Cisco Intercloud Fabric for Provider Installation Guide, Release 3.1.1

First Published: July 28, 2016

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

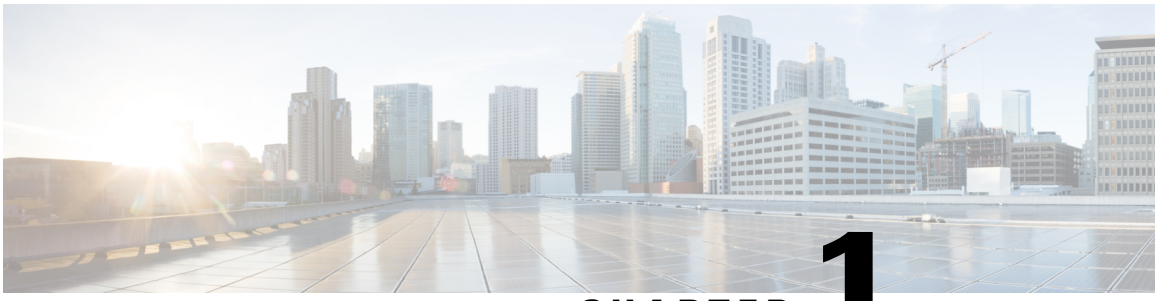

CHAPTER 1

Overview

Cisco Intercloud Fabric for Provider (ICFP) simplifies the complexity involved in working with a variety of public cloud APIs, and enables cloud API support for service providers who currently do not have API support. Cisco ICFP provides an extensible adapter framework that allows integration with a variety of provider cloud infrastructure management platforms, such as OpenStack, CloudStack, VMware vCloud Director, and any other API that can be integrated through a software development kit (SDK) provided by Cisco.

Currently, service providers have their own proprietary cloud APIs—such as Amazon EC2 and Windows Azure—that give customers limited choices and do not provide an easy method for moving from one provider to another. Cisco ICFP abstracts this complexity and translates Cisco Intercloud Fabric cloud API calls to the cloud platform APIs of different provider infrastructure platforms, giving customers the option of moving their workloads regardless of the cloud API used by the service provider.

Many service providers do not provide cloud APIs that Cisco Intercloud Fabric can use to deploy customers' workloads. One option for these providers is to provide direct access to the virtual machine (VM) manager's SDK or API, such as VMware vCenter or Microsoft System Center. However, this option exposes the provider environment and service providers prefer not to use it because of security concerns. Cisco ICFP, as the first point of authentication for the customer cloud when requesting cloud resources, enforces highly secure access to the provider environment. In addition, Cisco ICFP provides the cloud APIs that are required for service providers to be part of the provider ecosystem for Cisco Intercloud Fabric.

As the interface between Cisco Intercloud Fabric from customer private cloud environments and provider clouds (public and virtual private clouds), Cisco ICFP provides the following benefits:

- Standardizes and brings uniformity to cloud APIs, making it easier for Cisco Intercloud Fabric to consume cloud services from service providers that are a part of the Cisco Intercloud Fabric ecosystem.
- Helps secure access to a service provider's underlying cloud platform.
- Limits the utilization rate per customer or tenant environment.
- Provides northbound APIs for service providers for integration with existing management platforms.
- Supports multitenancy.
- Monitors resource usage for each tenant.
- Meters resource usage for each tenant.

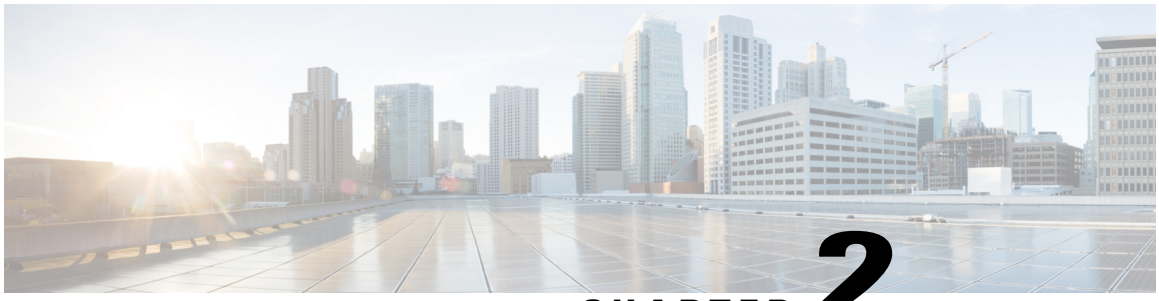

CHAPTER 2

Cisco ICFP Deployment Options

Deployment Options

You can deploy Cisco ICFP in the service provider data center in the following configurations:

- Standalone—Deployment on a single node.
- Multiple-node cluster—Deployment on multiple nodes including a high-availability (HA) pair and additional service nodes.

Cluster deployments are most effective when they are configured behind a load balancer. After these configurations are deployed, a provider-supplied load balancer is expected to manage cookie-based sessions and direct requests and responses appropriately.

The following topics describe these configuration options in more detail.

Standalone Configuration

In a standalone configuration, Cisco ICFP is deployed as a single virtual appliance that provides services and acts independently of other Cisco ICFP nodes. A standalone configuration is appropriate for environments in which redundancy is not a concern.

Cisco ICFP is installed in OpenStack environments using the Standalone role only. After installation, you can configure the Cisco ICFP virtual appliance as a primary node or service node as appropriate for your environment.

Cluster Configuration

In a cluster configuration, Cisco ICFP supports large-scale operations in provider environments by deploying multiple Cisco ICFP nodes. A provider-supplied load balancer distributes the load across the service nodes.

In a cluster configuration, Cisco ICFP is deployed on multiple virtual appliances using the Primary Node and Service Node roles, as shown in the following figure.

Figure 1: Cisco ICFP in a Multiple-Node Cluster Configuration

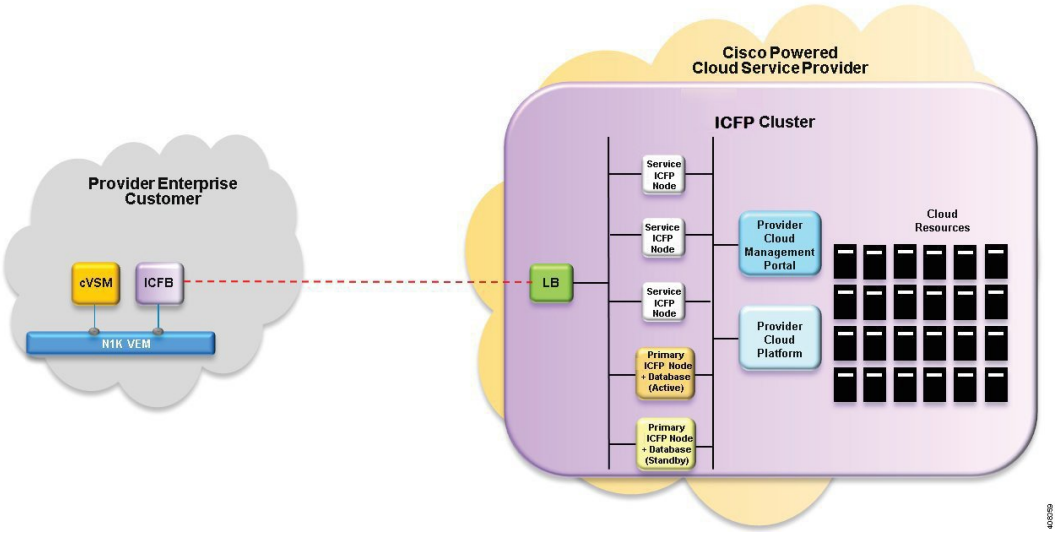

A multiple-node cluster contains the following components:

- Two Cisco ICFP virtual appliances with the Primary Node role that are deployed in an HA configuration.
- Two or more Cisco ICFP virtual appliances that are deployed as service nodes.
- A load balancer that forwards incoming traffic only to the service nodes.

HA Configuration in a Cluster

In an HA configuration, Cisco ICFP is deployed on two virtual appliances, both using the Primary Node role. Each virtual appliance in an HA pair includes a database for replication purposes. After both Cisco ICFP appliances are deployed, you specify which node is active and which node is standby.

The following concepts apply when Cisco ICFP is deployed in an HA configuration:

- The Cisco ICFP appliances in an HA pair have different management IP addresses.
- A single virtual IP address (VIP) is assigned to the active node.
- If the active node fails, the VIP is transferred to the standby node.
- When the original active node recovers, the VIP remains with the original standby node until that node fails.

Database replication works as follows:

- The active and standby nodes replicate each other's databases.
- At any time, only the database on the node with the VIP is used.
- When the database is updated on one node, the changes are replicated on the other node.

Communications in a Cluster

In a cluster configuration, Cisco ICFP manages communications as follows:

- Each service node establishes a connection with the database on the active node in the HA pair by using the virtual IP address for the HA pair.
- The load balancer accepts requests from Cisco Intercloud Fabric.
- The load balancer distributes the requests to the service nodes using a round-robin algorithm.
- Each new user session is directed to a different service node.
- Subsequent requests from the same session are sent to the same service node.
- The service node responds via the load balancer.

Session Persistence

Session persistence is managed by means of a PERSISTICFPP cookie that Cisco ICFP issues. The cookie, which is generated when a user logs in, ensures that all requests from that user session are directed to the same node. If a service node fails, the load balancer forwards requests for that service node to a different service node. The new receiving node first requires Cisco Intercloud Fabric to log in and then accepts new requests.

Note The service provider load balancer must be configured to persist sessions based on the PERSISTICFPP cookie.

Deployment Workflows

The deployment workflow that you use depends on whether you deploy Cisco ICFP in a VMware or OpenStack environment.

The following table describes the high-level tasks required to deploy Cisco ICFP in a multiple-node cluster in a VMware environment.

Table 1: Configuration Workflow for a Multiple-Node Cluster on VMware

Step	Task	Related Information
1.	Install two Cisco ICFP virtual appliances using the Primary Node role.	Installing Cisco ICFP on VMware, on page 14
2.	Install two or more Cisco ICFP virtual appliances using the Service Node role.	Installing Cisco ICFP on VMware, on page 14
3.	Configure additional storage.	Configuring NFS, on page 38

Step	Task	Related Information
4.	Configure HA on the appliances with the Primary Node role.	Configuring HA, on page 41
5.	Configure a load balancer for all 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
6.	Configure communications for the cluster with Cisco Intercloud Fabric.	*Cisco Intercloud Fabric Installation Guide*

The following table describes the high-level tasks required to deploy Cisco ICFP in a multiple-node cluster in an OpenStack environment.

