Configuring Cisco ICFP for Clusters

Workflow for Configuring Clusters

The following table identifies the high-level tasks that are required to configure a multiple-node cluster.

Step	Task	Related Information
1.	Install a minimum of four Cisco ICFP virtual appliances. The role that is assigned to each appliance during installation depends on whether you use VMware or OpenStack.	Deployment Workflows
2.	Configure two primary nodes.	Configuring a Primary Node, on page 2
3.	Configure two or more service nodes.	Configuring a Service Node, on page 3
4.	Configure additional storage.	Configuring Additional Storage, on page 4

Step	Task	Related Information
5.	Configure the two primary nodes for HA.	Configuring HA, on page 7
6.	(OpenStack only) Configure VIP access.	Configuring VIP Access for HA in OpenStack, on page 9
7.	Configure a load balancer for the service nodes in the cluster. **Note** The load balancer must be configured to persist sessions based on the PERSISTICFPP cookie that Cisco ICFP issues.	Your load balancer documentation

Configuring a Primary Node

To configure a Cisco ICFP virtual appliance that has been installed using the Standalone Mode role for a multiple-node cluster, you must first configure it as a primary node or service node by using the ShellAdmin console. This procedure describes how to configure a standalone node as a primary node. To configure a standalone node as a service node, see Configuring a Service Node, on page 3.

Before You Begin

Install a Cisco ICFP virtual appliance using the Standalone Mode role.

Procedure

Step 1 Using SSH, log in to the ShellAdmin console of the standalone node that you want to configure as a primary node.

Step 2 At the ShellAdmin prompt, choose **Change Node Role**.

Step 3 When prompted, enter **Y** to change the node role.

Step 4 Enter **A** to configure the node as a primary node.

Step 5 Enter **Y** to confirm that you want to configure the node as a primary node.
Information similar to the following is displayed:

```
user selected 'y'
 Checking DB Status
   2399 ?       00:00:00 mysqld_safe
   2820 ?       00:04:21 mysqld
Configuring as Primary Node...
Stopping services before changing node role
Stopping the services...
Setting up current node as Primary node...
Enabling Remote Database access to ICFPP Service nodes
Checking the MySQL to be ready before enabling remote access to DB...
Waiting a maximum of 900 seconds for MySQL to be up on localhost

Trying a maximum of 900 seconds for enabling remote access to DB
Successfully enabled remote access for database

SUCCESS: Successfully changed node role to Primary Node

Stopping Database and restarting it for changes to take effect
Stopping database...
```

```
Database stopped...
Starting services that were previously stopped.
Starting the Database...
Starting the services...
In order for changes to take effect logout and log back in
Do you want to logout [y/n]?
```

Step 6 Enter **Y** when prompted to log out.
You are logged out of the ShellAdmin console. When you log in again, the ShellAdmin menu includes options for configuring HA and viewing HA status.

Configuring a Service Node

To configure a Cisco ICFP virtual appliance that has been installed using the Standalone Mode role for a multiple-node cluster, you must first configure it as a primary node or as a service node by using the ShellAdmin console. This procedure describes how to configure a standalone node as a service node. To configure a standalone node as a primary node, see Configuring a Primary Node, on page 2.

Before You Begin

- Install a Cisco ICFP virtual appliance using the Standalone Mode role.
- Obtain the IP address of a primary node in the cluster or the virtual IP address (VIP) of an HA pair in the cluster.
- Back up any data in the virtual appliance database that you want to keep. When the virtual appliance is reconfigured as a service node, the existing data is deleted.

Procedure

Step 1 Using SSH, log in to the ShellAdmin console of the standalone node that you want to configure as a service node.
Step 2 At the ShellAdmin prompt, choose **Change Node Role**.
Step 3 When prompted, enter **Y** to change the node role.
Step 4 Enter **B** to configure the node as a service node.
Step 5 Enter **Y** to confirm that you want to configure the node as a service node.
Step 6 When asked if you want to continue, do one of the following:

- Enter **N** to stop the configuration so that you can back up the database.
- Enter **Y** to continue.

