermediate CA could then issue certificates to second-level intermediate CAs. This can continue for as many levels as the organization needs, although more than five are rarely necessary. The lowest-level CA issues certificates to others, such as websites or servers.

As long as one of the CAs in the certificate trust chain is trusted, certificates issued by the CAs are also trusted. For example, if the root CA is trusted, certificates issued by any of the CAs in the trust chain are also trusted.

Understanding Trusted Root Certification Authorities

You may be wondering how a CA is trusted. The primary method is by installing certificates from the root CA onto a system's trusted certificate authorities store. For example, [Figure 14-15](#) shows the Trusted Root Certification Authorities store for a Windows 7 system. Certificates in this store are from root CAs, and each certificate represents a certificate trust chain. Any certificate issued by any CA in this trust chain is automatically trusted.



The screenshot shows a Windows Explorer window titled "certlm - [Certificates - Local Computer\Trusted Root Certification Authorities\Certificates]". The window displays a list of certificates in the Trusted Root Certification Authorities store. The list has columns for Issued To, Issued By, Expiration Date, Intended Purposes, Friendly Name, and Status. The status for all certificates is "Trusted".

| Issued To | Issued By | Expiration Date | Intended Purposes | Friendly Name | Status |
|-------------------------------------|--|-----------------|---------------------------|-------------------------|---------|
| StartCom Certification Authority | StartCom Certification Authority | 9/17/2036 | Server Authenticati... | StartCom Certificati... | Trusted |
| StartCom Certification Authorit... | StartCom Certification Authority G2 | 12/31/2039 | Server Authenticati... | StartCom Certificati... | Trusted |
| Symantec Enterprise Mobile Ro... | Symantec Enterprise Mobile Root ... | 3/14/2032 | Code Signing | <None> | Trusted |
| Thawte Premium Server CA | Thawte Premium Server CA | 12/31/2020 | Server Authenticati... | thawte | Trusted |
| thawte Primary Root CA | thawte Primary Root CA | 7/16/2036 | Server Authenticati... | thawte | Trusted |
| thawte Primary Root CA - G3 | thawte Primary Root CA - G3 | 12/1/2037 | Server Authenticati... | thawte Primary Roo... | Trusted |
| Thawte Server CA | Thawte Server CA | 12/31/2020 | Server Authenticati... | thawte | Trusted |
| Thawte Timestamping CA | Thawte Timestamping CA | 12/31/2020 | Time Stamping | Thawte Timestamp... | Trusted |
| USERTrust RSA Certification Aut... | USERTrust RSA Certification Auth... | 1/18/2038 | Server Authenticati... | USERTrust Secure™ | Trusted |
| UTN-USERFirst-Hardware | UTN-USERFirst-Hardware | 7/9/2019 | IP security user, IP s... | USERTrust CA | Trusted |
| UTN-USERFirst-Object | UTN-USERFirst-Object | 7/9/2019 | Encrypting File Syst... | USERTrust (Code Si... | Trusted |
| VeriSign Class 3 Public Primary ... | VeriSign Class 3 Public Primary Ce... | 7/16/2036 | Code Signing, Serv... | VeriSign | Trusted |
| VeriSign Class 3 Public Primary ... | VeriSign Class 3 Public Primary Ce... | 1/18/2038 | Server Authenticati... | VeriSign Class 3 Pu... | Trusted |
| VeriSign Class 3 Public Primary ... | VeriSign Class 3 Public Primary Ce... | 7/16/2036 | Server Authenticati... | VeriSign | Trusted |
| VeriSign Commercial Software ... | VeriSign Commercial Software Pu... | 1/7/2004 | Secure Email, Code ... | VeriSign | Trusted |
| VeriSign Trust Network | VeriSign Trust Network | 8/1/2028 | Secure Email, Client... | VeriSign | Trusted |
| VeriSign Universal Root Certific... | VeriSign Universal Root Certificati... | 12/1/2037 | Server Authenticati... | VeriSign Universal R... | Trusted |

Trusted Root Certification Authorities store contains 61 certificates.

Figure 14-15 Trusted root certification authorities store

This is similar in concept to a driver's license issued by a Department of Motor Vehicles (DMV). If you need to prove your name or age, you can show your driver's license. Because people trust that the DMV has verified your information before giving you a driver's license, and your driver's license was issued by the DMV, most people trust the validity of your driver's license. On the other hand, if you bought an identification card from Darril's Cheap IDs, many people may question the validity of information on that ID card.

The root CA is similar to the DMV. Any certificates issued by the root CA are trusted because the root CA is trusted, just as licenses issued by the DMV are trusted because the DMV is trusted.



TIP There are other ways you can add certificates to a system so that they are trusted. However, the point here is that certificates will automatically be trusted if they're in a trust chain. E-commerce sites purchase certificates from CAs that are in the trust chain so that their certificates will automatically be trusted by customer computers.

Revoking Certificates

Most certificates will have a lifetime identified in the validity dates. However, in some situations the certificate is revoked before it expires. For example, the CA will revoke the certificate if the matching private key of a certificate is compromised. The CA maintains a list of all revoked certificates in a certification revocation list (CRL) and publishes the CRL as a V2 certificate. Most other certificates are V3 certificates. The CRL lists the serial numbers of certificates revoked by the CA.

Certificates include a list of CRL distribution points. Clients can request a copy of the CRL through these distribution points and verify that the certificate they are using is not listed on the CRL. This verification is often done as part of the normal process of validating a certificate.

Validating Certificates

When systems receive certificates from a website, they go through several checks to ensure that the certificate is valid. [Figure 14-16](#) shows the overall steps that occur, which are as follows:

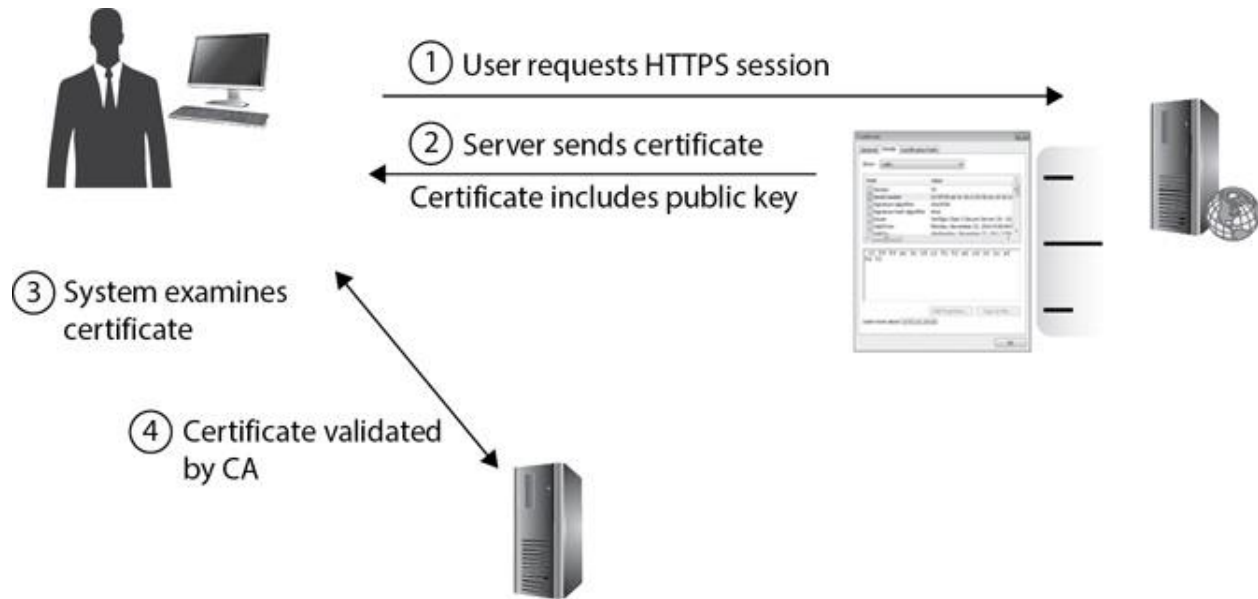


Figure 14-16 Validating a certificate

1. The user requests the HTTPS session.
2. The server responds with the certificate.
3. The user's system examines the certificate. It ensures that the certificate hasn't expired (based on the validity dates) and ensures that the website that issued the certificate is the same as the one identified in the certificate. It will also ensure that a CA in the certificate trust chain is trusted.
4. The user's system checks with the CA to ensure that the certificate has not been revoked. A common method is by requesting a copy of the CRL through a CRL distribution point. If the CRL includes the serial number of the certificate provided by the website, the client system knows the CA has revoked the certificate and it shouldn't be trusted.

Another method of checking to see whether the certificate is revoked is with the Online Certificate Status Protocol (OCSP). The client system first identifies the serial number of the certificate received from the website. Then, instead of requesting a copy of the CRL, the client sends the serial number to a server known as an OCSP responder (maintained by the CA). The OCSP responder identifies the health of the certificate with an answer of "good," "revoked," or "unknown." A response of "unknown" might indicate the certificate is a forgery. The benefit of OCSP over a CRL is that the entire CRL doesn't need to be transmitted.

If any of these checks fails, the user will see an error indicating what check failed. The errors are different based on the browser that the user is using.

Key Escrow

Key escrow refers to storing a copy of a private key for safekeeping or creating an alternative key to be used if necessary. The key can be retrieved if the original key is lost or unavailable. As an example, if an employee encrypts data and then leaves the organization, the data may not be accessible if the organization hasn't enabled a key escrow method.

Recovery agents are the individuals authorized to recover keys held in escrow. Typically, keys held in escrow are kept within the organization. However, there are some situations where the keys are held in escrow by a third party. If the keys are kept within the organization, the recovery agent is someone within the organization. If the keys are held by a third party, the recovery agent is typically someone in the third-party organization.

Alternative Certificate Trusts

The “Understanding the Certificate Trust Chain” section described a hierarchical trust chain, which provides centralized certificate management. A hierarchical trust chain allows CAs to issue certificates and for entities to verify the certificates are valid. For example, imagine a client trusts a root CA. The client will automatically trust certificates from the root CA or any intermediate CAs in the trust chain. Additionally, the client can validate the certificate by either retrieving the CRL or using OCSP. However, this isn't the only type of trust.

Two additional trust methods are a web of trust and a cross-certification trust.

Web of Trust

A web of trust (WOT) is a decentralized trust model allowing anyone to create self-signed certificates. Because a person who creates a self-signed certificate also signs the certificate, it isn't possible to easily validate its authenticity. A web of trust provides one way to solve this problem and is used with PGP and GPG (described earlier in this chapter) for encryption and digital signatures.

As a simple example, imagine that Lisa, Homer, Marge, and Bart each create their own self-signed certificates to encrypt e-mails and create digital signatures. They exchange these certificates with each other and ask everyone in their group to trust them. This is a web of trust.

Later, Lisa and Krusty exchange certificates and start exchanging secure e-mails. Lisa chooses to trust Krusty's certificate and could introduce Krusty's certificate to the group, encouraging everyone to trust it. If Bart (or anyone else in the web of trust) sends a certificate to Krusty, Krusty might not trust it because Bart is unknown to him. However, because Lisa trusts it, this might be enough for Krusty to trust it.

In the early days of PGP, its biggest challenge was the distribution of certificates. Phil Zimmermann, the creator of PGP, recognized that there wasn't a reliable method of proving identities. He introduced the web of trust concept in an early manual for PGP. Admittedly, the details are more technical than described in the example with Lisa's group, but the concept is

similar. Designated individuals are authorized to introduce certificates to the web of trust. If enough people trust a certificate, it provides a level of assurance. For example, if six of your friends trust a certificate, you might choose to trust it.



TIP PGP, GPG, and OpenPGP support a web of trust. Current implementations of PGP can also use a public key infrastructure to retrieve and validate certificates.

With OpenPGP, users can digitally sign other OpenPGP identity certificates to endorse them. As an example, if you know me, you can digitally sign my certificate to endorse it. Similarly, if I know you, I can digitally sign your certificate to endorse it. When a certificate has enough endorsements, other users can choose to trust it. In some cases, merely having the endorsement of a key individual may be enough to trust it. For example, if you are well-known in cybersecurity circles and you have digitally signed my certificate, someone else may choose to trust my certificate based on your endorsement.

Cross-Certification Trust

Cross-certification allows organizations with a private PKI and private CA to trust entities in another private PKI. As an example, imagine that Globex Corporation has created its own PKI to issue and manage certificates for Globex employees. Similarly, imagine that the Springfield University has created a private PKI for its students, staff, and faculty. By default, these certificates are not trusted outside the organization. If an employee at Globex sends a certificate to someone at Springfield University, it creates errors because the certificate isn't trusted and cannot be validated.

However, if the two organizations create a cross-certification agreement, it creates a mutual trust relationship. Users in both organizations automatically trust the certificates from the other organization. A cross-certification agreement will also implement methods allowing systems to validate the certificates.

Comparing Cryptanalysis Attacks

Cryptanalysis is the process of deciphering codes through analysis to compromise the confidentiality or integrity of data. Authorized users of encrypted data have the tools to decrypt and access data when needed. However, unauthorized users need to use cryptanalysis to decrypt and access the data.



EXAM TIP Both attackers and researchers use cryptanalysis techniques. Researchers use cryptanalysis techniques to search for weaknesses in cryptography with the goal of improving cryptographic methods. Attackers search for weaknesses with the goal of exploiting them.

Most cryptographic algorithms are published, allowing anyone to study them. Security experts often perform cryptanalysis on protocols to identify weaknesses based on new data available. As an example, SSL was thought to be secure for many years. However, a security team at Google discovered a significant vulnerability and published a paper on the POODLE attack in late 2014. This prompted many organizations to switch over to TLS sooner than they might have otherwise done. Similarly, the federal Canadian Cyber Incident Response Centre issued a security bulletin about the Heartbleed bug in April 2014 that affected many sites using specific versions of OpenSSL.

If vulnerabilities don't exist in the algorithm, attackers often look for weak keys or poorly implemented keys. For example, one of the core rules related to stream ciphers is to never reuse keys. If a protocol ignores this rule, the protocol is easier to crack. This emphasizes the importance of following fundamental key-management concepts.

Fundamental Key Management Concepts

While SSCPs are not expected to design encryption algorithms, (ISC)² does expect you to understand some of the underlying concepts, especially related to the use of keys. Some of these concepts were covered in more depth earlier in this chapter but are included here as a summary.

- **Key rotation** Key rotation refers to periodically replacing an old key with a new key. This is similar in concept to changing a password regularly. The length of time between using a different key varies for different algorithms. However, the important point here is that keys should be changed regularly.
- **Key composition** Key composition refers to creating strong keys. Strong keys are random and of sufficient length that they cannot be guessed. The length of the key varies for different algorithms. For example, AES-256 uses 256-bit keys and is very strong. However, the use of 256-bit keys for RSA is very weak. Instead, strong RSA keys are typically 2048 bits in length.
- **Key creation** Keys should be random. Most algorithms include code to ensure keys are random.
- **Key exchange** Keys should be exchanged in a secure manner. It's common to exchange symmetric keys using asymmetric cryptography.
- **Key revocation** The cryptography method should include a method of revoking compromised keys and informing users when a key has been revoked. Within a PKI, users can verify if a key is valid using a CRL or with OCSP.
- **Key escrow** Key escrow ensures that encrypted data can be retrieved even if the original key is somehow lost.

Known-Plaintext Attack

A known-plaintext (or known-text) attack is possible when the attacker has samples of both plaintext and ciphertext data. For example, if an attacker has plaintext from an encrypted message, the attacker can then use different methods to try to decrypt the ciphertext to the known text. If the known text is decrypted, the same method can be used to decrypt similar data.

The chosen-plaintext attack is a variant of the known-plaintext attack. In this attack, the cryptanalyst doesn't have access to all of the plaintext but either has access to some plaintext or is able to predict plaintext that is included in the ciphertext. For example, if messages are frequently preceded with the phrase "From:" then the analyst can assume that "From:" is the known plaintext and use this information in the attack.

Ciphertext-Only Attack

In contrast, a ciphertext-only attack is one where the attacker has only ciphertext for analysis without any useful information about the plaintext data. If the attacker can discover any of the plaintext by analyzing the ciphertext, the attack is successful.

This type of attack is usually successful only when weak cryptographic methods are used. For example, if the same key is used for too long a period, or the key isn't protected, or a weak key is used, it makes it much easier for an attacker to break the code. WEP was cracked using a ciphertext-only attack due to poor key-management techniques used with WEP (among other vulnerabilities).

Hashing and Collisions

A hash collision occurs when two different inputs to a hashing function produce the same hash. A hash collision is not possible in a sufficiently strong hashing algorithm. However, as computing power has grown, methods have been identified that create collisions with algorithms previously thought to be secure.

As an example, security experts have successfully used hash collision attacks on MD5 and SHA-1. This is a significant vulnerability and one of the reasons these two hashing functions are no longer recommended for use.

Researchers identified collisions with MD5 in 1996. In 2005, researchers were able to create pairs of documents and pairs of certificates with the same hash. In 2012, Microsoft reported that attackers used an MD5 collision to forge a Windows code-signing certificate.

SHA-1 was thought to be vulnerable to collisions since 2005, though researchers didn't successfully create a collision until 2017. They spent two years developing the technique and used multiple computers and processors working together performing over 9 quintillion computations (about 1 million trillion computations). It would equate to about 6,500 years of CPU time if a single computer were used. A well-funded attacker could use multiple computers each with multiple cores, to shorten the time required to create a SHA-1 collision.

Countermeasures Against Cryptanalysis Attacks

The primary protection against cryptanalysis attacks is to use strong algorithms and follow fundamental key management practices. Both are moving targets.

Strong algorithms from 2001 are not necessarily strong algorithms in 2018. Similarly, strong algorithms in 2018 won't necessarily be strong in 2030. When security experts identify vulnerabilities in cryptographic algorithms and recommend taking them out of use, IT personnel should do so as soon as possible.

The “Fundamental Key Management Concepts” section in this chapter covered the importance of using strong keys. It also covered management concepts such as rotating keys regularly, exchanging them securely, and ensuring there are methods to revoke them when they are compromised.

Chapter Review

Cryptography provides several important security benefits. You can encrypt data to prevent the loss of confidentiality, and you can use hashing algorithms to help ensure integrity. Many encryption and hashing protocols also provide authenticity, providing assurances that data is coming from a known source and is valid or reliable.

Organizations use cryptographic methods to secure sensitive data. In general, all nonpublic data should be encrypted to help ensure its protection. Additionally, regulations require organizations to protect certain data, such as personally identifiable information (PII) and protected health information (PHI). Encryption is a primary method of doing so.

Hashing provides one-way encryption to create a fixed-length hash (a number) from a file or message of any size. The created hash will always be the same as long as the file or message has not been modified. By comparing hashes at different times (such as before and after downloading a file), you can tell whether the original data has been modified or has lost integrity. If the hashes are the same, the data has not lost integrity. Message Digest 5 (MD5) and Secure Hashing Algorithm 1 (SHA-1) have been popular hashing algorithms, but both have been cracked and should not be used. NIST recommends the use of SHA-3, which can create hashes of 224, 256, 384, or 512 bits.

Encryption converts plaintext data into ciphertext data. The two primary categories of encryption are symmetric encryption and asymmetric encryption.

Symmetric encryption uses a single key to encrypt and decrypt data. Keys used with symmetric encryption need to be protected and changed often. If the key is not protected, an attacker can use it to decrypt the data.

Advanced Encryption Standard (AES) is a strong, efficient symmetric encryption algorithm. NIST identified it in 2002 as a standard to be used by the U.S. government agencies, and it has been adopted in many different commercial applications. AES can use key sizes of 128 bits

(AES or AES-128), 192 bits (AES-192), or 256 bits (AES-256). AES is the designated replacement for Data Encryption Standard (DES). DES has been cracked and is not recommended for use today. Other symmetric encryption algorithms include 3DES, Blowfish, International Data Encryption Algorithm (IDEA), and RC4 (though RC4 is not recommended for use today).

Asymmetric encryption uses two keys (a public key and a private key) created as matched pairs. Information encrypted by one of the keys can only be decrypted with the matching key. The public key is embedded in a certificate and freely shared, but the private key always remains private. RSA is used for asymmetric encryption and employs large prime numbers to create matching public and private keys. Transport Layer Security (TLS) uses asymmetric encryption to share a symmetric key privately, and then it uses symmetric encryption to encrypt session data.

Secure/Multipurpose Internet Mail Extensions (S/MIME) is a popular standard used to protect e-mail. It supports encryption and digital signatures. Digital signatures are created by hashing an e-mail message and then encrypting the hash with the sender's private key. The recipient uses the sender's public key to decrypt the hash, and if the decryption succeeds, the digital signature provides authentication, integrity, and nonrepudiation. DomainKeys Identified Mail (DKIM) helps detect e-mail spoofing by verifying the e-mail was sent by an authorized mail server.

Two other encryption schemes include steganography and Internet Protocol security (IPsec). Steganography is the practice of hiding data within data. IPsec is used to secure data going over a network and uses Authentication Header (AH) and Encapsulating Security Protocol (ESP). AH provides authentication and integrity, and ESP provides authentication, integrity, and confidentiality.

A public key infrastructure (PKI) includes all the components needed to issue and manage certificates, which follow the X.509 standard. The certificate authority (CA) issues the certificates. Certificates have expiration dates, but the CA can revoke the certificate earlier if necessary, such as if the matching private key is compromised. The CA publishes a list of revoked certificates by serial number in a certificate revocation list (CRL). Clients can request a copy of the CRL to validate a certificate. The Online Certificate Status Protocol (OCSP) is an alternative method used to validate certificates. The first CA in a certificate trust chain is called the root CA. If the root CA is trusted, all certificates issued by any CAs in the trust path are also trusted.

S/MIME depends on a centralized PKI for certificates. Pretty Good Privacy (PGP), GNU Privacy Guard (GPG), and OpenPGP are alternatives to S/MIME and can use a decentralized web of trust (WOT). A WOT allows individuals to create their own self-signed certificates and share them. When enough people endorse a certificate, it indicates a high level of trust that it is valid.

Cryptanalysis is the process of deciphering codes through analysis. Two common cryptanalysis attacks are a known-plaintext attack and a ciphertext-only attack. A known-plaintext (or known-text) attack is possible when the attacker has samples of both plaintext and ciphertext data. A ciphertext-only attack is one where the attacker has only ciphertext for analysis.

A hash collision occurs when two different inputs to a hashing function produce the same hash. When collisions are possible, it indicates the hashing function is not secure and should not be used.

Common countermeasures against cryptanalysis attacks start with using known strong algorithms. When flaws are found, organizations should replace older algorithms with recommended replacements. Additionally, organizations should ensure that they are implementing common key management concepts. This includes using strong keys, rotating them as recommended, exchanging them securely, and revoking them when they are compromised.

Questions

1. What basic security function does asymmetric encryption provide?

- A. Integrity
- B. Authentication
- C. Confidentiality
- D. Availability

2. Which of the following choices allows you to verify that a file has not been modified?

- A. AES
- B. SHA
- C. PKI
- D. IDEA

3. Which of the following choices provides one-way encryption of data?

- A. Symmetric
- B. Asymmetric
- C. Transport Layer Security
- D. Hashing

4. Which of the following uses a single key to encrypt and decrypt data?

- A. Symmetric

B. Asymmetric

C. Public key cryptography

D. SHA-1

5. Which of the following keys is changed the most often?

A. Public key

B. Private key

C. Symmetric key

D. Session key

6. Which of the following are used primarily for symmetric encryption? (Select all that apply.)

A. AES

B. Blowfish

C. 3DES

D. RSA

7. Which of the following is a symmetric 128-bit block cipher?

A. Data Encryption Standard (DES)

B. Triple Data Encryption Standard (3DES)

C. Advanced Encryption Standard (AES)

D. Blowfish

8. Of the following choices, which one is considered a strong, efficient symmetric encryption algorithm?

A. TLS

B. DES

C. 3DES

D. AES

9. What type of cryptography does public cryptography use?

- A.** Asymmetric encryption
- B.** Symmetric encryption
- C.** Steganography
- D.** One-way functions

10. Which of the following is an accurate statement related to asymmetric encryption?

- A.** It is used to privately share a private key.
- B.** It is used to privately share a public key.
- C.** It is used to privately share a secret key.
- D.** It is faster than symmetric encryption.

11. What is a common standard used to encrypt and digitally sign e-mail?

- A.** Symmetric encryption
- B.** S/MIME
- C.** TLS
- D.** Steganography

12. Of the following choices, what does a digital signature provide when it is used with e-mail?
(Select all that apply.)

- A.** Authentication
- B.** Integrity
- C.** Confidentiality
- D.** Nonrepudiation

13. What is used to create a digital signature used with e-mail?

- A.** The public key of the sender
- B.** The private key of the sender

C. The public key of the recipient

D. The private key of the recipient

14. Which of the following would most likely be used to encrypt data in an e-mail message before it is sent?

A. The public key of the sender

B. The private key of the sender

C. The public key of the recipient

D. The private key of the recipient

15. Someone has embedded a secret code within a picture used on a web page. What is the best description of this?

A. Symmetric encryption

B. Asymmetric encryption

C. Hashing

D. Steganography

16. How are public keys distributed to clients from Internet websites?

A. As e-mail attachments

B. Embedded in certificates

C. As cookies

D. Embedded in the HTML code for the page

17. How are public keys shared with other entities?

A. Published in a certificate

B. Encrypted by a private key

C. Encrypted by a session key

D. Public keys are not shared

18. A website sent a user a certificate to initiate a secure web session over the Internet. Of the following choices, what information would be in the certificate? (Select all that apply.)

- A. Name of the website
- B. Name of the issuing CA
- C. Private key
- D. Expiration date

19. Of the following choices, what is used to determine whether a certificate has been revoked?

- A. OCSP
- B. Digital signature
- C. CARL
- D. Trust chain

20. Researchers are attempting to discover weaknesses in an encryption algorithm using a known-plaintext attack. What is this called?

- A. Cryptography
- B. Cryptanalysis
- C. Criminal behavior
- D. Hashing

Answers

1. C. Encryption (any type of encryption, including both asymmetric and symmetric encryption) provides confidentiality for data. Hashing methods provide integrity. Authentication proves the identity of a user or system. Availability ensures that systems and data are available when needed.

2. B. Secure Hashing Algorithm (SHA) is a hashing algorithm, and hashing is a key method of ensuring integrity (or verifying a file has not been modified). SHA-3 should be used instead of SHA-1 and SHA-2. The hash is calculated at two different times, and if the hash is the same, the file has not been modified. Advanced Encryption Standard (AES) is a strong symmetric encryption protocol. A public key infrastructure (PKI) is used to support the creation, management, and distribution of certificates. International Data Encryption Algorithm (IDEA) is an older symmetric encryption protocol.

3. D. Hashing algorithms use a one-way encryption method to create a hash from the data. They are also known as one-way functions. The hashing algorithm creates a fixed-length hash from a file, but the hash cannot be used to re-create the file. Symmetric encryption and asymmetric encryption are both two-way, because data can be encrypted and decrypted. Transport Layer Security (TLS) is widely used to encrypt Internet traffic and uses both symmetric and asymmetric encryption.

4. A. Symmetric encryption uses a single key to encrypt and decrypt data. Asymmetric encryption uses two keys (a public key and a private key) to encrypt and decrypt information and is often referred to as public key cryptography. Secure Hashing Algorithm 1 (SHA-1) is a hashing algorithm and it doesn't use a key.

5. D. A session key is only used for a session (such as a web browsing session) and is changed more often than the other keys. Public and private keys typically last for a year or longer. Symmetric encryption uses a symmetric key (also called a secret key), which can stay the same for a specific piece of data as long as the data remains encrypted. While a session key is a symmetric key, all symmetric keys are not session keys.

6. A, B, C. Advanced Encryption Standard (AES), Blowfish, and Triple DES (3DES) are all used for symmetric encryption. RSA is an asymmetric encryption standard using public and private keys and is widely used with Transport Layer Security (TLS). The other choices are all symmetric encryption standards using a single key to encrypt and decrypt the data.

7. C. AES is a 128-bit block cipher. All of the other answers are 64-bit block ciphers.

8. D. The Advanced Encryption Standard (AES) is considered a strong, efficient symmetric encryption algorithm and it is widely used. DES is an older algorithm that has been cracked. 3DES is strong but takes more processing power and is less efficient than AES. Transport Layer Security (TLS) uses both symmetric and asymmetric encryption, and calling it a symmetric encryption algorithm is inaccurate.

9. A. Public key cryptography uses asymmetric encryption with two matched keys (a public key and a private key) to encrypt and decrypt information. Symmetric encryption uses a single key (often called a session key) to encrypt and decrypt data. Steganography hides data within data. Hashes are also known as one-way functions and they provide integrity.

10. C. Asymmetric encryption is used to privately share a secret key (or session key). Asymmetric encryption uses a matched pair of keys known as a private key and a public key. The private key is never shared, and the public key is publicly shared in a certificate. Symmetric encryption is faster than asymmetric encryption.

11. B. Secure/Multipurpose Internet Mail Extensions (S/MIME) is the standard used to encrypt and digitally sign e-mail. Symmetric encryption uses a single key to encrypt and decrypt data but cannot digitally sign e-mail. Transport Layer Security (TLS) encrypts data sent over a network and is used with HTTPS. Steganography is the practice of hiding data within data or in plain sight.

12. A, B, D. A digital signature provides authentication, integrity, and nonrepudiation. It does not provide confidentiality because the digital signature does not encrypt the data. It is possible to digitally sign an e-mail without encrypting it.

13. B. A digital signature is created by hashing a message and encrypting the hash with the sender's private key. The recipient can then decrypt the hash with the sender's public key. The recipient's keys are not used for a digital signature, but they are used to encrypt and decrypt e-mail.

14. C. E-mail is encrypted using the recipient's public key. The recipient's public key actually encrypts a symmetric key and uses the symmetric key to encrypt the e-mail. The recipient uses the recipient's private key to decrypt the symmetric key and then decrypts the message with the symmetric key. The sender's keys are not used to encrypt or decrypt e-mail.

15. D. Steganography is the practice of hiding data within data, such as embedding a secret code within a picture. Symmetric encryption uses a single key for encryption and decryption of data, while asymmetric encryption uses two keys (a public key and a private key) for encryption and decryption. Hashing creates a hash that can be used for integrity.

16. B. Public keys are embedded in certificates and distributed to clients in the certificate. Although users can send certificates to each other as e-mail attachments, a website does not use this method. Public keys are not included in cookies or in HTML code.

17. A. Public keys are published in certificates. They are never encrypted. They are shared so that other entities can use them for asymmetric encryption.

18. A, B, D. The certificate would include the name of the website, the name of the CA that issued the certificate, the expiration date of the certificate, and more. While the certificate would include a public key, it would not include the private key. The private key would be kept private on the server.

19. A. The Online Certificate Status Protocol (OCSP) is used to verify the health of a certificate. An OCSP responder will indicate whether a certificate has been revoked when queried with the certificate's serial number. A digital signature uses certificates but doesn't determine whether a certificate is revoked. CAs issue a certificate revocation list (CRL), but CRL isn't a valid acronym in the context of checking certificates. A trust chain determines if the CA that issued the certificate is trusted but it doesn't indicate if a certificate is revoked.

20. B. Cryptanalysis is the process of deciphering codes through analysis, and a known-plaintext attack is one method of cryptanalysis. Cryptography is the science of using different methods and techniques to encrypt data. It is not criminal to search for weaknesses, but the action taken after these weaknesses are discovered can be criminal. Hashing is the process of creating a hash from a file or a message with a hashing algorithm, and it is used to prevent the loss of integrity.

Find answers on the fly, or master something new. Subscribe today. [See pricing options.](#)

- [Settings](#)
- [Support](#)
- [Sign Out](#)

© 2020 [O'Reilly Media, Inc.](#) [Terms of Service](#) / [Privacy Policy](#)

-
-
-

[Skip to content](#)

-

- [Home](#)

- •

- [Your O'Reilly](#)

- [Profile](#)
- [History](#)
- [Playlists](#)
- [Highlights](#)

-

- [Featured](#)

- [Resource Centers](#)
- [Expert Playlists](#)

-

- [Practice](#)

- [Katacoda Scenarios](#)
- [Jupyter Notebooks](#)

-

- [Sandboxes](#)

- [Kubernetes](#)
- [Python](#)
- [TensorFlow](#)
- [Ubuntu](#)

-

- [Explore](#)

- [All Topics](#)
- [Most Popular Titles](#)
- [Recommendations](#)
- [Early Release](#)
- [Playlists](#)

-

- [Attend](#)

- [Live Training](#)
-
- [Certifications](#)
-
- [Newsletters](#)
-
- [Settings](#)
-
- [Support](#)
-
-
- [Sign Out](#)

Table of Contents for

SSCP Systems Security Certified Practitioner All-in-One Exam Guide, Third Edition, 3rd Edition

- [Search in book...](#)
-
- [Toggle Font Controls](#)
 -
 -
 -
 -
 -

[Prev Previous Chapter](#)

[Appendix About the Online Content](#)

[Next Next Chapter](#)

[Index](#)

GLOSSARY

3DES A symmetric encryption standard. It improves Data Encryption Standard (DES) by encrypting data in three passes with separate keys. It was one of the standards evaluated by the National Institute of Standards and Technology (NIST) with Advanced Encryption Standard (AES), but was not selected. Compared to AES, it is a slower and processor-intensive block cipher but is still strong and used in some applications. Also called *triple DES* and *three DES*.

3G Third generation of wireless technologies. 1G uses analog signals, and 2G uses digital signals. 3G provides higher transfer speeds using digital technologies.

4G Fourth generation of wireless technologies. It's an improvement over 3G and includes speeds as high as 100 Mbps for highly mobile users in trains and cars, and 1 Gbps for individuals walking or staying still.

802.1x An authentication standard used in some wireless and virtual private networks (VPNs). It uses an authentication server and requires users to provide valid credentials before accessing a wireless network or accessing the VPN resources.

AAAs of security Authentication, authorization, and accounting. Authentication ensures that entities can prove their identity. Authorization grants access to individuals based on their proven identity. Accounting tracks and records their activity in logs.

access control Mechanism used to restrict or control access to resources. Access controls can be logical (implemented with technology) and implemented by the security kernel, or physical, such as a locked door or security guard. Access controls allow subjects access to objects, such as allowing a user to access a file. Some relevant access control models are Mandatory Access Control (MAC), Discretionary Access Control (DAC), Role-based Access Control (RBAC or Role-BAC), Rule-based Access Control, and Attribute-based Access Control (ABAC). Biba, Bell-LaPadula, Clark-Wilson, and the Chinese Wall are specific MAC models.

access control list (ACL) A list of rules. ACLs are most commonly associated with routers and firewalls. The rules identify the traffic allowed in or out of a network. On host-based firewalls, the rules identify the traffic allowed in or out of a system.

access control matrix A list of objects along with the permissions granted for each object. You can think of an access control matrix as a group of ACLs. Each ACL represents a single object and lists all the permissions for that object. Compare to *capability table*.

accountability The ability of a system to track the activity of an individual. It depends on proper identification and authentication. If a system can identify individual users, track their actions, and monitor their behavior, it provides accountability.

accounting Logs that track the activity of a user through monitoring. One method of accounting is audit logs that create an audit trail.

accreditation The process of formally declaring that a system is approved to operate. Accreditation comes after system certification.

Address Resolution Protocol (ARP) Protocol that resolves Internet Protocol (IP) addresses to Media Access Control (MAC) addresses. ARP is used on the Data Link layer of the Open Systems Interconnection (OSI) Model. Compare to *Network Discovery Protocol (NDP)*, *Bootstrap Protocol (BootP)*, and *Reverse Address Resolution Protocol (RARP)*.

administrative controls Security controls that focus on the management of risk and the management of IT security. They include security policies, standards, procedures, guidelines, baselines, and other administrative practices that an organization uses to implement security. Compare to technical controls and physical security controls.

Advanced Encryption Standard (AES) A strong, efficient symmetric encryption algorithm. The National Institute of Standards and Technology (NIST) selected it in 2002 as a replacement for Data Encryption Standard (DES) as the standard used by the U.S. government. AES has since been adopted in both the public and private sectors and is widely used today in many applications. AES uses key sizes of 129, 192, and 256 bits and sometimes is listed based on the key size (such as AES-256).

advanced persistent threat (APT) A highly sophisticated group of attackers who have the capability and intent to carry out complex targeted attacks. Many governments are suspected of sponsoring groups known as APTs.

anomaly based A method of detection used by intrusion detection systems (IDSs) and intrusion prevention systems (IPSs). The IDS/IPS attempts to document normal behavior in the form of a baseline. It then monitors the activity and constantly compares it to the baseline. If the current activity differs significantly from the baseline, the IDS/IPS will issue an alert on the activity. It is sometimes called *behavior-based detection*.

antivirus (AV) software A primary method used to detect and prevent infections from malware. In addition to being able to detect and prevent infections, most AV software is able to remove the malware, restore the infected file to its original state, or quarantine the file.

armored virus A virus that uses code to make it difficult for AV researchers to reverse-engineer the code. Encryption is often combined with other methods to prevent reverse-engineering.

asymmetric encryption A process of encrypting and decrypting data using two matched keys known as a public key and a private key. It is also known as *public key cryptography*. Anything encrypted with the public key can be decrypted only with the matching private key. Anything encrypted with the private key can be decrypted only with the matching public key. The private key is always kept private and never shared. The public key is freely shared and publicly available.

attribute In a database, a table column. Data within the table is stored in rows, or tuples.

Attribute-based Access Control (ABAC) An access control model that uses attributes to determine access. It evaluates subject and object attributes and grants access based on the value of these attributes. Attributes can be almost any characteristic of a user, the environment, or the resource. Other access control models are Mandatory Access Control (MAC), Discretionary Access Control (DAC), Role-based Access Control (RBAC or Role-BAC), and Rule-based Access Control.

audit Inspection of activity and configurations within an organization to verify that personnel are following required policies and procedures. An audit can also be done by reviewing logs to reconstruct past events and create an audit trail.

audit trail A record of events occurring on a system or network, recorded in one or more logs. When you have access to all the logs, you can re-create the events that occurred leading up to an event and identify what occurred during an event.

authentication Proof of an identity, established by providing credentials. There are three types or factors of authentication: something you know (such as a username and password), something you have (such as a smart card), and something you are (using biometrics).

Authentication Header (AH) Used with Internet Protocol security (IPsec) to provide authentication between systems and integrity of the packets. The AH is created from a hash of the packet with additional authentication data, and the hash is encrypted to prevent tampering. The hash provides integrity. Packets protected with AH are identified with protocol ID 51. See also *Encapsulating Security Protocol (ESP)* and *Internet Protocol security (IPsec)*.

authenticity Assurance that data is coming from a known source and is valid or reliable. Validation processes, such as many encryption and hashing protocols, provide authenticity. Additionally, some processes, such as digital signatures, have authenticity as a key goal.

authorization A process that provides access to resources through the assignment of permissions. This process starts with authentication, where users claim an identity and prove their identity with a password or other credentials. Once users have been authenticated, authorization defines the resources that a user can access and the rights that a user can invoke.

availability One of the three main goals of information security known as the CIA security triad. Availability ensures that systems and data are up and operational when needed. The other two parts of the CIA triad are confidentiality and integrity.

backdoor A nontraditional method of accessing an application or system. It can be code embedded in an application that provides access to the application, the application's code, or data via a covert method. Attackers often try to install a backdoor onto a system after infecting it with malware that grants them remote access.

baseline A known starting point. Baselines are an important element of configuration control and are often implemented with images. If the baseline configuration is known, it's relatively simple to check the system to determine whether the configuration has been modified from the baseline. Anomaly-based intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) also use baselines by first documenting normal behavior in the form of a baseline. The system then monitors the activity and constantly compares it to the baseline.

bcrypt A key-stretching algorithm used to protect passwords on UNIX and Linux distributions stored in the shadow password file. It salts passwords before encrypting them with Blowfish, which helps thwart rainbow table attacks. Compare to *Password-Based Key Derivation Function 2 (PBKDF2)*.

Bell-LaPadula model An access control model used to ensure confidentiality. It uses two primary rules: no read up and no write down. Compare to *Biba model*.

Berkeley Internet Name Domain (BIND) A version of DNS software that runs on UNIX systems. It is freely available and runs on many DNS servers on the Internet.

Biba model An access control model used to ensure integrity. It uses two primary rules: no read down and no write up. Compare to *Bell-LaPadula model*.

biometrics Method of identifying unique characteristics of a person, such as a fingerprint or retina scan. Biometrics provide authentication in the "something you are" factor.

bit-copy tools Software used to capture the contents of a drive without modifying the data. The copy is known as a forensic duplicate image and can be examined without affecting the original.

blacklist A list of prohibited applications. Application blacklisting identifies specific applications that cannot run on a system. Compare to *whitelist*.

Bluetooth A wireless technology used for smaller wireless networks commonly called personal area networks (PANs). Devices must be paired with each other before they will communicate.

Bootstrap Protocol (BootP) A protocol that provides an IP address to clients and can be used to retrieve a bootable image for clients. It is similar to Reverse Address Resolution Protocol (RARP), although RARP only retrieves the IP address.

botnet A group of computers (often called *zombies*) controlled by an attacker. The term *botnet* is derived from *robot* and *network*. The attacker manages a command control center, and the computers in the botnet do the bidding of the attacker.

Brewer-Nash An access control model used to help prevent a conflict of interest. Data is classified based on conflict-of-interest classes. Users who have access to one class are denied access to data in conflicting classes. Sometimes called the *Chinese Wall model*.

bring your own device (BYOD) An organizational policy that allows users to bring their personally owned devices (such as smartphones and tablets) to work and connect them to the organization's network. Compare to *corporate-owned, personally enabled (COPE)* and *choose your own device (CYOD)*.

buffer overflow An error that can occur when a system receives more data than it expects and is unable to handle it gracefully. Attackers attempt to exploit buffer overflow errors to install malware on systems.

buffer overflow attack An attack on a system that has a buffer overflow vulnerability. Buffer overflow vulnerabilities can be reduced with input validation techniques and by keeping systems up to date.

bus Network configuration where all computing devices are connected directly to each other via a shared cable connection. Both ends of the bus must be terminated. If one of the terminators is not present or the cable is disconnected, communication with all devices on the bus stops. See also *mesh, star, token ring, and tree*.

business continuity plan (BCP) A written document that includes the processes and procedures to prevent mission-critical services from being interrupted or disrupted. A BCP includes disaster recovery elements used to restore the organization to fully functioning operations as quickly and efficiently as possible.

business impact analysis (BIA) A part of a BCP. It identifies the impact to the organization if any business functions are lost due to any type of incident. It helps an organization identify what business functions are critical to continued operations by identifying the impact to the business if a business function stops.

capability table A list of subjects, along with the capabilities granted to the subjects. These capabilities include rights and permissions. For example, a capability table can list several groups such as Project Managers, Project Team Leads, and Project Members. It then lists the rights and permissions granted to each of these groups. Compare to *access control matrix*.

certificate A file used for security purposes, such as authentication, encryption, nonrepudiation, protection of e-mail, and code signing. Certificate authorities issue and manage certificates. Public keys are distributed in certificates.

certificate authority (CA) An entity that issues and manages certificates throughout their lifetimes. CAs can be public (such as VeriSign) or private within an organization. Public CAs sell and validate certificates and provide assurances to users that certificates are valid. A CA is sometimes referred to as a *certification authority*.

certification A process that evaluates, describes, and tests a system and all of the controls that are in place to mitigate risks to the system. After a system has been certified, a separate accreditation process formally approves the system to operate.

chain of custody A document that shows exactly where a piece of evidence is from the point it is collected until it is disposed. A chain of custody provides proof that evidence has been protected. If the chain-of-custody document is not available, the validity of the evidence can be questioned and its usefulness negated.

change control A process that helps prevent unintended outages from changes. A change control process gives experts an opportunity to examine the change for potential problems before the change is implemented. Also called *change management*.

choose your own device (CYOD) An organizational policy that allows users to bring approved personally owned devices (such as smartphones and tablets) to work and connect them to the organization's network. Management approves a list of allowed mobile devices. Compare to *corporate-owned, personally enabled (COPE)* and *bring your own device (BYOD)*.

CIA security triad The three main goals of information security: confidentiality, integrity, and availability. Confidentiality ensures that unauthorized individuals cannot access data. Integrity prevents any unauthorized or unwanted modification of data or system configurations. Availability ensures that systems and data are available when needed.

ciphertext Data that is encrypted and not readable until it is decrypted into plaintext. Authorized personnel can decrypt ciphertext to return it to plaintext and read it. Ideally, unauthorized personnel cannot read ciphertext. Encryption techniques convert plaintext to ciphertext. Compare to *plaintext*.

ciphertext attack A cryptanalysis attack. Attackers use this when they only have ciphertext for analysis without any useful information about the plaintext data.

Clark-Wilson model An access control model used to ensure integrity; the model has stricter rules than the Biba model. These integrity rules enforce the principle of separation of duties.

cleartext Data that is not in an encrypted format. It is easily readable without any cryptographic or cryptanalysis techniques. Encryption techniques convert plaintext to ciphertext. Also called *plaintext*.

clipping level A predetermined threshold level. Many automated accounting systems use clipping levels to generate alerts after the system detects a preset number of events. An auditing system ignores events until the number of events reaches the clipping level.

code of ethics Statements and principles that individuals can use to guide their decisions and help in ethical dilemma situations. (ISC)²'s Code of Ethics includes a preamble and four canons describing ethical expectations from its certified practitioners. Candidates must commit to and

abide by them to earn and keep the SSCP certification. Organizations can also use ethics statements for internal employees.

cold site An alternative location used in business continuity planning. A cold site is a building with a roof, running water, and electricity. It doesn't include the necessary hardware, software, or personnel. In the event of an emergency, personnel move the required resources to the cold site location, hook them up, and configure them for operation. Compare to *hot site*, *warm site*, and *mobile site*.

collision, hash A collision that occurs when two different inputs to a hashing function produce the same hash. Hashing algorithms that are susceptible to hash collisions are considered compromised and should not be used.

collision, network A collision that occurs when two devices try to communicate at the same time over the same medium. Because both devices are sending traffic, the data from both systems collides with each other and none of the data is readable.

command injection An attack that attempts to inject commands into an application. In some cases, a command injection attack can inject operating system commands that would normally be executed at the command line. In other cases, it injects code such as JavaScript or SQL statements. Input validation techniques help mitigate command injection attacks.

Common Criteria A framework used to evaluate systems, formally known as Common Criteria for Information Technology Security Evaluation. It provides assurances that the specification, implementation, and evaluation of a system's security has gone through a rigorous and standardized process.

Common Vulnerability Scoring System (CVSS) An open standard that organizations can use to assess the severity of computer system security vulnerabilities. It outlines values for metrics that can then be plugged into a formula to determine the severity of any vulnerability.

community cloud A cloud-based service shared by two or more organizations. It is similar to a private cloud in that it is not available to the public. Compare to *public cloud*, *private cloud*, and *hybrid cloud*.

compensating control A control used if a primary control fails or is unavailable. When it isn't feasible or desirable to use the primary control, an organization can implement a compensating control.

confidentiality Ensures that unauthorized entities cannot access data. Access controls and encryption help protect against the loss of confidentiality. Confidentiality is one of the three main goals of information security known as the CIA security triad. The other two goals are integrity and availability.

configuration management A process that ensures that information about system configuration is available for any system and helps ensure that similar systems are configured similarly.

containment Step in the incident lifecycle that helps prevent a security incident from spreading. It is typically done by isolating an individual system by disconnecting it from the network, or isolating a network by disconnecting it from other networks. See also *incident lifecycle*.

control A means, method, action, technique, process, procedure, or device that reduces the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. Controls are risk management tools. The terms *control*, *countermeasure*, and *safeguard* are often used interchangeably.

control plane The part of a network that is responsible for identifying paths to other networks. Routers typically do this identification with routing protocols such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP). These routing protocols allow routers to share information with each other to create a virtual map of the network. Compare to *data plane* and *software defined network*.

corporate-owned, personally enabled (COPE) An organization purchases and issues devices (such as smartphones and tablets) to users instead of allowing them to connect their personally owned devices to the network. Compare to *bring your own device (BYOD)* and *choose your own device (CYOD)* policies.

corrective control A security control that takes action to reverse the effects of an event and/or restore a disabled or failed control. Many corrective controls often work in conjunction with detective controls. After detecting the event, the corrective control goes into action.

countermeasure A means, method, action, technique, process, procedure, or device that reduces the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. Controls are risk management tools. The terms *control*, *countermeasure*, and *safeguard* are often used interchangeably.

Crossover Error Rate (CER) The point where the False Acceptance Rate (FAR) and the False Rejection Rate (FRR) of a biometric system are equal or cross over. A lower CER indicates a better-performing biometric system. Compare to *False Acceptance Rate (FAR)* and *False Rejection Rate (FRR)*.

cross-site request forgery (CSRF or XSRF) An attack that results in websites executing unauthorized commands as if a user requested them. Attackers create specially crafted links and encourage users to click the links. For example, attackers can send the links in phishing e-mails or embed malicious code in a compromised website.

cross-site scripting (XSS) An attack that attempts to inject HTML or JavaScript into a web page. After the attack, the code executes on a user's system when the user visits the attacked

website. Successful attacks sometimes allow attackers to read cookies and use this data to launch session-hijacking attacks.

cryptanalysis The process of deciphering codes through analysis. Both attackers and researchers use cryptanalysis techniques. Researchers use them to search for weaknesses in cryptography, with the goal of improving cryptographic methods. Attackers use them to search for weaknesses, with the goal of exploiting them.

cryptography The science of scrambling, or encrypting, data. Cryptography includes the study of algorithms and other methods that help protect the confidentiality of data.

data at rest Any data that is in computer storage, such as on system hard drives, portable USB drives, flash drives, storage area networks, and backup tapes. Compare to *data in motion*.

Data Encryption Standard (DES) An older symmetric encryption standard using 56-bit keys. It has been broken and is rarely used anymore.

data in motion Any data being transmitted over a network (sometimes called data in transit). It includes data transmitted over an internal network using wired or wireless methods, as well as data transmitted over the Internet. Compare to *data at rest*.

data inference The ability of someone to gain knowledge by piecing together unclassified data to determine classified or sensitive information. It also includes the ability to gather large quantities of information and attempt to learn details of the information through deduction.

data loss prevention (DLP) Techniques used to monitor data usage and prevent the unauthorized use or transmission of sensitive data. Different types of DLP systems can monitor *data in motion* and *data at rest*.

data plane The part of a network that uses rules within an access control list (ACL) to identify whether the router will pass traffic or block traffic. These ACL rules are traditionally referred to as *packet-filtering rules*. The data plane decides if traffic should be forwarded or not. Sometimes called the forwarding plane. Compare to *control plane* and *software defined network*.

database view A virtual table that provides access to specific columns in one or more tables. A view doesn't hold any data but presents the data in the underlying table or tables. A database administrator can grant access to a view without granting access to a table to limit what a user can see and manipulate.

decryption The process of converting ciphertext data into plaintext data. Data is encrypted to prevent loss of confidentiality. Compare to *encryption*.

deduplication The process of keeping only a single copy of a file on a system instead of multiple identical files. Deduplication saves storage space.

defense diversity A defense-in-depth strategy using dissimilar technologies. Implementing a demilitarized zone (DMZ) with firewalls from two separate vendors is an example of defense diversity.

defense in depth A strategy that provides a layered approach to security. Instead of using one or two security controls, multiple controls are used. If one control fails, other controls continue to provide protection.

demilitarized zone (DMZ) A perimeter network used to host resources available on the Internet (such as web servers, e-mail servers, or FTP servers). The DMZ provides a layer of protection for the resources that would not be available if they were placed directly on the Internet.

denial of service (DoS) An attack that attempts to prevent a system from answering legitimate requests from users, directly affecting the availability portion of the CIA triad. The attack is launched by a single system. Compare to *distributed denial of service (DDoS)*.

deprovisioning The process of disabling and deleting inactive user accounts. Identity and access management (IAM) practices include both provisioning and deprovisioning activities.

detective control A security control that takes attempts to detect incidents either as they are occurring or after they've occurred. They can provide evidence that preventive controls are working (or not working).

deterrent control A control that attempts to deter would-be attackers from attempting an attack and/or deter users from trying to circumvent policies. It is similar to a preventive control. The primary difference is that a deterrent control encourages someone to decide not to take a specific action.

differential backup A type of backup used in full/differential backup strategies. Differential backups only back up data that has changed since the last full backup, without regard to any other differential backups. Compare to *incremental backup*.

digital signature A file used to provide authentication, integrity, and nonrepudiation security for e-mail. A digital signature is created by hashing an e-mail message and then encrypting the hash with the sender's private key. It is decrypted with the sender's public key.

disaster recovery plan (DRP) A document used to provide an organization with a plan to restore critical operations after a disaster. The overall goal is to provide employees with clear-cut steps on what to do and the order of these steps.

Discretionary Access Control (DAC) model An access control model commonly used to control access for file systems, such as New Technology File System (NTFS) or Network File System (NFS). Objects (such as files and folders) are owned by users. Users have full control over the objects and can grant others access. DAC provides the most granular level of control. Other access control models are Mandatory Access Control (MAC), Role-based Access Control

(RBAC or Role-BAC), Rule-based Access Control, and Attribute-based Access Control (ABAC).

distributed denial of service (DDoS) A DoS attack launched against a single system from multiple attackers. Botnets are often used in DDoS attacks. Compare to *denial of service (DoS)*.

DNS poisoning An attack that attempts to redirect traffic away from legitimate servers by corrupting Domain Name System (DNS) cache. Sometimes called DNS cache poisoning.

Domain Keys Identified Mail (DKIM) An email authentication mechanism that uses digital signatures to authenticate e-mails. The sending e-mail server attaches a digital signature to outgoing e-mail. The receiving e-mail server uses the digital signature to verify the e-mail has been sent by the sending e-mail server. Compare to *Sender Policy Framework (SPF)*.

Domain Name System (DNS) Provides name resolution services by resolving host names to IP addresses and IP addresses to host names. DNS uses port 53.

Domain Name System Security Extensions (DNSSEC) A suite of extensions that authenticate the source of Domain Name System (DNS) responses and authenticate the integrity of the response. Each response is signed with a Resource Record Signature (RRSIG), which is similar to a digital signature.

Domain-based Message Authentication, Reporting, and Conformance (DMARC) An extension of Sender Policy Framework (SPF) and Domain Keys Identified Mail (DKIM) that allows owners of a domain to publish a policy that identifies if they are using SPF, DKIM, or both.

dropper A type of Trojan horse that attempts to drop malware onto a computer. It might be from an infected file delivered via e-mail or delivered via a drive-by download.

dual factor authentication Authentication using two factors. There are three factors of authentication: something you know (such as a password), something you have (such as a smart card), and something you are (using biometrics). Compare to *single factor authentication* and *multifactor authentication*.

due care The practice of implementing security policies and practices to protect resources. Organizations are required by law to exercise due care or they can be found legally negligent.

due diligence The investigative steps that an organization takes prior to taking on something new, such as signing a contract or making a major purchase. An organization has an obligation to exercise due diligence to discover risks.

Dynamic Host Configuration Protocol (DHCP) A protocol used to assign TCP/IP configuration information to DHCP clients. This includes an IP address, a subnet mask, the address of a DNS server, and much more. DHCP uses UDP ports 67 and 68.

elliptic curve cryptography (ECC) A type of cryptography that creates public and private keys using a mathematically defined curve. It provides strong security with smaller key sizes when compared to other methods such as RSA. ECC is often used on smaller mobile devices because the smaller keys require less processing power.

Encapsulating Security Protocol (ESP) One of the underlying protocols of IPsec, ESP encrypts the data and also provides the same authentication services provided from Authentication Header (AH). Data encrypted with ESP cannot be read if captured with a sniffer. Packets protected with ESP are identified with protocol ID 50.

encryption The process of converting plaintext data into ciphertext data to prevent loss of confidentiality. The process is reversed by decrypting the ciphertext data to create the original plaintext data. Compare to *decryption*.

entitlement The process of granting privileges to users. The principle of least privilege (giving only the privileges needed for the job) is an important part of entitlement.

eradication Part of the incident lifecycle. It ensures that all elements of an attack or infection are completely removed from the affected systems.

escalation A process of reporting a security incident to other personnel in the organization who typically have more expertise with the incident. As an example, a help desk professional may report a malware infection to dedicated security personnel.

Ethernet Set of local area network (LAN) standards defined by the IEEE 802.3 documents. It is the most widely used standard for LANs.

event data analysis The process of reviewing and analyzing events to determine if they are potential incidents. Monitoring systems such as intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) automate event data analysis.

extranet An area of an organization's network used to host resources via the Internet but is available only to trusted entities. An extranet is available via the Internet, but only to a specific target audience. Compare to *intranet*.

