



Managing the F5 BIG-IP Load Balancer

- [F5 Load Balancing, page 2](#)
- [Understanding Load Balancing Terminology, page 3](#)
- [Adding a Network Element, page 4](#)
- [Viewing Application Services, page 5](#)
- [Creating an Application Service, page 5](#)
- [Virtual Servers, page 7](#)
- [Creating a Pool, page 12](#)
- [Adding Members to a Pool, page 13](#)
- [Viewing Pool Statistics, page 14](#)
- [Viewing Node Lists, page 15](#)
- [Viewing Node Statistics, page 16](#)
- [Viewing TCP Profiles, page 17](#)
- [Viewing UDP Profiles, page 17](#)
- [Partition, page 18](#)
- [Viewing Device Groups, page 19](#)
- [Creating a Device Group, page 19](#)
- [Viewing Traffic Groups, page 20](#)
- [Viewing Traffic Group Device Statistics, page 21](#)
- [Creating a Traffic Group, page 22](#)
- [Forcing a Traffic Group to a Standby State, page 23](#)
- [Viewing Device Information, page 23](#)
- [Viewing HA Groups, page 24](#)

F5 Load Balancing

Cisco UCS Director supports the creation and monitoring of F5 load balancers.

Although load balancing may be prevalent in the routing environment, it is also of growing importance in the virtual networking and VM environment. Server load balancing is a mechanism for distributing traffic across multiple virtual servers, offering high application and server resource utilization.

Server load balancing (SLB) is the process of deciding to which server a load-balancing device should send a client request for service. For example, a client request can consist of an HTTP GET for a web page or an FTP GET to download a file. The job of the load balancer is to select the server that can successfully fulfill the client request and do so in the shortest amount of time without overloading either the server or the server farm as a whole.

Depending on the load-balancing algorithm or predictor that you configure, the F5 BIG-IP performs a series of checks and calculations to determine the server that can best service each client request. F5 BIG-IP bases server selection on several factors, including the server with the fewest connections regarding load, source or destination address, cookies, URLs, or HTTP headers.

A high-level process flow of load balancing is as follows:

- 1 A client attempts to connect with a service on the load balancer.
- 2 The load balancer accepts the connection.
- 3 The load balancer decides which host should receive the connection and changes the destination IP address (or port) in order to match the service of the selected host.
- 4 The host accepts the load balancer's connection and responds to the original source, to the client (through its default route), and to the load balancer.
- 5 The load balancer acquires the return packet from the host and changes the source IP address (or port) to correspond to the virtual server IP address and port, and forwards the packet back to the client.
- 6 The client receives the return packet, assuming it came from the virtual server, and continues the rest of the process.

Cisco UCS Director enables the management, orchestration, and monitoring of the F5 load balancer. Following is a summary of the crucial processes:

- 1 Add the F5 load balancer. To add the F5 load balancer, choose **Administration > Physical Accounts**. On the **Physical Accounts** page, click **Managed Network Elements** and then click **Add Network Element**.
- 2 On adding the F5 load balancer as a managed element, Cisco UCS Director triggers Cisco UCS Director task inventory collection. The polling interval configured on the **System Tasks** specifies the frequency of inventory collection.
- 3 After the F5 load balancer is added to the Pod, it is listed with all other components of the pod environment at the account level. To see the F5 component information, choose **Physical > Network**. On the **Network** page, choose the Pod and click **Managed Network Elements**.

There are two ways to implement load balancing on an F5 device using Cisco UCS Director:

- 1 Use an iApps (BIG-IP) application service.

iApps application templates let you configure the BIG-IP system for your HTTP applications, by functioning as an interface to consistently deploy, manage, and monitor your servers. You can use default iApps templates or create and customize a template to implement load balancing on the F5 device.

- 2 Use Cisco UCS Director to:
 - Set up a managed element
 - Create a Pool
 - Add pool members
 - Create a virtual server

Understanding Load Balancing Terminology

This section describes common terms found in the F5 BIG-IP environment. In a load-balancing environment, a virtual server is a construct that allows multiple physical servers to appear as one for load-balancing purposes. A virtual server is bound to physical services running on real servers in a server farm and uses IP address and port information to distribute incoming client requests to the servers in the server farm according to a specified load-balancing algorithm.

