

BIG-IP® TMOS®: Implementations

Version 12.1



Table of Contents

Customizing the BIG-IP Dashboard.....	11
Overview: BIG-IP dashboard customization.....	11
Customizing the BIG-IP dashboard.....	11
 Creating an Active-Standby Configuration Using the Setup Utility.....	 13
Overview: Creating a basic active-standby configuration.....	13
Task summary.....	13
Licensing and provisioning the BIG-IP system.....	14
Configuring a device certificate.....	14
Configuring the management port and administrative user accounts.....	14
Enabling ConfigSync and high availability.....	15
Configuring the internal network.....	15
Configuring the external network.....	16
Configuring the network for high availability.....	17
Configuring a ConfigSync address.....	17
Configuring failover and mirroring addresses.....	17
Discovering a peer device.....	18
Implementation result.....	18
 Creating an Active-Active Configuration Using the Setup Utility.....	 19
Overview: Creating a basic active-active configuration.....	19
Task summary.....	20
Licensing and provisioning the BIG-IP system.....	21
Configuring a device certificate.....	21
Configuring the management port and administrative user accounts.....	21
Enabling ConfigSync and high availability.....	22
Configuring the internal network.....	22
Configuring the external network.....	23
Configuring the network for high availability.....	24
Configuring a ConfigSync address.....	24
Configuring failover and mirroring addresses.....	24
Establishing device trust.....	25
Creating a Sync-Failover device group.....	26
Creating an iApp application for the local device.....	26
Creating a traffic group for a remote device.....	27
Creating an iApp application for a remote device.....	27
Forcing a traffic group to a standby state.....	28
Syncing the BIG-IP configuration to the device group.....	28
Implementation Results.....	29

Creating an Active-Standby Configuration using the Configuration Utility.....	31
Overview: Creating an active-standby DSC configuration.....	31
About DSC configuration on a VIPRION system.....	31
DSC prerequisite worksheet.....	33
Task summary.....	34
Specifying an IP address for config sync.....	34
Specifying an IP address for connection mirroring.....	35
Specifying the HA capacity of a device.....	36
Establishing device trust.....	36
Creating a Sync-Failover device group.....	37
Syncing the BIG-IP configuration to the device group.....	38
Specifying IP addresses for failover communication.....	39
Syncing the BIG-IP configuration to the device group.....	40
Implementation result.....	40
 Creating an Active-Active Configuration using the Configuration Utility.....	 41
Overview: Creating an active-active DSC configuration.....	41
About DSC configuration on a VIPRION system.....	41
DSC prerequisite worksheet.....	43
Configurations using Sync-Failover device groups.....	44
Task summary.....	44
Specifying an IP address for config sync.....	45
Specifying an IP address for connection mirroring.....	45
Specifying the HA capacity of a device.....	46
Establishing device trust.....	47
Creating a Sync-Failover device group.....	48
Syncing the BIG-IP configuration to the device group.....	48
Specifying IP addresses for failover communication.....	49
Creating a second traffic group for the device group.....	50
Assigning traffic-group-2 to a floating virtual IP address.....	50
Assigning traffic-group-2 to a floating self IP address.....	51
Syncing the BIG-IP configuration to the device group.....	51
Forcing a traffic group to a standby state.....	52
Implementation result.....	52
 Configuring Load-aware Failover.....	 53
Overview: Implementing load-aware failover.....	53
About device utilization calculation.....	54
Task summary.....	54
Specifying the HA capacity of a device.....	54
Specifying an HA load factor for a traffic group.....	55
Implementation Results.....	56

Managing Traffic with Bandwidth Controllers.....	59
Overview: Bandwidth control management.....	59
Bandwidth controllers vs. rate shaping.....	59
About static bandwidth control policies.....	59
Task summary for creating a static bandwidth control policy.....	60
About dynamic bandwidth control policies.....	61
About the Maximum User Rate PPS setting.....	61
Task summary for creating a dynamic bandwidth control policy.....	61
Configuring Network Virtualization Tunnels.....	65
Overview: Configuring network virtualization tunnels.....	65
About network virtualization tunnels on the BIG-IP system.....	66
Virtualized network terminology.....	66
Centralized vs. decentralized models of network virtualization.....	67
About network virtualization tunnel types.....	68
About statically configured network virtualization tunnels.....	69
Considerations for statically configured network virtualization tunnels.....	69
Examples for manually populating L2 location records.....	69
Sample NVGRE configuration using tmsh.....	71
Sample VXLAN unicast configuration using tmsh.....	72
Sample command for virtual server to listen on a VXLAN tunnel.....	73
Commands for viewing tunnel statistics.....	73
About VXLAN multicast configuration.....	73
About bridging VLAN and VXLAN networks.....	73
Considerations for configuring multicast VXLAN tunnels.....	74
Task summary.....	74
About configuring VXLAN tunnels on high availability BIG-IP device pairs.....	77
About configuring VXLAN tunnels using OVSDb.....	77
Setting up the OVSDb management component.....	77
About configuring VXLAN-GPE tunnels.....	78
Creating a multicast VXLAN-GPE tunnel.....	78
Web Hosting Multiple Customers Using an External Switch.....	81
Overview: Web hosting multiple customers using an external switch.....	81
Illustration for hosting multiple customers using an external switch.....	81
Task summary for hosting multiple customers.....	81
Creating a VLAN with a tagged interface.....	82
Creating a load balancing pool.....	82
Creating a virtual server for HTTP traffic.....	83
Web Hosting Multiple Customers Using Untagged Interfaces.....	85
Overview: Web hosting multiple customers using untagged interfaces.....	85

Illustration for hosting multiple customers using untagged interfaces.....	85
Task summary for hosting multiple customers.....	85
Creating a VLAN with an untagged interface.....	86
Creating a load balancing pool.....	86
Creating a virtual server for HTTP traffic.....	87
Web Hosting Multiple Customers Using Route Domains.....	89
Overview: Use of route domains to host multiple web customers on the BIG-IP system.....	89
Illustration of sample BIG-IP configuration using route domains.....	90
Illustration of resulting route domain configuration.....	91
Task summary.....	91
Creating an administrative partition.....	91
Creating a VLAN with a tagged interface.....	92
Creating a self IP address for a default route domain in an administrative partition.....	93
Creating a route domain on the BIG-IP system.....	93
Creating a load balancing pool.....	94
Creating a virtual server.....	95
Configuring route advertisement for a virtual address.....	96
Adding routes that specify VLAN internal as the resource.....	96
Implementing the Link Layer Discovery Protocol.....	99
Overview: Implementing Link Layer Discovery Protocol.....	99
Task summary.....	100
Configuring global LLDP properties.....	100
Configuring LLDP settings for an individual interface.....	100
Implementation result.....	101
Configuring Remote High-Speed Logging.....	103
Overview: Configuring high-speed remote logging.....	103
About the configuration objects of high-speed remote logging.....	104
Creating a pool of remote logging servers.....	104
Creating a remote high-speed log destination.....	105
Creating a formatted remote high-speed log destination.....	105
Creating a publisher	106
Creating a logging filter.....	106
Disabling system logging	107
Troubleshooting logs that contain unexpected messages	107
Deploying Route Domains within a vCMP Guest.....	109
Overview: Deploying Route Domains within a vCMP Guest.....	109
Prerequisite configuration tasks.....	110

About VLAN and BIG-IP address configuration.....	110
Illustration of VLAN and BIG-IP address configuration.....	110
Tasks for the host administrator.....	111
Creating customer VLANs on the vCMP host.....	112
Assigning VLANs to the vCMP guest.....	113
Tasks for the guest administrator.....	113
Creating an administrative partition for each customer.....	113
About moving host-based VLANs to a customer partition.....	114
Creating a route domain for each administrative partition.....	116
Creating an empty traffic group for each customer.....	117
Assigning a traffic group to each administrative partition.....	118
Tasks for each customer administrator.....	118
Creating floating self IP addresses.....	118
Creating a pool.....	119
Creating a virtual server.....	120
Modifying a virtual IP address.....	121
Implementation results.....	121
Using Link Aggregation with Tagged VLANs for a One-network Topology.....	123
Overview: Configuring link aggregation using tagged VLANs on one network.....	123
Illustration of link aggregation for a one-network topology.....	124
Task summary.....	124
Creating a trunk.....	124
Adding a tagged interface to a VLAN.....	125
Creating a load balancing pool.....	125
Creating a virtual server with source address affinity persistence.....	126
Removing the self IP addresses from the default VLANs.....	126
Creating a VLAN group.....	127
Creating a self IP for a VLAN group.....	127
Using Link Aggregation with Tagged VLANs for a Two-network Topology.....	129
Overview: Configuring link aggregation of two interfaces using tagged VLANs on two networks.....	129
Illustration of link aggregation for a two-network topology.....	130
Task summary.....	130
Creating a trunk.....	130
Adding a tagged interface to a VLAN.....	131
Creating a load balancing pool.....	131
Creating a virtual server with source address affinity persistence.....	132
Configuring Packet Filtering.....	133
Overview: Setting up packet filtering.....	133
Task summary.....	133
Enabling SNAT automap for internal and external VLANs.....	133

Creating a default gateway pool.....	134
Creating a forwarding virtual server.....	134
Enabling packet filtering.....	135
Creating a packet filter rule.....	136
Referencing an External File from within an iRule.....	137
Overview: Referencing an external file from an iRule.....	137
iRule commands for iFiles.....	137
Task summary.....	138
Importing a file for an iRule.....	138
Creating an iFile.....	138
Writing an iRule that references an iFile.....	139
Implementation result.....	139
Configuring Remote User Authentication and Authorization.....	141
Overview: Remote authentication and authorization of BIG-IP user accounts.....	141
Task summary.....	141
Specifying LDAP or Active Directory server information.....	142
Specifying client certificate LDAP server information.....	143
Specifying RADIUS server information.....	145
Specifying TACACS+ server information.....	146
Configuring access control for remote user groups.....	147
Saving access control settings to a file.....	149
Importing BIG-IP configuration data onto other BIG-IP systems.....	149
Configuring Administrative Partitions to Control User Access.....	151
Overview: Administrative partitions for user access control.....	151
Task summary.....	151
Creating an administrative partition.....	151
Assigning roles to a user account.....	152
Working with Single Configuration Files.....	155
Overview: Working with single configuration files.....	155
tmsh commands for single configuration files (SCFs).....	155
Task summary.....	156
Creating and saving an SCF.....	156
Loading an SCF onto a target BIG-IP system.....	156
Using an SCF to restore a BIG-IP system configuration.....	157
Configuring a One-Arm Deployment Using WCCPv2.....	159
Overview: Configuring a one-arm deployment using WCCPv2.....	159
About WCCPv2 redirection on the BIG-IP system.....	159
Before you begin configuring an iSession connection.....	160

Task summary.....	161
Creating a VLAN for a one-arm deployment.....	161
Creating a self IP address for a one-arm deployment.....	162
Defining a route.....	163
Configuring WCCPv2.....	163
Verifying connectivity.....	166
Verifying WCCPv2 configuration for one-arm deployment.....	167
Creating an iSession connection.....	167
Validating iSession configuration in a one-arm deployment.....	169
Configuring the Cisco router for a one-arm deployment using WCCPv2.....	169
Viewing pertinent configuration details from the command line.....	171
Implementation result.....	176
 Legal Notices.....	 177
Legal notices.....	177

Customizing the BIG-IP Dashboard

Overview: BIG-IP dashboard customization

The BIG-IP® dashboard displays system statistics in selectable graphs, gauges, and tables. In addition to the pre-defined views, you can create custom combinations of the dashboard windows, called *views*, and save them in groups, called *view sets*. You can combine windows from different BIG-IP modules in a single view, or use just the windows you want for a single module. Windows are available only for those modules that you have licensed and provisioned.

Note: *The view set name for all pre-defined views is standard.*

Customizing the BIG-IP dashboard

You can create custom dashboard displays using the windows for any modules that are available on the BIG-IP® system.

1. On the Main tab, click **Statistics > Dashboard**.
A separate window opens for the BIG-IP dashboard.
2. On the Views control bar, click the Create custom view icon.
A blank canvas opens in design mode. The Dashboard Windows Chooser displays the available windows, grouped by module. You can click a module to display the available windows.
3. From the Dashboard Windows Chooser, drag and drop the windows you want onto the canvas.
After you drag a window to the canvas, you can resize it or change it to display the information you want by clicking a tab or filter.

Note: *The windows are not active when in design mode, so the data does not update in real time.*

4. When you have placed the windows you want onto the canvas, click the Save icon on the Custom Views control bar.
The Save View popup window opens.
5. Type a name for the view.
6. Type a new name for the view set, or select from the list.
7. Click **OK**.
The new view is saved, and appears in the **Views** list.
8. Click the double-gear icon on the Custom Views control bar to return to active mode.
The dashboard displays the custom view you just created, and updates the display with real-time data.

Creating an Active-Standby Configuration Using the Setup Utility

Overview: Creating a basic active-standby configuration

This implementation describes how to use the Setup utility to configure two new BIG-IP® devices that function as an active-standby pair. An *active-standby pair* is a pair of BIG-IP devices configured so that one device is actively processing traffic while the other device remains ready to take over if failover occurs. The two devices synchronize their configuration data and can fail over to one another in the event that one of the devices becomes unavailable.

Important: *The same version of BIG-IP system software must be running on all devices in the device group.*

First, you run the Setup utility on each device to configure base network components (that is, a management port, administrative passwords, and the default VLANs and their associated self IP addresses). Continue running it on each device to establish a trust relationship between the two devices, and create a Sync-Failover type of device group that contains two member devices.

After the Setup utility is run on both devices, each device contains the default traffic group that the BIG-IP system automatically created during setup. A *traffic group* represents a set of configuration objects (such as floating self IP addresses and virtual IP addresses) that process application traffic. This traffic group actively processes traffic on one of the two devices, making that device the active device. When failover occurs, the traffic group becomes active on (that is, floats to) the peer BIG-IP device.

By default, the traffic group contains the floating self IP addresses of the default VLANs. Whenever you create additional configuration objects such as self IP addresses, virtual IP addresses, and SNATs, the system automatically adds these objects to the default traffic group.

Task summary

The configuration process for a BIG-IP® system entails running the Setup utility on each of the two BIG-IP devices. When you run the Setup utility, you perform several tasks. Completing these tasks results in both BIG-IP devices being configured properly for an active-standby implementation.

Important: *After using the Setup utility to create an active-standby configuration, you can re-enter the utility at any time to adjust the configuration. Simply click the F5 logo in the upper-left corner of the BIG-IP Configuration utility, and on the Welcome screen, click **Run the Setup Utility**. Then page through the utility to find the appropriate screens.*

Licensing and provisioning the BIG-IP system

Configuring a device certificate

Configuring the management port and administrative user accounts

Enabling ConfigSync and high availability

Configuring the internal network

Configuring the external network

Configuring the network for high availability

Configuring a ConfigSync address

Configuring failover and mirroring addresses

Discovering a peer device

Licensing and provisioning the BIG-IP system

Using the Setup utility, you can activate the license and provision the BIG-IP® system.

1. From a workstation attached to the network on which you configured the management interface, type the following URL syntax where `<management_IP_address>` is the address you configured for device management:
`https://<management_IP_address>`
2. At the login prompt, type the default user name `admin`, and password `admin`, and click **Log in**.
The Setup utility screen opens.
3. Click **Next**.
4. Click **Activate**.
The License screen opens.
5. In the **Base Registration Key** field, paste the registration key.
6. Click **Next** and follow the process for licensing and provisioning the system.

***Note:** When you perform the licensing task so that you can run the F5 cloud ADC, you can accept the default provisioning values.*

7. Click **Next**.
This displays the screen for configuring general properties and user administration settings.

The BIG-IP system license is now activated, and the relevant BIG-IP modules are provisioned.

Configuring a device certificate

Import or verify the certificate for the BIG-IP device.

Do one of the following:

- Click **Import**, import a certificate, click **Import**, and then click **Next**.
- Verify the displayed information for the certificate and click **Next**.

Configuring the management port and administrative user accounts

Configure the management port, time zone, and the administrative user names and passwords.

1. On the screen for configuring general properties, for the **Management Port Configuration** setting, select **Manual** and specify the IP address, network mask, and default gateway.
2. In the **Host Name** field, type a fully-qualified domain name (FQDN) for the system.
The FQDN can consist of letters, numbers, and/or the characters underscore (`_`), dash (`-`), or period (`.`).
3. For the **Host IP Address** setting, retain the default value **Use Management Port IP Address**.

4. From the **Time Zone** list, select a time zone.
The time zone you select typically reflects the location of the F5® system.
5. For the **Root Account** setting, type and confirm a password for the `root` account.
The `root` account provides console access only.
6. For the **Admin Account** setting, type and confirm a password.
Typing a password for the `admin` account causes the system to terminate the login session. When this happens, log in to the F5 Configuration utility again, using the new password. The system returns to the appropriate screen in the Setup utility.
7. For the **SSH Access** setting, select or clear the check box.
8. Click **Next**.
9. In the Standard Network Configuration area of the screen, click **Next**.
This displays the screen for enabling configuration synchronization and high availability.

Enabling ConfigSync and high availability

When you perform this task, you set up config sync and connection mirroring, and you can specify the failover method (network, serial, or both).

1. For the **Config Sync** setting, select the **Display configuration synchronization options** check box.
This causes an additional ConfigSync screen to be displayed later.
2. For the **High Availability** setting, select the **Display failover and mirroring options** check box.
This displays the **Failover Method** list and causes additional failover screens to be displayed later.
3. From the **Failover Method** list, select **Network and serial cable**.
If you have a VIPRION® system, select **Network**.
4. Click **Next**.
This displays the screen for configuring the default VLAN **internal**.

Configuring the internal network

You can use the Setup utility to specify self IP addresses and settings for a VLAN on the BIG-IP® internal network. The default VLAN for the internal network is named `internal`.