False Acceptance Rate (FAR) A type 2 biometric error. It refers to the percentage of times a biometric system falsely identifies an unknown user as a known user. Compare to *False Rejection Rate (FRR)* and *Crossover Error Rate (CER)*.

false negative The lack of an indication of an attack or vulnerability when one exists. For example, if an attack is occurring, but an intrusion detection system does not detect it, it is a false negative. Similarly, if a vulnerability exists in a system or network, but a vulnerability scanner does not detect it, it is a false negative. Compare to *false positive*.

false positive A false indication of an attack or a vulnerability. Vulnerability scanners can give false positives indicating that a vulnerability exists even though it doesn't. Similarly, intrusion

detection and prevention systems can give an indication that a system is being attacked even when it isn't. Compare to *false negative*.

False Rejection Rate (FRR) A type 1 biometric error. It refers to the percentage of times a biometric system falsely rejects a known user and instead indicates that the user is unknown. Compare to *False Acceptance Rate (FAR)* and *Crossover Error Rate (CER)*.

fault tolerance The practice of adding and maintaining redundancy for disks, servers, connections, and sites. Fault tolerance controls help eliminate outages from single points of failure by adding redundancies.

federated access A single sign-on (SSO) technology that allows users in different networks to access multiple systems after logging on once. The systems can be using different operating systems owned and managed by different organizations.

Fiber Distributed Data Interface (FDDI) An extension of token ring networks that uses fiber-optic connections instead of copper. FDDI uses dual rings, and the second ring provides redundancy. FDDI token rings support speeds up to 100 Mbps over distances as far away as 120 miles.

file integrity checker An application that can verify that files have not been modified. In doing so, the application guards against a loss of *integrity*. File integrity checkers use *hashing* algorithms to capture hashes of files in a known good state. Later, they create hashes on the same file and compare the two hashes. If the file is unchanged, the hashes are the same, but if the file has been modified, the hashes are different.

File Transfer Protocol (FTP) A protocol used to transfer files from one system to another. Data is transmitted in cleartext using TCP ports 20 and 21. See also *Secure File Transfer Protocol (SFTP)* and *Trivial File Transfer Protocol (TFTP)*.

firewall Software or hardware used to filter traffic into or out of a network. A firewall can be a dedicated physical device or an additional application running on a system such as a desktop computer. A stateless firewall (also called a *packet-filtering firewall*) filters traffic based on individual packets. Compare to *stateful inspection firewall*.

first responder The first forensics-trained person on the scene of an information technology (IT) incident. It could be a highly trained forensics expert or a system administrator with only basic training. One of the primary responsibilities of a first responder is the preservation of the scene of the incident and all associated evidence.

foreign key A type of key used in databases. A foreign key in one table points to a unique primary key in another table to create a relationship between the two tables. See also *primary key*.

forensics The science of examining and inspecting crime scenes for evidence. Computer forensics is the science of examining and inspecting computer systems for evidence about an event or crime.

General Data Protection Regulation (GDPR) A European Union (EU) regulation that mandates the protection of personal data of individuals within the EU.

GNU Privacy Guard (GnuPG or GPG) A free, open source alternative to Pretty Good Privacy (PGP). It includes technologies that allow individuals to encrypt, decrypt, and digitally sign e-mails using self-signed certificates instead of a certificate issued by a certificate authority (CA). See also *web of trust (WOT)*.

guidelines Recommendations provided to members of an organization. Guidelines aren't mandatory or authoritative in nature. Guidelines (and procedures) are derived from policies.

hardening The practice of making a system more secure from its default configuration. Hardening often includes removing or disabling unused protocols and services, changing defaults, keeping systems up to date, enabling firewalls, and using AV software.

hardware token A device held by a user that displays a number or a password that changes frequently, such as every 60 seconds. The number is synchronized with a server and used as a one-time password.

Hash-based Message Authentication Code (HMAC) A hashing algorithm that works with other hashing algorithms. It uses a shared secret key that adds some randomness to the hash. It provides authenticity and nonrepudiation.

hashing The process of creating a hash (a number) by executing a hashing algorithm against a piece of data. If the original data is not changed, the hash will always be the same. You can execute the hashing algorithm on the file or message at one point to create a hash and then later execute the same hashing algorithm again. If the hashes are the same, the original data has not been modified and it has retained integrity.

Health Insurance Portability and Accountability Act (HIPAA) A U.S. law that mandates the protection of protected health information (PHI). See also *protected health information (PHI)*.

heuristics-based detection A method of virus detection used to detect previously unknown viruses. Heuristics attempts to detect a virus based on its behavior.

honeypot A computer set up to entice would-be attackers. It is often configured with weak security so that an attacker can easily hack into it, and usually has fake data that has no value to the company.

host-based firewall A firewall that provides protection for a single host. Many operating systems include host-based firewalls running as additional software. Compare to *network-based firewall*.

host-based IDS (HIDS) An intrusion detection system (IDS) installed on an individual system, such as a server or workstation. It can only monitor activity on the host, not network activity. Compare to *network-based IDS (NIDS)*.

hot site An alternative location used in business continuity planning. A hot site includes all the necessary resources to take over the operations of another location in a very short period of time, sometimes within minutes. It includes hardware such as servers and the network infrastructure, up-to-date data, and personnel to manage the functions of the alternative location. Compare to *cold site*, *warm site*, and *mobile site*.

hybrid cloud A cloud-based service that is a combination of any two or more clouds. Compare to *public cloud*, *community cloud*, and *private cloud*.

Hypertext Transfer Protocol (HTTP) The common protocol used to transfer web pages over the Internet. Most web browsers use HTTP by default. HTTP uses TCP port 80.

Hypertext Transfer Protocol Secure (HTTPS) An encrypted form of HTTP. Transport Layer Security (TLS) replaced Secure Sockets Layer (SSL) in most implementations of HTTPS. HTTPS uses TCP port 443.

identification The act of claiming an identity, typically by using a username. Users prove identity by providing proper credentials such as a password.

identity and access management (IAM) The process of managing accounts throughout their lifetime, including their access to resources. Provisioning includes creating accounts for users and granting appropriate privileges. Maintenance ensures that password policies and account policies are implemented properly and ensures that inactive accounts are disabled. Entitlement helps ensure that the principle of least privilege is enforced.

impact assessment An evaluation of the magnitude of harm that could result if a threat exploits a vulnerability. Impact assessments are often completed with risk assessments. See also *risk assessment*.

incident Any violation of policies or security practices that has the potential to result in an adverse event or that has resulted in an adverse event. Risk management practices attempt to prevent any incidents, detect incidents when they occur, and correct problems that occur from incidents. Also called a *security incident*.

incident lifecycle All the steps before, during, and after a security incident. It includes preparation, detection, analysis, escalation, containment, eradication, recovery, and lessons learned/implementation of new countermeasures.

incident response The process of responding to an incident. Many organizations have formal procedures for responding to incidents.

incremental backup A type of backup used in full/incremental backup strategies. Incremental backups back up only data that has changed since the last full or incremental backup. Compare to *differential backup*.

information rights management (IRM) Methods used to protect documents after they've been sent to an untrusted party. IRM can prevent recipients from forwarding, copying, modifying, printing, faxing, or pasting the contents from a file.

Infrastructure-as-a-Service (IaaS) A type of cloud computing that provides customers with hardware (such as servers) and infrastructure (such as routers or switches). Customers maintain all software on the hardware, but the provider owns and maintains the hardware. Also known as *Hardware-as-a-Service*. Compare to *Software-as-a-Service (SaaS)* and *Platform-as-a-Service (PaaS)*.

inline Device placement that requires all traffic to go through the device. This is often done with security appliances to ensure that all Internet traffic goes through the appliance before it hits devices on the internal network.

input validation The process of checking data before using it within an application. Web applications use input validation to prevent different types of injection attacks, such as SQL injection or cross-site scripting attacks. Input validation also helps prevent buffer overflow attacks.

insider threat A threat from someone who works within an organization, such as an employee, consultant, or vendor. Insiders have authorized access to the organization's assets and can potentially cause loss of confidentiality, integrity, and/or availability to these assets.

integrity Assurance that data or a system configuration has not been modified. Hashing and audit logs are two methods used to ensure integrity. Integrity is one of the three main goals of information security known as the CIA security triad. The other two goals are confidentiality and availability.

intellectual property Creations of the mind or intellect such as literary works, artistic works, and inventions. It includes industrial property and copyrights.

International Organization for Standardization (ISO) An independent international organization that publishes standards. Standards and related documents can be purchased from the ISO website. Compare to *National Institute of Standards and Technology (NIST)*.

Internet Control Message Protocol (ICMP) A protocol used to check or verify the health of a network or network device. ICMP is used with many diagnostic protocols such as ping, pathping, and tracert.

Internet Group Message Protocol (IGMP) A protocol used for IPv4 multicasting, where data is sent from one computer to multiple computers at the same time. IPv6 uses ICMPv6 messaging instead of IGMP.

Internet Message Access Protocol version 4 (IMAP4) A protocol used to manage e-mail in folders. IMAP4 servers store e-mail on the server and allow users to organize the e-mail in folders. It uses TCP port 143 or port 993 with Transport Layer Security (TLS). Compare to *Post Office Protocol version 3 (POP3)* and *Simple Mail Transfer Protocol (SMTP)*.

Internet Protocol security (IPsec) A security protocol used to provide security for IP traffic traveling over a network. IPsec includes Authentication Header (AH) and Encapsulating Security Protocol (ESP). AH provides authentication between the systems and verifies the integrity of the packets. ESP encrypts the data and provides the same authentication services provided by AH.

intranet A network that is internal to an organization. An intranet has private IP addresses, and clients within the intranet access the Internet with Network Address Translation (NAT), which translates private IP addresses to public and public IP addresses to private. Compare to *extranet*.

intrusion detection system (IDS) A system designed to provide continuous monitoring of networks and hosts to help protect them from attacks. The goal is to detect an attack as it is occurring and raise an alert. An IDS is not placed inline with traffic, but instead uses agents to capture traffic. Compare to *intrusion prevention system (IPS)*.

intrusion prevention system (IPS) Similar to an intrusion detection system (IDS), but with one primary difference. An IPS is placed inline with the traffic, preventing the malicious traffic from reaching the protected network.

job rotation The practice of moving employees between different jobs on a periodic basis. Job rotation is a security practice that can reduce the risk of fraud within a company. Organizations commonly combine it with a separation of duties policy.

key composition Refers to creating strong encryption keys. Strong keys are random and of sufficient length that they cannot be guessed. The length of keys varies with different algorithms.

key escrow The process of storing a copy of an encryption key or an alternate encryption key. It ensures that encrypted data can be retrieved even if the original encryption key is lost or unavailable.

key rotation Refers to periodically replacing encryption keys with new encryption keys. This is similar in concept to changing a password regularly. The length of time between using a different key varies for different algorithms.

key stretching A method that transforms a potentially weak password or encryption key into a more secure password or encryption key. Password-Based Key Derivation Function 2 (PBKDF2) and bcrypt are two key-stretching algorithms.

known-plaintext attack A cryptanalysis attack used when the attacker has samples of both plaintext and ciphertext data. For example, if an attacker has plaintext from an encrypted message, the attacker can then use different methods to try to decrypt the ciphertext to the known text. If the known text is decrypted, the same method can decrypt similar data.

Layer 2 Tunneling Protocol (L2TP) A tunneling protocol used with some virtual private networks (VPNs). It combines the strengths of Layer 2 Forwarding (L2F) and Point-to-Point Tunneling Protocol (PPTP) and is documented in RFC 2661. L2TP VPN traffic is commonly encrypted with Internet Protocol security (IPsec) (as L2TP/IPsec).

least privilege An important security principle that ensures that users are given only the rights and privileges that they need to perform their job, and no more.

lessons learned/implementation of new countermeasures The last step in the incident lifecycle. It includes a review of the incident, what caused it, what allowed it, and how the organization responded to the incident. This step often includes a review of countermeasures and a recommendation for new countermeasures. See also *incident lifecycle*.

load balancing A method of spreading the load of a service among several different servers. It is often used with web servers to serve the same web page content from several different servers. Load balancing can be accomplished with software or a dedicated hardware device. Two common configurations are *round robin* and *source address affinity*.

logic bomb Malicious code that executes in response to an event, such as a point in time, or in response to specific actions.

malware Malicious software or malicious code. It includes any code or software that can be harmful or destructive to computers, networks, or the computing environment as a whole. Some of the common malicious code types include viruses, worms, logic bombs, and Trojan horses.

Mandatory Access Control (MAC) model An access control model that provides the highest level of security. It is used by the U.S. military. Subjects and objects are assigned labels; subjects (such as users) can access objects (such as files) only when the labels match. Other access control models are Discretionary Access Control (DAC), Role-based Access Control (RBAC or Role-BAC), Rule-based Access Control, and Attribute-based Access Control (ABAC). Examples of models that support Mandatory Access Control include Bell-LaPadula and Biba.

mandatory vacation A security policy designed to reduce fraud within an organization. A mandatory vacation policy requires employees to take a vacation outside of the workplace for a specified period, such as at least five consecutive workdays. The goal is to require someone else to perform the employee's job functions, which increases the possibility of exposing any suspicious activities.

man-in-the-middle An attack that can capture traffic sent between two systems. It performs electronic eavesdropping by placing a system either logically or physically between the two other systems.

mantrap A security control designed to protect against piggybacking. It prevents more than one person from passing through a controlled entry at a time and ensures that each person uses credentials to gain entry.

maximum acceptable outage (MAO) The maximum amount of time a system can be down before critical business functions are affected. MAO is sometimes called maximum tolerable outage (MTO) or maximum tolerable downtime (MTD). The MAO helps an organization determine the recovery time objective (RTO).

Media Access Control (MAC) address A 48-bit or 12-digit hexadecimal number (such as 6c-62-6c-ba-73-6c) assigned to a network interface card (NIC). The MAC address uniquely identifies the computer in a network. Also called *physical address* or *hardware address*.

mesh A network topology that provides redundancy with multiple connections. Mesh topologies provide the highest availability of any of the topology methods. See also *bus*, *star*, *token ring*, and *tree*.

Message Digest 5 (MD5) A hashing algorithm used for integrity. MD5 creates a 128-bit hash. The U.S. government considers it cryptographically broken, so it is not recommended for use in many government applications. SHA-3 is a much stronger alternative and recommended for use instead.

mobile code Code that can be transferred between systems without the user taking explicit action to install it. It can include script code such as JavaScript, compiled code such as Java applets or ActiveX controls, documents with embedded macros, or malware that executes from a USB flash drive when the user inserts the drive into a system.

Mobile Device Management (MDM) Technologies used to manage mobile devices such as smartphones and tablets. MDM solutions can ensure devices meet minimum requirements before allowing them access to a corporate network.

mobile site An alternative location used in business continuity planning. A mobile site can easily be moved and is useful when an organization doesn't want to designate a specific alternative location. For example, it's possible to set up the inside of a storage container as an alternative location. These are the same storage containers that 18-wheeler tractor-trailers use to haul goods, and they are as easy to move as any other loads hauled over the highways.

multifactor authentication Authentication using more than one factor. There are three factors of authentication: something you know (such as a password), something you have (such as a smart card), and something you are (using biometrics). Compare to *single factor authentication* and *dual factor authentication*.

National Institute of Standards and Technology (NIST) An organization within the U.S. Department of Commerce that includes a Computer Security Division and an Information Technology Laboratory. NIST uses federal funds to perform research and regularly publishes Special Publications (SPs) related to cybersecurity that are freely available to anyone. Compare to *International Organization for Standardization (ISO)*.

near field communication (NFC) Technology in smartphones that allows a user to transfer information to other devices by waving a smartphone close to the other device or tapping the other device. NFC can be used to make purchases with credit cards registered on the phone.

Network Address Translation (NAT) A protocol that translates private IP addresses to public IP addresses and public IP addresses back to private IP addresses. NAT is often installed on a proxy server or a router that is on the edge of the network (between the Internet and the intranet).

Network Discovery Protocol (NDP) Resolves IPv6 addresses to EUI-64 addresses. It is similar to ARP, which resolves IPv4 addresses to MAC addresses. IPv6 uses a 64-bit modified EUI-64 address instead of the MAC address.

Network File System (NFS) Allows computers running different operating systems to access and share files over the network. It allows users running UNIX (and UNIX derivatives) to access files on Microsoft systems, and allows users running Microsoft systems to access files on UNIX-based systems. Sun Microsystems created NFS.

network-based firewall A firewall that provides protection for a network. Traffic to and from the network flows through the network-based firewall. Compare to *host-based firewall*.

network-based IDS (NIDS) An intrusion detection system (IDS) that monitors the network traffic for an attack. A NIDS typically has several nodes or agents stationed around the network connected to routers and possibly switches. Each of these nodes monitors the traffic and reports its findings to a NIDS management server.

non-Discretionary Access Control (non-DAC) model Alternative access control models used to control system file access. They include Mandatory Access Control (MAC), Role-based Access Control (RBAC or Role-BAC), Rule-based Access Control, and Attribute-based Access Control (ABAC).

nonpromiscuous mode A mode used by a sniffer that limits the traffic that the sniffer can capture. The sniffer only captures data sent directly to or from the IP address of the computer running the sniffer. Compare to *promiscuous mode*.

nonrepudiation The practice of ensuring that a party cannot believably deny (or repudiate) taking an action. Nonrepudiation is enforced through audit logging and with digital signatures.

object A resource accessed by a subject. For example, if a user accesses a file, the user is the subject and the file is the object.

one-time password A password used only once. Hardware tokens use synchronous one-time passwords with a hardware device held by the user showing a one-time password as a number or password. The password changes often and is synchronized with a server.

Online Certificate Status Protocol (OCSP) A method of validating certificates with a certificate authority (CA). Clients send the serial number of a certificate to a server known as an

OCSP responder. The OCSP responder identifies the health of the certificate with an answer of “good,” “revoked,” or “unknown.”

Open Systems Interconnection (OSI) Model A seven-layer model developed by the International Organization for Standardization (ISO) as a framework for connecting computers and networks together. The seven layers are Physical (layer 1), Data Link (layer 2), Network (layer 3), Transport (layer 4), Session (layer 5), Presentation (layer 6), and Application (layer 7).

packet-filtering firewall A firewall that filters traffic by examining the contents of a packet. A packet-filtering firewall can filter traffic based on IP addresses, subnet addresses, ports, some protocols, or any combination of these.

Password-Based Key Derivation Function 2 (PBKDF2) A key-stretching algorithm that adds a salt of at least 64 bits and then applies a cryptographic function multiple times. Salting helps thwart rainbow table attacks. Compare to *bcrypt*.

penetration test An assessment that starts with a vulnerability assessment. Instead of stopping after discovering vulnerabilities, a penetration test attempts to exploit the discovered vulnerabilities.

personally identifiable information (PII) Information that can be used to personally identify an individual. Many laws mandate the protection of PII.

pharming An attack that redirects users to bogus websites. It manipulates one of the host name resolution methods so that the host name resolves to a different website.

phishing An e-mail-based attack. Attackers send e-mails with the goal of tricking victims into clicking a link or providing sensitive information. The e-mail claims to be from a legitimate company and encourages the user to take an unsafe action.

physical controls Security controls that you can touch. These include any controls that control, restrict, and/or monitor physical access and provide environmental protection Compare to administrative controls and technical security controls.

piggybacking A social engineering tactic. Piggybacking occurs when someone passes through a controlled entry without providing credentials by following closely behind someone who has provided credentials. Also called *tailgating*.

plaintext Data that is not in an encrypted format. It is easily readable without any cryptographic or cryptanalysis techniques. Encryption techniques convert plaintext to ciphertext. Also called *cleartext*.

Platform-as-a-Service (PaaS) A type of cloud computing that provides customers with a preconfigured computing platform. Users have access to an operating system and one or more applications hosted on hardware that they can use over the Internet. Users don't have to purchase

the hardware and software, and the cloud provider maintains it. Compare to *Infrastructure-as-a-Service (IaaS)* and *Software-as-a-Service (SaaS)*.

Point-to-Point Tunneling Protocol (PPTP) A tunneling protocol used with some virtual private networks (VPNs). RFC 2637 defines it, and Microsoft uses its own version.

policy High-level documents used to provide guidance to members of an organization. A security policy provides direction to employees and is authoritative in nature.

polymorphic virus A virus that can morph or mutate each time it replicates to another machine or even each time it is run. When the virus mutates, it can be more difficult for antivirus (AV) software to detect it.

port number Numbers that identify protocols based on assigned well-known ports. For example, SMTP uses TCP port 25. This is not a protocol number, though protocol numbers also identify some protocols. Compare to *protocol number* and *well-known port*.

port scan attack An attempt to detect what ports are open on a system as part of an overall fingerprinting attack. Open ports indicate which services are running on the system.

Post Office Protocol version 3 (POP3) A protocol used to receive e-mail. POP3 servers send the e-mail to the client when the client connects. It uses TCP port 110 or port 995 with Transport Layer Security (TLS). Compare to *Simple Mail Transfer Protocol (SMTP)* and *Internet Message Access Protocol version 4 (IMAP4)*.

Pretty Good Privacy (PGP) Technologies that allow individuals to encrypt, decrypt, and digitally sign e-mails using self-signed certificates instead of a certificate authority (CA). PGP started as open source software freely available to anyone. However, it has gone through several iterations of ownership and is currently proprietary software owned by Symantec and can be used with a certificate issued by a CA. Compare to *GNU Privacy Guard (GnuPG or GPG)*.

preventive control A security control that attempts to prevent security incidents. The most basic type of preventive controls consists of written security policies, standards, and procedures.

primary key A key used in databases that ensures that each row or tuple within a table is unique.

private cloud A cloud-based service available only to users within an organization. As an example, a company can set up cloud-based storage for its employees. Compare to *public cloud*, *community cloud*, and *hybrid cloud*.

private IP address An IP address used within an organization's private network. RFC 1918 formally defines private IP addresses. Compare to *public IP address*.

private key A key used with public key cryptography for asymmetric encryption. A private key is part of a matched pair (matched to a public key). A private key is kept private and not shared with other entities. Compare to *public key* and *session key*.

privileges A combination of rights and permissions. Rights are actions that an account can take on a system, such as backing up files, changing the time, or rebooting the computer. Permissions identify what a user can do with resources, such as reading and writing to a file or printing to a printer.

procedure Specific action steps to accomplish tasks. Procedures (and guidelines) are derived from policies.

promiscuous mode A mode used by a sniffer that increases the amount of traffic that the sniffer can capture. The sniffer can capture all data that reaches the sniffer regardless of the destination IP address. Compare to *nonpromiscuous mode*.

proofing Part of the identify management lifecycle. Identity proofing verifies a person's identity before creating a user account. It can also be done to verify a user's identity before providing password hints, or allowing the user to change an account password.

protected health information (PHI) Personally identifiable information that includes an individual's medical or health data.

protocol number A number embedded in the first octet of TCP and UDP headers that identifies some protocols. For example, ICMP uses protocol number 1 and IGMP uses protocol number 2. This is not a port number, though port numbers also identify some protocols. Compare to *port number*.

provisioning The process of creating user accounts and granting access to the appropriate resources. Identity and access management (IAM) practices include both provisioning and deprovisioning activities.

proxy server A server used as an intermediary, or proxy, for internal clients accessing Internet resources. The proxy server retrieves the web page on the client's behalf and returns it to the client.

public cloud A cloud-based service provided by a third-party vendor and available to anyone. As an example, Apple provides cloud-based storage via its iCloud service. Compare to *private cloud*, *community cloud*, and *hybrid cloud*.

public IP address An IP address that is used on the public Internet. Compare to *private IP address*.

public key A key used with public key cryptography for asymmetric encryption. A public key is part of a matched pair (matched to a private key). A public key is freely shared and distributed within a certificate. Compare to *private key* and *session key*.

public key infrastructure (PKI) All the components necessary to create, manage, distribute, validate, and revoke certificates. A PKI is based on the X.509 protocol. The X.509 standard identifies many of the components and formats used by the PKI, certificate authorities (CAs), and certificates. Compare to *web of trust (WOT)*.

qualitative analysis A risk assessment method that uses subjective opinions from experts to identify the impact of a risk. It often categorizes a risk using words such as “low,” “medium,” and “high.” Compare to *quantitative analysis*.

quantitative analysis A risk assessment method that uses numerical-based data such as monetary figures to identify the actual cost associated with a risk. Compare to *qualitative analysis*.

quarantine Isolating a system or code. Network access controls isolate systems that don’t meet minimum health requirements. Antivirus software isolates potentially malicious software.

radio-frequency identification (RFID) A system used for identification, tracking, asset management, and inventory control. Products are “tagged” with an RFID tag, which is a small electronic device that marks the product. RFID readers can then read the tag to get information about the product.

RADIUS See *Remote Authentication Dial-In User Service (RADIUS)*.

rainbow table A table of passwords using a predefined character set and the hash for each of the passwords. Attackers use rainbow tables to perform rainbow table lookups and identify passwords from the hash.

ransomware A type of Trojan malware that takes control of a user’s computer and then demands a ransom from the user to get control back.

RC4 A symmetric encryption standard. This standard is sometimes called Ron’s Code or Rivest’s Cipher after its inventor, Ron Rivest. Security experts have identified multiple vulnerabilities in RC4 and it is no longer recommended for use.

recovery point objective (RPO) A recovery term associated with backups and databases that identifies the amount of data (in terms of time such as hours or days) that is acceptable to lose if a failure occurs. The RPO dictates the amount of resources needed to protect data to prevent loss in the case of a failure.

recovery time objective (RTO) The maximum amount of time that can be taken to restore a system or process to operation. If a failure takes longer than the RTO to restore, then it impacts the mission. The *maximum acceptable outage (MAO)* helps an organization determine the RTO.

remote access Trojan (RAT) A type of Trojan horse that provides an attacker with remote access to an infected system. RATs often create backdoors into systems.

Remote Authentication Dial-In User Service (RADIUS) A service that provides centralized authentication, authorization, and accounting (AAA) for remote clients. Remote users are authenticated against a credential's database, their account is checked to verify that they are authorized to use remote access, and their remote access activity is tracked in an accounting log.

residual risk The amount of risk that remains after steps have been taken to mitigate risk. In other words, residual risk = total risk – mitigated risk. Management is responsible for any losses that occur because of residual risk.

Reverse Address Resolution Protocol (RARP) You can think of this as a simplified version of DHCP. A client has a MAC address and uses RARP to get an IP address. It provides the clients some TCP/IP configuration information, but doesn't have the full capabilities of DHCP.

risk The probability or likelihood of a threat exploiting a vulnerability, resulting in a loss. A threat is any activity that can be a possible danger. A vulnerability is a weakness, and a loss represents a negative event for an organization. The level of loss represents the impact.

risk assessment A point-in-time evaluation of potential risks. It looks at the current situation and then attempts to determine what risks exist and how to address them.

risk management framework A structured process that organizations can use to manage risk for IT systems. The National Institute of Standards and Technology (NIST) and the International Organization for Standardization (ISO) have published risk management frameworks in SP 800-37, *Guide for Applying the Risk Management Framework to Federal Information Systems*, and ISO 31000, *Risk Management*, respectively.

risk register A central depository for known organizational risks and is a key method used to share threat intelligence. It includes details about the known risks and can consolidate data from multiple sources.

risk treatment Methods of managing risk such as accept, avoid, mitigate, recast, and transfer. Also known as risk management.

Role-based Access Control (RBAC or Role-BAC) model An access control model that uses roles to determine access. Subjects (such as users) are placed into roles, and access to objects (such as files) is granted to the roles. Other access control models are Mandatory Access Control (MAC), Discretionary Access Control (DAC), Rule-based Access Control, and Attribute-based Access Control (ABAC).

rootkit A set of programs that can run on a system largely undetected. It has root-level access to the system, similar to how a root-level administrator has full and complete control over a system.

round robin A configuration commonly used with load balancing in which the first connection hits the first server, the second connection hits the second server, and so on. The load balancer

continues to redirect the connections in the same rotating pattern. Compare to *source address affinity*.

RSA An algorithm used for encryption and decryption in public key cryptography. RSA uses large prime numbers to create secure, matching public and private keys. Its strength lies in the fact that it's computationally infeasible to factor the composite number created from these prime numbers as long as the prime numbers are sufficiently large. NIST recommends using keys of at least 2048 bits.

Rule-based Access Control model An access control model that uses rules to determine access. As an example, routers use rules within an access control list (ACL) to define traffic allowed in or out of a network. Other access control models are Mandatory Access Control (MAC), Discretionary Access Control (DAC), Role-based Access Control (Role-BAC), and Attribute-based Access Control (ABAC).

safeguard A means, method, action, technique, process, procedure, or device that reduces the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. Controls are risk management tools. The terms *control*, *countermeasure*, and *safeguard* are often used interchangeably.

salting A process that adds random bits to a password. Salting helps thwart rainbow table attacks and is used by bcrypt and Password-Based Key Derivation Function 2 (PBKDF2).

sandboxing A security control that runs computer programs in isolated areas of memory. Antivirus programs use sandboxing to check new or untrusted applications. Security professionals often use virtualization as a sandboxing technique to test applications.

scareware Malware that describes itself as free antivirus software to trick users into downloading it. It typically attempts to scare the user with a pop-up or other message that indicates the user's system is infected with a virus. Sometimes called rogueware.

Secure File Transfer Protocol (SFTP) A protocol that uses Secure Shell (SSH) to encrypt FTP traffic. SFTP functions similarly to FTP, but the encryption prevents sniffing attacks from capturing data in cleartext. Compare to *File Transfer Protocol (FTP)* and *Trivial FTP (TFTP)*.

Secure Hashing Algorithm 1 (SHA-1) A hashing algorithm used for integrity. SHA-1 creates a 160-bit hash. Due to vulnerabilities, it is no longer recommended for use.

Secure Hashing Algorithm 2 (SHA-2) A hashing algorithm used for integrity. SHA-2 is a hash of 224, 256, 384, or 512 bits. It is an improvement to SHA-1, though many experts believe it will be broken because it functions similarly to SHA-1. SHA-256 refers to SHA-2 with a 256-bit hash.

Secure Hashing Algorithm 3 (SHA-3) The winner of the NIST hash function competition from 2007 to 2012. It is not based on SHA-1 or SHA-2, but instead is a subset of the Keccak

algorithm. It creates hashes of 224, 256, 384, or 512 bits. It is recommended for use over SHA-1 and SHA-2.

Secure/Multipurpose Internet Mail Extensions (S/MIME) The standard used to encrypt and digitally sign e-mail. Almost all e-mail applications that support encryption and digital signatures use S/MIME.

Secure Real-time Transport Protocol (SRTP) A security protocol that provides confidentiality, message authentication, and replay protection for audio and video traffic, including Voice over Internet Protocol (VoIP).

Secure Shell (SSH) A basic encryption protocol used to create a secure session between two computers. SSH is used to encrypt other protocol traffic such as File Transfer Protocol (FTP) (called Secure FTP [SFTP]) and Secure Copy (SCP) traffic. It is a secure alternative to tools such as Telnet, rlogin, rsh, and rexec. SSH uses port 22.

Secure Sockets Layer (SSL) An older encryption protocol used to encrypt different types of traffic. SSL is susceptible to a POODLE (Padding Oracle On Downgraded Legacy Encryption) attack. Most organizations have stopped using SSL in favor of the designated replacement, Transport Layer Security (TLS). SSL commonly uses TCP port 443.

security policy A written document that provides the organization with a high-level view of its security goals. A security policy is authoritative in nature and provides direction for the creation of guidelines and procedures.

segmentation The process of breaking up a larger network into smaller interconnected networks. Physical segmentation uses routers to segment traffic. It's also possible to use virtual local area networks (VLANs) to logically segment traffic.

Sender Policy Framework (SPF) A specific Domain Name System (DNS) record that helps to reduce spam by identifying spoofed e-mail. SPF records identify e-mail servers authorized to send out e-mail for a domain. Compare to *Domain Keys Identified Mail (DKIM)*.

separation of duties A security principle that ensures that no single person has complete control over a process. When properly implemented, it can significantly reduce the risk of fraud within an organization.

service level agreement (SLA) An agreement between an organization and a vendor. It typically stipulates performance expectations, such as minimum uptime, and can include expectations for the availability and security of data

session key A key used with symmetric encryption. It is sometimes called a symmetric key. Compare to *private key*, *public key*, and *asymmetric key*.

shoulder surfing The practice of attempting to gain information by looking over a user's shoulder as he or she is entering data.

signature-based detection A method of virus detection used to detect known viruses. Viruses have specific characteristics that can be used to identify them uniquely. The signature can be a unique characteristic such as a specific byte pattern within the virus. Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) can also use signatures to detect known attack methods. It is sometimes called *knowledge-based detection*.

Simple Mail Transfer Protocol (SMTP) A protocol used to send e-mail. SMTP sends e-mail from e-mail clients to an e-mail server. E-mail servers use SMTP to send and receive e-mail between other SMTP servers. SMTP uses TCP port 25 or port 465 with Transport Layer Security (TLS). Compare to *Post Office Protocol version 3 (POP3)* and *Internet Message Access Protocol version 4 (IMAP4)*.

Simple Network Management Protocol (SNMP) A protocol used to manage network devices (such as routers and layer 3 switches) on a network. SNMP agents on the network devices send traps (error or event notifications) and other notification data back to a central server running SNMP management software. SNMP agents receive data on UDP 161 and send traps and notifications on UDP port 162.

single factor authentication Authentication using a single factor. There are three factors of authentication: something you know (such as a password), something you have (such as a smart card), and something you are (using biometrics). Compare to *dual factor authentication* and *multifactor authentication*.

single sign-on (SSO) The practice of allowing a subject (such as a user) to authenticate once and use the same credentials to access additional resources without authenticating again. This increases security because users have to remember only a single set of credentials and are less likely to write their credentials down. Federated access SSO systems allow users to access systems owned and managed by different organizations by logging on once using credentials recognized by the federated access system.

sniffer A protocol analyzer or packet sniffer. It can capture and analyze packets transmitted over a network. A popular sniffer is Wireshark.

social engineering The practice of using primarily nontechnical means to get people to give up sensitive data or to perform actions they wouldn't normally perform. A social engineer uses deception and fraud to trick or manipulate unsuspecting users.

software defined network (SDN) A method of replacing the data plane and the control plane on networks using virtualization technologies. An SDN can replace many routers and switches on a network.

Software-as-a-Service (SaaS) A type of cloud computing that provides one or more applications to users. Users access the applications with a web browser over the Internet. Web-based e-mail is an example of SaaS used by many people. Compare to *Infrastructure-as-a-Service (IaaS)* and *Platform-as-a-Service (PaaS)*.

something you are One of the three factors of authentication. This factor includes the use of biometrics to authenticate an individual based on a fingerprint, palm, retina, iris, or other unique characteristics.

something you have One of the three factors of authentication. This factor includes items such as proximity cards, smart cards, hardware tokens, and identification badges.

something you know One of the three factors of authentication. This factor includes knowledge such as passwords, personal identification numbers, mother's maiden name, or even personal information such as the name of your first pet.

source address affinity A configuration commonly used with load balancing in which the load balancer directs traffic to a server based on the user's IP address. Compare to *round robin*.

spam Unsolicited commercial e-mail (UCE) and other unwanted electronic messages sent in bulk. Spam is most often associated with e-mail, but can also be delivered via other means, such as instant messaging. It may include unwanted advertisements or malicious content launched from botnets.

spam filter An e-mail scanner that attempts to detect and block unsolicited e-mail (spam). Spam filters can be installed on end-user systems and e-mail servers and at the boundary of the network to scan all traffic from the Internet.

spyware Software that is installed on a user's system without the user's knowledge or consent. Spyware is considered a form of malware because it can cause damage to an individual.

SQL injection An attack that injects SQL code to read and manipulate databases. SQL injection attacks are mitigated with input validation techniques and stored procedures.

standard A proven norm or method. Standards are typically external to an organization but can influence the organization's policies, guidelines, and procedures.

star Network configuration where all computing devices are connected to a central device, such as a hub or a switch. Most networks use a star configuration, with all the devices connected to the central device with twisted pair cable. See also *bus*, *mesh*, *token ring*, and *tree*.

stateful inspection firewall A firewall that filters traffic based on the state of existing connections. It identifies active connections as they are created and monitors the status of these connections in a state table within the firewall.

steganography The practice of hiding data within data. When used to hide data within files, it modifies the least significant bits of bytes within the files.

subject An entity that can access a resource referred to as an object. For example, if a user accesses a file, the user is the subject and the file is the object.

symmetric encryption A process of encrypting and decrypting data using a single key. Symmetric encryption algorithms are constant, while the keys used to encrypt different sets of data are changed to prevent compromise. A popular symmetric encryption is Advanced Encryption Standard (AES).

systems development lifecycle (SDLC) A process used when planning, creating, testing, and deploying computing systems and computer applications. It includes five phases: design, implementation, maintenance, planning, and analysis

TACACS+ See *Terminal Access Controller Access Control System + (TACACS+)*.

TCP/IP Model A model developed by the Defense Advanced Research Projects Agency (DARPA). Some references list it with four layers (Application, Transport, Internet, and Link). Other references list it with five layers (Application, Transport, Network, Data Link, and Physical).

technical controls Security controls that are implemented through technical methods such as hardware, software, or firmware components. Compare to administrative controls and physical security controls.

Telnet An older protocol used to interact with remote systems using text-based commands. Telnet sends data (including login credentials) across the network in cleartext and is not recommended for use. Secure Shell (SSH) is a secure alternative.

Terminal Access Controller Access Control System + (TACACS+) An alternative to RADIUS to provide centralized authentication, authorization, and accounting (AAA) services for remote clients. TACACS+ provides improvements over TACACS, and both TACACS and TACACS+ use port 49. TACACS+ uses Transmission Control Protocol (TCP). TACACS uses User Datagram Protocol (UDP).

threat Any activity that can be a possible danger. When a threat exploits a vulnerability, an organization can suffer losses.

threat modeling A type of risk assessment that attempts to predict an attack. It uses various tools to model aspects of an attack against a system, a network, an application, or data.

token ring A network configuration where all computing devices are connected in a logical circle. Devices can transmit data onto the network only when they have a logical token that is passed from device to device. See also *bus*, *mesh*, *star*, and *tree*.

traffic shaping The process of delaying certain types of traffic (such as streaming audio and video traffic) so that other traffic performs better on a network.

Transmission Control Protocol (TCP) A connection-oriented protocol that provides guaranteed reliable communication for devices on a network. A three-way handshake establishes

a TCP connection. TCP uses packet sequencing, and the destination acknowledges every packet that it receives.

Transport Layer Security (TLS) The designated replacement for Secure Sockets Layer (SSL). It is based on SSL and is formally defined in RFC 5246. Almost all HTTPS sessions today are encrypted with TLS due to known vulnerabilities with SSL. TLS commonly uses TCP port 443.

trapdoor Code embedded in an application used to provide access to the application, the application's code, or its data via a covert method. Also called a *backdoor*.

tree A combination of a bus topology and a star topology. Instead of connecting multiple computer nodes in the bus configuration, it connects multiple star networks along a type of a bus network. See also *bus*, *mesh*, *star*, and *token ring*.

triage The process of prioritizing multiple events to determine the most severe. When faced with multiple injuries, medical personnel perform triage to prioritize patient treatment based on the severity of their medical conditions. When faced with multiple security incidents, IT first responders perform triage to prioritize responses based on the severity of each incident.

Trivial FTP (TFTP) A stripped-down version of FTP that provides very basic capabilities. It doesn't support authentication and doesn't use TCP to establish a session. However, it is very useful when transferring configuration files to and from network devices. TFTP uses UDP port 69.

Trojan horse Malware that looks like one thing but is something different. Users are often tricked into installing the malware, thinking that it will provide some benefit. When the application is installed, the Trojan horse also installs the malware.

Trusted Platform Module A hardware chip embedded in the computer's motherboard. It stores cryptographic keys, which the operating system can use to encrypt and decrypt data.

tuple A row storing data within a database table. Some database vendors call tuples rows, while other vendors call them tuples. Similarly, some vendors refer to database table columns as attributes.

User Datagram Protocol (UDP) A connectionless protocol that uses a best effort to send data without verification. Instead of checking to see whether a connection exists with another system before sending data, it simply sends it. In comparison, Transport Control Protocol (TCP) is a connection-oriented protocol that ensures a connection exists before sending data. As an example, Trivial File Transport Protocol (TFTP) uses UDP, but File Transport Protocol (FTP) uses TCP.

virtual desktop infrastructure (VDI) A virtual desktop hosted on a server and provided to users. A persistent desktop retains all changes made by a user. Each time a user logs on to a VDI system and makes changes, the changes are saved and retained. A non-persistent desktop is the same for all users.

virtual local area network (VLAN) A segmented LAN created on a switch. A VLAN allows an organization to segment traffic with a lot more flexibility when compared to segmenting traffic with a router.

virtual private network (VPN) A connection that provides access to a private network over a public network such as the Internet. VPNs use tunneling protocols (such as Point-to-Point Tunneling Protocol [PPTP] or Layer 2 Tunneling Protocol [L2TP]) to secure the traffic.

virus A form of malicious software (malware). It is an application or a piece of code that causes unexpected and usually negative results on computers. One of the key characteristics of a virus is that the infected file must be executed for the virus to run.

Voice over Internet Protocol (VoIP) The different technologies used to transmit multimedia and voice communications over IP networks, including the Internet.

volatile RAM Any memory that requires power to hold the data. When a system is powered down, any data (including potential forensic evidence) within volatile RAM is lost.

vulnerability A weakness. It can be a weakness in a system, a configuration, a process, hardware, software, or any other aspect of a system. If a threat can exploit a vulnerability, an organization can suffer losses.

vulnerability assessment A process used to discover vulnerabilities. Vulnerability assessments can use technical tools to scan networks and systems for vulnerabilities. They can also include nontechnical means such as social engineering tactics to determine whether employees are susceptible to social engineering attacks.

warm site An alternative location used in business continuity planning. A warm site is a compromise between a cold site and a hot site. The organization makes compromises with costs and time. Compare to *hot site*, *cold site*, and *mobile site*.

web of trust (WOT) A decentralized method of validating certificates. Individuals vouch for the trustworthiness of certificates. GNU Privacy Guard (GnuPG or GPG) and Pretty Good Privacy (PGP) can use a WOT for e-mail security. Compare to *public key infrastructure (PKI)*.

well-known port Logical port numbers from 0 to 1,023 mapped to specific protocols by the Internet Assigned Numbers Authority (IANA). For example, IANA assigned port 80 to the HTTP protocol. Port numbers are assigned as Transmission Control Protocol (TCP) ports and User Datagram Protocol (UDP) ports.

whitelist A list of authorized applications. Application whitelisting identifies specific applications that can run on a system and all other applications are blocked. Compare to *blacklist*.

Wi-Fi Protected Access (WPA/WPA2) Wi-Fi-related security protocols. WPA was introduced as a replacement for Wired Equivalent Privacy (WEP) to secure wireless networks.

WPA2 was created as a permanent replacement and is formalized as IEEE 802.11i. WPA and WPA2 can operate in either Personal mode or Enterprise mode. Personal mode uses a preshared key (PSK), and Enterprise mode uses an authentication server such as a Remote Dial-in User Authentication Service (RADIUS) or 802.1x server.

Wi-Fi Protected Access version 3 (WPA3) The next version of WPA announced by the Wi-Fi Alliance in January 2018. It is expected to provide a significant upgrade to wireless security when it is finalized and approved.

WiMAX Worldwide Interoperability for Microwave Access, the commercial name for 802.16-related products. WiMAX provides broadband wireless access over a large metropolitan area, such as a city, with speeds comparable to wired broadband access.

Wired Equivalent Privacy (WEP) The first security algorithm used for 802.11 wireless networks. WEP has several security issues and has been deprecated in favor of WPA and WPA2. It should not be used.

worm Malware that travels over the network looking for systems to infect. Worms do not require any type of interaction to execute.

zero day exploit an attack that takes advantage of unpublished vulnerabilities. In some cases, the vendor knows about the exploit but has not released a patch yet.

NOTES

The End

كتاب اسئلة المراجعة

SSCP[®]

Systems Security Certified Practitioner Practice Exams

Complete coverage of all seven domains of the SSCP exam

Includes updated exam objectives effective November 1, 2018

Filled with practice exam questions and answer explanations



Online content includes:

250 practice questions

Test engine that provides full-length practice exams and customized quizzes by chapter or exam domain



NICK MITROPOULOS

SSCP, CEH, CCNA Security

SSCP[®]

Systems Security Certified Practitioner Practice Exams

Nick Mitropoulos



New York Chicago San Francisco
Athens London Madrid Mexico City
Milan New Delhi Singapore Sydney Toronto

McGraw-Hill Education is an independent entity from (ISC)² and is not affiliated with (ISC)² in any manner. This study/training guide and/or material is not sponsored by, endorsed by, or affiliated with (ISC)² in any manner. This publication and accompanying media may be used in assisting students to prepare for the SSCP[®] exam. Neither (ISC)² nor McGraw-Hill Education warrant that use of this publication and accompanying media will ensure passing any exam. (ISC)², CISSP[®], CAP[®], ISSAP[®], ISSEP[®], ISSMP[®], SSCP[®], CSSLP[®], CCSP[®], and CBK[®] are trademarks or registered trademarks of (ISC)² in the United States and certain other countries. All other trademarks are trademarks of their respective owners.

CHAPTER 1

Access Controls

This chapter includes questions from the following objectives:

- 1.1 Implement and maintain authentication methods
- 1.2 Support internetwork trust architectures
- 1.3 Participate in the identity management lifecycle
- 1.4 Implement access controls

This chapter contains a complete set of questions and answers regarding various methods and best practices that can be used to implement user and device authentication (single/dual/multifactor authentication, SSO, federated access, centralized/decentralized authentication); different trust models (one-way, two-way, transitive trust); and network types (internet, intranet, and extranet) that can be utilized. Additionally, it contains questions regarding the identity management lifecycle (identity proofing, provisioning and authentication, maintenance and entitlement, de-provisioning), as well as the various access control models (DAC, MAC, RBAC, ABAC, subject/object based) that can be enforced.

Q QUESTIONS

1. An army officer needs to access a secure area. For that purpose, she enters her unique six-digit username, personal password, and answer to a security question. Which of the following describes the type of authentication being used?

- A. Multifactor authentication
- B. Strong authentication
- C. Single-factor authentication
- D. Two-factor authentication

2. Which of the following is indicative of a better biometric authentication system?

- A. Higher FAR
- B. Higher FRR
- C. Lower DER
- D. Lower CER

3. John is a security manager and is creating an Excel file containing his team's shift plan. He wants his team to be able to read the file in order to know what shift they are working in but doesn't want them to be able to edit the file and change shifts at will. Which of these access control models would be most suitable to accomplish this?

- A.** Mandatory
- B.** Discretionary
- C.** Rule-based
- D.** Role-based

4. The process of creating user accounts and granting them access to appropriate resources is called:

- A.** Provisioning
- B.** Authorization
- C.** Entitlement
- D.** Proofing

5. Bob wants to log on to a company system and is presented with a screen asking him to enter a username and password, followed by a second screen asking for an OTP. Which of these elements falls under identification?

- A.** Password
- B.** OTP
- C.** Username
- D.** Username and password combination

6. How many primary types of authentication factors are there?

- A.** 2
- B.** 3
- C.** 7
- D.** 4

7. Which of the following passwords is the strongest as per NIST's SP 800-63B?

A. YouCan'tTouchThi\$12

B. *iQ23!T

C. abcdefgh

D. Copenhagen

8. Sarah, a customer of Glober Bank, has successfully logged on to the bank's website and wants to complete a transaction. In order for that to be finalized, after inserting the destination account details she is required to use a device that generates an eight-digit number that she needs to provide to the website within 40 seconds of its creation time in order for the transaction to be verified and allowed. This is an example of:

A. Static password

B. Authentication

C. Synchronous dynamic password

D. Asynchronous dynamic password

9. Captain Cramer is reviewing a biometric authentication system that will be used at a nuclear missile silo. Which of the following is the most likely reason that would prevent him from approving the usage of this device?

A. Low CER

B. Increased amount of type 1 errors

C. Increased amount of type 2 errors

D. High FRR

10. What mechanism does Kerberos use to protect transmission confidentiality?

A. Symmetric encryption

B. Asymmetric encryption

C. TGT

D. KDC

11. What is the primary purpose of SSO?

- A. Authorization
- B. Confidentiality
- C. Availability
- D. Authentication

12. Which of the following is the best method of performing device authentication?

- A. MAC address filtering
- B. Dynamic IP addressing
- C. Cookies
- D. Certificate-based authentication

13. Josh is going away on a two-week holiday, but his Windows domain password expires in three days. He wants to still be able to use his laptop to access company network resources while he's away. What does he need to do?

- A. Change his domain password before it expires
- B. Call his company's IT team while away so they can have his password changed
- C. Do nothing, as he can always use his cached credentials after the first three days
- D. Log in to his laptop as a local administrator and reset his password

14. Which of these statements is true regarding DAC and non-DAC access control?

- A. In non-DAC models, users have ownership of their resources.
- B. In non-DAC models, security administrators control the access granted to users.
- C. DAC models are best for avoiding catastrophic system modifications.
- D. DAC models don't use any type of ACL.

15. Felly Corp. wants to improve the level of current security policies. As such, employees won't be allowed to connect to critical infrastructure using corporate machines from the guest network. This is an example of:

- A. Role-based Access Control (RBAC)
- B. MAC
- C. ABAC
- D. DAC

16. Nadia is on holiday and needs to access her banking website to make a transaction. She uses her friend's computer to log on to the bank's website, and after entering her credentials she is requested to enter her mother's maiden name. This is an example of:

- A. Passphrase
- B. OTP
- C. Cognitive password
- D. Static password

17. An attacker is using a brute-force tool in order to crack a user's password and gain access to online medical information. Which is the best option in order to protect against this type of attack?

- A. Account lockout policy
- B. Minimum password length
- C. Password history
- D. Maximum password age

18. Which is the best option for managing passwords?

- A. Write them down and save them in a safe to which only you have the key.
- B. Save them on a local text file.
- C. Store them on an external USB.
- D. Use KeePassXC.

19. Which is not a protocol for asynchronous dynamic password creation?

- A. HOTP

- B. OPIE**
- C. TOTP**
- D. S/KEY**

20. Elizabeth books a SSCP exam, and when she reaches the exam center she realizes that a palm vein scan is required. This is an example of which authentication factor?

- A. Something you know**
- B. Something you are**
- C. Something you have**
- D. Biometrics**

21. Which data classification scheme does the U.S. government use (listed from lowest to highest classification level)?

- A. Top Secret, Secret, Confidential, Unclassified**
- B. Unclassified, Confidential, Secret, Top Secret**
- C. Unclassified, Classified, Secret, Top Secret**
- D. Unclassified, Classified, Confidential, Top Secret**

22. Which of these statements is not correct regarding the Bell-LaPadula model?

- A. It uses the “no read up” rule.**
- B. It uses the “no write down” rule.**
- C. It is used to enforce confidentiality.**
- D. It is used to enforce integrity.**

23. Which of these models uses the “simple integrity axiom”?

- A. Biba**
- B. Clark-Wilson**
- C. Brewer-Nash**

D. Bell-LaPadula

24. A large law firm has a team (designated Alpha) representing a manufacturing client, while another team (designated Bravo) is working on a case relating to one of the manufacturing client's biggest competitors. Bravo team would greatly benefit from information that Alpha team has access to. Which of these models would be the most appropriate to ensure that a conflict of interest is avoided?

A. Bell-LaPadula

B. Biba

C. Clark-Wilson

D. Brewer-Nash

25. Dian signs in to her Gmail account and checks her e-mail. She then wants to access her YouTube channel to upload a new video, which she accesses without entering any further login details. This is an example of:

A. Federated access

B. SSO

C. Decentralized authentication

D. Two-step verification

26. Which of these combinations provides the strongest and most accurate authentication?

A. Username, cognitive password, PIN

B. Dynamic password, fingerprint

C. Retina scan, OTP

D. Symantec's VIP Access, Google Authenticator

27. Which of these statements regarding passwords is not true?

A. Use one strong password across various systems.

B. Change passwords often.

C. Never give out your password.

D. Use passphrases whenever possible.

28. Fine Investments Bank has a partnership with Gradie Security to share resources in order for the latter to manage their firewalls. They aim to exchange a lot of network documentation via a web-based portal, which will only be accessible by Gradie Security. This is an example of a(n):

- A.** DMZ
- B.** PAN
- C.** Intranet
- D.** Extranet

29. Mesquite Investment Group wants to allow their employees to browse Facebook only after their workday finishes. Which of these attributes is best to use?

- A.** Location
- B.** Group membership
- C.** Remote access
- D.** Temporal

30. A user accesses a document. Which of these statements is correct?

- A.** The user is the object.
- B.** The document is the object.
- C.** The document is the subject.
- D.** The user is both the subject and object.

31. Anemone Investment Group is using a two-step verification process (user's password combined with an OTP sent via SMS) to allow remote access to its secure file server. It is 2:00 A.M. on a Saturday, and one of their employees realizes his personal e-mail account has been compromised. He is using the same compromised credentials to access Anemone Investment Group's network remotely. The IT team only works Monday through Friday. How should he proceed?

- A.** Call the company CEO and report the issue.
- B.** Immediately drive to the office and change their password.

- C. Send an e-mail to the IT team and follow up with them on Monday morning.
- D. Delete their personal e-mail account.

32. Captain Jones has been granted top secret access. According to the Bell-LaPadula model, which rule ensures that she doesn't write information to a lower security level?

- A. Simple integrity axiom
- B. Star property
- C. Simple security property
- D. Star integrity axiom

33. Danny attempts to enter his credentials in order to log in to his corporate laptop. After entering a wrong password three times, he is unable to enter any credentials again. What is the most likely reason?

- A. Maximum password age.
- B. Password history.
- C. Account has been deleted.
- D. Account has been locked out.

34. An attacker managed to obtain the password file from Lara's laptop and is using a brute-force tool to guess her password. Which of the following is the best option in order to protect her account from this type of attack?

- A. Account lockout policy
- B. Password complexity
- C. Password history
- D. Password length

35. What is the greatest disadvantage of using a retina scanner?

- A. The scan can potentially identify and reveal the subject's medical conditions.
- B. They are often unreliable.
- C. The scan takes too long.

D. There is a large amount of false positives.

36. Which of the following isn't included in the identity and access management lifecycle?

A. Entitlement

B. Proofing

C. De-provisioning

D. Deletion

37. Which of the following is an example of a passphrase?

A. ILoveS@l@miAndCh33s3

B. P@ssw0rd

C. IUOIUQJHJH1987287!!dskUYQ62

D. passphrase

38. A home network serving the needs of a small family can be classified as a:

A. WAN

B. PAN

C. LAN

D. MAN

39. Which of the following is not an SSO technology?

A. SAML

B. SecureAuth IdP

C. Kerberos

D. OpenIDConnect

40. Which model is known as the "Chinese Wall" model?

A. Biba

B. Brewer-Nash

C. Bell-LaPadula

D. Clark-Wilson

QUICK ANSWER KEY

1. C

2. D

3. B

4. A

5. C

6. B

7. A

8. C

9. C

10. A

11. D

12. D

13. A

14. B

15. C

16. C

17. A

18. D

19. C

20. B

21. B

22. C

23. A

24. D

25. B

26. C

27. A

28. D

29. D

30. B

31. C

32. B

33. D

34. B

35. A

36. D

37. A

38. C

39. B

40. B

ANSWERS A

1. An army officer needs to access a secure area. For that purpose, she enters her unique six-digit username, personal password, and answer to a security question. Which of the following describes the type of authentication being used?

- A.** Multifactor authentication
- B.** Strong authentication
- C.** Single-factor authentication
- D.** Two-factor authentication

C. Both the previously mentioned authentication attributes (password, answer to a security question) belong to the same authentication category, which is “something you know.” As you may remember, the first step to answering any authentication factor–related question is to identify the different factors used to authenticate the user. In the example, there’s only one factor in use, therefore resulting in the system using single-factor authentication. Also note that the six-digit username is used for identification (not authentication).

A, B, and D are incorrect. **A** is incorrect because in order to have a multifactor authentication system in place, attributes would need to be present from at least two different types of factors. **B** is incorrect because there is no authentication factor known as “strong authentication.” **D** is incorrect because in order to have a two-factor authentication system in place, attributes would need to be present from exactly two different types of factors.

2. Which of the following is indicative of a better biometric authentication system?

- A.** Higher FAR
- B.** Higher FRR
- C.** Lower DER
- D.** Lower CER

D. CER (Crossover Error Rate) is where the FRR (False Rejection Rate) and FAR (False Acceptance Rate) cross over. A good biometric authentication system should have a low CER.

A, B, and C are incorrect. **A** is incorrect because FAR relates to the percentage of times a system falsely identifies an unknown user as a genuine system user who should be granted access to the system; thus, a larger FAR isn’t indicative of a good biometric system. **B** is incorrect because FRR relates to the percentage of time a system denies access to a perfectly valid system user. As in the previous case, a larger FRR is not indicative of a good biometric

system being in use. **C** is incorrect because DER is not a term used in biometric authentication systems.

3. John is a security manager and is creating an Excel file containing his team's shift plan. He wants his team to be able to read the file in order to know what shift they are working in but doesn't want them to be able to edit the file and change shifts at will. Which of these access control models would be most suitable to accomplish this?

- A. Mandatory
- B. Discretionary
- C. Rule-based
- D. Role-based

B. Discretionary access control (DAC) grants access to objects based on the identity of each subject. It also provides data owners the ability to grant permissions to subjects.