Table 2: Configuration Workflow for a Multiple-Node Cluster on OpenStack

Step	Task	Related Information
1.	Install four or more Cisco ICFP virtual appliances using the Standalone Node role.	Installing Cisco ICFP on OpenStack, on page 20
2.	Configure two appliances with the Primary Node role.	Configuring a Primary Node, on page 36
3.	Configure the remaining appliances with the Service Node role.	Configuring a Service Node, on page 37
4.	Configure additional storage.	Configuring a Cinder Volume, on page 39
5.	Configure HA on the appliances with the Primary Node role.	Configuring HA, on page 41
6.	Configure the HA nodes to permit network traffic via the VIP address.	Configuring VIP Access for HA in OpenStack, on page 43
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
8.	Configure communications for the cluster with Cisco Intercloud Fabric.	*Cisco Intercloud Fabric Installation Guide*

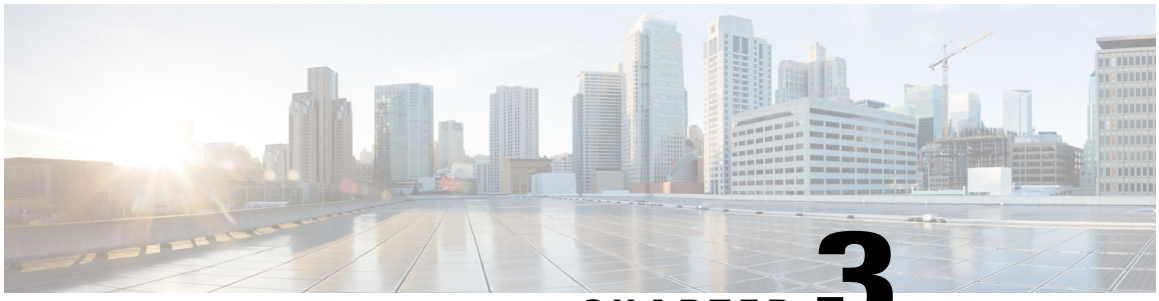

CHAPTER 3

Installation Requirements

System Requirements

You can deploy a Cisco ICFP virtual appliance on a system that meets the following requirements:

Table 3: Cisco ICFP System Requirements

Requirement	Description
Four virtual CPUs	1.8 GHz
Memory	8 GB RAM
Disk space	Disk space that is configured as follows: • Disk 1—100 GB for Cisco ICFP. • Disk 2—As much memory as required to support concurrent virtual machines being moved to the provider cloud. **Note** If additional storage is not configured, Cisco ICFP stores VM images uploaded from Cisco Intercloud Fabric on the local disk. For more information on configuring additional storage, see Configuring Additional Storage, on page 38.
One vNIC	Management network interface

Hypervisor Requirements

Cisco ICFP is a virtual appliance that can be deployed on the VMware vSphere Client or OpenStack KVM Hypervisor.

Table 4: Cisco ICFP Hypervisor Requirements

Hypervisor	Version
VMware	
VMware vSphere Client	ESXi 5.1, 5.5, and 6.0
OpenStack	
Cisco Intercloud Services OpenStack	Kilo
Red Hat Enterprise Linux OpenStack Platform	Kilo

Port Requirements

Ports must be configured as described in the following tables to ensure that Cisco ICFP can communicate effectively on the internal private network and the public network (Internet).

Table 5: Public Internet Inbound

Protocol	Port	Allow/Deny	Description
TCP	443	Allow	Allows inbound HTTPS traffic from the Internet so that Cisco Intercloud Fabric for Business can reach Cisco ICFP.

Table 6: Public Internet Outbound

Protocol	Port	Allow/Deny	Description
All	All	Deny	Cisco ICFP does not need to send outbound traffic to the Internet.

Table 7: Internal Network Inbound

Protocol	Port	Allow/Deny	Description
TCP	443	Allow	Allows inbound HTTPS traffic from the internal network, so that the Cisco ICFP web-based GUI can be accessed.

Protocol	Port	Allow/Deny	Description
TCP	22	Allow	Allows inbound SSH traffic from the internal network for Cisco ICFP administration.
TCP	3306	Allow	Allows inbound MySQL traffic from the internal network. Required if Cisco ICFP is configured in a multiple-node cluster.
TCP	8080	Allow	Allows inbound HTTP traffic for template uploads to CloudStack. Required if using the CloudStack adapter.

Note To ensure that the destination systems receive communications from Cisco ICFP, the ports in the following table must be open on any firewalls on the internal network between Cisco ICFP and the destination systems.

Table 8: Internal Network Outbound

Protocol	Firewall Port	Allow/Deny	Description
TCP	443	Allow	Allows HTTPS traffic to the internal network. Required to reach the cloud provider API/SDK gateway if it is running on HTTPS.
TCP	80	Allow	Allows HTTP traffic to the internal network. Required to reach the cloud provider API/SDK gateway if it is running on HTTP.
TCP	3306	Allow	Allows outbound MySQL traffic to other Cisco ICFP nodes on the internal network. Required if Cisco ICFP is configured in a multiple-node cluster.
TCP/UDP	514	Allow	Allows syslog traffic from Cisco ICFP to the syslog server.

Information Required for Configuration and Installation

Before installation, collect the following information:

Required Information	Mandatory / Optional	Your Information / Notes
For Preinstallation Configuration		

Required Information	Mandatory / Optional	Your Information / Notes
Cisco ICFP image location	Mandatory	
Cisco ICFP OVA or QCOW2 image name	Mandatory	
VM name	Mandatory	
VMware data store location	Mandatory for VMware	
Network / Port Profile for VM management	Mandatory	
KVM flavor name	Mandatory for OpenStack	
KVM Instance Security Group	Mandatory for OpenStack	
For Cisco ICFP Installation		
Installation type: Standalone, Primary, or Service Node For OpenStack environments, you can install only in Standalone mode.	Mandatory	
Hostname	Mandatory	
Admin / root / ShellAdmin account password	Mandatory	
Static IP address For OpenStack environments, this must be a public IP address.	Mandatory	
Subnet mask	Mandatory	
Gateway IP address	Mandatory	
Primary node IP address	Mandatory only for service node installations.	
NFS server IP address	Optional[1]	
NFS directory to mount	Optional	

Required Information	Mandatory / Optional	Your Information / Notes
Domain name	Optional	
DNS server IP address	Mandatory	
NTP server IP address or fully qualified domain name (FQDN)	Mandatory	
Cisco ICFP license	Optional	
Cisco ICFP Product Authorization Key (PAK)	Optional	

[1] If you do not configure NFS in a cluster deployment, template creation and VM migration can fail if a service node fails. If NFS is not available, you can configure a Cinder volume.

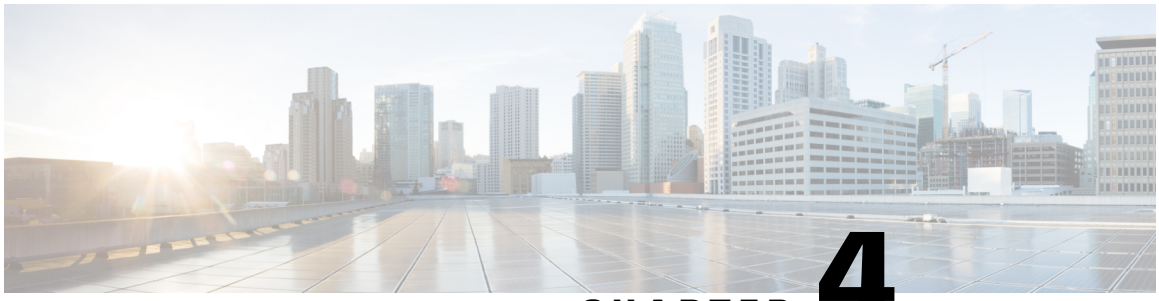

CHAPTER 4

Installing Cisco ICFP on VMware

Cisco ICFP Software

The Cisco ICFP software is available for download from Cisco.com. For assistance, contact your Cisco representative.

The Cisco ICFP software package (`icfp-dk9-3.1.1-pkg.zip`), contains the following files:

File	Description
`icfp-3.1.1.ova`	Cisco ICFP OVA file. Use this file to install Cisco ICFP in VMware environments. See Workflow for VMware Environments, on page 14.
`icfp-3.1.1.qcow2`	Cisco ICFP QCOW2 file. Use this file to install Cisco ICFP in OpenStack environments. See Workflow for OpenStack Environments, on page 19.
`README`	README file. This file contains information about installing and using Cisco ICFP.

The Cisco ICFP software includes a 60-day evaluation license with support for 20 hybrid cloud units (HCUs). To view the license details in the GUI after you install Cisco ICFP, choose **License**. The license details are displayed, including the license type, status, number of supported HCUs, and the term of the license. For more information, see Uploading Cisco ICFP Licenses, on page 25.

Workflow for VMware Environments

Cisco ICFP should be implemented by all service providers that interface with Cisco Intercloud Fabric for Business platforms. The only exceptions to this are Amazon EC2 and Windows Azure, which are available to Cisco Intercloud Fabric through their native public cloud APIs.

The following table identifies the high-level tasks involved in deploying Cisco ICFP in a VMware environment:

Step	Task	Related Information
1.	Confirm that you have met the installation requirements.	Installation Requirements, on page 7
2.	Gather the required information.	Information Required for Configuration and Installation, on page 9
3.	Install Cisco ICFP.	Installing Cisco ICFP on VMware, on page 14
4.	If needed after the installation, configure the Cisco ICFP IP address.	Configuring the IP Address for Network Access, on page 17
5.	(Optional) Upload the Cisco ICFP license file.	Uploading Cisco ICFP Licenses, on page 25
6.	(Optional) Configure Cisco ICFP virtual appliances for a multiple-node cluster.	Configuring Cisco ICFP for Clusters, on page 35
7.	Configure communications with Cisco Intercloud Fabric.	*Cisco Intercloud Fabric Installation Guide*

Installing Cisco ICFP on VMware

This procedure describes how to install Cisco ICFP in a VMware environment.

