



Subscriber Manager Failover

Published: December 23, 2013

Introduction

This chapter describes using the Subscriber Manager with clusters for redundancy.

The Subscriber Manager is part of the Cisco Service Control Application for Broadband (SCA BB) solution that is deployed in tier-one service provider environments. Subscriber Manager supports a fail-over mode. Fail-over mode minimizes system downtime that is caused by Subscriber Manager failure. (See the [“Information About Subscriber Manager Fail-Over”](#) section on page 3-1.)

This section describes using a cluster of two Subscriber Manager nodes in fail-over mode.



Note

For the purposes of this section, it is assumed that the reader is familiar with the Veritas Cluster technology.

- [Information About Subscriber Manager Fail-Over, page 3-1](#)
- [How to Recover from Fail-Over, page 3-5](#)

Information About Subscriber Manager Fail-Over

This section consists of these topics:

- [Overview, page 3-1](#)
- [Normal Operation, page 3-2](#)
- [Fail-Over Topology, page 3-3](#)
- [Fail-Over Operation, page 3-4](#)

Overview

The fail-over scheme that is implemented in the Subscriber Manager is based on the Veritas cluster technology. The cluster includes two machines, each running Subscriber Manager database and Veritas software. The Veritas Cluster Server (VCS) software consolidates the Subscriber Managers, which constitutes a single entity by providing a single virtual IP address for the entire cluster.

The cluster software distinguishes an active and a standby machine. The active machine *owns* the virtual IP address and all network connections. The standby machine is passive until a fail-over occurs. At fail-over, the IP address is passed from the failing server to the backup server, which becomes activated and re-establishes all network connections.

When a fail-over occurs, the login event generators (LEGs) lose their connection with the failed Subscriber Manager. The LEGs reconnect to the activated (backup) Subscriber Manager and retransmit their uncommitted messages. The activated Subscriber Manager connects to the SCE platforms and performs an SCE resynchronization.

The Subscriber Manager database replication agent constantly replicates the Subscriber Manager database from the active node to the standby node. This enables a fast fail-over from one Subscriber Manager to another because the subscriber data in the activated machine is always valid. The two Subscriber Manager nodes do not communicate except to pass the subscriber data.

The VCS uses software components called *cluster agents* to monitor and control the state of resources such as Network Interface Cards (NICs), disks, IP addresses, and processes. Cisco supplies cluster agents to monitor the Subscriber Manager and the Subscriber Manager database daemon and replication agent.

As part of the cluster operation, the Subscriber Manager database daemon and replication agents are operating regardless of the fail-over state. The Subscriber Manager Veritas agent monitors the daemon and the replication agent process. If one of them fails, a fail-over occurs.

**Note**

The Subscriber Manager software configuration on both the active and the standby machines must be *identical*. Apply the same configuration files to both machines.

Normal Operation

The two Subscriber Manager nodes operate in hot-standby mode. In this mode, the active node receives and processes all the Subscriber Manager events. The standby node waits and is ready to go into operation on fail-over. To enable seamless fail-over and to minimize the fail-over time, the two Subscriber Manager nodes operate without an external storage device.

During the normal operation of the cluster, the active node (selected by the cluster) does the following:

- Performs all Subscriber Manager functionality of a non-cluster environment
- Provides *health* information for the cluster agent
- Periodically replicates its subscriber database to the standby node


On the standby node, both the Subscriber Manager and the Subscriber Manager database software are running:

- The Subscriber Manager is fully configured. (It is applied with the same configuration files as the active node but does not interfere with the active node's operation.)
- The Subscriber Manager connects to the Subscriber Manager database but does not connect to the LEG and SCE devices.
- The Subscriber Manager database software operates as a replication client for the subscriber database, receiving and applying updates from the Subscriber Manager database software on the active node.

Fail-Over Topology

Figure 3-1 depicts a Subscriber Manager cluster configuration in a topology with a redundant AAA server and two SCE 2000 platforms that are cascaded for redundancy.

Figure 3-1 Subscriber Manager Cluster Configuration for Fail-Over Topology



The Subscriber Manager fail-over topology includes two Subscriber Manager nodes connected in a cluster topology.