A, C, and D are incorrect. **A** is incorrect because mandatory access control (MAC) uses labels to identify subjects and objects. It is often used in military environments, as it provides the maximum level of security between all access models. **C** is incorrect because Rule-based Access Control uses rules that establish access to resources. **D** is incorrect because Role-based Access Control defines access by using roles or groups and is often incorporated by organizations with large groups of users (i.e., a group of employees that belong to the security team could utilize a role named "security members," and each time new employees are hired they can be assigned to that particular group).

4. The process of creating user accounts and granting them access to appropriate resources is called:

- A. Provisioning
- B. Authorization
- C. Entitlement
- D. Proofing

A. The definition of provisioning is "The process of creating user accounts and assigning them access to suitable resources."

B, C, and D are incorrect. **B** is incorrect because authorization refers to providing proper account privileges (for the accounts created via the provisioning process). **C** is incorrect because entitlement refers to what privileges are provided to users and ensuring those align with their role

(i.e., an employee who left the HR department and moved to accounting shouldn't have access to other employees' bonuses, since they are now performing a different function, meaning they are not entitled to that data anymore). **D** is incorrect because proofing refers to establishing an individual's identity prior to account creation.

5. Bob wants to log on to a company system and is presented with a screen asking him to enter a username and password, followed by a second screen asking for an OTP. Which of these elements falls under identification?

- A.** Password
- B.** OTP
- C.** Username
- D.** Username and password combination

C. Bob first enters a username to claim his identity. Please note that at this point he just claims to be Bob, without providing any further evidence to support that (like his password).

A, B, and D are incorrect. **A** is incorrect because the password is used to authenticate the user and not identify him. **B** is incorrect because OTP (one-time password) is used as a second step of a two-step verification process. The user would typically encounter this screen only after entering a valid username and corresponding password (meaning he has claimed he is Bob and has provided Bob's correct password). **D** is a distractor, as a username and password pair can't be used in order to provide identification of a user. Passwords are used for authentication.

6. How many primary types of authentication factors are there?

- A.** 2
- B.** 3
- C.** 7
- D.** 4

B. There are three primary types of authentication factors. Those are "something you know," "something you have," and "something you are."

A, C, and D are incorrect because there are three primary types of authentication (as mentioned earlier). However, please do note that in some sources, you might find five authentication factors mentioned, which in essence includes two additional types: "somewhere you are" and "something you do." These are less reliable than the three primary types already noted. For example, "somewhere you are" usually requires some type of geolocation mechanism

to be active and allows the user to connect to a system only if he is originating from a specific location. However, depending on the source machine configuration, that can easily be defeated by using third-party proxy software in order to “trick” the system into accepting an incoming connection from a source location that it shouldn’t.

7. Which of the following passwords is the strongest as per NIST’s SP 800-63B?

A. YouCan’tTouchThi\$12

B. *iQ23!T

C. abcdefgh

D. Copenhagen

A. “YouCan’tTouchThi\$12” is a passphrase that is in accordance with NIST’s SP 800-63B password recommendations. Actually, it also includes a combination of lowercase and uppercase letters, a special character (\$), and two numbers to make it stronger.

B, C, and D are incorrect. **B** is incorrect because although “*iQ23!T” might seem strong at first glance (having a combination of lowercase and uppercase letters, as well as numbers and special characters), but has a length of seven characters, while NIST’s recommendation is a minimum of eight. **C** is incorrect because “abcdefgh” is considered a weak password. **D** is incorrect because “Copenhagen” would most likely be present in a dictionary of common words that could be used to brute-force the user’s password.

8. Sarah, a customer of Glober Bank, has successfully logged on to the bank’s website and wants to complete a transaction. In order for that to be finalized, after inserting the destination account details she is required to use a device that generates an eight-digit number that she needs to provide to the website within 40 seconds of its creation time in order for the transaction to be verified and allowed. This is an example of:

A. Static password

B. Authentication

C. Synchronous dynamic password

D. Asynchronous dynamic password

C. A synchronous dynamic password is generated by a token, which changes the password at given intervals (in this example, every 40 seconds).

A, B, and D are incorrect. **A** is incorrect because the question suggests that the password changes every 40 seconds; thus, it can’t be a static one, as that wouldn’t change. **B** is incorrect,

as the question states Sarah has already logged on successfully to the bank's website, meaning authentication has already taken place. **D** is incorrect because an asynchronous password wouldn't need to be provided within a specific time frame (i.e., 40 seconds).

9. Captain Cramer is reviewing a biometric authentication system that will be used at a nuclear missile silo. Which of the following is the most likely reason that would prevent him from approving the usage of this device?

- A. Low CER
- B. Increased amount of type 1 errors
- C. Increased amount of type 2 errors
- D. High FRR

C. A type 2 error (also known as FAR, False Acceptance Rate) can affect a system in a more negative way, as it means that it improperly allows access to illegitimate users, which would be detrimental in the case of a nuclear silo.

A, B, and D are incorrect. **A** is incorrect because ideally a good biometric system needs to have a low CER (Crossover Error Rate), as that means it is performing better. **B** is incorrect because type 1 errors refer to FRR (False Rejection Rate), meaning the system doesn't allow access to a legitimate user. Although this is a nuisance, it doesn't even come close to how worse things can become by allowing illegitimate users to enter the secure premises (which is what answer C depicts). **D** is incorrect, as it has exactly the same meaning as B but instead of the term "type 1 error," FRR is used.

10. What mechanism does Kerberos use to protect transmission confidentiality?

- A. Symmetric encryption
- B. Asymmetric encryption
- C. TGT
- D. KDC

A. Kerberos uses symmetric encryption in order to protect the transmission confidentiality.

B, C, and D are incorrect. **B** is incorrect because not all Kerberos versions use asymmetric encryption (for example, version 5 is able to use asymmetric encryption). **C** is incorrect because TGT is a term used to indicate the time-stamped ticket (known as ticket-granting ticket) that the Kerberos server will provide to the client. **D** is incorrect because the term KDC (also known as Key Distribution Center) is used to refer to the Kerberos server via another term.

11. What is the primary purpose of SSO?

- A. Authorization
- B. Confidentiality
- C. Availability
- D. Authentication

D. SSO (single sign-on) is mainly used to identify and authenticate users. Hence, from the available answers, authentication is the correct one.

A, B, and C are incorrect. **A** is incorrect because SSO doesn't provide authorization, as it is not responsible for the services the user can access after being authenticated to a system. **B** is incorrect because SSO doesn't provide confidentiality (which is usually achieved by using encryption). **C** is incorrect because SSO doesn't provide availability, which has to do with the system being available for usage.

12. Which of the following is the best method of performing device authentication?

- A. MAC address filtering
- B. Dynamic IP addressing
- C. Cookies
- D. Certificate-based authentication

D. Out of the proposed answers, certificate-based authentication is the best method, as it entails a digital certificate to be installed on the device in order to authenticate when used to connect to a network.

A, B, and C are incorrect. **A** is incorrect because MAC addresses can be easily spoofed. **B** is incorrect because dynamic IP addressing won't help uniquely identify a device, as the IP address would be different each time. **C** is incorrect because cookies can be deleted or manipulated and don't offer a fully reliable authentication mechanism.

13. Josh is going away on a two-week holiday, but his Windows domain password expires in three days. He wants to still be able to use his laptop to access company network resources while he's away. What does he need to do?

- A. Change his domain password before it expires
- B. Call his company's IT team while away so they can have his password changed

C. Do nothing, as he can always use his cached credentials after the first three days

D. Log in to his laptop as a local administrator and reset his password

A. The only way for Josh to maintain access to company resources while he's away on holiday is to ensure he changes his password before leaving. The question doesn't mention that a VPN client is configured, thus implying that Josh's laptop has no way of connecting to the corporate network once he's away.

B, C, and D are incorrect. **B** is incorrect because even if an IT administrator changes Josh's password at the Windows server level, his laptop will have no way of connecting to the corporate network (since there's no VPN present) to get the updated password. **C** is incorrect because although Josh can use locally cached credentials after the first three days, that won't allow him access to the corporate network resources but only to the local system ones. **D** is incorrect because a local administrator would only be able to administer local accounts, not Windows domain accounts.

14. Which of these statements is true regarding DAC and non-DAC access control?

A. In non-DAC models, users have ownership of their resources.

B. In non-DAC models, security administrators control the access granted to users.

C. DAC models are best for avoiding catastrophic system modifications.

D. DAC models don't use any type of ACL.

B. In non-DAC models (i.e., MAC), security administrators control the access provided to users, allowing for very robust but less granular control.

A, C, and D are incorrect. **A** is incorrect because users have ownership of their resources in DAC systems, thus allowing for access granularity. **C** is incorrect because non-DAC models are best for avoiding catastrophic system modifications, as they provide the system owner less control, which provides greater system protection. **D** is incorrect because DAC models do use discretionary ACLs (also known as DACs).

15. Felly Corp. wants to improve the level of current security policies. As such, employees won't be allowed to connect to critical infrastructure using corporate machines from the guest network. This is an example of:

A. RBAC (Role-based Access Control)

B. MAC

C. ABAC

D. DAC

C. ABAC works by evaluating subject and object attributes (i.e., a user attempting to log in to a finance server from the guest network) and allows/rejects accordingly.

A, B, and D are incorrect. **A** is incorrect because Role-based Access Control defines access by using roles or groups and isn't suitable for this scenario. **B** is incorrect because MAC uses labels for subject and object identification. **D** is incorrect because DAC models grant access to objects based on the identity of each subject (i.e., Alice has read and write permissions to a file).

16. Nadia is on holiday and needs to access her banking website to make a transaction. She uses her friend's computer to log on to the bank's website, and after entering her credentials she is requested to enter her mother's maiden name. This is an example of:

A. Passphrase

B. OTP

C. Cognitive password

D. Static password

C. Entering the mother's maiden name is a typical example of a cognitive password, which is something that only a valid system user would know (similar to a dog's name, first school that was attended, favorite color, or city of birth). It is worth noting that in real life, the answer to such a question doesn't need to be truthful, and actually this is a good way of not allowing an unauthorized user to discover accurate information and gain access to someone's account. For example, if a user's cognitive password is his favorite color (i.e., yellow), then an attacker can identify that in little time (often enough an attacker will use open source intelligence techniques, like browsing to someone's Facebook profile where a favorite color is being depicted). However, if an answer of "fever31" is set, the chance of someone identifying that as a favorite color is very slim.

A, B, and D are incorrect. **A** is incorrect because the mother's maiden name doesn't constitute a passphrase (although arguably, based on the earlier comment about not providing true answers to the system's questions, someone might indeed use a passphrase as an answer). **B** is incorrect because an OTP is generated by some type of hardware or software token. **D** is incorrect because although the maiden name doesn't change, it is not being used as a password here. Remember that Nadia already provided her credentials to the banking system but was requested to provide a cognitive password due to the fact that she is using her friend's machine, which is a new device that the banking system hasn't encountered before, so it needs to verify the legitimacy of the connection.

17. An attacker is using a brute-force tool in order to crack a user's password and gain access to online medical information. Which is the best option in order to protect against this type of attack?

- A. Account lockout policy
- B. Minimum password length
- C. Password history
- D. Maximum password age

A. Configuring a lockout policy with a low number of attempts to enter a password (commonly three to five) will mitigate the risk of an attacker brute forcing a user's password. That means that if a wrong password is entered five consecutive times, the user's account is locked, which will require an administrator to unlock it.

B, C, and D are incorrect. **B** is incorrect because a minimum password length policy doesn't guarantee a password won't be brute forced. For example, setting a password at a minimum of seven characters doesn't guarantee that the attacker won't guess it if they have enough time and processing power. Furthermore, it doesn't guarantee that the password is strong enough, as someone might use "abcdefg," which is considered very common. **C** is incorrect because password history (i.e., the system having a record of the most recently used passwords) doesn't guarantee a strong password is being used. Someone might switch "abcdefg" with "qwertyu," again not using a strong password. **D** is incorrect because maximum password age refers to how often users are compelled to alter their account passwords, which also doesn't guarantee a strong password is in use or an attacker can't brute force it.

18. Which is the best option for managing passwords?

- A. Write them down and save them in a safe to which only you have the key.
- B. Save them on a local text file.
- C. Store them on an external USB.
- D. Use KeePassXC.

D. Using a password manager (like KeePassXC or Windows Credential Manager) ensures that you store your passwords in an encrypted database, which can only be accessed via a main password (also known as a master password). As long as you ensure your master password is strong enough, then the data is fairly secure.

A, B, and C are incorrect. **A** is incorrect because passwords should never be written down. If they are, there's a possibility they are read, regardless of where they are stored. **B** is incorrect

because saving them on a local file presents the risk of that file being read, especially if the machine is compromised or stolen. **C** is incorrect because an external USB can be read, lost, or stolen, thus compromising your passwords.

19. Which is not a protocol for asynchronous dynamic password creation?

A. HOTP

B. OPIE

C. TOTP

D. S/KEY

C. TOTP (Time-based One Time Password) is a protocol that is used for creating time-based synchronous dynamic passwords.

A, B, and D are incorrect, as all three protocols relate to asynchronous dynamic password creation.

20. Elizabeth books a SSCP exam, and when she reaches the exam center she realizes that a palm vein scan is required. This is an example of which authentication factor?

A. Something you know

B. Something you are

C. Something you have

D. Biometrics

B. A palm vein scan falls under “something you are,” as it uses a person’s unique characteristics (palm vein feature) in order to provide authentication.

A, C, and D are incorrect. **A** is incorrect because “something you know” refers to a password or something else that the user has knowledge of. **C** is incorrect because “something you have” refers to an item that the user is in possession of, like a token or card. **D** is incorrect because although a palm vein scan is considered a biometric authentication method, the authentication factor that this falls under is “something you are.” There’s no authentication factor known as biometrics.

21. Which data classification scheme does the U.S. government use (listed from lowest to highest classification level)?

A. Top Secret, Secret, Confidential, Unclassified

B. Unclassified, Confidential, Secret, Top Secret

C. Unclassified, Classified, Secret, Top Secret

D. Unclassified, Classified, Confidential, Top Secret

B. The U.S. government classification order (from lowest to highest) is Unclassified, Confidential, Secret, and Top Secret.

A, C, and D are incorrect. **A** is incorrect because it presents the classification in an order of highest to lowest. **C** and **D** are incorrect, and note that they also contain a data classification level named “Classified” that doesn’t exist as part of the actual data classification scheme used. For the U.S. government, any information that is not categorized as unclassified is considered classified, and it is assigned a level of Confidential, Secret, or Top Secret.

22. Which of these statements is not correct regarding the Bell-LaPadula model?

A. It uses the “no read up” rule.

B. It uses the “no write down” rule.

C. It is used to enforce confidentiality.

D. It is used to enforce integrity.

C. The primary goal of the Bell-LaPadula model is to enforce confidentiality.

A, B, and D are incorrect. **A** is incorrect because Bell-LaPadula does use the “no read up” rule in order to ensure that a subject that has access to a security level can’t read an object found at a higher security level (as they are not authorized to access that information). **B** is incorrect because Bell-LaPadula also uses the “no write down” rule to ensure that subjects with access to a security level don’t write to objects at a lower security level. **D** is incorrect because the primary goal of this model is not to enforce integrity, but to enforce confidentiality, which is exactly why it uses the mentioned “no read up” and “no write down” rules.

23. Which of these models uses the “simple integrity axiom”?

A. Biba

B. Clark-Wilson

C. Brewer-Nash

D. Bell-LaPadula

A. The Biba model uses the “simple integrity axiom,” which states that subjects with access to a security level can’t read a lower security level object.

B, C, and D are incorrect. **B** is incorrect because Clark-Wilson uses certification and enforcement rules for separation of duties enforcement. **C** is incorrect because Brewer-Nash works by creating a logical barrier that prevents conflicts of interest but doesn’t have anything to do with the “simple integrity axiom.” **D** is incorrect because Bell-LaPadula uses the “no read up” and “no write down” rules but not the “simple integrity axiom.”

24. A large law firm has a team (designated Alpha) representing a manufacturing client, while another team (designated Bravo) is working on a case relating to one of the manufacturing client’s biggest competitors. Bravo team would greatly benefit from information that Alpha team has access to. Which of these models would be the most appropriate to ensure that a conflict of interest is avoided?

A. Bell-LaPadula

B. Biba

C. Clark-Wilson

D. Brewer-Nash

D. The Brewer-Nash model (also known as the Chinese Wall model) is primarily used to avoid conflicts of interest. According to that, an ethical screen model is used to classify data to distinct conflict-of-interest classes. In this case, a simple way of classifying data would be to ensure that the manufacturing client is placed in one class while their competitor is in another one. Hence, if Bravo team has access to the competitor class, it can’t access data from the conflicting class of the manufacturing client.

A, B, and C are incorrect. **A** is incorrect because Bell-LaPadula is used for enforcing confidentiality. **B** is incorrect because Biba is used for enforcing integrity. **C** is incorrect because Clark-Wilson is also used to enforce integrity. Hence, the most appropriate choice is the Brewer-Nash model.

25. Dian signs in to her Gmail account and checks her e-mail. She then wants to access her YouTube channel to upload a new video, which she accesses without entering any further login details. This is an example of:

A. Federated access

B. SSO

C. Decentralized authentication

D. Two-step verification

B. SSO refers to a single user authentication instance while the system allows for the same credentials to be used for the remainder of the session (typically until the user closes his browser or logs out of the application). Dian used her credentials to log in to her Gmail account, and when she browsed to YouTube, SSO was used, which is why she didn't have to re-enter her username and password.

A, C, and D are incorrect. **A** is incorrect because federated access refers to the ability of a single login to users in different networks, but Dian is only using Google's network so no federated access is being used. **C** is incorrect because in decentralized authentication, credentials to access different systems are stored in separate databases. That means that if two systems were accessed, two different login sessions would be required. **D** is incorrect because two-step verification relates to using a second factor to authenticate to a system (i.e., user enters his username and password and then a one-time code is sent to his mobile phone, which needs to be entered into the system to complete the authentication process).

26. Which of these combinations provides the strongest and most accurate authentication?

A. Username, cognitive password, PIN

B. Dynamic password, fingerprint

C. Retina scan, OTP

D. Symantec's VIP Access, Google Authenticator

C. Using a retina scan and OTP (one-time password) entails that a two-factor authentication system is in place (combining "something you are" and "something you have"). If you take a look at the other answers, A and D only use one factor (as explained later on). However, B also mentions two factors of authentication being in place: a dynamic password (which is another way of describing OTP) and a fingerprint. Note that a retina scan is considered one of the most accurate biometric authentication types and is certainly more accurate than a fingerprint, thus making it the best option.

A, B, and D are incorrect. **A** is incorrect because a cognitive password and PIN both belong to the same authentication factor, which is "something you know," making this a weak option (note that the username provides identification and not authentication). **B** is incorrect because it does entail two factors of authentication (dynamic password and fingerprint), which would be equivalent to what C mentions (OTP and retina scan). However, as a retina scan is more accurate, this is not the best answer. **D** is incorrect because Symantec's VIP Access and Google Authenticator both belong to the same authentication factor, that being "something you have," not making this the best answer.

27. Which of these statements regarding passwords is not true?

A. Use one strong password across various systems.

B. Change passwords often.

C. Never give out your password.

D. Use passphrases whenever possible.

A. Although the answer mentions a strong password is in use, that doesn't make it impossible to be intercepted, stolen, or provided unintentionally to an attacker. If that password is used across various systems, the malicious party will have access to everything, which could prove quite damaging for an organization.

B, C, and D are incorrect. B is incorrect because it is a good practice to change passwords often. C is incorrect, as it is highly recommended not to give your password out to anyone. D is incorrect because using passphrases is a great way of creating secure passwords.

28. Fine Investments Bank has a partnership with Gradie Security to share resources in order for the latter to manage their firewalls. They aim to exchange a lot of network documentation via a web-based portal, which will only be accessible by Gradie Security. This is an example of a(n):

A. DMZ

B. PAN

C. Intranet

D. Extranet

D. An extranet is a network that uses the Internet to host resources and would be the best way for Fine Investments Bank to provide documentation and other internal information to Gradie Security.

A, B, and C are incorrect. A is incorrect because a DMZ hosts resources that are available to everyone on the Internet, while an extranet is used to exchange information only with a trusted third party. B is incorrect because a PAN is a very small network, commonly used by an individual (i.e., creating a PAN between a laptop, tablet, and a phone). C is incorrect because an intranet is the internal network of an organization and would never be accessible by a third party (like Gradie Security).

29. Mesquite Investment Group wants to allow their employees to browse Facebook only after their workday finishes. Which of these attributes is best to use?

A. Location

B. Group membership

C. Remote access

D. Temporal

D. Assuming that the working day runs from 9:00 A.M. to 5:00 P.M., then Mesquite Investment Group wants to only allow Facebook usage out of those hours, which is something that can easily be done using the temporal attribute.

A, B, and C are incorrect. **A** is incorrect because the location is irrelevant, as the question doesn't limit this to one company branch (i.e., the Chicago branch). **B** is incorrect, as this refers to groups of subjects that are granted access to objects (i.e., admins, local admins, users). **C** is incorrect because the remote access attribute refers to allowing users to remotely connect to the company network.

30. A user accesses a document. Which of these statements is correct?

A. The user is the object.

B. The document is the object.

C. The document is the subject.

D. The user is both the subject and object.

B. An object is the resource being accessed; hence, the document is the object.

A, C, and D are incorrect. **A** is incorrect because the user is the one that accesses a resource, thus making him the subject. **C** is incorrect because a subject accesses a resource, but here the document is the resource being accessed. **D** is incorrect because the user is the one accessing the document, meaning he is strictly the subject, not both subject and object.

31. Anemone Investment Group is using a two-step verification process (user's password combined with an OTP sent via SMS) to allow remote access to its secure file server. It is 2:00 A.M. on a Saturday, and one of their employees realizes his personal e-mail account has been compromised. He is using the same compromised credentials to access Anemone Investment Group's network remotely. The IT team only works Monday through Friday. How should he proceed?

A. Call the company CEO and report the issue.

B. Immediately drive to the office and change their password.

C. Send an e-mail to the IT team and follow up with them on Monday morning.

D. Delete their personal e-mail account.

C. One of the greatest advantages of using a two-factor authentication scheme is that even if the user's password is compromised, the attacker would still need an OTP (which in this case he can only get if he obtains the user's mobile phone) to access the account. Based on that, the user can safely notify the IT team via e-mail at the time he was made aware of the issue and be diligent and follow up with them as soon as they are available.

A, B, and D are incorrect. **A** is incorrect because calling the company CEO to report an issue like this is out of his remit and a bit of an overkill. Reporting this to the IT team is the best way to move forward, as they are responsible for resetting or monitoring the user's account, checking access logs, and performing other similar activities. **B** is incorrect because (as mentioned earlier) a two-factor authentication system is in place; hence, there's no need to drive to the office in the middle of the night to change the account password. **D** is incorrect because even if the user deletes his personal e-mail account, this will have no effect on enhancing the corporate account's security.

32. Captain Jones has been granted top secret access. According to the Bell-LaPadula model, which rule ensures that she doesn't write information to a lower security level?

A. Simple integrity axiom

B. Star property

C. Simple security property

D. Star integrity axiom

B. The star property (also known as "no write down") ensures that subjects with access to a security level don't write to objects at a lower security level.

A, C, and D are incorrect. **A** is incorrect because it's a rule encountered in the Biba model, which doesn't allow the reading of objects at a lower security level. **C** is incorrect because the simple security property (also known as "no read up") ensures that a subject that has access to a security level can't read an object found at a higher security level. **D** is incorrect because the star integrity axiom is used in the Biba model and relates to not allowing subjects who are granted access at a security level to write to an object of a higher security level.

33. Danny attempts to enter his credentials in order to log in to his corporate laptop. After entering a wrong password three times, he is unable to enter any credentials again. What is the most likely reason?

A. Maximum password age.

B. Password history.

C. Account has been deleted.

D. Account has been locked out.

D. After entering a wrong password three times, Danny's account most likely was locked out in order to prevent a brute-force attack from taking place.

A, B, and C are incorrect. **A** is incorrect because the maximum password age defines how often users are obligated to change their account passwords (i.e., every ten days). **B** is incorrect because password history relates to changing a user's password. The system remembers previous passwords that a user has used (i.e., last eight) in order to prevent him from using the same password again. **C** is incorrect because an account wouldn't normally be deleted just because a user entered a wrong password.

34. An attacker managed to obtain the password file from Lara's laptop and is using a brute-force tool to guess her password. Which of the following is the best option in order to protect her account from this type of attack?

A. Account lockout policy

B. Password complexity

C. Password history

D. Password length

B. Password complexity is the best weapon against offline brute-force attacks. The more complex the password is, the less chance a brute-force attack has of identifying it.

A, C, and D are incorrect. **A** is incorrect because an account lockout policy would be very beneficial if this was an online brute-force attack. However, the attacker is using an offline tool while having the password file at his disposal, hence there's no account that can be locked out. **C** is incorrect because password history (as already mentioned in the previous question) relates to changing a user's password. **D** is incorrect because the password length doesn't ensure a strong password was used. For example, even if the password length is seven characters, if the password of choice is "hackers," then there's a higher chance of the password being identified by an attacker using any standard brute-force dictionary.

35. What is the greatest disadvantage of using a retina scanner?

A. The scan can potentially identify and reveal the subject's medical conditions.

B. They are often unreliable.

C. The scan takes too long.

D. There is a large amount of false positives.

A. The biggest disadvantage of a retina scanner is the fact that a scan can identify and reveal a subject's medical conditions, which makes it less popular to the majority of the population.

B, C, and D are incorrect. **B** is incorrect because retina scans are very reliable and are actually considered one of the best biometric authentication methods. **C** is incorrect because the subject's identity is verified very quickly. **D** is incorrect because retina scans have a very low rate of false positives (as well as false negatives).

36. Which of the following isn't included in the identity and access management lifecycle?

A. Entitlement

B. Proofing

C. De-provisioning

D. Deletion

D. There's no term known as deletion in the identity and access management lifecycle. Account deletion will normally take place at the de-provisioning phase.

A, B, and C are incorrect. **A** is incorrect because entitlement refers to what privileges are provided to users and ensuring those align with their role. **B** is incorrect because proofing refers to establishing an individual's identity prior to account creation. **C** is incorrect because de-provisioning relates to ensuring that inactive accounts are disabled and deleted.

37. Which of the following is an example of a passphrase?

A. ILoveS@l@miAndCh33s3

B. P@ssw0rd

C. IUOIUQHJH1987287!!dslkUYQ62

D. passphrase

A. A passphrase is defined as a lengthy character sequence that is meaningful to the user. ILoveS@l@miAndCh33s3 is a passphrase derived from "ILoveSalamiAndCheese" where the letter "a" has been replaced with the character "@" and the letter "e" with "3" in order to provide enhanced security.

B, C, and D are incorrect. **B** is incorrect because P@ssw0rd is just a password containing a combination of lowercase and uppercase letters, numeric digits, and special characters. That

doesn't constitute a passphrase. **C** is incorrect because IUOIUQJHJH1987287!!dskUYQ62 might again be a lengthy password (even longer than the passphrase used in option A); however, this doesn't have any meaning at all to the user. It is just random characters, and it would be impossible to remember. **D** is incorrect because, similarly to B, "passphrase" is just a word (which would be a poor password) and doesn't constitute a passphrase.

38. A home network serving the needs of a small family can be classified as a:

A. WAN

B. PAN

C. LAN

D. MAN

C. Home networks can be classified as LANs, as they are a group of connected devices within a small geographical area, which in this case is the user's residence.

A, B, and D are incorrect. **A** is incorrect because a WAN would be used for connecting devices across vast areas. **B** is incorrect because a PAN is a very small network, commonly used by an individual, and it wouldn't be suitable for a home network. **D** is incorrect because a MAN is a collection of various networks spanning over large geographical areas, like a large corporate complex spread across several miles.

39. Which of the following is not an SSO technology?

A. SAML

B. SecureAuth IdP

C. Kerberos

D. OpenIDConnect

B. SecureAuth IdP provides mobile device management and doesn't have anything to do with SSO.

A, C, and D are incorrect. **A** is incorrect because SAML is used to provide SSO. **C** is incorrect because Windows domains and Linux/UNIX realms do use Kerberos for SSO. **D** is incorrect because OpenIDConnect is an SSO technology.

40. Which model is known as the "Chinese Wall" model?

A. Biba

B. Brewer-Nash

C. Bell-LaPadula

D. Clark-Wilson

B. The Brewer-Nash model is also known as the Chinese Wall model.

A, C, and D are incorrect. All other models (Biba, Bell-LaPadula, Clark-Wilson) are unrelated to the Chinese Wall model.

CHAPTER 2

Security Operations and Administration

This chapter includes questions from the following objectives:

- 2.1 Comply with codes of ethics
- 2.2 Understand security concepts
- 2.3 Document, implement, and maintain functional security controls
- 2.4 Participate in asset management
- 2.5 Implement security controls and assess compliance
- 2.6 Participate in change management
- 2.7 Participate in security awareness and training
- 2.8 Participate in physical security operations (e.g., data center assessment, badging)

This chapter contains a complete set of questions and answers regarding the (ISC)² code of ethics; basic security concepts (confidentiality, integrity, availability); security control types (preventive, detective, corrective); asset management; security controls implementation and compliance assessment; change management; and participation in physical security operations. In addition, there's variety of questions focusing on security awareness and training (attack types, malicious code, performing security audits, and understanding security policies).

Q QUESTIONS

- 1.** Ben is a system administrator, and he was told that users are experiencing slow response times when trying to access the corporate website. Upon investigation, he notices that the web server shows a high number of established connections and thinks an attack might be taking place. Based on this, which of the following statements is correct?
- A.** The website is experiencing a DoS attack.
 - B.** A DDoS attack is underway.
 - C.** The server is experiencing increased traffic.
 - D.** A port scan is taking place.

2. Jason is the CFO of sales, and he just received an e-mail inviting him to access a Dropbox URL sharing financial projections for the next quarter. The URL is <http://xteyjsurls.com>. This is an example of:

- A.** A valid method of sharing company financial data
- B.** Spear phishing
- C.** Vishing
- D.** Whaling

3. Which of the following is classified as PHI?

- A.** Company name
- B.** Business address
- C.** Fingerprint
- D.** Healthcare provider

4. Next Gemini Bank is developing a new proprietary algorithm. What is the most suitable data classification label for that information?

- A.** Proprietary
- B.** Confidential
- C.** Public
- D.** Sensitive

5. Don is on his way to work and wants to review the material for the morning meeting. For that purpose, he powers on his laptop and opens the spreadsheet application to access the related file. What is this an example of?

- A.** Data in use
- B.** Data in motion
- C.** Data in transit
- D.** Data at rest

6. In which corporate policy are these statements most likely to be present?

Employees are not allowed to access social media sites.

Company information is not allowed to be stored in USB drives.

Client data must never be provided to any individual who is not a corporate employee.

It is not allowed to connect any personal devices to the company network.

- A. Backup policy
- B. Network access policy
- C. Guest access policy
- D. Acceptable use policy

7. A network administrator is considering placing a new firewall at the network perimeter. Which type of control is this?

- A. Detective
- B. Corrective
- C. Preventive
- D. Recovery

8. Which of these security practices is not a part of system hardening?

- A. Update operating system
- B. Enable appropriate services
- C. Enable firewall
- D. Remove unused protocols

9. James is searching for a way to apply critical patches across his company's network. Which of the following would be inappropriate for that purpose?

- A. TPM
- B. WSUS
- C. SCCM
- D. ConfigMgr

10. Dianne is working on a critical system implementation and needs to ensure that proper audit logging is in place to be able to account for any administrator action. This is an example of:

- A. Authentication
- B. Integrity
- C. Nonrepudiation
- D. Confidentiality

11. Which of these is considered a detective control?

- A. IPS
- B. IDS
- C. Security camera
- D. Disaster recovery plan

12. Jane is a system administrator and has chosen to allow company users to connect USB flash drives in read-only mode (i.e., users can only copy data from flash drives to company devices) so data loss can be avoided. What is the biggest disadvantage of that?

- A. Company employees don't have any way of sharing information.
- B. Malware will still be copied from any infected corporate machine to the USB flash drives.
- C. Standard users (non-administrator) can easily bypass the read-only setting.
- D. Company antivirus software won't be able to remove any infected files that exist on the flash drives.

13. Janet is a network administrator. Upon reviewing records of a recent network change, she realizes that one of her colleagues made a mistake and, in an attempt to cover his tracks, he deleted the related audit logs. How should Janet proceed?

- A. Do nothing in order not to create tension in the workplace.
- B. Confront her colleague by telling him she knows what happened so he should report it to the team's manager himself.
- C. Report the incident to her manager.

D. Help her colleague recover the audit logs from the related backups and ask him to never let this happen again.

14. A security analyst is investigating a possible network misconfiguration issue, resulting in increased internal network traffic. The offending machine's IP and MAC addresses have been identified, but the analyst needs to trace those to the machine owner in order to verify the legitimacy of the traffic. What is the best way of achieving this?

A. Search the company's incident management software to identify cases that involve the same IP/MAC combination.

B. Check the asset register.

C. Connect to the host in question via RDP to get the necessary information.

D. Block the host from the network to stop a further increase in traffic.

15. Which of these controls is the most appropriate to stop tailgating?

A. Security camera

B. Access badge

C. Mantrap

D. Physical security policy

16. How often do PCI merchants need to perform network scans according to PCI-DSS requirements?

A. Monthly

B. Quarterly

C. Biannually

D. Annually

17. Which of the following should a large bank's mergers and acquisitions team ensure has taken place before a major acquisition of an investment firm happens?

A. Due diligence

B. Due care

C. Separation of duties

D. Confidentiality

18. A company backup policy primarily uses magnetic tapes for backups. They also regularly back up data to the company cloud. Which type of control does the company cloud represent?

- A. Detective
- B. Deterrent
- C. Recovery
- D. Compensating

19. Which of the following is not a security policy characteristic?

- A. Mission statement
- B. Enforcement section
- C. Sanctions statement
- D. Accountability statement

20. Which of these is not included in asset management?

- A. Hardware
- B. People
- C. Data
- D. Software

21. What is the primary goal of change management?

- A. Review changes
- B. Minimize disruption
- C. Document changes
- D. Provide resource allocation

22. You are tasked with enhancing your company's physical security controls in order to ensure a proper audit trail exists. Which action would accomplish that goal?

- A. Place access card readers at the interior of secure operation rooms.
- B. Replace external 16-foot fences with taller ones.
- C. Increase security guard presence to achieve patrolling in pairs.
- D. Suggest hiring an external party to perform a physical security assessment.

23. What process would take place if an (ISC)² member doesn't comply with the (ISC)² code of ethics?

- A. Any violators will be reported to the company the member is working for.
- B. (ISC)² will revoke their certifications.
- C. A review from a peer panel will take place.
- D. They will get a one-off warning.

24. Which of the following is not a step of the change management process?

- A. Change approved or rejected
- B. Change resolved
- C. Change reviewed
- D. Change submitted

25. A system administrator is installing an IPS, a firewall, and a vulnerability scanner while at the same time drafting new security policies and training materials for his company. Which best describes his goal?

- A. Confidentiality
- B. Due diligence
- C. Accountability
- D. Defense in depth

26. In order for a merchant to comply with PCI-DSS, a network scan is required to be completed by which of the following?

- A. SAQ

B. ASV

C. QSA

D. ISA

27. Which of the following changes wouldn't require CAB approval?

A. Standard

B. Emergency

C. Normal

D. Expedited

28. An attacker manages to compromise a critical server, with no central logging capability enabled. After obtaining the information of interest, the audit logs are accessed and the related entries showing what actions the attacker performed are removed. Which of the following is being compromised?

A. Least privilege

B. Availability

C. Accountability

D. Authentication

29. Employee badges are primarily used for:

A. Identification

B. Authentication

C. Provisioning

D. Entitlement

30. John powers on his computer and successfully installs a new application. After a few minutes he realizes that he's unable to launch the newly installed application, as that is being blocked by the antivirus program. What would the antivirus most probably be using to block the application at that point?

A. Antivirus blacklisting

- B. Scheduled scan
- C. On-demand scan
- D. Heuristic-based detection

31. Which of the following is most suitable to protect against phishing attacks?

- A. Spam filtering provider
- B. User education
- C. E-mail gateway
- D. Antivirus

32. Which of the following wouldn't be used during an SQL injection attack?

- A. SELECT
- B. --
- C. <>
- D. ;

33. An attacker compromises a fashion retailer's website and places malware on its web page. The goal is for the malware to be downloaded by clients that visit the website to place future orders. This is an example of:

- A. Malvertising
- B. Trapdoor
- C. Drive-by download
- D. Rootkit

34. Which of these statements about configuration management is not accurate?

- A. Puppet uses Python.
- B. Microsoft offers "Group Policy" for system configuration.
- C. Imaging is a commonly used method.

D. Ansible can be used in a Linux environment.

35. In which of these control goal and class combinations does a motion sensor fall into?

A. Preventive, technical

B. Detective, technical

C. Preventive, physical

D. Detective, physical

36. Which of the following is not included in the (ISC)² code of ethics?

A. Act honorably.

B. Provide adequate service.

C. Advance and protect the profession.

D. Provide diligent service to principals.

37. Which of the following can be part of more than one control type (i.e., combine being detective and preventive)?

A. Security policy

B. IDS

C. IPS

D. Audit logs

38. A secure facility requires individuals to sign in at the reception desk and swipe their cards at the card readers. However, the security personnel do an end-of-shift check and realize that there's a mismatch between the two records, showing 150 people signing in at the reception desk but only 140 swiping in. What is the most likely reason for that?

A. People signing in at reception more than once.

B. Tailgating has occurred.

C. Some employees signed in at reception but entered the building via nonstandard entrances.

D. Card reader system experienced an error.

QUICK ANSWER KEY

1. C

2. D

3. D

4. B

5. A

6. D

7. C

8. B

9. A

10. C

11. B

12. D

13. C

14. B

15. C

16. B

17. A

18. D

19. C

20. B

21. B

22. A

23. C

24. B

25. D

26. B

27. A

28. C

29. A

30. D

31. B

32. C

33. C

34. A

35. D

36. B

37. C

38. B

ANSWERS A

1. Ben is a system administrator, and he was told that users are experiencing slow response times when trying to access the corporate website. Upon investigation, he notices that the web server shows a high number of established connections and thinks an attack might be taking place. Based on this, which of the following statements is correct?

- A.** The website is experiencing a DoS attack.
- B.** A DDoS attack is underway.
- C.** The server is experiencing increased traffic.
- D.** A port scan is taking place.

C. The question states that there is a high number of established connections. However, that is not necessarily an indicator that an attack is taking place. Without performing additional investigation, this would most likely be considered a time when the server is experiencing a peak traffic flow, which would account for the increased traffic.

A, B, and D are incorrect. **A** is incorrect because if a DoS (denial of service) attack was underway, users wouldn't be able to browse the website or their experience would be severely degraded, whereas at this time there's just a report of website slowness (also note that at this time, the slowness is just a user's subjective opinion, as no investigation has been performed to corroborate that claim). **B** is incorrect because for a DDoS (distributed denial of service) attack to take place, traffic would need to originate from different sources, which are all sending multiple packets to the web server. **D** is incorrect because a port scan would try to identify open ports on a target machine (while there's no indication of that here), and in addition, an attacker performing a port scan would typically try to remain undetected.

2. Jason is the CFO of sales, and he just received an e-mail inviting him to access a Dropbox URL sharing financial projections for the next quarter. The URL is <http://xteyjsurls.com>. This is an example of:

A. A valid method of sharing company financial data

B. Spear phishing

C. Vishing

D. Whaling

D. Whaling is a specific category of spear phishing where an attacker targets a C-level company executive (like CEO, CFO, or COO). In this example, the CFO (chief financial officer) was enticed to click a URL posing to be a Dropbox link. Upon a simple URL review, it can be noticed that <http://xteyjsurls.com> has nothing to do with Dropbox.

A, B, and C are incorrect. **A** is incorrect because even if the Dropbox link was genuine, it wouldn't be the preferred method of sharing sensitive company data. **B** is incorrect because "spear phishing" usually targets a team of people or a whole organization. Even if one specific individual is targeted via a spear-phishing e-mail, when that person is a high-level executive, the appropriate term used to describe this attack is "whaling." **C** is incorrect because "vishing" is a social engineering technique that uses the phone (not e-mail) as an attack vector.

3. Which of the following is classified as PHI?

A. Company name

B. Business address

C. Fingerprint

D. Healthcare provider

D. An individual's healthcare provider is considered PHI (protected health information), as it's classified as personal information about someone's health.

A, B, and C are incorrect. **A** is incorrect because a company name is publicly available. **B** is incorrect because a business address should be publicly available so potential clients can reach the business in question. **C** is incorrect because biometric information (like a fingerprint) is classified as PII (personally identifiable information).

4. Next Gemini Bank is developing a new proprietary algorithm. What is the most suitable data classification label for that information?

A. Proprietary

B. Confidential

C. Public

D. Sensitive

B. A newly designed proprietary algorithm would be classified as highly important information and would need to have the highest classification assigned. Given the possible answers, that would be "confidential."

A, C, and D are incorrect. **A** is incorrect because although the algorithm is proprietary in nature, an organization wouldn't typically have a label named "proprietary." Labels should be generic and depict the protection level the data requires. **C** is incorrect because highly sensitive information would never be classified as "public." **D** is incorrect because "sensitive" is not strong enough to depict the importance of proprietary data as well the information leakage impact at that level.

5. Don is on his way to work and wants to review the material for the morning meeting. For that purpose, he powers on his laptop and opens the spreadsheet application to access the related file. What is this an example of?

A. Data in use

B. Data in motion

C. Data in transit

D. Data at rest

A. When Don opened his spreadsheet application and accessed the related file, it was transferred into his computer's memory. That data (residing in memory or being processed) is "data in use."

B, C, and D are incorrect. **B** is incorrect because "data in motion" refers to a network being used for data transmission. **C** is incorrect because "data in transit" is just another term used to describe "data in motion." **D** is incorrect because the term "data at rest" in this example would be used to describe the data file stored on the hard drive before being opened and accessed by Don.

6. In which corporate policy are these statements most likely to be present?

Employees are not allowed to access social media sites.

Company information is not allowed to be stored in USB drives.

Client data must never be provided to any individual who is not a corporate employee.

It is not allowed to connect any personal devices to the company network.

A. Backup policy

B. Network access policy

C. Guest access policy

D. Acceptable use policy

D. An acceptable use policy defines what is and is not considered proper use of organizational hardware and information assets. All new employees must usually confirm they are aware of the policy and agree to adhere to it before commencing their work.

A, B, and C are incorrect. **A** is incorrect because a backup policy defines organizational backup procedures and would include items such as how often company assets need to be backed up, what type of backup method is to be used (e.g., incremental or full backup), and where the data should be backed up to (e.g., cloud or off-site location). **B** is incorrect because a network access policy refers to the rights and privileges of specific subjects when they gain access to the corporate network. This policy would most likely include information about company VPN connectivity, type of information that can be accessed through a workstation, authentication techniques, and access control mechanisms. **C** is incorrect because a guest access policy states the company's policy for allowing guests to connect to the corporate network. Additionally, the question mentions connecting devices to the company network and accessing client data, both of which wouldn't be related to a guest access policy.

7. A network administrator is considering placing a new firewall at the network perimeter. Which type of control is this?

A. Detective

B. Corrective

C. Preventive

D. Recovery

C. A firewall is classified as a preventive control, as it has the ability to prevent an attack by stopping offending traffic.

A, B, and D are incorrect. **A** is incorrect because a detective control would only be able to identify an attack when it's taking place or after it has happened. **B** is incorrect because a corrective control would perform an activity to counter the attack. **D** is incorrect because a recovery control aims to restore a system after an incident has happened.

8. Which of these security practices is not a part of system hardening?

A. Update operating system

B. Enable appropriate services

C. Enable firewall

D. Remove unused protocols

B. Enabling appropriate services wouldn't take place during a system hardening. On the contrary, during such an activity, services that are not absolutely needed would be disabled.

A, C, and D are incorrect. **A** is incorrect because updating the operating system is something highly recommended during a hardening process. **C** is incorrect because if a system does have an available firewall, this would definitely be enabled during this activity. **D** is incorrect because any unused protocols would be removed in order not to be exploited by an attacker.

9. James is searching for a way to apply critical patches across his company's network. Which of the following would be inappropriate for that purpose?

A. TPM

B. WSUS

C. SCCM

D. ConfigMgr

A. TPM (Trusted Platform Module) is a hardware chip used by various computer models (like Microsoft Surface Pro devices) for full disk encryption and has nothing to do with patching.

B, C, and D are incorrect. **B** is incorrect because WSUS (Windows Server Update Services) is often used for patch management. **C** is incorrect because SCCM (System Center Configuration Management) is also used for patching. **D** is incorrect because ConfigMgr (also known as Configuration Manager) provides patch management.

10. Dianne is working on a critical system implementation and needs to ensure that proper audit logging is in place to be able to account for any administrator action. This is an example of:

- A. Authentication
- B. Integrity
- C. Nonrepudiation
- D. Confidentiality

C. Nonrepudiation is the most suitable term in this case, as it refers to an individual not being able to refuse an action that was taken. Enabling audit logs will allow tracking of related administrator actions.

A, B, and D are incorrect. **A** is incorrect because authentication refers to verifying an individual's identity. **B** is incorrect because integrity refers to being able to detect any unauthorized data alteration. **D** is incorrect because confidentiality refers to information being accessible only by authorized individuals.

11. Which of these is considered a detective control?

- A. IPS
- B. IDS
- C. Security camera
- D. Disaster recovery plan

B. An IDS (intrusion detection system) is considered a detective control, as it only has the ability to provide event detection.

A, C, and D are incorrect. **A** is incorrect because an IPS (intrusion prevention system) is able to prevent an attack, as it can drop offending traffic before reaching the target network. **C** is incorrect because a security camera aims to prevent offenders from attempting to enter the secure premises. **D** is incorrect because a disaster recovery plan is considered a corrective control, as it's unable to detect an incident but can help restore the system in the state that it was prior to the incident taking place.

12. Jane is a system administrator and has chosen to allow company users to connect USB flash drives in read-only mode (i.e., users can only copy data from flash drives to company devices) so data loss can be avoided. What is the biggest disadvantage of that?

- A.** Company employees don't have any way of sharing information.
- B.** Malware will still be copied from any infected corporate machine to the USB flash drives.
- C.** Standard users (non-administrator) can easily bypass the read-only setting.
- D.** Company antivirus software won't be able to remove any infected files that exist on the flash drives.

D. Company antivirus software that identifies malicious files on any flash drives won't be able to delete those, as it would need "write permissions" on the flash drives to do so.

A, B, and C are incorrect. **A** is incorrect because although using flash drives is a very common way of file sharing, it is not the only one, or the most secure for that matter. Employees should be prompted to use alternative methods, such as secure company network shares or company cloud solutions supporting adequate data encryption. **B** is incorrect because no data from a corporate machine can be copied to the flash drives (since write permission has been disabled). **C** is incorrect because only administrators have adequate permissions to edit the related policy.

13. Janet is a network administrator. Upon reviewing records of a recent network change, she realizes that one of her colleagues made a mistake and, in an attempt to cover his tracks, he deleted the related audit logs. How should Janet proceed?

- A.** Do nothing in order not to create tension in the workplace.
- B.** Confront her colleague by telling him she knows what happened so he should report it to the team's manager himself.
- C.** Report the incident to her manager.
- D.** Help her colleague recover the audit logs from the related backups and ask him to never let this happen again.

C. According to the second cannon of the (ISC)² code of ethics, a member should "act honorably, honestly, justly, responsibly and legally." In accordance with that, Janet should report the incident to her manager as soon as possible in order for appropriate actions to be taken.

A, B, and D are incorrect. **A** is incorrect because doing nothing is not an option when employee misconduct has taken place. **B** is incorrect because taking matters into one's own hands and confronting others may have severe consequences even though the intention may be

correct. **D** is incorrect because being complicit and helping to cover up what has happened is not acting “honorably and justly,” not to mention it’s illegal.

14. A security analyst is investigating a possible network misconfiguration issue, resulting in increased internal network traffic. The offending machine’s IP and MAC addresses have been identified, but the analyst needs to trace those to the machine owner in order to verify the legitimacy of the traffic. What is the best way of achieving this?

A. Search the company’s incident management software to identify cases that involve the same IP/MAC combination.

B. Check the asset register.

C. Connect to the host in question via RDP to get the necessary information.

D. Block the host from the network to stop a further increase in traffic.

B. Checking the company’s asset register (or asset management software) is the best way to obtain information about any type of equipment an organization has, as well as to identify the owner of a device. Devices such as mobile phones, computers, routers, etc., should all be present in the asset register.

A, C, and D are incorrect. **A** is incorrect because searching past cases might not provide up-to-date information. For example, just because three months ago an incident was raised about the machine in question, which indicated it was owned by Jane, doesn’t make that accurate, as Jane might not be working for the company anymore. **C** is incorrect because it means taking a form of offensive action, as the analyst would be trying to connect to another individual’s machine, having access to all of the information residing in that machine. **D** is incorrect because an increase in traffic doesn’t provide enough evidence to warrant blocking the host machine. That could very well be a critical company server, and blocking its operation might have severe consequences for the business.

15. Which of these controls is the most appropriate to stop tailgating?

A. Security camera

B. Access badge

C. Mantrap

D. Physical security policy

C. A mantrap is one of the best mechanisms that can be used to avoid tailgating, as it can control the flow of people entering a building or secure area.

A, B, and D are incorrect. **A** is incorrect because although a security camera might record tailgating, the odds of someone monitoring that in real time at peak hours and acting each time a tailgater is present are very low. **B** is incorrect because access badges allow access to the building or area, but that doesn't stop a valid employee (in possession of an access badge) to tailgate one of his colleagues. **D** is incorrect because even if there's an appropriate organizational policy in effect, that is not enough to stop people from violating it. Only physical restrictions (i.e., a mantrap, turnstile) will be able to do that.

16. How often do PCI merchants need to perform network scans according to PCI-DSS requirements?

- A. Monthly
- B. Quarterly
- C. Biannually
- D. Annually

B. PCI defines four merchant levels: 4 (lowest) to 1 (highest). In all levels, quarterly scans are required.

A, C, and D are incorrect as they all present wrong answers. As mentioned, quarterly scans are mandatory for all four merchant levels.

17. Which of the following should a large bank's mergers and acquisitions team ensure has taken place before a major acquisition of an investment firm happens?

- A. Due diligence
- B. Due care
- C. Separation of duties
- D. Confidentiality

A. "Due diligence" refers to all the necessary actions the bank will need to take to verify that the investment firm and its profits are legal in order to ensure the legality of the acquisition. This is often a task outsourced to business intelligence firms, which perform an investigation and provide a full report to the requestor (i.e., the bank) raising anything that might need to be clarified or further investigated before the transaction happens.

B, C, and D are incorrect. **B** is incorrect because "due care" refers to ensuring appropriate actions are taken to protect company assets. **C** is incorrect because "separation of duties" refers

to not allowing one individual to possess full control of a critical process. **D** is incorrect because confidentiality refers to information being accessible only by authorized individuals.

18. A company backup policy primarily uses magnetic tapes for backups. They also regularly back up data to the company cloud. Which type of control does the company cloud represent?

- A. Detective
- B. Deterrent
- C. Recovery
- D. Compensating

D. The cloud provider is being used in case the backup tape (which is primarily used as a backup media) fails, thus making it a compensating control.

A, B, and C are incorrect. **A** is incorrect because a detective control has the ability to provide event detection, which is not true in the case of a backup to the cloud. **B** is incorrect because a deterrent control aims to deter an attacker from performing a hostile action. **C** is incorrect because a recovery control aims to restore a system after an incident has happened. Please note that the question focuses on what type of control the company cloud is, making the correct answer a “compensating” one. However, if the question focused on what type of control a backup policy is, then the correct answer would be a recovery control.

19. Which of the following is not a security policy characteristic?

- A. Mission statement
- B. Enforcement section
- C. Sanctions statement
- D. Accountability statement

C. A security policy doesn’t have a “sanctions statement.” Any related actions that can be pursued if the policy is not followed will be included in the “enforcement section.”

A, B, and D are incorrect. **A** is incorrect because a “mission statement” should be present and describe the overarching organizational vision. **B** is incorrect because an “enforcement section” would describe how a policy is going to be enforced and what steps might be taken in case personnel don’t adhere to it. **D** is incorrect because an “accountability statement” would describe the specific role types (i.e., security staff) that have a responsibility to follow the policy.

20. Which of these is not included in asset management?

- A. Hardware
- B. People
- C. Data
- D. Software

B. Although people might be referred to as organizational assets, they don't fall under asset management, which deals with how valuable company assets are managed.

A, C, and D are incorrect. **A** is incorrect because hardware should be managed from the moment that it first appears in a company up until the time that it is discarded. **C** is incorrect because "data" refers to how data is handled (i.e., categorized or labelled), as well as what media are used for storage (i.e., keeping a record of the serial numbers of all devices that store company data). **D** is incorrect because all software (including version and patch) should be properly listed for licensing and organizational security/awareness (i.e., if a new vulnerability is published about a software version the company is using, it should be able to search the asset register and identify what devices are using that software and require patching to be protected against that vulnerability).

21. What is the primary goal of change management?

- A. Review changes
- B. Minimize disruption
- C. Document changes
- D. Provide resource allocation

B. The primary goal of change management is to ensure any organizational disruption is kept to a minimum by carefully reviewing the proposed changes and their associated impact and approve or reject them accordingly.

A, C, and D are incorrect. **A** is incorrect because although change management will review a change thoroughly prior to approving/rejecting it, this is not the primary goal. The change is being reviewed with the goal of minimizing business impact. **C** is incorrect because documenting changes is part of change management process, but is not its primary goal. **D** is incorrect because resource allocation is not a part of change management. Change management is concerned about the change's impact, not about who will be performing it.

22. You are tasked with enhancing your company's physical security controls in order to ensure a proper audit trail exists. Which action would accomplish that goal?

- A. Place access card readers at the interior of secure operation rooms.
- B. Replace external 16-foot fences with taller ones.
- C. Increase security guard presence to achieve patrolling in pairs.
- D. Suggest hiring an external party to perform a physical security assessment.

A. Placing access card readers at the interior of secure areas would be a good starting point, as it will ensure that individuals who have swiped their cards to get access into the room will also do the same when exiting, creating a proper physical audit trail.

B, C, and D are incorrect. **B** is incorrect because replacing 16-foot fences with taller ones will not be an audit trail improvement. **C** is incorrect because although having more security guards is great, increasing presence to have them patrol in pairs doesn't improve the auditing capability. **D** is incorrect because hiring an external company to perform a physical assessment may lead to suggestions improving the organization's physical security stance but won't have any immediate impact on the existing audit trail process.

23. What process would take place if an (ISC)² member doesn't comply with the (ISC)² code of ethics?

- A. Any violators will be reported to the company the member is working for.
- B. (ISC)² will revoke their certifications.
- C. A review from a peer panel will take place.
- D. They will get a one-off warning.

C. In the event that a member violates the (ISC)² code of ethics, a peer review panel is formed so the case can be examined and a recommendation on how to proceed be made.

A, B, and D are incorrect. **A** is incorrect because (ISC)² doesn't have the remit to reach out to the company a member might be working for. **B** is incorrect because certifications aren't revoked straight away. That might be a recommendation from an appropriate panel, but a ruling will come out after it convenes. **D** is incorrect because a code of ethics violation is a very serious matter, which doesn't warrant any warnings.

24. Which of the following is not a step of the change management process?

- A. Change approved or rejected
- B. Change resolved

C. Change reviewed

D. Change submitted

B. The change management process doesn't include any step named "resolved," which is a term used to describe a ticket or incident status (in related ticketing/incident management tools). However, the term "change implemented" is used to describe a change request's implementation and associated documentation.

A, C, and D are incorrect. **A** is incorrect because at the "change approved or rejected" step, a decision will be made about whether this change will proceed to implementation or not. **C** is incorrect because at the "change reviewed" step, the change request can be thoroughly reviewed and discussed. **D** is incorrect because the "change submitted" phase is actually the first step of the change management process where an individual submits a specific change request regarding a change to be made on a device or system.

25. A system administrator is installing an IPS, a firewall, and a vulnerability scanner while at the same time drafting new security policies and training materials for his company. Which best describes his goal?

A. Confidentiality

B. Due diligence

C. Accountability

D. Defense in depth

D. "Defense in depth" is achieved by combining a variety of security devices and protocols to achieve a multilayered security approach. As such, installing an IPS, a firewall, and a vulnerability scanner and drafting new policies and training materials will aid in creating a multilayered security approach.

A, B, and C are incorrect. **A** is incorrect because confidentiality refers to information being accessible only by authorized individuals. **B** is incorrect because due diligence refers to a series of steps being undertaken by an organization or individual in order to minimize organizational risk. **C** is incorrect because accountability means being able to attribute an action to a specific individual.

26. In order for a merchant to comply with PCI-DSS, a network scan is required to be completed by which of the following?