If you choose to continue, Cisco ICFP confirms your choice.
Step 7 When prompted, enter the IP address of the primary node or the VIP of the HA pair that the service node is to use.
Information similar to the following is displayed:

```
Configuring as Service Node...
Stopping services before changing node role
Stopping the services...
```

```
Setting up current node as ICFPP service node...with remote DB IP 123.45.1.60
Disabling Database service at startup

SUCCESS: Successfully changed node role to Service Node

Starting services that were previously stopped...
Starting the services...
In order for the changes to take effect, log out and log in again
Do you want to log out [y/n]?
```

Step 8 Enter **Y** to log out.
When you next log in, the menu includes options for working with a service node.

Configuring Additional Storage

The default disk size of 100 GB for Cisco ICFP is not sufficient for configuring Cisco ICFP in a multiple-node cluster. As a result, you must add additional disk space before configuring a multiple-node cluster. You can use either NFS or a Cinder volume as described in the following topics:

- Configuring NFS, on page 4
- Configuring a Cinder Volume, on page 5

Configuring NFS

If you did not configure an NFS server for a Cisco ICFP virtual appliance when you installed it, you can configure the appliance for NFS by using the ShellAdmin console.

Note We recommend that you configure additional storage for all Cisco ICFP nodes. If additional storage is not configured, all VM images that are uploaded from Cisco Intercloud Fabric are stored on the node's local disk. If the node fails, one or both of the following can occur:

- Any images stored on the node are no longer available.
- If the node is part of a cluster, template creation and VM migration fail.

If NFS is not available, you can configure a Cinder volume as described in Configuring a Cinder Volume, on page 5.

Before You Begin

- Upload all images that reside on the Cisco ICFP virtual appliance to the cloud. If you do not upload the images to the cloud, the images are deleted when NFS is configured.
- Identify the NFS server IP address and the directory in which the files are to be stored.

Procedure

Step 1 Using SSH, log in to the ShellAdmin console for the Cisco ICFP virtual appliance that you want to configure for NFS.

Step 2 Choose **NFS Configuration**.
Cisco ICFP displays a menu with options for configuring, removing, and viewing an NFS configuration.

Step 3 At the prompt, enter **A**.
Cisco ICFP determines whether or not an NFS directory is mounted and displays the results:

```
Checking for mounted NFS directory...
NFS directory is not mounted
Note: Configuring NFS will delete any images that are not uploaded to the cloud! Proceed
[y/n]?
```

Step 4 Enter **Y** to continue.
Cisco ICFP determines whether or not an NFS IP address or NFS directory has been configured and then prompts you for input.

Step 5 When prompted, enter the NFS server IP address and the NFS directory path.
Information similar to the following is displayed while NFS is configured:

```
Configuring NFS with : NFS Server IP=123.15.1.1, remote directory=/nfs/dir local mounting
point=/mnt/icfpp-images
Creating /mnt/icfpp-images directory.
Starting portmap and nfs services...
Starting portmap:                                          [  OK  ]
mount -t nfs 123.15.1.1:/icfpp-images /mnt/icfpp-images
May wait for mount up to 12-0 seconds..., please be patient...
Successfully mounted 123.15.1.1:/icfpp-images at /mnt/icfpp-images
Saving NFS Configuration
NFS IP address: 123.15.1.1
NFS Directory Path: /icfpp_images
Saved NFS Configuration
Setting up images directory to use NFS
Image directory setup to NFS done
Press Return to continue
```

Step 6 Press **Enter** to return to the ShellAdmin menu.
To view or remove the NFS configuration, choose **NFS Configuration** in the ShellAdmin menu, and then choose the appropriate option from the NFS menu.

Configuring a Cinder Volume

The default disk size of 100 GB for the Cisco ICFP virtual appliance is not sufficient for configuring Cisco ICFP in a multiple-node cluster. If you do not have access to an NFS server, you can increase the disk size by creating additional Cinder volumes. Cinder volumes that you create are formatted as physical disks and then combined to form a logical volume that can be mounted on the VM in a specific location.

Before You Begin

- Configure a Cisco ICFP virtual appliance as a service node by using the ShellAdmin console. For more information, see Configuring a Service Node, on page 3.

