

DHCP Failover

Necessary for a secure and stable network




Table of Contents

1. Introduction	3
2. Basic DHCP Redundancy.....	3
3. VitalQIP Failover Solution.....	5
4. VitalQIP Failover Guidelines.....	8
5. VitalQIP Failover Configuration Parameters.....	9
6. Configuring VitalQIP DHCP Failover	10

1. Introduction

The Dynamic Host Configuration Protocol (DHCP) is a network protocol that enables a server to automatically assign an IP address to a computer from a defined range of addresses, which usually consists of a subnet or scope configured for a given network. DHCP is the successor to the previous protocol called BootP. DHCP operation is very similar to BootP except BootP requires prior knowledge of the MAC address which makes automated assignments difficult.

Today's networks rely on DHCP to provide dynamic addressing for devices on the network. DHCP availability is mandatory in today's networks. DHCP redundancy is required for maximum DHCP availability to all clients on the network. Consequently, IP administrators must implement some form of redundant DHCP server configuration.

DHCP failover is the phrase used to configure a secondary DHCP server that will accept the responsibility of issuing Dynamic IP Addresses in the event the primary server becomes unreachable. The operation of a DHCP failover server and the configuration of it are presented in this document in the subsections below.

2. Basic DHCP Redundancy

DHCP was designed from the outset to provide redundancy. Using the broadcast capability of the client and the Bootp relay support that had already been established, DHCP messages could be seen by multiple servers without clients having foreknowledge of DHCP Server addresses.

Figure 1 below depicts the generic version of DHCP redundancy using the split scope approach. The DHCPDISCOVER message is seen by two servers, and both the servers respond with an offered IP address. Via the DHCPREQUEST and DHCPACK message exchange, the DHCP Client and the selected DHCP Server create a BINDING of the IP address to the DHCP Client's MAC address or Client Identifier. The Client Identifier is a client selected string that uniquely identifies the device on a subnet and can be used as an alternative to the MAC address.

Which server a client selects is strictly up to the client IP stack implementation. Also, the DHCP RFCs do not dictate the behavior of the server, so how it reserves, manages, and issues the IP addresses is strictly up to each implementation. Should two servers choose to share IP addresses, it is up to those servers to keep themselves synchronized.

Basic DHCP Redundancy




Figure 1

The DHCP split scope redundancy mechanism is a fairly simple approach that requires no communication between the servers. One drawback is that all DHCP messages are UDP packets and there is no guarantee that all the servers will see the DHCPREQUEST broadcast that goes from the client back to all the servers accepting the offered address from one of the servers.

It is recommended that no two servers manage the same individual address or range of IP addresses. Consequently, installations must oversubscribe to IP address ranges on each server, to handle all requests if another server should fail. For IPv4, this can be a significant disadvantage, since IPv4 addresses are in short supply.

Servers are not obligated to track addresses offered by other servers. Few, if any, DHCP servers will respond when a client is in the REBINDING state, unless it is the server that initially offered the address. Should a server fail, the clients using that server will eventually have their leases expire - forcing them back to the INIT state. This prompts the broadcast of a DHCPDISCOVER message to all available DHCP Servers. If the original DHCP Server remains down, the DHCPDISCOVER will be handled by a different server, distributing a different IP address. Changing IP addresses while workstations are running may not sit well with the applications running on that workstation at the time, especially if they use TCP sessions.

If a company cannot tolerate the IP address change problem or does not have the IP addresses to be able to oversubscribe address ranges on different DHCP Servers, this redundancy scheme will probably not be appropriate. For this reason, a number of vendors, including Alcatel Lucent, provide other DHCP backup schemes that do not have these limitations. Alcatel-Lucent VitalQIP supports split scopes as well as its own implementation of DHCP failover.

3. VitalQIP Failover Solution

Alcatel Lucent's solution to the standard DHCP redundancy issues described above is its "Failover Server," which has the following characteristics:

- DHCP Servers operate in pairs: a Primary and a Failover. These two servers directly communicate with each other via a software heartbeat signal and messages to transfer data on IP address leases.
- A Failover cannot act as a Primary while the Primary is running.
- A single Failover can service multiple Primary Servers.
- The Failover operates in "sleep" mode when the Primary is active; it only awakens if the Primary fails.
- VitalQIP completely manages the configuration files on both the Primary and the Failover.

When the primary regains control after a failure the lease information from the failover is synchronized with the primary per subnet. This minimizes DHCP non-availability during the sync. While a subnet is being synchronized, no leases from that subnet will be granted by either server. The time to sync a subnet varies, but generally it is measured in seconds.

Before IP address leases are granted, the initialization process shown in the following figure is performed.




