



CHAPTER 4

Configuring High Availability

This chapter describes how to configure high availability on the MSE. The mobility services engine is a platform for hosting multiple mobility applications. Every active MSE is backed up by another inactive instance. The active MSE is called the primary MSE and the inactive MSE is called the secondary MSE.

The main component of high availability system is the health monitor. The health monitor configures, manages, and monitors the high availability setup. Heartbeat is maintained between the primary and secondary MSE. Health monitor is responsible for setting up the database, file replication, and monitoring the application. When the primary MSE fails and the secondary MSE takes over, the virtual address of the primary MSE is switched transparently.

This chapter contains the following sections:

- [Overview to the High Availability Architecture, page 4-1](#)
- [Pairing Matrix, page 4-2](#)
- [Guidelines and Limitations for High Availability, page 4-2](#)
- [Failover Scenario for High Availability, page 4-2](#)
- [Failback, page 4-3](#)
- [HA Licensing, page 4-3](#)
- [Configuring High Availability on the MSE, page 4-3](#)
- [Viewing Configured Parameters for High Availability, page 4-6](#)
- [Viewing High Availability Status, page 4-7](#)

Overview to the High Availability Architecture

This section provides an overview of the high availability architecture:

- Every active primary MSE is backed up by another inactive instance. The purpose of the secondary MSE is to monitor the availability and state of the primary MSE. The secondary MSE becomes active only after the failover procedure is initiated.
- The failover procedure can be manual or automatic.
- One secondary MSE can support two primary MSEs.
- There is one software and database instance for each registered primary MSE.

Pairing Matrix

Table 4-1 lists the server type pairing matrix information.

Table 4-1 Pairing Matrix

Primary Server Type	Secondary Server Type							
		3310	3350	3355	VA-2	VA-3	VA-4	VA-5
3310	Y	Y	Y	N	N	N	N	
3350	N	Y	Y	N	N	N	N	
3355	N	Y	Y	N	N	N	N	
VA-2	N	N	N	Y	Y	Y	Y	
VA-3	N	N	N	N	Y	Y	Y	
VA-4	N	N	N	N	N	Y	Y	
VA-5	N	N	N	N	N	N	Y	

Guidelines and Limitations for High Availability

- Both the health monitor IP and virtual IP should be accessible from NCS.
- The health monitor IP and virtual IP should always be different. The health monitor and virtual interface can be on the same interface or different interfaces.
- You can use either manual or automatic failover. Failover should be considered temporary. The failed MSE should be restored to normal as soon as possible, and failback will re-initiated. The longer it takes to restore the failed MSE, the longer the other MSEs sharing the secondary MSE must run without failover support.
- You can use either manual or automatic failback.
- Both the primary and secondary MSE should be running the same software version.
- High Availability over WAN is not supported.
- High availability over LAN is supported only when both the primary and secondary MSE are in the same subnet.
- The ports over which the primary and secondary MSEs communicate must be open (not blocked with network firewalls, application firewalls, gateways, and so on).

Failover Scenario for High Availability

When a primary MSE failure is detected, the following events take place:



Note

One secondary MSE can back up multiple primary MSEs.

- The primary MSE is confirmed as non-functioning (hardware fail, network fail, and so on) by the health monitor on the secondary MSE.

- If automatic failover has been enabled, the secondary MSE is started immediately and uses the corresponding database of the primary MSE. If automatic failover is disabled, an e-mail is sent to the administrator asking if they want to manually start failover.
- When the manual failover is configured, an e-mail is sent only if the email is configured for MSE alarms. When manual failover is configured and not invoked, there is no need for failback.
- Failback is invoked and the primary MSE assumes all the operations.
- The result of the failover operation is indicated as an event in the Health Monitor UI, and a critical alarm is sent to the administrator.

Failback

When the primary MSE is restored to its normal state if the secondary MSE is already failing over for the primary, then failback can be invoked.

Failback can occur only if the secondary MSE is in one of the following states for the primary instance:

- The secondary MSE is actually failing over for the primary MSE.
- If manual failover is configured but the administrator did not invoke it.
- The primary MSE failed but the secondary MSE cannot take over because it has encountered errors or it is failing over another primary MSE.
- Failback can occur only if the administrator starts up the failed primary MSE.