A. SAQ

B. ASV

C. QSA

D. ISA

B. An ASV (Approved Scanning Vendor) is an external security entity that uses authorized scanners to perform a scan. Merchants of all levels are required to use such a provider for any scans that are performed.

A, C, and D are incorrect. **A** is incorrect because an SAQ (Self-Assessment Questionnaire) can be used by an entity with a low number of credit card transactions to perform a compliance audit. **C** is incorrect because a QSA (Quality Security Assessor) is an external party that can perform an audit at a client site. **D** is incorrect because an ISA (Internal Security Assessor) is an internal company employee (who is ISA certified) able to perform specific client-site audits.

27. Which of the following changes wouldn't require CAB approval?

A. Standard

B. Emergency

C. Normal

D. Expedited

A. Not all changes require CAB (Change Authorization Board) approval. If a change is well documented and has taken place in the past (as in the case of a standard change), then no CAB approval is typically required.

B, C, and D are incorrect. **B** is incorrect because an emergency change will need CAB approval to ensure it has minimal business impact. **C** is incorrect because normal changes are not repetitive changes (like standard ones), so they also need to be reviewed by CAB. **D** is incorrect because expedited changes are normal changes that need to be implemented quickly (usually within four to six hours) and as such also require CAB review.

28. An attacker manages to compromise a critical server, with no central logging capability enabled. After obtaining the information of interest, the audit logs are accessed and the related entries showing what actions the attacker performed are removed. Which of the following is being compromised?

A. Least privilege

B. Availability

C. Accountability

D. Authentication

C. In the event that an attacker manages to access and manipulate the audit logs, the log integrity would be compromised, making accountability very difficult to achieve.

A, B, and D are incorrect. **A** is incorrect because “least privilege” refers to users having enough access to perform their role but not more than that. **B** is incorrect because availability refers to appropriate individuals being able to access necessary resources whenever they need them. **D** is incorrect because authentication refers to verifying an individual’s identity, which doesn’t relate to log manipulation.

29. Employee badges are primarily used for:

A. Identification

B. Authentication

C. Provisioning

D. Entitlement

A. Employee badges are primarily used for identification, as they typically contain an individual’s full name along with a picture. Badges are also used to provide authentication, but the primary purpose is still identification.

B, C, and D are incorrect. **B** is incorrect because (as mentioned earlier) authentication may be provided in addition to a badge’s primary purpose, which is identification. **C** is incorrect because “provisioning” refers to user account creation and proper permissions being assigned to it. **D** is incorrect because “entitlement” refers to ensuring a user has appropriate privileges for the role that is being performed at any given moment in time (i.e., if a user changes role, privileges from the previous role should be revoked accordingly).

30. John powers on his computer and successfully installs a new application. After a few minutes he realizes that he’s unable to launch the newly installed application, as that is being blocked by the antivirus program. What would the antivirus most probably be using to block the application at that point?

A. Antivirus blacklisting

B. Scheduled scan

C. On-demand scan

D. Heuristic-based detection

D. When heuristic detection is used, some activities might be suspended if the antivirus application considers certain behaviors suspicious. In this example, the antivirus agent considered that the newly installed application was performing actions that were deviating from what is expected, and in order to protect the operating system, it blocked further usage of the application in question.

A, B, and C are incorrect. **A** is incorrect because since the application was successfully installed, it couldn't have been blacklisted by the antivirus agent. **B** is incorrect because a scheduled scan would be performed at regular intervals that the user specified (which is not something that the question mentions). **C** is incorrect because on-demand scanning (similarly to a scheduled scan) would be performed at the user's request (but the question didn't mention the user requested such a scan).

31. Which of the following is most suitable to protect against phishing attacks?

- A. Spam filtering provider
- B. User education
- C. E-mail gateway
- D. Antivirus

B. User education is the most effective way to protect a company from phishing attacks. No matter how many technical controls are put in place, human nature is always going to be a primary target for any attacker.

A, C, and D are incorrect. **A** is incorrect because a spam filtering provider will deal with some (or even the majority) of the phishing e-mails an organization may receive, but the ones that do get past it will end up in user mailboxes. **C** is incorrect because an e-mail gateway will filter out some of the phishing e-mails (it will mainly identify various e-mails as spam) but not necessarily all of them. **D** is incorrect because an antivirus agent might detect some types of content as malicious (i.e., attachment contained in a phishing e-mail), but that would mean the attack has already penetrated the company's security mechanisms and managed to reach the end user (since an antivirus program would typically be running at a user's machine or a company server).

32. Which of the following wouldn't be used during an SQL injection attack?

- A. SELECT
- B. --
- C. <>

D. ;

C. The “<>” character sequence would typically be used during a cross-site scripting (XSS) attack in order to insert HTML or JavaScript code.

A, B, and D are incorrect. **A** is incorrect because “SELECT” is an SQL statement that could be used during an SQL injection attack. **B** is incorrect because two hyphens “--” indicate a comment in SQL, which could be used in an SQL attack. **D** is incorrect because a semicolon “;” is used to indicate an SQL command termination, which could be used in an SQL attack

33. An attacker compromises a fashion retailer’s website and places malware on its web page. The goal is for the malware to be downloaded by clients that visit the website to place future orders. This is an example of:

A. Malvertising

B. Trapdoor

C. Drive-by download

D. Rootkit

C. During a “drive-by download” attack, a web server will typically be compromised and malware will be inserted so victims that later visit the target website download that malware.

A, B, and D are incorrect. **A** is incorrect because “malvertising” uses seemingly legitimate ads (which are in fact malicious) to spread malware. **B** is incorrect because a “trapdoor” is usually created by software developers in order to gain later access to systems (usually to troubleshoot some issue or as an emergency access method). **D** is incorrect because a “rootkit” obtains full target system access and is very difficult to detect and remove. Although an attacker might have actually planted a rootkit on the website, there’s no evidence to support that claim based on the question’s information. However, the information provided does show typical signs of a drive-by attack being staged.

34. Which of these statements about configuration management is not accurate?

A. Puppet uses Python.

B. Microsoft offers “Group Policy” for system configuration.

C. Imaging is a commonly used method.

D. Ansible can be used in a Linux environment.

A. Puppet doesn't use Python but actually uses a proprietary language or Ruby DSL (domain-specific language).

B, C, and D are incorrect. **B** is incorrect because "Group Policy" is a common feature used for configuration management in large Microsoft environments. **C** is incorrect because "imaging" is a typical method of configuration management whereby the IT team creates a system image best fitting the needs of the organization and that is then used for deployment across the estate. **D** is incorrect because Ansible can indeed be used for configuration management in a Linux environment (similar to Puppet and Chef).

35. In which of these control goal and class combinations does a motion sensor fall into?

A. Preventive, technical

B. Detective, technical

C. Preventive, physical

D. Detective, physical

D. A motion sensor is a detective control, as its primary goal is to trigger an alarm if it detects any environmental movement. As it is something that can be touched, it is classified as a physical control (similar to a smoke detector or a door alarm).

A, B, and C are incorrect. **A** is incorrect because a motion sensor can't prevent someone from entering the area it monitors and is not a technical control. **B** is incorrect because although a motion sensor is a detective control, it doesn't fall into the technical control class. **C** is incorrect because although a motion sensor is classified as a physical control, it can't prevent someone from entering the area of supervision.

36. Which of the following is not included in the (ISC)² code of ethics?

A. Act honorably.

B. Provide adequate service.

C. Advance and protect the profession.

D. Provide diligent service to principals.

B. The (ISC)² code of ethics doesn't have a canon about providing "adequate service." The third canon, however, does clearly state that a member has to "provide diligent and competent service to principals."

A, C, and D are incorrect. **A** is incorrect, as the second canon states “act honorably, honestly, justly, responsibly and legally.” **C** is incorrect, as the fourth canon states “advance and protect the profession.” **D** is incorrect, as the third canon states “provide diligent and competent service to principals.”

37. Which of the following can be part of more than one control type (i.e., combine being detective and preventive)?

A. Security policy

B. IDS

C. IPS

D. Audit logs

C. An IPS (intrusion prevention system) is predominantly used as a preventive control. However, in principle, it combines both detective and preventive control categories, as it can detect and act on malicious traffic. In addition, an IPS can be configured to act as an IDS (like when the IPS is in learning mode, prior to allowing anomaly detection to take place).

A, B, and D are incorrect. **A** is incorrect because a security policy is classified as a preventive control. **B** is incorrect as an IDS is a detective control and can’t take preventive or corrective actions. **D** is incorrect because audit logs are a detective control (logging an action that has taken place on a particular system).

38. A secure facility requires individuals to sign in at the reception desk and swipe their cards at the card readers. However, the security personnel do an end-of-shift check and realize that there’s a mismatch between the two records, showing 150 people signing in at the reception desk but only 140 swiping in. What is the most likely reason for that?

A. People signing in at reception more than once.

B. Tailgating has occurred.

C. Some employees signed in at reception but entered the building via nonstandard entrances.

D. Card reader system experienced an error.

B. Tailgating is the most likely reason for the mismatch. People will typically sign in at the reception desk, as they need to go past the security guards. However, when they need to swipe in, they will often tailgate others, especially during morning peak hours.

A, C, and D are incorrect. **A** is incorrect because having multiple people signing in more than once is highly unlikely, especially when they are supervised by a security guard. **C** is

incorrect because most buildings have all entrances monitored or protected via an access card or some other form of access control. Since the question mentions employees have to sign in and swipe in to gain access, the odds of such a company leaving other entrances unmonitored are very low. **D** is incorrect because if the card reader experienced errors, the security team would most likely be notified of that fact by some type of alert.

CHAPTER 3

Risk Identification, Monitoring, and Analysis

This chapter includes questions from the following objectives:

- 3.1 Understand the risk management process
- 3.2 Perform security assessment activities
- 3.3 Operate and maintain monitoring systems (e.g., continuous monitoring)
- 3.4 Analyze monitoring results

This chapter contains a complete set of questions and answers regarding risk definition (threat sources, threat events, vulnerabilities, and threat intelligence); risk management (risk register, CVSS, RMF, and risk treatment); and risk assessment (threat modelling, vulnerability analysis, and risk assessment steps). In addition, it refers to operation and maintenance of monitoring systems (IDS/IPS, detection and prevention methods, monitoring result analysis, SIEM, and continuous monitoring) and security tests and assessments (vulnerability assessments and penetration tests).

Q QUESTIONS

1. Your organization's IT team just identified a new high-severity vulnerability in your systems. After careful review, you decide that it isn't as severe as initially mentioned and can actually be classified as medium. This is an example of:

- A. Risk mitigation
- B. Risk avoidance
- C. Risk transference
- D. Risk recasting

2. Each laptop costs \$1,000 to your company, and two machines need to be replaced annually per employee due to loss or theft. Based on these figures, the SLE is:

- A. \$500
- B. \$2,000
- C. \$1,000

D. 2

3. Which of the following focuses on using impact and likelihood for risk categorization?

- A.** Qualitative analysis
- B.** Quantitative analysis
- C.** CVSS
- D.** RMF

4. A script kiddie uses an FTP exploit to attack a company's FTP server, which uses default administrator credentials. In this scenario, the exploit is classified as a:

- A.** Threat event
- B.** Threat source
- C.** Vulnerability
- D.** Risk

5. An administrator wants to install a system on an e-mail server with the goal of preventing attacks. Which of the following would be most suitable for that purpose?

- A.** HIDS
- B.** HIPS
- C.** NIDS
- D.** NIPS

6. You are working as head of security for a high-profile investment bank, which just created a new web banking application. Prior to its release, you would like to hire a company to do a penetration test of that application to review how secure it is. Which of these types of test would you choose?

- A.** White box
- B.** Black box
- C.** Gray box
- D.** Internal testing

7. What is the first step of any vulnerability assessment?

- A.** Discovery.
- B.** Remediation.
- C.** Identify and recommend methods to reduce the vulnerabilities.
- D.** Gain permission from management.

8. Which of these statements about sharing threat intelligence is inaccurate?

- A.** The best method is to share as much internal information as possible.
- B.** It's recommended to set rules about what information can be shared.
- C.** One often-used standard for threat intelligence sharing is STIX.
- D.** Identify appropriate threat intelligence information sources.

9. What is the major advantage of using a honeynet versus using a honeypot?

- A.** Allows security personnel to observe how attackers act
- B.** Uses multiple physical machines
- C.** Simulates a live network
- D.** Diverts an attack from the production network

10. Which of these statements regarding honeypots is accurate?

- A.** A honeypot includes two or more servers.
- B.** Security administrators can use honeypots to learn how attackers work.
- C.** A honeypot contains about 40 percent valid company data and 60 percent dummy data.
- D.** Honeypot administrators try to make honeypot access more appealing by deactivating as much security as possible so attackers can access honeypots more easily.

11. Which of these tools would be most suitable to detect an attempt to modify a registry key on a web server?

- A.** Stateful inspection firewall

B. Web proxy

C. NIDS

D. HIPS

12. An SQL injection attack is launched against a server, but the IDS placed to protect it fails to generate an alert. This is an example of a:

A. True positive

B. True negative

C. False positive

D. False negative

13. An attacker replaces a system .dll file in a Windows machine with another .dll file that seemingly performs the same function but also logs user keystrokes. Which of these tools is more appropriate to protect from this type of attack?

A. File integrity checker

B. SIEM

C. Application whitelisting

D. HIDS

14. A system administrator wants to be able to ensure that host access and audit log files are still available for inspection after an attacker gains access to the company's infrastructure and is able to delete or alter the related files. Which of these tools/methods is more appropriate for that purpose?

A. File integrity checker

B. Host weekly backup

C. SIEM tool

D. Password-protect the host log files

15. Which of these risk treatment approaches is not a valid option?

A. Risk ignorance

- B. Risk avoidance
- C. Risk acceptance
- D. Risk transference

16. Which of these statements about NIDS is not accurate?

- A. Can be used to detect port scans.
- B. Require a lot of CPU power to function properly.
- C. Port mirroring is required for them to operate.
- D. Work transparently within the network.

17. Which of the following is an example of an accidental threat source?

- A. HVAC failure
- B. Cable fire at a data center
- C. Secretary deleting one of the CEO's board meeting e-mails
- D. Server room power outage

18. Which of the following is an evaluation that tries to predict when an attack will take place?

- A. Quantitative analysis
- B. Qualitative analysis
- C. Risk register
- D. Threat modelling

19. A company's single e-commerce server has gone down. The estimation of loss for each day it's not online is \$500, and a technician charges \$400 per day to repair it, with an initial estimation of repair time being one day. The monthly likelihood of such incidents is estimated at three. What is the ALE?

- A. \$32,400
- B. \$64,800
- C. \$15,600

D. \$53,700

20. The following definition, “a policy violation resulting in an adverse event,” defines which of these terms?

- A.** Malware
- B.** Security incident
- C.** Threat
- D.** Vulnerability

21. Events that an organization often wants to monitor for are described by the term:

- A.** Security alerts
- B.** Security incidents
- C.** Events of interest
- D.** Intrusions

22. An IPS is placed in front of the DMZ segment housing an FTP and a web server. Suddenly, the network administrator receives a report of traffic from a specific external IP address not being received by the web server. However, a packet capture right before the IPS suggests that web traffic is passing properly. What is the most likely reason for this?

- A.** Malformed web request.
- B.** Packet capture has not been set to capture web traffic properly.
- C.** Physical link failure.
- D.** IPS dropping the traffic in question.

23. John is working extended hours these last two days in order to prepare a PowerPoint presentation for an upcoming board meeting. A security analyst just received an alert from the IDS system that John’s machine is exhibiting suspicious behavior. Which of the following is the most likely reason for this?

- A.** Blacklisted application was executed.
- B.** IDS malicious signature was triggered.
- C.** Malicious macro was executed.

D. IDS anomaly-based detection was triggered.

24. Which of the following is the most important consideration for a new SIEM implementation?

A. Configure time synchronization.

B. Adjust appropriate alert thresholds.

C. Ensure all available logs are ingested.

D. Provide security analyst access to individual log sources.

25. Which of the following is a feature that only a vulnerability scanning tool (like Nessus or Qualys) would have in comparison to a port scanning tool (like Nmap or Unicornscan)?

A. ICMP sweep.

B. Perform password auditing.

C. Scan a system for specific services.

D. OS fingerprinting.

26. The process of disabling nonrequired services on a host is called:

A. Border security

B. Vulnerability scanning

C. Host hardening

D. Host boosting

27. A vulnerability scanner indicated an unpatched version of Windows running on one of the network hosts. However, when the system administrator inspected the machine in question, he noticed that the host had been properly patched. What is the most probable reason for this?

A. Misconfigured scan profile.

B. The host owner patched the machine after the vulnerability scan took place.

C. False positive report finding.

D. Scan signature needs to be updated.

28. Gemma has been asked to create a new Wi-Fi subnet at a medium-size law firm. This will be used as a guest network for external consultants and clients to use when coming into the office to attend meetings. Which of these security mechanisms would be most suitable for this new network?

- A. WIDS
- B. NIDS
- C. NIPS
- D. Web proxy

29. Which of these vulnerabilities constitutes the greatest risk for a security-mature organization?

- A. Exploit using a macro
- B. Uneducated user
- C. Antivirus that hasn't been updated for a week
- D. Vulnerability relating to CentOS

30. You are the head of security and your company is running a core finance application on a legacy Windows 2003 server. A recent vulnerability scan flagged the machine in question as being out of support (extended support for Windows 2003 ended in 2015). After contacting the application vendor, they alert you to the fact that the application is quite old and supported only on a Windows 2003 operating system. Which of these options do you choose?

- A. Upgrade the server to the latest version.
- B. Migrate the application to a new server.
- C. Identify a new application that can be safely used.
- D. Accept the risk.

31. Which of the following is a great advantage of qualitative risk analysis?

- A. Uses impact rating
- B. Demonstrates risk cost
- C. Clearly shows what risks are more important

D. Demonstrates control cost

32. What is the easiest and most accurate way for an organization to check if it's PCI compliant?

A. Perform a penetration test.

B. Use a vulnerability scanner.

C. Use nmap.

D. Check with the compliance team.

33. Which of the following would a vulnerability scanner use to detect the operating system of a target system?

A. TCP window size

B. Closed ports

C. Reply (or lack of reply) to an ICMP request

D. TTL incrementation

34. How can an organization best check if vulnerabilities identified by a vulnerability scan have been remediated or not?

A. Perform a penetration test.

B. Check with the system administrators.

C. Manually verify if they are present.

D. Run another vulnerability scan.

35. A penetration tester is trying to identify what level of patching a target system currently has. During which phase would that activity take place?

A. Discovery

B. Reconnaissance

C. Fingerprinting

D. Exploitation

36. Which of these terms indicates an attack?

- A. Security incident
- B. Security alert
- C. Event of interest
- D. Threshold

37. James is a security manager who has been tasked with identifying attack trends in his organization. He is looking for an easy and quick way to provide some constructive feedback as to what those are. Which of these would be best suited for this purpose?

- A. IDS
- B. UTM
- C. Honeypot
- D. Proxy

38. Which of the following can be classified as a vulnerability?

- A. Insider
- B. Tailgating
- C. Lack of turnstiles
- D. CCTV monitoring

QUICK ANSWER KEY

1. D

2. C

3. A

4. A

5. B

6. B

7. D

8. A

9. C

10. B

11. D

12. D

13. A

14. C

15. A

16. B

17. C

18. D

19. A

20. B

21. C

22. D

23. D

24. A

25. B

26. C

27. C

28. A

29. B

30. D

31. C

32. B

33. A

34. D

35. C

36. A

37. B

38. C

ANSWERS A

1. Your organization's IT team just identified a new high-severity vulnerability in your systems. After careful review, you decide that it isn't as severe as initially mentioned and can actually be classified as medium. This is an example of:

A. Risk mitigation

B. Risk avoidance

C. Risk transference

D. Risk recasting

D. Risk recasting refers to a risk being reclassified to a different severity level (which might be higher or lower than the initial one).

A, B, and C are incorrect. **A** is incorrect because risk mitigation can only be achieved by implementing appropriate controls that aim to reduce it. **B** is incorrect because risk avoidance refers to avoiding the risk by not pursuing the related activity (i.e., if malware analysis poses a high risk for a security team, the team manager might choose not to allow the team to perform that activity, thus avoiding the related risk). **C** is incorrect because risk transference entails having the risk transferred to another party and is usually achieved by purchasing insurance, which will offer coverage to an organization in the event of a related risk materializing (i.e., a company can purchase ransomware insurance so it can have coverage in case a ransomware attack takes place).

2. Each laptop costs \$1,000 to your company, and two machines need to be replaced annually per employee due to loss or theft. Based on these figures, the SLE is:

- A. \$500
- B. \$2,000
- C. \$1,000
- D. 2

C. The SLE (Single Loss Expectancy) expresses how much a single occurrence of a risk would cost, meaning how much it would cost the business each time a risk materializes (calculated per annum). Since each laptop costs \$1,000, each time one is lost or stolen, the business has to pay \$1,000 to acquire a new one.

A, B, and D are incorrect. **A** is incorrect because \$500 is the result of dividing \$1,000 (which is the cost of obtaining a new machine) by 2 (which is the number of machines stolen or lost each year, also known as ARO, or Annual Rate of Occurrence). But that's incorrect, as the SLE is used to express the cost of a single occurrence and doesn't need to be divided by the number of occurrences. **B** is incorrect because multiplying the SLE (which is \$1,000) with the ARO (which is 2) yields the ALE (also known as Annual Loss Expectancy), which is used to express the annual cost for the business if the risk in question materializes. **D** is incorrect because 2 is used to express the ARO, as two machines are lost or stolen each year.

3. Which of the following focuses on using impact and likelihood for risk categorization?

- A. Qualitative analysis
- B. Quantitative analysis
- C. CVSS
- D. RMF

A. Qualitative analysis mostly uses the impact and likelihood of a risk and usually classifies risks in three to four categories, such as low, medium, high, and critical.

B, C, and D are incorrect. **B** is incorrect because quantitative analysis uses cost figures and calculations to perform risk analysis. It includes terms like SLE (Single Loss Expectancy), ARO (Annual Rate of Occurrence), and ALE (Annual Loss Expectancy) to make decisions on what controls to implement in order to minimize organizational risk. **C** is incorrect because CVSS (Common Vulnerability Scoring System) doesn't focus on impact and likelihood for risk categorization. It's a standard referring to computer system vulnerabilities that is often used by companies and focuses on a variety of metrics such as attack vector, confidentiality impact,

attack complexity, and various others. **D** is incorrect because RMF (Risk Management Frameworks) refers to NIST and ISO documents (usually 800-37 and 31000:2018), which use a multistep approach for risk management but don't particularly focus on likelihood and impact.

4. A script kiddie uses an FTP exploit to attack a company's FTP server, which uses default administrator credentials. In this scenario, the exploit is classified as a:

- A. Threat event
- B. Threat source
- C. Vulnerability
- D. Risk

A. A threat event is initiated by a threat source. In this scenario, the threat source (i.e., the script kiddie) is trying to leverage an exploit towards the FTP server, thus constituting a threat event.

B, C, and D are incorrect. **B** is incorrect because the script kiddie (i.e., the attacker) is a threat source, as he is the one initiating a threat event. **C** is incorrect because a vulnerability is defined as a system flaw or weakness that could result in a security breach if properly exploited. In this case, the default administrator login credentials that the FTP server is configured to use are the vulnerability, which the attacker is trying to exploit to gain unauthorized entry to the system. **D** is incorrect because risk is related to how probable it is that a vulnerability will be exploited by a threat.

5. An administrator wants to install a system on an e-mail server with the goal of preventing attacks. Which of the following would be most suitable for that purpose?

- A. HIDS
- B. HIPS
- C. NIDS
- D. NIPS

B. In order for the e-mail server to be able to protect itself against various attacks, a host-based system will be required to have the ability to block any type of offending traffic it identifies. A HIPS (host intrusion prevention system) is the best answer, as it can be installed on the e-mail server and has the ability to not only identify but also drop any malicious traffic.

A, C, and D are incorrect. **A** is incorrect because although a HIDS (host intrusion detection system) can be installed on an e-mail server, it only offers the ability to alert when an attack

takes place but is not able to block offending traffic. **C** is incorrect because a NIDS (network intrusion detection system) would be placed at a specific network location and monitor traffic that passes through that point, in addition to not being able to block any attack that it identifies. **D** is incorrect because a NIPS (network intrusion prevention system) can be placed at a specific network location and monitor (and block) traffic that passes through that point but wouldn't be installed on an endpoint.

6. You are working as head of security for a high-profile investment bank, which just created a new web banking application. Prior to its release, you would like to hire a company to do a penetration test of that application to review how secure it is. Which of these types of test would you choose?

- A.** White box
- B.** Black box
- C.** Gray box
- D.** Internal testing

B. Prior to the launch of this new product, it would be highly recommended for the bank to perform a black box penetration test, as the testers would have no knowledge of the system. That way, a real-world scenario of the application being released to the public and an attacker obtaining it and attempting to exploit it can be simulated.

A, C, and D are incorrect. **A** is incorrect because a white box test entails the tester having full knowledge of the system, which wouldn't be the case with a banking application being offered to the public. **C** is incorrect because a gray box test implies at least some basic knowledge about the application, which isn't something that an external attacker would easily get. **D** is incorrect because "internal testing" is another term used for white box testing, which implies full knowledge of the application and wouldn't be suitable in this case.

7. What is the first step of any vulnerability assessment?

- A.** Discovery.
- B.** Remediation.
- C.** Identify and recommend methods to reduce the vulnerabilities.
- D.** Gain permission from management.

D. Before any type of vulnerability assessment commences, the most important step is to ensure that management has provided permission to begin the activity. A common mistake is for someone to be hosting a service or platform on a cloud provider (like Amazon AWS) and think that because of that fact, permission is not required to perform a vulnerability scan. However, a

review of the terms and conditions will verify that proper permission for any type of such engagement is required prior to a test being performed.

A, B, and C are incorrect. **A** is incorrect because discovery is the step that follows after proper permission has been provided. During this phase, related scans are performed with the aim of acquiring information about the various vulnerabilities that might exist within the target network. **B** is incorrect because remediation would take place at the end, after the related vulnerabilities have been identified and recommendations have been accepted by the target company about how to mitigate them. **C** is incorrect because identification and recommendation of methods to reduce the vulnerabilities would take place after running and analyzing the results of the vulnerability scan but certainly not as a first step within the process.

8. Which of these statements about sharing threat intelligence is inaccurate?

- A.** The best method is to share as much internal information as possible.
- B.** It's recommended to set rules about what information can be shared.
- C.** One often-used standard for threat intelligence sharing is STIX.
- D.** Identify appropriate threat intelligence information sources.

A. Sharing all related information is hardly ever useful for any given organization. Usually, a company will identify what internal information (if any) it is willing to share so it can be reduced to a minimum. Sharing internal information should be done very carefully so as to not release any proprietary or confidential organizational data to the wider public.

B, C, and D are incorrect. **B** is incorrect because the parties with which the organization's information is allowed to be shared should be very specific so any unintentional data leakage can be avoided. **C** is incorrect because STIX (Structured Threat Information eXpression) is indeed used for threat intelligence sharing, as it allows for a common language to share related data. **D** is incorrect because identifying appropriate sources (i.e., data sharing tools, procedures, platforms, etc.) is key to efficient data sharing.

9. What is the major advantage of using a honeynet versus using a honeypot?

- A.** Allows security personnel to observe how attackers act
- B.** Uses multiple physical machines
- C.** Simulates a live network
- D.** Diverts an attack from the production network

C. The greatest advantage of using a honeynet versus a honeypot is that it offers more machines and gives the attacker the impression of a full production network with multiple devices and services running. Sometimes, experienced attackers can identify honeypots by the way they are configured or the services they are running. This becomes more difficult in the case of a honeynet, where a whole network of devices seems to exist.

A, B, and D are incorrect. **A** is incorrect because observing how an attacker acts is something that can be achieved with both a honeypot and a honeynet, so it's not a unique advantage of one or the other. **B** is incorrect because a honeynet doesn't need to use multiple physical machines. One physical machine can be used (as long as it has enough resources), which will have multiple virtual machines to simulate a variety of other devices within the honeynet. **D** is incorrect because diverting an attack from the production network is something that would happen with both a honeypot and a honeynet.

10. Which of these statements regarding honeypots is accurate?

A. A honeypot includes two or more servers.

B. Security administrators can use honeypots to learn how attackers work.

C. A honeypot contains about 40 percent valid company data and 60 percent dummy data.

D. Honeypot administrators try to make honeypot access more appealing by deactivating as much security as possible so attackers can access honeypots more easily.

B. One of the primary goals of a honeypot is to observe attacker TTPs (tactics, techniques, and procedures) so that the organizational defenses are then adjusted appropriately against future attacks.

A, C, and D are incorrect. **A** is incorrect because a honeypot doesn't include two or more servers. That is actually the definition of a honeynet (which is composed of two or more honeypots). **C** is incorrect because a honeypot only contains dummy data to lure attackers into attacking it but doesn't have any real organizational data that would be in danger of getting stolen. **D** is incorrect because although honeypots need to be accessible by attackers, they also need to be appealing to them. In order for that to happen, they need to simulate a real system as much as possible. As such, some services need to be deactivated (mimicking a more hardened system), and access to a honeypot should be fairly straightforward but without being too direct, which might give away its true nature to a skilled attacker.

11. Which of these tools would be most suitable to detect an attempt to modify a registry key on a web server?

A. Stateful inspection firewall

B. Web proxy

C. NIDS

D. HIPS

D. A HIPS (host intrusion prevention system) would be the best tool to detect a registry key modification at the server level. A HIPS is installed on the host and aims to monitor for suspicious activity. Once something malicious is detected, it can take further steps to mitigate that activity.

A, B, and C are incorrect. **A** is incorrect because a stateful inspection firewall's goal is to filter inbound (and outbound) communications based on access control lists and a stateful inspection table. However, it won't be able to detect a registry key modification taking place at the server, as it is unable to monitor for that type of change. **B** is incorrect because a web proxy aims to filter web traffic and won't have any visibility of a change taking place on a particular host. **C** is incorrect because a NIDS (network intrusion detection system) is installed at a specific network location and monitors traffic that passes through that point, which wouldn't allow it to detect host registry key modifications.

12. An SQL injection attack is launched against a server, but the IDS placed to protect it fails to generate an alert. This is an example of a:

A. True positive

B. True negative

C. False positive

D. False negative

D. A "false negative" occurs when the IDS fails to generate an alert although a valid attack is taking place (meaning it is unable to detect it). In this case, the attacker launched an SQL injection attack but the IDS didn't generate any type of alert about that activity.

A, B, and C are incorrect. **A** is incorrect because a "true positive" occurs when the IDS correctly identifies an attack that is occurring and raises an alert to notify the security team or administrator of that fact (in the example of the SQL injection attack, the IDS should have raised an appropriate alert to notify of an attack taking place). **B** is incorrect because a "true negative" event means that there was no suspicious traffic worthy of raising an alert, and the IDS correctly didn't raise any alert. **C** is incorrect because a "false positive" event means that the IDS raises an alert without an attack being present (i.e., if a user sends the server some other type of request that the IDS considers an SQL injection attack and raises an alert for that traffic).

13. An attacker replaces a system .dll file in a Windows machine with another .dll file that seemingly performs the same function but also logs user keystrokes. Which of these tools is more appropriate to protect from this type of attack?

- A. File integrity checker
- B. SIEM
- C. Application whitelisting
- D. HIDS

A. File integrity checking tools create a list of file hashes, which they safely store on the local or remote system. Their goal is to generate hashes of the files they are tasked with protecting and then compare those with the values that exist on the file that was stored earlier. If the values match, it means no tampering took place.

B, C, and D are incorrect. **B** is incorrect because a SIEM (security information and event management) platform works by ingesting various log files and performing correlations across data but won't be able to check for any specific file manipulation. **C** is incorrect because application whitelisting aims to identify if a program or process that is executed on a system at any given moment in time matches a prepopulated list of applications and processes. If it's not in the list, then it won't be allowed. **D** is incorrect because although IDS systems sometimes offer file integrity checking, they usually do that by employing some file integrity checking tool, while the primary purpose of an IDS tool is to monitor the network or host (depending on if it's a NIDS or HIDS) for any suspicious traffic and alert upon it once it is identified.

14. A system administrator wants to be able to ensure that host access and audit log files are still available for inspection after an attacker gains access to the company's infrastructure and is able to delete or alter the related files. Which of these tools/methods is more appropriate for that purpose?

- A. File integrity checker
- B. Host weekly backup
- C. SIEM tool
- D. Password-protect the host log files

C. SIEM (security information and event management) tools store log data centrally (similar to any centrally logging platform), which means that once a copy of the logs is forwarded to a SIEM tool, that data is locally stored on the SIEM and remains there based on the retention period that has been configured. Typical log retention periods can be from three to five years (for highly critical or sensitive environments, like banking and military systems).

A, B, and D are incorrect. **A** is incorrect because a file integrity checker will alert when a file is modified but won't be able to provide an up-to-date copy of the log file so the administrator can review its content. **B** is incorrect because a system backup is a good way of obtaining a copy

of the host log files; however, as that takes place on a given period (i.e., weekly), it's not enough to ensure that up-to-date information exists, as it might be even a week old in the case where an attacker gains access close to when a backup was taken. **D** is incorrect, as password protecting the host log files will ensure that a file won't be tampered with (assuming the attacker is not able to brute force the file password) but doesn't offer any protection against file deletion, so the attacker can still delete the file altogether.

15. Which of these risk treatment approaches is not a valid option?

- A. Risk ignorance
- B. Risk avoidance
- C. Risk acceptance
- D. Risk transference

A. Sometimes companies tend to ignore risks rather than following various valid risk treatment approaches. However, risk ignorance is not a valid risk strategy. If a company doesn't wish to mitigate or otherwise handle an identified risk, it would be prudent to accept that risk.

B, C, and D are incorrect. **B** is incorrect because risk avoidance is an option that can be followed in the event that an activity is not performed. That is, if there's a high risk associated with selling laptops in an area with high levels of theft, not opening a shop in that area will avoid the associated risk. **C** is incorrect because sometimes accepting the risk is the best option, especially in cases where the cost associated with risk reduction is greater than the asset value. For example, assume that a home office includes two laptops with a total cost of \$4,000. The best way to protect it is to install a state-of-the-art security system and physical guards that would cost \$8,000. In that case, the owner might decide to accept the risk, since applying appropriate protection would cost double the amount of the assets (without, of course, considering the value of the data of those two machines). **D** is incorrect, as transferring the risk (usually to an insurance company) is a commonly accepted risk treatment approach. Buying cyber insurance is very common (especially for medium to large companies) in order to be protected in the event of a cyber attack.

16. Which of these statements about NIDS is not accurate?

- A. Can be used to detect port scans.
- B. Require a lot of CPU power to function properly.
- C. Port mirroring is required for them to operate.
- D. Work transparently within the network.

B. NIDS don't require substantial CPU power to operate, as that is a characteristic requirement of HIDS. Because a HIDS is installed on a given host, it requires significant amounts of processing power to be able to monitor for all host activities and generate appropriate alerts.

A, C, and D are incorrect. **A** is incorrect because a NIDS can be used to detect port scans. (Traffic generated by different types of port scans is fairly distinctive, so most NIDS tools have predefined signatures that will identify it. However, attackers attempt to use methods to go undetected, like trying to fragment packets or perform slow scans to evade IDS systems). **C** is incorrect, as port mirroring is required for a NIDS to be able to monitor traffic. With port mirroring enabled, the NIDS will get a copy of the network traffic and be able to monitor it for suspicious patterns, which will cause it to generate appropriate alerts. **D** is incorrect because NIDS are placed at specific network segments and monitor traffic passing through them. That function is transparent to an attacker that manages to gain access to a file server, as the NIDS is not installed on the server and thus can't be switched off or manipulated by the attacker.

17. Which of the following is an example of an accidental threat source?

- A.** HVAC failure
- B.** Cable fire at a data center
- C.** Secretary deleting one of the CEO's board meeting e-mails
- D.** Server room power outage

C. A secretary deleting one of the CEO's e-mails relating to a board meeting is an example of an accidental threat source. In this case, the secretary was able to accidentally perform this action due to the level of access she has to her manager's mailbox.

A, B, and D are incorrect. **A** is incorrect because an HVAC (heating, ventilation, and air conditioning) failure would be classified as a structural threat source, similar to any IT, software, or other environmental control failure. **B** is incorrect because a cable fire at a data center would be classified as an environmental threat source, like a hurricane or a tornado. **D** is incorrect because a server room power outage would be classified as an environmental threat source (external service failures, such as power outages or satellite disk failures, can be classified as environmental threat sources).

18. Which of the following is an evaluation that tries to predict when an attack will take place?

- A.** Quantitative analysis
- B.** Qualitative analysis
- C.** Risk register

D. Threat modelling

D. Threat modelling tries to estimate when an attack will take place by using various methods with which it models attack parameters about any given system.

A, B, and C are incorrect. **A** is incorrect because quantitative analysis uses cost figures and calculations to perform risk analysis. **B** is incorrect because qualitative analysis mostly uses the impact and likelihood of a risk and usually classifies risks in three to four categories. **C** is incorrect because a risk register is used to record all organizational risks and related impact, as well as the progress of placing controls to mitigate those.

19. A company's single e-commerce server has gone down. The estimation of loss for each day it's not online is \$500, and a technician charges \$400 per day to repair it, with an initial estimation of repair time being one day. The monthly likelihood of such incidents is estimated at three. What is the ALE?

A. \$32,400

B. \$64,800

C. \$15,600

D. \$53,700

A. In order to calculate the ALE, the following formula will be required: $ALE = SLE \times ARO$. Furthermore, the SLE can be found by adding all related server costs, which are daily loss and repair cost. Based on that, the SLE will be $\$500 + (\$400 \times 1 \text{ day}) = \$500 + \$400 = \900 . The ARO is $3 \text{ monthly occurrences} \times 12 \text{ months} = 36 \text{ occurrences per year}$. $ALE = SLE \times ARO = \$900 \times 36 = \$32,400$.

B, C, and D are incorrect. All those figures are distractors, as none of them is correct.

20. The following definition, "a policy violation resulting in an adverse event," defines which of these terms?

A. Malware

B. Security incident

C. Threat

D. Vulnerability

B. A security incident can be defined as "a policy violation resulting in an adverse event" (i.e., privilege escalation or data exfiltration).

A, C, and D are incorrect. **A** is incorrect because the term “malware” is used to describe malicious software and is one example of a security incident. However, there are various examples of security incidents that don’t have any relation to malware, like tailgating, identity theft, and many more. **C** is incorrect because a threat is any action that can pose a danger to an organizational asset. **D** is incorrect because a vulnerability is a weakness in a system (which is something that attackers often try to exploit to gain unauthorized access).

21. Events that an organization often wants to monitor for are described by the term:

- A.** Security alerts
- B.** Security incidents
- C.** Events of interest
- D.** Intrusions

C. The term “events of interest” is used to describe any event of particular interest (usually logged in a database or registered in some type of application/system log) that an organization wants to keep track of. It is worth mentioning that it doesn’t necessarily refer to a suspicious action, as sometimes a system administrator might be interested in receiving confirmations of activities that have been successfully completed (i.e., antivirus updates or a new Windows policy push).

A, B, and D are incorrect. **A** is incorrect because the term “security alerts” usually refers to some type of activity that warrants the attention of a security analyst or system administrator and is usually produced when a related rule is triggered within some type of security tool (usually an IDS, firewall, router, SIEM, or similar device). **B** is incorrect because a security incident can be defined as a policy violation resulting in an adverse event (i.e., privilege escalation or data exfiltration). **D** is incorrect because the term “intrusion” refers to some type of unwarranted access into a system, database, or tool.

22. An IPS is placed in front of the DMZ segment housing an FTP and a web server. Suddenly, the network administrator receives a report of traffic from a specific external IP address not being received by the web server. However, a packet capture right before the IPS suggests that web traffic is passing properly. What is the most likely reason for this?

- A.** Malformed web request.
- B.** Packet capture has not been set to capture web traffic properly.
- C.** Physical link failure.
- D.** IPS dropping the traffic in question.

D. Since the packet capture shows web traffic being received by the web server without traffic from a specific source passing through the IPS, it seems that the traffic in question is being identified as malicious and thus is blocked from traversing the rest of the network.

A, B, and C are incorrect. **A** is incorrect because although a malformed web request could have been sent, there's not enough information to suggest that this is the case and is the specific reason for dropping traffic from the IP in question. **B** is incorrect because there's no actual indication that web traffic is not being captured properly. In addition, the fact that only traffic from a specific IP is not seen by the web server suggests that packet capturing has been properly configured but some part of the traffic is just not passing through the IPS. **C** is incorrect because a physical link failure would result in no traffic reaching any of the servers in the DMZ, which is not the case here.

23. John is working extended hours these last two days in order to prepare a PowerPoint presentation for an upcoming board meeting. A security analyst just received an alert from the IDS system that John's machine is exhibiting suspicious behavior. Which of the following is the most likely reason for this?

A. Blacklisted application was executed.

B. IDS malicious signature was triggered.

C. Malicious macro was executed.

D. IDS anomaly-based detection was triggered.

D. Other than the fact that John started working extended hours, there doesn't seem to be any other deviation in his behavior nor any suspicious activity. However, it is likely that the IDS (using anomaly-based detection) has considered the nonstandard hours a malicious deviation for his machine, thus triggering the related alert.

A, B, and C are incorrect. **A** is incorrect because John is using a presentation program (PowerPoint), which in all likelihood he has used in the past (the question doesn't state this was the first use and doesn't mention anything that could be considered an illegitimate application execution); hence, this doesn't suggest that a blacklisted application was executed. **B** is incorrect because there's no indication of a malicious IDS signature being triggered or an activity that could suggest malicious activity taking place. **C** is incorrect because there's no indication that any macro was used.

24. Which of the following is the most important consideration for a new SIEM implementation?

A. Configure time synchronization.

B. Adjust appropriate alert thresholds.

C. Ensure all available logs are ingested.

D. Provide security analyst access to individual log sources.

A. In order for proper analysis to take place, a SIEM needs accurate data synchronization. If an analyst tries to track an attacker's actions through various device logs that are in different time zones, the analysis can become unmanageable. In addition, it becomes very difficult for cooperating third parties to track any reports that are produced when data from different time zones is referenced in various parts of the report.

B, C, and D are incorrect. **B** is incorrect because although adjusting the alert thresholds is important to ensure that related teams are only engaged in the event of a valid security incident, this is something that can happen during the initial stages of the SIEM operation and usually happens gradually as more log sources are added and analyst teams become more familiar with its operation. It's difficult to ensure that thresholds are properly adjusted right at the beginning, as this takes substantial effort and time. **C** is incorrect because adding all available log sources to the SIEM (depending on their amount and volume) would result in overwhelming the tool as well as the analyst team supporting it. Organizations usually start by on-boarding the basic log sources that are required to achieve their goals and gradually adding more devices as they grow, while at the same time performing fine-tuning. **D** is incorrect because a SIEM works as a central logging and correlation platform. It allows analysts to view all the basic information to perform their tasks. Although having access to individual log sources can sometimes be beneficial (i.e., the analyst can review the raw logs that haven't been normalized within the SIEM, meaning all initial information is present), the whole point of the SIEM is to offer log storage and correlation without needing to access the end devices.

25. Which of the following is a feature that only a vulnerability scanning tool (like Nessus or Qualys) would have in comparison to a port scanning tool (like Nmap or Unicornscan)?

A. ICMP sweep.

B. Perform password auditing.

C. Scan a system for specific services.

D. OS fingerprinting.

B. A password audit can be performed by vulnerability scanning tools, as those have the ability to identify accounts that use default or empty passwords or a password that matches the username of the account and other similar conditions that indicate a weak or noncompliant password being used.

A, C, and D are incorrect. **A** is incorrect because both a port scanner and a vulnerability analysis tool can perform an ICMP sweep to check what hosts are active within a given network. **C** is incorrect because a service scan can be performed by both categories of tools. It is common

for attackers to use nmap to perform reconnaissance and identify specific services that might be running on a target host with the intention of finding an exploit that can be used against those. **D** is incorrect because OS fingerprinting can be performed by both vulnerability and port scanners. Attackers often use open port scanners to identify any old or unpatched OS versions running on a target host, while security analysts use vulnerability scanning tools to identify the same hosts and report them to the IT team for patching or upgrading.

26. The process of disabling nonrequired services on a host is called:

- A. Border security
- B. Vulnerability scanning
- C. Host hardening
- D. Host boosting

C. The term “host hardening” is used to describe the process during which a host is configured to only run the required set of services that support its role (i.e., a web server doesn’t necessarily need to be running an FTP service). Nonessential services should be disabled so they can’t be exploited by an attacker, thus reducing the host’s attack surface.

A, B, and D are incorrect. **A** is incorrect because “border security” entails reviewing the various routes to the network from any external source. Any device that is placed in the network path should be checked (i.e., switches, routers, firewalls, IDS/IPS, etc.). **B** is incorrect because vulnerability scanning entails using a tool like Qualys or Nessus to scan a network or a specific host for any vulnerabilities and flag those to the system administrators so they can be appropriately dealt with. **D** is incorrect because the term “host boosting” is a distractor and doesn’t have any real meaning.

27. A vulnerability scanner indicated an unpatched version of Windows running on one of the network hosts. However, when the system administrator inspected the machine in question, he noticed that the host had been properly patched. What is the most probable reason for this?

- A. Misconfigured scan profile.
- B. The host owner patched the machine after the vulnerability scan took place.
- C. False positive report finding.
- D. Scan signature needs to be updated.

C. Vulnerability scanners may sometimes produce inaccurate results, which is the most probable reason why lack of a certain patch or system update would be reported by the scanner without it actually being absent in reality.

A, B, and D are incorrect. **A** is incorrect because there's no indication that the scan profile was somehow misconfigured. **B** is incorrect because the system administrator would be responsible for deploying new patches within the environment, and he would know what the patching schedule is. Hence, it is improbable that such an activity took place after the initial scan without him realizing. **D** is incorrect because although a problem with the signature might be causing an issue, there's not enough information in the question to deduce that. Furthermore, the question does state "What is the most probable reason," which again suggests a false positive is the correct answer.

28. Gemma has been asked to create a new Wi-Fi subnet at a medium-size law firm. This will be used as a guest network for external consultants and clients to use when coming into the office to attend meetings. Which of these security mechanisms would be most suitable for this new network?

A. WIDS

B. NIDS

C. NIPS

D. Web proxy

A. A WIDS (wireless intrusion detection system) would be most suitable to monitor traffic to this new subnet. Since a new wireless segment is being implemented, installing a WIDS will allow Gemma to monitor the wireless network for suspicious activity that she can be alerted for.

B, C, and D are incorrect. **B** is incorrect because even though a NIDS can also be used to monitor traffic to this new subnet, since it's a wireless one, the best option is to use a wireless IDS (aimed specifically at Wi-Fi), not a traditional NIDS. **C** is incorrect because although a NIPS could also be employed for the purposes of monitoring and acting upon malicious traffic that is present with the network in question, a WIDS is still the best option (aiming at the wireless aspect). In addition, since this is a guest network and not a critical internal network, there's a small probability that automatically blocking traffic will be required (which is what an IPS can do). Gemma can review any suspicious activity and take appropriate action, thus leaving the option of a WIDS as the best one. **D** is incorrect because a web proxy is a great solution for filtering web traffic but is unable to monitor any other type of traffic that traverses the network, which is why using an IDS is highly recommended.

29. Which of these vulnerabilities constitutes the greatest risk for a security-mature organization?

A. Exploit using a macro

B. Uneducated user

C. Antivirus that hasn't been updated for a week

D. Vulnerability relating to CentOS

B. An uneducated user is usually the greatest risk, as an attacker can use social engineering tactics to gain valuable information just by asking for it over a phone or an SMS message. No matter how efficient and elaborate the company's security architecture might be, all of that can be easily defeated by a user who provides a username and a password for remotely accessing a system to someone posing as a valid IT employee.

A, C, and D are incorrect. **A** is incorrect because a macro exploit can have severe consequences, but in order for it to run properly, macro execution needs to be allowed and on top of that, any endpoint monitoring tool needs to allow that action. In a security-mature organization, macro execution would usually be disabled, and often some type of endpoint protection tool would be enabled to monitor for malicious macro execution. **C** is incorrect because an antivirus that has not been updated does pose a risk, but judging by the fact that only a week has gone by with no update, an uneducated user still poses a greater risk. **D** is incorrect because a vulnerability relating to a specific OS (no matter how severe it might be) can only have serious side effects if that OS is being actively used in the organization. If the company doesn't have any CentOS machines (something that the question doesn't mention), then this is not a concern.

30. You are the head of security and your company is running a core finance application on a legacy Windows 2003 server. A recent vulnerability scan flagged the machine in question as being out of support (extended support for Windows 2003 ended in 2015). After contacting the application vendor, they alert you to the fact that the application is quite old and supported only on a Windows 2003 operating system. Which of these options do you choose?

- A. Upgrade the server to the latest version.
- B. Migrate the application to a new server.
- C. Identify a new application that can be safely used.
- D. Accept the risk.

D. Accepting the risk may sometimes be the only way to move forward (that's why risk acceptance is a valid risk strategy, although it is not recommended to follow it unless all other avenues have been exhausted). As this is a core finance application, meaning its operation is crucial to the business, a way to continue using it must be identified. If upgrading the server (due to being out of support) or reverting to the application's vendor for a solution (i.e., if the application is only supported on a Windows 2003 system) are both nonviable options, accepting the risk seems the only viable solution.

A, B, and C are incorrect. **A** is incorrect because upgrading the server to the latest version (which is several versions after the current one that the server is running) might disrupt the application's normal operation or even disrupt the functionality of the legacy application (only

supported in Windows 2003). **B** is incorrect because migrating the application to a new server (judging by the fact that the application is old and not supported anymore) might again result in not being able to use it further. **C** is incorrect because although identifying a new application (given the age of the current one being used) might take some time, it seems like a viable way moving forward. However, this will possibly take several weeks or months, as the new application will need to be identified, then reviewed from the finance team to ensure that it fulfills all related needs, and then tested and accepted. However, that won't solve the challenge at hand (i.e., the machine being flagged as out of support and a core application being run on it), so this can only be a viable long-term approach.

31. Which of the following is a great advantage of qualitative risk analysis?

- A.** Uses impact rating
- B.** Demonstrates risk cost
- C.** Clearly shows what risks are more important
- D.** Demonstrates control cost

C. The greatest advantage of using qualitative analysis is that it can provide a clear depiction of the most important organizational risks that need immediate attention.

A, B, and D are incorrect. **A** is incorrect because although qualitative analysis uses impact ratings, that doesn't constitute its greatest advantage, but rather is a part of how it operates. The benefit of using impact rating (as well as likelihood rating) is being able to prioritize organizational risks appropriately. **B** is incorrect because the risk of a cost is depicted using quantitative risk analysis. **D** is incorrect because control cost (similar to risk cost) is also depicted using quantitative risk analysis.

32. What is the easiest and most accurate way for an organization to check if it's PCI compliant?

- A.** Perform a penetration test.
- B.** Use a vulnerability scanner.
- C.** Use nmap.
- D.** Check with the compliance team.

B. A vulnerability scanner offers the option of configuring an appropriate profile, allowing you to scan for specific types of compliance (like PCI-DSS, HIPAA, etc.). Depending on the results, a report usually can be generated showing a compliance percentage (e.g., 60 percent), thus tracking how compliant the organization is and what steps are required to reach full compliance.

A, C, and D are incorrect. **A** is incorrect because a penetration test would be more suitable to check if vulnerabilities can be exploited by attackers. **C** is incorrect because nmap would commonly be used to check for active hosts or open ports within a given network or specific host. **D** is incorrect because checking with the compliance team might not produce accurate results. Even though someone might think that the company is compliant, some network, hardware, or any other type of change might have created an issue that needs to be addressed before the organization is fully compliant.

33. Which of the following would a vulnerability scanner use to detect the operating system of a target system?

- A.** TCP window size
- B.** Closed ports
- C.** Reply (or lack of reply) to an ICMP request
- D.** TTL incrementation

A. Most OS fingerprinting tools use the TCP window size to gain information about the OS a remote system is running. Depending on the remote system response, a result can be reached as to what OS it's running.

B, C, and D are incorrect. **B** is incorrect because closed ports don't help identify an OS, as some might be closed because they are not used. **C** is incorrect because an ICMP request might not be replied to due to the fact that a firewall is filtering that type of traffic and doesn't allow for a reply to be sent. **D** is incorrect because TTL is decremented as packets pass through various systems (also note that tools commonly use the initial TTL to perform OS fingerprinting, as that varies in different systems).

34. How can an organization best check if vulnerabilities identified by a vulnerability scan have been remediated or not?

- A.** Perform a penetration test.
- B.** Check with the system administrators.
- C.** Manually verify if they are present.
- D.** Run another vulnerability scan.

D. Running a new vulnerability scan is the easiest and most accurate way to check if previously identified vulnerabilities are still present within a given system. The same scanning profile as before can be used in order to verify which of the vulnerabilities are still present.

A, B, and C are incorrect. **A** is incorrect because a penetration test should only be performed after any vulnerabilities have been successfully remediated, as the penetration testers would exploit any existing ones. **B** is incorrect because checking with a system administrator doesn't guarantee accurate results (the individual might not be in a position to properly verify the vulnerability's existence or sometimes might feel like being questioned about the level of his work). **C** is incorrect because manual verification might not be possible in large environments where a vulnerability scan might show 1,000 vulnerabilities in need of remediation.

35. A penetration tester is trying to identify what level of patching a target system currently has. During which phase would that activity take place?

- A. Discovery
- B. Reconnaissance
- C. Fingerprinting
- D. Exploitation

C. Trying to identify patching levels, open ports, and OS versions would all take place during the "fingerprinting" phase.

A, B, and D are incorrect. **A** is incorrect because "discovery" entails identifying what hosts are active within the target network, as well as obtaining as much information about the target as possible that can be later used to exploit its systems. **B** is incorrect because "reconnaissance" is just a different term for "discovery." **D** is incorrect because during "exploitation," the penetration tester would try to exploit any of the vulnerabilities that were identified in previous steps.

36. Which of these terms indicates an attack?

- A. Security incident
- B. Security alert
- C. Event of interest
- D. Threshold

A. The term "security incident" can be defined as "a policy violation resulting in an adverse event." Usually, some type of monitoring will be in place to notify a team or a security administrator of an event of interest or an alert. After further investigation, that may prove to be nothing substantial or may be an actual security incident (i.e., privilege escalation or data exfiltration) that requires more in-depth investigation and possibly escalation to other organizational teams or third parties.

B, C, and D are incorrect. **B** is incorrect because a “security alert” isn’t an indication of an actual attack. However, it does show that a rule has been triggered in a tool or system and has generated an alert for someone to investigate. For example, an administrator entering a wrong remote access password five consecutive times might generate an alert but that doesn’t mean that an actual brute-force attack is taking place. **C** is incorrect because an “event of interest” describes an event (usually logged in a database or registered in some type of application/system log) that an organization wants to keep track of. **D** is incorrect because a “threshold” is a limit that, when surpassed, an alert is generated to draw attention to the activity in question. However, more investigation is required to decide if an actual attack is taking place.

37. James is a security manager who has been tasked with identifying attack trends in his organization. He is looking for an easy and quick way to provide some constructive feedback as to what those are. Which of these would be best suited for this purpose?

- A. IDS
- B. UTM
- C. Honeypot
- D. Proxy

B. A UTM (Unified Threat Management) device combines a variety of functions, which may include IDS/IPS, spam filtering, proxy, firewall, and antivirus. Performing trend analysis based on data from such a device would provide James with a wealth of information that can be filtered to show attack patterns, failed attack attempts versus successful ones, attacker IPs of interest, targeted services, and a lot more.

A, C, and D are incorrect. **A** is incorrect because IDS logs (although quite valuable) will only show a part of the picture (and that would heavily depend on how logging has been configured and what type of IDS is being used as well as what signatures it is currently employing). **C** is incorrect because a honeypot is great for gathering data for trend analysis and allows security personnel to monitor attacker activities and obtain a lot of detail on how they operate. However, efficiently setting up a honeypot may take a substantial amount of time. (In large organizations, this will require a request to be submitted and may take a lot of time to be reviewed—a separate subnet or even a whole network might need to be set up to maintain appropriate segregation from the corporate network, and several tests might need to be performed before this actually goes live, so this approach can’t be classified as quick and easy.) **D** is incorrect because a proxy would only provide insight to the specific type of data that is being proxied (i.e., if it’s a web proxy, then trend analysis around web browsing patterns can be provided).

38. Which of the following can be classified as a vulnerability?

- A. Insider

B. Tailgating

C. Lack of turnstiles

D. CCTV monitoring

C. A vulnerability is defined as “a system flaw or weakness that could result in a security breach if properly exploited.” In this case, the lack of turnstiles can be exploited by an attacker, who can perform tailgating due to this fact.

A, B, and D are incorrect. **A** is incorrect because an insider would be classified as the threat source, which is someone who can exploit a vulnerability. **B** is incorrect because tailgating is the threat event that a threat source (i.e., insider or outsider) can initiate. **D** is incorrect because the CCTV monitoring option is a distractor, as it doesn't pose a vulnerability, but rather enhances physical security at any given situation.