1. Specify the **Self IP** setting for the internal network:
 - a) In the **Address** field, type a self IP address.
 - b) In the **Netmask** field, type a network mask for the self IP address.
 - c) For the **Port Lockdown** setting, retain the default value.
2. Specify the **Floating IP** setting:
 - a) In the **Address** field, type a floating IP address.
This address should be distinct from the address you type for the **Self IP** setting.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services and the device needs to failover to a device group peer, use the second, Secondary Private IP address for the floating IP address.*

- b) For the **Port Lockdown** setting, retain the default value.

3. For the **VLAN Tag ID** setting, retain the default value, **auto**.
This is the recommended value.
4. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) Click **Add**.
5. Click **Next**.
This completes the configuration of the internal self IP addresses and VLAN, and displays the screen for configuring the default VLAN **external**.

Configuring the external network

You can use the Setup utility to specify self IP addresses and settings for a VLAN on the BIG-IP® external network. The default VLAN for the external network is named `external`.

1. Specify the **Self IP** setting for the external network:
 - a) In the **Address** field, type a self IP address.
 - b) In the **Netmask** field, type a network mask for the self IP address.
 - c) For the **Port Lockdown** setting, retain the default value.
2. In the **Default Gateway** field, type the IP address that you want to use as the default gateway to VLAN **external**.
3. Specify the **Floating IP** setting:
 - a) In the **Address** field, type a floating IP address.
This address should be distinct from the address you type for the **Self IP** setting.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services and the device needs to failover to a device group peer, use the second, Secondary Private IP address for the floating IP address.*

- b) For the **Port Lockdown** setting, retain the default value.
4. For the **VLAN Tag ID** setting, retain the default value, **auto**.
This is the recommended value.
5. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) Click **Add**.
6. Click **Next**.
This completes the configuration of the external self IP addresses and VLAN, and displays the screen for configuring the default VLAN **HA**.

Configuring the network for high availability

To configure a network for high availability, specify self IP addresses and settings for VLAN **HA**, which is the VLAN that the system will use for failover and connection mirroring.

1. For the **High Availability VLAN** setting, retain the default value, **Create VLAN HA**.
2. Specify the **Self IP** setting for VLAN **HA**:
 - a) In the **Address** field, type a self IP address.
 - b) In the **Netmask** field, type a network mask for the self IP address.
3. For the **VLAN Tag ID** setting, retain the default value, **auto**.
This is the recommended value.
4. For the **Interfaces** setting,
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Untagged**.
 - c) Click **Add**.
5. Click **Next**.
This configures the self IP address and VLAN that the system will use for high availability and displays the default IP address that the system will use for configuration synchronization.

Configuring a ConfigSync address

Use this task to specify the address that you want the system to use for configuration synchronization.

1. From the **Local Address** list, select a self IP address.
Do not select a management IP address.
2. Click **Next**.
This displays the screen for configuring unicast and multicast failover addresses.

Configuring failover and mirroring addresses

Follow these steps to specify the local unicast and mirroring addresses that you want the BIG-IP® system to use for high availability. During the final step of running the Setup utility, the system exchanges these addresses with its trusted peer. If you are configuring a VIPRION® system, configure a multicast failover address as well.

1. Locate the Failover Unicast Configuration area of the screen.
2. Under Local Address, confirm that there are entries for the self IP addresses that are assigned to the **HA** and **internal** VLANs and for the local management IP address for this device. If these entries are absent, click the **Add** button to add the missing entries to the list of Failover Unicast Addresses.
 - a) For the **Address** setting, select the address for the VLAN you need to add (either **HA** or **internal**).
 - b) In the **Port** field, type a port number or retain the default port number, 1026.
 - c) Click **Repeat** to add additional self IP addresses, or click **Finished**.
 - d) Repeat these steps to add a management IP address.

3. Click **Next**.
4. From the **Primary Local Mirror Address** list, retain the default value, which is the self IP address for VLAN `HA`.
5. From the **Secondary Local Mirror Address** list, select the address for VLAN `internal`.
6. Click **Finished**.

Discovering a peer device

You can use the Setup utility to discover a peer device for the purpose of exchanging failover and mirroring information.

1. Under **Standard Pair Configuration**, click **Next**.
2. If this is the first device of the pair that you are setting up, then under **Configure Peer Device**, click **Finished**.
To activate device discovery, you must first run the Setup utility on the peer device.
3. If this is the second device of the pair that you are setting up:
 - a) Under **Discover Configured Peer Device**, click **Next**.
 - b) Under **Remote Device Credentials**, specify the `Management IP` address, `Administrator Username`, and `Administrator Password`.
 - c) Click **Retrieve Device Information**.
4. Click **Finished**.

After the second device has discovered the first device, the two devices have a trust relationship and constitute a two-member device group. Also, each device in the pair contains a default traffic group named `Traffic-Group-1`. By default, this traffic group contains the floating IP addresses that you defined for VLANs `internal` and `external`.

Implementation result

To summarize, you now have the following BIG-IP® configuration on each device of the pair:

- A management port, management route, and administrative passwords defined.
- A VLAN named `internal`, with one static and one floating IP address.
- A VLAN named `external`, with one static and one floating IP address.
- A VLAN named `HA` with a static IP address.
- Configuration synchronization, failover, and mirroring enabled.
- Failover methods of serial cable and network (or network-only, for a VIPRION® platform).
- A designation as an authority device, where trust was established with the peer device.
- A Sync-Failover type of device group with two members defined.
- A default traffic group that floats to the peer device to process application traffic when this device becomes unavailable. This traffic group contains two floating self IP addresses for VLANs `internal` and `external`.

On either device in the device group, you can create additional configuration objects, such as virtual IP addresses and SNATs. The system automatically adds these objects to `Traffic-Group-1`.

Creating an Active-Active Configuration Using the Setup Utility

Overview: Creating a basic active-active configuration

This implementation describes how to use the Setup utility to configure two new BIG-IP® devices that function as an active-active pair. An *active-active* pair is a pair of BIG-IP devices configured so that both devices are actively processing traffic and are ready to take over one another if failover occurs. The two devices synchronize their configuration data to one another.

Note: *Access Policy Manager (APM) is not supported in an Active-Active configuration. APM is supported in an Active-Standby configuration with two BIG-IP systems only.*

Important: *The same version of BIG-IP system software must be running on all devices in the device group.*

Using this implementation, you begin by running the Setup utility on each device to configure its base network components. Base network components include a management port, administrative passwords, and default VLANs and their associated self IP addresses. You also use Setup to configure configuration synchronization and high availability.

You then use the BIG-IP® Configuration utility to:

- Establish trust between the two devices
- Create a Sync-Failover type of device group that contains two member devices
- Create a second traffic group
- Create two iApp™ application services

In this configuration, both devices actively process application traffic, each for a different application. One device processes its application traffic using the configuration objects associated with the default floating traffic group, `traffic-group-1`. By default, this traffic group contains the floating self IP addresses of the default VLANs. The other device processes its application traffic using a second traffic group that you create.

If one of the devices becomes unavailable for any reason, the other device automatically begins processing traffic for the unavailable peer, while continuing to process the traffic for its own application.

This illustration shows an example of the device group that this implementation creates, named `Device Group A`. This device group contains two BIG-IP devices, `Device 1` and `Device 2`.

The configuration shows two traffic groups, `traffic-group-1` and `traffic-group-2`, each containing failover objects. For `traffic-group-1`, `Device 1` is the default device. For `traffic-group-2`, `Device 2` is the default device. If `Device 1` becomes unavailable, `traffic-group-1` floats to `Device 2`. If `Device 2` becomes unavailable, `traffic-group-2` floats to `Device 1`.

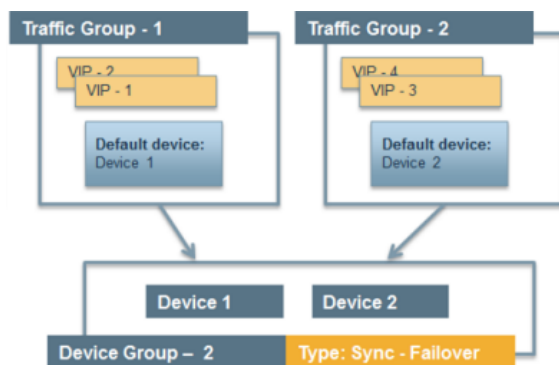


Figure 1: Device group with active-active configuration

By implementing this configuration, you ensure that:

- Each device has base network components configured.
- Any objects on a BIG-IP device that you configure for synchronization remain synchronized between the two devices.
- Failover capability and connection mirroring are enabled on each device.

Important: For active-active configurations, you must enable network failover instead of hard-wired serial failover.

Task summary

The BIG-IP® configuration process begins with running the Setup utility on each of the two BIG-IP devices. Once you have completed that task, you can log into either of the BIG-IP devices and perform all of the remaining tasks, on that device only. This results in both BIG-IP devices being configured properly for an active-active implementation.

Important: After using the Setup utility to create a redundant system configuration, you can re-enter the utility at any time to adjust the configuration. Simply click the F5 logo in the upper-left corner of the BIG-IP Configuration utility, and on the Welcome screen, click **Run the Setup Utility**. Then page through the utility to find the appropriate screens.

Licensing and provisioning the BIG-IP system

Configuring a device certificate

Configuring the management port and administrative user accounts

Enabling ConfigSync and high availability

Configuring the internal network

Configuring the external network

Configuring the network for high availability

Configuring a ConfigSync address

Configuring failover and mirroring addresses

Establishing device trust

Creating a Sync-Failover device group

Creating an iApp application for the local device

Creating a traffic group for a remote device

Creating an iApp application for a remote device

Forcing a traffic group to a standby state
Syncing the BIG-IP configuration to the device group

Licensing and provisioning the BIG-IP system

Using the Setup utility, you can activate the license and provision the BIG-IP® system.

1. From a workstation attached to the network on which you configured the management interface, type the following URL syntax where `<management_IP_address>` is the address you configured for device management:
`https://<management_IP_address>`
2. At the login prompt, type the default user name `admin`, and password `admin`, and click **Log in**.
The Setup utility screen opens.
3. Click **Next**.
4. Click **Activate**.
The License screen opens.
5. In the **Base Registration Key** field, paste the registration key.
6. Click **Next** and follow the process for licensing and provisioning the system.

***Note:** When you perform the licensing task so that you can run the F5 cloud ADC, you can accept the default provisioning values.*

7. Click **Next**.
This displays the screen for configuring general properties and user administration settings.

The BIG-IP system license is now activated, and the relevant BIG-IP modules are provisioned.

Configuring a device certificate

Import or verify the certificate for the BIG-IP device.

Do one of the following:

- Click **Import**, import a certificate, click **Import**, and then click **Next**.
- Verify the displayed information for the certificate and click **Next**.

Configuring the management port and administrative user accounts

Configure the management port, time zone, and the administrative user names and passwords.

1. On the screen for configuring general properties, for the **Management Port Configuration** setting, select **Manual** and specify the IP address, network mask, and default gateway.
2. In the **Host Name** field, type a fully-qualified domain name (FQDN) for the system.
The FQDN can consist of letters, numbers, and/or the characters underscore (`_`), dash (`-`), or period (`.`).
3. For the **Host IP Address** setting, retain the default value **Use Management Port IP Address**.
4. From the **Time Zone** list, select a time zone.

The time zone you select typically reflects the location of the F5® system.

5. For the **Root Account** setting, type and confirm a password for the `root` account.
The `root` account provides console access only.
6. For the **Admin Account** setting, type and confirm a password.
Typing a password for the `admin` account causes the system to terminate the login session. When this happens, log in to the F5 Configuration utility again, using the new password. The system returns to the appropriate screen in the Setup utility.
7. For the **SSH Access** setting, select or clear the check box.
8. Click **Next**.
9. In the Standard Network Configuration area of the screen, click **Next**.
This displays the screen for enabling configuration synchronization and high availability.

Enabling ConfigSync and high availability

When you perform this task, you set up config sync and connection mirroring, and you can specify the failover method (network, serial, or both).

1. For the **Config Sync** setting, select the **Display configuration synchronization options** check box.
This causes an additional ConfigSync screen to be displayed later.
2. For the **High Availability** setting, select the **Display failover and mirroring options** check box.
This displays the **Failover Method** list and causes additional failover screens to be displayed later.
3. From the **Failover Method** list, select **Network and serial cable**.
If you have a VIPRION® system, select **Network**.
4. Click **Next**.
This displays the screen for configuring the default VLAN **internal**.

Configuring the internal network

You can use the Setup utility to specify self IP addresses and settings for a VLAN on the BIG-IP® internal network. The default VLAN for the internal network is named `internal`.

1. Specify the **Self IP** setting for the internal network:
 - a) In the **Address** field, type a self IP address.
 - b) In the **Netmask** field, type a network mask for the self IP address.
 - c) For the **Port Lockdown** setting, retain the default value.
2. Specify the **Floating IP** setting:
 - a) In the **Address** field, type a floating IP address.
This address should be distinct from the address you type for the **Self IP** setting.

Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services and the device needs to failover to a device group peer, use the second, Secondary Private IP address for the floating IP address.

- b) For the **Port Lockdown** setting, retain the default value.

3. For the **VLAN Tag ID** setting, retain the default value, **auto**.
This is the recommended value.
4. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) Click **Add**.
5. Click **Next**.
This completes the configuration of the internal self IP addresses and VLAN, and displays the screen for configuring the default VLAN **external**.

Configuring the external network

You can use the Setup utility to specify self IP addresses and settings for a VLAN on the BIG-IP® external network. The default VLAN for the external network is named `external`.

1. Specify the **Self IP** setting for the external network:
 - a) In the **Address** field, type a self IP address.
 - b) In the **Netmask** field, type a network mask for the self IP address.
 - c) For the **Port Lockdown** setting, retain the default value.
2. In the **Default Gateway** field, type the IP address that you want to use as the default gateway to VLAN **external**.
3. Specify the **Floating IP** setting:
 - a) In the **Address** field, type a floating IP address.
This address should be distinct from the address you type for the **Self IP** setting.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services and the device needs to failover to a device group peer, use the second, Secondary Private IP address for the floating IP address.*

 - b) For the **Port Lockdown** setting, retain the default value.
4. For the **VLAN Tag ID** setting, retain the default value, **auto**.
This is the recommended value.
5. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) Click **Add**.
6. Click **Next**.
This completes the configuration of the external self IP addresses and VLAN, and displays the screen for configuring the default VLAN **HA**.

Configuring the network for high availability

To configure a network for high availability, specify self IP addresses and settings for VLAN **HA**, which is the VLAN that the system will use for failover and connection mirroring.

1. For the **High Availability VLAN** setting, retain the default value, **Create VLAN HA**.
2. Specify the **Self IP** setting for VLAN **HA**:
 - a) In the **Address** field, type a self IP address.
 - b) In the **Netmask** field, type a network mask for the self IP address.
3. For the **VLAN Tag ID** setting, retain the default value, **auto**.
This is the recommended value.
4. For the **Interfaces** setting,
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Untagged**.
 - c) Click **Add**.
5. Click **Next**.
This configures the self IP address and VLAN that the system will use for high availability and displays the default IP address that the system will use for configuration synchronization.

Configuring a ConfigSync address

Use this task to specify the address that you want the system to use for configuration synchronization.

1. From the **Local Address** list, select a self IP address.
Do not select a management IP address.
2. Click **Next**.
This displays the screen for configuring unicast and multicast failover addresses.

Configuring failover and mirroring addresses

Follow these task steps to specify the unicast IP addresses of the local device that you want the system to use for failover. Typically, you specify the self IP address for the local VLAN **HA**, as well as the IP address for the management port of the local device. If you are configuring a VIPRION[®] system, configure a multicast failover address as well.

Important: When configuring failover and mirroring IP addresses, you select addresses of the local device only. Later, during the process of device discovery, the two devices in the device group discover each other's addresses.

1. Locate the Failover Unicast Configuration area of the screen.
2. Under Local Address, confirm that there are entries for the self IP addresses that are assigned to the **HA** and **internal** VLANs and for the local management IP address for this device. If these entries are absent, click the **Add** button to add the missing entries to the list of Failover Unicast Addresses.
 - a) For the **Address** setting, select the address for the VLAN you need to add (either **HA** or **internal**).

- b) In the **Port** field, type a port number or retain the default port number, 1026.
- c) Click **Repeat** to add additional self IP addresses, or click **Finished**.
- d) Repeat these steps to add a management IP address.

3. Click **Next**.

4. From the **Primary Local Mirror Address** list, retain the default value, which is the self IP address for VLAN `HA`.

5. From the **Secondary Local Mirror Address** list, select the address for VLAN `internal`.

6. Click **Finished**.

This causes you to leave the Setup utility.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `Bigip_1`, `Bigip_2`, and `Bigip_3` each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `Bigip_1` and add devices `Bigip_2` and `Bigip_3` to the local trust domain; there is no need to repeat this process on devices `Bigip_2` and `Bigip_3`.

1. On the Main tab, click **Device Management > Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
 - If the BIG-IP device is an appliance, type the management IP address for the device.
 - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
 - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
 - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the management IP address and name of the remote device are correct.
7. Click **Finished**.

After you perform this task, the local device is now a member of the local trust domain. Also, the BIG-IP system automatically creates a special Sync-Only device group for the purpose of synchronizing trust information among the devices in the local trust domain, on an ongoing basis.

Repeat this task to specify each device that you want to add to the local trust domain.

Creating a Sync-Failover device group

This task establishes failover capability between two BIG-IP® devices. If the active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You can perform this task on any authority device within the local trust domain.

1. On the Main tab, click **Device Management > Device Groups**.
2. On the Device Groups list screen, click **Create**.
The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Includes** list.

The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.

5. For the **Network Failover** setting, select or clear the check box:
 - Select the check box if you want device group members to handle failover communications by way of network connectivity. This choice is required for active-active configurations.
 - Clear the check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

For active-active configurations, you must select network failover, as opposed to serial-cable (hard-wired) connectivity.

6. Click **Finished**.

You now have a Sync-Failover device group containing two BIG-IP devices as members.

Creating an iApp application for the local device

Use this procedure to create a set of related configuration objects on the system (that is, an application).