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 39.

Before You Begin

- Set your keyboard to United States English.
- Unzip the Cisco ICFP software package to obtain the OVA file and the README file.
- Review the README file for information related to Cisco ICFP installation and operation.
- Copy the Cisco ICFP OVA image to a location that is available from the VMware vSphere Client.
- Make sure that all requirements are met as specified in System Requirements, on page 7.
- Collect the information required for the installation. See Information Required for Configuration and Installation, on page 9.

Procedure

Step 1 Using the **VMware vSphere Client**, log in to the vCenter server.

Step 2 Choose the host on which to deploy the Cisco ICFP virtual appliance.

Step 3 Choose **File** > **Deploy OVF Template**.

Step 4 In the wizard, provide the information as described in the following table:

Screen	Action
Source	Choose the Cisco ICFP OVA using one of the following methods: • Browse to the location, choose the file, and click **Open**. • Deploy from a URL on your local area network. Replace **FQDN** with the IP address or the fully qualified domain name, and click **Next**.
OVF Template Details	Verify the details.
End User License Agreement	Read the agreement and click **Accept**.
Name and Location	1 Enter a name for the virtual appliance. 2 Choose the VMware data center or host where Cisco ICFP will reside.
Deployment Configuration	Choose the type of deployment: • **Standalone—**Used for single-node deployments. • **Primary Node—**Used for HA deployment in a multiple-node cluster. • **Service Node—**Used in cluster deployments for handling requests.
Storage	Choose the location in which to store the Cisco ICFP files.

Screen	Action
Disk Format	Choose the required format for the virtual appliance disks: • **Thick Provision Lazy Zeroed—**Allocates storage immediately in thick format. • **Thick Provision Eager Zeroed—**Allocates storage in thick format. Creating disks might take longer using this option. • **Thin Provision—**Allocates storage on demand as data is written to disk.
Network Mapping	Choose the required network.
Properties Address any errors that are indicated in red-colored text below a selection box.	
Node Mode	Choose the type of deployment for this node: Standalone, Primary Node, or Service Node. The mode you choose should match the deployment type in the **Deployment Configuration** screen.
ICFPP Hostname	Enter the hostname for the Cisco ICFP node.
ICFPP Password	Enter and confirm the password to use for admin, root, and ShellAdmin account access.
Static IP Address	Enter the static IP address to use for the Cisco ICFP node.
Static IP Subnet Mask	Enter the subnet mask to apply to the node IP address.
IP Gateway	Enter the gateway IP address.
Primary Node IP Address for Service Node	For service nodes only, enter the IP address of the primary node or the virtual IP address (VIP) of the HA pair for database access.
NFS Server IP Address	Enter the IP address for an NFS server. **Note** If you do not configure NFS in a multiple-node cluster deployment, template creation and VM migration can fail if a service node fails.
NFS Server Directory to Mount	NFS server directory to be mounted.
Domain Name	Enter the domain name for the node, such as cisco.com.
DNS Server IP Address	Enter the DNS server IP address.
NTP Server IP (FQDN or IP Address)	Enter the NTP server IP address or fully qualified domain name.
Ready to Complete	Review the deployment settings for accuracy.

Step 5 Click **Finish**. A progress indicator displays the task status until Cisco ICFP is deployed. For additional information, right-click the VM in the VMware vSphere Client and choose **Open Console**.

Step 6 After Cisco ICFP is successfully deployed, power on the virtual appliance.

What to Do Next

If needed, configure the Cisco ICFP IP address for network address. For more information, see Configuring the IP Address for Network Access, on page 17.

Configuring the IP Address for Network Access

After installing Cisco ICFP in a VMware environment, you might need to configure the Cisco ICFP IP address for network access.

The Cisco ICFP IP address is configured during installation by using Open Virtualization Format (OVF) parameters. However, if the IP address is not configured correctly, you must configure the static IP address by using the ShellAdmin console options as described in this procedure.

Procedure

Step 1 Using SSH, connect to the Cisco ICFP ShellAdmin console by using the following information:

- Cisco ICFP IP address
- Username—shelladmin
- Password—The password that you set when you installed Cisco ICFP

Step 2 At the ShellAdmin prompt, choose **Configure Network Interface** to configure the static IP address.

Step 3 Enter **S** to configure a static IP address.

Step 4 Enter the Ethernet interface that you want to configure, such as eth0 or eth1.

Step 5 When prompted for the IP version, choose **IPv4**.

Step 6 Enter the static IP address, netmask, and gateway IP address.

Step 7 Enter **Y** to confirm the information.
The Cisco ICFP virtual appliance reboots and displays a screen with the URL for accessing Cisco ICFP.

Step 8 (Optional) To verify that the change has been applied, log in to the ShellAdmin console and choose **Display Network Details**.

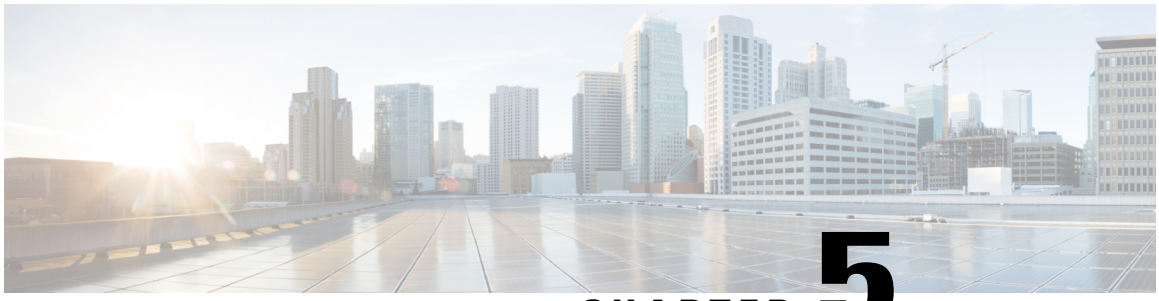

CHAPTER 5

Installing Cisco ICFP on OpenStack

Workflow for OpenStack Environments

Cisco ICFP should be implemented by all service providers that interface with Cisco Intercloud Fabric for Business platforms. The only exceptions to this are Amazon EC2 and Windows Azure, which are available to Cisco Intercloud Fabric through their native public cloud APIs.

The following table identifies the high-level tasks involved in installing and configuring Cisco ICFP in an OpenStack environment.

Step	Task	Related Information
1.	Confirm that you have met the installation requirements.	Installation Requirements, on page 7
2.	Gather the required information.	Information Required for Configuration and Installation, on page 9
3.	Configure OpenStack for Cisco ICFP and launch a Cisco ICFP instance.	Installing Cisco ICFP on OpenStack, on page 20
4.	(Optional) Upload a Cisco ICFP license file.	Uploading Cisco ICFP Licenses, on page 25
5.	Configure Cisco ICFP for use with Cisco Intercloud Fabric.	Configuring Cisco ICFP for Cisco Intercloud Fabric, on page 21
6.	(Optional) Configure Cisco ICFP virtual appliances for a multiple-node cluster.	Configuring Cisco ICFP for Clusters, on page 35
7.	Configure Cisco Intercloud Fabric for use with Cisco ICFP.	*Cisco Intercloud Fabric Installation Guide*

Installing Cisco ICFP on OpenStack

To install Cisco ICFP on OpenStack, you must import an image, create a flavor, and launch an instance. This procedure describes how to complete these tasks.

The amount of time required for this procedure depends on the platform:

- If the platform does not support QCOW2, the procedure can take up to two hours to complete, depending on the amount of time it takes to upload the image and convert it from QCOW2 format to RAW.
- If the platform supports QCOW2, no conversion is required, and the procedure takes less time.

Before You Begin

- Download the Cisco ICFP software package from Cisco.com. For assistance, contact your Cisco representative.
- Unzip the downloaded file to obtain the QCOW2 file and the README file. For more information, see Cisco ICFP Software, on page 13.
- Review the README file for information related to installing and using Cisco ICFP with OpenStack.
- Confirm that you have met the requirements in System Requirements, on page 7.
- Gather the information identified in Information Required for Configuration and Installation, on page 9.
- In OpenStack:
  - Confirm that you have admin privileges.
  - Create an OpenSource RC file (*name*-openrc.sh) in which you define your environmental variables and login credentials.
  - Create a project on which to install Cisco ICFP.
  - Confirm that the Cinder service is up and running.
  - Configure a security group that allows traffic on ports 22, 80, 443, and 3306.

  For more information about performing these operations in OpenStack, see docs.openstack.org.

Procedure

Step 1 In the shell from which you will enter **glance** commands, enter the following command:

```
source name-openrc.sh
```

Step 2 Copy the Cisco ICFP image to the system running the **glance** CLI.

Step 3 Using the **glance** CLI, upload an image to the OpenStack server by entering the following command:

```
glance image-create --name icfp-n.n.n --disk-format qcow2 --container-format bare --file
```

```
./icfp-n.n.n.qcow2
```

where *icfp-n.n.n* is the name of the Cisco ICFP image, such as `icfp-3.1.1`.

After the image has been uploaded, it appears in the OpenStack Dashboard Images table at **Admin** > **Images** or ***project*** > **Manage Compute** > **Images & Snapshots**.

Step 4 In the OpenStack Dashboard, choose **Admin** > **Flavors**, and click **Create Flavor**.

Step 5 In the **Create Flavor** dialog box, enter the following information, and click **Create Flavor**:

- Name—Enter a flavor name.
- vCPUs—Enter **4**.
- RAM MB—Enter **8192**.
- Root Disk—Enter the desired disk size in gigabytes.
- Ephemeral Disk—Enter **0**.
- Swap Disk—Enter **0**.

Step 6 Choose **Project** > ***project*** > **Manage Compute** > **Volumes**, and click **Create Volume**.

Step 7 In the **Create Volume** dialog box, add a volume with the size 100 GB, and click **Create Volume**.

Step 8 In OpenStack, obtain the following information:

- Flavor ID
- Image ID
- Network ID

Step 9 At the command line, enter the following command to launch Cisco ICFP:

```
nova boot --image image-id --flavor flavor-id
--nic net-id=network-id --block-device-mapping vdb=volume-id
icfp-instance-name
```

A Cisco ICFP instance is launched.