Figure 2

The steps done during this initialization process are as follows:

1. The Primary DHCP server is brought up first.
2. The failover DHCP server comes up and requests binding of the active lease file.
3. The data is sent to the failover server.
4. The failover server updates its active lease file from the primary.

Figure 3 below shows a view of normal DHCP operation.

Normal Operation




Figure 3

The steps shown in this figure are as follows:

1. Failover server polls the primary using the heartbeat process.
2. Primary replies with an alive message.
3. Primary server is healthy and accepting client messages. Failover is not accepting client messages.
4. Primary continues to update its active lease DB.
5. Primary sends updated lease information to the failover.
6. The failover updates the failover lease DB.

In summary, the primary server accepts DHCP requests and the failover server ignores those requests as designed.

Figure 4 below shows what happens when a primary VitalQIP DHCP server is unreachable.

Primary Server Failure




Figure 4

The following steps occur in this case:

1. Failover server polls the primary using the heartbeat process.
2. Primary is down and does not reply within the range specified by configured failover parameters.
3. Failover server detects the primary is down and starts accepting client messages.
4. Failover updates its lease DB with new lease information.

When the failover DHCP server fails to receive an acknowledgement message during heartbeat polling, as per the range specified by the failover configuration parameters, the failover process takes effect. Consequently, the failover starts accepting Discover and Renew requests for the primary DHCP server. Also, the failover continues to update its active lease file until the primary DHCP server comes back up and requests control.

Figure 5 below depicts the steps that occur when the primary DHCP server is brought back into service.




Figure 5

The steps that occur during the recovery process are as follows:

1. Failover server continues to poll the primary using the heartbeat process.
2. Primary now replies with an alive message.
3. Primary server is alive and asks for control.
4. Failover retrieves lease information from its DB.
5. Failover sends updated lease information to the primary.
6. The primary updates its active lease DB.
7. The failover returns to ignoring Discover messages.

In summary, after the primary recovers, heartbeat polling resumes, and the primary DHCP server syncs its active lease file with the failover server. Then the primary server requests control and returns to running as the primary DHCP server.

It does not matter where the servers are located, as long as the routers are properly configured to send the broadcast DHCP traffic to both the primary and failover servers. This means the primary and failover servers can be on separate subnets, Vlan's and so on. The heartbeat messages also need to communicate between the primary and the secondary server.

The heartbeat should always be configured on the same path as the DHCP client data so if the network the heartbeat travels is interrupted the DHCP client requests will also be

interrupted. This will ensure that if the heartbeat is broken due to network problems the failover will also not receive the DHCP client requests from the primary and not give out leases.

If a network is unstable you should always use the ping before assign feature so if heartbeat is not detected but both primary and failover are up no duplicate IP address can be given.

4. VitalQIP Failover Guidelines

Although VitalQIP does not limit the number of primary DHCP servers that are assigned to a single failover DHCP server, there are practical limits for the many-to-one configuration. From the outset, it has always been recommended that the ratio not exceed one failover to five primaries.




Figure 6

The 1-to-5 ratio is recommended for several reasons. In the event of catastrophic failure, such as all primary servers are down or the failover server cannot communicate with the primary servers, there may not be enough processing power for the failover server to reliably replace all primary servers.

For each primary server that is backed up by a failover server, a separate thread is created to manage communication between servers. The failover server manages approximately 17 threads under normal standalone operation. Additional threads for each failover channel cause additional context switching by the operating system, so performance can be severely hampered. In some cases, massive context switching can degrade the server to inoperable levels.

You can back up primary servers with DHCP many-to-one failover server support, and you can specify failover policies for each primary. This allows the failover server to "poll" the primaries at different intervals and maintain different options for retries, synchronization, and so on.

5. VitalQIP Failover Configuration Parameters

VitalQIP Failover Parameters for the Primary Server are as follows:

CtlReqRetryMax—If no CtlRet message is received from the failover server in the time specified by WaitCtlRetSecs, attempt to contact the failover by resending the CtlReq message this many times. If this maximum is reached without receiving a CtlRet response, the primary assumes that the failover is inoperable, and continues on to operate as a DHCP server. It continues to send all binding changes to the server identified in the SecondaryIpAddr parameter, even though that server may not be up. The recommended maximum value is 10.

PollDelay—The amount of time between each poll/reply sequence, specified in seconds. Upon startup, after synchronization (if specified), the failover sends a Poll message to the primary server. It waits for the amount of time specified by the WaitPollRplSecs parameter for a reply. If a reply is received, the failover server "sleeps" for the amount of time specified by this parameter before sending another Poll message to the primary server. The maximum value is 86400 (1 day).

SyncBindRetryMax —The number of times the server should resend a binding update if an Ack is not received within WaitSyncBindAckSecs.

SyncBindingBufSize—Size of the "options" area for synchronizing the binding information between the primary and secondary servers.