1. On the Main tab, click **iApp > Application Services**.
2. Click **Create**.
3. In the **Name** field, type the name for your application service.
4. From the **Template** list, select a template.
5. From the Template Selection list, select **Advanced**.
This causes additional settings to appear.
6. For the **Configure Sync and/or Failover for this application?** setting, select **Yes**.
7. For the **Traffic Group** setting, ensure that the **Inherit traffic group from current partition / path** field and **traffic-group-1** are selected.
8. Configure remaining settings as needed.
9. At the bottom of the screen click **Finished** to save your changes.

You now have an iApp application service, which is associated with the traffic group assigned to the **root** folder, `traffic-group-1`.

Creating a traffic group for a remote device

Prerequisite: If you intend to specify a MAC masquerade address when creating a traffic group, you must first create the address, using an industry-standard method for creating a locally-administered MAC address.

Perform this procedure to create a traffic group to run on the remote BIG-IP® device. You create this traffic group on the local device. Later, you move the traffic group to the remote device by forcing this traffic group on the local device to a standby state.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. On the lower half of the screen, verify that the list shows the default floating traffic group (traffic-group-1) for the local device.
3. On the Traffic Groups screen, click **Create**.
4. Type the name `traffic-group-2` for the new traffic group.
5. Type a description of the new traffic group.
6. Click **Next**.
7. In the **MAC Masquerade Address** field, type a MAC masquerade address.
When you specify a MAC masquerade address, you reduce the risk of dropped connections when failover occurs. This setting is optional.
8. Click **Next**.
9. Select or clear the check box for the **Auto Failback** option:
 - Select the check box to cause the traffic group, after failover, to fail over again to the first device in the traffic group's ordered list when that device (and only that device) is available.
 - Clear the check box to cause the traffic group, after failover, to remain active on its current device until failover occurs again.
10. Click **Next**.
11. Make sure that the displayed traffic group settings are correct.
12. Click **Finished**.

You now have a floating traffic group for which the default device is the peer device.

Creating an iApp application for a remote device

Use this procedure when you want to create an application to run on a remote device and associate it with the traffic group named `traffic-group-2` that you previously created.

1. On the Main tab, click **iApp > Application Services**.
2. Click **Create**.
3. From the **Template** list, select a template.
4. From the Template Selection list, select **Advanced**.
This causes additional settings to appear.
5. In the **Name** field, type the name for your application service.
6. For the **Configure Sync and/or Failover for this application?** setting, select **Yes**.
7. For the **Traffic Group** setting, clear the **Inherit traffic group from current partition / path** field and from the list, select **traffic-group-2**.
8. Configure remaining settings as needed.
9. At the bottom of the screen click **Finished** to save your changes.

You now have an iApp application associated with `traffic-group-2`.

Forcing a traffic group to a standby state

You perform this task when you want the selected traffic group on the local device to fail over to another device (that is, switch to a `Standby` state). Users typically perform this task when no automated method is configured for a traffic group, such as auto-failback or an HA group. By forcing the traffic group into a `Standby` state, the traffic group becomes active on another device in the device group. For device groups with more than two members, you can choose the specific device to which the traffic group fails over.

1. Log in to the device on which the traffic group is currently active.
2. On the Main tab, click **Device Management > Traffic Groups**.
3. In the Name column, locate the name of the traffic group that you want to run on the peer device.
4. Select the check box to the left of the traffic group name.
If the check box is unavailable, the traffic group is not active on the device to which you are currently logged in. Perform this task on the device on which the traffic group is active.
5. Click **Force to Standby**.
This displays target device options.
6. Choose one of these actions:
 - If the device group has two members only, click **Force to Standby**. This displays the list of traffic groups for the device group and causes the local device to appear in the Next Active Device column.
 - If the device group has more than two members, then from the **Target Device** list, select a value and click **Force to Standby**.

The selected traffic group is now in a standby state on the local device and active on another device in the device group.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: *You perform this task on either of the two devices, but not both.*

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, from the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, from the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

After performing this task, all BIG-IP configuration data that is eligible for synchronization to other devices is replicated on each device in the device group.

Implementation Results

To summarize, you now have the following BIG-IP® configuration on each device of the pair:

- A management port, management route, and administrative passwords defined
- A VLAN named `internal`, with one static and one floating IP address
- A VLAN named `external`, with one static and one floating IP address
- A VLAN named `HA` with a static IP address
- Configuration synchronization, failover, and mirroring enabled
- Failover methods of serial cable and network
- Local IP addresses defined for failover and connection mirroring
- A designation as an authority device, where trust is established with the peer device
- A Sync-Failover type of device group with two members
- The default traffic group named `traffic-group-1` with `Device 1` as the default device
- An iApp application associated with `traffic-group-1`
- A traffic group named `traffic-group-2` with `Device 2` as the default device
- An iApp application associated with `traffic-group-2`

Creating an Active-Standby Configuration using the Configuration Utility

Overview: Creating an active-standby DSC configuration

The most common TMOS[®] device service clustering (DSC[®]) implementation is an *active-standby* configuration, where a single traffic group is active on one of the devices in the device group and is in a standby state on a peer device. If failover occurs, the standby traffic group on the peer device becomes active and begins processing the application traffic.

To implement this DSC implementation, you can create a Sync-Failover device group. A Sync-Failover device group with two or more members and one traffic group provides configuration synchronization and device failover, and optionally, connection mirroring.

If the device with the active traffic group goes offline, the traffic group becomes active on a peer device, and application processing is handled by that device.

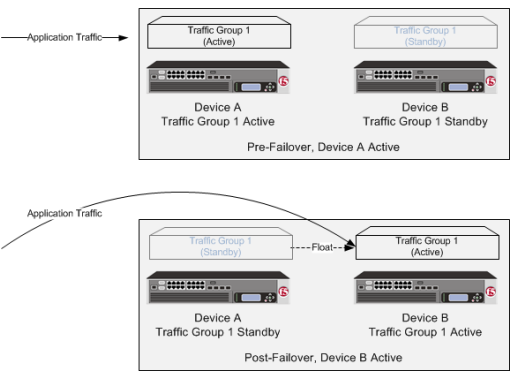


Figure 2: A two-member Sync-Failover device group for an active-standby configuration

About DSC configuration on a VIPRION system

The way you configure device service clustering (DSC[®]) (also known as redundancy) on a VIPRION[®] system varies depending on whether the system is provisioned to run the vCMP[®] feature.

For non-vCMP systems

For a device group that consists of VIPRION systems that are not licensed and provisioned for vCMP, each VIPRION cluster constitutes an individual device group member. The following table describes the IP addresses that you must specify when configuring redundancy.

Table 1: Required IP addresses for DSC configuration on a non-vCMP system

Feature	IP addresses required
Device trust	The primary floating management IP address for the VIPRION cluster.
ConfigSync	The unicast non-floating self IP address assigned to VLAN <code>internal</code> .

Feature	IP addresses required
Failover	<ul style="list-style-type: none"> Recommended: The unicast non-floating self IP address that you assigned to an internal VLAN (preferably VLAN <code>HA</code>), as well as a multicast address. Alternative: All unicast management IP addresses that correspond to the slots in the VIPRION cluster.
Connection mirroring	For the primary address, the non-floating self IP address that you assigned to VLAN <code>HA</code> . The secondary address is not required, but you can specify any non-floating self IP address for an internal VLAN..

For vCMP systems

On a vCMP system, the devices in a device group are virtual devices, known as *vCMP guests*. You configure device trust, config sync, failover, and mirroring to occur between equivalent vCMP guests in separate chassis.

For example, if you have a pair of VIPRION systems running vCMP, and each system has three vCMP guests, you can create a separate device group for each pair of equivalent guests. Table 4.2 shows an example.

Table 2: Sample device groups for two VIPRION systems with vCMP

Device groups for vCMP	Device group members
Device-Group-A	<ul style="list-style-type: none"> Guest1 on chassis1 Guest1 on chassis2
Device-Group-B	<ul style="list-style-type: none"> Guest2 on chassis1 Guest2 on chassis2
Device-Group-C	<ul style="list-style-type: none"> Guest3 on chassis1 Guest3 on chassis2

By isolating guests into separate device groups, you ensure that each guest synchronizes and fails over to its equivalent guest. The following table describes the IP addresses that you must specify when configuring redundancy:

Table 3: Required IP addresses for DSC configuration on a VIPRION system with vCMP

Feature	IP addresses required
Device trust	The cluster management IP address of the guest.
ConfigSync	The non-floating self IP address on the guest that is associated with VLAN <code>internal</code> on the host.
Failover	<ul style="list-style-type: none"> Recommended: The unicast non-floating self IP address on the guest that is associated with an internal VLAN on the host (preferably VLAN <code>HA</code>), as well as a multicast address. Alternative: The unicast management IP addresses for all slots configured for the guest.
Connection mirroring	For the primary address, the non-floating self IP address on the guest that is associated with VLAN <code>internal</code> on the host. The secondary address is not required, but you can specify any non-floating self IP address on the guest that is associated with an internal VLAN on the host.

DSC prerequisite worksheet

Before you set up device service clustering (DSC®), you must configure these BIG-IP® components on each device that you intend to include in the device group.

Table 4: DSC deployment worksheet

Configuration component	Considerations
Hardware, licensing, and provisioning	Devices in a device group must match with respect to product licensing and module provisioning. Heterogeneous hardware platforms within a device group are supported.
BIG-IP software version	Each device must be running BIG-IP version 11.x. This ensures successful configuration synchronization.
Management IP addresses	Each device must have a management IP address, a network mask, and a management route defined.
FQDN	Each device must have a fully-qualified domain name (FQDN) as its host name.
User name and password	Each device must have a user name and password defined on it that you will use when logging in to the BIG-IP Configuration utility.
root folder properties	The platform properties for the root folder must be set correctly (Sync-Failover and traffic-group-1).
VLANs	<p>You must create these VLANs on each device, if you have not already done so:</p> <ul style="list-style-type: none"> • A VLAN for the internal network, named <code>internal</code> • A VLAN for the external network, named <code>external</code> • A VLAN for failover communications, named <code>HA</code>
Self IP addresses	<p>You must create these self IP addresses on each device, if you have not already done so:</p> <ul style="list-style-type: none"> • Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>internal</code>. • Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>external</code>. • A non-floating self IP address on the internal subnet for VLAN <code>HA</code>. <hr/> <p>Note: When you create floating self IP addresses, the BIG-IP system automatically adds them to the default floating traffic group, <code>traffic-group-1</code>. To add a self IP address to a different traffic group, you must modify the value of the self IP address Traffic Group property.</p> <hr/> <p>Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the IP address you specify must be the floating IP address for high availability fast failover that you configured for the EC2 instance.</p> <hr/>
Port lockdown	For self IP addresses that you create on each device, you should verify that the Port Lockdown setting is set to Allow All , All Default , or Allow Custom . Do not specify None .

Configuration component	Considerations
Application-related objects	You must create any virtual IP addresses and optionally, SNAT translation addresses, as part of the local traffic configuration. You must also configure any iApps [®] application services if they are required for your application. When you create these addresses or services, the objects automatically become members of the default traffic group, <code>traffic-group-1</code> .
Time synchronization	The times set by the NTP service on all devices must be synchronized. This is a requirement for configuration synchronization to operate successfully.
Device certificates	Verify that each device includes an x509 device certificate. Devices with device certificates can authenticate and therefore trust one another, which is a prerequisite for device-to-device communication and data exchange.

Task summary

Use the tasks in this implementation to create a two-member device group, with one active traffic group, that syncs the BIG-IP[®] configuration to the peer device and provides failover capability if the peer device goes offline. Note that on a vCMP[®] system, the devices in a specific device group are vCMP guests, one per chassis.

Important: When you use this implementation, F5 Networks recommends that you synchronize the BIG-IP configuration twice, once after you create the device group, and again after you specify the IP addresses for failover.

Task list

[Specifying an IP address for config sync](#)
[Specifying an IP address for connection mirroring](#)
[Specifying the HA capacity of a device](#)
[Establishing device trust](#)
[Creating a Sync-Failover device group](#)
[Syncing the BIG-IP configuration to the device group](#)
[Specifying IP addresses for failover communication](#)
[Syncing the BIG-IP configuration to the device group](#)

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP[®] system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.

3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
F5 Networks recommends that you use the default value, which is the self IP address for the internal VLAN. This address must be a non-floating (static) self IP address and not a management IP address.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the internal self IP address that you select must be an internal private IP address that you configured for this EC2 instance as the **Local Address**.*

6. Click **Update**.

After performing this task, the other devices in the device group can synchronize their configurations to the local device whenever a sync operation is initiated.

Specifying an IP address for connection mirroring

You can specify the local self IP address that you want other devices in a device group to use when mirroring their connections to this device. Connection mirroring ensures that in-process connections for an active traffic group are not dropped when failover occurs. You typically perform this task when you initially set up device service clustering (DSC®).

Note: *You must perform this task locally on each device in the device group.*

1. Confirm that you are logged in to the device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Mirroring.
5. For the **Primary Local Mirror Address** setting, retain the displayed IP address or select another address from the list.
The recommended IP address is the self IP address for either VLAN `HA` or VLAN `internal`.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the self IP address you specify must be one of the private IP addresses that you configured for this EC2 instance as the **Primary Local Mirror Address**.*

6. For the **Secondary Local Mirror Address** setting, retain the default value of **None**, or select an address from the list.
This setting is optional. The system uses the selected IP address in the event that the primary mirroring address becomes unavailable.
7. Click **Update**.

In addition to specifying an IP address for mirroring, you must also enable connection mirroring on the relevant virtual servers on this device.

Specifying the HA capacity of a device

Before you perform this task, verify that this device is a member of a device group and that the device group contains three or more devices.

You perform this task when you have more than one type of hardware platform in a device group and you want to configure load-aware failover. *Load-aware failover* ensures that the BIG-IP® system can intelligently select the next-active device for each active traffic group in the device group when failover occurs. As part of configuring load-aware failover, you define an HA capacity to establish the amount of computing resource that the device provides relative to other devices in the device group.

Note: *If all devices in the device group are the same hardware platform, you can skip this task.*

1. On the Main tab, click **Device Management > Devices**.

This displays a list of device objects discovered by the local device.

2. In the Name column, click the name of the device for which you want to view properties.

This displays a table of properties for the device.

3. In the **HA Capacity** field, type a relative numeric value.

You need to configure this setting only when you have varying types of hardware platforms in a device group and you want to configure load-aware failover. The value you specify represents the relative capacity of the device to process application traffic compared to the other devices in the device group.

Important: *If you configure this setting, you must configure the setting on every device in the device group.*

If this device has half the capacity of a second device and a third of the capacity of a third device in the device group, you can specify a value of 100 for this device, 200 for the second device, and 300 for the third device.

When choosing the next active device for a traffic group, the system considers the capacity that you specified for this device.

4. Click **Update**.

After you perform this task, the BIG-IP system uses the **HA Capacity** value to calculate the current utilization of the local device, to determine the next-active device for failover of other traffic groups in the device group.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `Bigip_1`, `Bigip_2`, and `Bigip_3` each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `Bigip_1` and add devices `Bigip_2` and `Bigip_3` to the local trust domain; there is no need to repeat this process on devices `Bigip_2` and `Bigip_3`.

1. On the Main tab, click **Device Management > Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
 - If the BIG-IP device is an appliance, type the management IP address for the device.
 - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
 - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
 - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the management IP address and name of the remote device are correct.
7. Click **Finished**.

After you perform this task, the local device is now a member of the local trust domain. Also, the BIG-IP system automatically creates a special Sync-Only device group for the purpose of synchronizing trust information among the devices in the local trust domain, on an ongoing basis.

Repeat this task to specify each device that you want to add to the local trust domain.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP® devices. If an active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management > Device Groups**.
2. On the Device Groups list screen, click **Create**.
The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. From the **Configuration** list, select **Advanced**.
5. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Includes** list.
The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.
6. For the **Network Failover** setting, select or clear the check box:
 - Select the check box if you want device group members to handle failover communications by way of network connectivity. This choice is required for active-active configurations.
 - Clear the check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

For active-active configurations, you must select network failover, as opposed to serial-cable (hard-wired) connectivity.

7. For the **Automatic Sync** setting, specify whether configuration synchronization occurs manually or automatically:
 - Select the check box when you want the BIG-IP system to automatically sync the BIG-IP configuration data whenever a config sync operation is required. In this case, the BIG-IP system syncs the configuration data whenever the data changes on any device in the device group.
 - Clear the check box when you want to manually initiate each config sync operation. In this case, F5 networks recommends that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.
8. For the **Full Sync** setting, specify whether the system synchronizes the entire configuration during synchronization operations:
 - Select the check box when you want all sync operations to be full syncs. In this case, every time a config sync operation occurs, the BIG-IP system synchronizes all configuration data associated with the device group. This setting has a performance impact and is not recommended for most customers.
 - Clear the check box when you want all sync operations to be incremental (the default setting). In this case, the BIG-IP system syncs only the changes that are more recent than those on the target device. When you select this option, the BIG-IP system compares the configuration data on each target device with the configuration data on the source device and then syncs the delta of each target-source pair.

If you enable incremental synchronization, the BIG-IP system might occasionally perform a full sync for internal reasons. This is a rare occurrence and no user intervention is required.

9. In the **Maximum Incremental Sync Size (KB)** field, retain the default value of 1024, or type a different value.

This value specifies the total size of configuration changes that can reside in the incremental sync cache. If the total size of the configuration changes in the cache exceeds the specified value, the BIG-IP system performs a full sync whenever the next config sync operation occurs.

10. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: You perform this task on either of the two devices, but not both.

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, from the Name column, select the name of the relevant device group.

The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, from the Sync Status column, select the device that shows a sync status of `Changes Pending`.

4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

After performing this task, all BIG-IP configuration data that is eligible for synchronization to other devices is replicated on each device in the device group.

Specifying IP addresses for failover communication

You perform this task to specify the local IP addresses that you want other devices in the device group to use for continuous health-assessment communication with the local device. You must perform this task locally on each device in the device group.