- **Virtual servers**—In a load-balancing environment, a virtual server is a construct that allows multiple physical servers to appear as one for load-balancing purposes. A virtual server is bound to physical services running on real servers in a server farm and uses IP address and port information to distribute incoming client requests to the servers in the server farm according to a specified load-balancing algorithm.
- **Pools**—A pool is a collection of virtual servers that provide similar services available on multiple hosts. (See the pool members or nodes entry for additional information.)
- **Pool members or nodes**—When creating a pool, you assign one or more pool members to it. A pool member or node is a logical object that represents a physical node (and a service) on the network. When you add a virtual server to a pool, it becomes a pool member. A member or node includes the TCP port of the actual application that is receiving traffic.



Tip A virtual server can be a member of multiple pools. In a different pool, it can have different attributes and play a different role. For example, a virtual server could be a backup resource for a different type of requests, such as requests from a different part of the world.

- **Nodes**—Physical servers that receive traffic from a load balancer.
- **Profiles**—A profile can be either local or roaming. We recommended that you operate using roaming profiles rather than local profiles. Using roaming profiles assures you that your settings are always available to you at all times.

Adding a Network Element

In order to create a virtual server that supports load balancing, first add a network element in Cisco UCS Director. After a Load Balancer is added as a network element in Cisco UCS Director, it appears on the **Managed Network Element** screen.

Before You Begin

You must be logged in to the appliance to complete this task.

Step 1 Choose **Administration > Physical Accounts**.

Step 2 On the **Physical Accounts** page, click **Managed Network Elements**.

Step 3 Click **Add Network Element**.

Step 4 On the **Add Network Element** screen, complete the following fields:

Name	Description
Pod drop-down list	Choose the pod to which the network element belongs.
Device Category drop-down list	Choose the device category for this network element. For example: F5 Load Balancer .
Device IP field	The IP address for this device.
Protocol drop-down list	Choose the protocol to be used. The list may include the following: <ul style="list-style-type: none"> • Telnet • SSH • HTTP • HTTPS <p>Note When working with an F5 load balancer device, HTTP and HTTPS are the only valid selections.</p>
Port field	The port to use.
Login field	The login name.
Password field	The password associated with the login name.

Step 5 Click **Submit**.

Adding the F5 Load Balancer triggers the system task inventory collection. The polling interval configured on the **System Tasks** screen specifies the frequency of inventory collection.

What to Do Next

To modify or edit a virtual server, choose the server, and then click **Modify**. To remove a virtual server, choose the server, and then click **Delete**.

Viewing Application Services

The **Application Services** screen provides information on the following items:

- Pod Name
- Name
- Template
- Partition/Path
- Device Group
- Traffic Group
- Active Device



Note F5 BIG-IP devices with versions lower than 12.0 show a blank Active Device column.

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **Application Services** to view information for the existing application services.
-

Creating an Application Service

This topic describes how to create an application service for a selected pod.

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose a pod that includes the F5 network, expand the directory as necessary, and click the F5 network.
- Step 3** Click **Application Services**.
- Step 4** Click **Create**.
- Step 5** On the **Create Application Service** screen, complete the following fields:

Name	Description
Select F5 Partition	The partition name.
Application Service Name field	The virtual server name.
Template field	<p>The name of the template.</p> <p>Note When working with an F5 network, this value should be set as f5.http.</p>
Virtual Server IP field	The IP address of the destination device.
FQDN names of Virtual Server field	<p>The fully qualified domain name (FQDN) names of the virtual server. Separate each FQDN name with a comma.</p> <p>Note Clients can use the FQDNs that you enter here to access the web servers. For each FQDN, your DNS Administrator must configure a DNS entry that resolves to the IP address you entered for the BIG-IP virtual server.</p>
Nodes List	<p>Select a node from the Nodes list and click Submit.</p> <p>If a node that you want to associate with the Virtual Server does not appear in the Nodes list:</p> <ul style="list-style-type: none"> • Click + to add it. The Add Entry to Nodes lists screen appears. • Provide the Node IP address, the Port, and the Connection limit; then click Submit.
Inherit Device Group from current Partition/Path check box	Check to automatically inherit the device group from the current partition or path. This option is selected by default.
Device Group	Select the device group to associate with the application service.
Inherit Traffic Group from current Partition/Path check box	Check to automatically inherit the traffic group from the current partition or path. This option is selected by default
Traffic Group	Select the traffic group to associate with the application service.

Step 6 Click **Submit**.

Virtual Servers

In a load-balancing environment, a virtual server is a construct that allows multiple physical servers to appear as one for load-balancing purposes. A virtual server is bound to physical services running on real servers in a server farm. The virtual server uses IP address and port information to distribute incoming client requests to the servers in the server farm according to a specified load-balancing algorithm.