Configuring Cisco ICFP for Cisco Intercloud Fabric

After you have installed Cisco ICFP on an OpenStack server and launched a Cisco ICFP instance, you can configure Cisco ICFP for use with Cisco Intercloud Fabric.

Before You Begin

Confirm the following:

- Cisco ICFP has been installed on an OpenStack server and an instance has been launched.
- You know the Cisco ICFP public IP address.
- If Keystone V3 Identity Service is enabled on the cloud instance, the authentication domain.

- If group-based policies are enabled on the cloud instance, the external segment name and the name of the external group that is used to connect internal groups to the Internet.

Procedure

Step 1 In a browser, enter the public IP address assigned to the Cisco ICFP instance and log in to the Cisco ICFP GUI. The default credentials are:

- Username: admin
- Password: changeme

Step 2 In the OpenStack dashboard, choose **Project** > ***project*** > **Access & Security**, and click the **API Access** tab.

Step 3 In the **API Endpoints** table, locate and note the service endpoint Uniform Resource Identifier (URI) for the **Identity** service.

Step 4 In the Cisco ICFP GUI, choose **Cloud Instances**, and click the **Add Cloud Instance** icon.

Step 5 In the **New Cloud Instance** dialog box, provide the following information, and click **Create**:

Field	Description
Cloud Instance Name	Name of the cloud instance.
Select Cloud	The cloud instance type: Cisco or Custom.
Select Module	For a Cisco cloud instance type, choose the module type. For example, choose OSP for an OpenStack Platform cloud. For a custom cloud instance, enter the custom module name.
Endpoint URI	The endpoint hostname or IP address of the cloud instance.
Parameters The parameters that are displayed depend on the selected module.	
Image Conversion Support on Cloud	For OSP modules, indicate whether or not image conversion on the cloud is required.
First Boot Image Conversion Support	For OSP modules, indicate whether or not image conversion during VM boot on the cloud is required.
Enable Group-Based Policy Support	For OSP modules, indicate whether or not the provider OpenStack cloud uses a group-based policy framework.
Enable Keystone V3 API Support	For OSP modules, indicate whether or not OpenStack Keystone V3 Identity Service is used for authentication in the provider OpenStack cloud.
Enable Boot from Volume Support	For OSP modules, indicate whether or not the cloud instance is to be booted from a Cinder volume.

Step 6 In the Cisco ICFP GUI, choose **Tenant Accounts**, and click the **Add Tenant Account** icon.

Step 7 In the **New Tenant Account** dialog box, provide the following information, and click **Create**:

Field	Description
Tenant Name	Enter the tenant name.
Select Cloud	Choose the cloud instance. You cannot change the cloud instance after adding the tenant.
Max Servers	Enter the maximum of servers provisioned for the tenant, including stopped VMs.
Username	Enter the tenant account username.
Email	Enter the tenant account email address.
Parameters The parameters that are displayed depend on the selected cloud.	
External Segment Name	For an OpenStack cloud that uses a group-based policy framework, enter the external segment name.
Domain Name	For an OpenStack cloud with Keystone V3 Identity Service enabled, enter the domain name.
External Group Name	For an OpenStack cloud that uses a group-based policy framework, enter the name of the external group that is used to connect internal groups to the Internet.

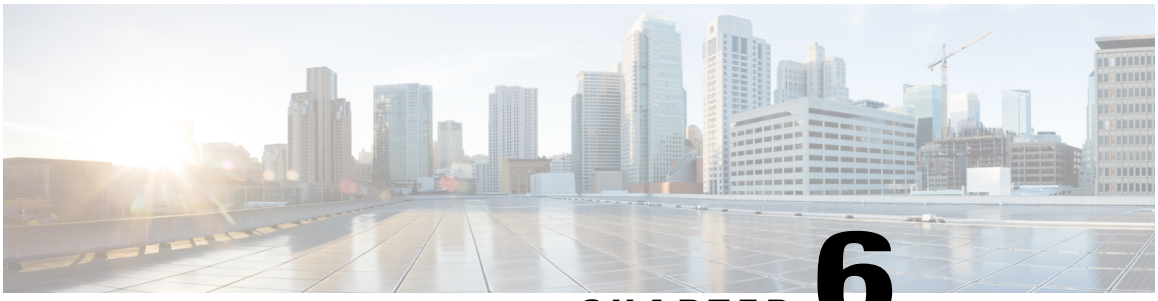

CHAPTER 6

Uploading Cisco ICFP Licenses

Cisco ICFP Licensing

A Cisco ICFP license is based on hybrid cloud units (HCUs). One or more HCUs are used for each VM running in the public cloud. A powered-off VM does not use any HCUs.

For Amazon Web Services and Microsoft Azure, two HCUs are used for each VM. For example, if the HCU count is ten, five VMs can run in the public cloud. For Cisco-powered providers, one HCU is used for each VM. For example, if the HCU count is ten, ten VMs can run in the public cloud.

Cisco ICFP includes the following types of licenses:

- Evaluation License (ICFP-EVAL-EBD)—Cisco ICFP includes a 60-day, 20-HCU evaluation license that lets you try the software before you purchase permanent licenses. The evaluation period begins when you install the software and expires within 60 days of installation.
- Permanent License (ICFP-CPC)—Permanent licenses have an expiration date. The license file specifies the number of licenses that you purchased. Contact your Cisco representative to purchase permanent licenses.
- Partner License (ICFP-NFR-EBDS)—Partner (not for resale) licenses are available only to Cisco partners for demonstration and lab purposes. Partner licenses have an expiration date. The license file specifies the number of licenses that you purchased. Contact your Cisco representative to purchase partner licenses.

Cisco ICFP Licensing Workflow

This workflow applies to all Cisco ICFP licenses except for the Cisco ICFP evaluation license. This workflow is not required for the 60-day evaluation license that is included with Cisco ICFP.

1. Before installing Cisco ICFP, locate your Cisco ICFP license and Product Authorization Key (PAK).
   To purchase a license, contact your Cisco representative.
2. Register the PAK on the Cisco software license site.
   For more information, see Generating a License Using a PAK, on page 26.
3. Install Cisco ICFP.
   For more information, see the following topics:
   - Installing Cisco ICFP on VMware, on page 13
   - Installing Cisco ICFP on OpenStack, on page 19
4. Upload the license in Cisco ICFP.
   For more information, see Uploading a License, on page 27.
5. Check license status.
   For more information, see Viewing License Details, on page 27.
6. Update the license.
   To update an existing license, use the procedure for uploading a license.

Generating a License Using a PAK

This procedure describes how to generate a Cisco ICFP license by using a PAK.

Before You Begin

Obtain the Cisco ICFP PAK.

Procedure

Step 1 In a browser, go to the Cisco Product License Registration page.
This page offers training on assigning PAKs or tokens, and a link to the Cisco Product License Registration tool.

Step 2 Click **Continue to Product License Registration**.

Step 3 In the **Product License Registration** screen, enter the PAK number in the **Get New Licenses** field, and click **Fulfill**.

Step 4 In the **Get New Licenses** dialog box, provide the required information and click **Submit**.
The status of your request is displayed, and a digital license agreement and a zipped license file are sent to the email address that you specified.

What to Do Next

Upload the license file in Cisco ICFP. For more information, see Uploading a License, on page 27.

Uploading a License

Use this procedure to upload a new license in Cisco ICFP or to update an existing license. To ensure continuous operation, be sure to update the license before the current license expires.

Before You Begin

- Obtain the Cisco ICFP license file. For more information, see Generating a License Using a PAK, on page 26.
- If you received a zipped license file by email, extract and save the `.lic` file to your local machine.

Procedure

Step 1 In the Cisco ICFP GUI, choose **License** and click the **Upload License** icon.

Step 2 In the dialog box, select the Cisco ICFP `.lic` file and click **Upload**.

Step 3 When prompted, click **Yes** to confirm the upload.
After the license file is successfully processed, a success message is displayed.

Viewing License Details

To view license details in the Cisco ICFP GUI, choose **License**.

The license details are displayed, including the license type, the license status, the number of HCUs supported, and the term of the license.

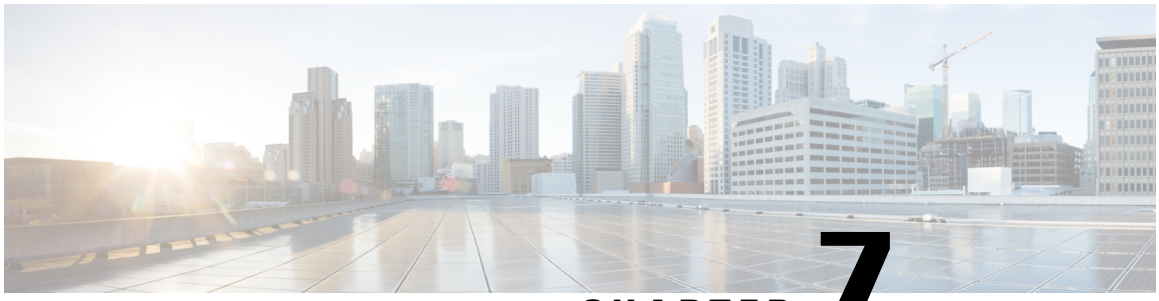

CHAPTER 7

Upgrading Cisco ICFP

Upgrading Standalone Nodes or Multiple-Node Clusters

Cisco ICFP enables you to upgrade standalone appliances and multiple-node clusters for bug fixes and updated adapters. For more information, see the following topics:

- Supported Upgrade Paths, on page 29
- Upgrading a Standalone Node, on page 30
- Upgrading a Multiple-Node Cluster, on page 31

Supported Upgrade Paths

Cisco ICFP 3.1.1 supports the following upgrade paths:

- OpenStack—Cisco ICFP 2.3.1 to 3.1.1.
- VMware—Cisco ICFP 2.3.1 to 3.1.1.

Restarting Services Automatically

Beginning with version 2.3.1, Cisco ICFP includes a feature that automatically restarts Infra services when you upgrade Cisco ICFP to a newer version.

When you upgrade Cisco ICFP from 2.3.1 to 3.1.1 or higher, the service restart feature is automatically enabled and you do not need to restart Infra services.

Upgrading a Standalone Node

This procedure enables you to upgrade Cisco ICFP to a newer version and apply Cisco bug fixes on a standalone node. To upgrade a multiple-node cluster, see Upgrading a Multiple-Node Cluster, on page 31.