Note: The IP addresses that you specify must belong to route domain 0.

1. Confirm that you are logged in to the device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Failover Network.
5. For the Failover Unicast Configuration settings, click **Add** for each IP address on this device that other devices in the device group can use to exchange failover messages with this device. The unicast IP addresses you specify depend on the type of device:

Platform	Action
Appliance without vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN HA) and the static management IP address currently assigned to the device.
Appliance with vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN HA) and the unique management IP address currently assigned to the guest.
VIPRION without vCMP®	Type a static self IP address associated with an internal VLAN (preferably VLAN HA). If you choose to specify unicast addresses only (and not a multicast address), you must also type the existing, static management IP addresses that you previously configured for all slots in the cluster. If you choose to specify one or more unicast addresses and a multicast address, then you do not need to specify the existing, per-slot static management IP addresses when configuring addresses for failover communication.
VIPRION with vCMP	Type a self IP address that is defined on the guest and associated with an internal VLAN on the host (preferably VLAN HA). If you choose to specify unicast failover addresses only (and not a multicast address), you must also type the existing, virtual static management IP addresses that you previously configured for all slots in the guest's virtual cluster. If you choose to specify one or more unicast addresses and a multicast address, you do not need to specify the existing, per-slot static and virtual management IP addresses when configuring addresses for failover communication.

Important: Failover addresses should always be static, not floating, IP addresses.

6. To enable the use of a failover multicast address on a VIPRION® platform (recommended), then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.

7. If you enabled **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.
If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

After you perform this task, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: *You perform this task on either of the two devices, but not both.*

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, from the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, from the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

After performing this task, all BIG-IP configuration data that is eligible for synchronization to other devices is replicated on each device in the device group.

Implementation result

You now have a Sync-Failover device group set up with an active-standby DSC™ configuration. This configuration uses the default floating traffic group (named `traffic-group-1`), which contains the application-specific floating self IP and virtual IP addresses, and is initially configured to be active on one of the two devices. If the device with the active traffic group goes offline, the traffic group becomes active on the other device in the group, and application processing continues.

Creating an Active-Active Configuration using the Configuration Utility

Overview: Creating an active-active DSC configuration

A common TMOS® device service clustering (DSC®) implementation is an active-standby configuration, where a single traffic group is active on one of the devices in the device group, and is in a standby state on a peer device. Alternatively however, you can create a second traffic group and activate that traffic group on a peer device. In this *active-active* configuration, the devices each process traffic for a different application simultaneously. If one of the devices in the device group goes offline, the traffic group that was active on that device fails over to a peer device. The result is that two traffic groups can become active on one device.

To implement this DSC implementation, you create a Sync-Failover device group. A Sync-Failover device group with two or more members provides configuration synchronization and device failover, and optionally, connection mirroring.




Figure 3: A two-member Sync-Failover group for an active-active configuration

About DSC configuration on a VIPRION system

The way you configure device service clustering (DSC®) (also known as redundancy) on a VIPRION® system varies depending on whether the system is provisioned to run the vCMP® feature.

For non-vCMP systems

For a device group that consists of VIPRION systems that are not licensed and provisioned for vCMP, each VIPRION cluster constitutes an individual device group member. The following table describes the IP addresses that you must specify when configuring redundancy.

Table 5: Required IP addresses for DSC configuration on a non-vCMP system

Feature	IP addresses required
Device trust	The primary floating management IP address for the VIPRION cluster.
ConfigSync	The unicast non-floating self IP address assigned to VLAN <i>internal</i> .

Feature	IP addresses required
Failover	<ul style="list-style-type: none"> Recommended: The unicast non-floating self IP address that you assigned to an internal VLAN (preferably VLAN <code>HA</code>), as well as a multicast address. Alternative: All unicast management IP addresses that correspond to the slots in the VIPRION cluster.
Connection mirroring	For the primary address, the non-floating self IP address that you assigned to VLAN <code>HA</code> . The secondary address is not required, but you can specify any non-floating self IP address for an internal VLAN..

For vCMP systems

On a vCMP system, the devices in a device group are virtual devices, known as *vCMP guests*. You configure device trust, config sync, failover, and mirroring to occur between equivalent vCMP guests in separate chassis.

For example, if you have a pair of VIPRION systems running vCMP, and each system has three vCMP guests, you can create a separate device group for each pair of equivalent guests. Table 4.2 shows an example.

Table 6: Sample device groups for two VIPRION systems with vCMP

Device groups for vCMP	Device group members
Device-Group-A	<ul style="list-style-type: none"> Guest1 on chassis1 Guest1 on chassis2
Device-Group-B	<ul style="list-style-type: none"> Guest2 on chassis1 Guest2 on chassis2
Device-Group-C	<ul style="list-style-type: none"> Guest3 on chassis1 Guest3 on chassis2

By isolating guests into separate device groups, you ensure that each guest synchronizes and fails over to its equivalent guest. The following table describes the IP addresses that you must specify when configuring redundancy:

Table 7: Required IP addresses for DSC configuration on a VIPRION system with vCMP

Feature	IP addresses required
Device trust	The cluster management IP address of the guest.
ConfigSync	The non-floating self IP address on the guest that is associated with VLAN <code>internal</code> on the host.
Failover	<ul style="list-style-type: none"> Recommended: The unicast non-floating self IP address on the guest that is associated with an internal VLAN on the host (preferably VLAN <code>HA</code>), as well as a multicast address. Alternative: The unicast management IP addresses for all slots configured for the guest.
Connection mirroring	For the primary address, the non-floating self IP address on the guest that is associated with VLAN <code>internal</code> on the host. The secondary address is not required, but you can specify any non-floating self IP address on the guest that is associated with an internal VLAN on the host.

DSC prerequisite worksheet

Before you set up device service clustering (DSC®), you must configure these BIG-IP® components on each device that you intend to include in the device group.

Table 8: DSC deployment worksheet

Configuration component	Considerations
Hardware, licensing, and provisioning	Devices in a device group must match with respect to product licensing and module provisioning. Heterogeneous hardware platforms within a device group are supported.
BIG-IP software version	Each device must be running BIG-IP version 11.x. This ensures successful configuration synchronization.
Management IP addresses	Each device must have a management IP address, a network mask, and a management route defined.
FQDN	Each device must have a fully-qualified domain name (FQDN) as its host name.
User name and password	Each device must have a user name and password defined on it that you will use when logging in to the BIG-IP Configuration utility.
root folder properties	The platform properties for the root folder must be set correctly (Sync-Failover and traffic-group-1).
VLANs	<p>You must create these VLANs on each device, if you have not already done so:</p> <ul style="list-style-type: none"> • A VLAN for the internal network, named <code>internal</code> • A VLAN for the external network, named <code>external</code> • A VLAN for failover communications, named <code>HA</code>
Self IP addresses	<p>You must create these self IP addresses on each device, if you have not already done so:</p> <ul style="list-style-type: none"> • Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>internal</code>. • Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>external</code>. • A non-floating self IP address on the internal subnet for VLAN <code>HA</code>. <hr/> <p>Note: When you create floating self IP addresses, the BIG-IP system automatically adds them to the default floating traffic group, <code>traffic-group-1</code>. To add a self IP address to a different traffic group, you must modify the value of the self IP address Traffic Group property.</p> <hr/> <p>Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the IP address you specify must be the floating IP address for high availability fast failover that you configured for the EC2 instance.</p> <hr/>
Port lockdown	For self IP addresses that you create on each device, you should verify that the Port Lockdown setting is set to Allow All , All Default , or Allow Custom . Do not specify None .

Configuration component	Considerations
Application-related objects	You must create any virtual IP addresses and optionally, SNAT translation addresses, as part of the local traffic configuration. You must also configure any iApps [®] application services if they are required for your application. When you create these addresses or services, the objects automatically become members of the default traffic group, <code>traffic-group-1</code> .
Time synchronization	The times set by the NTP service on all devices must be synchronized. This is a requirement for configuration synchronization to operate successfully.
Device certificates	Verify that each device includes an x509 device certificate. Devices with device certificates can authenticate and therefore trust one another, which is a prerequisite for device-to-device communication and data exchange.

Configurations using Sync-Failover device groups

This illustration shows two separate Sync-Failover device groups. In the first device group, only **LTM1** processes application traffic, and the two BIG-IP devices are configured to provide active-standby high availability. This means that **LTM1** and **LTM2** synchronize their configurations, and the failover objects on **LTM1** float to **LTM2** if **LTM1** becomes unavailable.

In the second device group, both **LTM1** and **LTM2** process application traffic, and the BIG-IP devices are configured to provide active-active high availability. This means that **LTM1** and **LTM2** synchronize their configurations, the failover objects on **LTM1** float to **LTM2** if **LTM1** becomes unavailable, and the failover objects on **LTM2** float to **LTM1** if **LTM2** becomes unavailable.

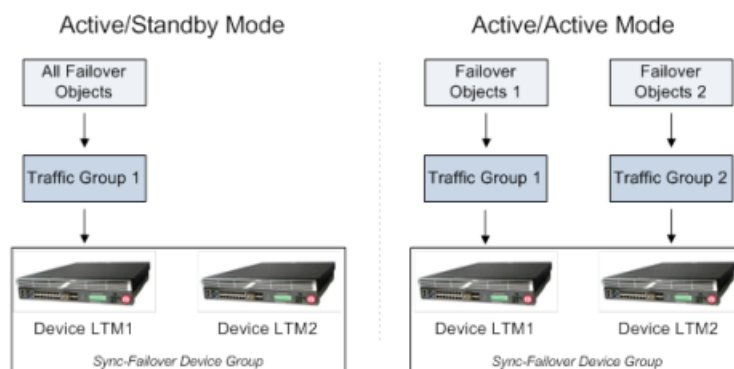


Figure 4: Comparison of Active-Standby and Active-Active device groups

Task summary

Use the tasks in this implementation to create a two-member device group, with two active traffic groups, that syncs the BIG-IP[®] configuration to the peer device and provides failover capability if the peer device goes offline. Note that on a vCMP[®] system, the devices in a specific device group are vCMP guests, one per chassis.

Important: When you use this implementation, F5 Networks recommends that you synchronize the BIG-IP configuration twice, once after you create the device group, and again after you specify the IP addresses for failover.

Task list

Specifying an IP address for config sync
Specifying an IP address for connection mirroring
Specifying the HA capacity of a device
Establishing device trust
Creating a Sync-Failover device group
Syncing the BIG-IP configuration to the device group
Specifying IP addresses for failover communication
Creating a second traffic group for the device group
Assigning traffic-group-2 to a floating virtual IP address
Assigning traffic-group-2 to a floating self IP address
Syncing the BIG-IP configuration to the device group
Forcing a traffic group to a standby state

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP® system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: *You must perform this task locally on each device in the device group.*

1. Confirm that you are logged in to the device you want to configure.
2. On the Main tab, click **Device Management** > **Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
F5 Networks recommends that you use the default value, which is the self IP address for the internal VLAN. This address must be a non-floating (static) self IP address and not a management IP address.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the internal self IP address that you select must be an internal private IP address that you configured for this EC2 instance as the **Local Address**.*

6. Click **Update**.

After performing this task, the other devices in the device group can synchronize their configurations to the local device whenever a sync operation is initiated.

Specifying an IP address for connection mirroring

You can specify the local self IP address that you want other devices in a device group to use when mirroring their connections to this device. Connection mirroring ensures that in-process connections for an active traffic group are not dropped when failover occurs. You typically perform this task when you initially set up device service clustering (DSC®).

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the device you want to configure.
 2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
 3. In the Name column, click the name of the device to which you are currently logged in.
 4. From the Device Connectivity menu, choose Mirroring.
 5. For the **Primary Local Mirror Address** setting, retain the displayed IP address or select another address from the list.
The recommended IP address is the self IP address for either VLAN `HA` or VLAN `internal`.
-

Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the self IP address you specify must be one of the private IP addresses that you configured for this EC2 instance as the **Primary Local Mirror Address**.

6. For the **Secondary Local Mirror Address** setting, retain the default value of **None**, or select an address from the list.
This setting is optional. The system uses the selected IP address in the event that the primary mirroring address becomes unavailable.
7. Click **Update**.

In addition to specifying an IP address for mirroring, you must also enable connection mirroring on the relevant virtual servers on this device.

Specifying the HA capacity of a device

Before you perform this task, verify that this device is a member of a device group and that the device group contains three or more devices.

You perform this task when you have more than one type of hardware platform in a device group and you want to configure load-aware failover. *Load-aware failover* ensures that the BIG-IP® system can intelligently select the next-active device for each active traffic group in the device group when failover occurs. As part of configuring load-aware failover, you define an HA capacity to establish the amount of computing resource that the device provides relative to other devices in the device group.

Note: If all devices in the device group are the same hardware platform, you can skip this task.

1. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
 2. In the Name column, click the name of the device for which you want to view properties.
This displays a table of properties for the device.
 3. In the **HA Capacity** field, type a relative numeric value.
You need to configure this setting only when you have varying types of hardware platforms in a device group and you want to configure load-aware failover. The value you specify represents the relative capacity of the device to process application traffic compared to the other devices in the device group.
-

Important: If you configure this setting, you must configure the setting on every device in the device group.

If this device has half the capacity of a second device and a third of the capacity of a third device in the device group, you can specify a value of 100 for this device, 200 for the second device, and 300 for the third device.

When choosing the next active device for a traffic group, the system considers the capacity that you specified for this device.

4. Click **Update.**

After you perform this task, the BIG-IP system uses the **HA Capacity** value to calculate the current utilization of the local device, to determine the next-active device for failover of other traffic groups in the device group.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `Bigip_1`, `Bigip_2`, and `Bigip_3` each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `Bigip_1` and add devices `Bigip_2` and `Bigip_3` to the local trust domain; there is no need to repeat this process on devices `Bigip_2` and `Bigip_3`.

1. On the Main tab, click **Device Management > Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
 - If the BIG-IP device is an appliance, type the management IP address for the device.
 - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
 - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
 - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the management IP address and name of the remote device are correct.
7. Click **Finished**.

After you perform this task, the local device is now a member of the local trust domain. Also, the BIG-IP system automatically creates a special Sync-Only device group for the purpose of synchronizing trust information among the devices in the local trust domain, on an ongoing basis.

Repeat this task to specify each device that you want to add to the local trust domain.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP devices that you intend to run in an active-active configuration. If an active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management > Device Groups**.
2. On the Device Groups list screen, click **Create**.
The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Includes** list.
The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.
5. For the **Network Failover** setting, verify that network failover is enabled.
Network failover must be enabled for active-active configurations (that is, device groups that will contain two or more active traffic groups).
6. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members. This device group is configured for environments that require the use of two or more active traffic groups to process application traffic.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP[®] configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: You perform this task on either of the two devices, but not both.

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, from the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, from the Sync Status column, select the device that shows a sync status of **Changes Pending**.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

After performing this task, all BIG-IP configuration data that is eligible for synchronization to other devices is replicated on each device in the device group.

Specifying IP addresses for failover communication

You perform this task to specify the local IP addresses that you want other devices in the device group to use for continuous health-assessment communication with the local device. You must perform this task locally on each device in the device group.

Note: The IP addresses that you specify must belong to route domain 0.

1. Confirm that you are logged in to the device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Failover Network.
5. For the Failover Unicast Configuration settings, click **Add** for each IP address on this device that other devices in the device group can use to exchange failover messages with this device. The unicast IP addresses you specify depend on the type of device:

Platform	Action
Appliance without vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN HA) and the static management IP address currently assigned to the device.
Appliance with vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN HA) and the unique management IP address currently assigned to the guest.
VIPRION without vCMP®	Type a static self IP address associated with an internal VLAN (preferably VLAN HA). If you choose to specify unicast addresses only (and not a multicast address), you must also type the existing, static management IP addresses that you previously configured for all slots in the cluster. If you choose to specify one or more unicast addresses and a multicast address, then you do not need to specify the existing, per-slot static management IP addresses when configuring addresses for failover communication.
VIPRION with vCMP	Type a self IP address that is defined on the guest and associated with an internal VLAN on the host (preferably VLAN HA). If you choose to specify unicast failover addresses only (and not a multicast address), you must also type the existing, virtual static management IP addresses that you previously configured for all slots in the guest's virtual cluster. If you choose to specify one or more unicast addresses and a multicast address, you do not need to specify the existing, per-slot static and virtual management IP addresses when configuring addresses for failover communication.

Important: Failover addresses should always be static, not floating, IP addresses.

6. To enable the use of a failover multicast address on a VIPRION® platform (recommended), then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
7. If you enabled **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.
If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

After you perform this task, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Creating a second traffic group for the device group

This task creates a second active floating traffic group to process application traffic. The default floating traffic group (traffic-group-1) processes application traffic for the local device.

Note: For this implementation, name this traffic group **traffic-group-2**.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. On the Traffic Groups screen, click **Create**.
3. Type the name `traffic-group-2` for the new traffic group.
4. In the **HA Load Factor** field, specify a value that represents the application load for this traffic group relative to other active traffic groups on the local device.

Important: If you configure this setting, you must configure the setting on every traffic group in the device group.

5. In the **MAC Masquerade Address** field, type a MAC masquerade address.
When you specify a MAC masquerade address, you reduce the risk of dropped connections when failover occurs. This setting is optional.
6. Select or clear the check box for the **Auto Failback** option:
 - Select the check box to cause the traffic group, after failover, to fail over again to the first device in the traffic group's ordered list when that device (and only that device) is available.
 - Clear the check box to cause the traffic group, after failover, to remain active on its current device until failover occurs again.
7. For the **Failover Order** setting, in the **Available** box, select a device name and using the Move button, move the device name to the **Enabled** box. Repeat for each device that you want to include in the ordered list.

This setting is optional. Only devices that are members of the relevant Sync-Failover device group are available for inclusion in the ordered list. If you have enabled the auto-failback feature on the traffic group, ensure that the first device in the ordered list is the device to which you want this traffic group to fail back to when that first device becomes available.