HA Licensing

A separate license is not required to set up an MSE HA system. A virtual appliance secondary does not need an activation license.

Configuring High Availability on the MSE

Configuring high availability on the MSE involves the following two steps:

- During the installation of the MSE software, you must perform certain configurations using the command-line client.
- Pair up the primary and secondary MSE from the NCS UI.

**Note**

If you do not want high availability support and if you are upgrading from an older release, you can continue to use the old IP address for the MSE. If you want to set up high availability, then you must configure the health monitor IP address. The health monitor then becomes a virtual IP address.

**Note**

By default, all MSEs are configured as primary.

To configure high availability on the primary MSE, follow these steps:

- Step 1** Ensure that the network connectivity between the primary and secondary is functioning and that all the necessary ports are open.
- Step 2** Install the correct version of MSE on the primary MSE.
- Step 3** Make sure that the same MSE release version that is loaded on the other primary MSE and secondary MSE is also loaded on the new primary MSE.
- Step 4** On the intended primary MSE, enter the following command:

```
/opt/mse/setup/setup.sh
```

```
-----
Welcome to the appliance setup.
Please enter the requested information. At any prompt,
enter ^ to go back to the previous prompt. You may exit at
any time by typing <Ctrl+C>.
You will be prompted to choose whether you wish to configure a
parameter, skip it, or reset it to its initial default value.
Skipping a parameter will leave it unchanged from its current
value.
Changes made will only be applied to the system once all the
information is entered and verified.
-----
```

- Step 5** Configure the hostname:

```
Current hostname=[mse]
Configure hostname? (Y)es/(S)kip/(U)se default [Skip]:
```

The hostname should be a unique name that can identify the device on the network. The hostname should start with a letter, end with a letter or number, and contain only letters, numbers, and dashes.

- Step 6** Configure the domain name:

Enter a domain name for the network domain to which the device belongs. The domain name should start with a letter, and it should end with a valid domain name suffix such as *.com*. It must contain only letters, numbers, dashes, and dots.

```
Current domain=[]
Configure domain name? (Y)es/(S)kip/(U)se default [Skip]:
```

- Step 7** Configure the HA role:

```
Current role=[Primary]
Configure High Availability? (Y)es/(S)kip/(U)se default [Skip]:
High availability role for this MSE (Primary/Secondary):
Select role [1 for Primary, 2 for Secondary] [1]: 1
```

Health monitor interface holds physical IP address of this MSE server.
This IP address is used by Secondary, Primary MSE servers and NCS to communicate among themselves

```
Select Health Monitor Interface [eth0/eth1] [eth0]:eth0
```

```
-----
Direct connect configuration facilitates use of a direct cable connection between the
primary and secondary MSE servers.
This can help reduce latencies in heartbeat response times, data replication and failure
detection times.
```

Please choose a network interface that you wish to use for direct connect. You should appropriately configure the respective interfaces.

```
\ "none\" implies you do not wish to use direct connect configuration.
-----
```

Step 8 Configure Ethernet interface parameters:

```

Select direct connect interface [eth0/eth1/none] [none]: eth0
Enter a Virtual IP address for first this primary MSE server:
Enter Virtual IP address [172.31.255.255]:
Enter the network mask for IP address 172.31.255.255.
Enter network mask [255.255.255.0]:
Current IP address=[172.31.255.255]
Current eth0 netmask=[255.255.255.0]
Current gateway address=[172.31.255.256]
Configure eth0 interface parameters? (Y)es/(S)kip/(U)se default [Skip]:

```

Step 9 When prompted for “eth1” interface parameters, enter Skip to proceed to the next step. A second NIC is not required for operation:

```

Configure eth1 interface parameters? (Y)es/(S)kip/(U)se default [Skip]:

```

Follow [Step 10](#) through [Step 13](#) to configure the secondary MSE.