CHAPTER 4

Incident Response and Recovery

This chapter includes questions from the following objectives:

- 4.1 Support incident lifecycle
- 4.2 Understand and support forensic investigations
- 4.3 Understand and support Business Continuity Plan (BCP) and Disaster Recovery Plan (DRP) activities

This chapter contains a complete set of questions and answers regarding support of an incident's lifecycle (preparation, detection, analysis and escalation, containment, eradication, recovery, lessons learned/implementation of new countermeasure), understanding fault tolerance (RAID, failover clusters, load balancing, redundant connections), and backups. It also contains questions and answers about understanding BCP and DRP activities (BIA, DRP, BCP, restoration planning, testing plans, and establishing alternative locations), as well as computer forensics (evidence handling and computer forensics investigations).

Q QUESTIONS

1. In which of these phases of the incident response lifecycle would the identification of an appropriate tool for memory forensics fall into?

- A. Containment
- B. Detection, Analysis and Escalation
- C. Preparation
- D. Recovery

2. Which of the following provides no fault tolerance?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 6

3. Which of these statements regarding a "hot site" is inaccurate?

- A. It is commonly usable within a few hours.
- B. It includes network infrastructure.
- C. Cost is quite significant.
- D. Contains up-to-date organizational data.

4. If a company server is unavailable for more than 45 minutes, core business functions are affected. That means that 45 minutes is the:

- A. RPO
- B. BIA
- C. DRP
- D. MAO

5. A security analyst identified outbound command and control traffic towards a malicious IP address. Which of these steps would most likely follow?

- A. Investigate why there's traffic towards the malicious IP address.
- B. Block the IP at the firewall.
- C. Take a forensic image of the affected machine's memory.
- D. Reimage the asset using a previous backup.

6. A forensic investigator needs to capture data for later analysis. Which one of these types of data should be captured first?

- A. Remote logs
- B. Local drive
- C. Swap file
- D. Backup tape

7. Which of the following does forensic software use to ensure no data tampering has taken place?

- A. Serial numbers

- B.** Hashing
- C.** Digital signatures
- D.** Encryption

8. A company's client database is stored on a server so that online purchases can take place. The database needs to be accessible constantly even if the server goes down. Which of the following would be more suitable for that purpose?

- A.** Failover cluster
- B.** Load balancing
- C.** Full backup
- D.** Differential backup

9. A company is hosting an SQL, cloud, and web server at its own premises. Which of these backup locations would you recommend as the safest option for backup storage and later retrieval of the web server's data?

- A.** A backup server placed in the server room, along with the web server
- B.** A magnetic tape stored on the premises but in a different building
- C.** An encrypted copy of the data backed up on the company cloud server
- D.** Backup of the data stored on Amazon AWS

10. A DRP's goal is to:

- A.** Identify monetary losses that a critical organizational function may suffer because of an incident.
- B.** Detail how to restore critical functions after an incident.
- C.** Provide a plan on how a company can continue running once an incident takes place.
- D.** Provide a risk evaluation at a specific moment in time.

11. An organization wants to test the existing DRP plan. Which of these methods is less intrusive but able to identify gaps in the plan?

- A.** Parallel test

- B.** Live test
- C.** Tabletop exercise
- D.** Warm site

12. Using a proprietary forensic tool for investigation relates to which of these reliability factors?

- A.** Clarity
- B.** Error rate
- C.** Credibility
- D.** Testability

13. Which of these tools is not an open source forensic application/suite?

- A.** SIFT
- B.** Sleuth Kit
- C.** FTK
- D.** Coroner's Toolkit

14. An online sports retailer is experiencing an increased volume of client web requests, as it just released new products with a heavy discount. The company is considering some options to cope with these requests. Which of these options is the best one?

- A.** WAF
- B.** RAID
- C.** Failover cluster
- D.** Load balancing cluster

15. A company's backup policy states that a full backup is performed every Monday at 9:00 P.M. Incremental backups are performed daily at 12:00 P.M. If a server goes down at 8:30 P.M. on Thursday evening, how many backups in total will be required to restore the server?

- A.** 1
- B.** 2

C. 3

D. 4

16. Which of the following is not one of the three forensic investigation phases?

A. Acquire

B. Validate

C. Analyze

D. Authenticate

17. A security administrator is reviewing corporate SIEM logs regarding possible incidents for a company using BYOD. Which of these activities is most likely a valid security incident?

A. Playing an online game

B. Browsing to Facebook

C. Downloading a movie

D. Using Google Drive

18. A type of malware has been identified on the corporate network. It has the ability to record keystrokes, and it is known to open TCP port 52145 on an infected system and receive command-and-control traffic from its creator using IP address 10.10.10.10. (Please note that this IP is private and is used specifically for the purposes of the question. In a real-life scenario, the attacker would be using a public IP address.) Which of these actions would most likely be taken during the “containment” phase of an incident response plan?

A. Block communication to 10.10.10.10 at the perimeter firewall.

B. Close TCP port 52145 on all infected systems.

C. Create an IDS rule for traffic to 10.10.10.10.

D. Use a web proxy to block offending traffic.

19. Source address affinity is used to direct traffic according to the:

A. User’s source address

B. Server’s source address

- C. User's source address and port combination
- D. User's geographical source location

20. In which phase would chain of custody be present?

- A. Acquire
- B. Authenticate
- C. Analyze
- D. All of the above

21. A company's backup policy states that a full backup is performed every Friday at 9:00 P.M. Differential backups are performed daily at 8:00 P.M. If a server goes down at 11:00 P.M. on Tuesday evening, how many backups in total will be required to restore the server?

- A. 1
- B. 2
- C. 3
- D. 4

22. What is one of the most important challenges of a "lessons learned" session?

- A. Ensure all systems are working normally after an incident.
- B. Verify no traces of malicious activity are present within the company network.
- C. Refrain from blaming specific individuals about the incident or its investigation.
- D. Ensure that processes are improved to deal with future incidents.

23. A company wants to select a dedicated alternative location for continuing its operations in the event of an incident, while minimizing operational downtime. Which of the following would be most appropriate for that purpose?

- A. Hot site
- B. Warm site
- C. Cold site

D. Mobile site

24. Which of the following is the last phase of incident response?

A. Recovery

B. Eradication

C. Containment

D. Implementation of new countermeasures

25. A company's backup policy states that a full backup is performed every Tuesday at 8:00 A.M. Incremental backups are performed daily at 6:00 P.M. If a server goes down at 4:30 P.M. on the following Sunday evening, how many full backups will be required to restore the server?

A. 1

B. 2

C. 3

D. 4

26. An organization performs a full weekly backup and is considering combining that with a differential or incremental one. Why would a full/differential backup be preferred over a full/incremental backup?

A. A full/differential backup takes less time to complete during the week.

B. When restoring a system, recovery time is reduced.

C. A full/incremental backup takes less time to complete during the week.

D. Differential backups are smaller in size.

27. If a company wants to ensure it maintains constant Internet connectivity, which of these options would be most appropriate?

A. Failover cluster

B. Load balancing

C. Redundant connection

D. Optical fiber

28. Which of the following is the definition of RTO?

- A.** How much data (at most) a business can sustain losing (i.e., three days of data)
- B.** How much time it takes (at most) for a system or service to be restored after an outage (i.e., three days to restore the web server after crashing)
- C.** How much time (at most) a crucial application or function can be unavailable without severe business impact
- D.** Identifying what the business impact of an incident is to the organization

29. Which of the following isn't a phase of the incident lifecycle?

- A.** Analyze evidence
- B.** Preparation
- C.** Eradication
- D.** Containment

30. Anything of particular interest that may be identified in a system is classified as a(n):

- A.** Alert
- B.** Threshold
- C.** Event
- D.** Baseline

31. A company's RPO is set at three days, while the MAO is four days. Which of these values indicates a realistic RTO?

- A.** 3
- B.** 5
- C.** 7
- D.** 9

32. A company wants to select a dedicated alternative location for continuing its operations in the event of an incident. They have some budget limitations and can cope with operating from an

alternative location within 24 to 48 hours of any major incident that might compromise the primary location. Which of the following would be most appropriate for that purpose?

- A. Hot site
- B. Warm site
- C. Cold site
- D. Mobile site

33. A police chief appointed John to lead a high-profile case involving organized crime, which requires the investigation of 100 machines with various forensic artifacts. John has been working as a senior forensic investigator for the last six years and has been teaching advanced forensics at the local police college. Which of these factors is the chief trying to support by appointing John on this case?

- A. Testability
- B. Error rate
- C. Acceptance
- D. Credibility

QUICK ANSWER KEY

- 1.** C
- 2.** A
- 3.** A
- 4.** D
- 5.** B
- 6.** C
- 7.** B
- 8.** A
- 9.** D
- 10.** B

11. C

12. D

13. C

14. D

15. D

16. B

17. C

18. A

19. A

20. B

21. B

22. C

23. A

24. D

25. A

26. B

27. C

28. B

29. A

30. C

31. A

32. B

33. D

ANSWERS A

1. In which of these phases of the incident response lifecycle would the identification of an appropriate tool for memory forensics fall into?

- A. Containment
- B. Detection, Analysis and Escalation
- C. Preparation
- D. Recovery

C. Software and hardware identification would normally take place in the preparation phase. During that phase, the team would highlight the appropriate hardware and software tools (i.e., write blockers, USB flash drives, evidence bags, field kits, incident handling platforms, memory forensics and data erasure software) that are required to adequately respond to any given incident.

A, B, and D are incorrect. **A** is incorrect because “containment” refers to using appropriate methods, tools, and processes (already identified and/or created during the preparation phase) to contain an incident already taking place in an organization. **B** is incorrect because “Detection, Analysis, and Escalation” would follow the preparation phase. At this point, an incident would be discovered (usually by some alert or report that the security team might receive) and then further investigated and escalated as appropriate. **D** is incorrect because “Recovery” refers to the processes and steps that are required in order to mitigate the incident and return to business as usual. In the case of a crypto-ransomware encrypting 100 machines, during the “Recovery” phase, those would most likely be reimaged and restored from backups so they can once again be used by the company employees.

2. Which of the following provides no fault tolerance?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 6

A. RAID 0 uses two or more disks and is primarily used for read/write operation improvement but is unable to provide any fault tolerance.

B, C, and D are incorrect. **B** is incorrect because RAID 1 has two disks using mirror configuration so if one of the disks fails, data is still retrievable from the copy that exists on the

second disk. **C** is incorrect because RAID 5 is able to offer fault tolerance as long as only a single drive fails. **D** is incorrect because RAID 6 provides fault tolerance. It is similar to RAID 5, but the main difference is that parity data is written on two drives (using at least four drives in total), extending the fault tolerance of the model to two drives.

3. Which of these statements regarding a “hot site” is inaccurate?

A. It is commonly usable within a few hours.

B. It includes network infrastructure.

C. Cost is quite significant.

D. Contains up-to-date organizational data.

A. A hot site is available in a very short time. Usually that is within one hour or even less, but not a few hours, as the organization using a hot site wants to be able to switch from the primary site to the alternative location in as little time as possible (often even in minutes).

B, C, and D are incorrect. **B** is incorrect because a “hot site” includes network infrastructure, water, and heating, allowing it to be usable immediately in case of any failure to the organization’s primary site. **C** is incorrect because cost is very high for a hot site. If a company wants to lower that, it could go with the option of a warm site, which would be more cost-efficient but takes more time to transfer operations to the alternative location. **D** is incorrect because the data contained in a “hot site” is usually constantly updated so it can be immediately available for use when it is required.

4. If a company server is unavailable for more than 45 minutes, core business functions are affected. That means that 45 minutes is the:

A. RPO

B. BIA

C. DRP

D. MAO

D. The MAO (maximum acceptable outage) is the time that a crucial application or service can be unavailable without severe business impact. For example, a private health provider might determine that patients queuing up to be served at their branch usually leave for a competitor if a doctor doesn’t triage them in the first 15 minutes of a visit. In that case, the MAO might be set to 15 minutes, as the provider starts losing business after that.

A, B, and C are incorrect. **A** is incorrect because the RPO (recovery point objective) is used to express how much data a business can sustain losing. For example, in the case of a military environment, the RPO may be set to 0 (or something close to it) if the related system controls troop deployment and personnel tracking. **B** is incorrect because BIA (business impact analysis) refers to identifying what the business impact of an incident is to the organization. **C** is incorrect because DRP (disaster recovery plan) details how to restore key business functions after an incident (or outage) takes place.

5. A security analyst identified outbound command and control traffic towards a malicious IP address. Which of these steps would most likely follow?

- A.** Investigate why there's traffic towards the malicious IP address.
- B.** Block the IP at the firewall.
- C.** Take a forensic image of the affected machine's memory.
- D.** Reimage the asset using a previous backup.

B. After identifying and analyzing the threat, containment would follow. That can be achieved by blocking the offending IP at the firewall level so no outbound or inbound malicious traffic can pass through the firewall.

A, C, and D are incorrect. **A** is incorrect because investigating why there is traffic towards the malicious IP should happen after the threat is contained. Once the block is implemented, additional investigation can take place to identify the reason for that traffic being generated. **C** is incorrect because containing the threat should be the first step. Since a malicious IP address has been identified, blocking traffic to and from that is key to ensure that no additional malicious traffic is present on the network. In addition, if that IP address is being used to exfiltrate data, then implementing a block is crucial so no organizational information is exfiltrated. **D** is incorrect because reimaging the asset would normally take place during the recovery stage. Once the investigation has been concluded, the assets can be reimaged and returned to the employees so they can be used again.

6. A forensic investigator needs to capture data for later analysis. Which one of these types of data should be captured first?

- A.** Remote logs
- B.** Local drive
- C.** Swap file
- D.** Backup tape

C. The swap file data is the most volatile from the options provided in the question. A swap file is a form of memory (an extension of the machine's RAM) that is stored on the machine's hard drive. When it's powered off, that file is deleted.

A, B, and D are incorrect. **A** is incorrect because any remote logs are stored at a remote location and shouldn't be affected by any activity taking place on the primary machine being investigated. **B** is incorrect because the machine's local hard disk contains data that can survive a system shutdown, which is not the case with the swap file. **D** is incorrect because a backup tape is a form of removable backup media and is certainly less volatile than the swap file data.

7. Which of the following does forensic software use to ensure no data tampering has taken place?

- A.** Serial numbers
- B.** Hashing
- C.** Digital signatures
- D.** Encryption

B. Forensic tools use hashing algorithms (commonly MD5 and SHA) to create a unique signature of the captured data. This can be later used to prove no data tampering has occurred since the hash of the original data has remained unaltered.

A, C, and D are incorrect. **A** is incorrect because a serial number can't be used to prove data tampering hasn't taken place. Forensic tools use serial numbers to identify the different devices being analyzed so an analyst can switch between those during an investigation. **C** is incorrect because forensic tools don't use digital signatures for file tampering purposes. **D** is incorrect because encryption isn't used to ensure evidence hasn't been altered. Encryption can only be used to allow for safe data storage by ensuring that if data is obtained by an unauthorized third party, it will remain inaccessible.

8. A company's client database is stored on a server so that online purchases can take place. The database needs to be accessible constantly even if the server goes down. Which of the following would be more suitable for that purpose?

- A.** Failover cluster
- B.** Load balancing
- C.** Full backup
- D.** Differential backup

A. A failover cluster is the most appropriate solution, as it can allow the database to be accessible if the server goes down. For that purpose, at least two servers are required, which can then be configured into a pair. One of them can be set to operate as the primary instance, while the other will be the secondary one. Once the secondary server detects any type of failure on the primary instance, it can assume its place as primary, thus guaranteeing minimal database downtime.

B, C, and D are incorrect. **B** is incorrect because load balancing is primarily used for performance improvement, not for fault tolerance. **C** is incorrect because a full backup will not guarantee that the database will be accessible all the time. If a server goes down, an administrator would need to intervene and manually restore that backup in addition to containing only accurate information up until the moment the backup was taken. If a full backup was taken a week ago, then after a server failure, the backup that is restored will contain data that is a week old, so all other transactions would have been lost. **D** is incorrect because a differential backup won't be able to ensure constant access to the server database, as any type of restoration would need to be manually performed by an administrator, with substantial associated time if any failure occurs out of normal business hours, in addition to any transactions that are not contained in the backup that was lost.

9. A company is hosting an SQL, cloud, and web server at its own premises. Which of these backup locations would you recommend as the safest option for backup storage and later retrieval of the web server's data?

- A.** A backup server placed in the server room, along with the web server
- B.** A magnetic tape stored on the premises but in a different building
- C.** An encrypted copy of the data backed up on the company cloud server
- D.** Backup of the data stored on Amazon AWS

D. A backup of the data on a third-party cloud is the safest option, as that ensures the data survives any physical disaster that the main server room faces.

A, B, and C are incorrect. **A** is incorrect because if the backup is placed in the server room, any physical disaster will destroy both the primary server and its backup. **B** is incorrect because a magnetic tape stored in a different building provides some assurance, but if a major physical disaster takes place (i.e., an earthquake or a tornado), that would most likely destroy the backup as well. **C** is incorrect because although the encrypted copy of the data stored on the company cloud might initially look like a good solution, the question mentions that the cloud is hosted on a server located in the server room, meaning the primary location of the web server. Hence, any physical disaster would result in losing that data.

10. A DRP's goal is to:

A. Identify monetary losses that a critical organizational function may suffer because of an incident.

B. Detail how to restore critical functions after an incident.

C. Provide a plan on how a company can continue running once an incident takes place.

D. Provide a risk evaluation at a specific moment in time.

B. The primary goal of a DRP (disaster recovery plan) is to provide a plan on how key business functions of an organization can continue to operate after an incident takes place.

A, C, and D are incorrect. **A** is incorrect because the identification of monetary losses from an incident is a primary deliverable of the BIA (business impact analysis). **C** is incorrect because a BCP (business continuity plan) depicts how a company can continue to operate once an incident takes place. **D** is incorrect because an evaluation of organizational risks at a specific moment in time is the goal of risk assessment.

11. An organization wants to test the existing DRP plan. Which of these methods is less intrusive but able to identify gaps in the plan?

A. Parallel test

B. Live test

C. Tabletop exercise

D. Warm site

C. A tabletop exercise is the least intrusive method but still quite capable of identifying any procedural gaps that may exist. During such an activity, personnel from various teams are gathered in a room, and a hypothetical incident is presented and analyzed, while each team describes the actions that would be performed. If there's any issue, it can be flagged for later review so it's appropriately addressed.

A, B, and D are incorrect. **A** is incorrect because a parallel test requires a lot of preparation as well as infrastructure resources. During such a test, primary business systems are still responsible for handling all traffic, while recovery systems are built and tested to review if production traffic can properly flow through them. **B** is incorrect because a live test requires the most preparation and is the riskiest approach. During such a test, primary systems are not used to handle traffic, as that will be directed to any redundant systems that exist. This is the most accurate method of determining how well a DRP plan can perform, but the risk of downtime is present since the secondary systems might not work as expected. **D** is incorrect because a warm site test entails substantial preparation, and if an organization intends to fully test the secondary

location, it might decide to divert all (or the majority) of the production traffic through that location, which can be quite intrusive should something not go as planned.

12. Using a proprietary forensic tool for investigation relates to which of these reliability factors?

- A. Clarity
- B. Error rate
- C. Credibility
- D. Testability

D. Testability relates to how much a tool or technique has been tested before being used in an investigation. Using a proprietary tool will not allow a variety of people from different industries to test the tool, hence its testability would be very low.

A, B, and C are incorrect. **A** is incorrect because clarity refers to how easy it is to explain the tool or technique and its results to the court. **B** is incorrect because error rate refers to the tool's statistics regarding errors. **C** is incorrect because credibility refers to the individual handling a tool or using a method and how credible he or she is, as well as how reproducible the technique is by similar specialists.

13. Which of these tools is not an open source forensic application/suite?

- A. SIFT
- B. Sleuth Kit
- C. FTK
- D. Coroner's Toolkit

C. FTK (Forensic Toolkit) is a tool developed by Access Data for forensic investigations. It provides the ability to capture a device (i.e., server, laptop, desktop, mobile) and acquire a forensic image of it for later analysis.

A, B, and D are incorrect. **A** is incorrect because SIFT (SANS Investigative Forensic Toolkit) is an open source suite of forensic tools developed by the SANS Institute that is based on Ubuntu 14.04. **B** is incorrect because the Sleuth Kit is an open source collection of CLI tools used for disk imaging and analysis. **D** is incorrect because the Coroner's Toolkit is a collection of open source tools aiding data recovery and system data analysis.

14. An online sports retailer is experiencing an increased volume of client web requests, as it just released new products with a heavy discount. The company is considering some options to cope with these requests. Which of these options is the best one?

- A. WAF
- B. RAID
- C. Failover cluster
- D. Load balancing cluster

D. A load balancing cluster is the best option for this scenario, as it is completely scalable and allows for large load optimization. When the organization experiences increased web requests, it can add more web servers to handle those accordingly. If traffic decreases with time, some servers can be powered down to reduce operational costs.

A, B, and C are incorrect. **A** is incorrect because a WAF (web application firewall) would be used for security (i.e., to filter web traffic), not for load optimization. **B** is incorrect because any type of RAID (Redundant Array of Independent Disks) would mainly be used for fault tolerance, not for performance optimization. **C** is incorrect because a failover cluster would be used to ensure continuous operation of the web server, allowing for an additional server (or more) to be used and handle the traffic in case the primary machine is unavailable.

15. A company's backup policy states that a full backup is performed every Monday at 9:00 P.M. Incremental backups are performed daily at 12:00 P.M. If a server goes down at 8:30 P.M. on Thursday evening, how many backups in total will be required to restore the server?

- A. 1
- B. 2
- C. 3
- D. 4

D. As per the organizational backup policy, a full backup will take place on Monday at 9:00 P.M. Following that, an incremental backup will take place on Tuesday at 12:00 P.M., as well as Wednesday and Thursday at the same time. After the last incremental backup (Thursday at 12:00 P.M.), the server goes down (8:30 P.M. on Thursday). That means all previously mentioned backups will need to be used to restore the server to its most current state, meaning all four (one full and three incremental ones).

A, B, and C are incorrect, as these are distractors.

16. Which of the following is not one of the three forensic investigation phases?

- A. Acquire
- B. Validate
- C. Analyze
- D. Authenticate

B. There is no forensic phase known as “validate.” Any type of required data validation would be performed during the “authenticate” phase.

A, C, and D are incorrect. **A** is incorrect because “acquire” is the phase where evidence from the scene is gathered (i.e., hard disk/memory capture, mobile phone acquisition). **C** is incorrect because “analyze” is the last phase of the process (following acquisition and authentication) where all captured data is analyzed in detail. It is worth noting that analysis is performed on copies of the original data (which should have been obtained during the “acquire” phase) so that no tampering of the original evidence takes place. **D** is incorrect because the “authenticate” phase includes evidence control and proof that it hasn’t been modified (i.e., by using proper chain of custody forms).

17. A security administrator is reviewing corporate SIEM logs regarding possible incidents for a company using BYOD. Which of these activities is most likely a valid security incident?

- A. Playing an online game
- B. Browsing to Facebook
- C. Downloading a movie
- D. Using Google Drive

C. Depending on how strict the organizational policies are, it is possible that all mentioned activities constitute a security incident (although this is highly unlikely in the case of BYOD, or bring your own device). However, regardless of that, most companies would consider any user downloading a movie while in the corporate network a genuine security incident indicating a policy violation, as the user is most probably downloading copyrighted material.

A, B, and D are incorrect. **A** is incorrect because playing an online game is not necessarily a security incident (especially in the case of BYOD), as that may be allowed by the corporate policy. **B** is incorrect because browsing to social media is often allowed by most companies (although sometimes access hours may be restricted). In addition, when using BYOD, companies are aware that users using their personal devices would most likely be accessing their social media or using personal applications to some extent but tend to focus on the benefits of BYOD.

D is incorrect because using Google Drive doesn't necessarily indicate a security incident, especially because some companies use Google Drive for business purposes.

18. A type of malware has been identified on the corporate network. It has the ability to record keystrokes, and it is known to open TCP port 52145 on an infected system and receive command-and-control traffic from its creator using IP address 10.10.10.10. (Please note that this IP is private and is used specifically for the purposes of the question. In a real-life scenario, the attacker would be using a public IP address.) Which of these actions would most likely be taken during the "containment" phase of an incident response plan?

A. Block communication to 10.10.10.10 at the perimeter firewall.

B. Close TCP port 52145 on all infected systems.

C. Create an IDS rule for traffic to 10.10.10.10.

D. Use a web proxy to block offending traffic.

A. The primary goal of the "containment" phase is to stop further malicious activity. Since the malware is communicating with its creator via IP address 10.10.10.10, blocking that traffic at the perimeter firewall will ensure that all such communication ceases. Even if there are multiple malware instances within different machines in the company network, none of them will be able to communicate with the malware's creator to exfiltrate information or receive instructions.

B, C, and D are incorrect. **B** is incorrect because closing TCP port 52145 will not necessarily stop communication with the malicious IP address, which might also be used as a source for other attacks (not necessarily using the same port). **C** is incorrect because creating an IDS rule for traffic to 10.10.10.10 is beneficial for receiving alerts regarding offending traffic but will not contain it, as the IDS doesn't have an ability to block that communication. **D** is incorrect because a web proxy will only be able to filter web traffic, which will have no effect in this case, as the traffic in question is not reported to be web related.

19. Source address affinity is used to direct traffic according to the:

A. User's source address

B. Server's source address

C. User's source address and port combination

D. User's geographical source location

A. Source address affinity uses the source IP address of a user. It is used for load balancing and allows the load balancer to redirect requests originating from the same client IP address to be served by the same server.

B, C, and D are incorrect. **B** is incorrect because requests will originate from the client, not the server side. **C** is incorrect because source affinity only works by identifying a request from the same client IP, but the client source port isn't relevant. **D** is incorrect because the user's geographical location doesn't change anything, as source affinity only uses the client IP address, regardless of its physical location.

20. In which phase would chain of custody be present?

- A. Acquire
- B. Authenticate
- C. Analyze
- D. All of the above

B. Chain of custody refers to controlling evidence from the moment of collection, which is something that would happen during the "authenticate" phase. The purpose of a chain of custody is to demonstrate that adequate control has been enforced to ensure that the evidence hasn't been tampered with and has been properly handled.

A, C, and D are incorrect. **A** is incorrect because "acquire" is the phase where evidence from the scene is gathered (i.e., hard disk/memory capture, mobile phone acquisition). **C** is incorrect because "analyze" is the last phase of the process (following acquisition and authentication), where all captured data is analyzed in detail. It is worth noting that the analysis is performed on copies of the original data (which should have been obtained during the "acquire" phase) so that no tampering of the original evidence takes place. **D** is incorrect because chain of custody starts being used at the "authenticate" phase.

21. A company's backup policy states that a full backup is performed every Friday at 9:00 P.M. Differential backups are performed daily at 8:00 P.M. If a server goes down at 11:00 P.M. on Tuesday evening, how many backups in total will be required to restore the server?

- A. 1
- B. 2
- C. 3
- D. 4

B. As per the company policy, a full backup will first be taken on Friday at 9:00 P.M. Following that, differential backups will be taken on Saturday, Sunday, Monday, and Tuesday at 8:00 P.M. Finally, after those are taken, the server goes down. As you may remember, when a differential backup is taken, all changes from the previous full backup are recorded. Since the

server went down at 11:00 P.M. on Tuesday and a differential backup was taken three hours prior to that, a restoration process will only require that backup in addition to the full backup taken last Friday. Hence, a total of two backups will be required.

A, C, and D are incorrect, as these are distractors.

22. What is one of the most important challenges of a “lessons learned” session?

- A.** Ensure all systems are working normally after an incident.
- B.** Verify no traces of malicious activity are present within the company network.
- C.** Refrain from blaming specific individuals about the incident or its investigation.
- D.** Ensure that processes are improved to deal with future incidents.

C. After any given incident, most company employees involved with the investigation always have concerns about how they performed and if the incident could have been avoided altogether if they had acted differently. However, a “lessons learned” session should focus on how to improve things for future responses and not dwell on blaming specific people for how they performed and if they were solely responsible for the incident taking place.

A, B, and D are incorrect. **A** is incorrect because during the “recovery” phase, all systems would be brought back to normal either by reverting any damage caused by the incident or by restoring them from a previous backup (which is a common method to ensure that nothing might have been overlooked). **B** is incorrect because ensuring all threats have been neutralized would normally take place during the “eradication” phase, where the goal is to ensure the components of any given threat are removed from the environment. **D** is incorrect because ensuring that processes are improved to deal with future incidents is the actual goal of a “lessons learned” session, but the challenge is to do that without focusing on blaming specific individuals for the incident or how it was investigated.

23. A company wants to select a dedicated alternative location for continuing its operations in the event of an incident, while minimizing operational downtime. Which of the following would be most appropriate for that purpose?

- A.** Hot site
- B.** Warm site
- C.** Cold site
- D.** Mobile site

A. A “hot site” is the best choice for this purpose, as it offers a dedicated alternative location for continuing operations in the event of any incident rendering the primary location unusable. Such a site has everything in place (i.e., HVAC, network infrastructure, hardware) so the organization can continue operating.

B, C, and D are incorrect. **B** is incorrect because a “warm site” offers some of the capabilities of a hot site but not all of them (i.e., commonly the data from the primary location is not up to date and a synchronization needs to take place before commencing operations). In addition, it takes some time before a warm site is fully activated for use, meaning that it won’t offer minimal operational downtime, which is what is required here. **C** is incorrect because a “cold site” will take substantial resources and time to activate, as it doesn’t have any hardware or software set up. **D** is incorrect because a mobile site is used when an organization doesn’t require a dedicated alternative location. In addition, it usually consists of necessary hardware, but software and data are not kept up to date, which means that some time for setting it up is required, thus it does not offer minimal operational downtime.

24. Which of the following is the last phase of incident response?

A. Recovery

B. Eradication

C. Containment

D. Implementation of new countermeasures

D. The last incident response phase is called “lessons learned/implementation of new countermeasure” (most commonly known as “lessons learned”). The goal of this phase is to review if current processes are adequate or need to be amended in order to respond to any future incidents more efficiently. In addition, evaluation of current countermeasures takes place, as new ones might need to be introduced to protect the environment from future attacks.

A, B, and C are incorrect. **A** is incorrect because during “recovery,” actions are taken to restore system operation to a known good state, like the one prior to the incident. **B** is incorrect because during “eradication,” the root cause of the issue is identified and eliminated. **C** is incorrect because during “containment,” the incident will be contained and further damage to the business stopped.

25. A company’s backup policy states that a full backup is performed every Tuesday at 8:00 A.M. Incremental backups are performed daily at 6:00 P.M. If a server goes down at 4:30 P.M. on the following Sunday evening, how many full backups will be required to restore the server?

A. 1

B. 2

C. 3

D. 4

A. Full backups are performed on a weekly basis (every Tuesday). Since the server went down on the following Sunday, only one full backup is available for the restoration. Note that the question refers to the number of full backups, not the total number of backups required.

B, C, and D are incorrect, as these are distractors.

26. An organization performs a full weekly backup and is considering combining that with a differential or incremental one. Why would a full/differential backup be preferred over a full/incremental backup?

A. A full/differential backup takes less time to complete during the week.

B. When restoring a system, recovery time is reduced.

C. A full/incremental backup takes less time to complete during the week.

D. Differential backups are smaller in size.

B. Using a full/differential backup results in reduced recovery time, as the differential backup captures any changes from the last full backup. That means that at any given moment, a total of two backups will be required to perform system restoration: the most recent full weekly backup and the most recent differential one.

A, C, and D are incorrect. **A** is incorrect because a full/differential backup takes more time to be completed during the week, as all changes from the moment the full weekly backup was taken need to be recorded. **C** is incorrect because a full/incremental backup does take less time to complete during the week (i.e., if taken daily, it records daily database transactions, which takes a relatively short time). However, that won't be a reason for choosing a full/differential backup solution over a full/incremental one. It's rather a statement regarding how the full/differential backup method works. **D** is incorrect because differential backups are larger in size (not smaller) because they include all changes since the last full weekly backup was taken, which may require substantial space.

27. If a company wants to ensure it maintains constant Internet connectivity, which of these options would be most appropriate?

A. Failover cluster

B. Load balancing

C. Redundant connection

D. Optical fiber

C. The best method to maintain “close to constant” Internet connectivity is to have multiple connections (at least two) in order to be able to use an alternative connection in the event of any issue. Many organizations use leased lines to get constant connectivity at very good upload and download speeds. Having a second leased line as a redundant connection is a great option, but since the cost of those is quite high, more cost-effective options are usually preferred (like a very low-end leased line or very good-rate ADSL option, which is significantly cheaper) so they can efficiently cope when the primary line goes down.

A, B, and D are incorrect. **A** is incorrect because a failover cluster is used to provide fault tolerance (usually relating to servers and services), not connection redundancy. **B** is incorrect because load balancing is used to distribute load efficiently among different servers (i.e., file or web servers) and doesn’t aim to address network connection redundancy. Although it may offer redundancy from an aspect of a server or service going down (i.e., when having three servers and one goes down, the load balancer would distribute remaining requests to the two remaining servers), network connectivity to the load balancer is still required in order to make that happen. **D** is incorrect because using optical fiber would offer great speeds, would support longer distances, would be resistant to electromagnetic interference, and would offer great security, but all those advantages don’t have anything to do with maintaining constant connectivity. This means that an optical fiber may still become unavailable, hence the company needs to have an alternative connection available when that happens.

28. Which of the following is the definition of RTO?

A. How much data (at most) a business can sustain losing (i.e., three days of data)

B. How much time it takes (at most) for a system or service to be restored after an outage (i.e., three days to restore the web server after crashing)

C. How much time (at most) a crucial application or function can be unavailable without severe business impact

D. Identifying what the business impact of an incident is to the organization

B. The definition of RTO (recovery time objective) is the time it takes (at most) for a system or service to be restored after an outage (i.e., three days to restore the web server after going down). This is what usually drives business processes and administrators to architect proper systems that can support service restoration at the level that is required for a company to properly function.

A, C, and D are incorrect. **A** is incorrect because the amount of data a business can sustain losing (i.e., three days of data) expresses the RPO (recovery point objective). This may vary but in critical environments can be quite challenging to achieve (i.e., a SCADA [supervisory control and data acquisition] system controlling water pressure might have an RPO of minutes in order

for operations to properly continue). **C** is incorrect because the time a crucial application or function can be unavailable without severe business impact taking place refers to the MAO (maximum acceptable outage), also known as MTD (maximum tolerable downtime). **D** is incorrect because identification of the business impact of an incident to the organization is the definition of BIA (business impact analysis).

29. Which of the following isn't a phase of the incident lifecycle?

- A. Analyze evidence
- B. Preparation
- C. Eradication
- D. Containment

A. "Analyze evidence" (also known as "analyze") is one of the phases of a computer forensics investigation, during which all previously acquired data is fully analyzed to investigate the incident that took place.

B, C, and D are incorrect. **B** is incorrect because during "preparation" the team would highlight the appropriate tools that are required to adequately respond to any given incident and review any processes the organization has in place, as well as enforce the company defenses to ensure adequate preparation is in place prior to an incident. **C** is incorrect because "eradication" is when the root cause of the issue is identified and eliminated. **D** is incorrect because "containment" refers to using appropriate methods, tools, and processes (which were already identified and/or created during the preparation phase) to contain an incident already taking place in an organization.

30. Anything of particular interest that may be identified in a system is classified as a(n):

- A. Alert
- B. Threshold
- C. Event
- D. Baseline

C. An "event" can be classified as any activity that may be identified in a system. A user named James logging on a system might be classified as an event.

A, B, and D are incorrect. **A** is incorrect because an "alert" (usually identified on an IDS/IPS, SIEM, or some other similar system or tool) is a notification for a particular activity of interest that someone (i.e., security analyst or administrator) would like to be informed of and

usually is created after a rule of some kind is triggered, which is set to look for the activity in question. **B** is incorrect because a “threshold” can be defined as a specific limit configured in a rule or tool that, once surpassed, an appropriate action is taken. **D** is incorrect because a “baseline” is a known initial state and is commonly used in configuration management to describe an initial image of an operating system that an organization may be using, or it may refer to security tools that require a baseline to be set before performing related analysis (i.e., anomaly-based IDS, which needs to have a baseline of the network created in order to alert on what events might be valid security incidents).

31. A company’s RPO is set at three days, while the MAO is four days. Which of these values indicates a realistic RTO?

- A. 3
- B. 5
- C. 7
- D. 9

A. The time a crucial application or function can be unavailable without severe business impact refers to the MAO (maximum acceptable outage), which is stated as four days. The RTO is the time it takes (at most) for a system or service to be restored after an outage. Since the business can cope for up to four days (MAO set at four days), the RTO needs to be less than the MAO in order for the business to successfully recover and continue operations. From the provided options, only the first one (three days) allows for a viable RTO.

B, C, and D are incorrect, as these are distractors.

32. A company wants to select a dedicated alternative location for continuing its operations in the event of an incident. They have some budget limitations and can cope with operating from an alternative location within 24 to 48 hours of any major incident that might compromise the primary location. Which of the following would be most appropriate for that purpose?

- A. Hot site
- B. Warm site
- C. Cold site
- D. Mobile site

B. A “warm site” would be most suitable in this case. It offers some of the capabilities of a “hot site” but not all of them (i.e., commonly the data from the primary location is not up to date and synchronization needs to take place before commencing operations). However, the cost of maintaining a “warm site” is significantly lower than that of a “hot site.” In addition, such a site

would be usually available within a matter of hours, commonly up to 24, which aligns with the organization's requirements.

A, C, and D are incorrect. **A** is incorrect because a “hot site,” although available in a very short time, would require a substantial budget, and the company in this scenario only has a limited budget to spend, thus making this option not suitable for what is required. **C** is incorrect because a “cold site” (although quite affordable) will take substantial resources and time to activate, as it doesn't have any hardware or software set up. **D** is incorrect because a “mobile site” is used when an organization doesn't require a dedicated alternative location.

33. A police chief appointed John to lead a high-profile case involving organized crime, which requires the investigation of 100 machines with various forensic artifacts. John has been working as a senior forensic investigator for the last six years and has been teaching advanced forensics at the local police college. Which of these factors is the chief trying to support by appointing John on this case?

- A.** Testability
- B.** Error rate
- C.** Acceptance
- D.** Credibility

D. The police chief is trying to ensure that the individual who handles the case has adequate credibility, which will aid John when the case goes to court. Credibility refers to the individual handling a tool or using a method and how credible he or she is, as well as how reproducible the technique is by similar specialists. Since John has been working at the police force as a senior forensic investigator for six years and has been teaching forensics at the police academy, he should be credible enough to lead the investigation in question.

A, B, and C are incorrect. **A** is incorrect because testability relates to how much a tool or technique has been tested before being used in an investigation, but the focus of the question is the investigator, not a tool or technique. **B** is incorrect because the error rate refers to a tool's statistics regarding errors, but the question focuses on the investigator and his qualifications. **C** is incorrect because acceptance refers to a scientific theory or technique and how much that has been peer reviewed, not the actual investigator handling a case.

CHAPTER 5

Cryptography

This chapter includes questions from the following objectives:

- 5.1 Understand fundamental concepts of cryptography
- 5.2 Understand reasons and requirements for cryptography
- 5.3 Understand and support secure protocols
- 5.4 Understand Public Key Infrastructure (PKI) systems

This chapter contains a complete set of questions and answers regarding the fundamental principles of cryptography (necessary terminology, hashing, symmetric and asymmetric encryption, steganography, IPSEC) and why cryptography is used, as well as its key requirements (types of data, handling techniques, data protection). It also includes questions about secure protocol operation (basic protocol ports and operation, remote access solutions) and PKI (certificates, certificate authorities, key escrow, and certificate trust).

Q QUESTIONS

1. Consider two text files: The first one contains the sentence “this is an SSCP example” and the second one contains the sentence “this is another SSCP example with different text”. The MD5 hash of the first file is 063ba1995095c95c2e2184837fd47067, and the MD5 hash of the second file is 063ba1995095c95c2e2184837fd47067. What is this an example of?

- A. Work factor
- B. Collusion
- C. Rainbow tables
- D. Collision

2. Which of these data labels would be most suitable for client information that a company is storing?

- A. Confidential
- B. Private
- C. Sensitive
- D. Public

3. James is the head of sales and needs to send a proposal regarding a new product to the CEO. He changes its value by adding 10 percent with the intent of getting that for himself afterwards. Which of these terms describes this scenario?

- A. Data inference
- B. Deduplication
- C. Data diddling
- D. Data interference

4. Consider the command shown in the following illustration. Which of these statements is accurate?

```
[nick:~ Nick$ ping www.google.com
PING www.google.com (216.58.204.4): 56 data bytes
64 bytes from 216.58.204.4: icmp_seq=0 ttl=53 time=4.352 ms
64 bytes from 216.58.204.4: icmp_seq=1 ttl=53 time=3.001 ms
64 bytes from 216.58.204.4: icmp_seq=2 ttl=53 time=3.739 ms
64 bytes from 216.58.204.4: icmp_seq=3 ttl=53 time=4.043 ms
64 bytes from 216.58.204.4: icmp_seq=4 ttl=53 time=4.080 ms
64 bytes from 216.58.204.4: icmp_seq=5 ttl=53 time=4.420 ms
64 bytes from 216.58.204.4: icmp_seq=6 ttl=53 time=3.031 ms
^C
--- www.google.com ping statistics ---
7 packets transmitted, 7 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 3.001/3.809/4.420/0.542 ms
```

- A. The operating system in use is a version of Windows.
- B. The command shows echo reply packets.
- C. The remote host is unreachable.
- D. The command completed its execution and terminated automatically.

5. A company wants to ensure bandwidth is preserved for business-related traffic (related to the internal company network) while employees access the corporate network remotely. Which of the following would be the most appropriate option?

- A. Split tunneling
- B. Tunnel mode

C. Transport mode

D. RADIUS

6. Which of the following is the most accurate statement regarding what TLS uses for key exchange and session encryption?

A. Symmetric encryption

B. Asymmetric encryption

C. Both symmetric and asymmetric encryption

D. SSL

7. Which of the following can't be achieved with digital signatures?

A. Integrity

B. Confidentiality

C. Nonrepudiation

D. Authentication

8. A key escrow is used for storing:

A. Public keys

B. Digital certificates

C. WOT certificates

D. Private keys

9. Which of these algorithms provides integrity in addition to authenticity?

A. MD5

B. RSA

C. SHA-3

D. HMAC-MD5

10. Jane is a system administrator and wants to download a new tool from a third-party website (the official vendor's server doesn't have enough bandwidth, and Jane achieves a very slow download speed, and the file is quite large). After downloading the tool, she calculates an MD5 hash of it. When attempting to compare the calculated MD5 hash with the official vendor's hash as stated on the vendor website, she notices they are different. After repeating the whole process, she still doesn't see the hashes matching. What would be the most probable reason for this?

- A. The file that was downloaded from the third-party server is corrupted.
- B. The file that was downloaded from the third-party server contains malware.
- C. The file that was downloaded from the third-party server doesn't seem to be the official vendor file.
- D. This is an example of a collision.

11. Which of the following would be most suitable to protect files stored on a hard drive?

- A. AES-128
- B. TLS 1.1
- C. RSA-512
- D. 3-DES

12. Which of these ports would be used when a user attempts to send an e-mail?

- A. TCP 110
- B. TCP 25
- C. UDP 53
- D. TCP 443

13. Which person is responsible for defining a company's data classification scheme?

- A. Data owner
- B. Data custodian
- C. User
- D. Management

14. Which of the following would be used for a new certificate registration request?

- A. WOT
- B. RA
- C. CRL
- D. OCSP

15. Which of these authentication protocols would you use to support mutual authentication?

- A. CHAP
- B. PAP
- C. MS-CHAPv2
- D. MS-CHAPv3

16. You want to ensure your company's e-mails are encrypted to maintain communication confidentiality. Which of these options would you choose for that purpose?

- A. DKIM
- B. SPF
- C. S/MIME
- D. DMARC

17. Consider this cipher suite:

TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384

What key exchange algorithm is being used?

- A. TLS
- B. ECDHE_RSA
- C. CAMELIA_256_CBC
- D. SHA384

18. Your company wants to simplify the process of employees remotely connecting to the company network via VPN. The main requirement is to be able to connect using a web browser. Which of these methods would be most appropriate for that goal?

- A. TLS VPN
- B. Transport mode
- C. Tunnel mode
- D. SSH

19. Which of these attacks would be most suitable to an attacker who has samples of the plaintext and ciphertext messages?

- A. Rainbow table attack
- B. Dictionary attack
- C. Chosen plaintext
- D. Known plaintext

20. An attacker is using a text file's spaces and tabs to store information. Which of the following is this an example of?

- A. Encoding
- B. Hashing
- C. Steganography
- D. Encryption

21. You have a file containing the sentence "This is the plaintext". You use a key of "This is my first key" and generate the following ciphertext: "yu9+m0/mY7rF/6dv6SXNVC7D69Q7b1N4". You encrypt the same file again but with a different key: "This is my second key". The result is "yu9+m0/mY7rF/6dv6SXNVC7D69Q7b1N4". What is this an example of?

- A. Key clustering
- B. Collision
- C. Collusion

D. Work factor

22. You just received an e-mail from your bank account manager that your account has been breached. She kindly asks you to reset your password and emphasizes not to disclose it to anyone who may be posing as a bank employee over the phone. You follow the URL contained in her e-mail, and your browser presents you with the following message: “This connection is untrusted.” What is the most likely reason for that?

- A.** Certificate has expired.
- B.** You have received a phishing e-mail.
- C.** Bank website experiencing technical issues.
- D.** A self-signed certificate is being used.

23. Which of the following would you use to check if a certificate has been revoked, using the least amount of network resources?

- A.** CRL
- B.** CLR
- C.** RA
- D.** OCSP

24. James wants to communicate with Dianne. For that purpose, he encrypts a message with Dianne’s public key before sending it to her. Which of the following is he trying to enforce?

- A.** Integrity
- B.** Authentication
- C.** Confidentiality
- D.** Nonrepudiation

25. If a group of eight people uses symmetric key cryptography, how many keys are required to ensure that everyone’s communications are secured while each pair uses a different encryption key?

- A.** 8
- B.** 28

C. 32

D. 64

QUICK ANSWER KEY

1. D

2. B

3. C

4. B

5. A

6. C

7. B

8. D

9. D

10. C

11. A

12. B

13. D

14. B

15. C

16. C

17. B

18. A

19. D

20. C

21. A

22. B

23. D

24. C

25. B

ANSWERS A

1. Consider two text files: The first one contains the sentence “this is an SSCP example” and the second one contains the sentence “this is another SSCP example with different text”. The MD5 hash of the first file is 063ba1995095c95c2e2184837fd47067, and the MD5 hash of the second file is 063ba1995095c95c2e2184837fd47067. What is this an example of?

A. Work factor

B. Collusion

C. Rainbow tables

D. Collision

D. The two MD5 file hashes are identical; however, the content of the two files is different. This is known as a collision and relates to one of the basic principles of a hashing algorithm, which is to be collision free, meaning that two different files should never generate the same hash value. Even a slight change in a file that is used as input in a hashing algorithm should generate a different output.

A, B, and C are incorrect. **A** is incorrect because the work factor relates to cryptography and expresses the time and effort that it would take to decrypt an encrypted message. **B** is incorrect, as it refers to the act of one or more individuals or companies conspiring to commit fraud. Although the syntax might be similar to the correct answer, this term has actually nothing to do with hashing algorithms. **C** is incorrect, as rainbow tables contain precomputed hash values intended to provide a reverse lookup resource for hash values.

2. Which of these data labels would be most suitable for client information that a company is storing?

A. Confidential

B. Private

C. Sensitive

D. Public

B. Client information would most probably be labeled private, which is a label reserved for data that is to be used strictly internally. If such data is released to the public domain, serious damage may be caused to the organization. In addition, some client data would most probably be PII (personally identifiable information), and disclosure of it would be an offense that needs to be reported to the appropriate regulating body, which may impose related fines.

A, C, and D are incorrect. **A** is incorrect because data labeled confidential would most likely be used to indicate important information that is crucial to the organization and its operation (i.e., may include internal processes, proprietary research information about new products, incident response protocols, etc.). **C** is incorrect because sensitive information does require a degree of protection (because any sensitive data leak might result in damage for the company) but its nature is not as crucial as private or confidential information. Examples may include company presentations and organizational structure information. **D** is incorrect because public data is publicly available information, and client information wouldn't be classified as such.

3. James is the head of sales and needs to send a proposal regarding a new product to the CEO. He changes its value by adding 10 percent with the intent of getting that for himself afterwards. Which of these terms describes this scenario?

A. Data inference

B. Deduplication

C. Data diddling

D. Data interference

C. Data diddling is a term used to describe any unauthorized data change before or while information is entered into a system. In this scenario, James altered the product price before the CEO reviewed it, with the goal of getting 10 percent for himself once that purchase is finally approved.

A, B, and D are incorrect. **A** is incorrect because data inference is used to describe the case where an individual is able to gain knowledge about restricted data (usually labelled secret or classified) by accessing public information. **B** is incorrect because deduplication is used to describe a file being stored in a system without any duplicates existing. It's a concept commonly used to ensure that only one file copy is kept and anyone requiring access to its contents is always able to access the most up-to-date information, while saving valuable storage space. **D** is incorrect because there's no associated term named data interference. This option is a distractor.

4. Consider the command shown in the following illustration. Which of these statements is accurate?

```
[nick:~ Nick$ ping www.google.com
PING www.google.com (216.58.204.4): 56 data bytes
64 bytes from 216.58.204.4: icmp_seq=0 ttl=53 time=4.352 ms
64 bytes from 216.58.204.4: icmp_seq=1 ttl=53 time=3.001 ms
64 bytes from 216.58.204.4: icmp_seq=2 ttl=53 time=3.739 ms
64 bytes from 216.58.204.4: icmp_seq=3 ttl=53 time=4.043 ms
64 bytes from 216.58.204.4: icmp_seq=4 ttl=53 time=4.080 ms
64 bytes from 216.58.204.4: icmp_seq=5 ttl=53 time=4.420 ms
64 bytes from 216.58.204.4: icmp_seq=6 ttl=53 time=3.031 ms
^C
--- www.google.com ping statistics ---
7 packets transmitted, 7 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 3.001/3.809/4.420/0.542 ms
```

- A. The operating system in use is a version of Windows.
- B. The command shows echo reply packets.
- C. The remote host is unreachable.
- D. The command completed its execution and terminated automatically.

B. The output is from the execution of a ping command where echo request packets are being sent to www.google.com, which in turn replies with a number of echo reply messages.

A, C, and D are incorrect. A is incorrect because, by default, a ping command execution in Windows will send four packets to the target host. However, here there are seven echo reply packets, thus indicating that another operating system is in use. C is incorrect because the remote host isn't unreachable. ICMP echo reply messages are being successfully received, indicating the remote host is up and running and responding with echo reply. If it wasn't reachable, a response of "Request timeout" would be present in the response messages. Also, at the end of the command output there's a statement of "0.00% packet loss." If the host was unreachable, that would be stating some percentage of packet loss (in the case where the host was completely unreachable to echo request packets, that would state "100.00% packet loss"). D is incorrect because the command didn't terminate automatically, as it can be seen that the user used ^C to manually terminate its execution.

5. A company wants to ensure bandwidth is preserved for business-related traffic (related to the internal company network) while employees access the corporate network remotely. Which of the following would be the most appropriate option?

- A. Split tunneling
- B. Tunnel mode

C. Transport mode

D. RADIUS

A. A split tunnel is a common method of preserving bandwidth for any organization, as it ensures only the interesting traffic traverses the corporate VPN tunnel. Any traffic that doesn't need to be encrypted will reach its destination directly, without going through that VPN tunnel. A classic example is websites like social media. When a user connects to the company VPN, any organization-related traffic is encrypted via the VPN tunnel. However, when the user browses to Facebook, that traffic doesn't go through the VPN tunnel.

B, C, and D are incorrect. **B** is incorrect because tunnel mode is a IPSEC mode in which the whole packet is encrypted. This can't be configured in such a way as to allow only business-related traffic to go through the VPN, which is something that split tunneling can achieve. However, split tunneling can be configured with different tunneling protocols. **C** is incorrect because transport mode is a IPSEC mode in which only the packet payload is encrypted (not the packet header). Split tunneling can also be configured for use with this tunneling protocol, but it's not transport mode that fulfils the requirement, but rather split tunneling. **D** is incorrect because RADIUS provides authentication, authorization, and accounting but wouldn't be suitable in this scenario, as the question focuses on bandwidth consumption when the traffic relates to business.

6. Which of the following is the most accurate statement regarding what TLS uses for key exchange and session encryption?

A. Symmetric encryption

B. Asymmetric encryption

C. Both symmetric and asymmetric encryption

D. SSL

C. TLS (Transport Layer Security) uses both symmetric and asymmetric encryption. More specifically, it initially uses asymmetric encryption to share a secret (symmetric) key. Once that has been achieved, it uses that symmetric key to encrypt all further communications.

A, B, and D are incorrect. **A** is incorrect because it uses symmetric encryption to encrypt the session after the secret key has been exchanged. **B** is incorrect because it uses asymmetric encryption for key exchange, but after that, symmetric encryption is used to encrypt all further communications. **D** is incorrect because TLS doesn't use SSL. SSL was the first name that the protocol had, and after SSLv3 the name changed to TLS (with version 1.1 being the first one out).

7. Which of the following can't be achieved with digital signatures?

- A. Integrity
- B. Confidentiality
- C. Nonrepudiation
- D. Authentication

B. A digital signature is a message's encrypted hash. The purpose of using a digital signature is to provide message integrity as well as sender authentication and nonrepudiation. However, in order to enforce confidentiality, the message would need to be encrypted, but as stated earlier it is the hash of the message that is being encrypted, not the message itself. A third party can still intercept and read that message.

A, C, and D are incorrect. **A** is incorrect because a digital signature provides integrity. Since the e-mail's content is hashed prior to being sent, the recipient would be able to identify any unauthorized modification that a third party may have made if the message was intercepted and altered. **C** is incorrect because nonrepudiation is something that a digital signature can provide. That is because in order for the digital signature to be created, the sender's private key needs to be used, which is something that only the sender should have. **D** is incorrect because digital signatures provide authentication, since they are a form of proof to the recipient that any message actually came from a valid sender (being the digital signature's owner).

8. A key escrow is used for storing:

- A. Public keys
- B. Digital certificates
- C. WOT certificates
- D. Private keys

D. A key escrow is used for safely keeping private keys in order to be able to retrieve them in the event the company's private key is lost or damaged. Only specific people usually have permission to retrieve private keys from a key escrow facility due to the sensitivity of their nature.

A, B, and C are incorrect. **A** is incorrect because public keys are publicly available and are contained in an individual's digital certificate. **B** is incorrect because any digital certificates would be issued and managed by the certificate authority. Please note that a CA (certificate authority) may be external (i.e., Symantec or VeriSign) or internal to the organization (i.e., when a certificate is needed for a system that is used only internally). **C** is incorrect because there's no such term as WOT certificate. WOT (web of trust) is a decentralized trust model with the goal of allowing a system to establish authenticity between a public key and its owner.

9. Which of these algorithms provides integrity in addition to authenticity?

- A. MD5
- B. RSA
- C. SHA-3
- D. HMAC-MD5

D. HMAC (Hash-based Message Authentication Code) uses a secret key only known to the sender and recipient of a message. As such, it is used in combination with other hashing algorithms (like MD5 or SHA). When a sender wants to send a message (hashed with MD5), the first step is to create that message's MD5 hash. After that, the secret key is used to create the HMAC-MD5, which is what will be sent to the recipient. Since only the sender and recipient know that key, it is possible to follow the process at the recipient's side and establish that the message originated from the sender it was expected from in addition to not having been modified.

A, B, and C are incorrect. **A** is incorrect because MD5 can only provide message integrity, but no authenticity is supported (unless combined with HMAC). **B** is incorrect because RSA (Rivest Shamir Adleman) is an asymmetric encryption algorithm and is not used to enforce integrity or authenticity. **C** is incorrect because SHA-3 can only provide message integrity, but no authenticity is supported. That can only be achieved by combining it with HMAC.

10. Jane is a system administrator and wants to download a new tool from a third-party website (the official vendor's server doesn't have enough bandwidth, and Jane achieves a very slow download speed, and the file is quite large). After downloading the tool, she calculates an MD5 hash of it. When attempting to compare the calculated MD5 hash with the official vendor's hash as stated on the vendor website, she notices they are different. After repeating the whole process, she still doesn't see the hashes matching. What would be the most probable reason for this?

- A. The file that was downloaded from the third-party server is corrupted.
- B. The file that was downloaded from the third-party server contains malware.
- C. The file that was downloaded from the third-party server doesn't seem to be the official vendor file.
- D. This is an example of a collision.

C. A mismatch of the two hashes would commonly point to the files being different. If the file that was downloaded from the third-party website was the original file that is hosted on the official vendor's web server, the same hash should have been calculated. Since that is not the case, it mostly likely points to two different files being present.