If auto-failback is enabled and the first device in the **Failover Order** list is unavailable, no auto-failback occurs and the traffic group continues to run on the current device. Also, if none of the devices in the **Failover Order** list is currently available when failover occurs, the BIG-IP system ignores the **Failover Order** setting and performs load-aware failover instead, using the **HA Load Factor** setting.
8. Click **Finished**.

You now have a second floating traffic group on the local device (in addition to the default floating traffic group) so that once the traffic group is activated on the remote devices, devices in the device group can process traffic for different applications.

Assigning traffic-group-2 to a floating virtual IP address

This task assigns a floating traffic group to a virtual IP address on a device.

1. On the Main tab, click **Local Traffic > Virtual Servers > Virtual Address List**.
The Virtual Address List screen opens.
2. In the Name column, click the virtual address that you want to assign to the traffic group.
This displays the properties of that virtual address.
3. From the **Traffic Group** list, select **traffic-group-2 (floating)**.
4. Click **Update**.

The device's floating virtual IP address is now a member of your second traffic group. The virtual IP address can now fail over to other devices in the device group.

Assigning traffic-group-2 to a floating self IP address

This task assigns your floating self IP address to traffic-group-2.

1. On the Main tab, click **Network > Self IPs**.
2. In the Name column, click the floating self IP address assigned to VLAN `internal`.
This displays the properties of that self IP address.
3. From the **Traffic Group** list, select **traffic-group-2 (floating)**.
4. Click **Update**.

The device's floating self IP address is now a member of your second traffic group. The self IP address can now fail over to other devices in the traffic group.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: You perform this task on either of the two devices, but not both.

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, from the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, from the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

After performing this task, all BIG-IP configuration data that is eligible for synchronization to other devices is replicated on each device in the device group.

Forcing a traffic group to a standby state

You perform this task when you want the selected traffic group on the local device to fail over to another device (that is, switch to a `Standby` state). Users typically perform this task when no automated method is configured for a traffic group, such as auto-failback or an HA group. By forcing the traffic group into a `Standby` state, the traffic group becomes active on another device in the device group. For device groups with more than two members, you can choose the specific device to which the traffic group fails over.

1. Log in to the device on which the traffic group is currently active.
2. On the Main tab, click **Device Management > Traffic Groups**.
3. In the Name column, locate the name of the traffic group that you want to run on the peer device.
4. Select the check box to the left of the traffic group name.
If the check box is unavailable, the traffic group is not active on the device to which you are currently logged in. Perform this task on the device on which the traffic group is active.
5. Click **Force to Standby**.
This displays target device options.
6. Choose one of these actions:
 - If the device group has two members only, click **Force to Standby**. This displays the list of traffic groups for the device group and causes the local device to appear in the Next Active Device column.
 - If the device group has more than two members, then from the **Target Device** list, select a value and click **Force to Standby**.

The selected traffic group is now in a standby state on the local device and active on another device in the device group.

Implementation result

You now have a Sync-Failover device group set up with an active-active DSC[®] configuration. In this configuration, each device has a different active traffic group running on it. That is, the active traffic group on one device is the default traffic group (named `traffic-group-1`), while the active traffic group on the peer device is a traffic group that you create. Each traffic group contains the floating self IP and virtual IP addresses specific to the relevant application.

If one device goes offline, the traffic group that was active on that device becomes active on the other device in the group, and processing for both applications continues on one device.

Configuring Load-aware Failover

Overview: Implementing load-aware failover

Load-aware failover is a BIG-IP® feature designed for use in a Sync-Failover device group. Configuring *load-aware failover* ensures that the traffic load on all devices in a device group is as equivalent as possible, factoring in any differences in device capacity and the amount of application traffic that traffic groups process on a device.

For example, suppose you have a heterogeneous three-member device group in which one device (Bigip_C) has twice the hardware capacity of the other two devices (Bigip_A and Bigip_B).

If the device group has four active traffic groups that each process the same amount of application traffic, then the load on all devices is equivalent when devices Bigip_A and Bigip_B each contain one active traffic group, while device Bigip_C contains two active traffic groups.

The figure shows a Sync-Failover device group where application traffic is directed to the device with the most capacity relative to the other device group members.




Figure 5: Sync-Failover device group with traffic groups of equal load:

The BIG-IP system implements load-aware failover by calculating a numeric, current utilization score for each device, based on numeric values that you specify for each device and traffic group relative to the other devices and traffic groups in the device group. The system then uses this current utilization score to determine which device is the best device in the group to become the next-active device when failover occurs for a traffic group.

The overall result is that the traffic load on each device is as equivalent as possible in a relative way, that is, factoring in individual device capacity and application traffic load per traffic group.

About device utilization calculation

The BIG-IP® system on each device performs a calculation to determine the device's current level of utilization. This utilization level indicates the ability for the device to be the next-active device in the event that an active traffic group on another device must fail over within a heterogeneous device group.

The calculation that the BIG-IP performs to determine the current utilization of a device is based on these factors:

Device capacity

A local device capacity relative to other device group members.

Active local traffic groups

The number of active traffic groups on the local device.

Active remote traffic groups

The number of remote active traffic groups for which the local device is the next-active device.

A multiplying load factor for each active traffic group

A multiplier value for each traffic group. The system uses this value to weight each active traffic group's traffic load compared to the traffic load of each of the other active traffic groups in the device group.

The BIG-IP system uses all of these factors to perform a calculation to determine, at any particular moment, a score for each device that represents the current utilization of that device. This utilization score indicates whether the BIG-IP system should, in its attempt to equalize traffic load on all devices, designate the device as a next-active device for an active traffic group on another device in the device group.

The calculation that the BIG-IP performs for each device is:

```
(The sum of local active traffic group loads + The sum of remote active
traffic group loads) / device capacity
```

Task summary

To implement load-aware failover you specify a value representing the relative traffic load for a traffic group and, optionally, a value representing the relative capacity of the BIG-IP® device.

Task list

Specifying the HA capacity of a device

Specifying an HA load factor for a traffic group

Specifying the HA capacity of a device

Before you perform this task, verify that this device is a member of a device group and that the device group contains three or more devices.

You perform this task when you have more than one type of hardware platform in a device group and you want to configure load-aware failover. *Load-aware failover* ensures that the BIG-IP® system can intelligently

select the next-active device for each active traffic group in the device group when failover occurs. As part of configuring load-aware failover, you define an HA capacity to establish the amount of computing resource that the device provides relative to other devices in the device group.

Note: *If all devices in the device group are the same hardware platform, you can skip this task.*

1. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
2. In the Name column, click the name of the device for which you want to view properties.
This displays a table of properties for the device.
3. In the **HA Capacity** field, type a relative numeric value.
You need to configure this setting only when you have varying types of hardware platforms in a device group and you want to configure load-aware failover. The value you specify represents the relative capacity of the device to process application traffic compared to the other devices in the device group.

Important: *If you configure this setting, you must configure the setting on every device in the device group.*

If this device has half the capacity of a second device and a third of the capacity of a third device in the device group, you can specify a value of 100 for this device, 200 for the second device, and 300 for the third device.

When choosing the next active device for a traffic group, the system considers the capacity that you specified for this device.

4. Click **Update**.

After you perform this task, the BIG-IP system uses the **HA Capacity** value to calculate the current utilization of the local device, to determine the next-active device for failover of other traffic groups in the device group.

Specifying an HA load factor for a traffic group

You perform this task when you want to specify the relative application load for an existing traffic group, for the purpose of configuring load-aware failover. *Load-aware failover* ensures that the BIG-IP® system can intelligently select the next-active device for each active traffic group in the device group when failover occurs. When you configure load-aware failover, you define an application traffic load (known as an *HA load factor*) for a traffic group to establish the amount of computing resource that an active traffic group uses relative to other active traffic groups.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, click the name of a traffic group.
This displays the properties of the traffic group.
3. From the **Failover Method** list, select **Load Aware**.
This displays the **HA Load Factor** setting.
4. In the **HA Load Factor** field, specify a value that represents the application load for this traffic group relative to other active traffic groups on the local device.

Important: *If you configure this setting, you must configure the setting on every traffic group in the device group.*

5. Click **Update**.

After performing this task, the BIG-IP system uses the **HA Load Factor** value as a factor in calculating the current utilization of the local device, to determine whether this device should be the next-active device for failover of other traffic groups in the device group.

Implementation Results

For this implementation example, the load-aware configuration now consists of both a user-specified relative high availability (HA) hardware capacity for each device and a relative load factor for each active traffic group.

Using the example in the overview, devices Bigip_A and Bigip_B are the same hardware platform and therefore have the same HA capacity, while Bigip_C has twice the HA capacity of the other two devices. Also, devices Bigip_A and Bigip_B currently have one active traffic group each, while Bigip_C has two active traffic groups. All three traffic groups process the same amount of application traffic.




Figure 6: Device utilization scores based on device capacity and traffic group load

The device utilization score that the BIG-IP® system calculates in this implementation is the sum of all traffic load values on a device divided by the device capacity.

Table 9: Calculating the utilization score for Bigip_A

HA capacity	Active traffic group	HA load factor	Potential active traffic group	HA load factor	Device utilization score
10	Traffic-group-1	1	Traffic-group-2	1	$2/10 = .2$

Table 10: Calculating the utilization score for Bigip_B

HA capacity	Active traffic group	HA load factor	Potential active traffic group	HA load factor	Device utilization score
10	Traffic-group-2	1	Traffic-group-3	1	$2/10 = .2$

Table 11: Calculating the utilization score for Bigip_C

HA capacity	Active traffic group	HA load factor	Potential active traffic group	HA load factor	Device utilization score
20	Traffic-group-3 and Traffic-group-4	1 and 1	Traffic-group-1	1	$3/20=.15$

This example shows the results of the calculations that the BIG-IP system performs for each device in the device group. The example shows that although device `Bigip_C` currently has the two active traffic groups, the device has the most available resource due to having the lowest utilization score of .15. In this case, `Bigip_C` is most likely the next-active device for the other two devices in the device group.

Managing Traffic with Bandwidth Controllers

Overview: Bandwidth control management

Fine-grained bandwidth control is essential to service providers, large enterprises, and remote access services (RAS) solutions. Bandwidth controllers on the BIG-IP® system can scale easily, work well in a distributed environment, and are easy to configure for various networks. Depending on the type of policy you configure, you can use bandwidth controllers to apply specified rate enforcement to traffic flows or mark traffic that exceeds limits.

Bandwidth control policies can be static or dynamic. Through the user interface (browser or `tmsm` command-line utility), when you apply a bandwidth control policy to a virtual server, packet filter, or route domain, you can apply only one policy at a time, and that is a static policy. Using iRules®, you can combine static and dynamic bandwidth control policies up to eight policies on a connection, but only one of the eight policies can be a dynamic policy. A packet is transmitted only when all the attached policies allow it. The system as a whole supports a maximum of 1024 policies.

Important: *Applying a bandwidth controller policy to a route domain affects all traffic transmitted by the BIG-IP system to VLANs in the route domain, including health monitors and DNS queries.*

Note: *Only static bandwidth control policies support SNMP queries.*

Bandwidth controllers vs. rate shaping

Bandwidth controller is the updated version of rate shaping on the BIG-IP® system. These features are mutually exclusive. You can configure and use either rate shaping or bandwidth controllers, but not both. Bandwidth controllers include distributed control, subscriber fairness, and support for a maximum rate of 320 Gbps. Rate shaping is hierarchical and supports minimum bandwidth (committed information rate), priority, and flow fairness.

About static bandwidth control policies

A *static* bandwidth control policy controls the aggregate rate for a group of applications or a network path. It enforces the total amount of bandwidth that can be used, specified as the maximum rate of the resource you are managing. The rate can be the total bandwidth of the BIG-IP® device, or it might be a group of traffic flows.

Task summary for creating a static bandwidth control policy

This procedure includes the steps for assigning a static bandwidth control policy to traffic, using a virtual server. Alternatively, you can assign a static bandwidth control policy to a packet filter or a route domain.

Task list

Creating a static bandwidth control policy

Adding a static bandwidth control policy to a virtual server

Creating a static bandwidth control policy

You can create a static bandwidth control policy to limit the bandwidth that traffic uses on the BIG-IP® system.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type a name for the bandwidth control policy.
4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for the resource you are managing.
The number must be in the range from 1 Mbps to 320 Gbps. This value is the amount of bandwidth available to all the connections going through this static policy.
5. Click **Finished**.

For the bandwidth control policy to take effect, you must apply the policy to traffic, using a virtual server, packet filter, or route domain.

Adding a static bandwidth control policy to a virtual server

Adding a static bandwidth control policy to a virtual server is one way to apply the policy to traffic. Alternatively, you can add the bandwidth control policy to a packet filter or a route domain.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. From the **Configuration** list, select **Advanced**.
4. From the **Bandwidth Controller** list, select a bandwidth control policy.
Only static bandwidth control policies are available in this list.
5. Click **Update** to save the changes.

The BIG-IP® system now applies rate enforcement to the traffic intercepted by this virtual server, according to the static bandwidth policy you selected. A static bandwidth policy associated with a virtual server applies only to client-initiated flows, and not to bandwidth for traffic flowing toward the client.

About dynamic bandwidth control policies

You can create dynamic bandwidth control policies to restrict bandwidth usage per subscriber or group of subscribers, per application, per network egress link, or any combination of these. A *dynamic* bandwidth control policy provides fairness on traffic flows, according to configurable parameters, within an upper bandwidth limit. The BIG-IP® system activates the dynamic bandwidth control policy for each user only when the user participates. When you create a dynamic bandwidth control policy, it acts as a policy in waiting, until the system detects egress traffic that matches the traffic you want to control and creates an instance of the policy. At that moment, the system applies the bandwidth control policy limits, as specified. No bandwidth control occurs until the system detects traffic and creates an instance of the policy. With this feature, an Internet service provider (ISP) can create and revise a single policy that can apply to millions of users.

The BIG-IP system can enforce multiple levels of bandwidth limits through the dynamic policy. For example, a user could be limited by the maximum rate, a per user rate, and a per category rate (such as for an application), all from the same dynamic policy. When the total of the maximum user rate for all the instances exceeds the maximum rate specified in the dynamic policy, the BIG-IP system maintains fairness among all users and spreads the limitation equally among users belonging to a dynamic policy. In addition, you can specify per instance the maximum number of packets per second, which functions as a DoS (Denial of Service) limiter without fair share allocation.

You can also configure a dynamic bandwidth control policy to mark packets that exceed the maximum per-user rate for a specified session. The WAN router typically handles the marked packets. When marking is enabled, enforcement is implicitly disabled. You configure marking by using the **IP Marking (TOS/DSCP)** or **L2 Marking (802.1p)** setting. For example, a common use of QoS marking is for Voice over IP (VoIP) traffic. VoIP is usually assigned to the Expedited Forwarding (EF) class by using the DSCP value of 46, thus prioritized according to importance and sensitivity to loss/latency. You can mark packets per policy or per category (within a policy). Category marking supersedes policy marking.

The bandwidth controller is only an enforcer. For a dynamic bandwidth control policy, you also can use a virtual server (through iRules®), Policy Enforcement Manager™, or Access Policy Manager® to identify users and apply dynamic bandwidth control policies to traffic.

About the Maximum User Rate PPS setting

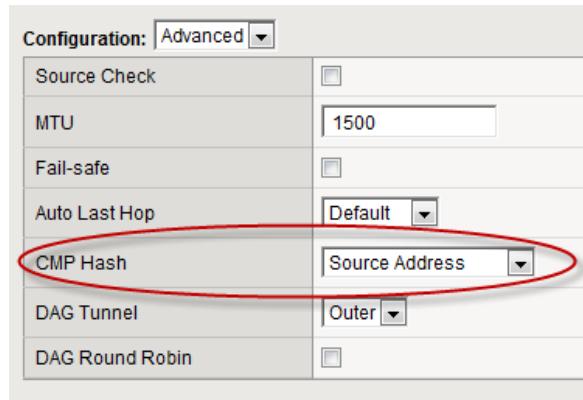
When you specify the value for the **Maximum User Rate PPS** setting for a dynamic bandwidth control policy, you are specifying the packets per second. Unlike the **Maximum Rate Per User**, this setting is not applied across a BIG-IP® system. This value depends on the packet sizes (MTU) configured in your network, and you need to tune the value accordingly. Although you can specify a lower value, the lowest recommended value is 256 KBPS. For a maximum user rate lower than 256 KBPS, the packet sizes in the network would need to be smaller than 1514 bytes. The issue is that for large packet sizes and a low maximum user rate, not enough bytes would be recharged for the bucket, and most packets would be dropped. If you are configuring categories, keep in mind that the maximum user rate is shared among the categories associated with a policy.

Task summary for creating a dynamic bandwidth control policy

Before you create a dynamic bandwidth control policy, F5 recommends that you select the **Source Address** for the **CMP Hash** setting on the VLAN properties screen for the VLAN that carries the traffic you want to manage. The BIG-IP® system uses source and destination hashes to control the way incoming traffic is distributed among the instances of the Traffic Management Microkernel (TMM) service. Subscriber-based

bandwidth control depends on having a unique one-to-one relationship between bandwidth control policy and subscriber. Subscribers are commonly identified using a unique IP address, and, therefore, load distribution among the instances of TMM service must use the source IP address as the key.

This screen snippet highlights the proper setting.



The screenshot shows a configuration window with a 'Configuration:' dropdown set to 'Advanced'. Below it is a table of settings:

Source Check	<input type="checkbox"/>
MTU	1500
Fail-safe	<input type="checkbox"/>
Auto Last Hop	Default
CMP Hash	Source Address
DAG Tunnel	Outer
DAG Round Robin	<input type="checkbox"/>

The 'CMP Hash' row is circled in red, indicating the correct setting for dynamic bandwidth control.

Figure 7: CMP Hash setting for dynamic bandwidth control

This procedure describes the steps for attaching a dynamic bandwidth control policy to a traffic flow, and then applying the policy to traffic, using a virtual server. For information about using Policy Enforcement Manager™ to implement the policy, refer to the F5 documentation for Policy Enforcement Manager.