Although the virtual server is of primary importance because it is used to administer pools and pool members, the practical flow of performing the setup is as follows:

- 1 Create a pool
- 2 Add members to pool
- 3 Create a virtual server that uses the pool



Remember

Before you can create a virtual server that supports load balancing, you must add the F5 load balancer as a network element.

Creating a Virtual Server



Tip

To get inventory information about the F5 device, navigate to **Converged > Pod**, then click the middle of the large Pod icon. In the row of Network images, double-click BIG-IP to see current information about it.

Before You Begin

An account with the F5 BIG-IP server.

Step 1 Choose **Physical > Network**.

Step 2 Click a pod in which one of the Managed Network Elements is the F5 BIG-IP device you want to use, then click the network.

With a Pod highlighted under **Physical > Network** and the Managed Network Elements selected, a row of information in the table include the IP address for the F5 BIG-IP device you want to. In the left column, click the network associated with this device.

Step 3 Click **Virtual Servers** and then click **Create**.

Step 4 On the **Create Virtual Server** complete the following fields:

Name	Description
Virtual Server Name field	The virtual server name.
Virtual Server Description field	A unique description of this virtual server.

Name	Description
Virtual Server Type field	The type of virtual server (preselected).
Destination IP field	The IP address of the destination device, the virtual server. This is a network address, which the system uses this network address with Mask to represent a range of IP addresses.
Destination IP Address Mask field	The IP address mask of the destination device.
Service Port drop-down list	<p>The data transfer protocol associated with the service port</p> <p>Choose a protocol from the following list:</p> <ul style="list-style-type: none"> • FTP • HTTP • HTTPS • TELNET • PPTP • SMTP • SNMP • SNMP-TRAP • SSH • Other
Port Number	The port number to be assigned to this server.
Protocol drop-down list	<p>Choose a protocol for high-speed data transfer from the following list:</p> <ul style="list-style-type: none"> • TCP • UDP <p>Note TCP is the default value for Protocol.</p>

Name	Description
Protocol Profile (Client) button	<p>Choose a client protocol profile with the correct Parent Profile for the Client. The parent profiles are TCP and UDP.</p> <p>The Client Protocol here specifies that the selected profile is a client-side profile. The drop-down list contains entries for each client protocol profile that has been defined.</p> <p>Remember The profile selection applies to TCP and UDP connections only. If TCP is set as Protocol, the TCP-related profile should be selected in client and server-side profiles. Likewise, UDP-related profiles should be selected if UDP is selected as the protocol. Hence parent profiles are TCP and UDP.</p>
Protocol Profile (Server) button	<p>Choose a server protocol profile with the correct Parent Profile for the Server.</p> <p>Server Protocol: Specifies that the selected profile is a server-side profile. Options are: (Use Client Profile) - the default - and entries for each defined server protocol profile.</p>
Default Pool button	<p>Choose a default pool.</p> <p>Note If you have created a virtual server for load balancing, you must assign a default load balancing pool to this virtual server. A default pool is the pool to which the BIG-IP system sends traffic if no iRule exists specifying a different pool. When you first create the virtual server, you assign an existing default pool to it. From then on, the virtual server automatically directs traffic to that default pool. It is not strictly mandatory to add default pool while creating a virtual server, so there is no default.</p>

Step 5 Click **Submit**.

Viewing Virtual Servers

The **Virtual Servers** screen provides information on the following items:

- Pod Name
- Virtual Server Name
- Status

- Destination IP Address
- Service Port
- Type
- Partition/Path
- Active Device



Note F5 BIG-IP devices with versions lower than 12.0 show a blank Active Device column.

Before You Begin

Create a Virtual Server.

-
- Step 1** Choose **Physical > Network**.
 - Step 2** Choose a pod that includes a load balancing network.
 - Step 3** Choose the load balancing server.
 - Step 4** Click **Virtual Servers** to view the F5 server.
-

Viewing Virtual Server Statistics

The **Virtual Servers Statistics** screen provides information on the following items:

- Pod Name
- Virtual Server
- Status
- Partition/Path
- Bits In
- Bits Out
- Packets In
- Packets Out
- Current Connections
- Maximum Connections
- Total Request
- CPU Utilization

Before You Begin

Create a Virtual Server.