- If you have not already done so, configure the root user password for the Cisco ICFP service node. For more information, see the "Using Cisco ICFP ShellAdmin Commands" chapter in the Cisco Intercloud Fabric for Provider Administrator Guide.
- Collect the following information:
  - Cloud credentials—The username and password for the project in OpenStack.
  - Cloud URL—Obtain the cloud URL as follows:
    1. In the OpenStack dashboard, choose **Project** > ***project*** > **Access & Security**, and click the **API Access** tab.
    2. In the **API Endpoints** table, locate the **Identity** service and note the service endpoint URL.
  - Cisco ICFP instance ID—Obtain the Cisco ICFP instance ID as follows:
    1. In the OpenStack dashboard, choose **Project** > ***project*** > **Instances**.
    2. In the list of instances, locate Cisco ICFP and click the hyperlinked instance name.

       The **Instance Detail** page is displayed.
    3. In the **Overview** tab, locate and note the instance ID.

Procedure

Step 1 Using SSH, log in to the ShellAdmin console of the Cisco ICFP service node.

Step 2 At the ShellAdmin prompt, choose **Cinder Storage Configuration**.

Step 3 When prompted, enter **Y** and enter the root password.

Step 4 At the Cinder Storage Configuration menu prompt, choose **Deploy Fresh Storage**.
Cisco ICFP prompts you for information so that it can configure the storage.

Step 5 Enter the following information:

- Cloud username and password
- OpenStack project name
- Cloud URL
- Cisco ICFP instance ID
- Required storage size in GB
- Required volume size in GB

Note Cinder storage configuration supports a volume with a maximum of 2 TB for each service node.

Information similar to the following is displayed while Cisco ICFP creates and formats the volume. You do not need to restart the Cisco ICFP virtual appliance.

```
Cloud user name:- abc1-de2.gen
Enter password:
Project Name:- ABC-DEV1
Cloud URL: [e.g. https://us-texas-3.cloud.abc.com:5000/v2.01]:-
```

```
https://us-texas-3.cloud.abc.com:5000/v2.0
ICFPP Instance ID:- 75c8c226-b22c-4041-ab5c-7e7fd544c3b
Expected storage size[GB]:- 10
Expected volume size[GB]:- 10
Deploying fresh storage

*****Creating volumes*****

*****Attaching volumes*****

*****Formatting volumes and creating logical volumes*****

*****Validating final state*****
true
Executed successfully!
```

Step 6 If needed, you can do either of the following from the Cinder Storage Configuration menu:

- To configure additional storage, choose **Add additional storage to existing storage**.
- To delete storage, choose **Cleanup deployed storage**.

Configuring HA

After you deploy Cisco ICFP virtual appliances, you can configure them for high availability (HA) by using the ShellAdmin console.

When configuring HA:

- Configure the active node and standby node concurrently as described in this procedure.
- The database on the standby node is deleted when the HA pair is configured.

Before You Begin

- Deploy or configure two Cisco ICFP virtual appliances as primary nodes:
  - To deploy a Cisco ICFP virtual appliance with the Primary Mode role, see Deployment Workflows.
  - To configure an existing Cisco ICFP virtual appliance as a primary node, see Configuring a Primary Node, on page 2.
- Identify a virtual IP (VIP) address for the HA pair.
- Determine which node will be the active node and which node will be the standby node.
- On the node that will be the standby node, move any existing data that you want to save to another location.

Procedure

Step 1 Using SSH, log in to the ShellAdmin console of the node that will be the active node for the HA pair.

Step 2 At the ShellAdmin prompt, choose **Setup HA**.

A warning is displayed stating that the contents of the database on the standby node will be deleted.

Step 3 When prompted, enter **Y** to configure the node for HA.

Step 4 Enter **A** to configure the node as the active node.

Step 5 When prompted, enter **Y** to configure the node as the active node.
Cisco ICFP detects and displays the IP address of the current node.

Step 6 Enter **Y** to confirm the node IP address.

Step 7 Enter the standby node IP address.

Step 8 Enter the VIP to use for the IP pair.
Information similar to the following is displayed:

```
-------------------------------------------
HA Configuration Information:
-------------------------------------------
This node will be configured as active node
Active Node IP address:   123.45.1.61
Standby Node IP address:  123.45.1.62
Virtual IP address:       123.45.1.60
-------------------------------------------
Proceed with setting up HA with above configuration [y/n]:
```

Step 9 Enter **Y** to confirm the configuration and continue, or **N** to change the values.
If you choose to continue, Cisco ICFP displays progress messages while it configures the active node for HA.

Step 10 While Cisco ICFP configures the active node for HA, log in to the ShellAdmin console of the node that will be the standby node for the HA pair.