Upgrading from Cisco ICFP version 2.3.1 to 3.1.1 automatically resets the admin account password to **changeme**. For information on changing the admin account password to another password, see Changing the Admin Account Password, on page 31.

Before You Begin

- Obtain the Cisco ICFP upgrade file (`icfp-upgrade-3.1.1.tar.gz`) from Cisco.com. For assistance, contact your Cisco representative.
- Confirm that the upgrade file is accessible from the Cisco ICFP virtual appliance.

Procedure

Step 1 In the Cisco ICFP GUI, choose **Upgrade Software**, and click the **Upgrade Adapter** icon.

Step 2 In the **Upgrade Adapter** dialog box, provide the following information:

Field	Description
Adapter Type	Choose **Cisco**.
Adapter Name	*Display only.* This field displays CAPI by default.
Adapter Description	Enter the desired description.
Adapter Version	Enter the new version.
Select File to Upload	Browse to the Cisco ICFP upgrade file and click **Upload**.

Step 3 When prompted, click **Yes** to confirm the upload.
When upgrading from Cisco ICFP 2.3.1 to 3.1.1 or higher, a message is displayed stating that the upgrade will start in 2 minutes. After approximately 2 minutes, the upgrade is installed, the services automatically restart, and the GUI becomes unresponsive.

Step 4 Finish the upgrade by refreshing the browser and logging in to the Cisco ICFP GUI.

Step 5 To verify that the upgrade was successful, click **About** in the GUI toolbar and confirm that the correct version is displayed.
Cisco ICFP displays the version, build number, build date, the last modification date, and the version hash value.

What to Do Next

If required, change the admin account password as described in Changing the Admin Account Password, on page 31.

Changing the Admin Account Password

Use this procedure to change the password for the Cisco ICFP admin account for standalone and multiple-node clusters as follows:

- To change the password for a standalone node, log in to the Cisco ICFP GUI for that node.
- To change the password for a multiple-node cluster, log in to the Cisco ICFP GUI for the active primary node in the cluster.

Before You Begin

You must have admin account access to perform this task.

Procedure

Step 1 In the Cisco ICFP toolbar, choose **Admin**.

Step 2 In the **Admin Log In** dialog box, enter the current credentials for logging in to the Cisco ICFP admin account and click **Login**.
If you have recently upgraded to Cisco ICFP 3.1.1, the password is **changeme**.

Step 3 In the **Admin Panel** dialog box, click the **Password** tab.

Step 4 Enter the new password in the **New Password** and **Confirm New Password** fields, and click **Apply**.
A success message indicates that the password has been successfully updated.

Step 5 Click **Close**.

Step 6 Log out of the Cisco ICFP GUI and log in again with the new password.

Upgrading a Multiple-Node Cluster

Use this procedure to upgrade a multiple-node cluster for bug fixes and updated adapters. To upgrade a standalone Cisco ICFP virtual appliance, see Upgrading a Standalone Node, on page 30.

Upgrading from Cisco ICFP version 2.3.1 to 3.1.1 automatically resets the admin account password to **changeme**. For information on changing the admin account password to another password, see Changing the Admin Account Password, on page 31.

This procedure applies to multiple-node clusters with the following components and configuration:

- An HA pair that:
  - Consists of two Cisco ICFP virtual appliances configured with the Primary Node role.
  - Is configured with one active node and one standby node.

- Additional Cisco ICFP virtual appliances that are configured as service nodes.

The workflow for upgrading a cluster includes the following high-level tasks:

1. Stop the virtual IP (VIP) service on the HA active node.
2. Monitor status while services fail over to the HA standby node.
3. Upgrade the current HA active node (originally the standby node).
4. Start the VIP service on the current HA standby node (originally the active node).
5. Stop the VIP service on the upgraded HA active node.
6. Monitor status while services fail over to the current HA standby node, making it the active node again.
7. Upgrade the current HA active node.
8. Start the VIP service on the current HA standby node.
9. Upgrade each service node.
10. If required, change the admin account password.

The following procedure describes how to perform these tasks.

Before You Begin

- Obtain the Cisco ICFP upgrade file (**`icfp-upgrade-3.1.1.tar.gz`**) from Cisco.com. For assistance, contact your Cisco representative.
- Confirm that the upgrade file is accessible from the Cisco ICFP virtual appliance.
- Confirm that HA has been configured on two Cisco ICFP virtual appliances that are configured with the Primary Node role.

Procedure

Step 1 Stop the VIP service on the HA active node as follows:

a) Log in to the ShellAdmin console for the HA 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.

Step 2 Log in to the ShellAdmin console for the HA standby node.

Step 3 In the ShellAdmin console for the standby node, choose **Display Services Status** to monitor the following events as they occur:

- HA services fail over to the standby node in the HA pair.
- Infra services start running on the standby node.
- The GUI for the standby node becomes available for logging in.

It can take a few minutes for the services to start and for the GUI of the standby node to be accessible from the browser.

Note The node that was originally the HA standby node becomes the HA active node.

Step 4 Upgrade the currently active node of the HA pair as follows:

a) Log in to the Cisco ICFP GUI for the active node of the HA pair by using the management IP address of the node.
b) In the GUI, choose **Upgrade Software** and click the **Upgrade Adapter** icon.
c) In the **Upgrade Adapter** dialog box, provide the required information.
For information about the fields in this dialog box, see Upgrading a Standalone Node, on page 30.
d) Click **Upload**.
e) When prompted, click **Yes** to confirm that you want to upload the selected file.

When upgrading from Cisco ICFP 2.3.1 to 3.1.1 or higher, the Infra services restart automatically and you can log in to Cisco ICFP after approximately 2 minutes.

Step 5 Verify that the HA active node was successfully upgraded as follows:

a) Log in to the Cisco ICFP GUI of the active node by using the management IP address of the node.
b) Click **About** in the Cisco ICFP toolbar.
c) Confirm that the correct version is displayed.

Step 6 Restart the VIP service on the current HA standby node as follows:

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

Step 7 Stop the VIP service on the currently active node that was upgraded in Step 4 as follows:

a) Log in to the Shell Admin console for the currently active node in the HA pair.
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 8 Log in to the ShellAdmin console for the standby node in the HA pair.

Step 9 In the ShellAdmin console for the standby node, choose **Display Services Status** to monitor the following events as they occur:

- HA services fail over to the standby node in the HA pair.
- Infra services start running on the standby node.
- The GUI for the standby node becomes available for logging in.

It can take a few minutes for the services to start and for the GUI of the standby node to be accessible from the browser.

Note The node that was previously the HA standby node becomes the HA active node.

Step 10 Upgrade the HA active node as follows:

a) Using the management IP address instead of the virtual IP address for the HA pair, log in to the Cisco ICFP GUI for the HA active node.
b) Upgrade the node as described in Step 4.
c) Verify that the upgrade was successful as described in Step 5.

Step 11 Restart the VIP service on the HA standby node as follows:

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

Step 12 Upgrade each service node in the cluster as follows:

a) Log in to the Cisco ICFP GUI for the service node.
b) Upgrade the service node by uploading the upgrade package as described in Step 4.
When upgrading from Cisco ICFP 2.3.1 to 3.1.1 or higher, the Infra services restart automatically and you can log in to the upgraded service node after approximately 2 minutes.

Step 13 Verify that each service node upgraded successfully as follows:

a) For each service node, refresh the browser and log in to the Cisco ICFP GUI for the service node.
b) Click **About** in the Cisco ICFP toolbar and confirm that the correct version is displayed.

Step 14 (Optional) If required, change the admin account password as described in Changing the Admin Account Password, on page 31.

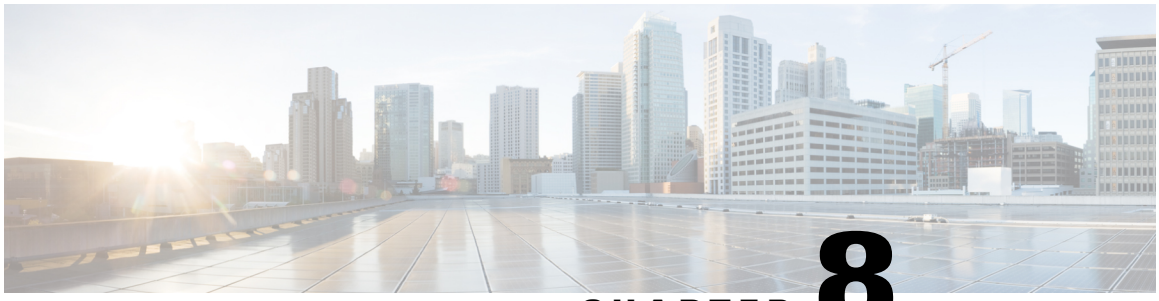

CHAPTER 8

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, on page 5
2.	Configure two primary nodes.	Configuring a Primary Node, on page 36
3.	Configure two or more service nodes.	Configuring a Service Node, on page 37
4.	Configure additional storage.	Configuring Additional Storage, on page 38

Step	Task	Related Information
5.	Configure the two primary nodes for HA.	Configuring HA, on page 41
6.	(OpenStack only) Configure VIP access.	Configuring VIP Access for HA in OpenStack, on page 43
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 37.

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 36.

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 38
- Configuring a Cinder Volume, on page 39

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 39.

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 37.

- 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, on page 5.
  - To configure an existing Cisco ICFP virtual appliance as a primary node, see Configuring a Primary Node, on page 36.
- 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 43.

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 35.

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, on page 14.
- 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, on page 20.
- Configure existing Cisco ICFP virtual appliances using the Standalone Node role as primary nodes. For more information, see Configuring a Primary Node, on page 36.

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 41.

Step 5 Configure service nodes for the cluster. For more information, see Configuring a Service Node, on page 37.

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.