Note: This value *must* be same on *both* servers. The minimum value is 64. The maximum value is 4096.

Use Failover Server —If set to True, the Failover Server parameter appears, where an administrator can select the fully qualified host name of the failover server to be used by the primary server. If set to False, no failover server is associated with the primary server.

WaitCtlRetSecs—The number of seconds that this primary server should wait for a response to the CtlReq (request for control) message sent to the failover server at start-up. The maximum number of seconds is 3600 (1 hour). If a value of zero is supplied, the default is used.

WaitSyncBindAckSecs —The number of seconds the server waits for an Ack to a binding update packet when operating on the sending side of synchronizing with the other server.

WaitSyncBindUpdateSecs—The number of seconds that the server should wait for subsequent binding update packets when operating on the receiving side of synchronizing with the other server.

VitalQIP Failover Parameters for Secondary Server are as follows:

PollRetryMax—If no reply to the Poll message is received from the primary, the failover retries sending the Poll message and waiting for a reply for this number of times before assuming that the primary has crashed and becomes active in handling DHCP client requests. The maximum is 10.

SyncBindingBufSize—Size of the "options" area for synchronizing the binding information between the primary and secondary servers.

Note: This value *must* be the same on *both* servers. The minimum value is 64. The maximum value is 4096.

SyncBindings—If this parameter is True, this failover server requests all current binding information from the primary at startup.

SyncBindRetryMax—The number of times the server should resend a binding update if an Ack is not received within WaitSyncBindAckSecs.

SyncFailCritical—If True, the failover server terminates upon failure to synchronize bindings with the primary server. Note that this parameter only applies if the SyncBindings parameter is True.

SyncReqRetryMax—If no SyncStart message is received from the primary server in the time specified by WaitSyncStartSecs, attempt to contact the primary by resending the SyncReq message this many times. If this maximum is reached without receiving a SyncStart response, this failover checks the value of the SyncFailCritical parameter. If that parameter is True, this secondary server terminates. Otherwise, this secondary server initiates polling of the primary server. The maximum value is 10.

WaitPollRplSecs—The number of seconds that the failover should wait for a Poll reply from the primary. Note that if the failover receives a binding update from the primary during this period, the primary is assumed operational, and the binding update message is taken as a response to the poll. The maximum is 3600 (1 hour).

WaitSyncBindAckSecs—The number of seconds the server waits for an Ack to a binding update packet when operating on the sending side of synchronizing with the other server.


WaitSyncBindUpdateSecs—The number of seconds that the server should wait for subsequent binding update packets when operating on the receiving side of synchronizing with the other server.

WaitSyncStartSecs —The number of seconds that the failover server should wait for a response to the SyncReq (request for bindings) message sent to the primary server at start-up. A value of zero causes the server to wait indefinitely. The maximum value is 3600 (1 hour).

6. Configuring VitalQIP DHCP Failover

To configure DHCP failover, always review the current Administrator Guide for the release of VitalQIP you are using. Below are the basic steps to configure DHCP failover in VitalQIP 7.2.

In the VitalQIP GUI, click on the **DHCP** tab, then **Add Server** from the **Actions** button. Enter the required data and then configure the parameters for the DHCP server on the default DHCP server tab. Click on the **Failover Server Type** and then choose **Failover/Secondary** from the list at the bottom. At this time the system will add lines through all parameters that are associated with a primary DHCP server to show they are not applicable. Click on each parameter one by one and set the correct value. When done click on the **Submit** button. Figure 7 below contains a screenshot of the GUI that shows the parameters for a failover DHCP server.



My View Address Management DNS **DHCP** Infrastructure Reports Task

DHCP Servers

Refresh Add +

- Standalone/Primary
 - dhcp1.beef.com
 - dhcp1.drink.com
 - dhcp1_to.beef.com
 - dhcp5.beer.com
 - drinkdhcp.drink.com
 - tebeast.beef.com
- Failover/Secondary

Save Delete Actions Refresh

* DHCP Server ActiveLease MAC Pool Reports

Server Properties

Server Address: 10.1.0.4 *Server Type: LUCENT DHCP 5.5

Host Name: dhcp1 *Domain Name: beef.com

Server Parameter

Name	Value
Failover Server Type	Failover/Secondary
ClientReqRetryMax	3
PollDelay	60
PollRetryMax	3
SyncBindRetryMax	3
SyncBindingBufSize	1024
SyncBindings	True
SyncFailCritical	False
SyncReqRetryMax	3
Use Failover Server	False
WaitClRetSecs	5
WaitPollRptSecs	5

Failover Server Typ

Failover/Secondary

Standalone/Primary

Failover/Secondary

Figure 7

For further information on DHCP failover and VitalQIP please visit <http://enterprise.alcatel-lucent.com/?product=VitalQIP&page=overview>.