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **Virtual Servers Statistics** to view the statistics.
-

Viewing Virtual Addresses

The **Virtual Address** screen provides information on the following items:

- Pod Name
- Name
- Status
- Address
- Partition/Path
- Traffic Group
- Active Device



Note F5 BIG-IP devices with versions lower than 12.0 show a blank Active Device column.

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **Virtual Address** to view information for the virtual addresses associated with the pod.
-

Updating Traffic Groups for Virtual Addresses

You can update the traffic group for a virtual address associated with the pod.

-
- Step 1** Choose **Physical > Network**.
 - Step 2** Choose a pod.
 - Step 3** Choose the load balancing server.
 - Step 4** Click **Virtual Address**.
 - Step 5** Choose the virtual address for which you want to update the traffic group, and click **Update Traffic Group**.
 - Step 6** On the **Update Traffic Group** screen, you can select the new traffic group, or check **Inherit Traffic Group from current Partition/Path**.
 - Step 7** Click **Submit**.
-

Creating a Pool

The **Pool list** screen enables you to create a new pool to the list of pools associated with an F5 load balancer. It also allows you to delete a pool, or to modify a pool by adding or deleting pool members.

The **Pool list** screen lists the following data for each pool:

- Pod Name
- Server Name
- Status
- Members
- Partition/Path

-
- Step 1** Choose **Physical > Network**.
 - Step 2** Click the pod that supports the F5 BIG-IP network that you want to use, then click that network.
 - Step 3** Click **Pool List** to view the existing list of pools.
Using controls that appear in the screen, you can also delete a pool from this list, see the members of any selected pool, or add or delete members of an existing pool.
 - Step 4** Click **Create**.
 - Step 5** On the **Create Pool** screen, complete the following fields:

Name	Description
Pool Name field	The Pool name.

Name	Description
Load Balancing Method drop-down list	Choose a load balancing method from the following list: <ul style="list-style-type: none"> • Round Robin • Ratio (member) • Lease Connections (member) • Observed (member) • Predictive (member) • Ratio (node)
Use Existing Nodes List selector	From the Existing Nodes list, add a node to the named pool. If the node you want to use does not yet appear in the Existing Nodes List, click + to add it. A screen appears, titled Add Entry to Existing Node List . Provide the Node Name and the port that you want this node to use, then click Submit . Choose the node. Then click Submit . Note For the purposes of setting up an F5 pool, the node name here does not have to use IPO address format. If a node is created by a virtual server creation service or by an application creation service, then the name of the node is set as the IP address of the node. If a node is created using the Create Node option, a name of another format may be displayed.

Step 6 Click **Submit**.

Step 7 To see additional details about available pools and nodes, see **Pool Statistics**, **Node list**, and **Node Statistics** screens.

Adding Members to a Pool

Before You Begin

Create a Pool.

Step 1 Choose **Physical > Network**.

Step 2 Click the pod that supports the F5 BIG-IP network of interest to you, then click that network.

Step 3 Click **Pool List** to view the existing list of pools.

Using controls that appear in this screen, you can also create or delete a pool from this list, or see the members of any selected pool.

Step 4 Click the pool in which you want to add, delete, or modify members, then click **View Details**. The **Pool Members** screen appears, listing the members of the pool you most recently selected.

Step 5 To add a member to the selected pool, click **Add**. The **Add Pool Member** screen appears.

Step 6 If you want to use an existing Node, then check the box labeled "Do you want to use existing node?" Otherwise, provide the Address and Service port.

- a) Check the box labeled "Do you want to use existing node?"
- b) Provide the following information.
- c) On the **Add Pool Member** screen, complete the following fields:

Name	Description
Node Name selector.	<ul style="list-style-type: none"> • Click Select... to open a list of existing nodes associated with the Pool-related account. • Checkmark a node to select it, then click Select. <p>The selected node appears in the Add Pool Member screen.</p>
Service Port	Enter the service port for the node.

d) Click **Submit**.

A message confirms that the member was added to the pool successfully.

Step 7 If you do not want to use an existing Node, then provide the Address and Service Port data and click Submit. A message confirms that the member was added to the pool successfully and the new member is listed.

Step 8 To return to the **Pool List** screen, click the **Back** button.

Viewing Pool Statistics

The **Pool Statistics** screen provides the following information:

- Pod Name
- Pool Name
- Status
- Partition/Path
- Bits In
- Bits Out

- Packets In
- Packets Out
- Current Connections
- Maximum Connections
- Total Connections
- Total Requests
- Request Queue Depth
- Request Queue Max Age

Before You Begin

Create a Virtual Server.