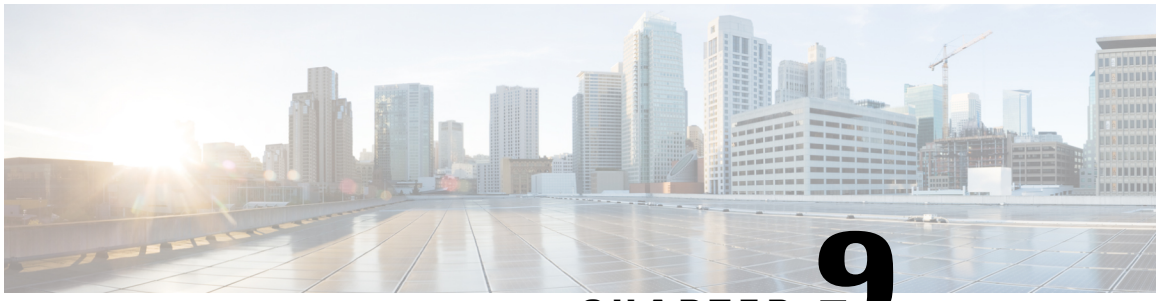

CHAPTER 9

Configuring VMware vCloud Director for Cisco ICFP

Configuring VMware vCloud Director

Installing Cisco ICFP at a cloud provider site enables you to support a hybrid cloud environment with Cisco Intercloud Fabric for Business. For VMware vCloud Director (VCD) environments, Cisco ICFP includes a built-in VCD adapter that enables Cisco ICFP to integrate with the VCD platform. This VCD-Cisco ICFP integration can be viewed as the infrastructure that binds the enterprise virtualization platform, such as VMware vCenter, to the provider cloud platform, VCD.

The following illustration depicts how Cisco Intercloud Fabric interfaces with the provider VCD platform through Cisco ICFP.

Figure 2: VCD and Cisco Intercloud Fabric Integration

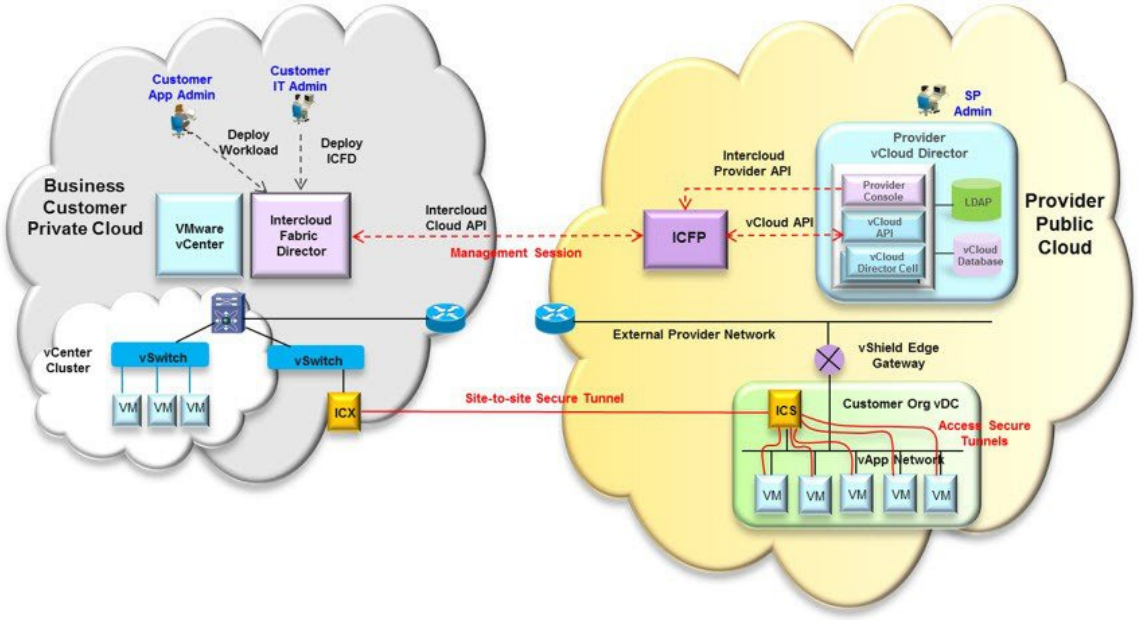

The secure site-to-site tunnel illustrated in the image is created between an Intercloud Fabric Switch (ICS) on the provider cloud and an Intercloud Fabric Extender (ICX) on the private cloud. In addition to providing secure communications between the private and provider clouds, this site-to-site tunnel enables Cisco Intercloud Fabric Secure Extender to integrate with VCD for each tenant network.

Before the ICS and ICX can communicate via the Internet, you must:

- Assign a public IP address to the ICS so that the ICX can reach the ICS.
- Ensure that the vShield Edge Gateway provides NAT functionality so that the ICS can connect to the Internet.

The following figure shows an example deployment:

Figure 3: vShield Edge Gateway Deployment Example

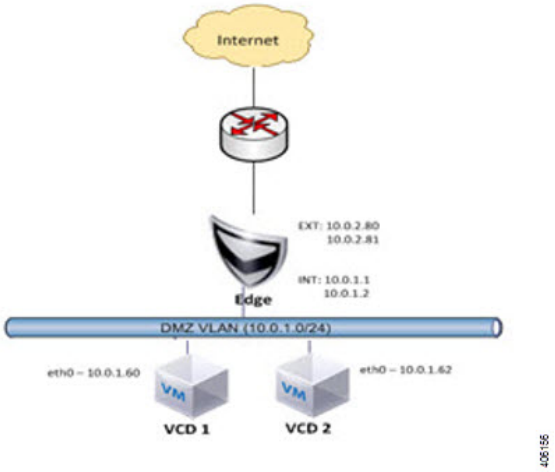

A vShield Edge Gateway is an interconnecting appliance that provides many edge network service features, including:

- DHCP
- Firewall
- IPsec VPN
- Load-balancer
- NAT
- Static route

The following figure shows how Organization X connects the Org Network to an external network through a vShield Edge Gateway and directly to vApp networks.

Figure 4: VCD Networking Model

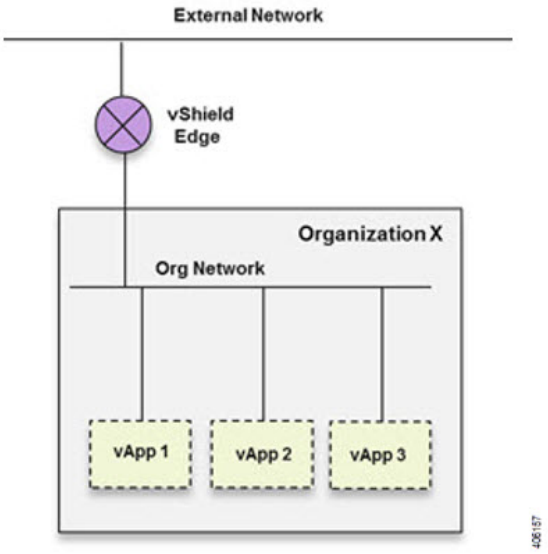

Workflow for Integrating VCD with Cisco ICFP

To integrate VCD with Cisco ICFP, you must provision certain infrastructure resources in the target VCD platform. The following table identifies the tasks required to provision these resources:

Step	Task	Related Information
1.	Ensure that the following prerequisites are met: • VCD version 5.5 is installed. • You have access to the VCD system administrator account.	VMware VCD documentation
2.	Create an external network.	Creating an External Network, on page 55
3.	Deploy the vShield Edge Gateway.	Adding a vShield Edge Gateway on an Org VDC, on page 56
4.	Create an Org VDC network.	Creating an Org VDC Internal Network, on page 57
5.	Create a catalog.	Creating a Catalog, on page 59
6.	Ensure that NAT and firewall services are configured on the vShield Edge Gateway.	Verifying NAT and Firewall Service Configuration, on page 59

For additional information on any of these topics, see your VMware documentation.

After you have successfully integrated VCD with Cisco ICFP, you can configure a cloud instance and add a tenant as described in Configuring Cisco ICFP for Cisco Intercloud Fabric, on page 61.

Creating an External Network

This procedure describes how to create an external network in a virtual data center (VDC).

Procedure

Step 1 Log in to the VCD GUI as system administrator.

Step 2 Choose **System** > **Manage & Monitor** > **Cloud Resources** > **External Networks**.

Step 3 In the **External Networks** pane, click **Add**.
The **New External Network** wizard opens, guiding you through the configuration process.

Step 4 In the **Select vSphere Network** screen, choose the VDC vCenter and the DVS port group created for the vSphere management network, and click **Next**.

Step 5 In the **Configure External Network** screen, click **Add**.

Step 6 In the **Add Subnet** dialog box, enter the following information for the external network:

- Gateway IP address
- Network mask
- DNS server IP address
- Static IP address or IP address range

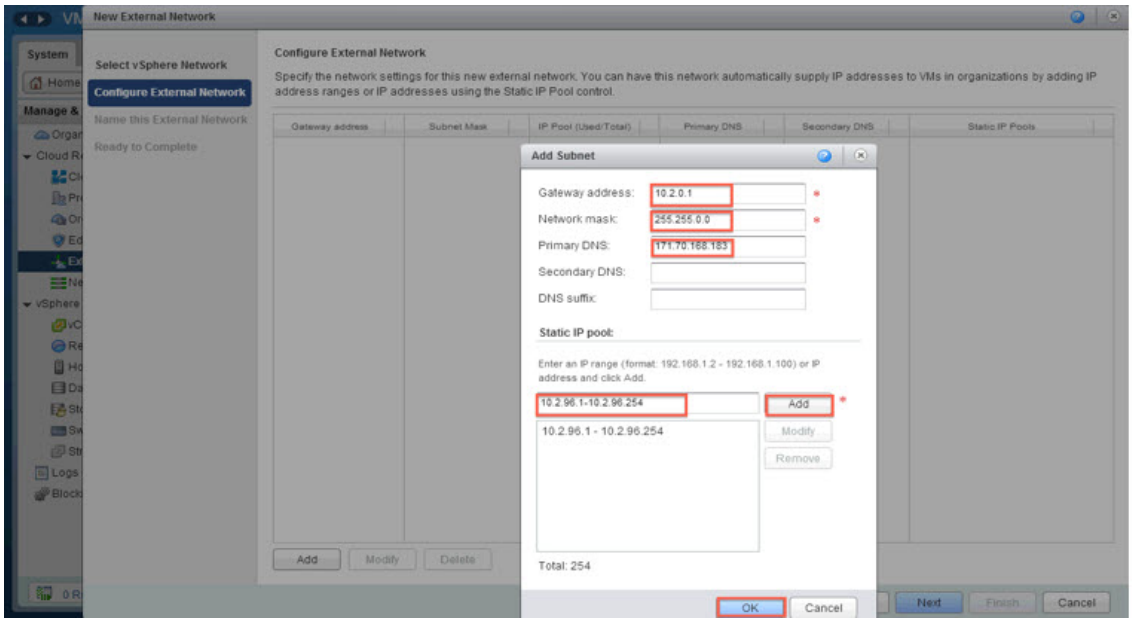

Step 7 In the **Name this External Network** screen, enter a name for the external network, and click **Next**.