A, B, and D are incorrect. **A** is incorrect because Jane downloaded the file for a second time, as the question states the whole process was repeated. The odds of the file being corrupted in transit two times in a row are fairly slim, and as such, it is more likely that the two files are different. **B** is incorrect because although this may be a possibility, there's not enough information in the question to be certain that the file contains malware. The only thing that can be deduced with relative certainty is that the file from the third-party website seems to be different from the one on the vendor's website. **D** is incorrect because a collision would only take place if two different files had the same file hash. However, the hashes here are different so this isn't an example of collision.

11. Which of the following would be most suitable to protect files stored on a hard drive?

- A. AES-128
- B. TLS 1.1
- C. RSA-512
- D. 3-DES

A. From the available options, AES (Advanced Encryption Standard) is the most suitable one. AES-128 is stronger than both 3-DES and RSA-256.

B, C, and D are incorrect. **B** is incorrect because TLS 1.1 is used for encrypting data in motion (commonly to establish a secure communication between a client and a server for online transactions), not data at rest (like hard drive files). Even for the purpose of protecting data in motion, version 1.1 is considered obsolete and a later version should be used. The latest version is 1.3 (released in August 2018), which various companies have started adopting in their systems. **C** is incorrect because RSA is commonly used for key exchange. In addition, a key of 512 bits is not considered secure enough to provide adequate encryption (researchers suggest that an RSA-2048-bit key will be secure enough until 2030). **D** is incorrect because AES is considered 3-DES's successor, as it performs very well, whereas 3-DES is very slow, especially when implemented in software. In addition, AES offers 192- and 256-bit keys, which provide stronger encryption than the 168-bit key that 3-DES offers.

12. Which of these ports would be used when a user attempts to send an e-mail?

- A. TCP 110
- B. TCP 25
- C. UDP 53
- D. TCP 443

B. Outgoing e-mail would commonly use the SMTP protocol (using TCP port 25). If its secure version was used (SMTP over TLS/SSL), then TCP port 465 would be used.

A, C, and D are incorrect. **A** is incorrect because TCP port 110 (POP3) would be used for receiving e-mails. If its secure version was used (POP3 over TLS/SSL), then TCP port 995 would be used. **C** is incorrect because UDP port 53 is used by the DNS protocol, which facilitates the resolution of host names to IP addresses. **D** is incorrect because TCP port 443 is used by HTTPS (which is HTTP over TLS/SSL).

13. Which person is responsible for defining a company's data classification scheme?

A. Data owner

B. Data custodian

C. User

D. Management

D. Management is responsible for defining an organization's data classification scheme. Not only that, the management team is also responsible for clarifying the security requirements that data of each classification level requires.

A, B, and C are incorrect. **A** is incorrect because data owners have the responsibility of ensuring that data in each classification level is adequately protected as per what is stated in the related security policy. **B** is incorrect because a data custodian would commonly have administrative tasks relating to the data (i.e., performing regular data backups or granting data access to appropriate individuals). **C** is incorrect because a user would commonly be the one accessing data. However, note that users do have the responsibility of protecting the data they handle.

14. Which of the following would be used for a new certificate registration request?

A. WOT

B. RA

C. CRL

D. OCSP

B. An RA (registration authority) can be used for an organization's certificate registration request. It is tasked with data verification, and after that is completed the request would be passed to the CA (certificate authority). Note that it is not mandatory to use an RA, but large companies do often use one. However, a request can also still just reach the CA directly.

A, C, and D are incorrect. **A** is incorrect because a WOT (web of trust) is a decentralized trust model with the goal of allowing a system to establish authenticity between a public key and its owner. **C** is incorrect because a CRL (certificate revocation list) is a list of revoked certificates that the CA maintains and is used by clients to check a certificate's validity. **D** is incorrect because OCSP (Online Certificate Status Protocol) is another method of checking if a certificate has been revoked. The difference is that instead of a revocation list being provided (like in the case of CRL), the client would just send a request to the OCSP server asking if a certificate of interest is valid or not, and the server would return a verdict.

15. Which of these authentication protocols would you use to support mutual authentication?

A. CHAP

B. PAP

C. MS-CHAPv2

D. MS-CHAPv3

C. MS-CHAPv2 is a CHAP (Challenge-Handshake Authentication Protocol) implementation by Microsoft that supports mutual authentication between client and server. This ensures that the client doesn't provide login credentials to a malicious entity that might be posing as the authentication server.

A, B, and D are incorrect. **A** is incorrect because CHAP only supports verification of the client. **B** is incorrect because PAP uses a two-way handshake to only identify the client, so no mutual authentication is supported. **D** is incorrect because there's no MS-CHAPv3. This option is a distractor.

16. You want to ensure your company's e-mails are encrypted to maintain communication confidentiality. Which of these options would you choose for that purpose?

A. DKIM

B. SPF

C. S/MIME

D. DMARC

C. S/MIME (Secure/Multipurpose Internet Mail Extensions) is a protocol that allows one to encrypt and digitally sign e-mails. This works by using the recipient's public key to encrypt an outgoing message. After that has been received, only the recipient can decrypt it (as that would require the recipient's private key), which ensures the security of that transaction. Furthermore,

as S/MIME also supports digital signatures, a sender can sign the message that is sent so that the recipient knows it originated from a valid entity.

A, B, and D are incorrect. **A** is incorrect because DKIM (Domain Keys Identified Mail) is a method of authenticating e-mail and identifying any attempt to spoof a valid sender. It doesn't support e-mail encryption; hence, it wouldn't be used to secure e-mail communications. **B** is incorrect because SPF (Sender Policy Framework) is used to authenticate e-mails and identify e-mail spoofing, not to encrypt them. Assume that you receive an e-mail originating from your bank (which is called amazing bank), sourced from e-mail account security@amazing-bank.com, asking you to change your password to a new value that is provided in the e-mail. If that e-mail is not actually sourced from "amazing bank" but is just an attempt by an attacker to spoof the source header so you are tempted to change your password to what is provided, that would be identified as forged e-mail. **D** is incorrect because DMARC (Domain-based Message Authentication, Reporting, and Conformance) is a protocol that allows organizations to specify if SPF or DKIM (or neither of them) is used and what action to take if validation fails (i.e., mark the message as spam) by publishing the relevant policy through DNS records. However, this protocol wouldn't be used for encrypting messages.

17. Consider this cipher suite:

TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384

What key exchange algorithm is being used?

- A. TLS
- B. ECDHE_RSA
- C. CAMELIA_256_CBC
- D. SHA384

B. A cipher suite is a collection of algorithms that ensure a network connection is secure. It consists of the communication protocol (SSL or TLS), a key exchange algorithm (accompanied by an authentication algorithm), an encryption algorithm, and an algorithm for message authentication. According to this, the key exchange algorithm is ECDHE, which means that elliptic curve Diffie-Hellman is being used for key exchange (while RSA is being used as an authentication algorithm to sign the server's ephemeral keys). A full description of ECDHE_RSA can be found in RFC 4492.

A, C, and D are incorrect. **A** is incorrect because TLS states the communication protocol being used. **C** is incorrect because CAMELIA_256_CBC is the encryption algorithm being used. Camelia is a symmetric encryption block cipher with possible keys of 128, 192, and 256 bits. Here, a 256-bit key is specified. Also note that CBC indicates the encryption algorithm's mode of operation. Here, Cipher Block Chaining (CBC) is being used. **D** is incorrect because SHA384

indicates the message authentication algorithm being used, which is an SHA algorithm with a 384-bit key.

18. Your company wants to simplify the process of employees remotely connecting to the company network via VPN. The main requirement is to be able to connect using a web browser. Which of these methods would be most appropriate for that goal?

- A. TLS VPN
- B. Transport mode
- C. Tunnel mode
- D. SSH

A. A TLS VPN would be the simplest way of connecting to a VPN via a web browser. Some of the advantages of this approach are the ease of use (an employee can use the company VPN from any system that has a web browser), ease of configuration (it doesn't require anything specific to be configured, since most organizations already allow HTTPS [TCP 443] traffic within their environment) and cost-efficiency (doesn't tend to require a lot of support due to the simple operation).

B, C, and D are incorrect. **B** is incorrect because transport mode is a IPSEC mode in which only the packet payload is encrypted (not the packet header). As such, it wouldn't be used in a TLS VPN; a IPSEC VPN is a different implementation. **C** is incorrect because tunnel mode is also a IPSEC mode of operation in which the whole packet is encrypted and would only be used in a IPSEC VPN, not in a TLS VPN. **D** is incorrect because SSH (Secure Shell) isn't an option commonly used for VPNs. It is often used to encrypt communications between a client and a server (i.e., for administrative purposes) and to encrypt other protocols that are not secure by design (i.e., FTP).

19. Which of these attacks would be most suitable to an attacker who has samples of the plaintext and ciphertext messages?

- A. Rainbow table attack
- B. Dictionary attack
- C. Chosen plaintext
- D. Known plaintext

D. A known plaintext attack would be preferred once an attacker is able to get samples of both the plaintext and ciphertext. During this attack, the attacker would commonly try to identify the original key used for encrypting communications.

A, B, and C are incorrect. **A** is incorrect because a rainbow table attack uses precomputed hashes and is mostly used for offline attacks where the attacker has managed to obtain a copy of a system's hashed password file. **B** is incorrect because a dictionary attack would most commonly be used to try and break a user's password by trying a variety of common words contained in a dictionary. This technique is often used to access system accounts by using automated tools, which can utilize various dictionary files and try multiple combinations until a match has been identified. **C** is incorrect because during a chosen-plaintext attack, an attacker would select random plaintexts that would need to be encrypted so the corresponding ciphertexts can be obtained (but the question states that the attacker has both the plaintext and ciphertext messages already).

20. An attacker is using a text file's spaces and tabs to store information. Which of the following is this an example of?

- A. Encoding
- B. Hashing
- C. Steganography
- D. Encryption

C. Steganography can be defined as a technique allowing information to be hidden within a file. In this scenario, the attacker is attempting to take advantage of the natural properties of a text file (i.e., the spaces and tabs) in order to hide information within the file without drastically changing its format. If that file is intercepted and viewed by any third party during transit, it would be very difficult at first glance to determine there's anything hidden within it. Usually, special-purpose tools are used (i.e., steganalysis suites), which can scan a file and try to identify if there's any information hidden within it.

A, B, and D are incorrect. **A** is incorrect because encoding means that the initial information is represented in a different way after being encoded. Consider base64 encoding; the phrase "This is an SSCP book" becomes "VGhpcyBpcyBhbiBTU0NQIGJvb2s=". As you can see, the encoded output bears no resemblance to the initial phrase. **B** is incorrect because hashing uses a one-way function to generate a fixed-length output and is used for data integrity, not for concealing information. The MD5 output of the previous phrase is "efdc6672bef132dd47a4a7bc0eb8f89d". **D** is incorrect because encryption entails scrambling data in such a way that only specifically authorized individuals (i.e., that possess the decryption key) can decode it. Although this is sometimes confused with encoding, the difference is that when a simple encoding technique is used (i.e., base64), the original information is still retrievable by using the same technique to reverse the process. With encryption, the technique becomes much more complicated, as the ultimate goal is to thwart anyone from uncovering the original information. In addition to encryption being more complicated, even if the algorithm used is known, retrieving the original information without the encryption key would be very difficult. When encoding is used, decoding information when the encoding scheme is known (i.e., base64) is fairly trivial.

21. You have a file containing the sentence “This is the plaintext”. You use a key of “This is my first key” and generate the following ciphertext: “yu9+m0/mY7rF/6dv6SXNVC7D69Q7b1N4”. You encrypt the same file again but with a different key: “This is my second key”. The result is “yu9+m0/mY7rF/6dv6SXNVC7D69Q7b1N4”. What is this an example of?

- A. Key clustering
- B. Collision
- C. Collusion
- D. Work factor

A. Key clustering takes place when the same plaintext is encrypted with two different keys and yields the same ciphertext. In this case, the plaintext is “This is the plaintext”. The two keys are “This is my first key” and “This is my second key”. They both yield the same ciphertext of “yu9+m0/mY7rF/6dv6SXNVC7D69Q7b1N4”.

B, C, and D are incorrect. **B** is incorrect because collision refers to hashing and happens when two different files generate the same file hash. MD5 is known to generate a lot of collisions, which makes it a weak hashing algorithm. **C** is incorrect because collusion refers to the act of one or more individuals or companies conspiring to commit fraud. **D** is incorrect because work factor relates to cryptography and expresses the time and effort it would take to decrypt an encrypted message.

22. You just received an e-mail from your bank account manager that your account has been breached. She kindly asks you to reset your password and emphasizes not to disclose it to anyone who may be posing as a bank employee over the phone. You follow the URL contained in her e-mail, and your browser presents you with the following message: “This connection is untrusted.” What is the most likely reason for that?

- A. Certificate has expired.
- B. You have received a phishing e-mail.
- C. Bank website experiencing technical issues.
- D. A self-signed certificate is being used.

B. The account manager should refrain from placing URL links in the e-mail and should state something like “please log in to your account and ensure that you change your password.” As the question states, once the link was followed, an error about the legitimacy of the secure connection was displayed (which shouldn’t normally happen with your bank’s valid website). Although you can choose to disregard that error and continue to the target website (sometimes attackers will even include details of how to do that in the phishing e-mail, stating known errors

are being experienced currently with a series of steps on how to bypass the browser warnings), but you should refrain from taking such actions with your banking website. If you are not sure, you should report this to your bank and ask for further guidance.

A, C, and D are incorrect. **A** is incorrect because the odds of a bank having its certificate being expired are very slim. This is the primary means of verifying the validity of the bank to the clients and then establishing an encrypted connection for data transaction, so the bank should know better and not allow its certificate to expire. **C** is incorrect because if the bank's website was experiencing technical issues, that wouldn't translate to the connection being untrusted. The message points to a certificate issue or an illegitimate connection. Also, consider that when a bank faces any technical issues, it would usually alert its clients so they know about it (either directly or via an announcement on the website). **D** is incorrect because no legitimate bank would ever use a self-signed certificate for interacting with its clients. Organizations sometimes use self-signed certificates for interorganization activities, but if such activities are publicly facing, it is a poor policy to have self-signed certificates in use.

23. Which of the following would you use to check if a certificate has been revoked, using the least amount of network resources?

- A. CRL
- B. CLR
- C. RA
- D. OCSP

D. From the options provided, only two relate to certificate revocation, which are A and D. OCSP (Online Certificate Status Protocol) is a protocol used to check if a certificate has been revoked, which consumes very few network resources. The client sends a request to the OCSP server containing a certificate's serial number, asking if the certificate of interest is valid or not, and the server comes back with a verdict (i.e., revoked).

A, B, and C are incorrect. **A** is incorrect because when a CRL (certificate revocation list) is used, a full list of revoked certificates that the CA maintains would need to be sent to the client, taking more time and network resources than using OCSP. **B** is incorrect because there's no term known as CLR in use in the context of digital certificates. **C** is incorrect because an RA (registration authority) is used when an organization needs to submit a certificate registration request for a new certificate to be issued.

24. James wants to communicate with Dianne. For that purpose, he encrypts a message with Dianne's public key before sending it to her. Which of the following is he trying to enforce?

- A. Integrity
- B. Authentication

C. Confidentiality

D. Nonrepudiation

C. When James encrypts his message with Dianne's public key, he ensures that the only person who can read the message is Dianne. As such, he's protecting the message from unauthorized disclosure, which means confidentiality is being enforced.

A, B, and D are incorrect. A is incorrect because integrity can't be ensured by encrypting the message with Dianne's public key. Any third party that intercepts this message can still alter it in transit, and Dianne will have no way of verifying its integrity. However, integrity could be maintained if James used HMAC. B is incorrect because authentication can't be ensured, as anyone can send Dianne a message and encrypt it with her public key. That doesn't mean James was the one who sent her the message. If a digital certificate was used, Dianne would know this message came from James. D is incorrect because nonrepudiation can't be ensured, as there's no way to prove that James took the action of sending the message to Dianne. If required, a digital signature could be used in order to provide nonrepudiation.

25. If a group of eight people uses symmetric key cryptography, how many keys are required to ensure that everyone's communications are secured while each pair uses a different encryption key?

A. 8

B. 28

C. 32

D. 64

B. Each pair of people will need a different key to secure all communications and ensure there's no key repetition. Here's a quick formula to find how many keys are required: If k is the number of people in the group, then the total number of keys is: $\text{Total} = [k \times (k - 1)]/2$. Since $k = 8$, that means: $\text{Total} = 8 \times (8 - 1)/2 = 8 \times 7/2 = 28$.

A, C, and D are incorrect. These are all distractors.

CHAPTER 6

Network and Communications Security

This chapter includes questions from the following objectives:

- 6.1 Understand and apply fundamental concepts of networking
- 6.2 Understand network attacks and countermeasures (e.g., DDoS, man-in-the-middle, DNS poisoning)
- 6.3 Manage network access controls
- 6.4 Manage network security
- 6.5 Operate and configure network-based security devices
- 6.6 Operate and configure wireless technologies (e.g., Bluetooth, NFC, Wi-Fi)

This chapter contains a complete set of questions and answers regarding basic and advanced networking and communications (OSI model, network topologies, basic protocols and ports, remote access solutions, and virtual environments); attacks and defenses (attack types and related countermeasures in addition to monitoring and analysis); and network access controls (network access control management and remote access solutions). It also contains questions regarding network security management (LAN security management and firewall concepts), monitoring systems, and network-based detection, as well as wireless security (types of wireless technologies and security protocols).

Q QUESTIONS

1. What is a datagram at the Data Link layer called?

- A. Bit
- B. Packet
- C. Frame
- D. Segment

2. Which of the following is not a part of the TCP/IP model?

- A. Application
- B. Network

C. Transport

D. Link

3. Which topology is the least resilient?

A. Star

B. Tree

C. Mesh

D. Bus

4. If you want to receive information about a domain's mail servers, which of these DNS records would you use?

A. MX

B. PTR

C. A

D. AAAA

5. Which of these commands doesn't use ICMP (by default)?

A. Ping

B. Tracert

C. Pathping

D. Traceroute

6. You have recently been asked to manage the network security infrastructure of your organization. As a first step, you want to implement a mechanism that can be used to obtain performance-related information (CPU and memory status) from your routers and firewalls. Which of the following would you use for that purpose?

A. ICMP

B. SNMP

C. Telnet

D. FTP

7. When an attacker sends a phishing e-mail, which of these fields would be spoofed to make it resemble a legitimate message?

A. From

B. To

C. Source headers

D. MAC address

8. During a SYN flood, an attacker first sends multiple SYN packets, while the target system responds with SYN/ACK. What does the source (attacker) respond with after that?

A. SYN

B. ACK

C. SYN/ACK

D. Nothing

9. During which of these attacks would the attacker try to capture data from a specific system?

A. Salami attack

B. Spoofing

C. Active MiTM

D. Passive MiTM

10. You receive a text message on your mobile phone stating your new phone bill has been issued, but you don't use the mobile provider that was mentioned in the message. What type of attack is this?

A. Vishing

B. Whaling

C. Smishing

D. Spear phishing

11. Which of these statements about the benefits of VLANs is inaccurate?

- A. Increased security
- B. Excellent physical segmentation
- C. Enhanced performance
- D. No additional equipment required for configuration

12. Which of these attacks uses UDP packets and requires spoofing the source IP address so it becomes the address of the victim?

- A. Fraggle
- B. Smurf
- C. LAND
- D. Ping of death

13. You are a network administrator and want to ensure connectivity between two of your internal subnets so that traffic can flow properly. Which of these devices would you most likely use?

- A. Hub
- B. Repeater
- C. Switch
- D. Router

14. A company wants to protect against employees accessing suspicious URLs. Which of the following would be the most suitable solution?

- A. WAF
- B. AV
- C. Web proxy
- D. E-mail gateway

15. You are a security architect designing a security solution for your network. You want to use a combination of firewalls from different vendors to defend the internal network and separate it

from the rest of the infrastructure. Based on this description, which of the following are you trying to enforce?

- A. Defense in depth
- B. Defense diversity
- C. Multivendor approach
- D. Next-generation firewalls

16. Which of the following wouldn't be possible when using a stateful inspection firewall?

- A. Allow access to 10.10.10.0/24.
- B. Allow access to 192.168.2.5 only over port 80.
- C. Deny traffic from a host initialing a connection with an ACK packet.
- D. Deny traffic relating to video.

17. Consider a firewall that has the following ruleset (note that rules have been simplified to make it easier for the reader):

```
deny icmp any any
permit tcp any 10.2.2.2 eq 22
deny tcp any 10.2.2.2 eq www
```

Which of these statements is inaccurate?

- A. Telnet is allowed to 10.2.2.2.
- B. ICMP is not allowed.
- C. Web traffic is not allowed to 10.2.2.2.
- D. FTP is not allowed to 10.2.2.2.

18. Which of the following would be performed by NAC?

- A. Authenticate a remote client.
- B. Create audit logs about the actions a machine performed while on the network.
- C. Check a client machine against an organizational health policy.

D. Allow a client machine to issue specific commands.

19. Which of the following would you use to adequately secure the wireless network of a small office with ten employees, without any excessive administrative burden?

A. WEP (with AES)

B. WPA2 (with AES)

C. WEP-Enterprise

D. WPA2-Enterprise

20. You enter the subway and want to use your mobile phone to touch the reader and pay for your fare. Which protocol would you most likely use for that transaction?

A. 4G

B. WiMAX

C. RFID

D. NFC

21. Which of the following is the least effective method of securing a wireless network?

A. Hiding the SSID

B. Using MAC address filtering

C. Using WEP for security

D. Using NFC

22. Which of these attacks would allow an attacker to intercept communications?

A. Bluesnarfing

B. Bluejacking

C. Bluebugging

D. Bluetagging

23. Which of these methods for mobile device protection wouldn't be performed before a device is lost or stolen?

- A. MDM
- B. Full device encryption
- C. Enable PIN protection
- D. Remote deletion

24. You have a network composed of a hundred endpoints, ten firewalls, five routers, five IPS sensors, and three proxies across two geographical regions. Data from all of those devices is being forwarded to a SIEM platform. What would be the most crucial element for any type of security investigation you perform?

- A. Log integrity
- B. Verbose logging across all devices
- C. Appropriate alert thresholds
- D. Time synchronization

25. Which of these media types is most susceptible to electromagnetic interference?

- A. Shielded twisted pair
- B. UTP
- C. Optical fiber
- D. STP cat 6

26. In a full mesh topology consisting of ten devices, how many connections would be required?

- A. 10
- B. 20
- C. 45
- D. 90

27. Which of the following is a valid IPv6 address representation?

- A. 0:0:0:0:ffff:a0a:a0a

B. 0:0:0:0:gggg:a0a:a0a

C. 0:0:0:0:ffff:a0a

D. ::ffff:a0a::

28. As a network administrator of a small law firm you get a call from John stating that he's unable to access the law firm's website, <http://www.which-is-best-law-firm-ever.com>, hosted at 10.10.10.10. (Please note that the scenario uses a private IP address but in a real-world situation this address would be a public one.) Hannah, one of John's colleagues sitting next to him, is able to access the web page just fine. You ask John to try accessing <http://10.10.10.10>, and the webpage loads fine, and an attempt to ping 10.10.10.10 is also successful. What is most likely the problem?

A. DHCP server issue.

B. Website is temporarily unavailable.

C. The web server is blocking John's IP address.

D. DNS server issue.

29. For management purposes, you need to ensure that administrators only use secure protocols to log-in to devices. Which of the following would be most appropriate for that?

A. SFTP

B. HTTPS

C. SSH

D. Telnet

30. You use your browser to navigate to a web page located at internal IP address 192.168.1.2, which is a web server listening on TCP port 8080. Your machine has an IP address of 10.10.10.16 and a source port of TCP 500. During the first stage of the communication (when your machine is sending the initial SYN packet to the web server), which of this information given is most probably wrong?

A. Source IP address of 10.10.10.16

B. Source TCP port 500

C. Destination IP address of 192.168.1.2

D. Destination TCP port 8080

31. Which of these methods would you choose to protect a Windows 2003 physical server that needs to directly communicate with critical infrastructure?

- A. Hardening
- B. VM isolation
- C. Patching
- D. User awareness

32. You have enhanced your network security by adding MAC address filtering. Which of these attacks can be used against that in order to gain unauthorized network access?

- A. Sniffing
- B. MiTM
- C. Spoofing
- D. Session hijacking

33. Which of the following would you use to prevent an SQL injection attack?

- A. NIDS
- B. Stateful inspection firewall
- C. Disallow stored procedures
- D. Input validation

34. Which of the following devices uses MAC addresses?

- A. Hub
- B. Switch
- C. Repeater
- D. Router

35. In which of the following would you configure ERSPAN so you mirror traffic to your NIDS?

- A. Riverbed switch.

- B.** Cisco Catalyst switch.
- C.** MikroTik router.
- D.** None of the above support ERSPAN.

36. As a network administrator you are asked to select the most cost-effective security device that can filter based on the following criteria:

- IP addresses/subnets
- Specific ports

Which of these devices would be more suitable in this case?

- A.** Router
- B.** Stateful inspection firewall
- C.** Packet-filtering firewall
- D.** IDS

37. Which of these types of attackers would probably target a military installation to obtain new research and development (R&D) plans?

- A.** APT
- B.** Script kiddie
- C.** Hacktivist
- D.** White hat

38. During a DDoS attack, who would be responsible for issuing commands?

- A.** Botnet
- B.** Zombie
- C.** Victim
- D.** Bot herder

39. Which of the following needs to take place so your machine can capture all the wireless transmissions within your device's range?

- A. An external USB-wireless adapter is required.
- B. Wireless NIC set to nonpromiscuous mode.
- C. Wireless NIC set to promiscuous mode.
- D. Wireless network needs to be open (i.e., not have any security protocol enabled).

40. Which of the following is a valid WPA-PSK preshared key?

- A. KsIA349
- B. AK09EYei
- C. £jh85418
- D. RΣ98*6Te

QUICK ANSWER KEY

1. C

2. B

3. D

4. A

5. D

6. B

7. A

8. D

9. C

10. C

11. B

12. A

13. D

14. C

15. B

16. D

17. A

18. C

19. B

20. D

21. A

22. C

23. D

24. D

25. B

26. C

27. A

28. D

29. C

30. B

31. A

32. C

33. D

34. B

35. B

36. C

37. A

38. D

39. C

40. B

ANSWERS A

1. What is a datagram at the Data Link layer called?

A. Bit

B. Packet

C. Frame

D. Segment

C. Each OSI layer adds more information to the original data. When data initially arrives at a machine, it comes in the form of binary 0 and 1 (reaching the Physical layer, where it is known as “bit”), but once that is passed to the Data Link layer, it is then called a “frame.” At that layer MAC (media access control) addresses are used, which uniquely identify a NIC (network interface card).

A, B, and D are incorrect. **A** is incorrect because a datagram is called a bit at the Physical layer, where data is converted from binary 0 and 1 to streams of bits. **B** is incorrect because a datagram is called a packet at the Network layer. In this layer, IP addresses are used and routing takes place to ensure packets are appropriately routed to their final destination. **D** is incorrect because a datagram is known as a segment at the Transport layer, which is where error detection and establishment of reliable communications take place.

2. Which of the following is not a part of the TCP/IP model?

A. Application

B. Network

C. Transport

D. Link

B. The Network layer is a term used in the OSI model, not TCP/IP. TCP/IP uses Internet layer to refer to the respective Network layer of the OSI model.

A, C, and D are incorrect. **A** is incorrect because TCP/IP uses the Application layer to map three of the OSI model layers (Application, Presentation, and Session). **C** is incorrect because OSI's Transport layer is named Transport in the TCP/IP model (note that it is also known as Host-to-Host layer). **D** is incorrect because TCP/IP uses the Link layer (also known as Network Access layer) to map OSI's Physical and Data Link layers.

3. Which topology is the least resilient?

- A.** Star
- B.** Tree
- C.** Mesh
- D.** Bus

D. When using a bus topology, any issue with one of the connected devices will bring the whole network to a halt. A bus topology uses T connectors for connectivity. That means that one end of a computer is connected to the bus and the other is connected to the next computer. If one of the computers or T connectors experiences an issue, the whole network stops working. In addition, if a lot of machines are connected to that network, identifying the issue may require vast amounts of time. This means the network will not be usable until the issue is resolved.

A, B, and C are incorrect. **A** is incorrect because a star topology commonly uses a switch to allow the various machines to connect, which means that any issue with one machine doesn't affect the network's overall operation. **B** is incorrect because a tree topology uses a combination of a bus and a star, where multiple star networks connect over a common bus. If the bus experiences an issue, all the star networks are disconnected until the issue is resolved. **C** is incorrect because in a mesh topology, all machines are individually connected with each other, meaning that any issues with one connection don't affect the operation of the network.

4. If you want to receive information about a domain's mail servers, which of these DNS records would you use?

- A.** MX
- B.** PTR
- C.** A
- D.** AAAA

A. "MX" (mail exchange) is the type of record responsible for providing information about mail servers. A security analyst will often use a command like "nslookup," in which a specific

parameter for MX records can be specified (which is set $q = \text{MX}$) to only return mail server-related information.

B, C, and D are incorrect. **B** is incorrect because “PTR” (pointer) records are used to resolve an IP address to a corresponding host name, which is what a reverse lookup would use. **C** is incorrect because an “A” (host) record would normally be used for a DNS request trying to resolve a host name to an IPv4 address. **D** is incorrect because an “AAAA” (host) record would normally be used for a DNS request trying to resolve a host name to an IPv6 address.

5. Which of these commands doesn’t use ICMP (by default)?

- A. Ping
- B. Tracert
- C. Pathping
- D. Traceroute

D. Traceroute is a Linux/UNIX command that uses UDP by default. It is used for checking the route of the traffic towards a destination.

A, B, and C are incorrect. **A** is incorrect because ping uses ICMP packets and is used to check if there’s connectivity between two hosts. **B** is incorrect because tracert is used for checking the route of the traffic towards a destination, but the difference is that tracert uses ICMP to accomplish this. **C** is incorrect because pathping also uses ICMP to track the intermediate hops between a source and destination and provides network latency information.

6. You have recently been asked to manage the network security infrastructure of your organization. As a first step, you want to implement a mechanism that can be used to obtain performance-related information (CPU and memory status) from your routers and firewalls. Which of the following would you use for that purpose?

- A. ICMP
- B. SNMP
- C. Telnet
- D. FTP

B. SNMP (Simple Network Management Protocol) is used extensively for network device management. It can be used to communicate with a monitored device for a variety of tasks, like to query about its current state (i.e., current CPU utilization or free disk and memory levels).

A, C, and D are incorrect. **A** is incorrect because ICMP is mostly used for diagnostic purposes by commands like ping and tracert. **C** is incorrect because Telnet is commonly used for connecting remotely to a system using basic commands via command line. **D** is incorrect because FTP (File Transfer Protocol) is used for file transfers.

7. When an attacker sends a phishing e-mail, which of these fields would be spoofed to make it resemble a legitimate message?

- A.** From
- B.** To
- C.** Source headers
- D.** MAC address

A. In order for any attacker to be able to resemble a legitimate message, the “From” field would need to be spoofed to include a legitimate sender. Think of an example in which you receive an e-mail from someone claiming to be your account manager at Magic Bank. The “From” field might be something like “account-management@magic-bank.com,” which would entice you to open it and follow any instructions (like navigating to a URL) it contains, thinking it’s from your legitimate account manager.

B, C, and D are incorrect. **B** is incorrect because there’s no point in spoofing the “To” field, as the recipient would be the true victim of the phishing attack. **C** is incorrect because a message’s headers will have information about what “path” the message used to reach you. In this example, the actual “Received” header field would not state “account-management@magic-bank.com,” but rather the actual e-mail account that was used to send this message. **D** is incorrect because the MAC address isn’t displayed to the victim so it wouldn’t make any difference if it’s real or not.

8. During a SYN flood, an attacker first sends multiple SYN packets, while the target system responds with SYN/ACK. What does the source (attacker) respond with after that?

- A.** SYN
- B.** ACK
- C.** SYN/ACK
- D.** Nothing

D. A SYN flood aims to overwhelm the target system with half-open connections. In order for that to happen, after a source sends a SYN packet and a destination responds with SYN/ACK, no other packet is sent. However, the target system will still have the connection information in

its connection table, thus occupying space, and will wait for a completion of that connection with a final ACK from the source (which never comes).

A, B, and C are incorrect. **A** is incorrect because a SYN packet would have been sent when the connection was initiated. **B** is incorrect because an ACK packet would normally be sent from the source to complete the three-way handshake in the case that this wasn't a SYN flood but a legitimate connection request. **C** is incorrect because a SYN/ACK would only be sent from the target after receiving a request for a connection (SYN packet) from the source.

9. During which of these attacks would the attacker try to capture data from a specific system?

A. Salami attack

B. Spoofing

C. Active MiTM

D. Passive MiTM

C. During an active MiTM (man in the middle) attack, an attacker would use a sniffing tool (i.e., Wireshark) to capture traffic of interest from one specific system (or systems).

A, B, and D are incorrect. **A** is incorrect because a salami attack aims at taking advantage of small actions that may lead to a large outcome when combined—for example, an attacker using DNS to exfiltrate small chunks of data and after several days or weeks of effort managing to extract a rather large file with proprietary organizational data. **B** is incorrect because during a spoofing attack, the attacker would attempt to spoof the IP or MAC address he or she has been assigned to make it look like his or her machine is a legitimate network device. **D** is incorrect because during a passive MiTM attack, the attacker would try to capture traffic from any system (contrary to an active MiTM, where specific systems would be targeted).

10. You receive a text message on your mobile phone stating your new phone bill has been issued, but you don't use the mobile provider that was mentioned in the message. What type of attack is this?

A. Vishing

B. Whaling

C. Smishing

D. Spear phishing

C. Smishing is a type of attack that uses text messages as an attack vector. Typically, a text message would be sent to the victim about a utility bill or online purchase confirmation or even a

contest with a prize. Regardless of the context, the message will usually contain a URL, which the victim is requested to navigate to for an action to be completed, which, of course, is used to have the victim enter personal details or run a possible exploit.

A, B, and D are incorrect. **A** is incorrect because vishing doesn't use text messages, but rather the actual phone to place a call to the victim (or notify the victim to call the attacker) and convince them to provide personal details or perform a specific action. **B** is incorrect because whaling is a type of phishing attack where an executive is targeted and a phishing e-mail is sent to that victim. However, no text messages are used to deliver the attack. **D** is incorrect because a spear phishing attack uses e-mail as an attack vector. This message is sent to various individuals (commonly working in the same company or team) in the hope that a URL that is included in the message will be followed and credentials provided or a specific file (attached to the message) run for an action to be invoked.

11. Which of these statements about the benefits of VLANs is inaccurate?

- A.** Increased security
- B.** Excellent physical segmentation
- C.** Enhanced performance
- D.** No additional equipment required for configuration

B. VLANs (virtual local area networks) work by virtually segmenting the network and not providing physical segmentation. The way a VLAN operates is by allowing a network administrator to configure specific devices (connected to one or more switches) as parts of a VLAN. The devices can be in the same room or in different floors of a building and still be able to communicate with each other as if they were physically connected.

A, C, and D are incorrect. **A** is incorrect because VLANs offer increased security, as they separate broadcast domains (similar to a router). That means no other devices can intercept the traffic of a specific VLAN, allowing for greater security. **C** is incorrect because VLANs offer enhanced performance because less broadcast traffic is required, since broadcast traffic will now be sent to the related VLAN only. **D** is incorrect because if existing switches support VLAN configuration, no additional equipment is required to configure them.

12. Which of these attacks uses UDP packets and requires spoofing the source IP address so it becomes the address of the victim?

- A.** Fraggle
- B.** Smurf
- C.** LAND

D. Ping of death

A. A fraggle attack works by using UDP packets while spoofing the source IP address so it becomes the address of the victim. It targets UDP 7 (echo protocol, which will make the operation work similarly with ping) or UDP 19 (character generation protocol, which will make the target system respond with a random character when it receives traffic over that port).

B, C, and D are incorrect. **B** is incorrect because a smurf attack uses ICMP packets while sending traffic to various network hosts with a spoofed IP address to match the victim. That way, any ICMP reply packets will go to the victim, thus overwhelming it with traffic. **C** is incorrect because a LAND attack makes a target machine continuously send traffic to itself, since the packets have both the source and destination IP address of the victim. **D** is incorrect because a “ping of death” attack changes the size of a ping packet to 64KB or more, which has been known to cause crashes, especially to older machines.

13. You are a network administrator and want to ensure connectivity between two of your internal subnets so that traffic can flow properly. Which of these devices would you most likely use?

A. Hub

B. Repeater

C. Switch

D. Router

D. A router can be used so that traffic between two different subnets can be properly routed to any machine. Routers are used for forwarding packets within different networks by use of internal routing tables (either using static routes or dynamic routing protocols).

A, B, and C are incorrect. **A** is incorrect because a hub is unable to route traffic to different networks and is commonly used to connect machines over a common network. **B** is incorrect because a repeater is used to regenerate the signal in order to allow it to reach longer distances, so it wouldn't be suitable in this scenario. **C** is incorrect because a switch is used to connect devices over a local area network and wouldn't be appropriate for routing traffic between different networks. Please note that there are also layer 3 switches that have a lot of functionality and combine functions of a switch and a router. Those types of devices would be able to route traffic accordingly, but the question doesn't mention layer 3 switches. In addition, the question asks what you would “most likely use,” and commonly a router would be used to route traffic between different networks.

14. A company wants to protect against employees accessing suspicious URLs. Which of the following would be the most suitable solution?

- A. WAF
- B. AV
- C. Web proxy
- D. E-mail gateway

C. A web proxy would be the most suitable solution, as it has the ability to filter all web requests (i.e., the URLs that the employees will navigate to), and if any of them are found to be suspicious, it can block the content from specific websites. Furthermore, browsing categories can be specified for the staff, and if an individual is not authorized to browse a specific category, that won't be allowed, regardless if the destination page is malicious or not.

A, B, and D are incorrect. A is incorrect because a WAF (web application firewall) would be suitable to protect the business against any application attack, like an attacker trying to use cross-site scripting or an SQL injection attack, but that wouldn't be used for filtering user web traffic. B is incorrect because an AV (antivirus) tool would be useful to identify any malicious activity at an endpoint level (e.g., suspicious file execution or malicious processes running). Although a lot of modern antivirus applications do offer web filtering, that's a different component, and the question doesn't specify if the AV has web filtering capabilities in place, thus not making this the most suitable answer. D is incorrect because an e-mail gateway would be used to filter any e-mails for suspicious content (including malicious URLs), but that can only filter e-mail content. It won't help to filter any user's web traffic since that is not associated with e-mail (i.e., user opening a browser and navigating to a malicious website).

15. You are a security architect designing a security solution for your network. You want to use a combination of firewalls from different vendors to defend the internal network and separate it from the rest of the infrastructure. Based on this description, which of the following are you trying to enforce?

- A. Defense in depth
- B. Defense diversity
- C. Multivendor approach
- D. Next-generation firewalls

B. The term defense diversity describes the use of different vendors in order to provide resiliency in case of any issues with one (or some) of them. That means that if for some reason one of these firewalls fails to provide adequate protection or it is affected by a possible zero-day exploit, the other one(s) can still protect the corporate environment.

A, C, and D are incorrect. **A** is incorrect because defense in depth describes the use of a layered defense in which several defenses are in place that the attacker would need to penetrate to access any data. For example, there might be a border router at the perimeter (with properly configured access lists), followed by a firewall, an IPS, and other devices, all of which would need to be defeated if an attacker wanted to obtain access to sensitive data. **C** is incorrect because there's no such term as multivendor approach used in security. The correct terms are defense in depth and defense diversity. This option is a distractor. **D** is incorrect because all of the firewalls could very well be next-generation devices but that doesn't mean you are trying to enforce that. You are still trying to enforce defense diversity, and in order to do that you are using next-generation firewalls.

16. Which of the following wouldn't be possible when using a stateful inspection firewall?

- A.** Allow access to 10.10.10.0/24.
- B.** Allow access to 192.168.2.5 only over port 80.
- C.** Deny traffic from a host initialing a connection with an ACK packet.
- D.** Deny traffic relating to video.

D. A stateful inspection firewall can work as a traditional packet-filtering firewall, plus it offers the added benefit of being able to track the state of connections via a stateful inspection table. That means traffic will only be allowed if it relates to an already established session or a valid new connection (assuming the ruleset allows it). However, it's not able to perform protocol inspection, which is what would be required in order to inspect traffic and classify it as video while afterwards taking an appropriate action regarding it (i.e., permit or deny).

A, B, and C are incorrect. **A** is incorrect because it is very easy to allow access to a specific subnet using just features from a traditional packet-filtering firewall. **B** is incorrect because allowing access to a specific host over a specific port is again something that a stateful inspection firewall can perform. **C** is incorrect because a stateful inspection firewall can monitor traffic and drop any malformed packets (i.e., if a client sends a request to start a connection with something other than a SYN packet, the firewall can deny that traffic).

17. Consider a firewall that has the following ruleset (note that rules have been simplified to make it easier for the reader):

```
deny icmp any any
permit tcp any 10.2.2.2 eq 22
deny tcp any 10.2.2.2 eq www
```

Which of these statements is inaccurate?

- A.** Telnet is allowed to 10.2.2.2.

B. ICMP is not allowed.

C. Web traffic is not allowed to 10.2.2.2.

D. FTP is not allowed to 10.2.2.2.

A. Based on the ruleset, Telnet is not allowed to 10.2.2.2. There's no rule relating to Telnet (only one for SSH: TCP port 22). This means that once the firewall applies the current rules to find a match, it will reach the end of the ruleset, where there is an implicit deny rule. That means that whatever traffic was not explicitly allowed earlier, will now be denied.

B, C, and D are incorrect. **B** is incorrect because ICMP is not allowed as per the first rule. **C** is incorrect because web traffic is not allowed to 10.2.2.2 as per the third rule (denying traffic to 10.2.2.2 over port www, which is web traffic). **D** is incorrect because FTP is not allowed to 10.2.2.2. Because there's no rule allowing FTP, after the existing rules are examined for a match, the implicit deny rule will be applied, which will deny any FTP traffic (since it wasn't explicitly allowed earlier).

18. Which of the following would be performed by NAC?

A. Authenticate a remote client.

B. Create audit logs about the actions a machine performed while on the network.

C. Check a client machine against an organizational health policy.

D. Allow a client machine to issue specific commands.

C. A NAC (Network Access Control) device is primarily used to check if a client that is requesting to connect to the corporate network fulfills the minimum requirements to be allowed to join the network. If that is the case, a NAC device will allow the client to connect. If not, the client will usually be restricted to a quarantined network until it achieves the minimum company security requirements.

A, B, and D are incorrect. **A** is incorrect because the authentication of a remote client would happen by an authentication server (i.e., RADIUS or TACACS+). **B** is incorrect because a NAC device wouldn't be able to perform auditing of a network client. That is a task that needs to be performed by the authentication server (i.e., a server supporting AAA, like RADIUS, TACACS+, or Diameter). **D** is incorrect because reviewing a client's permissions to issue specific commands takes place during the authorization process (which typically happens when a client logs on using an AAA server).

19. Which of the following would you use to adequately secure the wireless network of a small office with ten employees, without any excessive administrative burden?

- A. WEP (with AES)
- B. WPA2 (with AES)
- C. WEP-Enterprise
- D. WPA2-Enterprise

B. The best solution to provide adequate security without added administrative requirements would be WPA2 (with AES). It allows the use of AES encryption, while it ensures that the old and insecure WEP protocol is not used but without going to the extent of using WPA2-Enterprise, which requires an authentication server to be implemented.

A, C, and D are incorrect. **A** is incorrect because WEP is an insecure wireless protocol that doesn't support AES but actually uses RC-4. **C** is incorrect, as there's no enterprise mode in WEP. This option is a distractor. **D** is incorrect because WPA2-Enterprise could be used as the most secure option from the ones provided (as it supports the use of an authentication server, like RADIUS, which would be used to authenticate clients to the network). However, that comes with an added administrative cost of configuring and maintaining an authentication server, which would possibly be too much overhead for a small office of ten people.

20. You enter the subway and want to use your mobile phone to touch the reader and pay for your fare. Which protocol would you most likely use for that transaction?

- A. 4G
- B. WiMAX
- C. RFID
- D. NFC

D. NFC (near field communication) is a protocol that allows various devices (including smartphones that support it) to make contactless payments. This is the equivalent of using a contactless debit or credit card. Note that this protocol also allows for devices in close proximity (usually around 3 to 4 inches) to connect and exchange data.

A, B, and C are incorrect. **A** is incorrect because 4G is a protocol that allows mobile devices (like phones and tablets) to perform tasks, like use mobile web browsing, make video calls, and access social and streaming media. **B** is incorrect because WiMAX (Worldwide Interoperability for Microwave Access) is a wireless protocol (also known as Wireless MAN) that allows a wireless network to span large areas (with an approximate radius of 30 miles). **C** is incorrect because RFID (radio frequency identification) is commonly used for tagging and inventory control. You can tag your suitcase with an RFID chip to ensure you are able to locate it in the event that it gets lost or stolen while you travel.

21. Which of the following is the least effective method of securing a wireless network?

- A.** Hiding the SSID
- B.** Using MAC address filtering
- C.** Using WEP for security
- D.** Using NFC

A. Hiding the SSID (service set identifier) means disabling the network's SSID broadcasting feature, with the goal of hiding the network's existence. However, this doesn't really offer actual protection against an attacker with even elementary skills. That's because the SSID name is still part of the control frames being sent, so if anyone attempts to use a packet analyzer (like Wireshark) to capture the traffic, they can identify the network's SSID.

B, C, and D are incorrect. **B** is incorrect because using MAC address filtering is better than hiding the network's SSID. In order for an attacker to gain access to the network, the first thing to do would be to identify a MAC address that is allowed on the network and then take the additional step of spoofing his or her MAC address to pose as one of the legitimate machines. **C** is incorrect because using WEP doesn't offer optimal security (as it is easily defeated) but at least the network doesn't remain open so anyone can connect to it; in addition, in order for someone to crack WEP, basic skills are still required. **D** is incorrect because NFC (near field communication) is a protocol that allows various devices (including smartphones that support it) to make contactless payments or perform data transactions. However, it's unable to secure a wireless network, thus making it an invalid option.

22. Which of these attacks would allow an attacker to intercept communications?

- A.** Bluesnarfing
- B.** Bluejacking
- C.** Bluebugging
- D.** Bluetagging

C. Bluebugging would be used by an attacker in order to control a Bluetooth-enabled device (often used to target phones). Common tasks that can be performed are remotely controlling a device, sending/reading text messages, placing/intercepting calls, and setting call forwarding.

A, B, and D are incorrect. **A** is incorrect because bluesnarfing takes advantage of the fact that some Bluetooth-enabled devices are set to be discovered by default, and if such a device is susceptible to bluesnarfing, the attacker may be able to connect to that device by pairing with it and obtain information like the phone book or calendar files but not intercept communications. **B**

is incorrect because bluejacking (being the most harmless Bluetooth attack in comparison to bluesnarfing and bluebugging) allows an attacker to use Bluetooth to send short messages to any neighboring devices susceptible to bluejacking. However, other than sending a short message, this attack doesn't really offer the attacker anything else (certainly not the ability to intercept communications). **D** is incorrect because there's no associated Bluetooth attack known as bluetagging. This term is a distractor.

23. Which of these methods for mobile device protection wouldn't be performed before a device is lost or stolen?

- A.** MDM
- B.** Full device encryption
- C.** Enable PIN protection
- D.** Remote deletion

D. From the provided options, remote deletion is the only method that would be performed after the device is lost or stolen. Obviously, you can remotely delete your device at any moment in time, but clearly that would only make sense after that device has been verified as lost or stolen.

A, B, and C are incorrect. **A** is incorrect because an MDM (Mobile Device Management) tool would be used to enroll the device under its supervision once the device is initially handed over to the user (typically an employee of a company implementing the MDM solution). **B** is incorrect because the device's full encryption option (if such an option is supported by the device) would need to be configured before that is lost or stolen. Once the configuration has been made, the device starts being fully encrypted. **C** is incorrect because a PIN (personal identification number) would need to be configured to protect the device before that is lost or stolen. If a company policy is in place, it will commonly force the user to set a four- to six-digit PIN to ensure that the phone never remains unlocked so any sensitive data remains protected at all times.

24. You have a network composed of a hundred endpoints, ten firewalls, five routers, five IPS sensors, and three proxies across two geographical regions. Data from all of those devices is being forwarded to a SIEM platform. What would be the most crucial element for any type of security investigation you perform?

- A.** Log integrity
- B.** Verbose logging across all devices
- C.** Appropriate alert thresholds
- D.** Time synchronization

D. For any effective investigation to take place, appropriate time synchronization needs to be present for all the devices across the network. Even with a handful of devices, like five to six, analysis can become very complicated when those devices are not in a common time zone (i.e., UTC). Imagine you receive an alert from the SIEM with a time stamp of 14:00 UTC-8 (which might be the time zone of the tool or the location where it is installed) and you log in to individual devices to investigate further. Some of them are in UTC+4 and others in UTC-2. How are you going to be able to track a specific attacker's activities across all those time zones?

A, B, and C are incorrect. **A** is incorrect because log integrity is crucial, but in this scenario the logs are being forwarded to a SIEM tool, which stores copies of them, so a good level of integrity is present. **B** is incorrect because verbose logging across various devices might be too intense (depending on the amount of traffic and device type) and cause issues with them, while it also consumes a lot of network bandwidth. SIEM platforms use EPS (events per second), and as that increases, the cost of the tool grows exponentially—not to mention that a more robust architecture will be required in order to cope with the excessive traffic being forwarded to the SIEM tool. However, if there's no time synchronization, even with an abundance of data, it will be very difficult to identify crucial artifacts of any attack. **C** is incorrect because for any SIEM to work effectively, appropriate alert thresholds need to be configured; however, those can be gradually defined as the tool is being used. Lack of time synchronization will pose a greater challenge when investigating malicious activities across logs of various devices in different regions (and various time zones).

25. Which of these media types is most susceptible to electromagnetic interference?

A. Shielded twisted pair

B. UTP

C. Optical fiber

D. STP cat 6

B. UTP (unshielded twisted pair) is a cable type that is very commonly used in computer networks (especially for short to medium distances), CCTV circuits, and telephone systems, as it's very cost effective. However, its greatest disadvantage is that it is susceptible to both electromagnetic and radio frequency interference.

A, C, and D are incorrect. **A** is incorrect because shielded twisted pair (STP) uses additional shielding to reduce the effect of electromagnetic interference. Although it is more expensive as a result, it is often used for outdoors implementations and signal transmissions over large distances. **C** is incorrect because optical fiber cabling uses light to transmit data and as such is not susceptible to electromagnetic interference. It's often used in high-speed networks and can allow high data transmission speeds over large distances, while also offering enhanced security in comparison to traditional media, like copper cables. **D** is incorrect because STP Cat 6 (also

known as category 6) is a type of STP cable that is less prone to electromagnetic interference than common UTP cables.

26. In a full mesh topology consisting of ten devices, how many connections would be required?

- A. 10
- B. 20
- C. 45
- D. 90

C. In a full mesh topology, each device uses multiple connections and is individually connected to every other device, thus offering enhanced availability. In order to identify the appropriate number of connections, you can use one of two options. Either create a diagram with ten nodes and interconnect those while counting how many lines are required for the connections or use a mathematical formula to obtain the result faster. Assuming h is the number of the devices in the model, the formula is $h \times (h - 1)/2$. Since $h = 10$, that makes the total: $10 \times (10 - 1)/2 = 45$ connections.

A, B, and D are incorrect. These are distractors.

27. Which of the following is a valid IPv6 address representation?

- A. 0:0:0:0:0:ffff:a0a:a0a
- B. 0:0:0:0:0:gggg:a0a:a0a
- C. 0:0:0:0:0:ffff:a0a
- D. ::ffff:a0a::

A. A valid IPv6 address representation uses only hexadecimal characters in eight groups of four characters each. For display purposes, leading (not trailing) zeros can be omitted, and a group of consecutive zeros can be represented by a “::” character only once. The only answer that fulfills all these requirements is option A.

B, C, and D are incorrect. **B** is incorrect, as it includes an invalid segment (“gggg”) instead of only using appropriate hexadecimal numbers. **C** is incorrect because it uses seven groups of characters, not eight (without any legitimate display compression taking place). **D** is incorrect because the display compression is wrong, as it includes two sets of “::” when only one is allowed.

28. As a network administrator of a small law firm you get a call from John stating that he's unable to access the law firm's website, <http://www.which-is-best-law-firm-ever.com>, hosted at 10.10.10.10. (Please note that the scenario uses a private IP address but in a real-world situation this address would be a public one.) Hannah, one of John's colleagues sitting next to him, is able to access the web page just fine. You ask John to try accessing <http://10.10.10.10>, and the webpage loads fine, and an attempt to ping 10.10.10.10 is also successful. What is most likely the problem?

- A. DHCP server issue.
- B. Website is temporarily unavailable.
- C. The web server is blocking John's IP address.
- D. DNS server issue.

D. Let's summarize the information provided in the question. John is unable to access the law firm's website, but Hannah is able to access that just fine, likely pointing to the fact there's no website issue. At the same time, John is able to access the website using its IP address and manages to ping that IP successfully, indicating communication between John's machine and the web server is also working properly. The difference between accessing the page on a browser and using the web server's IP address in the browser is that in the first case John's machine will use DNS to do a lookup and identify the web server that the web page is hosted at. If the DNS server is not working or the DNS server settings on John's machine are not correct, then this lookup will fail, thus not allowing the web page to be accessed.

A, B, and C are incorrect. **A** is incorrect because there's no indication of a DHCP server issue. Although the question doesn't specify if a DHCP server is in use, the fact that John is able to ping the web server successfully means that John's machine has a properly configured IP address, so a DHCP issue doesn't seem to be present. **B** is incorrect because Hannah managed to access the website, so there doesn't seem to be any issues with its availability. **C** is incorrect because if something was blocking John's IP address, he wouldn't be able to ping the server successfully. Of course, it's technically possible for a device (i.e., a firewall) to only block specific traffic from John's machine (i.e., web traffic), but it is highly unlikely that the web server would block web traffic from John's machine, as there is no information in the question to support that.

29. For management purposes, you need to ensure that administrators only use secure protocols to log-in to devices. Which of the following would be most appropriate for that?

- A. SFTP
- B. HTTPS
- C. SSH

D. Telnet

C. SSH (Secure Shell) is most commonly used for remote administration of devices, as it supports encrypted communication and is capable of protecting any related data transaction (i.e., password exchange or command output).

A, B, and D are incorrect. **A** is incorrect because SFTP does offer encryption, but it only allows for file transfers to take place, so it's not very useful for device management, which would most likely require various commands to be issued. **B** is incorrect because HTTPS is commonly used to secure a browser communication between a client and server, like in the case of a user logging in to his or her bank account. Although some devices support management over HTTPS, not all of them do, and administrators tend not to use HTTPS to access a device's GUI, as it's often unstable and doesn't allow them to run command scripts to allow for easy maintenance. **D** is incorrect because Telnet is an insecure protocol that doesn't offer encryption, so any communication over Telnet can be easily intercepted by an attacker, who would then be able to obtain sensitive data.

30. You use your browser to navigate to a web page located at internal IP address 192.168.1.2, which is a web server listening on TCP port 8080. Your machine has an IP address of 10.10.10.16 and a source port of TCP 500. During the first stage of the communication (when your machine is sending the initial SYN packet to the web server), which of this information given is most probably wrong?

A. Source IP address of 10.10.10.16

B. Source TCP port 500

C. Destination IP address of 192.168.1.2

D. Destination TCP port 8080

B. TCP port 500 falls within the well-known port range 0 to 1023 and as such, it wouldn't be used by a system as a source port. Only a port from the dynamic range (49152 through 65535) would be used for that purpose. Various operating systems may have a slightly different pool of ports that are used as source ports, but they would certainly not use any of the well-known ones.

A, C, and D are incorrect. **A** is incorrect because a source IP address of 10.10.10.16 is a valid private IP address. **C** is incorrect because a destination IP address of 192.168.1.2 is also a valid address for an internal web server. **D** is incorrect because TCP port 8080 may not be the standard TCP 80 port that a web server listens to by default, but 8080 can be configured for that purpose, as it's not one of the well-known ports.

31. Which of these methods would you choose to protect a Windows 2003 physical server that needs to directly communicate with critical infrastructure?

- A. Hardening
- B. VM isolation
- C. Patching
- D. User awareness

A. The most effective method to protect a legacy system that requires direct critical infrastructure communication is to perform system hardening. That way, although the operating system is quite old, the attack surface will be reduced to only include required services.

B, C, and D are incorrect. **B** is incorrect because the question states that this machine needs to directly communicate with critical infrastructure, so isolating it by using a virtual machine will not help in that regard. In addition, the question states the server is a physical machine. As this is a legacy device and no support is provided anymore, trying to migrate it to a virtual machine for isolation might create additional issues. **C** is incorrect because there's no support for Windows Server 2003 anymore (extended support ended in 2015), which means that no new patches are available for this operating system. **D** is incorrect because user awareness would entail educating users about browsing habits and what activities to avoid performing, but since this is a server and not a user machine, that will have little effect.