Task list

- Creating a dynamic bandwidth control policy*
- Adding categories to a dynamic bandwidth control policy*
- Creating an iRule for a dynamic bandwidth control policy*
- Adding a dynamic bandwidth control policy to a virtual server*

Creating a dynamic bandwidth control policy

You can create a dynamic bandwidth control policy to shape the traffic to which you apply the policy. You can configure the policy to mark packets per policy or per category. You can also specify the maximum number of packets per second per instance. Adding categories to a bandwidth control policy is a separate task, which you perform after you have created and saved the policy.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type a name for the bandwidth control policy.
4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for all the instances created for this dynamic policy.
The number must be in the range from 1 Mbps to 320 Gbps.
5. From the **Dynamic** list, select **Enabled**.
The screen displays additional settings.
6. In the **Maximum Rate Per User** field, type a number and select the unit of measure to indicate the most bandwidth that each user or session associated with the bandwidth control policy can use.
The number must be in the range from 1 Mbps to 2 Gbps.
7. In the **Maximum User Rate PPS** field, type a number and select the unit of measure to specify the limit in packets per second that traffic is allowed per instance.

This is an optional setting, which functions as a DoS (Denial of Service) limiter without fair share allocation. You must also specify the **Maximum Rate Per User**. The system applies whichever value is lower to the traffic flow rate. When both values are specified, both must pass for packets to go through.

8. Enable the **Measure** setting, if you want to measure bandwidth on all future instances of this bandwidth control policy.

The system measures bandwidth with the frequency you specify in the **Log Period** setting, and sends it to the log publisher you specify using the **Log Publisher** setting.

9. From the **IP Marking (TOS/DSCP)** list, select **Specify** and type a number between 0 and 63 to assign a Type of Service (ToS) level to packets that exceed the maximum per-user rate.

If you do not want to set a ToS level, maintain the default setting, **Pass Through**.

10. From the **L2 Marking (802.1p)** list, select **Specify** and type a number between 0 and 7 to assign a Quality of Service (QoS) level to packets that exceed the maximum per-user rate.

If you do not want to set a QoS level, maintain the default setting, **Pass Through**.

11. Click **Finished**.

For the dynamic bandwidth control policy to take effect, you must attach the policy to a traffic flow, and then apply the policy to traffic, using a virtual server (through iRules®), Policy Enforcement Manager™, or Access Policy Manager®.

Adding categories to a dynamic bandwidth control policy

Before you can add categories, you must create a bandwidth control policy.

After you create a bandwidth control policy, you can add up to 32 categories of traffic for the policy to control. All the categories share the bandwidth specified for the bandwidth control policy, in accordance with the rate specified for each category.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click the name of the bandwidth control policy to which you want to add categories.
3. In the Categories area, click **Add**.
4. In the **Category Name** field, type a descriptive name for the category.
5. In the **Max Category Rate** field, type a value to indicate the most bandwidth that this category of traffic can use, and select the unit of measure from the list, or select **%** and type a percentage from 1 to 100.
If you specify a rate, the number must be in the range from 500 Kbps to the rate specified for the **Maximum Rate Per User** setting. A percentage indicates that this category can use up to the specified percentage of the maximum per-user rate. These values are upper limits (not minimum or guaranteed), so the sum can exceed the value you specified for the **Maximum Rate Per User** setting.
6. From the **IP Marking (TOS/DSCP)** list, select **Specify** and type a number between 0 and 63 to assign a Type of Service (ToS) level to packets that exceed the **Max Category Rate**.
If you do not want to set a ToS level, maintain the default setting, **Pass Through**.
7. From the **L2 Marking (802.1p)** list, select **Specify** and type a number between 0 and 7 to assign a Quality of Service (QoS) level to packets that exceed the **Max Category Rate**.
If you do not want to set a QoS level, maintain the default setting, **Pass Through**.
8. Click **Finished**.

Creating an iRule for a dynamic bandwidth control policy

To implement a dynamic bandwidth control policy, you can use iRules® to attach the policy to a user.

Note: For complete and detailed information iRules syntax, see the F5 Networks DevCentral web site (<http://devcentral.f5.com>).

1. On the Main tab, click **Local Traffic > iRules**.
The iRule List screen opens, displaying any existing iRules.
2. Click **Create**.
The New iRule screen opens.
3. In the **Name** field, type a unique name for the iRule.
The full path name of the iRule cannot exceed 255 characters.
4. In the **Definition** field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.
For example, to apply the dynamic bandwidth policy `dynamic_bwc_policy200` to a user session, type the following iRule, where `set mycookie` defines a user session. A *session* is a combination of client IP address and port.

```
when CLIENT_ACCEPTED {  
    set mycookie [IP::remote_addr]:[TCP::remote_port]  
    BWC::policy attach dynamic_bwc_policy200 $mycookie  
}
```

5. Click **Finished**.
The new iRule appears in the list of iRules on the system.

You have now identified the user for a dynamic bandwidth control policy.

You must then apply the iRule to the virtual server that intercepts the traffic you want to manage.

Adding a dynamic bandwidth control policy to a virtual server

After you create a dynamic bandwidth control policy and attach it to a flow or flows using iRules[®], you must apply the policy to traffic by adding the iRule to a virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Resources**.
4. In the iRules area, click **Manage**.
5. From the **Available** list, select the name of the iRule that you want to assign, and using the Move button, move the name to the **Enabled** list.
6. Click **Finished**.

The BIG-IP[®] system now manages bandwidth for the traffic intercepted by this virtual server, according to the dynamic bandwidth policy specified in the assigned iRule.

Configuring Network Virtualization Tunnels

Overview: Configuring network virtualization tunnels

Large data centers and cloud service providers are benefiting from large scale network virtualization. Network Virtualization provides connectivity in cloud environments by overlaying Layer 2 segments over a Layer 3 infrastructure. The overlay network can be dynamically extended with multiple virtualized networks without affecting the Layer 3 infrastructure. This number of virtualized networks is typically much larger than the number of VLANs the infrastructure can support.

You can configure a BIG-IP® system to function as a gateway in a virtualized network, bridging the data center virtualized networks with the physical network (L2 gateway), or performing routing and higher L4-L7 functionality among virtual networks of different types (L3 gateway). Connecting these networks allows for expansion, and provides a mechanism to streamline the transition of data centers into a virtualized model, while maintaining connectivity.

This illustration shows the BIG-IP system as a network virtualization gateway.




Figure 8: The BIG-IP system as a network virtualization gateway

In a virtualized network, the BIG-IP system needs to learn about other virtualization tunnel endpoints. Each hypervisor has a tunnel endpoint. The hypervisor needs to locate the virtual machines it manages, by maintaining a form of the L2 location records, typically, IP addresses and MAC addresses, virtual network identifiers, and virtual tunnel endpoints.

About network virtualization tunnels on the BIG-IP system

When you configure a BIG-IP® system as a network virtualization gateway, the system represents the connection as a tunnel, which provides a Layer 2 interface on the virtual network. You can use the tunnel interface in both Layer 2 and Layer 3 configurations. After you create the network virtualization tunnels, you can use the tunnels like you use VLANs on a BIG-IP system, such as for routing, assigning self IP addresses, and associating with virtual servers.

Creating a network virtualization tunnel

Creating a network virtualization tunnel on a BIG-IP® system provides an L2 gateway to connect the physical underlay network with a virtual overlay network.

1. On the Main tab, click **Network > Tunnels > Tunnel List > Create** or **Carrier Grade NAT > Tunnels > Create**.

The New Tunnel screen opens.

2. In the **Name** field, type a unique name for the tunnel.
3. From the **Profile** list, select the tunnel profile you created for network virtualization.
This selection must be a profile based on either the `gre` or `vxlان` parent profile, depending on your virtualized network environment.
4. In the **Local Address** field, type the self IP address of the VLAN through which the remote hypervisor is reachable.
5. For the **Remote Address** list, retain the default selection, **Any**.
6. In the **Key** field, type the VNI (Virtual Network Identifier) to use for a VXLAN tunnel or the Virtual Subnet Identifier (VSID) to use for a NVGRE tunnel.
This field appears above the **Profile** field when you select a profile that requires this setting.
7. Click **Finished**.

This tunnel is now available to use in virtualized network routing configurations, depending on how you configure your network.

Virtualized network terminology

These terms are associated with virtualized networks.

forwarding database (FDB)

The *FDB* is the database that contains mappings between the MAC address of each virtual machine and the IP address of the hypervisor machine on which it resides.

L2 gateway

The Layer 2 gateway performs the bridge functionality between VLAN and virtual segments in a virtualized network.

L3 gateway

The Layer 3 gateway performs routing and higher L4-L7 functionality among virtualized network segments of different types.

overlay network

The *overlay network* is a virtual network of VMs built on top of a stable L2-L3 structure. The view from one VM to another is as if they were on the same switch, but, in fact, they could be far afield.

tunnel endpoint

A *tunnel endpoint* originates or terminates a tunnel. In a virtualized network environment, the tunnel IP addresses are part of the L2 underlay network. The same local IP address can be used for multiple tunnels.

underlay network

The *underlay network* is the L2 or L3 routed physical network, a mesh of tunnels.

virtualized network

A *virtualized network* is when you create a virtual L2 or L3 topology on top of a stable physical L2 or L3 network. Connectivity in the virtual topology is provided by tunneling Ethernet frames in IP over the physical network.

VNI

The *Virtual Network Identifier (VNI)* is also called the VXLAN segment ID. The system uses the VNI to identify the appropriate tunnel.

VSID

The *Virtual Subnet Identifier (VSID)* is a 24-bit identifier used in an NVGRE environment that represents a virtual L2 broadcast domain, enabling routes to be configured between virtual subnets.

VTEP

The *VXLAN Tunnel Endpoint (VTEP)* originates or terminates a VXLAN tunnel. The same local IP address can be used for multiple tunnels.

VXLAN

Virtual eXtended LAN (VXLAN) is a network virtualization scheme that overlays Layer 2 over Layer 3. VXLAN uses Layer 3 multicast to support the transmission of multicast and broadcast traffic in the virtual network, while decoupling the virtualized network from the physical infrastructure.

VXLAN gateway

A *VXLAN gateway* bridges traffic between VXLAN and non-VXLAN environments. The BIG-IP® system uses a VXLAN gateway to bridge a traditional VLAN and a VXLAN network, by becoming a network virtualization endpoint.

VXLAN header

In addition to the UDP header, encapsulated packets include a *VXLAN header*, which carries a 24-bit VNI to uniquely identify Layer 2 segments within the overlay.

VXLAN segment

A *VXLAN segment* is a Layer 2 overlay network over which VMs communicate. Only VMs within the same VXLAN segment can communicate with each other.

Centralized vs. decentralized models of network virtualization

Using the BIG-IP® system as a network virtualization gateway, you can set up virtualized network segments using either a centralized or decentralized model.

Centralized model

In a centralized model, a network orchestrator or controller manages the virtualized network segments. The orchestrator has full view of VTEPs, L2, and L3 information in the overlay, and is responsible for pushing this information to hypervisors and gateways. Microsoft Hyper-V and VMware NSX environments use this model.




Figure 9: Centralized model of network virtualization

Decentralized model

A decentralized model of network virtualization does not require a network orchestrator or controller. In this model, the router learns the tunnel endpoint and MAC address locations by flooding broadcast, multicast, and unknown destination frames over IP multicast. VMware vSphere 5.1 environments use this model.




Figure 10: Decentralized model of network virtualization

About network virtualization tunnel types

The BIG-IP® system supports multiple network virtualization tunnel types. You can even combine virtualized network segments based on different tunnel types. This table offers a quick comparison of the tunnel types.

VXLAN (Multicast)	VXLAN (Unicast)	NVGRE	Transparent Ethernet Bridging
Decentralized	Centralized	Centralized	Centralized

VXLAN (Multicast)	VXLAN (Unicast)	NVGRE	Transparent Ethernet Bridging
VMware vSphere 5.1	VMware NSX	Microsoft SCVMM/Hyper-V	OpenStack
VXLAN UDP Encapsulation	VXLAN UDP Encapsulation	GRE-based Encapsulation	GRE-based Encapsulation
24-bit ID	24-bit ID	24-bit ID	32-bit ID
Endpoints discovered dynamically	Endpoints statically configured	Endpoints statically configured	Endpoints statically configured
Floods unknown and broadcast frames using IP multicast.	Can flood using unicast replication.	Does not flood (completely static).	Floods using unicast replication.

In addition to the above types of tunnels, the BIG-IP system supports the creation of Geneve Network Virtualization tunnels. Like VXLAN, Geneve tunnels support multicast and multipoint flooding.

Note: IPv4 multicast addresses in the local network control block (224.0.0/24) [RFC 5771] should not be used for configuring the remote address of the VXLAN/Geneve tunnels with multicast flooding.

About statically configured network virtualization tunnels

For the centralized model, you can use VXLAN (Unicast), NVGRE, or Transparent Ethernet Bridging, depending on the cloud environment. Using an agent or plug-in, or the `tmsh` command-line utility, you can statically configure the FDB and ARP forwarding table entries. Using the `tmsh` command-line utility or browser interface, you can create the network virtualization tunnels, which are managed by the network controller.

Considerations for statically configured network virtualization tunnels

As you configure a BIG-IP® system to be an L2 or L3 gateway for statically configured network virtualization tunnels, keep these considerations in mind.

- The BIG-IP system must be licensed for SDN Services.
- If you have over 2000 tunnels, set the **Management (MGMT)** setting on the Resource Provisioning screen is to **Large (System > Resource Provisioning)**.

Examples for manually populating L2 location records

Using the `tmsh` command-line utility, you can add static FDB records and ARP entries for each virtual tunnel endpoint.

- Add static FDB (forwarding database) entries to associate MAC addresses with specified tunnel endpoints. For example, the following command creates an FDB entry that associates the MAC address 00:01:02:03:04:05 with the tunnel endpoint 10.1.1.1 of the tunnel vxlan0.

```
# tmsh modify net fdb tunnel vxlan0 records add {  
    00:01:02:03:04:05 { endpoint 10.1.1.1 } }
```

- Delete a MAC address from an FDB entry.

```
# tmsh modify net fdb tunnel vxlan0 records delete { 00:01:02:03:04:05 }
```

- Delete a static ARP.

```
# tmsh delete net arp 10.3.3.1
```

- Add an IP address to a MAC address in the ARP table.

```
# tmsh create net arp 10.3.3.1 { ip-address 10.3.3.1 mac-address  
    00:01:02:03:04:05 }
```

Using the iControl/REST API, you can program a network controller to build and maintain network virtualization tunnels. This example adds an entry to the FDB table that associates the MAC address 00:01:02:03:04:05 with the tunnel endpoint 10.1.1.2 of the tunnel vxlan0-tunnel.

```
$ curl -u admin:f5site02 -H "Content-Type:=application/json" -k -X PUT  
'https://172.30.69.69/mgmt/tm/net/fdb/tunnel/~Common~vxlan0-tunnel' -d  
'{"kind":"tm:net:fdb:tunnel:tunnelstate","name":"vxlan0-tunnel","partition":"Common",  
"fullPath":"/Common/vxlan0-tunnel","generation":1,  
"selfLink":"https://localhost/mgmt/tm/net/fdb/tunnel/~Common~vxlan0-tunnel?  
ver=11.5.0","records":[{"name":"00:01:02:03:04:05",  
"endpoint":"10.1.1.2"}]}' |python -m json.tool  
{  
    "fullPath": "/Common/vxlan0-tunnel",  
    "generation": 1,  
    "kind": "tm:net:fdb:tunnel:tunnelstate",  
    "name": "vxlan0-tunnel",  
    "partition": "Common",  
    "records": [  
        {  
            "endpoint": "10.1.1.2",  
            "name": "00:01:02:03:04:05"  
        }  
    ],  
    "selfLink":  
    "https://localhost/mgmt/tm/net/fdb/tunnel/~Common~vxlan0-tunnel?ver=11.5.0"  
}
```

Sample NVGRE configuration using tmsh

This listing example illustrates the steps for creating a routing configuration that includes an NVGRE tunnel on the BIG-IP® system. F5 Networks provides an API for you to configure the F5 SCVMM Gateway Provider plug-in to build and manage NVGRE tunnels.

```
create net vlan wan {
    interfaces add { 1.1 }
    mtu 1550
}
create net self 10.1.1.1/24 {
    address 10.1.1.1/24
    vlan wan
}
create net tunnels gre nvgre {
    encapsulation nvgre
}
create net tunnels tunnel nvgre5000 {
    local-address 10.1.1.1
    remote-address any
    profile nvgre
    key 5000
}
create net vlan legacy5000 {
    interfaces add { 2.1 }
}
create net route-domain 5000 {
    id 5000
    vlans add { nvgre5000 legacy5000 }
}
create net self 10.3.3.1%5000/24 {
    address 10.3.3.1%5000/24
    vlan nvgre5000
}
create net self 10.4.4.1%5000/24 {
    address 10.4.4.1%5000/24
    vlan legacy5000
}
create net route 10.5.5.0%5000/24 {
    network 10.5.5.0%5000/24
    gw 10.3.3.2%5000
}
create net route 10.6.6.0%5000/24 {
    network 10.6.6.0%5000/24
    gw 10.3.3.3%5000
}
modify net fdb tunnel nvgre5000 {
    records add {
        00:FF:0A:03:03:02 { endpoint 10.1.2.1 }
        00:FF:0A:03:03:03 { endpoint 10.1.3.1 }
    }
}
create net arp 10.3.3.2%5000 {
    mac-address 00:FF:0A:03:03:02
}
create net arp 10.3.3.3%5000 {
    mac-address 00:FF:0A:03:03:03
}
```

Sample VXLAN unicast configuration using tmsh

This example listing illustrates the steps for creating a routing configuration that includes a VXLAN tunnel on the BIG-IP® system. This configuration adds the tunnel to a route domain. You can use the iControl/REST API to configure a network controller to build and manage VXLAN (unicast) tunnels.

```
create net vlan wan {
    interfaces add { 1.1 }
    mtu 1550
}
create net self 10.1.1./24 {
    address 10.1.1.1/24
    vlan wan
}
create net tunnels vxlan vxlan-static {
    flooding-type none
}
create net tunnels tunnel vxlan5000 {
    local-address 10.1.1.1
    remote-address any
    profile vxlan-static
    key 5000
}
create net vlan legacy5000 {
    interfaces add { 2.1 }
}
create net route-domain 5000 {
    id 5000
    vlans add { vxlan5000 legacy5000 }
}
create net self 10.3.3.1%5000/24 {
    address 10.3.3.1%5000/24
    vlan vxlan5000
}
create net self 10.4.4.1%5000/24 {
    address 10.4.4.1%5000/24
    vlan legacy5000
}
create net route 10.5.5.0%5000/24 {
    network 10.5.5.0%5000/24
    gw 10.3.3.2%5000
}
create net route 10.6.6.0%5000/24 {
    network 10.6.6.0%5000/24
    gw 10.3.3.3%5000
}
modify net fdb tunnel vxlan5000 {
    records add {
        00:FF:0A:03:03:02 { endpoint 10.1.2.1 }
        00:FF:0A:03:03:03 { endpoint 10.1.3.1 }
    }
}
create net arp 10.3.3.2%5000 {
    mac-address 00:FF:0A:03:03:02
}
create net arp 10.3.3.3%5000 {
    mac-address 00:FF:0A:03:03:03
}
```


Sample command for virtual server to listen on a VXLAN tunnel

An alternative for including a network virtualization tunnel in a routing configuration is to create a virtual server that listens for the tunnel traffic, such as in the following example.

```
# tmsh create ltm virtual http_virtual destination 10.3.3.15%5000:http
ip-protocol tcp vlans add { vxlan5000 }
```

The code in this example creates a virtual server `http_virtual` that listens for traffic destined for the IP address `10.3.3.15` on the tunnel named `vxlan5000`.