Step 8 In the **Ready to Complete** screen, review the content for accuracy and click **Finish**.
The newly created external network is displayed in the **External Networks** pane.

Adding a vShield Edge Gateway on an Org VDC

You must add a vShield Edge Gateway to integrate the Provider VDC and Org VDC with Cisco ICFP.

Before You Begin

Confirm that the following have been configured:

- A Provider VDC
- An Org VDC
- An external network

Procedure

Step 1 In the VCD GUI, choose **System** > **Manage & Monitor** > **Cloud Resources** > **Organization VDCs**.

Step 2 In the **Organization VDCs** table, double-click the Org VDC where the vShield Edge Gateway is to be added.
The screen is refreshed with information about the selected VDC.

Step 3 Choose the **Edge Gateways** tab and click **Add**.
The **New Edge Gateway** wizard opens, guiding you through the configuration process.

Step 4 In the **Configure Edge Gateway** screen, configure the vShield Edge Gateway for connectivity with the external network as follows, and then click **Next**:

a) Choose the required edge gateway configuration: Compact, Full, or Full-4.
b) If the edge gateway is to be configured for HA, check the **Enable High Availability** check box.
c) In the **Advanced Options** section, check the **Sub-Allocate IP Pools** check box.

Step 5 In the **External Networks** screen, choose the external network that you created in Creating an External Network, on page 55 and click **Add**. If the external network is not listed, create a new external network.

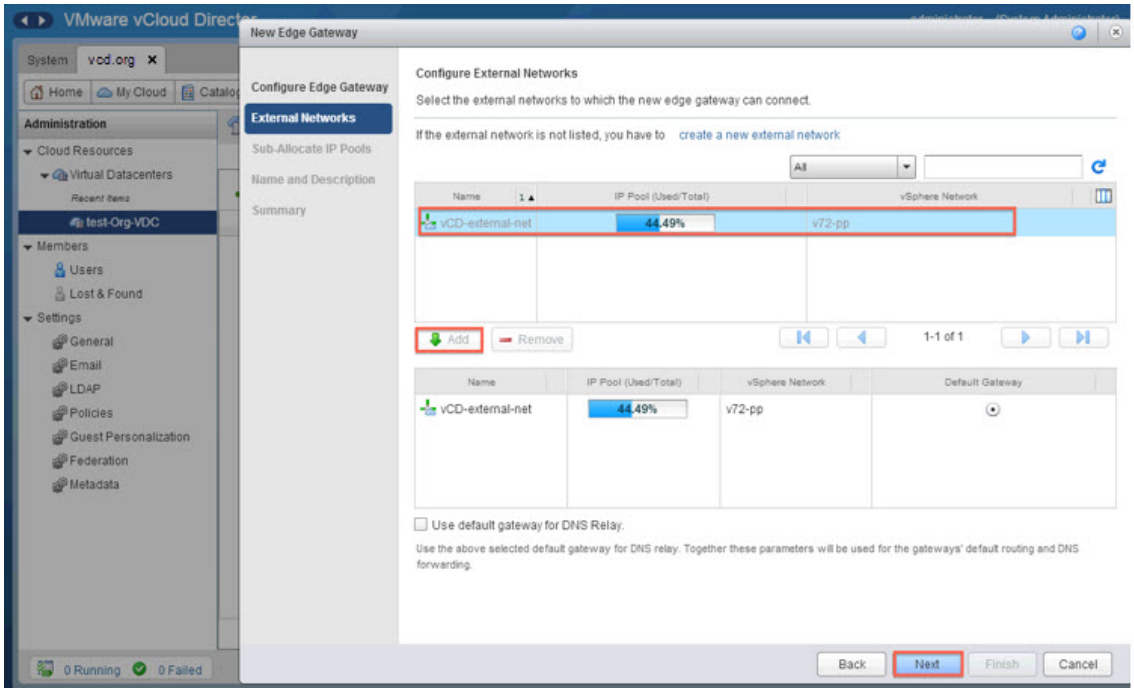

Step 6 After the external network is added to the list of networks in the lower portion of the screen, click **Next**.

Step 7 In the **Sub-Allocate IP Pools** screen, identify the range of IP addresses allocated for each externally-connected interface on the external network, and click **Next**.

Step 8 In the **Name and Description** screen, enter the edge gateway name and description, and then click **Next**.

Step 9 In the **Summary** screen, review the information for accuracy and click **Finish**.

Creating an Org VDC Internal Network

Use this procedure to create an internal network for the Org VDC.

Procedure

Step 1 In the VCD GUI, choose **System** > **Manage & Monitor** > **Cloud Resources** > **Organization VDCs**.

Step 2 In the **Organization VDCs** table, double-click the Org VDC where you want to create the internal network. The screen is refreshed with information about the selected VDC.

Step 3 In the **Org VDC Networks** tab, in the toolbar, choose **Actions** > **Add Network**.

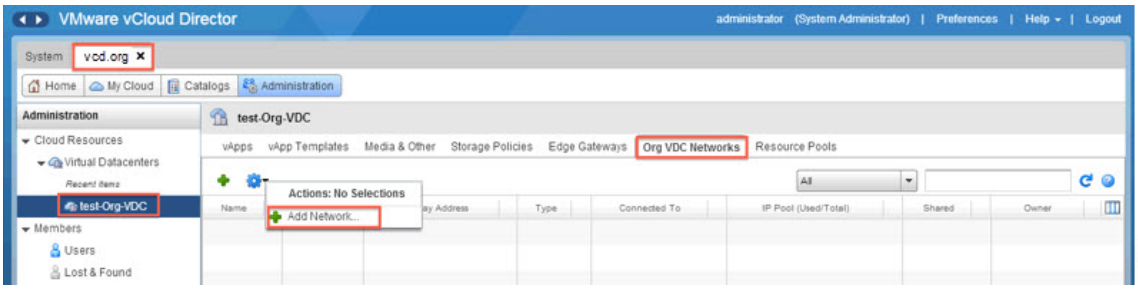

The **New Organization Network** wizard opens, guiding you through the configuration process.

Step 4 In the **Select Network Type** screen:

a) Choose **Create a routed network by connecting to an existing edge gateway**.

b) Choose the vShield Edge Gateway that you created in Adding a vShield Edge Gateway on an Org VDC, on page 56.

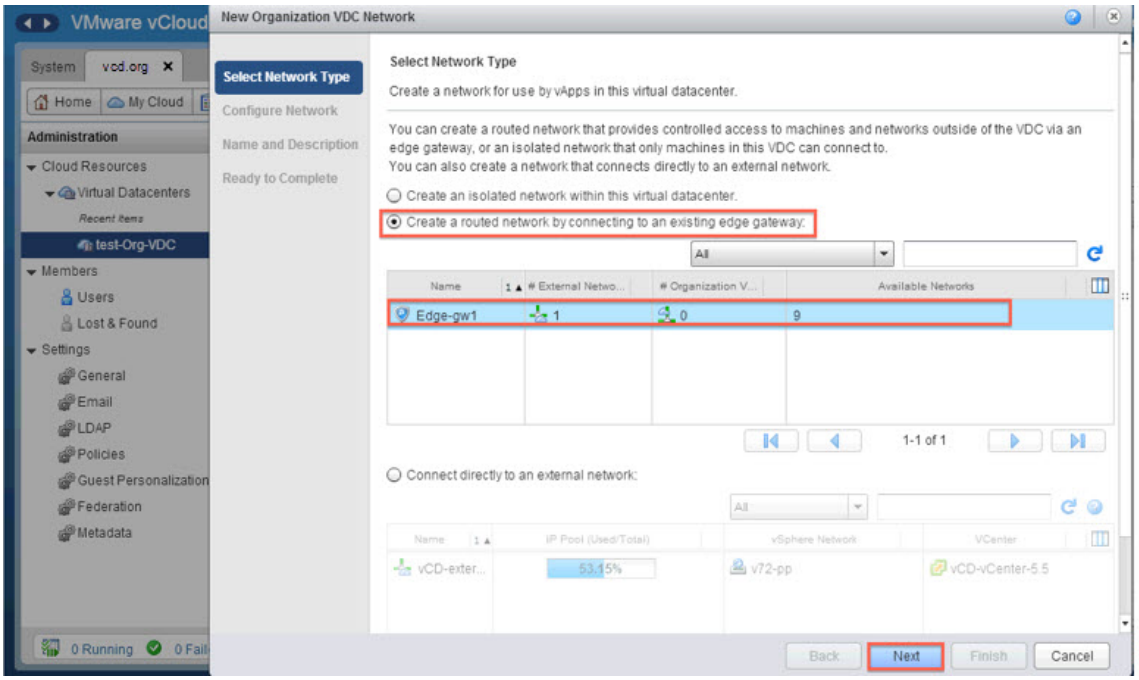

Step 5 In the **Configure Network** screen:

a) Enter the following information:

- Gateway IP address
- Network mask
- DNS server IP address

b) In the Static IP pool area, enter an IP address or an IP address range and click **Add**.

Step 6 In the **Name and Description** screen, enter a name and description (optional) for the Org VDC internal network.

Step 7 In the **Ready to Complete** screen, review the information for accuracy and click **Finish**.

Creating a Catalog

A catalog enables you to upload images from Cisco ICFP to VCD.

For additional information about creating catalogs and selecting options, see your VMware vCloud Director documentation.

Procedure

Step 1 In the VCD GUI, choose **System** > **Manage & Monitor** > **Cloud Resources** > **Organization VDCs**.

Step 2 In the **Organization VDCs** table, double-click the Org VDC in which to add the catalog.
The screen is refreshed with information about the selected VDC.

Step 3 Choose the **Catalogs** tab and, in the toolbar, choose **Actions** > **Add Catalog**.
A dialog box opens with multiple tabs so that you can configure the catalog and user access.