32. You have enhanced your network security by adding MAC address filtering. Which of these attacks can be used against that in order to gain unauthorized network access?

- A. Sniffing
- B. MiTM
- C. Spoofing
- D. Session hijacking

C. An attacker can use a packet analyzer and identify a legitimate machine's MAC address. Following that, the attacker would spoof the attacking machine's MAC address to a legitimate MAC address in order to be allowed access to the network.

A, B, and D are incorrect. **A** is incorrect because a sniffing attack would commonly be used to obtain data being transferred in an unencrypted format (i.e., Telnet, FTP, and other similar protocols). Although an attacker would typically use sniffing to intercept the traffic and identify a legitimate MAC address that can be used for spoofing the attacking machine's address, it is the actual spoofing that allows the MAC address filtering mechanism to be defeated, not the initial sniffing that would be performed. **B** is incorrect because an MiTM (man in the middle) attack allows data between systems to be intercepted, but that wouldn't be used to defeat MAC address filtering. **D** is incorrect because a session hijacking attack entails an attacker intercepting traffic

between systems (similar to an MiTM attack), with the intent of obtaining information about an established session that will be later used to impersonate one of the communication ends. If an attacker manages to steal a user's session cookie during a communication with a server, that could be used later to establish a session with the server and try to impersonate the original user.

33. Which of the following would you use to prevent an SQL injection attack?

- A. NIDS
- B. Stateful inspection firewall
- C. Disallow stored procedures
- D. Input validation

D. Enforcing input validation is a very good countermeasure against SQL injection attacks. For example, when a website is requesting the age of a user, that field should be restricted to only contain numbers, since nothing else would be included in a user's age. That way, no special characters can be used for SQL injection (i.e., "--", ";").

A, B, and C are incorrect. **A** is incorrect because a NIDS has the ability to identify an SQL injection attack but is unable to prevent it. It can only provide an alert so a security administrator or analyst can review the event and take further action. **B** is incorrect because a stateful inspection firewall doesn't have the ability to filter web application requests, meaning it wouldn't be able to detect or prevent such an attack. **C** is incorrect because stored procedures can also be used to prevent SQL injection attacks by allowing the stored procedure to accept the user's data as input and then use predefined SQL statements to have that data inserted, thus not allowing a user to manipulate the SQL fields of a website's form, for example. As such, this feature should be enabled in order to prevent an SQL injection attack, but the answer states it shouldn't be allowed, thus making this option incorrect.

34. Which of the following devices uses MAC addresses?

- A. Hub
- B. Switch
- C. Repeater
- D. Router

B. A switch is used to connect devices over a local area network and is able to direct traffic based on device MAC addresses. That's because the switch maintains a table with device MAC addresses. When two devices on the network need to communicate, the switch does a lookup and forwards traffic from the source to the appropriate destination based on the table entries.

A, C, and D are incorrect. **A** is incorrect because a hub is commonly used to connect machines over a common network but doesn't have any context of MAC addresses. It is just used to provide device connectivity by copying the packets received and forwarding them on all other interfaces. **C** is incorrect because a repeater is used to regenerate a signal in order to allow it to reach longer distances but is not able to use MAC addresses, as it works at layer 1 (MAC addresses work at layer 2). **D** is incorrect because a router is used for forwarding packets within different networks by using static or dynamic routing. It uses IP addresses for that purpose and works at layer 3.

35. In which of the following would you configure ERSPAN so you mirror traffic to your NIDS?

- A.** Riverbed switch.
- B.** Cisco Catalyst switch.
- C.** MikroTik router.
- D.** None of the above support ERSPAN.

B. From the provided options, the only device that supports ERSPAN (Encapsulated Remote Switch Port Analyzer) is Cisco Catalyst switches. ERSPAN is a Cisco-proprietary protocol that uses a GRE tunnel and allows traffic to be mirrored to another port of the switch or to another switch within the network.

A, C, and D are incorrect. **A** is incorrect because Riverbed switches don't have the ability to support ERSPAN, as it is a Cisco-proprietary protocol. **C** is incorrect because MikroTik routers don't support ERSPAN (it's a Cisco-proprietary protocol). **D** is incorrect as it's a distractor.

36. As a network administrator you are asked to select the most cost-effective security device that can filter based on the following criteria:

- IP addresses/subnets
- Specific ports

Which of these devices would be more suitable in this case?

- A.** Router
- B.** Stateful inspection firewall
- C.** Packet-filtering firewall
- D.** IDS

C. Since the requirements only describe filtering based on IP addresses/subnets and specific ports, a packet-filtering firewall would be adequate and offer the most cost-effective option.

A, B, and D are incorrect. **A** is incorrect because a router is not considered a security device. Although quite a few models can support access lists, which might fulfill the criteria, the question specifically mentions a “security device able to filter,” while a router’s purpose is to route traffic between various networks. **B** is incorrect because although a stateful inspection firewall can fulfill the requirements, it would typically be more expensive than a packet-filtering one, and as such it doesn’t constitute the most cost-effective option. **D** is incorrect because an IDS would be used to inspect traffic for malicious patterns but wouldn’t be used to filter network traffic based on IP addresses/subnets and ports.

37. Which of these types of attackers would probably target a military installation to obtain new research and development (R&D) plans?

- A. APT
- B. Script kiddie
- C. Hacktivist
- D. White hat

A. An APT (advanced persistent threat) is a well-funded and highly skilled category of attackers, usually nation-state sponsored, that have the ability to launch sophisticated attacks towards rival government and military organizations.

B, C, and D are incorrect. **B** is incorrect because a script kiddie wouldn’t have the skill or the tools required to be able to target a military installation. **C** is incorrect because hacktivists launch attacks usually with political or environmental motives and they don’t tend to focus on profit. Hence, any information that would be obtained from such an attack would most likely be leaked online for everyone to view. **D** is incorrect because the term white hat is used to describe a penetration tester (also known as an ethical hacker) with no intent to do any harm. Hence, such a person wouldn’t target a military installation with the goal of stealing and then selling information.

38. During a DDoS attack, who would be responsible for issuing commands?

- A. Botnet
- B. Zombie
- C. Victim
- D. Bot herder

D. The bot herder is responsible for controlling the network of victim machines and issuing related instructions to them (i.e., direct them to launch a DDoS attack or send spam e-mails).

A, B, and C are incorrect. **A** is incorrect because the term botnet is used to describe a collection of victim machines under the attacker's control. **B** is incorrect because a zombie is just another term used to describe a compromised machine that is under the control of an attacker. **C** is incorrect because the term "victim" refers to a compromised machine that is now functioning as a part of the bot network (botnet) and is under the bot herder's control.

39. Which of the following needs to take place so your machine can capture all the wireless transmissions within your device's range?

A. An external USB-wireless adapter is required.

B. Wireless NIC set to nonpromiscuous mode.

C. Wireless NIC set to promiscuous mode.

D. Wireless network needs to be open (i.e., not have any security protocol enabled).

C. In order to intercept all wireless transmissions within the machine's range, the wireless NIC (network interface controller) needs to be switched to promiscuous mode. If that doesn't happen, it is not possible to capture wireless traffic from various network devices.

A, B, and D are incorrect. **A** is incorrect because there's no need for an external USB-wireless adapter. Any wireless NIC would do, as long as it's able to support operation in promiscuous mode. **B** is incorrect because if the wireless NIC is not set in promiscuous mode, only traffic towards your machine can be captured. **D** is incorrect because it doesn't matter if the wireless network is open or a security protocol is in place (like WEP or WPA). Any traffic can still be intercepted. For example, the first step to cracking WEP's encryption is to capture enough wireless traffic for further manipulation.

40. Which of the following is a valid WPA-PSK preshared key?

A. KsIA349

B. AK09EYei

C. £jh85418

D. RΣ98*6Te

B. WPA-PSK keys need to be at least 8 and up to 64 ASCII characters in length. The only answer that fulfills those requirements is option B.

A, **C**, and **D** are incorrect. **A** is incorrect because it's seven characters long, which is less than the minimum of eight. **C** is incorrect because although it's eight characters long, it contains the £ character, which is a non-ASCII character. **D** is incorrect because although it's eight characters long, it contains the Σ character, which is a Greek alphabet letter and not an ASCII character.

CHAPTER 7

Systems and Application Security

This chapter includes questions from the following objectives:

- 7.1 Identify and analyze malicious code and activity
- 7.2 Implement and operate endpoint device security
- 7.3 Operate and configure cloud security
- 7.4 Operate and secure virtual environments

This chapter contains a complete set of questions and answers regarding attacks (attack categories and attacker types, as well as related countermeasures); malicious code (types of malware, delivery methods, and countermeasures); virtual environments; and cloud computing security.

Q QUESTIONS

1. In which of these categories would a penetration tester belong?

- A.** White hacker
- B.** Black hat
- C.** Gray hat
- D.** White hat

2. A pharmaceutical company was made aware that a competitor is developing a new drug that will result in massive business loss. The lead scientist requires in-depth information about this drug and wants to identify someone who can covertly obtain data from the competitor. Which of these attackers would most likely be able to perform this task?

- A.** Script kiddie
- B.** Commercial hacker
- C.** Hactivist
- D.** Nation-state actor

3. You have submitted a suspicious file for analysis to an AV vendor, but they are experiencing difficulties in understanding the malware's operation. Based on this fact, what type of malware is most likely present?

- A.** Polymorphic virus
- B.** Worm
- C.** Armored virus
- D.** Boot sector virus

4. You just received an e-mail from someone stating he has managed to access all your files and get screenshots from your camera. Which of the following is the attacker most likely using?

- A.** RAT
- B.** Trojan horse
- C.** Ransomware
- D.** Keylogger

5. During a drive-by download attack, where would the attacker most likely host malware?

- A.** Target machine
- B.** Attacking machine
- C.** Attacker's website
- D.** Compromised third-party website

6. If you want to ensure a system is protected against a user who might execute a malicious application from a USB disk, which of these scan types would you choose?

- A.** On-demand
- B.** Real-time
- C.** Scheduled
- D.** Signature-based

7. Your company is mainly using Windows and Apple devices and has recently implemented spam filtering and web proxy filtering and installed endpoint antivirus software on all devices. Which of the following presents the greatest risk to the company?

- A. AV with no update for the last month
- B. Unpatched RHEL 6.0 OS
- C. Fully patched Windows Vista OS
- D. Phishing e-mails

8. You browse to a website and receive a pop-up message stating your computer is vulnerable and in immediate need of a missing patch. Which of the following might be present on that website?

- A. PUA
- B. Spyware
- C. Virus
- D. Scareware

9. Which of these types of hardware would most commonly be virtualized to ensure the business reduces the related cost of physical equipment?

- A. Server
- B. Firewall
- C. Router
- D. VLAN

10. Which of these terms is unrelated to virtualization?

- A. Host
- B. VPN
- C. Hypervisor
- D. Guest

11. Which of the following is not characteristic of a worm?

- A. Spreads over the network
- B. Consumes bandwidth
- C. Requires human intervention
- D. May bring a system to a halt

12. Which of these statements regarding virtualization is inaccurate?

- A. Improved elasticity
- B. Easier migration
- C. Better resource management
- D. Less robust system restoration

13. How can you best protect against VM escape?

- A. Patch the hypervisor software.
- B. Patch the guest OS.
- C. Use a next-generation firewall on the host.
- D. Use a host antivirus.

14. Which of the following would be most crucial to investigate and address as soon as possible when attempting to harden a corporate e-mail server that also acts as a DNS server?

- A. Open TCP port 53
- B. OS license expiring soon
- C. Open TCP port 80
- D. Open UDP port 161

15. One of the employees at your company is accessing social media websites. As the information security officer, you have been tasked with forming an appropriate document that prohibits that behavior. Which of the following would be best for that purpose?

- A. Security policy
- B. AUP

C. Allowed application policy

D. Related guideline

16. Which of the following provides hardware disk encryption and remote attestation?

A. BYOD

B. CYOD

C. MDM

D. TPM

17. A friend of yours had her mobile phone stolen. Which of these actions would you recommend next?

A. Remote wipe

B. Remote lock

C. MDM

D. Encrypt the phone

18. Which of the following would provide an attacker with high-level system permissions and often appear as a legitimate process?

A. Logic bomb

B. Trojan horse

C. Rootkit

D. Macro virus

19. Which of these attacks allows an attacker access to forbidden file system locations and often targets a system's password file?

A. XSS

B. Directory traversal

C. Buffer overflow

D. CSRF

20. A new company employee received an e-mail containing a Google Drive URL, which is supposedly used for exchanging data with the company's clients. When he tries to browse to it and log in, his browser presents him with a security warning stating that the connection to that site is not secure. Which term describes this attack more accurately?

- A. Pharming
- B. Phishing
- C. Social network attack
- D. Vishing

21. Which of the following tries to disable or attack the AV software?

- A. Stealth virus
- B. Multipartite virus
- C. Polymorphic virus
- D. Retrovirus

22. Which of the following statements regarding Java applets is incorrect?

- A. Require JRE to be present so they can run.
- B. They're commonly used in web applications.
- C. They execute at the server.
- D. Work on various operating systems.

23. What is the best method to protect against shoulder surfing?

- A. Clean desktop
- B. Privacy filter
- C. User education
- D. Screensaver

24. Which of the following provides the least benefit against phishing e-mails?

- A. Install spam filters on all web servers.

B. Users shouldn't open unsolicited e-mail attachments.

C. Scan all e-mails for viruses.

D. Maintain phishing campaign awareness.

25. Which of the following can't be provided by a TPM?

A. Boot protection

B. Encryption key storage

C. Remote wiping

D. Device identification

26. A cloud service can be accessed through a variety of operating systems. Which cloud characteristic does that relate to?

A. Measured service

B. Resource pooling

C. Elasticity

D. Broad network access

27. Which of these cloud operation models would be better for easy storage and access, for a European company that needs to be in full control of sensitive data as well as conform with GDPR?

A. Public cloud

B. Private cloud

C. Community cloud

D. Hybrid cloud

28. Your company is thinking of migrating to a cloud model for a variety of services. However, security is of paramount importance, so all responsibility for OS and application security should reside with your company. Which of these cloud models would be more suitable?

A. PaaS

B. SaaS

C. IaaS

D. DaaS

29. Which of these types of attackers would most likely create and use a zero-day exploit?

A. APT

B. Insider threat

C. Script kiddie

D. Social engineering

30. Which of the following is the best option to ensure secure deletion of data stored on a public cloud?

A. Use specialized software.

B. Ask cloud provider to provide the hard disk.

C. Use cloud provider's tool to securely delete data.

D. Encrypt data prior to uploading and delete when no longer needed.

31. Which of these attacks would most probably be used by a remote attacker who wants to steal sensitive files from a CEO's machine?

A. XSS

B. Trojan

C. ARP spoofing

D. SQL injection

32. Which of the following exclusively refers to an illegal operation being present?

A. Botnet

B. Backdoor

C. Adware

D. None of the above

33. A newly formed company wants to use cloud computing in order to save budget on hardware and administration. Which of the following responsibility models would be most appropriate for that purpose?

- A. PaaS
- B. Private cloud
- C. Public cloud
- D. IaaS

34. Which of these types of data residing in a cloud server would be protected by HIPAA?

- A. Patient first name
- B. Last four digits of patient credit card
- C. Patient year of birth
- D. Healthcare provider

35. Your company requires data to be labeled appropriately so it can be placed in different cloud servers, depending on its importance. How would you classify data relating to company employees?

- A. Confidential
- B. Private
- C. Sensitive
- D. Public

36. Which of these tasks would you avoid performing on a corporate host's virtual machine?

- A. Run another OS.
- B. Develop and test applications.
- C. Reverse-engineer malware.
- D. Host a file server.

37. Which of the following would prevent a zero-day exploit from running on your server?

- A. Signature-based AV
- B. Behavioral-based AV
- C. Network IPS
- D. None of the above

38. Your company wants to develop a new application that will be distributed in a preconfigured VM (containing the application running on a Linux OS). Which of these terms best describes this?

- A. Host
- B. Guest
- C. Virtual appliance
- D. Hypervisor

QUICK ANSWER KEY

- 1.** D
- 2.** B
- 3.** C
- 4.** A
- 5.** D
- 6.** B
- 7.** C
- 8.** D
- 9.** A
- 10.** B
- 11.** C
- 12.** D

13. A

14. C

15. B

16. D

17. A

18. C

19. B

20. B

21. D

22. C

23. B

24. A

25. C

26. D

27. B

28. C

29. A

30. D

31. B

32. C

33. A

34. D

35. B

36. C

37. D

38. C

ANSWERS A

1. In which of these categories would a penetration tester belong?

A. White hacker

B. Black hat

C. Gray hat

D. White hat

D. A penetration tester would be classified as a white hat, which is a term used to describe individuals who perform tests on systems to discover vulnerabilities and alert the system owners to mitigate any associated risks before attackers have a chance to exploit security gaps. Note that another term used is ethical hacker, and sometimes another term used for the process of penetration testing is ethical hacking.

A, B, and C are incorrect. **A** is incorrect because there's no term known as white hacker, as the correct term is white hat. This option is a distractor. **B** is incorrect because a black hat is an individual who intends to identify and exploit vulnerabilities and use that knowledge to perform malicious activities on the target system. **C** is incorrect because a gray hat is someone who has advanced skills in computing (and often in networking, scripting, and penetration testing) and although not primarily using that skillset nefariously, there's an element of going outside the boundaries of the law sometimes, for example, scanning a web server without permission only to alert the system owner of possible open ports that are used by vulnerable services and need to be disabled.

2. A pharmaceutical company was made aware that a competitor is developing a new drug that will result in massive business loss. The lead scientist requires in-depth information about this drug and wants to identify someone who can covertly obtain data from the competitor. Which of these attackers would most likely be able to perform this task?

A. Script kiddie

B. Commercial hacker

C. Hacktivist

D. Nation-state actor

B. A commercial hacker is someone who is usually hired to target a particular organization with a very distinct goal. Common goals include data theft, industrial espionage or sabotage, and data alteration or destruction.

A, C, and D are incorrect. **A** is incorrect because a script kiddie is someone who has basic knowledge of attacks and is mostly limited to using already existing tools but is not a sophisticated attacker. As such, a script kiddie wouldn't be suitable for performing a covert operation against a pharmaceutical company. **C** is incorrect because a hacktivist is a type of attacker with a specific ideology, that is, in favor of world peace or against weapons. Stealing data from a pharmaceutical company to provide it to a competitor is not something that a hacktivist would commonly be interested in doing. **D** is incorrect because a nation-state actor refers to an individual or group of people working on behalf of a government to perform hostile activities, commonly against other governments. Such actors possess a high degree of technical skill, have enormous budgets, and perform highly disruptive operations (i.e., attacks on nuclear power plants or military installations).

3. You have submitted a suspicious file for analysis to an AV vendor, but they are experiencing difficulties in understanding the malware's operation. Based on this fact, what type of malware is most likely present?

A. Polymorphic virus

B. Worm

C. Armored virus

D. Boot sector virus

C. An armored virus is specifically developed to thwart reverse-engineering attempts. Common techniques include several layers of code obfuscation, use of encryption, and adding code sections that don't serve any real purpose but are just present to create confusion about what the real operation is.

A, B, and D are incorrect. **A** is incorrect because a polymorphic virus is a type of virus that can change a malicious file in order to make signature detection more difficult; however, the replication and payload mechanisms remain unaltered. **B** is incorrect because a worm's main characteristic is its ability to propagate through a network without any user interaction. When it is first installed on a system, it will usually try to find a way to self-replicate to other neighboring machines, which can result in an increase in network traffic and resource consumption. **D** is incorrect because a boot sector virus is installed on the operating system's MBR (master boot record) and is then loaded into memory when the system starts. If any removable media are inserted on the system (i.e., external hard disks or USB flash drives), the virus can copy itself to those and continue spreading.

4. You just received an e-mail from someone stating he has managed to access all your files and get screenshots from your camera. Which of the following is the attacker most likely using?

- A. RAT
- B. Trojan horse
- C. Ransomware
- D. Keylogger

A. A RAT (remote access Trojan) can give an attacker full control of a target machine, including the ability to record keystrokes, access all files, use the camera to take screenshots, and install files.

B, C, and D are incorrect. **B** is incorrect because a Trojan horse is an application that is supposed to have a legitimate function but actually performs other tasks in addition to that. Assume you are looking for software that can compress audio files. You find one through a third-party website and start using it. If this was a Trojan horse, it would compress audio files but do something additional to that, possibly without even your knowing, like log your keystrokes or collect passwords from your browser and then send that to the application's author. **C** is incorrect because ransomware commonly works by not allowing a user access to his or her files before an amount is paid to the attacker. Furthermore, the term crypto-ransomware is used to describe a scenario where the user's files are encrypted and the attacker requires an amount to be paid in order to provide the encryption key to the victim. **D** is incorrect because a keylogger will only record victim keystrokes, which will then be passed to the attacker.

5. During a drive-by download attack, where would the attacker most likely host malware?

- A. Target machine
- B. Attacking machine
- C. Attacker's website
- D. Compromised third-party website

D. During a drive-by attack, an attacker will place malicious code that the victim needs to download and execute so an action takes place, that is, allow the attacker access to the victim's machine by opening a backdoor or making the target machine part of a botnet. The attacker will need to store the code somewhere that is easily accessible by a victim, so a common practice is to compromise third-party websites and store malicious code there and then redirect an unsuspecting victim to download and execute it.

A, B, and C are incorrect. **A** is incorrect because a target machine (meaning the victim's personal machine) wouldn't allow uninterrupted access to any other future victim. For example, even if the attacker finds a way to store malware on a target machine and installs a web server on it so future victims can download malicious code, that doesn't guarantee availability, as the victim might power off the computer and go on holiday for a month. **B** is incorrect because if the attacker hosts malware on his or her machine, that would be identifiable and easily blocked by law enforcement. Also, it is much more possible that someone will browse to a legitimate third-party website (i.e., a known newspaper or online store) that the attacker has compromised, as the victim would hardly ever consider that malicious. **C** is incorrect because (similarly to the previous point) a website owned by the attacker can be identified and taken down by law enforcement, so it would commonly only be available for a short period.

6. If you want to ensure a system is protected against a user who might execute a malicious application from a USB disk, which of these scan types would you choose?

- A.** On-demand
- B.** Real-time
- C.** Scheduled
- D.** Signature-based

B. A real-time scan is appropriate for continuous system protection. The antivirus will be able to inspect any application that the user attempts to execute (including the ones run from external media, like portable applications) and constantly monitor activity for malicious indicators.

A, C, and D are incorrect. **A** is incorrect because an on-demand scan will only run if the user initiates it. However, that doesn't ensure protection against a malicious application that is installed or run at any moment in time. This feature is mostly useful if the user suspects malware is present on a system or a new device has been plugged in (i.e., external storage) and that needs to be scanned for malicious artifacts. **C** is incorrect because a scheduled scan would be preferred if the user would like to ensure a system is periodically scanned for malware. However, that won't have any effect on an application that the user is running before that scan is performed. **D** is incorrect because there's no scan type known as signature-based. This refers to the type of detection an antivirus is using. Signature-based means that the antivirus works by having a signature database, which contains signatures of known malware.

7. Your company is mainly using Windows and Apple devices and has recently implemented spam filtering and web proxy filtering and installed endpoint antivirus software on all devices. Which of the following presents the greatest risk to the company?

- A.** AV with no update for the last month
- B.** Unpatched RHEL 6.0 OS

C. Fully patched Windows Vista OS

D. Phishing e-mails

C. Finding a machine running Windows Vista in a company actively using Windows presents a major risk, as that operating system's support ended in April 2017. That means that although it has been fully patched, no patches or security updates have been released from April 2017 onwards, leaving that host vulnerable to any exploits that have been released since then.

A, B, and D are incorrect. A is incorrect because antivirus solutions were recently installed on all company endpoints, so even if they haven't been updated for a month, they pose a limited risk to the business. B is incorrect because an unpatched version of RHEL 6.0 wouldn't have any impact on a business not running Red Hat Linux. D is incorrect because even if company employees receive phishing e-mails on a regular basis, spam filtering and web proxy filtering have been implemented, which will provide a fairly good amount of protection against phishing e-mails.

8. You browse to a website and receive a pop-up message stating your computer is vulnerable and in immediate need of a missing patch. Which of the following might be present on that website?

A. PUA

B. Spyware

C. Virus

D. Scareware

D. Scareware is a type of malware that prompts the user to install some type of security application, trying to entice an individual who visits a website where scareware is hosted by stating there's a missing update or patch or no antivirus present on the victim's machine. As soon as someone downloads and installs the application in question, various malicious tasks may start to take place, depending on what the goal of the scareware is (i.e., join the machine to a botnet or steal credentials).

A, B, and C are incorrect. A is incorrect because a PUA (potentially unwanted application) would commonly not have a malicious goal but is mostly used to display annoying advertisements to the user or might perform some other trivial task like change the browser's default home page or search engine. B is incorrect because any software that is classified as spyware aims to obtain as much information about a user as possible by a variety of methods like logging keystrokes, obtaining usernames and passwords from browsers, or tracking user browsing behavior and purchasing habits. C is incorrect because a virus is a type of program that can cause a variety of symptoms on a target system like slowness and increased network or resource consumption, but in order for it to execute it first needs to be run by a user, and the

question doesn't state anything about a file being downloaded and run. It just mentions a pop-up that was present when the user browsed to the website.

9. Which of these types of hardware would most commonly be virtualized to ensure the business reduces the related cost of physical equipment?

- A. Server
- B. Firewall
- C. Router
- D. VLAN

A. Most companies have multiple servers within their environment in order to support a variety of tasks, like e-mail, DHCP, Active Directory, FTP, database, web, authentication, and licensing. Due to the cost of those devices and the effort required to maintain them, they tend to be virtualized quite often, as that significantly limits the associated cost and at the same time makes administration and troubleshooting easier. It also offers smaller downtimes because if a virtual server is damaged, it can be reconstituted from an image in a relatively short time. Whereas if a physical machine is experiencing issues, a technician would need to go to the data center to troubleshoot and possibly replace faulty hardware, which can take several hours at best.

B, C, and D are incorrect. **B** is incorrect because although some vendors offer virtualized firewall versions, a firewall commonly requires substantial resources to be able to cope with network traffic and apply related rules from its ruleset (which might have thousands of rules that need to be evaluated in very short time). In addition, there are various concerns about virtual environment security and how easy it may be for an attacker that has compromised an ESXi server, for example, to gain access to VMs of security devices (i.e., firewalls) and alter related security settings to allow for access to the network. **C** is incorrect because networking devices mostly rely on specialized hardware to be able to perform decisions on how traffic should be routed in minimal time. Virtualizing such functions would take significant CPU power, not to mention it would also mean having virtual network interfaces, which would make the challenge even greater. **D** is incorrect because a VLAN (virtual local area network) is not a hardware device but a virtual component.

10. Which of these terms is unrelated to virtualization?

- A. Host
- B. VPN
- C. Hypervisor
- D. Guest

B. A VPN (virtual private network) uses the word “virtual” in its title but has nothing to do with virtualization (i.e., using a virtual server). A VPN facilitates the creation of a secure connection over an insecure network (i.e., the Internet). It is called “virtual,” as there’s no actual physical tunnel between two remote ends, but because of the VPN’s operation (providing enhanced security with encryption support), it’s like building an isolated path between those devices for the duration of the session.

A, C, and D are incorrect. **A** is incorrect because a host is the device where virtualization will be implemented, as the related virtualization software (i.e., Virtual Box or VMware) will be installed there. **C** is incorrect because a hypervisor is the actual software (i.e., Virtual Box or VMware) that offers virtualization. **D** is incorrect because any virtual machine that may be running on a host system is called a guest. Note that a physical host can support several guests, depending on its physical resource limitations.

11. Which of the following is not characteristic of a worm?

- A. Spreads over the network
- B. Consumes bandwidth
- C. Requires human intervention
- D. May bring a system to a halt

C. A worm doesn’t require any human intervention in order to execute and propagate. It has the ability to self-replicate to other machines and infect multiple hosts within a network. Note that a virus requires human intervention in order to be able to run.

A, B, and D are incorrect. **A** is incorrect because worms tend to use networks to spread to as many hosts as possible. **B** is incorrect because worms consume a lot of network bandwidth when trying to locate possible hosts to infect. At the same time, they attempt to take advantage of any protocol that will allow them to replicate (like SMB or FTP), thus creating major spikes in bandwidth consumption. **D** is incorrect because a worm’s execution is known to bring systems to a halt or even force them to restart frequently, as they are unable to cope with the amount of processing cycles and memory that a worm uses.

12. Which of these statements regarding virtualization is inaccurate?

- A. Improved elasticity
- B. Easier migration
- C. Better resource management
- D. Less robust system restoration

D. Using a virtual machine offers more robust restoration (not less). That's because the process of restoring the virtual machine when something happens is very straightforward, since a previous snapshot can be used for restoration. Even if a machine goes down and needs to be rebuilt, that can be a matter of minutes or perhaps a few hours, while in the case of a physical server, it may even take days or at least several hours. That's especially true if a physical server malfunction requires a hardware replacement before restoration can start.

A, B, and C are incorrect. **A** is incorrect because virtualization offers improved elasticity. If you have a virtualized file server and various clients start accessing files, you can always change the provisioning to account for the increased demand. If for some reason that stops, then the machine can be scaled down to address that, offering a high level of elasticity. **B** is incorrect because if migration to another machine is required, that can happen easily, as a set of files (which is what a virtual machine is composed of) will need to be copied to the new device so that restoration can take place there. **C** is incorrect because (as also mentioned earlier) a virtual machine can be provisioned according to the specific needs of an organization. If a robust web server is needed, it can be provisioned accordingly. Similarly, if an administrator notices that the machine stays idle too much, resources can be reduced to account for that. In addition, if more storage is required, that can be added with minimal interruption to the service being provided.

13. How can you best protect against VM escape?

A. Patch the hypervisor software.

B. Patch the guest OS.

C. Use a next-generation firewall on the host.

D. Use a host antivirus.

A. VM escape is a type of exploit that can take place when an attacker manages to break out of the context of a virtual machine and interact directly with the host. This is made possible due to a potential vulnerability within the hypervisor (the software that provisions the guest OS and controls how that interacts with the host). In order to ensure the risk is mitigated as much as possible, the hypervisor should always run the most up-to-date versions with the latest patches so that attackers can't use any known exploits to perform VM escape.

B, C, and D are incorrect. **B** is incorrect because although patching the guest OS is quite important, it's the vulnerability within the hypervisor that ultimately allows an attacker to perform a VM escape. Please note that patching the guest OS will aid a lot in avoiding any risks from vulnerabilities that can lead to that machine being compromised. In a typical VM escape scenario, even if the hypervisor is vulnerable, the attacker will start by first compromising the guest OS, and after that is achieved, a hypervisor exploit will follow. However, if the hypervisor is sufficiently patched, then the only thing an attacker can achieve is to compromise the guest OS (assuming the attacker is not using a zero-day hypervisor exploit). **C** is incorrect because using a next-generation firewall on the host will most likely not achieve anything, as the attacker would

perform this attack by compromising the guest OS and work through that to exploit the hypervisor. **D** is incorrect because using a host AV will not necessarily protect against this attack, as the attacker will initially compromise the guest OS and then exploit the hypervisor. It would be more beneficial to install an AV on the guest OS in the hope that the initial attack will be detected and stopped before the attacker manages to exploit the hypervisor.

14. Which of the following would be most crucial to investigate and address as soon as possible when attempting to harden a corporate e-mail server that also acts as a DNS server?

- A. Open TCP port 53
- B. OS license expiring soon
- C. Open TCP port 80
- D. Open UDP port 161

C. Having TCP port 80 open indicates that there's possibly a web server running on the machine (assuming it's using the default port). However, this machine is classified as an e-mail and DNS server, not a web server, so there wouldn't be any reason for having TCP port 80 open. As such, this would need to be further investigated (possibly verify with the IT team if there's any reason for that port to be open) to ensure any risk is mitigated accordingly.

A, B, and D are incorrect. **A** is incorrect because TCP port 53 is used by DNS for zone transfers and any requests over 512 bytes, so since the server is acting as a DNS server, this is required. **B** is incorrect because addressing a licensing issue wouldn't fall within the tasks performed during hardening, as it's not an immediate security concern. However, this would be something to report to the IT team so the issue is appropriately resolved. **D** is incorrect because UDP port 161 is used by default by SNMP (Simple Network Management Protocol), which would commonly be utilized to report the server's health status to a management console and alert the administrator of any issues.

15. One of the employees at your company is accessing social media websites. As the information security officer, you have been tasked with forming an appropriate document that prohibits that behavior. Which of the following would be best for that purpose?

- A. Security policy
- B. AUP
- C. Allowed application policy
- D. Related guideline

B. An AUP (acceptable use policy) would be the most appropriate document to specify what actions are permitted or not allowed for corporate network users. An AUP document is

something that all company employees have to adhere to. As such, including a statement about social media access not being allowed throughout the organization is something that all corporate employees will need to abide by.

A, C, and D are incorrect. **A** is incorrect because a security policy is an overarching document containing corporate security goals and high-level objectives. It references specific documents (i.e., AUP, incident response policy, or physical security policy), which the organization will use to achieve its goals. **C** is incorrect because there's no type of policy known as "allowed application policy." This option is a distractor. **D** is incorrect because a guideline is a type of document used to provide related recommendations but is not obligatory. An AUP document, on the other hand, is authoritative and usually mentions possible sanctions in the event the related policy is not adhered to.

16. Which of the following provides hardware disk encryption and remote attestation?

- A.** BYOD
- B.** CYOD
- C.** MDM
- D.** TPM

D. A TPM (Trusted Platform Module) is a hardware module that provides hardware encryption to devices supporting it. It works by storing the cryptographic keys that the OS uses for data encryption and decryption. It also provides remote attestation, meaning it provides an authorized party (like an OS vendor) the ability to verify if any tampering has taken place on the machine.

A, B, and C are incorrect. **A** is incorrect because BYOD (bring your own device) is a policy regulating the use of personal devices (like laptops, phones, and tablets) on a corporate network and doesn't relate to encryption. **B** is incorrect because CYOD (choose your own device) is very similar to the concept of BYOD, but here the organization defines a list of preapproved devices from which the employee can select one that will be purchased and managed by the company. However, this doesn't relate to hardware encryption. **C** is incorrect because MDM (mobile device management) supports full disk encryption of enrolled devices but not remote attestation.

17. A friend of yours had her mobile phone stolen. Which of these actions would you recommend next?

- A.** Remote wipe
- B.** Remote lock
- C.** MDM

D. Encrypt the phone

A. If a phone is stolen, you can report the theft to the police. However, if there's no immediate threat to human life, the odds of someone recovering that phone are very slim, unless the perpetrator is caught and the phone has not yet been sold to another party. When a device is stolen, it is more crucial to protect the data present on the owner's phone (e.g., stored credit cards) than hope to recover the device. Commonly, the best course of action would be to remotely wipe all data on the phone so that no third party can use it without the owner's knowledge.

B, C, and D are incorrect. **B** is incorrect because using a remote lock would lock the phone and only allow someone who knows the related lock password to unlock it (note that the remote lock password is not the device's touch screen password but is placed when the owner activates the remote lock). However, that doesn't do anything to delete the phone data, and there's always a chance that a thief will be able to recover that. In addition, there's no point in locking the device and leaving the data intact if there's no real chance of recovering it. A lock would be a prudent step to take if you have left your phone in a coffee shop or a bus or a cinema, for example, and you are in the process of retrieving it. However, if that's not likely to happen (like in the case of theft), then the best course of action would be to remove all data from the device. **C** is incorrect because it is highly unlikely that your friend is using an MDM solution, as that's commonly used by companies since it comes with an associated cost. Even if she does, though, the solution would still be able to remotely wipe that phone, which is an option that most modern MDM solutions would offer. **D** is incorrect because encrypting the phone is a step that should have been performed prior to the theft in order not to allow a third party to be able to access the device's data.

18. Which of the following would provide an attacker with high-level system permissions and often appear as a legitimate process?

A. Logic bomb

B. Trojan horse

C. Rootkit

D. Macro virus

C. The term rootkit is used to describe any application that acquires elevated operating system access and tries to hide its presence by a variety of methods (i.e., injecting itself into benign processes, appearing as a valid system application or process, intercepting OS API calls, or hiding in unallocated space).

A, B, and D are incorrect. **A** is incorrect because a logic bomb doesn't necessarily require elevated privileges to run, but even if it does, it won't usually try to hide its execution and usually won't have a rootkit's sophistication. Logic bombs are commonly used by developers or

malicious insiders and are executed when specific conditions are met, that is, an employee removed from the HR database or a user account deletion. **B** is incorrect because a Trojan horse is a type of malware that performs a known function while executing malicious tasks without the user's knowledge. **D** is incorrect because a macro virus is a type of malware that is embedded in a file format that supports macro execution (i.e., in Word or Excel documents) and, once opened, runs a specific set of instructions. It is worth noting that in order for a macro virus to run, macros need to have been enabled on the target computer.

19. Which of these attacks allows an attacker access to forbidden file system locations and often targets a system's password file?

- A. XSS
- B. Directory traversal
- C. Buffer overflow
- D. CSRF

B. A directory traversal attack is a type of web attack in which an attacker tries to access forbidden file system parts, that is, on a web server. The attacker would most commonly target the password file so user credentials can be harvested, but the term is used to describe access to any directory where access is not allowed.

A, C, and D are incorrect. **A** is incorrect because an XSS (cross-site scripting) attack would inject malicious code (i.e., JavaScript or HTML) to a web page, with the goal of bypassing existing access control and executing the malicious code. **C** is incorrect because a buffer overflow attack works by allowing the attacker to write into specific memory areas that shouldn't be normally available (i.e., an application using a buffer to hold data of a specific size, which is being overrun when more data is supplied and due to the application being vulnerable to a buffer overflow attack, it tries to store that information in the predefined buffer, thus resulting in writing to other memory areas). **D** is incorrect because a CSRF (cross-site request forgery) attack aims to utilize the user's browser to send a malicious request to a website without any indication to the user. Assume you are logged in to an online gaming website and want to transfer some of your credits to a friend. A CSRF exploit can be used while you are logged in to the website to have credits sent to the attacker's account instead.

20. A new company employee received an e-mail containing a Google Drive URL, which is supposedly used for exchanging data with the company's clients. When he tries to browse to it and log in, his browser presents him with a security warning stating that the connection to that site is not secure. Which term describes this attack more accurately?

- A. Pharming
- B. Phishing

C. Social network attack

D. Vishing

B. This is a typical example of a phishing e-mail being used to obtain user credentials. The employee received what was supposedly a standard e-mail invitation to the company's Google Drive location to which he would probably log in using his corporate credentials, thus passing that to an attacker. There are two things that companies need to address to avoid this. One thing is to ensure employees get adequate training to distinguish illegitimate e-mails, and the second thing is to be cognizant of security warnings. If one is not certain what the reason for the displayed message is, it's best to reach out to the IT team than to just proceed with the activity.

A, C, and D are incorrect. **A** is incorrect because pharming is a type of attack in which the victim tries to browse to a particular website but is instead getting redirected to a fake website because of DNS being exploited (i.e., a vulnerability in DNS itself or manual modification of the user's machine resolution ability, for example, by modifying the hosts file). **C** is incorrect because a social network attack would be performed using some social network invitation or e-mail supposedly originating from social media, that is, a LinkedIn connection invitation, to entice the user to click the URL contained in that e-mail and accept the new connection. **D** is incorrect because vishing is a type of social engineering attack that uses phone or VoIP and engages the victim in a discussion in order to extract useful personal information; for example, someone calling and posing as an electricity provider's employee and asking to confirm your full name, address, and date of birth to give you details about a new offer.

21. Which of the following tries to disable or attack the AV software?

A. Stealth virus

B. Multipartite virus

C. Polymorphic virus

D. Retrovirus

D. A retrovirus is a type of virus that tries to attack a device's antivirus program. It may attempt to warn the victim that supposedly unsafe software has been found on the machine and then uninstall legitimate AV applications. Sometimes it has the ability to modify the virus signature database and corrupt or delete it altogether.

A, B, and C are incorrect. **A** is incorrect because a stealth virus doesn't attack the AV application but actually tries to evade detection from it. Sometimes it does that by changing the offending file's name or size to make it seem like the infected file is the actual legitimate host file. **B** is incorrect, as a multipartite virus works by infecting various parts of a target system with the intent of making removal very difficult. **C** is incorrect because a polymorphic virus can

change a malicious file in order to make signature detection more difficult; however, the replication and payload mechanisms remain unaltered.

22. Which of the following statements regarding Java applets is incorrect?

- A.** Require JRE to be present so they can run.
- B.** They're commonly used in web applications.
- C.** They execute at the server.
- D.** Work on various operating systems.

C. Java applets are interactive components of a web page (i.e., buttons, text) that are executed on the client's web browser and use a VM (commonly known as JVM, or Java Virtual Machine) for execution. A typical example is capturing a user's mouse input so the program can manipulate it afterwards.

A, B, and D are incorrect. **A** is incorrect because Java applets require JRE (Java Runtime Environment) to be installed in the client's browser so they can function properly. If that is not installed, a blank window is usually displayed stating that JRE is required in order for the applet to execute, which is something that attackers commonly try to exploit and redirect users to malicious websites, where they are supposedly downloading the latest JRE version while in fact malware is being delivered. **B** is incorrect because Java applets are used by many web applications, for example, to transfer the processing load to the client. **D** is incorrect because Java applets work on a variety of operating systems (i.e., Windows, Linux, UNIX, macOS), as they work on the client browser and are not dependent on the OS architecture.

23. What is the best method to protect against shoulder surfing?

- A.** Clean desktop
- B.** Privacy filter
- C.** User education
- D.** Screensaver

B. A privacy filter (also known as a visual filter) can be quite effective against any shoulder surfing attack. When such a filter is used, the visible area of the screen is limited to only the person placed directly in front of the monitor. This can be extremely beneficial for people who travel a lot and tend to work in crowded places (i.e., train stations, buses, and airports).

A, C, and D are incorrect. **A** is incorrect because a clean desktop entails clearing out the files and folders residing on a user's desktop, but that won't help in the case where individuals tend to

work extensively with a variety of open documents (i.e., working all the time through presentations, spreadsheets, and various documents). **C** is incorrect because user education always goes a long way to limiting or preventing various types of social engineering attacks. However, the education aspect of this case would be to encourage individuals to use a privacy filter while working in public places so they can protect themselves from shoulder surfing attacks. **D** is incorrect because a screensaver won't help discourage shoulder surfing when a user is working on a computer. A screensaver is useful if the machine is left unattended for a short time. However, that will not have any effect if a password isn't set when the screen is locked, which is a common mistake.

24. Which of the following provides the least benefit against phishing e-mails?

- A.** Install spam filters on all web servers.
- B.** Users shouldn't open unsolicited e-mail attachments.
- C.** Scan all e-mails for viruses.
- D.** Maintain phishing campaign awareness.

A. No e-mail accounts would be configured on a web server because that device should only be serving client web requests. Hence, there's no real benefit of having spam filters installed on it.

B, C, and D are incorrect. **B** is incorrect because users shouldn't open e-mail attachments from people they don't know. This is, of course, a general rule of thumb, as sometimes it is very hard to do in reality, especially from people in HR, sales, and similar roles. **C** is incorrect because AV clients should be configured to scan all e-mail content for malicious URLs and all attachments for any suspicious files. **D** is incorrect because it is highly recommended for every individual to maintain vigilance about any large phishing campaign so appropriate attention can be demonstrated. For example, if you work at a fashion retailer, it would be great to know there's a large phishing campaign targeting fashion retailers these days so you maintain appropriate awareness.

25. Which of the following can't be provided by a TPM?

- A.** Boot protection
- B.** Encryption key storage
- C.** Remote wiping
- D.** Device identification

C. A TPM (Trusted Platform Module) can't provide device remote wiping, which is a feature commonly supported by MDM (mobile device management) solutions.

A, B, and D are incorrect. **A** is incorrect because TPMs offer boot protection by storing hash values of specific system files and other OS key metrics in order to allow the system to detect any alterations by unauthorized software. **B** is incorrect because the TPM is used to store any hard disk encryption keys. **D** is incorrect because device identification is achieved by a unique identification key, which is embedded in the TPM module by the manufacturer during its creation.

26. A cloud service can be accessed through a variety of operating systems. Which cloud characteristic does that relate to?

A. Measured service

B. Resource pooling

C. Elasticity

D. Broad network access

D. Having the ability to access a cloud service from a variety of operating systems is known as broad network access. Usually, this feature uses two main methods. One is making an application available via a web browser, which is seamless to any operating system being used, as the application runs over the client's web browser. Another method is to allow a user to download the related software locally to a machine. But most cloud providers that use this feature make clients available for a variety of platforms, like UNIX, Linux, Windows, Android, and macOS.

A, B, and C are incorrect. **A** is incorrect because a measured service refers to being able to get charged only for what is being specifically used. If a business has 50 cloud servers but is currently only using 20 of them while the rest are powered off, they have to be able to pay according to that. **B** is incorrect because resource pooling refers to being able to effectively manage resources so they are used as much as possible. If an eight-core server is only using two cores for client A, then use the remaining processing power to serve client B (if that is allowed). **C** is incorrect because elasticity refers to being able to scale up or down, according to the needs of the business, at any given moment in time. If you are an online travel agency and it's a school holiday period, you might find that you need much more resources than other periods of the year, so you should be able to increase your resources accordingly for that period and then scale down again with no service disruption.

27. Which of these cloud operation models would be better for easy storage and access, for a European company that needs to be in full control of sensitive data as well as conform with GDPR?

A. Public cloud

B. Private cloud

C. Community cloud

D. Hybrid cloud

B. Since the company wants to maintain easy storage and access but also needs to conform with GDPR, the best option would be to use a private cloud. This will provide ease of storage (and associated administration) while the organization will have full control of all sensitive data, which will be residing in the company's private cloud instance. Of course, since GDPR conformity is a prerequisite, the instance can be located in any European company branch.

A, C, and D are incorrect. **A** is incorrect because using a public cloud will not give the company full control over the data nor provide ease of access in case a public cloud instance is damaged, is under maintenance, or has its data stolen. **C** is incorrect because a community cloud is commonly implemented when different users want to share information over a common media; for example, two different universities with researchers participating in a common research project. However, the company is a sole organization, and the data in question would only be accessible by its employees and not anyone else. **D** is incorrect because a hybrid cloud combines two clouds (or more). As such, it wouldn't be suitable in this case because only a private cloud would provide full control over the data while maintaining easy storage and access.

28. Your company is thinking of migrating to a cloud model for a variety of services. However, security is of paramount importance, so all responsibility for OS and application security should reside with your company. Which of these cloud models would be more suitable?

A. PaaS

B. SaaS

C. IaaS

D. DaaS

C. Since the company wants to remain responsible for its own security while using a cloud model of choice, the best option is to select a model in which the client has the greatest degree of control, which is IaaS (Infrastructure as a Service). In that model, the cloud provider is responsible for providing the hardware (i.e., servers, storage, networking) and the client is responsible for everything else (i.e., operating system and application installation and associated security).

A, B, and D are incorrect. **A** is incorrect because in PaaS (Platform as a Service) the cloud provider is responsible for providing an operating system and associated applications (in addition to hardware) so the security of those resides with the cloud provider. **B** is incorrect because SaaS (Software as a Service) provides the least amount of control to a client since the cloud provider is responsible for everything and the client is just using a particular service (i.e., Dropbox). **D** is incorrect because there's no cloud model known as DaaS. This is a distractor.

29. Which of these types of attackers would most likely create and use a zero-day exploit?

- A.** APT
- B.** Insider threat
- C.** Script kiddie
- D.** Social engineering

A. The term APT (advanced persistent threat) refers to individuals or, most commonly, groups of attackers who aim to compromise target networks or devices and stay undetected for long periods while persistency is maintained. They are usually politically motivated and often use very sophisticated attack mechanisms, including zero-day exploits, as they possess the manpower and technical skillset to create them.

B, C, and D are incorrect. **B** is incorrect because an insider threat would not have a need for a zero-day exploit to gain access to a system, but would instead usually take advantage of the already existing access due to being a company employee and possibly place some type of backdoor or try to perform privilege escalation. **C** is incorrect because a script kiddie wouldn't be technically skilled enough to create a zero-day exploit and would most commonly be restricted to using freely available tools or already prepared scripts. **D** is incorrect because social engineering is not a type of attacker but actually a type of attack. The individual using such an attack is called a social engineer.

30. Which of the following is the best option to ensure secure deletion of data stored on a public cloud?

- A.** Use specialized software.
- B.** Ask cloud provider to provide the hard disk.
- C.** Use cloud provider's tool to securely delete data.
- D.** Encrypt data prior to uploading and delete when no longer needed.

D. Since the data is stored on a public cloud, it is very difficult to be certain it has been fully removed since the hard disk is owned by a public cloud provider. As such, the best option is to ensure data is encrypted before being uploaded, using a strong encryption algorithm. Even if your attempt to remove it from the public cloud results in some being left behind, it will still be encrypted and won't be easily usable by any third party managing to get a copy of it.

A, B, and C are incorrect. **A** is incorrect because it doesn't matter what software you use or how good it is—you can never be certain that data has been fully removed if you don't own the storage media being used to store it. **B** is incorrect because there's a slim chance of obtaining a

hard disk from a public cloud provider, especially since it might hold other client data in addition to yours. Even if this was feasible, it's not a practical solution, as data may be spread across a variety of disks in different geographical locations. **C** is incorrect because using a tool provided by your cloud provider doesn't guarantee full deletion of the data, especially if you don't know how the tool exactly works or how credible it is, but also due to the fact that you have no way of verifying its operation.

31. Which of these attacks would most probably be used by a remote attacker who wants to steal sensitive files from a CEO's machine?

- A.** XSS
- B.** Trojan
- C.** ARP spoofing
- D.** SQL injection

B. There are various categories of Trojans, like downloaders (created for downloading additional malicious files to a target machine), remote access Trojans (aiming to provide the attacker full control of a machine), backdoor Trojans (which will create a backdoor to a system so the attacker can gain remote access), and infostealer Trojans (which perform information theft from the machine they infect). An infostealer Trojan would be the best option for an attacker who wants to acquire sensitive data from a CEO's machine.

A, C, and D are incorrect. **A** is incorrect because an XSS (cross-site scripting) attack would inject code (i.e., JavaScript or HTML) into a web page but wouldn't be suitable in the case of attacking the CEO's machine, as that's not running any web service. **C** is incorrect because ARP spoofing requires access to the local network in order to be able to launch a successful attack. That's because during an ARP spoofing attack, the attacker needs to be able to attach his or her machine to the local network so he or she can send ARP messages with the goal of impersonating a legitimate network host. However, as the question mentions a remote attacker, the attack wouldn't be useful in this scenario. **D** is incorrect because an SQL injection attack is an application attack commonly launched against web servers. As in the previous case of the XSS attack, this wouldn't be useful against a CEO's machine, which wouldn't be running any web service that can be exploited.

32. Which of the following exclusively refers to an illegal operation being present?

- A.** Botnet
- B.** Backdoor
- C.** Adware
- D.** None of the above

C. Adware displays advertisements about a variety of products. The author either gains money by individuals viewing the ads or by enticing people to click a link, box, banner, or other advertisement element. As such, it is never considered something that a user would authorize as being installed, especially since some adware applications are known to have keylogging and spyware capabilities.

A, B, and D are incorrect. **A** is incorrect because although a botnet may usually be referred to as something malicious, it doesn't necessarily make it the case. A botnet can be defined as a collection of machines working towards a common task. When that network of devices is used for nefarious purposes (i.e., to perform a DoS attack), then the botnet is definitely deemed malicious and illegal. But there are also legitimate purposes for such a network; for example, machines joined into the SETI (Search for Extraterrestrial Intelligence) program, which "borrow" computing cycles so that science is advanced by searching for extraterrestrial life. **B** is incorrect because backdoors (also known as programming hooks) are often implemented by developers so as to have emergency access in the event that the application or system experiences issues. However, if they aren't removed promptly, an attacker may discover them and gain system access, but that's more of an oversight or poor practice than something relating to an illegal operation or malware. However, if an attacker, disgruntled developer, or company employee was placing a backdoor to retain uninterrupted access to an application or system, this would be an illegal action. **D** is incorrect because this is a distractor.

33. A newly formed company wants to use cloud computing in order to save budget on hardware and administration. Which of the following responsibility models would be most appropriate for that purpose?

A. PaaS

B. Private cloud

C. Public cloud

D. IaaS

A. The only two responsibility-related models that are mentioned as possible answers are PaaS and IaaS. The question mentions that the company is newly formed and needs to reserve hardware and administration budget, hence it would make sense to choose the model that would offload as many functions as possible to the cloud provider, which is PaaS (Platform as a Service). In PaaS, the cloud provider would be tasked with providing the necessary hardware, OS, and associated applications.

B, C, and D are incorrect. **B** is incorrect because the term "private cloud" is used to refer to a cloud operation model, not a responsibility model. **C** is incorrect because the term public cloud is also used to refer to a cloud operation model and not a responsibility model. **D** is incorrect because IaaS (Infrastructure as a Service) entails a significant administration burden (as the

customer is the one responsible for maintaining the OS and associated applications, including upgrading and patching), which is something that the company wants to avoid.

34. Which of these types of data residing in a cloud server would be protected by HIPAA?

- A. Patient first name
- B. Last four digits of patient credit card
- C. Patient year of birth
- D. Healthcare provider

D. HIPAA (Health Insurance Portability and Accountability Act) relates to protection of PHI (protected health information), which is any information about an individual's health care, status, or related treatment provisioning. As such, an individual's healthcare provider is classified as PHI and should remain securely protected under HIPAA.

A, B, and C are incorrect. **A** is incorrect because a patient's first name is not distinctive to the individual and doesn't require special handling. However, please note that the full name warrants protection under HIPAA. **B** is incorrect, as the last four digits of a credit card don't provide enough information to uniquely identify the card holder. In addition, even a full credit card number doesn't constitute PHI. That information falls under PII (personally identifiable information). **C** is incorrect, as a patient's year of birth is not classified as PHI since it doesn't provide enough information to uniquely identify an individual.

35. Your company requires data to be labeled appropriately so it can be placed in different cloud servers, depending on its importance. How would you classify data relating to company employees?

- A. Confidential
- B. Private
- C. Sensitive
- D. Public

B. Any data relating to company staff would be classified as private since it's only intended to be used internally, and any intentional or unintentional release of the information would cause a very negative impact to the company.

A, C, and D are incorrect. **A** is incorrect because a label of confidential would commonly refer to data that has grave importance to the company (like an acquisition of a new organization) and if leaked would have tremendous impact. **C** is incorrect because sensitive data refers to

information that should be protected (i.e., salary ranges for different types of positions) but wouldn't have a detrimental effect if disclosed. **D** is incorrect because public data is available to the wider public, and employee data would never be freely available.

36. Which of these tasks would you avoid performing on a corporate host's virtual machine?

- A.** Run another OS.
- B.** Develop and test applications.
- C.** Reverse-engineer malware.
- D.** Host a file server.

C. Virtual machines can be used to perform a variety of tasks, and several individuals do use them for reverse engineering. However, that is not recommended to be performed on the corporate network. If you intend to use a virtual machine to analyze malware, that should be done in an air-gap environment, with no possibility of communicating with the corporate network. Also, you need to consider the fact that some types of malware don't run on virtual environments, as they have the ability to detect them and stay dormant. Finally, another consideration would be to ensure your host system is adequately protected from common attacks and exploits (i.e., VM escape, insecure configuration of hypervisor and guest OS).

A, B, and D are incorrect. **A** is incorrect because running another operating system is a common task for a VM. In fact, sometimes individuals (like developers or system administrators) tend to use them as an alternative to their primary operating system. **B** is incorrect because VMs can be used to develop and test applications. If a developer is creating a new application that is supposed to be running on four different operating systems, then there's no point in having four different physical machines. It would be much easier and cost-effective to just use four VMs. **D** is incorrect because hosting a file server can be easily achieved by a VM. Several companies have started to host various types of servers in virtualized environments in order to achieve more efficiency, less administration, and reduced cost.

37. Which of the following would prevent a zero-day exploit from running on your server?

- A.** Signature-based AV
- B.** Behavioral-based AV
- C.** Network IPS
- D.** None of the above

D. The key takeaway from this question is that no security tool can guarantee prevention of a zero-day exploit. The challenge with those type of exploits is that they are, by definition, new to

the security vendor and the vendor is vulnerable to the related exploit. As such, none of the options are guaranteed to prevent a zero-day exploit.

A, B, and C are incorrect. **A** is incorrect because a signature-based AV works by using a signature database, which wouldn't include a zero-day exploit. **B** is incorrect because a behavioral-based AV wouldn't necessarily prevent a zero-day exploit. However, please note that any behavioral analysis is your best chance at identifying any suspicious activity, as it will analyze the activity patterns so there's a chance that an active exploit is detected. However, the question states what would prevent such an exploit, and this method is not guaranteed to prevent it. **C** is incorrect because a NIPS (network intrusion prevention system) would only be able to inspect and act on offending network traffic, but not at something taking place on your server.