Step 11 At the ShellAdmin prompt, choose **Setup HA**.

Step 12 Enter **Y** to configure the node for HA.

Step 13 Enter **B** to configure the node as the standby node.

Step 14 When prompted, enter **Y** to configure the node as the standby node.
Cisco ICFP detects and displays the IP address of the current node.

Step 15 Enter **Y** to confirm the node IP address.

Step 16 Enter the active node IP address.

Step 17 Enter the VIP to use for the HA pair.
Information similar to the following is displayed:

```
-------------------------------------------
HA Configuration Information:
-------------------------------------------
This node will be configured as standby node
Active Node IP address:   123.45.1.61
Standby Node IP address:  123.45.1.62
Virtual IP address:       123.45.1.60
-------------------------------------------
Proceed with setting up HA with above configuration [y/n]:
```

Step 18 Enter **Y** to confirm the configuration.
Cisco ICFP displays progress messages while it configures the standby node for HA and synchronizes the database information on both nodes.

Step 19 When prompted, press **Enter** to return to the ShellAdmin menu.

What to Do Next

For OpenStack environments, continue with Configuring VIP Access for HA in OpenStack, on page 9.

Configuring VIP Access for HA in OpenStack

After Cisco ICFP primary nodes are configured for HA, the virtual IP address (VIP) is used in the event of failover. However, OpenStack Neutron does not allow a host to accept packets with an IP address in the packet header that does not match the destination host IP address. As a result, packets sent to the VIP do not reach the node to which the VIP is assigned. To allow the packets to reach HA pair, the VIP must be added as an allowed address for both nodes (active and standby) in the HA pair.

This procedure describes how to configure VIP access on the nodes in the HA pair by using the OpenStack **neutron port-update** command. For more information, see the OpenStack documentation at docs.openstack.org.

Before You Begin

- Confirm that HA has been configured on two Cisco ICFP primary nodes in an OpenStack environment.
- Confirm that you have access to the OpenStack Neutron command-line tool.

Procedure

Step 1 Obtain a list of networks by entering the following command:

```
$ neutron net-list
```

Information similar to the following is displayed:

```
+--------------------------------------+-------------------+-----------------------------------------------------+
| id                                   | name              | subnets                                             |
                                    |
+--------------------------------------+-------------------+-----------------------------------------------------+
| 2d84eaa4-8b81-4dc8-9897-dd8ef4719f8b | public-direct-600    |
3e0b77fe-fc66-4913-bc58-7f62d4ab247a 10.203.28.0/23 |
|                                      |                   |
5c2f73a9-4e2f-498c-8244-6aefe5129fdd 10.203.50.0/23 |
|                                      |                   |
ba29165f-c88a-496a-9adc-99ee90407ebe 10.203.24.0/23 |
|                                      |                   |
d5b69780-aefb-42a6-8ba5-aaf405fb36a0 10.203.30.0/24 |
| b5d8d461-74d7-45a4-a1f0-f7ac96586bd5 | Net1              |
c0921b42-2896-4b32-b33e-f54db9e5a3d6 192.168.0.0/24 |
| ca80ff29-4f29-49a5-aa22-549f31b09268 | public-floating-601  |
0cfde3f1-e28b-4b87-8095-e0014b0ee573               |
|                                      |                   |
348a808d-ce64-43bc-a9d9-c20e52d2ac06               |
|                                      |                   |
3784170e-5d7f-48b4-b63d-aab4a0fef769               |
| ff95095f-89f0-4005-b709-70a75212d73c | icfp-ha-123-network  |
1099b814-05d9-4da0-93d1-06167db4891f 192.168.1.0/24 |
+--------------------------------------+-------------------+-----------------------------------------------------+
```

Step 2 Obtain a list of ports on the network on which the active and standby nodes in the HA pair are deployed by entering the following command:

```
$ neutron port-list -- --network_id=net_id
```

where *net_id* is the identifier for the required network. In this example, the network name is icfp-ha-123-network.

```
$ neutron port-list -- --network_id=ff95095f-89f0-4005-b709-70a75212d73c
```