Step 4 In the **General** tab, enter a name and a description (optional) for the catalog.

Step 5 In the **Sharing** tab:

a) Click **Add Members**.

b) Choose the users or groups of users who can access the catalog.

c) In the **Access Level** field, choose the level of access for each user or group of users: Read-only, Read/Write, or Full Control.

Step 6 In the **Storage** tab, choose the type of storage.

Step 7 In the **Metadata** tab:

a) From the **Type** drop-down list, choose the metadata type.

b) In the **Name** field, enter a name for this metadata entry.

c) In the **User access of metadata** field, choose the level of access for the metadata: Read/Write, Read-only, or Hidden.

d) In the **Value** field, enter a text value for the metadata entry.

Step 8 After you have configured the catalog, click **OK**.

Verifying NAT and Firewall Service Configuration

When VCD is integrated with Cisco ICFP, NAT and firewall services are configured automatically, enabling the vShield Edge Gateway to communicate with the external network. This procedure enables you to confirm that NAT and firewall services have been configured on the vShield Edge Gateway as expected.

Procedure

Step 1 In the VCD GUI, choose **System** > **Manage & Monitor** > **Cloud Resources** > **Organization VDCs**.

Step 2 In the **Organization VDCs** table, double-click the Org VDC where you created the vShield Edge Gateway (Adding a vShield Edge Gateway on an Org VDC, on page 56).
The screen is refreshed with information about the selected VDC.

Step 3 In the **Edge Gateways** tab, right-click the required edge gateway and choose **Edge Gateway Services**.

Step 4 In the **Configure Services** dialog box, confirm the following:

a) In the **NAT** tab, confirm that Source NAT and Destination NAT rules are displayed, as shown in the following example:

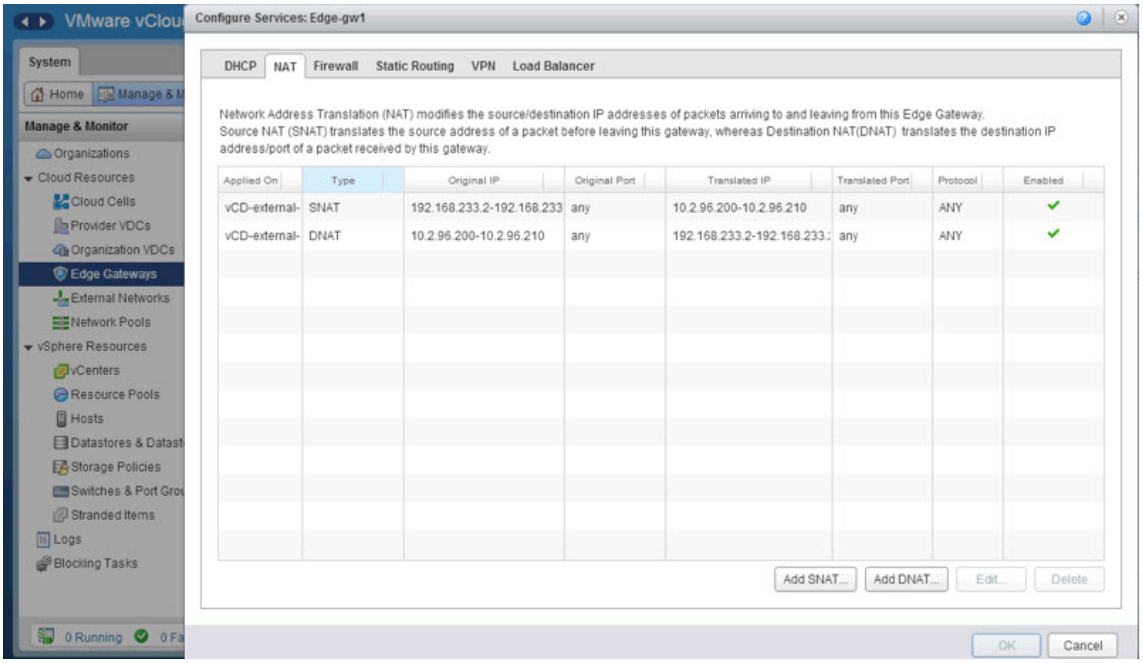

b) In the **Firewall** tab, confirm that inbound traffic is allowed for the following destination ports and protocols:

- 22—TCP
- 443—TCP
- 500—TCP, UDP
- 4500—TCP, UDP
- 6644—TCP, UDP
- 6646—TCP, UDP

The information should resemble the following example:

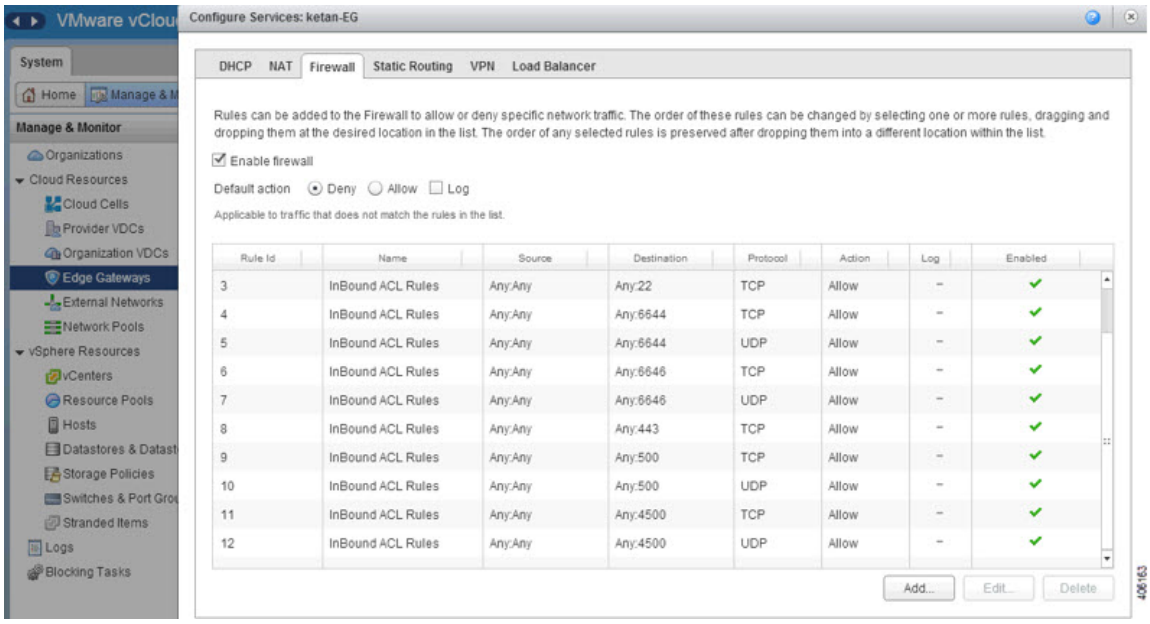

Configuring Cisco ICFP for Cisco Intercloud Fabric

After you have installed Cisco ICFP on a VMware server and launched a Cisco ICFP instance, you can configure Cisco ICFP for use with Cisco Intercloud Fabric.

Before You Begin

Confirm the following:

- Cisco ICFP has been installed on a VMware server and an instance has been launched.
- You know the Cisco ICFP public IP address.

Procedure

Step 1 In the Cisco ICFP GUI, choose **Cloud Instances**, and click the **Add Cloud** icon.

Step 2 In the **New Cloud Instance** dialog box, provide the following information, and click **Create**:

Field	Description
Cloud Instance Name	Name of the cloud instance.
Type	The cloud instance type: Cisco or Custom.

Field	Description
Module Name	For a Cisco cloud instance type, choose the module name, such as **VCDP** for VMware vCloud Director Platform. For a custom cloud instance, enter the custom module name.
Endpoint URI	The endpoint hostname or IP address of the cloud instance.

Step 3 In the Cisco ICFP GUI, choose **Tenant Accounts**, and click the **Add Tenant Account** icon.

Step 4 In the **New Tenant Account** dialog box, provide the following information, and click **Create**:

Field	Description
Tenant Name	Enter the tenant name. You cannot change the name after adding the tenant.
Select Cloud	Choose the name of the cloud instance that you created in the previous steps. You cannot change the cloud instance name after adding the tenant.
Org Name	For VMware vCloud Director clouds, enter the name of the organization to which the tenant belongs.
Max Servers	Enter the maximum number of servers provisioned for the tenant, including stopped VMs.
Username	Enter the tenant account username.
Email	Enter the tenant account email address.

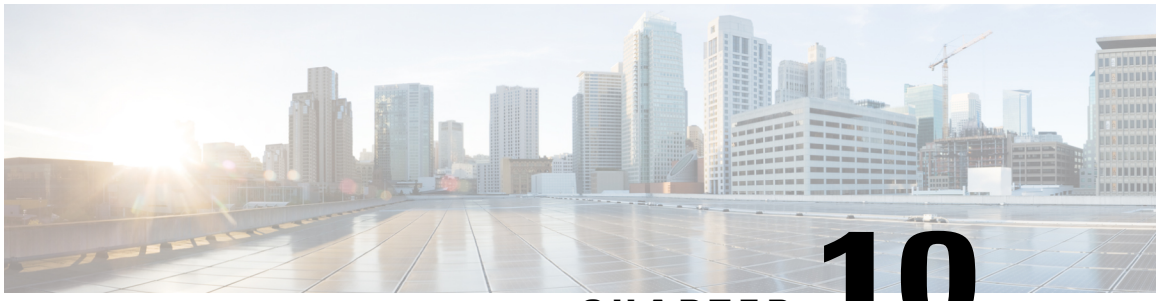

CHAPTER 10

Additional Information

Related Documentation

Cisco Intercloud Fabric for Provider

The Cisco Intercloud Fabric for Provider documentation is available at the following URL:

http://www.cisco.com/c/en/us/support/cloud-systems-management/intercloud-fabric/tsd-products-support-series-home.html

Cisco Intercloud Fabric for Business

The Cisco Intercloud Fabric for Business documentation is available at the following URL:

http://www.cisco.com/c/en/us/support/cloud-systems-management/intercloud-fabric/tsd-products-support-series-home.html

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the What's New in Cisco Product Documentation RSS feed. RSS feeds are a free service.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to: intercloud-fabric-doc-feedback@cisco.com.

We appreciate your feedback.