38. Your company wants to develop a new application that will be distributed in a preconfigured VM (containing the application running on a Linux OS). Which of these terms best describes this?

- A. Host
- B. Guest
- C. Virtual appliance
- D. Hypervisor

C. Your company aims to provide an easy way of delivering the application to its clients for use in a virtual environment. Hence, the application will be installed on a virtual machine running Linux, which will be ready for distribution to any client. That way, any client can obtain this preconfigured virtual machine, install it on their host operating system, and start using it. A virtual appliance is a VM with a preconfigured OS and one available application.

A, B, and D are incorrect. **A** is incorrect because, as stated earlier, a host is the machine where virtualization will take place. It can be a server or a desktop or whatever other type of device the hypervisor will run on. **B** is incorrect because guest is a generic term used to describe any virtual machine running on a host. **D** is incorrect because a hypervisor is the software that offers virtualization (i.e., Virtual Box or VMware). Note that in the case of a virtual appliance, the goal is to create an image that can be deployed immediately from the client. But in order for that to happen, a specific hypervisor will need to be used. If the virtual appliance has been created to run on Virtual Box, then that's the hypervisor that needs to be used. If another hypervisor is used, the virtual appliance won't function.

APPENDIX A

Pre-assessment Test

Typically, people either love pre-assessment tests or they hate them. I used to belong to the second category. However, as I participated in many certification exams and had the privilege of teaching quite a large number of people over the years, I gradually realized that students don't actually hate pre-assessments but they do tend to be afraid of them (or assessments in general, for that matter).

The goal of this pre-assessment test is to allow you to adjust your studies according to the areas you're not familiar with at the level that the exam demands you to be. For example, most technical people know whether AES is a symmetric or asymmetric encryption algorithm. But do you know the possible key sizes it uses? Can you distinguish if a key of 56 bits can be an actual AES key?

That is the type of the thing the assessment can help you identify so your efforts are focused accordingly. However, if you don't know the answers to any of those AES questions, or you get a very low test score after taking this pre-assessment test, or you are not even in the mood to attempt the test, that really doesn't mean anything in terms of predicting your success. I have met people who didn't have any idea of what penetration testing is and as soon as they compromised their first machine they got hooked. They started putting in endless hours to master the techniques, eventually landed a junior ethical hacker role, and began winning their first capture the flag competitions within two to three years. The point is that if you are diligent and thorough in your approach, nothing will stop you from passing the exam and obtaining your SSCP certification.

Start reading the SSCP theory in order to get the hang of things, then create a study plan and stick to it. Even if you don't feel like taking the test to start with, you can always brush through the theory and take it at a later time. However, if you do decide to take it, you can greatly benefit from the fact that it might help you identify at the outset some areas that require improvement. This test contains a collection of 30 questions across all seven domains. The specific questions that relate to each domain are provided in the following table:

| Domain | Exam Weight | Question numbers in Pre-assessment Test |
|---|--------------------|--|
| 1.0 Access Controls | 16% | 1–5 |
| 2.0 Security Operations and Administration | 15% | 6–10 |
| 3.0 Risk Identification, Monitoring, and Analysis | 15% | 11–14 |
| 4.0 Incident Response and Recovery | 13% | 15–18 |
| 5.0 Cryptography | 10% | 19–22 |
| 6.0 Network and Communications Security | 16% | 23–26 |
| 7.0 Systems and Application Security | 15% | 27–30 |
| Total | 100% | |

For example, if you take the test and find that you get the answers to the first four questions wrong, then it would be highly advisable to review the theory relating to domain 1.0 in detail. However, sometimes you might get a question wrong because you didn't understand it properly or because the answers related to items you are not familiar with.

On the other hand, a good score does not necessarily indicate that you are ready to take the exam. Don't forget that this pre-assessment test only comprises of 30 questions while the actual exam contains 125. Also, students tend to respond instinctively to questions they don't know by selecting random answers and sometimes those are correct. The time for random guessing is not when you take this pre-assessment test as that will not provide accurate results and will not help you study efficiently for your exam. If you put enough hours into preparing for the exam and testing your knowledge as you go along, you will very soon find out that you don't really need to guess on any answers.

Stick to the question facts. Read everything carefully. If you don't remember the theory behind the question, go back to supplement your knowledge and then give the question another go. Don't give up. Keep at it. Plan your studies. Assess yourself as you go along. Adjust your study plan accordingly (spend more time on those chapters that cover topics you are less familiar with). Don't rush. When you are studying, time is something you are working with, but when you are taking a real exam, it's what you are working against. Oh, and because I mentioned time, to make this a more realistic assessment, remember to only allocate about 43 minutes to complete this assessment. That is proportional to the three hours that is provided to complete the 125 questions of the actual exam.

Best of luck in your preparation journey, although I am sure you will do great!

Q QUESTIONS

1. Bob wants to log on to a company system and is presented with a screen asking him to enter a username and password, followed by a second screen asking for an OTP. Which of these elements falls under identification?

- A. Password
- B. OTP
- C. Username
- D. Username and password combination

2. Captain Cramer is reviewing a biometric authentication system that will be used at a nuclear missile silo. Which of the following is the most likely reason that would prevent him from approving the usage of this device?

- A. Low CER

B. Increased amount of type 1 errors

C. Increased amount of type 2 errors

D. High FRR

3. Felly Corp. wants to improve the level of current security policies. As such, employees won't be allowed to connect to critical infrastructure using corporate machines from the guest network. This is an example of:

A. RBAC (Role-based Access Control)

B. MAC

C. ABAC

D. DAC

4. A large law firm has a team (designated Alpha) representing a manufacturing client, while another team (designated Bravo) is working on a case relating to one of the manufacturing client's biggest competitors. Bravo team would greatly benefit from information that Alpha team has access to. Which of these models would be the most appropriate to ensure that a conflict of interest is avoided?

A. Bell-LaPadula

B. Biba

C. Clark-Wilson

D. Brewer-Nash

5. Which of the following is not an SSO technology?

A. SAML

B. SecureAuth IdP

C. Kerberos

D. OpenIDConnect

6. A network administrator is considering placing a new firewall at the network perimeter. Which type of control is this?

A. Detective

B. Corrective

C. Preventive

D. Recovery

7. A company backup policy primarily uses magnetic tapes for backups. They also regularly back up data to the company cloud. Which type of control does the company cloud represent?

A. Detective

B. Deterrent

C. Recovery

D. Compensating

8. Which of the following is not a security policy characteristic?

A. Mission statement

B. Enforcement section

C. Sanctions statement

D. Accountability statement

9. What process would take place if an (ISC)² member doesn't comply with the (ISC)² code of ethics?

A. Any violators will be reported to the company the member is working for.

B. (ISC)² will revoke their certifications.

C. A review from a peer panel will take place.

D. They will get a one-off warning.

10. In order for a merchant to comply with PCI-DSS, a network scan is required to be completed by which of the following?

A. SAQ

B. ASV

C. QSA

D. ISA

11. A company's single e-commerce server has gone down. The estimation of loss for each day it's not online is \$500, and a technician charges \$400 per day to repair it, with an initial estimation of repair time being one day. The monthly likelihood of such incidents is estimated at three. What is the ALE?

- A. \$32,400
- B. \$64,800
- C. \$15,600
- D. \$53,700

12. An IPS is placed in front of the DMZ segment housing an FTP and a web server. Suddenly, the network administrator receives a report of traffic from a specific external IP address not being received by the web server. However, a packet capture right before the IPS suggests that web traffic is passing properly. What is the most likely reason for this?

- A. Malformed web request.
- B. Packet capture has not been set to capture web traffic properly.
- C. Physical link failure.
- D. IPS dropping the traffic in question.

13. Which of the following is the most important consideration for a new SIEM implementation?

- A. Configure time synchronization.
- B. Adjust appropriate alert thresholds.
- C. Ensure all available logs are ingested.
- D. Provide security analyst access to individual log sources.

14. You are the head of security and your company is running a core finance application on a legacy Windows 2003 server. A recent vulnerability scan flagged the machine in question as being out of support (extended support for Windows 2003 ended in 2015). After contacting the application vendor, they alert you to the fact that the application is quite old and supported only on a Windows 2003 operating system. Which of these options do you choose?

- A. Upgrade the server to the latest version.

- B.** Migrate the application to a new server.
- C.** Identify a new application that can be safely used.
- D.** Accept the risk.

15. A company's client database is stored on a server so that online purchases can take place. The database needs to be accessible constantly even if the server goes down. Which of the following would be more suitable for that purpose?

- A.** Failover cluster
- B.** Load balancing
- C.** Full backup
- D.** Differential backup

16. Using a proprietary forensic tool for investigation relates to which of these reliability factors?

- A.** Clarity
- B.** Error rate
- C.** Credibility
- D.** Testability

17. An organization performs a full weekly backup and is considering combining that with a differential or incremental one. Why would a full/differential backup be preferred over a full/incremental backup?

- A.** A full/differential backup takes less time to complete during the week.
- B.** When restoring a system, recovery time is reduced.
- C.** A full/incremental backup takes less time to complete during the week.
- D.** Differential backups are smaller in size.

18. A company's RPO is set at three days, while the MAO is four days. Which of these values indicates a realistic RTO?

- A.** 3

B. 5

C. 7

D. 9

19. Consider two text files: The first one contains the sentence “this is an SSCP example” and the second one contains the sentence “this is another SSCP example with different text”. The MD5 hash of the first file is 063ba1995095c95c2e2184837fd47067, and the MD5 hash of the second file is 063ba1995095c95c2e2184837fd47067. What is this an example of?

A. Work factor

B. Collusion

C. Rainbow tables

D. Collision

20. A key escrow is used for storing:

A. Public keys

B. Digital certificates

C. WOT certificates

D. Private keys

21. Which of the following would be used for a new certificate registration request?

A. WOT

B. RA

C. CRL

D. OCSP

22. You want to ensure your company’s e-mails are encrypted to maintain communication confidentiality. Which of these options would you choose for that purpose?

A. DKIM

B. SPF

C. S/MIME

D. DMARC

23. Which topology is the least resilient?

A. Star

B. Tree

C. Mesh

D. Bus

24. Which of the following would be performed by NAC?

A. Authenticate a remote client.

B. Create audit logs about the actions a machine performed while on the network.

C. Check a client machine against an organizational health policy.

D. Allow a client machine to issue specific commands.

25. You have a network composed of a hundred endpoints, ten firewalls, five routers, five IPS sensors, and three proxies across two geographical regions. Data from all of those devices is being forwarded to a SIEM platform. What would be the most crucial element for any type of security investigation you perform?

A. Log integrity

B. Verbose logging across all devices

C. Appropriate alert thresholds

D. Time synchronization

26. Which of these media types is most susceptible to electromagnetic interference?

A. Shielded twisted pair

B. UTP

C. Optical fiber

D. STP cat 6

27. How can you best protect against VM escape?

- A.** Patch the hypervisor software.
- B.** Patch the guest OS.
- C.** Use a next-generation firewall on the host.
- D.** Use a host antivirus.

28. Which of these attacks allows an attacker access to forbidden file system locations and often targets a system's password file?

- A.** XSS
- B.** Directory traversal
- C.** Buffer overflow
- D.** CSRF

29. Which of the following tries to disable or attack the AV software?

- A.** Stealth virus
- B.** Multipartite virus
- C.** Polymorphic virus
- D.** Retrovirus

30. Which of these types of data residing in a cloud server would be protected by HIPAA?

- A.** Patient first name
- B.** Last four digits of patient credit card
- C.** Patient year of birth
- D.** Healthcare provider

QUICK ANSWER KEY

1. C

2. C

3. C

4. D

5. B

6. C

7. D

8. C

9. C

10. B

11. A

12. D

13. A

14. D

15. A

16. D

17. B

18. A

19. D

20. D

21. B

22. C

23. D

24. C

25. D

26. B

27. A

28. B

29. D

30. D

ANSWERS A

1. Bob wants to log on to a company system and is presented with a screen asking him to enter a username and password, followed by a second screen asking for an OTP. Which of these elements falls under identification?

A. Password

B. OTP

C. Username

D. Username and password combination

C. Bob first enters a username to claim his identity. Please note that at this point he just claims to be Bob, without providing any further evidence to support that (like his password).

A, B, and D are incorrect. A is incorrect because the password is used to authenticate the user and not identify him. B is incorrect because OTP (one-time password) is used as a second step of a two-step verification process. The user would typically encounter this screen only after entering a valid username and corresponding password (meaning he has claimed he is Bob and has provided Bob's correct password). D is a distractor, as a username and password pair can't be used in order to provide identification of a user. Passwords are used for authentication.

2. Captain Cramer is reviewing a biometric authentication system that will be used at a nuclear missile silo. Which of the following is the most likely reason that would prevent him from approving the usage of this device?

A. Low CER

B. Increased amount of type 1 errors

C. Increased amount of type 2 errors

D. High FRR

C. A type 2 error (also known as FAR, False Acceptance Rate) can affect a system in a more negative way, as it means that it improperly allows access to illegitimate users, which would be detrimental in the case of a nuclear silo.

A, B, and D are incorrect. **A** is incorrect because ideally a good biometric system needs to have a low CER (Crossover Error Rate), as that means it is performing better. **B** is incorrect because type 1 errors refer to FRR (False Rejection Rate), meaning the system doesn't allow access to a legitimate user. Although this is a nuisance, it doesn't even come close to how worse things can become by allowing illegitimate users to enter the secure premises (which is what answer C depicts). **D** is incorrect, as it has exactly the same meaning as B but instead of the term "type 1 error," FRR is used.

3. Felly Corp. wants to improve the level of current security policies. As such, employees won't be allowed to connect to critical infrastructure using corporate machines from the guest network. This is an example of:

A. RBAC (Role-based Access Control)

B. MAC

C. ABAC

D. DAC

C. ABAC works by evaluating subject and object attributes (i.e., a user attempting to log in to a finance server from the guest network) and allows/rejects accordingly.

A, B, and D are incorrect. **A** is incorrect because Role-based Access Control defines access by using roles or groups and isn't suitable for this scenario. **B** is incorrect because MAC uses labels for subject and object identification. **D** is incorrect because DAC models grant access to objects based on the identity of each subject (i.e., Alice has read and write permissions to a file).

4. A large law firm has a team (designated Alpha) representing a manufacturing client, while another team (designated Bravo) is working on a case relating to one of the manufacturing client's biggest competitors. Bravo team would greatly benefit from information that Alpha team has access to. Which of these models would be the most appropriate to ensure that a conflict of interest is avoided?

A. Bell-LaPadula

B. Biba

C. Clark-Wilson

D. Brewer-Nash

D. The Brewer-Nash model (also known as the Chinese Wall model) is primarily used to avoid conflicts of interest. According to that, an ethical screen model is used to classify data to distinct conflict-of-interest classes. In this case, a simple way of classifying data would be to ensure that the manufacturing client is placed in one class while their competitor is in another one. Hence, if Bravo team has access to the competitor class, it can't access data from the conflicting class of the manufacturing client.

A, B, and C are incorrect. **A** is incorrect because Bell-LaPadula is used for enforcing confidentiality. **B** is incorrect because Biba is used for enforcing integrity. **C** is incorrect because Clark-Wilson is also used to enforce integrity. Hence, the most appropriate choice is the Brewer-Nash model.

5. Which of the following is not an SSO technology?

A. SAML

B. SecureAuth IdP

C. Kerberos

D. OpenIDConnect

B. SecureAuth IdP provides mobile device management and doesn't have anything to do with SSO.

A, C, and D are incorrect. **A** is incorrect because SAML is used to provide SSO. **C** is incorrect because Windows domains and Linux/UNIX realms do use Kerberos for SSO. **D** is incorrect because OpenIDConnect is an SSO technology.

6. A network administrator is considering placing a new firewall at the network perimeter. Which type of control is this?

A. Detective

B. Corrective

C. Preventive

D. Recovery

C. A firewall is classified as a preventive control, as it has the ability to prevent an attack by stopping offending traffic.

A, B, and D are incorrect. **A** is incorrect because a detective control would only be able to identify an attack when it's taking place or after it has happened. **B** is incorrect because a

corrective control would perform an activity to counter the attack. **D** is incorrect because a recovery control aims to restore a system after an incident has happened.

7. A company backup policy primarily uses magnetic tapes for backups. They also regularly back up data to the company cloud. Which type of control does the company cloud represent?

- A.** Detective
- B.** Deterrent
- C.** Recovery
- D.** Compensating

D. The cloud provider is being used in case the backup tape (which is primarily used as a backup media) fails, thus making it a compensating control.

A, B, and C are incorrect. **A** is incorrect because a detective control has the ability to provide event detection, which is not true in the case of a backup to the cloud. **B** is incorrect because a deterrent control aims to deter an attacker from performing a hostile action. **C** is incorrect because a recovery control aims to restore a system after an incident has happened. Please note that the question focuses on what type of control the company cloud is, making the correct answer a “compensating” one. However, if the question focused on what type of control a backup policy is, then the correct answer would be a recovery control.

8. Which of the following is not a security policy characteristic?

- A.** Mission statement
- B.** Enforcement section
- C.** Sanctions statement
- D.** Accountability statement

C. A security policy doesn't have a “sanctions statement.” Any related actions that can be pursued if the policy is not followed will be included in the “enforcement section.”

A, B, and D are incorrect. **A** is incorrect because a “mission statement” should be present and describe the overarching organizational vision. **B** is incorrect because an “enforcement section” would describe how a policy is going to be enforced and what steps might be taken in case personnel don't adhere to it. **D** is incorrect because an “accountability statement” would describe the specific role types (i.e., security staff) that have a responsibility to follow the policy.

9. What process would take place if an (ISC)² member doesn't comply with the (ISC)² code of ethics?

A. Any violators will be reported to the company the member is working for.

B. (ISC)² will revoke their certifications.

C. A review from a peer panel will take place.

D. They will get a one-off warning.

C. In the event that a member violates the (ISC)² code of ethics, a peer review panel is formed so the case can be examined and a recommendation on how to proceed be made.

A, B, and D are incorrect. **A** is incorrect because (ISC)² doesn't have the remit to reach out to the company a member might be working for. **B** is incorrect because certifications aren't revoked straight away. That might be a recommendation from an appropriate panel, but a ruling will come out after it convenes. **D** is incorrect because a code of ethics violation is a very serious matter, which doesn't warrant any warnings.

10. In order for a merchant to comply with PCI-DSS, a network scan is required to be completed by which of the following?

A. SAQ

B. ASV

C. QSA

D. ISA

B. An ASV (Approved Scanning Vendor) is an external security entity that uses authorized scanners to perform a scan. Merchants of all levels are required to use such a provider for any scans that are performed.

A, C, and D are incorrect. **A** is incorrect because an SAQ (Self-Assessment Questionnaire) can be used by an entity with a low number of credit card transactions to perform a compliance audit. **C** is incorrect because a QSA (Quality Security Assessor) is an external party that can perform an audit at a client site. **D** is incorrect because an ISA (Internal Security Assessor) is an internal company employee (who is ISA certified) able to perform specific client-site audits.

11. A company's single e-commerce server has gone down. The estimation of loss for each day it's not online is \$500, and a technician charges \$400 per day to repair it, with an initial estimation of repair time being one day. The monthly likelihood of such incidents is estimated at three. What is the ALE?

A. \$32,400

B. \$64,800

C. \$15,600

D. \$53,700

A. In order to calculate the ALE, the following formula will be required: $ALE = SLE \times ARO$. Furthermore, the SLE can be found by adding all related server costs, which are daily loss and repair cost. Based on that, the SLE will be $\$500 + (\$400 \times 1 \text{ day}) = \$500 + \$400 = \900 . The ARO is 3 monthly occurrences \times 12 months = 36 occurrences per year. $ALE = SLE \times ARO = \$900 \times 36 = \$32,400$.

B, C, and D are incorrect. All those figures are distractors, as none of them is correct.

12. An IPS is placed in front of the DMZ segment housing an FTP and a web server. Suddenly, the network administrator receives a report of traffic from a specific external IP address not being received by the web server. However, a packet capture right before the IPS suggests that web traffic is passing properly. What is the most likely reason for this?

A. Malformed web request.

B. Packet capture has not been set to capture web traffic properly.

C. Physical link failure.

D. IPS dropping the traffic in question.

D. Since the packet capture shows web traffic being received by the web server without traffic from a specific source passing through the IPS, it seems that the traffic in question is being identified as malicious and thus is blocked from traversing the rest of the network.

A, B, and C are incorrect. **A** is incorrect because although a malformed web request could have been sent, there's not enough information to suggest that this is the case and is the specific reason for dropping traffic from the IP in question. **B** is incorrect because there's no actual indication that web traffic is not being captured properly. In addition, the fact that only traffic from a specific IP is not seen by the web server suggests that packet capturing has been properly configured but some part of the traffic is just not passing through the IPS. **C** is incorrect because a physical link failure would result in no traffic reaching any of the servers in the DMZ, which is not the case here.

13. Which of the following is the most important consideration for a new SIEM implementation?

- A. Configure time synchronization.
- B. Adjust appropriate alert thresholds.
- C. Ensure all available logs are ingested.
- D. Provide security analyst access to individual log sources.

A. In order for proper analysis to take place, a SIEM needs accurate data synchronization. If an analyst tries to track an attacker's actions through various device logs that are in different time zones, the analysis can become unmanageable. In addition, it becomes very difficult for cooperating third parties to track any reports that are produced when data from different time zones is referenced in various parts of the report.

B, C, and D are incorrect. **B** is incorrect because although adjusting the alert thresholds is important to ensure that related teams are only engaged in the event of a valid security incident, this is something that can happen during the initial stages of the SIEM operation and usually happens gradually as more log sources are added and analyst teams become more familiar with its operation. It's difficult to ensure that thresholds are properly adjusted right at the beginning, as this takes substantial effort and time. **C** is incorrect because adding all available log sources to the SIEM (depending on their amount and volume) would result in overwhelming the tool as well as the analyst team supporting it. Organizations usually start by on-boarding the basic log sources that are required to achieve their goals and gradually adding more devices as they grow, while at the same time performing fine-tuning. **D** is incorrect because a SIEM works as a central logging and correlation platform. It allows analysts to view all the basic information to perform their tasks. Although having access to individual log sources can sometimes be beneficial (i.e., the analyst can review the raw logs that haven't been normalized within the SIEM, meaning all initial information is present), the whole point of the SIEM is to offer log storage and correlation without needing to access the end devices.

14. You are the head of security and your company is running a core finance application on a legacy Windows 2003 server. A recent vulnerability scan flagged the machine in question as being out of support (extended support for Windows 2003 ended in 2015). After contacting the application vendor, they alert you to the fact that the application is quite old and supported only on a Windows 2003 operating system. Which of these options do you choose?

- A. Upgrade the server to the latest version.
- B. Migrate the application to a new server.
- C. Identify a new application that can be safely used.
- D. Accept the risk.

D. Accepting the risk may sometimes be the only way to move forward (that's why risk acceptance is a valid risk strategy, although it is not recommended to follow it unless all other

avenues have been exhausted). As this is a core finance application, meaning its operation is crucial to the business, a way to continue using it must be identified. If upgrading the server (due to being out of support) or reverting to the application's vendor for a solution (i.e., if the application is only supported on a Windows 2003 system) are both nonviable options, accepting the risk seems the only viable solution.

A, B, and C are incorrect. **A** is incorrect because upgrading the server to the latest version (which is several versions after the current one that the server is running) might disrupt the application's normal operation or even disrupt the functionality of the legacy application (only supported in Windows 2003). **B** is incorrect because migrating the application to a new server (judging by the fact that the application is old and not supported anymore) might again result in not being able to use it further. **C** is incorrect because although identifying a new application (given the age of the current one being used) might take some time, it seems like a viable way moving forward. However, this will possibly take several weeks or months, as the new application will need to be identified, then reviewed from the finance team to ensure that it fulfills all related needs, and then tested and accepted. However, that won't solve the challenge at hand (i.e., the machine being flagged as out of support and a core application being run on it), so this can only be a viable long-term approach.

15. A company's client database is stored on a server so that online purchases can take place. The database needs to be accessible constantly even if the server goes down. Which of the following would be more suitable for that purpose?

A. Failover cluster

B. Load balancing

C. Full backup

D. Differential backup

A. A failover cluster is the most appropriate solution, as it can allow the database to be accessible if the server goes down. For that purpose, at least two servers are required, which can then be configured into a pair. One of them can be set to operate as the primary instance, while the other will be the secondary one. Once the secondary server detects any type of failure on the primary instance, it can assume its place as primary, thus guaranteeing minimal database downtime.

B, C, and D are incorrect. **B** is incorrect because load balancing is primarily used for performance improvement, not for fault tolerance. **C** is incorrect because a full backup will not guarantee that the database will be accessible all the time. If a server goes down, an administrator would need to intervene and manually restore that backup in addition to containing only accurate information up until the moment the backup was taken. If a full backup was taken a week ago, then after a server failure, the backup that is restored will contain data that is a week old, so all other transactions would have been lost. **D** is incorrect because a differential backup won't be able to ensure constant access to the server database, as any type of restoration would need to be

manually performed by an administrator, with substantial associated time if any failure occurs out of normal business hours, in addition to any transactions that are not contained in the backup that was lost.

16. Using a proprietary forensic tool for investigation relates to which of these reliability factors?

- A. Clarity
- B. Error rate
- C. Credibility
- D. Testability

D. Testability relates to how much a tool or technique has been tested before being used in an investigation. Using a proprietary tool will not allow a variety of people from different industries to test the tool, hence its testability would be very low.

A, B, and C are incorrect. **A** is incorrect because clarity refers to how easy it is to explain the tool or technique and its results to the court. **B** is incorrect because error rate refers to the tool's statistics regarding errors. **C** is incorrect because credibility refers to the individual handling a tool or using a method and how credible he or she is, as well as how reproducible the technique is by similar specialists.

17. An organization performs a full weekly backup and is considering combining that with a differential or incremental one. Why would a full/differential backup be preferred over a full/incremental backup?

- A. A full/differential backup takes less time to complete during the week.
- B. When restoring a system, recovery time is reduced.
- C. A full/incremental backup takes less time to complete during the week.
- D. Differential backups are smaller in size.

B. Using a full/differential backup results in reduced recovery time, as the differential backup captures any changes from the last full backup. That means that at any given moment, a total of two backups will be required to perform system restoration: the most recent full weekly backup and the most recent differential one.

A, C, and D are incorrect. **A** is incorrect because a full/differential backup takes more time to be completed during the week, as all changes from the moment the full weekly backup was taken need to be recorded. **C** is incorrect because a full/incremental backup does take less time to

complete during the week (i.e., if taken daily, it records daily database transactions, which takes a relatively short time). However, that won't be a reason for choosing a full/differential backup solution over a full/incremental one. It's rather a statement regarding how the full/differential backup method works. **D** is incorrect because differential backups are larger in size (not smaller) because they include all changes since the last full weekly backup was taken, which may require substantial space.

18. A company's RPO is set at three days, while the MAO is four days. Which of these values indicates a realistic RTO?

A. 3

B. 5

C. 7

D. 9

A. The time a crucial application or function can be unavailable without severe business impact refers to the MAO (maximum acceptable outage), which is stated as four days. The RTO is the time it takes (at most) for a system or service to be restored after an outage. Since the business can cope for up to four days (MAO set at four days), the RTO needs to be less than the MAO in order for the business to successfully recover and continue operations. From the provided options, only the first one (three days) allows for a viable RTO.

B, C, and D are incorrect, as these are distractors.

19. Consider two text files: The first one contains the sentence "this is an SSCP example" and the second one contains the sentence "this is another SSCP example with different text". The MD5 hash of the first file is 063ba1995095c95c2e2184837fd47067, and the MD5 hash of the second file is 063ba1995095c95c2e2184837fd47067. What is this an example of?

A. Work factor

B. Collusion

C. Rainbow tables

D. Collision

D. The two MD5 file hashes are identical; however, the content of the two files is different. This is known as a collision and relates to one of the basic principles of a hashing algorithm, which is to be collision free, meaning that two different files should never generate the same hash value. Even a slight change in a file that is used as input in a hashing algorithm should generate a different output.

A, B, and C are incorrect. **A** is incorrect because the work factor relates to cryptography and expresses the time and effort that it would take to decrypt an encrypted message. **B** is incorrect, as it refers to the act of one or more individuals or companies conspiring to commit fraud. Although the syntax might be similar to the correct answer, this term has actually nothing to do with hashing algorithms. **C** is incorrect, as rainbow tables contain precomputed hash values intended to provide a reverse lookup resource for hash values.

20. A key escrow is used for storing:

- A. Public keys
- B. Digital certificates
- C. WOT certificates
- D. Private keys

D. A key escrow is used for safely keeping private keys in order to be able to retrieve them in the event the company's private key is lost or damaged. Only specific people usually have permission to retrieve private keys from a key escrow facility due to the sensitivity of their nature.

A, B, and C are incorrect. **A** is incorrect because public keys are publicly available and are contained in an individual's digital certificate. **B** is incorrect because any digital certificates would be issued and managed by the certificate authority. Please note that a CA (certificate authority) may be external (i.e., Symantec or VeriSign) or internal to the organization (i.e., when a certificate is needed for a system that is used only internally). **C** is incorrect because there's no such term as WOT certificate. WOT (web of trust) is a decentralized trust model with the goal of allowing a system to establish authenticity between a public key and its owner.

21. Which of the following would be used for a new certificate registration request?

- A. WOT
- B. RA
- C. CRL
- D. OCSP

B. An RA (registration authority) can be used for an organization's certificate registration request. It is tasked with data verification, and after that is completed the request would be passed to the CA (certificate authority). Note that it is not mandatory to use an RA, but large companies do often use one. However, a request can also still just reach the CA directly.

A, C, and D are incorrect. **A** is incorrect because a WOT (web of trust) is a decentralized trust model with the goal of allowing a system to establish authenticity between a public key and its owner. **C** is incorrect because a CRL (certificate revocation list) is a list of revoked certificates that the CA maintains and is used by clients to check a certificate's validity. **D** is incorrect because OCSP (Online Certificate Status Protocol) is another method of checking if a certificate has been revoked. The difference is that instead of a revocation list being provided (like in the case of CRL), the client would just send a request to the OCSP server asking if a certificate of interest is valid or not, and the server would return a verdict.

22. You want to ensure your company's e-mails are encrypted to maintain communication confidentiality. Which of these options would you choose for that purpose?

A. DKIM

B. SPF

C. S/MIME

D. DMARC

C. S/MIME (Secure/Multipurpose Internet Mail Extensions) is a protocol that allows one to encrypt and digitally sign e-mails. This works by using the recipient's public key to encrypt an outgoing message. After that has been received, only the recipient can decrypt it (as that would require the recipient's private key), which ensures the security of that transaction. Furthermore, as S/MIME also supports digital signatures, a sender can sign the message that is sent so that the recipient knows it originated from a valid entity.

A, B, and D are incorrect. **A** is incorrect because DKIM (Domain Keys Identified Mail) is a method of authenticating e-mail and identifying any attempt to spoof a valid sender. It doesn't support e-mail encryption; hence, it wouldn't be used to secure e-mail communications. **B** is incorrect because SPF (Sender Policy Framework) is used to authenticate e-mails and identify e-mail spoofing, not to encrypt them. Assume that you receive an e-mail originating from your bank (which is called amazing bank), sourced from e-mail account security@amazing-bank.com, asking you to change your password to a new value that is provided in the e-mail. If that e-mail is not actually sourced from "amazing bank" but is just an attempt by an attacker to spoof the source header so you are tempted to change your password to what is provided, that would be identified as forged e-mail. **D** is incorrect because DMARC (Domain-based Message Authentication, Reporting, and Conformance) is a protocol that allows organizations to specify if SPF or DKIM (or neither of them) is used and what action to take if validation fails (i.e., mark the message as spam) by publishing the relevant policy through DNS records. However, this protocol wouldn't be used for encrypting messages.

23. Which topology is the least resilient?

A. Star

- B.** Tree
- C.** Mesh
- D.** Bus

D. When using a bus topology, any issue with one of the connected devices will bring the whole network to a halt. A bus topology uses T connectors for connectivity. That means that one end of a computer is connected to the bus and the other is connected to the next computer. If one of the computers or T connectors experiences an issue, the whole network stops working. In addition, if a lot of machines are connected to that network, identifying the issue may require vast amounts of time. This means the network will not be usable until the issue is resolved.

A, B, and C are incorrect. **A** is incorrect because a star topology commonly uses a switch to allow the various machines to connect, which means that any issue with one machine doesn't affect the network's overall operation. **B** is incorrect because a tree topology uses a combination of a bus and a star, where multiple star networks connect over a common bus. If the bus experiences an issue, all the star networks are disconnected until the issue is resolved. **C** is incorrect because in a mesh topology, all machines are individually connected with each other, meaning that any issues with one connection don't affect the operation of the network.

24. Which of the following would be performed by NAC?

- A.** Authenticate a remote client.
- B.** Create audit logs about the actions a machine performed while on the network.
- C.** Check a client machine against an organizational health policy.
- D.** Allow a client machine to issue specific commands.

C. A NAC (Network Access Control) device is primarily used to check if a client that is requesting to connect to the corporate network fulfills the minimum requirements to be allowed to join the network. If that is the case, a NAC device will allow the client to connect. If not, the client will usually be restricted to a quarantined network until it achieves the minimum company security requirements.

A, B, and D are incorrect. **A** is incorrect because the authentication of a remote client would happen by an authentication server (i.e., RADIUS or TACACS+). **B** is incorrect because a NAC device wouldn't be able to perform auditing of a network client. That is a task that needs to be performed by the authentication server (i.e., a server supporting AAA, like RADIUS, TACACS+, or Diameter). **D** is incorrect because reviewing a client's permissions to issue specific commands takes place during the authorization process (which typically happens when a client logs on using an AAA server).

25. You have a network composed of a hundred endpoints, ten firewalls, five routers, five IPS sensors, and three proxies across two geographical regions. Data from all of those devices is being forwarded to a SIEM platform. What would be the most crucial element for any type of security investigation you perform?

- A. Log integrity
- B. Verbose logging across all devices
- C. Appropriate alert thresholds
- D. Time synchronization

D. For any effective investigation to take place, appropriate time synchronization needs to be present for all the devices across the network. Even with a handful of devices, like five to six, analysis can become very complicated when those devices are not in a common time zone (i.e., UTC). Imagine you receive an alert from the SIEM with a time stamp of 14:00 UTC-8 (which might be the time zone of the tool or the location where it is installed) and you log in to individual devices to investigate further. Some of them are in UTC+4 and others in UTC-2. How are you going to be able to track a specific attacker's activities across all those time zones?

A, B, and C are incorrect. **A** is incorrect because log integrity is crucial, but in this scenario the logs are being forwarded to a SIEM tool, which stores copies of them, so a good level of integrity is present. **B** is incorrect because verbose logging across various devices might be too intense (depending on the amount of traffic and device type) and cause issues with them, while it also consumes a lot of network bandwidth. SIEM platforms use EPS (events per second), and as that increases, the cost of the tool grows exponentially—not to mention that a more robust architecture will be required in order to cope with the excessive traffic being forwarded to the SIEM tool. However, if there's no time synchronization, even with an abundance of data, it will be very difficult to identify crucial artifacts of any attack. **C** is incorrect because for any SIEM to work effectively, appropriate alert thresholds need to be configured; however, those can be gradually defined as the tool is being used. Lack of time synchronization will pose a greater challenge when investigating malicious activities across logs of various devices in different regions (and various time zones).

26. Which of these media types is most susceptible to electromagnetic interference?

- A. Shielded twisted pair
- B. UTP
- C. Optical fiber
- D. STP cat 6

B. UTP (unshielded twisted pair) is a cable type that is very commonly used in computer networks (especially for short to medium distances), CCTV circuits, and telephone systems, as it's very cost effective. However, its greatest disadvantage is that it is susceptible to both electromagnetic and radio frequency interference.

A, C, and D are incorrect. **A** is incorrect because shielded twisted pair (STP) uses additional shielding to reduce the effect of electromagnetic interference. Although it is more expensive as a result, it is often used for outdoors implementations and signal transmissions over large distances. **C** is incorrect because optical fiber cabling uses light to transmit data and as such is not susceptible to electromagnetic interference. It's often used in high-speed networks and can allow high data transmission speeds over large distances, while also offering enhanced security in comparison to traditional media, like copper cables. **D** is incorrect because STP Cat 6 (also known as category 6) is a type of STP cable that is less prone to electromagnetic interference than common UTP cables.

[27.](#) How can you best protect against VM escape?

A. Patch the hypervisor software.

B. Patch the guest OS.

C. Use a next-generation firewall on the host.

D. Use a host antivirus.

A. VM escape is a type of exploit that can take place when an attacker manages to break out of the context of a virtual machine and interact directly with the host. This is made possible due to a potential vulnerability within the hypervisor (the software that provisions the guest OS and controls how that interacts with the host). In order to ensure the risk is mitigated as much as possible, the hypervisor should always run the most up-to-date versions with the latest patches so that attackers can't use any known exploits to perform VM escape.

B, C, and D are incorrect. **B** is incorrect because although patching the guest OS is quite important, it's the vulnerability within the hypervisor that ultimately allows an attacker to perform a VM escape. Please note that patching the guest OS will aid a lot in avoiding any risks from vulnerabilities that can lead to that machine being compromised. In a typical VM escape scenario, even if the hypervisor is vulnerable, the attacker will start by first compromising the guest OS, and after that is achieved, a hypervisor exploit will follow. However, if the hypervisor is sufficiently patched, then the only thing an attacker can achieve is to compromise the guest OS (assuming the attacker is not using a zero-day hypervisor exploit). **C** is incorrect because using a next-generation firewall on the host will most likely not achieve anything, as the attacker would perform this attack by compromising the guest OS and work through that to exploit the hypervisor. **D** is incorrect because using a host AV will not necessarily protect against this attack, as the attacker will initially compromise the guest OS and then exploit the hypervisor. It

would be more beneficial to install an AV on the guest OS in the hope that the initial attack will be detected and stopped before the attacker manages to exploit the hypervisor.

28. Which of these attacks allows an attacker access to forbidden file system locations and often targets a system's password file?

- A. XSS
- B. Directory traversal
- C. Buffer overflow
- D. CSRF

B. A directory traversal attack is a type of web attack in which an attacker tries to access forbidden file system parts, that is, on a web server. The attacker would most commonly target the password file so user credentials can be harvested, but the term is used to describe access to any directory where access is not allowed.

A, C, and D are incorrect. **A** is incorrect because an XSS (cross-site scripting) attack would inject malicious code (i.e., JavaScript or HTML) to a web page, with the goal of bypassing existing access control and executing the malicious code. **C** is incorrect because a buffer overflow attack works by allowing the attacker to write into specific memory areas that shouldn't be normally available (i.e., an application using a buffer to hold data of a specific size, which is being overrun when more data is supplied and due to the application being vulnerable to a buffer overflow attack, it tries to store that information in the predefined buffer, thus resulting in writing to other memory areas). **D** is incorrect because a CSRF (cross-site request forgery) attack aims to utilize the user's browser to send a malicious request to a website without any indication to the user. Assume you are logged in to an online gaming website and want to transfer some of your credits to a friend. A CSRF exploit can be used while you are logged in to the website to have credits sent to the attacker's account instead.

29. Which of the following tries to disable or attack the AV software?

- A. Stealth virus
- B. Multipartite virus
- C. Polymorphic virus
- D. Retrovirus

D. A retrovirus is a type of virus that tries to attack a device's antivirus program. It may attempt to warn the victim that supposedly unsafe software has been found on the machine and then uninstall legitimate AV applications. Sometimes it has the ability to modify the virus signature database and corrupt or delete it altogether.

A, B, and C are incorrect. **A** is incorrect because a stealth virus doesn't attack the AV application but actually tries to evade detection from it. Sometimes it does that by changing the offending file's name or size to make it seem like the infected file is the actual legitimate host file. **B** is incorrect, as a multipartite virus works by infecting various parts of a target system with the intent of making removal very difficult. **C** is incorrect because a polymorphic virus can change a malicious file in order to make signature detection more difficult; however, the replication and payload mechanisms remain unaltered.

30. Which of these types of data residing in a cloud server would be protected by HIPAA?

A. Patient first name

B. Last four digits of patient credit card

C. Patient year of birth

D. Healthcare provider

D. HIPAA (Health Insurance Portability and Accountability Act) relates to protection of PHI (protected health information), which is any information about an individual's health care, status, or related treatment provisioning. As such, an individual's healthcare provider is classified as PHI and should remain securely protected under HIPAA.

A, B, and C are incorrect. **A** is incorrect because a patient's first name is not distinctive to the individual and doesn't require special handling. However, please note that the full name warrants protection under HIPAA. **B** is incorrect, as the last four digits of a credit card don't provide enough information to uniquely identify the card holder. In addition, even a full credit card number doesn't constitute PHI. That information falls under PII (personally identifiable information). **C** is incorrect, as a patient's year of birth is not classified as PHI since it doesn't provide enough information to uniquely identify an individual.

The End

نقاط مهمة حول الاختبار

من هنا يجب التركيز

سأذكر نقاط مهمة لمنهج SSCP قد تسهل عليكم فهم الاسئلة في الاختبار

حتى اذا لم تتخذ قرار لتجربة الاختبار **يجب** على كل متخصص في مجال أمن المعلومات أو الأمن السيبراني معرفتها سواء كان ممارس/مراقب/ محلل :-

1- معرفة الفرق بين DAC/MAC/ABAC/RBAC
لم يتم التعمق فيها بمادة سكيورتي بلس في دومين Access Control

شرحها لكم بشكل مفصل ومراح تتسونها ان شاء الله

شرح انواع التحكم بالصلاحيات MAC - DAC - N-DAC

2- ان تكون لغتك ممتازة

3- التحلي بالصبر والعزيمة لانها ماده تحتاج الى تركيز وفهم وليس الحفظ

4- بالنسبة لبرامج الحماية من الفايروسات والبرمجيات الضارة يجب معرفة الفرق بين Knowledge - and Behavior-based Detection

5- الفرق بين Anomaly based Vs Huristic based

باختصار:

Knowledge based /Signature based/pattern based

نفس الهدف وهي تكتشف الفايروسات عن طريق تحديث تواقيع برنامج الفايروسات ويتم اضافة هاشات جديدة.

Behavioral detection based

باختصار كشف بعض ثغرات الزيرو دي يتم عن طريق السلوك للملف لانه يقوم بعملية Record/Learning وينشئ بروفایل ويعمل مقارنه في حال اكتشاف سلوك غريب او مشبوه

ولكن عيبه انه يعطي انذارات كثيرة false positive

Behavioral +Anomaly خليط بين Huristic based

6-معرفة الفرق بين incident Response و IRT+ DRP

7-معرفة Risk management framework 800-37

8-الفرق بين النسخة الاحتياطية عندما تكون on-site و off-site

9-التعامل مع الادلة في الجرائم الإلكترونية ولماذا يتم اخذ نسخة احتياطية من الادلة وماهي اول خطوة بعد الحصول على الهاردسك المتعلق بالجريمة؟

10-المحافظة Preserving على الادلة من اهم الخطوات

11- ماذا تعرف عن threat vector

threat vector is a path that an attacker might take to take advantage of a vulnerability and do harm. For example, fire is a threat that could do harm to your physical environment or resources

12-انواع privilege escalation

13-الفرق بين quantitative and qualitative

15- الفرق بين Prudent Man, Due Diligence, and Due Care

16-التعامل مع المخاطر risk control

17-هل DRP جزء من BCP او العكس وماذا يندرج تحت BCP ?

18-الفرق بين False positive , True positive, False negative, True negative

19-أهمية البولسي في المنظمة او الشركة. مثلا انها تنص على توعية الموظفين ومعاينة كل من يخالفها

20-ماهي الفائدة من BCP ولماذا نعمل له Test كل فتره وماهي المدة اللازمة لاجراء الفحص لضمان الخطة؟كل سنة او عند التغيير؟

21-الفرق بين hot site و multi-site

22-تذكروا دائم ان hot site هو نسخة من موقعك الحالي ويحتوي على العتاد والاجهزة

23- hot site مكلف جداً ولكنه فعال في حال الانتقال اليه بسبب تعطل الموقع الرئيسي

24- worm site يحتوي على بعض الاجهزة والتكليف والكيابل ولايحتوي على السيرفرات واتصال النت متوسط التكلفة.

cold site-25 صاله فارغة تحتوي على راکات وتکييف فقط وغير مکلفة

26-hot site 100% - warm site 50% - cold site 0%

27- MAC is a static access control method. Resources are classified using labels. Clearance labels are assigned to users who need to work with resources

MAC supports a security requirement of confidentiality more so than the others.

MAC:makes decisions based upon "labeling and then permissions"

MAC:All access permissions are controlled solely by the system administrator.

28-DAC supports the security requirement of availability more so than the others.

DAC model the data owner can decide who has access (if he has that permission on the data) and add or remove people from the list.

29-Rule-based access control is based on rules to deny or allow access to resources

Examples of Rules Based Access Control include situations such as permitting access for an account or group to a network connection at certain hours of the day or days of the week.

Rules Based Access Control, access is allowed or denied to resource objects based on a set of rules defined by a system administrator.

With router ACLs we determine which IPs or port numbers are allowed through the router, and this is done using rules.

30- ماهي اول خطوة في حالة الاستجابة للحدث بناء على NIST 800-61 بعد التأكد من اصابة الجهاز بالفايروس

Eradication

Recovery

Containment /isolated

Preparation

سؤال بصيغة اخرى

-Once a security incident has been contained, the next step is:

Eradication

Recovery

Containment /isolated

Preparation

ابحثوا عن الجواب من هنا

After an incident has been contained, eradication may be necessary to eliminate components of the incident, such as deleting malware and disabling breached user accounts,

لازم تعرف معنى هذه

المصطلحات Impersonation, maintaining, countermeasures, abuse, Phase, Compromised

31- الفرق بين DMZ/Extranet

32- الفرق بين SOD و Dual man control

33- الفرق بين Countermeasure & Safeguard

Safeguard—A built-in **proactive** security control implemented to provide protection against threats

Countermeasure—An added-on **reactive** security controls

34-Vulnerability-A system weakness

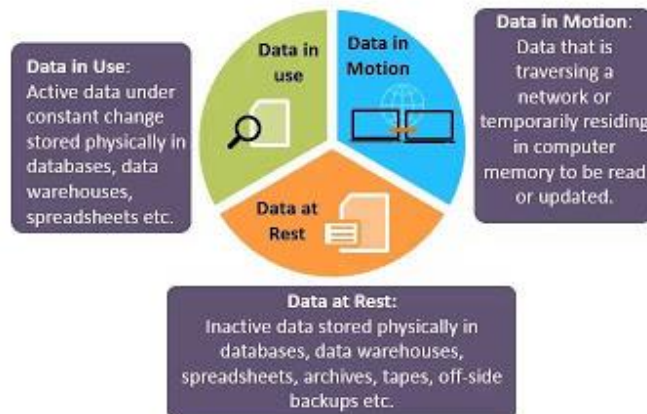
35-ماهو الفرق بين Adverse events/ incident/ event

data states-36

Cloud System Security – Cryptography (Asset Security)

• Data States

- In Process (**Data in Use**)
- Passive (**Data at Rest**)
- In Transit (**Data in Motion**)



37-يجب فهم بروتوكول TLS بشكل مفصل وكيف تتم عملية تبادل المفاتيح وتشفير المفتاح الخاص ولماذا يستخدم نوعين من asymmetric + symmetric ونوع الخوارزمية المستخدمة في symmetric

is the change in normal behavior of the port due Deviation Network-38 to port scanning

business continuity plan (BCP) includes multiple elements, such as A39 .business impact analysis (BIA) and a disaster recovery plan (DRP) a

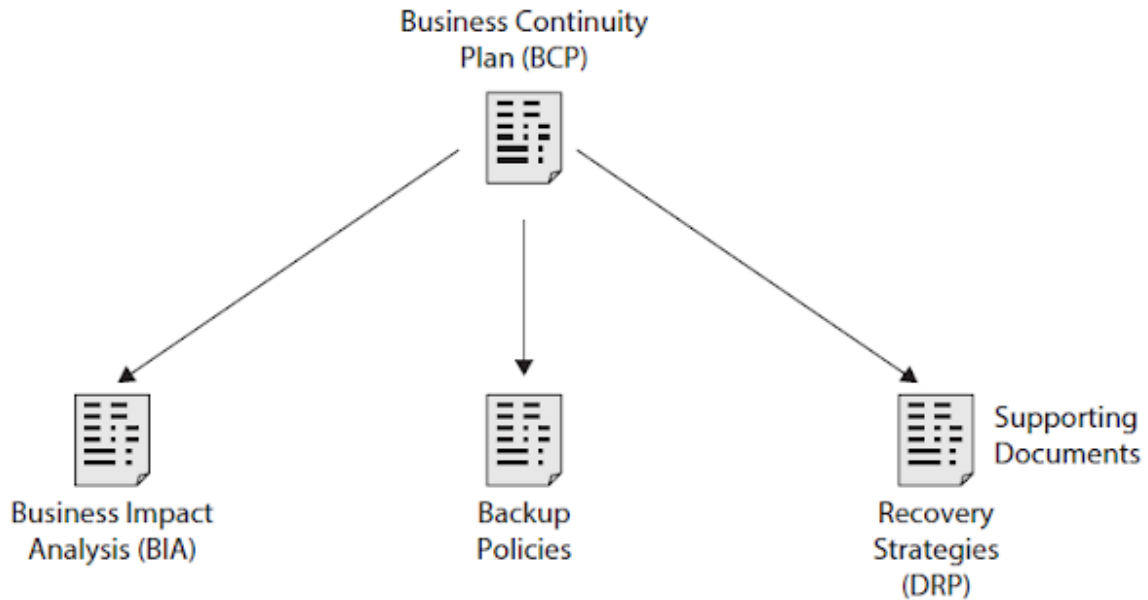


Figure 12-4 BCP and supporting elements

?What is the purpose of a BIA

- A. To identify recovery plans**
- B. To drive the creation of the BCP**
- C. To test recovery plans**
- D. To identify critical business functions**

a BIA aids in the identification of critical organizational functions and helps determine recovery time objectives, which are both prerequisites for developing the business continuity strategy and related DRP

? stakeholders من هم -40

Which type of network segment is created by a switch-41

VLAN

Subnet

Domain

Community

سؤال سهل جداً

**War dialing is the concept of dialing phone numbers seeking to -42
discover answering modems**

**War driving is the activity where hackers travel around an area in -43
search for wireless network signals**

-44

in IT security manager is struggling to keep the organization's computers An installed onto working order. He is testing updates and configuring them to be to various systems systems and making tweaks to the configuration settings discovers systems which do not as business tasks require. However, he often using out-of-date settings. This may have the necessary updates or which are disconnected from the company network when be caused by systems being .used for special offline projects taken into the field or when

What technology should the IT security manager implement to help handle ?this complex issue

IEEE 802.1x

IEEE 802.x1

incident response

NAC

tricky سؤال

-45

an attacker gains remote control access over a system, they want to Once-45 service retain this illicit access. Some attackers will block the system update needed to maintain to prevent new updates from fixing vulnerabilities that are prevent such a remote control over the compromised system. To resources on the compromised system from allowing the attacker to access network, what security mechanism should be implemented

Complex password authentication
Intrusion detection system
Network access control
A Web security gateway

إذا تعاملت مع ISE او clear pass ستعرف الاجابة

46-Eavesdropping vs Interception
47-Delphi technique - anonymously
48-NIST Risk Assessment Process
49-Defense in Depth
50-Administering the Identity Management Life cycle

51-HTTPS ضروري ومهم تعرفون تفاصيل مرحلة اتصال-

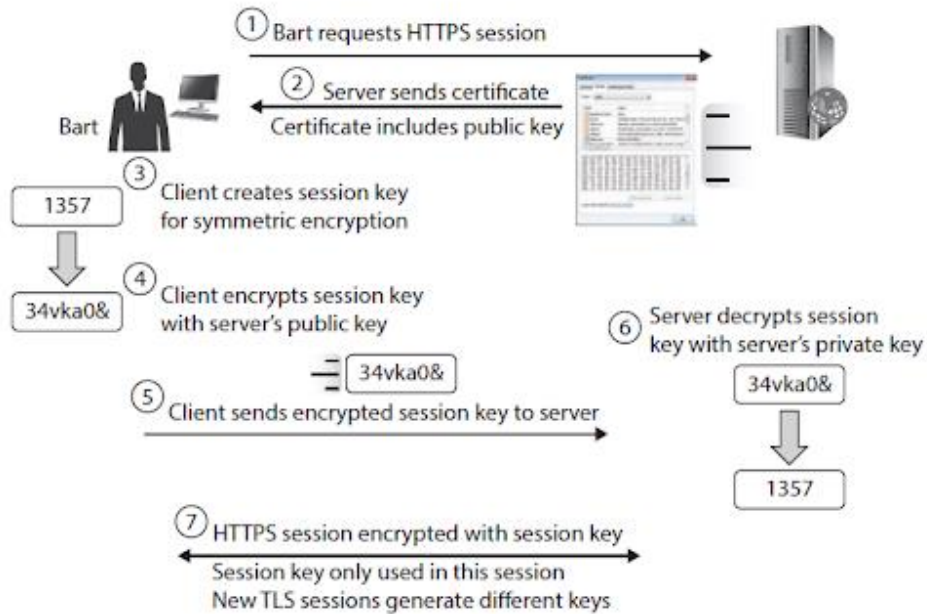


Figure 14-6 Creating a secure HTTPS session with TLS

Transport Layer Security (TLS) is the designated **replacement** for Secure Sockets Layer (SSL), and is based on SSL
SSL is **vulnerable** to **POODLE** attacks

52- ماهي أنواع البولييسي

53- PCI DSS industry standard يجب على البنوك ان تلتزم ب

54- يجب على المختص في هذا المجال معرفة الفرق بين-
Guidelines,Procedures,Standards,Policy

55- للمحافظة على بيانات المرضى HIPAA يجب على القطاع الصحي الإلتزام بقانون

Individuals should also receive specialized security training in compliance-based areas such as HIPAA medical information regulations and patient information privacy

56-Active FTP vs Passive FTP

57-explicit deny / Implicit deny

58-clipping level / threshold

59- acceptance risk Action معناه لم يتم اتخاذ

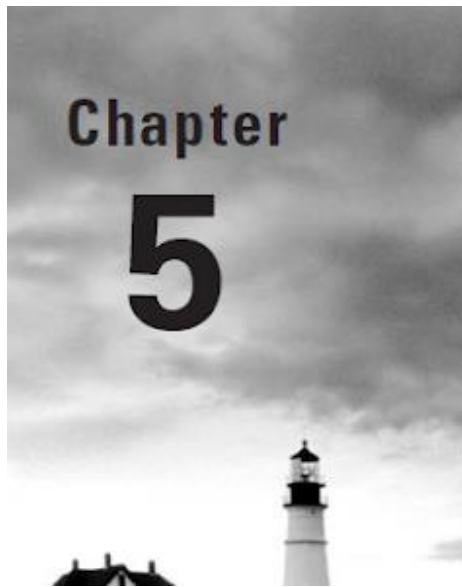
60-BCP الحدث بعد يكون

61-DRP الحدث قبل يكون

62-risk register=risk logs+date+action

63-normalization/correlation الفرق بين-

64-تشابتر 5 في الكتاب الاول ان شاء الله راح يعبجكم بصراحه صعب شوي- شدوا حيلكم-👁️



Domain 3: Risk Identification, Monitoring, and Analysis

SSCP EXAM OBJECTIVES COVERED IN THIS CHAPTER:

تشابتر 7 ايضاً يحتاج لفهم وتركيز وتلخيص على ورقة وفهم كل خطوة-65

Risk, Response, and Recovery

In this chapter, you will learn about

- The definition of risk, threat sources, threat events, and vulnerabilities
- Primary methods of managing risk
- The definition of residual risk
- Risk assessment steps
- Differences between quantitative and qualitative analysis
- Steps in incident response

66-Analysis vs. Assessment

67-threat intelligence

68- بما انك ممارس في امن المعلومات في حال عندك جهاز مصاب في الشبكة اكيد اول خطوة تبلك البورت وتغزل الجهاز عن الشبكة ، كيف تحدد ان الجهاز يحتاج؟

eradication or reimaging

يعمل على الطبقة السابعة والبعض عتقد على الرابعه او الخامسة اليكم التفصيل TLS الاغلب يعتقد ان بروتوكول-69

RFCs 2246, 4346, 5246, and 6176 formally define TLS. The following list describes operations performed by TLS on **different layers**:

- **Transport layer** TLS creates the initial connection on the Transport layer (layer 4). With the name Transport Layer Security, it's implied that TLS operates on the Transport layer.
- **Session layer** The Handshake Protocol within TLS operates on the Session layer (layer 5) when negotiating details for the TLS session. This includes details such as what specific encryption algorithm to use.
- **Application layer** The TLS Record Protocol within TLS operates on the Application layer (layer 7). It provides confidentiality with an encryption algorithm and integrity with a hashing algorithm.