Commands for viewing tunnel statistics

You can use the `tmsh` command-line utility to view tunnel statistics, listing either all the tunnels on the BIG-IP® system or statistics about a particular tunnel.

View per-tunnel statistics:

```
# tmsh show net tunnels tunnel
```

View static and dynamic FDB entries:

```
# tmsh show net fdb tunnel
```

About VXLAN multicast configuration

In a VMware vSphere 5.1 environment, you can configure VXLAN without knowing all the remote tunnel endpoints. The BIG-IP® system uses multicast flooding to learn unknown and broadcast frames. VXLAN can extend the virtual network across a set of hypervisors, providing L2 connectivity among the hosted virtual machines (VMs). Each hypervisor represents a VXLAN tunnel endpoint (VTEP). In this environment, you can configure a BIG-IP system as an L2 VXLAN gateway device to terminate the VXLAN tunnel and forward traffic to and from a physical network.

About bridging VLAN and VXLAN networks

You can configure Virtual eXtended LAN (VXLAN) on a BIG-IP® system to enable a physical VLAN to communicate with virtual machines (VMs) in a virtual network.




Figure 11: The VXLAN gateway

When you configure a BIG-IP system as an L2 VXLAN gateway, the BIG-IP system joins the configured multicast group, and can forward both unicast and multicast or broadcast frames on the virtual network. The BIG-IP system learns about MAC address and VTEP associations dynamically, thus avoiding unnecessary transmission of multicast traffic.

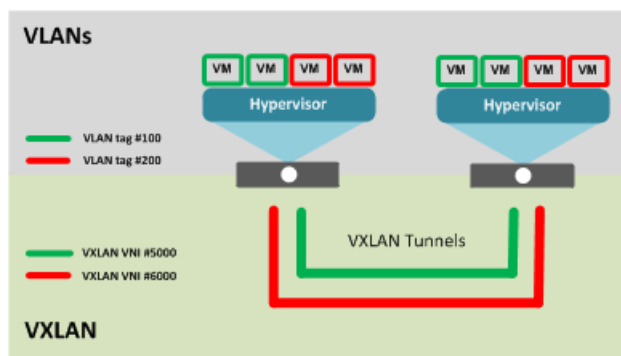


Figure 12: Multiple VXLAN tunnels

Considerations for configuring multicast VXLAN tunnels

As you configure VXLAN on a BIG-IP[®] system, keep these considerations in mind.

- If you configure the BIG-IP device as a bridge between physical VLANs and a VXLAN tunnel, the number of virtualized network segments in the overlay is limited to the maximum number of physical VLANs (4094). This limitation does not apply to Layer 3 configurations.
- You need to configure a separate tunnel for each VNI. The tunnels can have the same local and remote endpoint addresses.
- For the Layer 2 network, you must ensure a loop-free topology.
- Do not modify the configuration of a VXLAN tunnel after it is created. Instead, delete the existing tunnel and create a new one.

Task summary

Before you configure VXLAN, ensure that these conditions are met:

- The BIG-IP® system must be licensed for SDN Services.
- Network connectivity exists between the BIG-IP system and the hypervisors.
- If you have over 2000 tunnels, the **Management (MGMT)** setting on the Resource Provisioning screen is set to **Large (System > Resource Provisioning)**.

Task list

Specifying a port number

Modifying a VLAN for disaggregation of VXLAN tunnel traffic

Creating a multicast VXLAN tunnel

Creating a bridge between VXLAN and non-VXLAN networks

Specifying a port number

Before you perform this task, confirm that you have enabled the **DAG Tunnel** setting on the relevant VLAN.

When you enable the DAG tunnel feature on a VLAN, you must also configure a `bigdb` variable that specifies a port number so that associated tunnels can disaggregate based on the inner header of a packet.

1. Open the Traffic Management Shell (tmsh).

```
tmsh
```

2. Specify a port number to be used.

```
modify sys db iptunnel.vxlan.udpport value <port_number>
```

The value that you specify with this `bigdb` variable applies to all VLANs on which the **DAG Tunnel** setting is enabled.

Important: Typically, a tunnel uses port 4789. If you choose to use a different port number, you must ensure that the port number specified in the relevant VXLAN profile matches the value you set with this command.

Modifying a VLAN for disaggregation of VXLAN tunnel traffic

You perform this task when you want to use an existing VLAN with a VXLAN tunnel that disaggregates traffic based on the inner header of the packet (hardware-disaggregated or DAG tunnel).

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the Name column, click the relevant VLAN name.
This displays the properties of the VLAN.
3. From the **DAG tunnel** list, select **Inner**.
This disaggregates encapsulated packets based on the inner headers.
4. Click **Finished**.

Creating a multicast VXLAN tunnel

Creating a VXLAN multicast tunnel on a BIG-IP® system provides an L2 VXLAN gateway to connect the physical network with a virtualized network.

1. On the Main tab, click **Network > Tunnels > Tunnel List > Create or Carrier Grade NAT > Tunnels > Create**.

The New Tunnel screen opens.

2. In the **Name** field, type a unique name for the tunnel.
3. From the **Profile** list, select **vxlان**.

This setting tells the system which tunnel profile to use. The system-supplied VXLAN profile specifies port 4789. To change the port number, you can create a new VXLAN profile, which then appears in this list.
4. In the **Local Address** field, type the self IP address of the VLAN through which the remote hypervisor is reachable.
5. In the **Remote Address** field, type the multicast group address associated with the VXLAN segment.
6. For the **Mode** list, retain the default selection, **Bidirectional**.
7. In the **MTU** field, type the maximum transmission unit of the tunnel.

The default value is 0. The valid range is from 0 to 65515.
8. For the **Use PMTU** (Path MTU) setting, select or clear the check box.

If enabled and the tunnel MTU is set to 0, the tunnel will use the PMTU information. If enabled and the tunnel MTU is fixed to a non-zero value, the tunnel will use the minimum of PMTU and MTU. If disabled, the tunnel will use fixed MTU or calculate its MTU using tunnel encapsulation configurations.
9. From the **TOS** list, select **Preserve**, or select **Specify** and type a Type of Service (TOS) value.

The valid range is from 0 to 255.
10. From the **Auto-Last Hop** list, select a value.
 - Choose **Default** if you want the system to use the global **Auto Last Hop** setting (if enabled).
 - Choose **Enabled** if you want this setting to take precedence over the global **Auto Last Hop** setting, for this tunnel only.
 - Choose **Disabled** if you want to disable auto last hop behavior for this tunnel only.
11. From the **Traffic Group** list, select the traffic group that includes the local IP address for the tunnel.
12. Click **Finished**.

Creating a bridge between VXLAN and non-VXLAN networks

Before you begin this task, verify that a VXLAN multicast tunnel exists on the BIG-IP® system.

You can create a VLAN group to bridge the traffic between a VXLAN overlay network (Layer 3) and a non-VXLAN (Layer 2) network.

1. On the Main tab, click **Network > VLANs > VLAN Groups**.

The VLAN Groups list screen opens.
2. Click **Create**.

The New VLAN Group screen opens.
3. In the **Name** field, type a unique name for the VLAN group.
4. For the **VLANs** setting, select the VLAN that connects to the non-VXLAN Layer-2 network and the VXLAN tunnel you created, and using the Move button (<<), move your selections from the **Available** list to the **Members** list.
5. Click **Finished**.

About configuring VXLAN tunnels on high availability BIG-IP device pairs

By default, the BIG-IP® system synchronizes all existing tunnel objects in its config sync operation. This operation requires that the local IP address of a tunnel be set to a floating self IP address. In a high availability (HA) configuration, any tunnel with a floating local IP address would be available only on the active device, which would prevent some features, such as health monitors, from using the tunnel on the standby device. To make a tunnel available on both the active and standby devices, you need to set the local IP address to a non-floating self IP address, which then requires that you exclude tunnels from the config sync operation. To disable the synchronization of tunnel objects, you can set a `bigdb` variable on both devices.

Disabling config sync for tunnels

In certain cases, you might want to disable config sync behavior for tunnels, such as when you need to make VXLAN tunnels functional on all devices in a BIG-IP® device group configured for high availability. The tunnel config sync setting applies to all tunnels created on the BIG-IP device.

Important: *Disable config sync on both the active and standby devices before you create any tunnels.*

1. Log in to the `tmsh` command-line utility for the BIG-IP system.
2. Determine whether the variable is already disabled, by typing this command.

```
tmsh list sys db iptunnel.configsync value
```
3. Disable the variable.

```
tmsh modify sys db iptunnel.configsync value disable
```
4. Save the configuration.

```
tmsh save sys config
```
5. F5 recommends that you reboot both the active and standby devices.

Now you can create tunnels with non-floating local IP addresses on both the active and standby devices.

About configuring VXLAN tunnels using OVSDB

You can configure the OVSDB management component of the BIG-IP® system to create VXLAN tunnels using an SDN controller. SDN Controllers use the OVSDB (Open vSwitch Database) management protocol to pass overlay configuration information to hypervisors and gateways. The BIG-IP system supports using the Hardware VTEP schema with the OVSDB management protocol to communicate with OVSDB-capable controllers. The system uses the OVSDB management component to set up this communication. After the BIG-IP system establishes communication with an OVSDB controller, it can begin to receive tunnel endpoint configuration information. The system uses this information to create VXLAN tunnel objects and their associated FDB entries. The created objects allow the system to communicate with other VXLAN endpoints in an overlay segment.

Setting up the OVSDB management component

You can set up the OVSDB management component to communicate with OVSDB-capable controllers to create VXLAN tunnel objects.

1. On the Main tab, click **System > Configuration > OVSDDB**.
2. From the **OVSDDB** list, select **Enable**.
3. In the **Controller Addresses** field, type an OVSDDB-capable controller IP address and click **Add**.
If the BIG-IP system is communicating with a controller cluster, repeat this step for each controller IP address.
4. In the **Tunnel Local Address** field, type the IP address of the local endpoint of the tunnel.
The value must be an IPv4 address. The OVSDDB management component uses the value in this field to configure the **Local Address** setting when creating VXLAN tunnel objects.
5. In the **Port** field, type the controller's port or retain the default.
6. Configure the SSL certificate settings.
 - a) From the **Certificate File** list, select a certificate file to be presented to the controller.
 - b) From the **Certificate Key File** list, select the certificate key file that has the private key.
 - c) From the **CA Certificate File** list, select the CA certificate file.
This is the file containing the CA certificate used to validate the certificates presented by the controller.
7. Click **Update**.

About configuring VXLAN-GPE tunnels

You can configure a VXLAN Generic Protocol Extension (GPE) tunnel when you want to add fields to the VXLAN header. One of these fields is `Next Protocol`, with values for Ethernet, IPv4, IPv6, and Network Service Header (NSH).

Creating a multicast VXLAN-GPE tunnel

Creating a VXLAN Generic Protocol Extension (GPE) multicast tunnel on a BIG-IP® system provides an L2 VXLAN gateway to connect the physical network with a virtualized network. Unlike a standard VXLAN tunnel, this tunnel type supports the processing of VXLAN GPE-encapsulated Ethernet frames.

1. On the Main tab, click **Network > Tunnels > Tunnel List > Create or Carrier Grade NAT > Tunnels > Create**.
The New Tunnel screen opens.
2. In the **Name** field, type a unique name for the tunnel.
3. From the **Profile** list, select **vxlan-gpe**.
This setting tells the system which tunnel profile to use. The system-supplied `vxlan-gpe` profile specifies port 4790. To change the port number, you can create a new VXLAN-GPE profile, which then appears in this list.
4. In the **Local Address** field, type the self IP address of the VLAN through which the remote hypervisor is reachable.
5. In the **Remote Address** field, select **Any**, or select **Specify** and type the multicast group address associated with the VXLAN-GPE segment.
6. For the **Mode** list, retain the default selection, **Bidirectional**.
7. In the **MTU** field, type the maximum transmission unit of the tunnel.
The default value is 0. The valid range is from 0 to 65515.

8. For the **Use PMTU** (Path MTU) setting, select or clear the check box.
If enabled and the tunnel MTU is set to 0, the tunnel will use the PMTU information. If enabled and the tunnel MTU is fixed to a non-zero value, the tunnel will use the minimum of PMTU and MTU. If disabled, the tunnel will use fixed MTU or calculate its MTU using tunnel encapsulation configurations.
9. From the **TOS** list, select **Preserve**, or select **Specify** and type a Type of Service (TOS) value.
The valid range is from **0** to **255**.
10. From the **Auto-Last Hop** list, select a value.
 - Choose **Default** if you want the system to use the global **Auto Last Hop** setting (if enabled).
 - Choose **Enabled** if you want this setting to take precedence over the global **Auto Last Hop** setting, for this tunnel only.
 - Choose **Disabled** if you want to disable auto last hop behavior for this tunnel only.
11. From the **Traffic Group** list, select the traffic group that includes the local IP address for the tunnel.
12. Click **Finished**.

Web Hosting Multiple Customers Using an External Switch

Overview: Web hosting multiple customers using an external switch

You can use the BIG-IP® system to provide hosting services, including application delivery, for multiple customers.

To host multiple web customers, you can incorporate an external switch into the configurations. In this illustration, the BIG-IP system has an interface (5.1) assigned to three VLANs on a network. The three VLANs are **vlanA**, **vlanB**, and **vlanB**. Interface **5.1** processes traffic for all three VLANs. Note that each VLAN contains two servers, and serves a specific customer.

Tip: An alternate way to implement web hosting for multiple customers is to use the route domains feature.

Illustration for hosting multiple customers using an external switch




Figure 13: Hosting multiple customers using an external switch

Task summary for hosting multiple customers

Perform these tasks to host multiple customers using an external switch.

Task list

Creating a VLAN with a tagged interface

Creating a load balancing pool

Creating a virtual server for HTTP traffic

Creating a VLAN with a tagged interface

When you create a VLAN with tagged interfaces, each of the specified interfaces can process traffic destined for that VLAN.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. Click **Create**.
The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.
The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged**.
 - c) Click **Add**.
6. If you want the system to verify that the return route to an initial packet is the same VLAN from which the packet originated, select the **Source Check** check box.
7. In the **MTU** field, retain the default number of bytes (**1500**).
8. From the **Configuration** list, select **Advanced**.
9. If you want to base redundant-system failover on VLAN-related events, select the **Fail-safe** check box.
10. From the **Auto Last Hop** list, select a value.
11. From the **CMP Hash** list, select a value.
12. To enable the **DAG Round Robin** setting, select the check box.
13. Click **Finished**.
The screen refreshes, and displays the new VLAN in the list.

The new VLAN appears in the VLAN list.

Creating a load balancing pool

You can create a *load balancing pool* (a logical set of devices such as web servers that you group together to receive and process traffic) to efficiently distribute the load on your server resources.

Note: *You must create the pool before you create the corresponding virtual server.*

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click << to move the monitor to the **Active** list.

Tip: Hold the Shift or Ctrl key to select more than one monitor at a time.

5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.
The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
 - Select **Disabled** to disable priority groups. This is the default option.
 - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.
 - c) In the **Service Port** field, type a port number, or select a service name from the list.
 - d) (Optional) In the **Priority** field, type a priority number.
 - e) Click **Add**.
8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server for HTTP traffic

This task creates a destination IP address for application traffic. As part of this task, you must assign the relevant pool to the virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address in CIDR format.
The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is 10.0.0.1 or 10.0.0.0/24, and an IPv6 address/prefix is ffe1::0020/64 or 2001:ed8:77b5:2:10:10:100:42/64. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Note: The IP address you type must be available and not in the loopback network.

5. In the **Service Port** field, type 80, or select **HTTP** from the list.
6. From the **HTTP Profile** list, select **http**.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. Click **Finished**.

You now have a virtual server to use as a destination address for application traffic.

Web Hosting Multiple Customers Using Untagged Interfaces

Overview: Web hosting multiple customers using untagged interfaces

One way to implement web hosting for multiple customers is to use multiple interfaces on the BIG-IP® system to directly host traffic for multiple customers, without the need for an external switch. With this scenario, you must configure the VLANs with untagged instead of tagged interfaces. As shown in the following illustration, two BIG-IP system interfaces are assigned to each VLAN. For example, interfaces 1.1 and 1.2 are assigned to VLAN **vlanA**. Each interface is assigned to a VLAN as an untagged interface.

