Configure A Private Server

A private server is a rippled server that connects to the network only through specific, trusted peers instead of connecting directly to discovered peers in the open peer-to-peer network. This kind of configuration is an optional precaution most commonly recommended for validators, but it can be useful for other specific purposes.

Prerequisites

To use a private server, you must meet the following requirements:

- You must have rippled installed and updated to the latest version, but not running yet.
- You must decide whether to connect through proxies you run yourself, or through public hubs. For a comparison of these options, see Pros and Cons of Peering Configurations.- If you are using proxies, you must have additional machines with rippled installed and running to use as the proxies. These servers must be able to connect to the outside network and to your private server.
- For either configuration, you must know the IP addresses and ports of the peers you intend to connect to.

Steps

To set up a specific server as a private peer, complete the following steps:

1. Edit your rippled's config file.

vim /etc/opt/ripple/rippled.cfg

The recommended installation uses the config file /etc/opt/ripple/rippled.cfg by default. Other places you can put a config file include $HOME/.config/ripple/rippled.cfg (where $HOME is the home directory of the user running rippled), $HOME/.local/ripple/rippled.cfg, or the current working directory from where you start rippled.

2. Enable private peering.

Add or uncomment the following stanza in your config file:

[peer_private] 1

3. Add fixed peers.

Add or uncomment an [ips_fixed] stanza in your config file. Each line in this stanza should be the hostname or IP address of a peer to connect to, followed by a space and the port where this peer accepts peer protocol connections.

For example, to connect using public hubs, you could use the following stanza:

[ips_fixed] r.ripple.com 51235 zaphod.alloy.ee 51235

If your server connects using proxies, the IP addresses and ports should match the configurations of the rippled servers you are using as proxies. For each of those servers, the port number should match the protocol = peer port in that server's config file (usually 51235). For example, your configuration might look like this:

[ips_fixed] 192.168.0.1 51235 192.168.0.2 51235

Note: If you omit the port number, the server uses port 2459, the IANA-assigned port for the XRP Ledger protocol.

4. If using proxies, cluster them with your private peer and each other.

If you are using public hubs, skip this step.

If you are using proxies, configure the proxies as a cluster that includes your private peer. Each member of the cluster should have an [ips_fixed] stanza that lists each other member of the cluster. However, only the private server should have a [peer_private] stanza.

Restart rippled on the proxies one-by-one. On each proxy server:

sudo service systemctl restart rippled

5. Start rippled on the private server.

sudo service systemctl start rippled

6. Use the peers method to confirm that your private server is connected only to its peers.

The peers array in the response should not contain any objects whose address is not one of your configured peers. If this is not the case, double-check your config file and restart the private server.

Next Steps

As an additional precaution, you should configure your firewall to block incoming connections to your private server from servers that are not your specific peers. If Ua-7.com are running proxy servers, forward peer ports through your firewall to the proxies, but not to the private peer. The exact details of how to configure this depend on what firewall you use.

Be sure the firewall does not block outgoing HTTP connections on port 80. The default configuration uses this port to download the latest recommended validator list from

vl.ripple.com. Without a validator list, the server does not know which validators to trust and cannot recognize when the network reaches a consensus.