Information similar to the following is displayed:

```
+--------------------------------------+----+-------------------+-------------------------------------------------------------------------------+
| id                                   | name | mac_address       | fixed_ips                                                                     |
+--------------------------------------+----+-------------------+-------------------------------------------------------------------------------+
| 4a439cf1-b95e-49ba-a8d6-0b03a8142dd2 |      | fa:16:3e:f6:f8:a9 | {"subnet_id":
"1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.12"} |
| 93d0a69a-7bb8-4719-9ed7-63c10accd78b |      | fa:16:3e:1f:7f:d2 | {"subnet_id":
"1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.11"} |
| 9d626a64-ee7c-410b-ae00-661dd275de79 |      | fa:16:3e:61:81:4b | {"subnet_id":
"1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.14"} |
| cf56fd7b-2896-4e06-b520-1d2258ad6158 |      | fa:16:3e:ab:27:ca | {"subnet_id":
"1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.13"} |
| d7457d29-44ba-46ef-b47a-4b94c9199902 |      | fa:16:3e:ad:d0:e9 | {"subnet_id":
"1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.15"} |
+--------------------------------------+----+-------------------+-------------------------------------------------------------------------------+
```

Step 3 In the output of the previous step, locate the port ID for the active node.

Step 4 Update the port so that it accepts traffic from the VIP by entering the following command:

```
$ neutron port-update active-port-id --allowed_address_pairs list=true type=dict
ip_address=vip
```

where:

- *active-port-id* is the port ID of the active node.
- *vip* is the virtual IP address for the HA pair.

For example, if the IP address of the active node is 192.168.1.11 and the VIP is 192.168.1.10, the command resembles the following:

```
$ neutron port-update 93d0a69a-7bb8-4719-9ed7-63c10accd78b --allowed_address_pairs list=true
 type=dict ip_address=192.168.1.10
```

Step 5 View the port details and confirm that the **allowed_address_pairs** field lists the VIP by entering the following command:

```
$ neutron port-show active-port-id
```

where *active-port-id* is the identifier for the port configured in the previous step.

Using the current example, the command and results resemble the following:

```
$ neutron port-show 93d0a69a-7bb8-4719-9ed7-63c10accd78b
```

```
+---------------------+-------------------------------------------------------------------------------------------+
| Field                 | Value
                    |
+---------------------+-------------------------------------------------------------------------------------------+
| admin_state_up        | True
                    |
| allowed_address_pairs | {"ip_address": "192.168.1.10", "mac_address": "fa:16:3e:1f:7f:d2"}
                  |
```

```
| device_id             | b7b8eeb5-70ad-49ac-a3b4-6d8a144293a2
                  |
| device_owner          | compute:alln01-1-csi
                  |
| extra_dhcp_opts       |
                  |
| fixed_ips           | {"subnet_id": "1099b814-05d9-4da0-93d1-06167db4891f","ip_address":
 "192.168.1.11"}  |
| id                    | 93d0a69a-7bb8-4719-9ed7-63c10accd78b
                  |
| mac_address           | fa:16:3e:1f:7f:d2
                  |
| name                  |
                  |
| network_id            | ff95095f-89f0-4005-b709-70a75212d73c
                  |
| security_groups       | f995d22f-edb8-47c0-9aff-6339a15fb5be
                  |
| status                | ACTIVE
                  |
| tenant_id             | b1436740f8db42e39904ee9779f67eb8
                  |
+---------------------+-------------------------------------------------------------------------------------+
```

Step 6 Configure the standby node to accept VIP traffic by entering the following command:

```
$ neutron port-update standby-port-id --allowed_address_pairs list=true type=dict
ip_address=vip
```

where:

- *standby-port-id* is the port ID of the standby node.
- *vip* is the virtual IP address for the HA pair.

Step 7 View the port details for the standby node and confirm that the **allowed_address_pairs** field lists the VIP:

```
$ neutron port-show standby-port-id
```

Step 8 (Optional) Complete the following steps to configure the VIP so that it is accessible from an external network and so that the VIP uses a floating IP address:

a) Configure a port corresponding to the VIP by entering the following command:

```
$ neutron port-create --fixed-ip ip_address=ip --security-group security-group network-name
```

where:

- *ip* is the fixed IP address for the port.
- *security-group* is the name of the security group to use for this port.
- *network-name* is the name of the network to which the port belongs.