Two dedicated (private) redundant networks interconnect the two nodes:

- Heartbeat network—Used by the Veritas Cluster Server to perform cluster monitoring and control.
- Replication network—Used by the replication process to pass the subscriber records.

The two nodes should be located at the same site. The heartbeat network is implemented by using back-to-back connectivity between the two nodes or by using redundant switches. Each node in the cluster has redundant network paths (NICs) connecting it to all of the external entities with which the Subscriber Manager communicates (AAA, LEG, SCE).

Each node in the cluster has a minimum of six Ethernet NICs deployed as follows:

- Two NICs are used for the private heartbeat network
- Two NICs are used for the private replication network
- Two NICs are used for the public network (connectivity to SCEs and LEGs, and management of the Subscriber Manager)


The cluster has a virtual IP (VIP) address for communication with the external entities. Each node in the cluster also has an IP address for administration of the node/cluster, as well as an IP address for replication.

If the primary NIC of the public network fails, a failover occurs to the secondary NIC on the same node, keeping the same IP addresses (VIP1), with no failover of the cluster. If the primary NIC of the replication or heartbeat network fails, a failover occurs to the secondary NIC on the same node, keeping the same IP addresses (VIP2 and VIP3), with no failover of the cluster.

Figure 3-2 illustrates the regular and virtual IP addresses used in a cluster configuration:

- Administration of the nodes uses IP1/IP2 and IP3/IP4 respectively.
- The cluster IP address for external clients over the public network uses VIP1.

Figure 3-2 Regular and Virtual IP in Cluster Configuration



157083

For more information about replication IP configuration, see [Appendix E, “Veritas Cluster Server.”](#)

Fail-Over Operation

During normal operation, the Veritas Cluster Server automatically selects one of the Subscriber Manager servers to be active and the other to be standby.

The active Subscriber Manager performs all the normal Subscriber Manager functions. The two servers maintain the heartbeat mechanism between them, and the active server continuously replicates the subscriber database to the standby server’s database.

The standby Subscriber Manager server acts as a hot-standby machine, so that it is ready to take over (become activated) in minimal fail-over time.

The following types of failures trigger the fail-over mechanism:

- Subscriber Manager application failure, including failure of the Subscriber Manager database.
- Failure of the Subscriber Manager database daemon of the Subscriber Manager database replication process.
- SUN server failure, due to failure of one of the resources of the server; for example, failure of both of the public network NICs.
- Manual activation of fail-over.



Note

Communication failure does *not* cause a fail-over if there is a redundant NIC. Because each SUN machine has two NICs for connecting to external devices, a failure of one NIC causes a switch to the redundant NIC, without activating the fail-over mechanism.

After detecting a failure, the standby Subscriber Manager becomes activated, and the following occurs:

- The activated Subscriber Manager takes over the IP resources of the virtual IP mechanism.
- The LEGs reconnect to the activated Subscriber Manager.

- The activated Subscriber Manager creates IP connections with the SCEs and resynchronizes with them.
- The activated Subscriber Manager starts processing information that is sent from the LEGs and forwards it to the SCEs.

How to Recover from Fail-Over

Different types of failures require different triggering for the recovery procedure. Some failures, such as intra-node ports link-failure, recover automatically when the link revives. Others failures might need manual intervention.

Recovery might occur when a Subscriber Manager that experienced a failure recovers or after it is replaced (if needed). The recovery procedure is intended to take the cluster back to a fully functional mode. When the recovery procedure concludes, the behavior is the same as it was after installation.

A failed Subscriber Manager server is recovered manually or automatically, according to the type of failure that occurred. The following sections describe the recovery procedures:

- [Machine Reboot, page 3-5](#)
- [Replacing the Server, page 3-5](#)
- [Database Duplication Recovery, page 3-6](#)
- [Database Duplication Recovery Management, page 3-7](#)

Machine Reboot

Recovery from a machine reboot is a fully automatic recovery process, whereby the failed Subscriber Manager server reboots. After reestablishing a connection with the other server and synchronizing the databases, the cluster of the two Subscriber Manager servers is ready again for fail-over operation.