-
- Step 1** Choose **Physical > Network**.
 - Step 2** In the left pane, choose a pod that includes an F5 BIG-IP network.
 - Step 3** Expand the pod, then choose the F5 BIG-IP load balancer account.
 - Step 4** Click **Pool Statistics** to view the existing information.
-

Viewing Node Lists

The **Node list** screen provides information on the following items:

- Pod Name
- Server Name
- Status
- Description
- IP Address
- Partition/Path

Before You Begin

Create a Virtual Server.

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **Node List** to view the existing node list information.
-

Viewing Node Statistics

The **Node Statistics** screen provides the following information:

- Pod Name
- Node Name
- Status
- Partition/Path
- Bits In
- Bits Out
- Packets In
- Packets Out
- Current Connections
- Maximum Connections
- Total Connections
- Total Requests

Before You Begin

Create a Virtual Server.

-
- Step 1** Choose **Physical > Network**.
- Step 2** In the left pane, choose a pod that includes an F5 BIG-IP network.
- Step 3** Expand the pod, then choose the F5 BIG-IP load balancer account.
- Step 4** Click **Node Statistics** to view information for the existing nodes.
-

Viewing TCP Profiles

The **TCP Profiles** screen provides information on the following items:

- Pod Name
- Server Name
- Parent Profile
- Partition/Path

Before You Begin

Create a Virtual Server.

-
- Step 1** Choose **Physical > Network**.
- Step 2** In the left pane, choose a pod that includes an F5 BIG-IP network.
- Step 3** Expand the pod, then choose the F5 BIG-IP load balancer account.
- Step 4** Click **TCP Profiles** to view the existing list of TCP profiles.
-

Viewing UDP Profiles

The **UDP Profiles** screen provides information on the following items:

- Pod Name
- Server Name
- Parent Profile
- Partition/Path

-
- Step 1** Choose **Physical > Network**.
- Step 2** In the left pane, choose a pod that includes an F5 BIG-IP network.
- Step 3** Expand the pod, then choose the F5 BIG-IP load balancer account.
- Step 4** Click **UDP Profiles** to view the existing list of UDP profiles.
-

Partition

A partition is a logical container that you can create to contain a defined set of BIG-IP system objects. As an Administrator to the BIG-IP system, you can create administrative partitions to control other users' access to BIG-IP objects. When a specific set of objects resides in a partition, you can give certain users the authority to view and manage the objects in that partition only, rather than to all objects on the BIG-IP system.

Creating a Partition

In order to create a partition, perform the following procedure:

Before You Begin

You must be logged in to the appliance to complete the task.

-
- Step 1** On the menu bar, choose **Physical > Network**.
- Step 2** Choose **Managed Network Elements**.
- Step 3** Select a device from the **Unassigned Pods** in the left pane.
- Step 4** Double click the element or select it from F5.
- Step 5** Click **Partition** and a list of previously created partitions are displayed.
- Step 6** Click **Create** to add a partition.
- Step 7** On the **Create Partition** screen, complete the following fields:

Name	Description
Partition Name field	A unique name for the partition.
Description field	A detailed description of the partition.

- Step 8** Click **Submit**.
The partition is created successfully.

- Step 9** Click **OK**.

- Note**
- The report displays the name or path of the partition you created in the **Partition/Path** column.
 - You can also create partitions from the F5 appliance. Go to **System > Users > Partition List** and click **Create**.
-

What to Do Next

You can delete partitions you have created. To delete partitions, click **Delete** under **Partition** or in the F5 appliance, go to **System > Users > Partition List**, select the partition, and click **Delete**.

Viewing Device Groups

The **Device Group** report provides information on the following items:

- Account Name
- Pod Name
- Name
- App Service
- ASM Synchronize
- Auto Synchronize
- Full Load On Synchronize
- Incremental Config Synchronize
- Network Failover
- Save On Auto Synchronize

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **Device Group** to view information for the existing device groups.
-

Creating a Device Group

Cisco UCS Director lets you create a device group that runs on the F5 BIG-IP server. You can create two types of device groups. A sync-failover device group contains devices that synchronize configuration data and support traffic groups for failover purposes. A sync-only device group contains devices that synchronize configuration data, but do not synchronize failover objects.

Before You Begin

An account with the F5 BIG-IP server.