Step 10 Configure the hostname for the secondary MSE:

```

Current hostname=[]
Configure hostname? (Y)es/(S)kip/(U)se default [Skip]:

```

Step 11 Configure the domain name:

```

Current domain=
Configure domain name? (Y)es/(S)kip/(U)se default [Skip]:

```

Step 12 Configure the HA role:

```

Current role=[Primary]
Configure High Availability? (Y)es/(S)kip/(U)se default [Skip]:
High availability role for this MSE (Primary/Secondary)
Select role [1 for Primary, 2 for Secondary] [1]: 2
Health monitor interface holds physical IP address of this MSE server.
This IP address is used by Secondary, Primary MSE servers and NCS to communicate among
themselves
Select Health Monitor Interface [eth0/eth1] [eth0]:[eth0/eth1]
-----
Direct connect configuration facilitates use of a direct cable connection between the
primary and secondary MSE servers.
This can help reduce latencies in heartbeat response times, data replication and failure
detection times.
Please choose a network interface that you wish to use for direct connect. You should
appropriately configure the respective interfaces.
\"none\" implies you do not wish to use direct connect configuration.
-----

```

Step 13 Configure ethernet interface parameters:

```

Select direct connect interface [eth0/eth1/none] [none]: eth1
Enter a Virtual IP address for first this primary MSE server
Enter Virtual IP address [172.19.35.61]:
Enter the network mask for IP address 172.19.35.61:
Enter network mask [255.255.254.0]:
Current IP address=[172.19.35.127]
Current eth0 netmask=[255.255.254.0]
Current gateway address=[172.19.34.1]
Configure eth0 interface parameters? (Y)es/(S)kip/(U)se default [Skip]:

```

- Step 14** Once you have configured both the primary MSE and secondary MSE, the NCS UI should be used to set up a pairing between the primary and secondary MSE.
- Step 15** Once the primary MSE is added successfully, choose **Services > High Availability** or click the primary MSE device in the **Services > Mobility Services Engine** page, and choose **HA Configuration > Service High Availability** from the left sidebar menu.
- The HA Configuration page appears.
- Step 16** Enter the secondary device name with which you want to pair the primary MSE.
- Step 17** Enter the secondary IP address which is the health monitor IP address of the secondary MSE.
- Step 18** Enter the secondary password. This is the NCS communication password configured on the MSE.
- Step 19** Specify the failover type. You can choose either Manual or Automatic from the Failover Type drop-down list. After 10 seconds, the system fails over. The secondary server waits for a maximum of 10 seconds for the next heartbeat from the primary server. If it does not get the heartbeat in 10 seconds, it declares a failure.
- Step 20** Specify the failback type by choosing either **Manual** or **Automatic** from the Failback Type drop-down list.
- Step 21** Specify the Long Failover Wait in seconds.
- After 10 seconds, the system fails over. The maximum failover wait is 2 seconds.
- Step 22** Click **Save**.
- The pairing and the synchronization happens automatically.
- Step 23** To check whether the heartbeat is received from the primary MSE or not, choose **Services > Mobility Services Engine**, and click **Device Name** to view the configured parameters.
- Step 24** Choose **HA Configuration > Service High Availability** from the left sidebar menu.
- Check whether the heartbeat is received from the primary MSE or not.
-

Viewing Configured Parameters for High Availability

To view the configured parameters for high availability, follow these steps:

- Step 1** Choose **Services > High Availability**.
- Step 2** Click **Device Name** to view its configured fields.
- The HA configuration page appears.
- Step 3** Choose **Services High Availability > HA Configuration** from the left sidebar menu. The HA Configuration page shows the following information:
- Primary Health Monitor IP
 - Secondary Device Name
 - Secondary IP Address
 - Secondary Password
 - Failover Type
 - Failback Type

- Long Failover Wait
-

Viewing High Availability Status

To view the high availability status, follow these steps:

-
- Step 1** Choose **Services > High Availability**.
- Step 2** Click **Device Name** to view the desired status.
The HA Configuration page appears.
- Step 3** Choose **Services High Availability > HA Status** from the left sidebar menu. The HA Configuration page shows the following information:
- Current high Availability Status
 - Status—Shows whether the primary and secondary MSE instances are correctly synchronized or not.
 - Heartbeats—Shows whether the heartbeat is received from the primary MSE or not.
 - Data Replication—Shows whether the data replication between the primary and secondary databases is happening or not.
 - Mean Heartbeat Response Time—Shows the mean heartbeat response time between the primary and secondary MSE instance.
 - Event Log—Shows all the events generated by the MSE. The last 20 events can be viewed.
-