Note The steps in this procedure are automatic.

Step 1 The reboot process runs on the node.

Step 2 VCS makes the node standby.

Step 3 The node boots.

Step 4 VCS establishes intra-node communication and the new node joins the cluster.

Step 5 The Subscriber Manager database replication process is started from the point before the reboot.

The Subscriber Manager in the recovered server is ready after the database recovery process runs and the Subscriber Manager moves from INIT state to standby state.

Replacing the Server

You must replace the server when the machine suffers an unrecoverable physical failure. You must install a new machine with a fresh Subscriber Manager, Subscriber Manager database, and VCS.

Replacing the server is a manual recovery procedure. You must physically replace the failed Subscriber Manager server. After you connect the new Subscriber Manager server to the network, configure it, and synchronize the two databases, the cluster of two Subscriber Manager servers is once again capable of fail-over operation.

Step 1 Connect a new server to the inter-node ports and intra-node ports (but leave the network ports disconnected).

Step 2 Perform the basic network and cluster configurations manually (the first time).

Step 3 Copy the configuration files from the active node.

Use the following CLU command if you need to copy only the **p3sm.cfg** file:

```
p3sm --load-config --remote=NEW-SM_IP
```

Step 4 Perform the Subscriber Manager database duplication operation.

See the “[Database Duplication Recovery](#)” section on page 3-6.

Step 5 Start the VCS operation on the recovered node.

Step 6 Connect the network ports.

The Subscriber Manager in the recovered server is ready after the database recovery process completes and the Subscriber Manager moves from INIT state to standby state.

Database Duplication Recovery

Database duplication recovery is a manual recovery procedure, which is required when the standby node database loses synchronization with the active node database. Loss of synchronization can occur when one of the Subscriber Manager machines is replaced or when the replication process on the active node fails to replicate all of the data inserted into its database (replication NICs were disconnected).

Step 1 Stop the cluster server (VCS) monitoring of the resources.

Enter the VCS CLU **hastop -local** command to stop the VCS.

Step 2 Stop the Subscriber Manager so that it will not be affected by clearing the database.

Enter the CLU command **p3sm --stop**.

Step 3 Stop the replication agent.

Enter the CLU command **p3db --rep-stop**.

Step 4 Destroy the database.

Enter the CLU command **p3db --destroy-rep-db**.

Step 5 Duplicate the remote database on the local machine.

Enter the CLU command **p3db --duplicate**.

Step 6 Start the cluster server monitoring of the resources.

Enter the VCS CLU **hastart** command, which automatically starts the replication process and the Subscriber Manager.

Database Duplication Recovery Management

Configure the two Subscriber Manager servers by using Command-Line Utilities and a configuration file (see the “[Configuring the Subscriber Management Solution](#)” section on page 5-4). Perform the configuration for the active Subscriber Manager and then manually replicate the configuration for the standby Subscriber Manager.

Step 1 Establish an FTP connection between the active and standby machines.

Step 2 Copy the configuration files.

Copy all the configuration files from the `~pcube/sm/server/root/config/` folder on the active node to the standby node. Apply the Subscriber Manager configuration file by using the CLU command `p3sm --load-config`.

Step 3 Copy the database-related configuration files to the required location.

If you changed the database-related configuration files, copy the files to `/etc/system` (for Solaris) or to `/etc/sysctl.conf` (for Linux), and `sys.odbc.ini` from the active node to the standby node.



Note If you perform this step, you must reboot the standby node.



Note If the database is located in different directories in the two nodes, the `sys.odbc.ini` files in the nodes are not identical and the actual parameter changed in the file must be copied.

Step 4 Configure and administer the Veritas Cluster Server by using Veritas tools.

Notifications are enabled through SNMP traps that the Veritas Cluster Server provides. The Veritas Cluster Server supports SNMP traps such as the following:

- Fatal failure detected (local or remote)
 - Secondary node starts fail-over procedure
 - Secondary node is operational (end of fail-over)
-