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose pod in which one of the Managed Network Elements is the F5 BIG-IP device you want to use, then click that network.
- Step 3** Click **Device Group** and then click **Create**.
- Step 4** On the **Create Device Group** screen, complete the following fields:

Name	Description
Select F5 Partition	The partition name.
Name field	A unique name for the device group.
Description field	A detailed description of the device group.
ASM Synchronize check box	Check to enable automatic synchronization of security policies and Application Security Manager (ASM) configurations.
Group Type drop-down list	Choose the device group type from the following list: <ul style="list-style-type: none"> • sync-only • sync-failover <p>Note Sync-only is the default value for the device group type.</p>
Auto Synchronize check box	Check to automatically synchronize configuration changes between devices in the same device group.
Full Load On Synchronize check box	Check to enable full load synchronization.
Incremental Config Synchronize Size Max field	This value specifies the total size of configuration changes that can reside in the incremental sync cache. 1024 (KB) is the default value.
Save On Auto Synchronize check box	Check to save the configuration to file during the automatic synchronization.
Members	Choose accounts to add to the device group.

Step 5 Click **Submit**.

What to Do Next

You can modify a device group or delete a device group by selecting the device group and clicking **Modify** or **Delete** respectively.

Viewing Traffic Groups

The **Traffic Group** report provides information on the following items:

- Account Name

- Pod Name
- Name
- Auto Failback Enabled
- Auto Failback Time
- Default Device
- HA Group
- HA Load Factor
- MAC
- Unit ID
- HA Order

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **Traffic Group** to view information for the existing traffic groups.
-

Viewing Traffic Group Device Statistics

The **Traffic Group Device Stats** report provides information on the following items:

- Account Name
- Pod Name
- Partition
- Device Name
- Failover State
- Next Active
- Traffic Group

-
- Step 1** Choose **Physical > Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **Traffic Group**.
- Step 5** Choose the traffic group for which you want to view the traffic group device statistics and click **View Details**.
-

Creating a Traffic Group

Cisco UCS Director lets you create a traffic group that runs on the F5 BIG-IP server.

Before You Begin

An account with the F5 BIG-IP server.

Step 1 Choose **Physical > Network**.

Step 2 Choose a pod in which one of the Managed Network Elements is the F5 BIG-IP device you want to use, then click that network.

Step 3 Click **Traffic Group** and then click **Create**.

Step 4 On the **Create Traffic Group** screen, complete the following fields:

Name	Description
Select F5 Partition	The partition name.
Name field	A unique name for the traffic group.
Description field	A detailed description of the traffic group.
Failover Method drop-down list	<p>Choose a failover method from the following list:</p> <ul style="list-style-type: none"> • HA Group • HA Order • Load Aware <p>Note Load Aware is the default value for the failover method. If you select HA Group, select the HA Group to failover to. If you select HA Order, specify the failover order in the Failover Order table.</p>
HA Load Factor field	The value that represents the application load for this traffic group relative to other traffic groups on the local device. The HA load factor range must be between 1-1000.
Auto Failback Time field	The number of seconds after which auto failback expires. The default value is 60.
MAC Address field	A unique MAC address that functions as the MAC masquerade address and floats on failover to minimize dropped connections

Step 5 Click **Submit**.

What to Do Next

You can modify a traffic group or delete a traffic group by selecting the traffic group and clicking **Modify** or **Delete** respectively.

Forcing a Traffic Group to a Standby State

Cisco UCS Director lets you force a traffic group to a standby state. By forcing the traffic group into a standby state, the traffic group becomes idle on the device, and becomes active on another device in the device group.

Step 1 Choose **Physical > Network**.

Step 2 Choose a pod in which one of the Managed Network Elements is the F5 BIG-IP device you want to use, then click that network.

Step 3 Click **Traffic Group**.

Step 4 Choose the traffic group that you want to force into a standby state and click **Force Standby**.

Step 5 On the **Force Traffic Group To Standby** screen, click **Submit**.

Viewing Device Information

The **Device Info** report provides information on the following items:

- Account Name
- Pod Name
- Status
- Name
- IP Address
- Host Name
- Product

- Version

-
- Step 1** Choose **Physical** > **Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **Device Info** to view information for the existing devices associated with the pod.
-

Viewing HA Groups

The **HA Group** report provides information on the following items:

- Account Name
- Pod Name
- Full Path
- Name

-
- Step 1** Choose **Physical** > **Network**.
- Step 2** Choose a pod.
- Step 3** Choose the load balancing server.
- Step 4** Click **HA Group** to view information for the existing HA groups associated with the pod.
-