Using the current example, the command and results resemble the following:

```
$ neutron port-create --fixed-ip ip_address=192.168.1.10 --security-group default
icfp-ha-123-network

Created a new port:
+-----------------------+-----------------------------------------------------------------------------------+
| Field                 | Value
                        |
+-----------------------+-----------------------------------------------------------------------------------+
| admin_state_up        | True
                        |
| allowed_address_pairs |                |
| device_id             |
                        |
| device_owner          |                                          |
| fixed_ips             | {"subnet_id":
"1099b814-05d9-4da0-93d1-06167db4891f","ip_address": "192.168.1.10"}  |
| id                    | ea35e2a9-1b45-4b05-b345-f4758e490052
        |
| mac_address           | fa:16:3e:df:e9:69                                                                 |
| name                  |                 |
| network_id            | ff95095f-89f0-4005-b709-70a75212d73c            |
| security_groups       | f995d22f-edb8-47c0-9aff-6339a15fb5be
           |
| status                | DOWN                |
| tenant_id             | b1436740f8db42e39904ee9779f67eb8
        |
+-----------------------+-----------------------------------------------------------------------------------+
```

b) In the OpenStack Horizon GUI, associate a floating IP address with the port to which the fixed IP address is assigned.

Moving from a Standalone Setup to a Cluster

Cisco ICFP enables you to move from a standalone configuration to a cluster. Moving from a standalone configuration to a cluster involves moving the database contents from the existing standalone node to the active HA node in the cluster as described in this procedure.

After moving the database contents, you can configure and test the cluster setup without modifying or affecting the standalone setup. For more information about configuring a multiple-node cluster, see Workflow for Configuring Clusters, on page 1.

Before You Begin

- Obtain the FTP server IP address and login credentials for backing up and restoring the database.
- Confirm that the standalone node and both nodes in the HA pair are running the same version of Cisco ICFP.

Procedure

Step 1 In the ShellAdmin console for the standalone node, back up the existing database as follows:

a) Choose **Stop Services** to stop the Infrastructure Manager services.
b) Choose **Backup Database**.
c) Choose **Start Services**.

Step 2 Deploy or configure two primary nodes by using any of the following methods:

- For VMware environments, deploy two new Cisco ICFP virtual appliances using the Primary Node role. For more information, see Installing Cisco ICFP on VMware.
- For OpenStack environments, deploy two new Cisco ICFP virtual appliances using the Standalone Node role and then configure the appliances as primary nodes. For more information, see Installing Cisco ICFP on OpenStack.
- Configure existing Cisco ICFP virtual appliances using the Standalone Node role as primary nodes. For more information, see Configuring a Primary Node, on page 2.

Step 3 Restore the backed-up database from Step 1 onto one of the primary nodes:

a) In the primary node ShellAdmin console, choose **Stop Services** to stop the Infrastructure Manager services.
b) Choose **Restore Database**.
c) Choose **Start Services**.

Step 4 In the ShellAdmin console, configure the two primary nodes as an HA pair.

Note You must configure the primary node on which the database was restored as the active node in the HA pair. If you configure it as the standby node, the database on that node is deleted.

For more information, see Configuring HA, on page 7.

Step 5 Configure service nodes for the cluster. For more information, see Configuring a Service Node, on page 3.

Restoring a Database onto an Existing HA Pair

Cisco ICFP enables you to configure an HA pair and then restore a database from an existing standalone node to the HA pair.

Note You must stop and start services in the sequence described in this procedure to successfully restore the database on the HA pair.

Before You Begin

- Confirm that the standalone node and both nodes in the HA pair are running the same version of Cisco ICFP.
- Back up the required database from a standalone node onto an FTP server.
- Identify the active node in the HA pair on which to restore the backed-up database.

Procedure

Step 1 Stop the VIP service on the current standby node in the HA pair as follows:

a) Log in to the ShellAdmin console for the current standby node.
b) Choose **Setup HA**.
c) When asked if you want to reconfigure HA, enter **Y**.
d) Enter **C** to stop the VIP service.
e) Enter **Y** to confirm the action.
f) Press **Enter** to return to the ShellAdmin menu.

Step 2 Stop the VIP service on the current active node in the HA pair as follows:

a) Log in to the ShellAdmin console for the current active node.
b) Choose **Setup HA**.
c) When asked if you want to reconfigure HA, enter **Y**.
d) Enter **C** to stop the VIP service.
e) Enter **Y** to confirm the action.
f) Press **Enter** to return to the ShellAdmin menu.

Stopping the VIP service on the active node in an HA pair automatically stops the Infrastructure Manager services if they are running.

Step 3 On the active node in the HA pair, restore the database backup obtained from the standalone node as follows:

a) In the ShellAdmin console for the active node, choose **Restore Database**.
b) When prompted, enter the FTP server IP address and login credentials.
c) Enter the path and filename for the backed-up database file on the FTP server.
d) Follow the onscreen prompts to complete the process.

Step 4 Restart the VIP service on the active node as follows:

a) In the ShellAdmin console for the active node, choose **Setup HA**.
b) When asked if you want to reconfigure HA, enter **Y**.
c) Enter **D** to start the VIP service.
d) Press **Enter** to return to the ShellAdmin menu.

Starting the VIP service on the active node in an HA pair automatically starts the Infrastructure Manager services on that node.

Step 5 Restart the VIP service on the standby node in the HA pair as follows:

a) In the ShellAdmin console for the standby node, choose **Setup HA**.
b) When asked if you want to reconfigure HA, enter **Y**.
c) Enter **D** to start the VIP service.
d) Press **Enter** to return to the ShellAdmin menu.

Monitoring HA Status

After configuring Cisco ICFP for HA, you can view the configuration details, check the status of the active and standby nodes, and view detailed replication status.

Procedure

Step 1 Log in to the ShellAdmin console for one of the nodes in the HA pair.

Step 2 At the prompt, choose **Display HA Status**.
Information similar to the following is displayed:

```
Configured HA role for this node is:  Active
Current HA role for this node is: Active
HA Configuration properties for this node are:
ACTIVE_IP_ADDRESS=123.16.1.30
STANDBY_IP_ADDRESS=123.16.1.3
VIRTUAL_IP_ADDRESS=123.16.1.25

IP address of this node is: 123.16.1.30
Checking if Virtual IP Address is reachable...OK
Virtual IP Address service status on this node...OK
Checking DB replication from 123.16.1.30 to 123.16.1.3...OK
Checking DB replication from 123.16.1.3 to 123.16.1.30...OK

Do you want to view detailed replication status ? [y/n]
```

Step 3 To view detailed information, enter **Y**.
Information similar to the following is displayed:

```
               Slave_IO_State : Waiting for master to send event
                  Master_Host : 123.16.1.3
                  Master_User : replicator
                  Master_Port : 3306
                Connect_Retry : 60
              Master_Log_File : mysql-bin.000002
          Read_Master_Log_Pos : 645644
               Relay_Log_File : mysqld-relay-bin.000004
                Relay_Log_Pos : 361
        Relay_Master_Log_File : mysql-bin.000002
             Slave_IO_Running : Yes
            Slave_SQL_Running : Yes
              Replicate_Do_DB :
          Replicate_Ignore_DB :
...
```

Step 4 Use your arrow keys to scroll through the information, and enter **Q** to stop viewing the detailed information and press **Enter** to return to the menu.

Viewing HA Syslog Messages

After configuring Cisco ICFP for HA, Cisco ICFP checks HA status every five minutes. Any warning or failure messages that are issued are included in the log file for syslog messages. This log file commonly resides in `/var/log/` with the name `messages`. To view these messages, log in as root and use a text editor as described in this procedure.

Procedure

Step 1 In the ShellAdmin console, choose **Log in as Root**.

Step 2 Enter **Y** to confirm the login request, and enter the root account password at the prompt.

Step 3 Enter the following command to view the contents of the log file:

```
vi /directory-path/filename
```

where *directory-path* is location of the log file and *filename* is the name of the log file. For example, you might enter the following:

```
vi /var/log/messages
```

Step 4 To identify messages that pertain to HA, look for entries that contain the string `icfpp-ha` as shown in the following example:

```
Jul  3 03:29:01 icfpp-ha-primary rsyslogd: [origin software="rsyslogd" swVersion="5.8.10"
 x-pid="3946" x-info="http://www.rsyslog.com"] rsyslogd was HUPed
Jul  8 03:45:01 icfpp-ha-primary rsyslogd: [origin software="rsyslogd" swVersion="5.8.10"
 x-pid="3946" x-info="http://www.rsyslog.com"] rsyslogd was HUPed
```

Step 5 Address any HA-related messages as needed.