Tip: An alternate way to implement web hosting for multiple customers is to use the route domains feature.

Illustration for hosting multiple customers using untagged interfaces




Figure 14: Hosting multiple customers using untagged interfaces

Task summary for hosting multiple customers

Perform these tasks to host multiple customers using tagged interfaces on VLANs.

Task list

Creating a VLAN with an untagged interface

Creating a load balancing pool

Creating a virtual server for HTTP traffic

Creating a VLAN with an untagged interface

You can create a VLAN that uses untagged interfaces.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. Click **Create**.
The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.
The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. For the **Interfaces** setting,
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Untagged**.
 - c) Click **Add**.
6. Click **Finished**.
The screen refreshes, and displays the new VLAN in the list.

The interfaces that you specified in this task process traffic for this VLAN only.

Creating a load balancing pool

You can create a *load balancing pool* (a logical set of devices such as web servers that you group together to receive and process traffic) to efficiently distribute the load on your server resources.

Note: You must create the pool before you create the corresponding virtual server.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click << to move the monitor to the **Active** list.

Tip: Hold the *Shift* or *Ctrl* key to select more than one monitor at a time.

5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.
The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
 - Select **Disabled** to disable priority groups. This is the default option.
 - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:

- a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
- b) In the **Address** field, type an IP address.
- c) In the **Service Port** field, type a port number, or select a service name from the list.
- d) (Optional) In the **Priority** field, type a priority number.
- e) Click **Add**.

8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server for HTTP traffic

This task creates a destination IP address for application traffic. As part of this task, you must assign the relevant pool to the virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address in CIDR format.

The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is 10.0.0.1 or 10.0.0.0/24, and an IPv6 address/prefix is ffe1::0020/64 or 2001:ed8:77b5:2:10:10:100:42/64. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

***Note:** The IP address you type must be available and not in the loopback network.*

5. In the **Service Port** field, type 80, or select **HTTP** from the list.
6. From the **HTTP Profile** list, select **http**.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. Click **Finished**.

You now have a virtual server to use as a destination address for application traffic.

Web Hosting Multiple Customers Using Route Domains

Overview: Use of route domains to host multiple web customers on the BIG-IP system

Using the *route domains* feature of the BIG-IP® system, you can provide hosting service for multiple customers by isolating each type of application traffic within a defined address space on the network. This enhances security and dedicates BIG-IP resources to each application.

Using route domains, you can also use duplicate IP addresses on the network, provided that each of the duplicate addresses resides in a separate route domain and is isolated on the network through a separate VLAN. For example, if you are processing traffic for two different customers, you can create two separate route domains. The same node address (such as 10.0.10.1) can reside in each route domain, in the same pool or in different pools, and you can assign a different monitor to each of the two corresponding pool members.

A good example of the use of traffic isolation on a network is an ISP that services multiple customers, where each customer deploys a different application. The first illustration shows two route domain objects on a BIG-IP system, where each route domain corresponds to a separate customer, and thus, resides in its own partition. Within each partition, the ISP created the network objects and local traffic objects required for that customer's application (AppA or AppB).

The sample configuration results in the BIG-IP system segmenting traffic for two different applications into two separate route domains. The routes for each application's traffic cannot cross route domain boundaries because cross-routing restrictions are enabled on the BIG-IP system by default. The second illustration shows the resulting route isolation for AppA and AppB application traffic.

Illustration of sample BIG-IP configuration using route domains

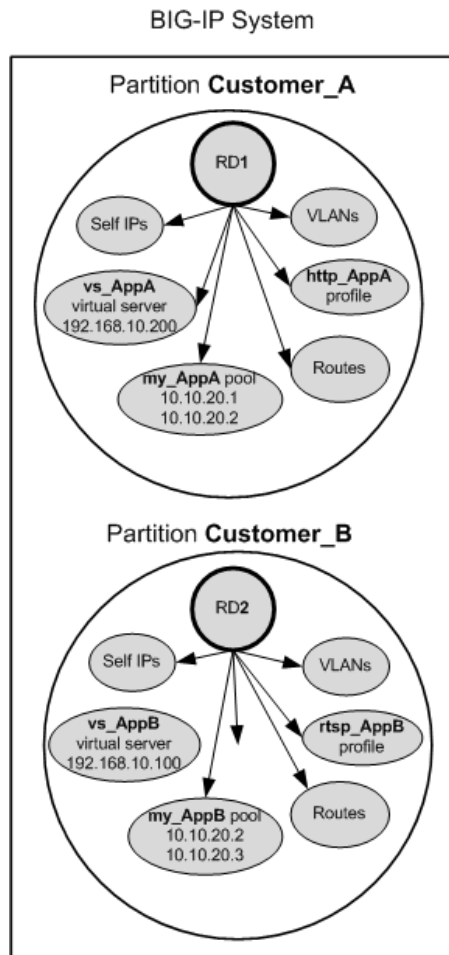


Figure 15: Sample BIG-IP configuration using route domains

Illustration of resulting route domain configuration

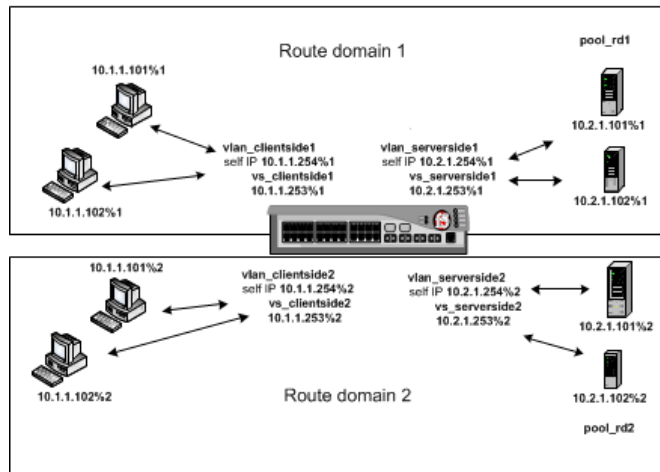


Figure 16: Resulting route domain configuration

Task summary

Perform these tasks to host multiple web customers using route domains.

Task list

Creating an administrative partition

Creating a VLAN with a tagged interface

Creating a self IP address for a default route domain in an administrative partition

Creating a route domain on the BIG-IP system

Creating a load balancing pool

Creating a virtual server

Configuring route advertisement for a virtual address

Adding routes that specify VLAN internal as the resource

Creating an administrative partition

You perform this task to create an administrative partition. An *administrative partition* creates an access control boundary for users and applications.

1. On the Main tab, expand **System** and click **Users**.
The Users List screen opens.
2. On the menu bar, click **Partition List**.
3. Click **Create**.
The New Partition screen opens.
4. In the **Partition Name** field, type a unique name for the partition.
An example of a partition name is `Spanned_VIP`.
5. Type a description of the partition in the **Description** field.

This field is optional.

6. For the **Device Group** setting, choose an action:

Action	Result
Retain the default value.	Choose this option if you want the folder corresponding to this partition to inherit the value of the device group attribute from folder <code>root</code> .
Clear the check box and select the name of a device group.	Choose this option if you do not want the folder corresponding to this partition to inherit the value of the device group attribute from folder <code>root</code> .

7. For the **Traffic Group** setting, choose an action:

Action	Result
Retain the default value.	Choose this option if you want the folder corresponding to this partition to inherit the value of the traffic group attribute from folder <code>root</code> .
Clear the check box and select the name of a traffic group.	Choose this option if you do not want the folder corresponding to this partition to inherit the value of the traffic group attribute from folder <code>root</code> .

8. Click **Finished**.

The new partition appears in the partition list.

Creating a VLAN with a tagged interface

When you create a VLAN with tagged interfaces, each of the specified interfaces can process traffic destined for that VLAN.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. Click **Create**.
The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.
The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged**.
 - c) Click **Add**.
6. If you want the system to verify that the return route to an initial packet is the same VLAN from which the packet originated, select the **Source Check** check box.
7. In the **MTU** field, retain the default number of bytes (**1500**).
8. From the **Configuration** list, select **Advanced**.
9. If you want to base redundant-system failover on VLAN-related events, select the **Fail-safe** check box.
10. From the **Auto Last Hop** list, select a value.

11. From the **CMP Hash** list, select a value.
12. To enable the **DAG Round Robin** setting, select the check box.
13. Click **Finished**.
The screen refreshes, and displays the new VLAN in the list.

The new VLAN appears in the VLAN list.

Creating a self IP address for a default route domain in an administrative partition

Before creating a self IP address, ensure that you have created an internal VLAN and an external VLAN on the BIG-IP system.

Using this procedure, you must create two self IP addresses on the BIG-IP system. One self IP address is associated with the internal VLAN, and the other is associated with the external VLAN. Self IP addresses enable the BIG-IP system and other devices on the network to route application traffic through the associated VLAN.

1. On the Main tab, click **Network > Self IPs**.
2. Click **Create**.
The New Self IP screen opens.
3. In the **IP Address** field, type an IP address.
This IP address should represent the address space of the VLAN that you specify with the **VLAN** setting. Because the route domain that you previously created is the default route domain for the administrative partition, you do not need to append the route domain ID to this IP address.
The system accepts IP addresses in both the IPv4 and IPv6 formats.
4. In the **Netmask** field, type the network mask for the specified IP address.
For example, you can type 255.255.255.0.
5. From the **VLAN/Tunnel** list, select the VLAN to associate with this self IP address.
 - On the internal network, select the internal or high availability VLAN that is associated with an internal interface or trunk.
 - On the external network, select the external VLAN that is associated with an external interface or trunk.
6. Click **Finished**.
The screen refreshes, and displays the new self IP address.

The BIG-IP system has a self IP address that is associated with the internal or external network.

Creating a route domain on the BIG-IP system

Before you create a route domain:

- Ensure that an external and an internal VLAN exist on the BIG-IP® system.
- If you intend to assign a static bandwidth controller policy to the route domain, you must first create the policy. You can do this using the BIG-IP Configuration utility.
- Verify that you have set the current partition on the system to the partition in which you want the route domain to reside.

You can create a route domain on BIG-IP system to segment (isolate) traffic on your network. Route domains are useful for multi-tenant configurations.

1. On the Main tab, click **Network > Route Domains**.
The Route Domain List screen opens.
2. Click **Create**.
The New Route Domain screen opens.
3. In the **Name** field, type a name for the route domain.
This name must be unique within the administrative partition in which the route domain resides.
4. In the **ID** field, type an ID number for the route domain.
This ID must be unique on the BIG-IP system; that is, no other route domain on the system can have this ID.
An example of a route domain ID is 1.
5. In the **Description** field, type a description of the route domain.
For example: This route domain applies to application traffic for Customer A.
6. For the **Strict Isolation** setting, select the **Enabled** check box to restrict traffic in this route domain from crossing into another route domain.
7. For the **Parent Name** setting, retain the default value.
8. For the **VLANs** setting, from the **Available** list, select a VLAN name and move it to the **Members** list.
Select the VLAN that processes the application traffic relevant to this route domain.
Configuring this setting ensures that the BIG-IP system immediately associates any self IP addresses pertaining to the selected VLANs with this route domain.
9. For the **Dynamic Routing Protocols** setting, from the **Available** list, select one or more protocol names and move them to the **Enabled** list.
You can enable any number of listed protocols for this route domain.

Important: You must enable the BGP protocol.

10. From the **Bandwidth Controller** list, select a static bandwidth control policy to enforce a throughput limit on traffic for this route domain.
11. From the **Partition Default Route Domain** list, select either **Another route domain (0) is the Partition Default Route Domain** or **Make this route domain the Partition Default Route Domain**.
This setting does not appear if the current administrative partition is partition `Common`.
When you configure this setting, either route domain 0 or this route domain becomes the default route domain for the current administrative partition.
12. Click **Finished**.
The system displays a list of route domains on the BIG-IP system.

You now have another route domain on the BIG-IP system.

Creating a load balancing pool

You can create a *load balancing pool* (a logical set of devices such as web servers that you group together to receive and process traffic) to efficiently distribute the load on your server resources.

Note: You must create the pool before you create the corresponding virtual server.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.

3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click << to move the monitor to the **Active** list.

***Tip:** Hold the Shift or Ctrl key to select more than one monitor at a time.*

5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.
The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
 - Select **Disabled** to disable priority groups. This is the default option.
 - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.
 - c) In the **Service Port** field, type a port number, or select a service name from the list.
 - d) (Optional) In the **Priority** field, type a priority number.
 - e) Click **Add**.
8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server

A virtual server represents a destination IP address for application traffic.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address in CIDR format.
The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is 10.0.0.1 or 10.0.0.0/24, and an IPv6 address/prefix is fe1::0020/64 or 2001:ed8:77b5:2:10:10:100:42/64. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

***Note:** The IP address you type must be available and not in the loopback network.*

5. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
6. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.

Configuring route advertisement for a virtual address

Before configuring route advertisement on a virtual address, verify that you have enabled one or more dynamic routing protocols on the route domain pertaining to this virtual address. Also verify that you have configured the relevant dynamic routing protocols for route redistribution.

Perform this task to advertise a route for this virtual address to other routers on your network.

Important: *This task pertains only to configurations for which you have enabled dynamic routing protocols on the relevant route domain. If you have not enabled dynamic routing protocols on the relevant route domain, you can skip this task.*

1. On the Main tab, click **Local Traffic > Virtual Servers > Virtual Address List**.
The Virtual Address List screen opens.
2. In the Name column, click the virtual address for which you want to advertise a route.
This displays the properties of that virtual address.
3. Verify that the **ARP** field is selected.
4. From the **Advertise Route** list, choose one of these options:

Option	Description
When any virtual server is available	Specifies that the system advertises a route for this virtual IP address whenever any virtual server associated with this virtual IP address is available.
When all virtual servers(s) are available	Specifies that the system advertises a route for this virtual IP address whenever all virtual servers associated with this virtual IP address is available.
Always	Specifies that the system always advertises a route for this virtual IP address.
5. For the **Route Advertisement** setting, select the box.
This makes it possible for the BIG-IP system to advertise this virtual IP address when you have enabled any dynamic routing protocols.
6. Click **Update**.
7. Repeat this task for each virtual address for which you want to advertise a route.

The BIG-IP system advertises a route for this virtual address to other routers when one or more dynamic routing protocols are enabled and are configured for route redistribution.

Adding routes that specify VLAN internal as the resource

Ensure that you set the current administrative partition to the partition in which you want a specific customer's configuration to reside.

You must add a route for each destination IP address pertaining to the route domain. A destination address in this case is typically a node address for a pool member.

1. On the Main tab, click **Network > Routes**.
2. Click **Add**.
The New Route screen opens.
3. In the **Name** field, type a unique user name.

This name can be any combination of alphanumeric characters, including an IP address.

4. In the **Destination** field, type either the destination IP address for the route, or IP address 0.0.0.0 for the default route.

This address can represent either a host or a network. Also, if you are using the route domains and the relevant route domain is the partition default route domain, you do not need to append a route domain ID to this address.

5. In the **Netmask** field, type the network mask for the destination IP address.
6. From the **Resource** list, select **Use VLAN/Tunnel**.

A VLAN represents the VLAN through which the packets flow to reach the specified destination.

7. From the **VLAN** list, select **Internal**.
8. Click **Finished**.

The BIG-IP system now includes routes to the nodes in the load balancing pool for a specific route domain.

Implementing the Link Layer Discovery Protocol

Overview: Implementing Link Layer Discovery Protocol

The BIG-IP® system supports Link Layer Discovery Protocol (LLDP). LLDP is a Layer 2 industry-standard protocol (IEEE 802.1AB) that gives a network device such as the BIG-IP system the ability to advertise its identity and capabilities to multi-vendor neighbor devices on a network. The protocol also enables a network device to receive information from neighbor devices.

LLDP transmits device information in the form of LLDP messages known as LLDP Packet Data Units (LLDPDUs).

In general, this protocol:

- Advertises connectivity and management information about the local BIG-IP device to neighbor devices on the same IEEE 802 LAN.
- Receives network management information from neighbor devices on the same IEEE 802 LAN.
- Operates with all IEEE 802 access protocols and network media.

Using the BIG-IP Configuration utility or `tmsh`, you can use this particular implementation to configure BIG-IP system interfaces to transmit LLDPDUs to neighbor devices. More specifically, you can:

- Specify the exact content of LLDPDUs that a BIG-IP system interface transmits to a neighbor device. You specify this content by configuring the **LLDP Attributes** setting on each individual interface.
- Globally specify the frequencies of various message transmittal properties, and specify the number of neighbors from which interfaces can receive messages. These properties apply to all interfaces on the BIG-IP system.

The following illustration shows a BIG-IP system that transmits LLDP messages to three neighbor devices: another BIG-IP system, an external switch, and an external router. Note that LLDP is enabled on all of the devices.




Figure 17: The BIG-IP system and LLDP transmittal

Task summary

Perform these tasks to implement Link Layer Discovery Protocol (LLDP) on selected BIG-IP system interfaces.

Task list

Configuring global LLDP properties

Configuring LLDP settings for an individual interface

Configuring global LLDP properties

You can configure a set of general LLDP properties that apply to all interfaces on the BIG-IP system. These settings pertain to LLDP message transmission frequencies and the maximum number of neighbors to which each interface can send LLDP messages.

Note: Although you use this procedure to globally enable the LLDP feature on the BIG-IP system, you can also disable LLDP for any individual interface. You do this by configuring the specific properties of that interface.

1. On the Main tab, click **Network > Interfaces > LLDP > General**.
This displays the general LLDP properties that you can configure on the system.
2. From the **LLDP** list, select **Enabled**.
3. For the remainder of the settings, retain or change the default values.
4. Click the **Update** button.

This task activates support for the LLDP protocol on the BIG-IP system, and configures the system to transmit LLDPDUs according to the specified frequencies.

Configuring LLDP settings for an individual interface

You can use this procedure to configure the settings for an individual interface on the BIG-IP system.

1. On the Main tab, click **Network > Interfaces > Interface List**.
The Interface List screen displays the list of interfaces on the system.
2. In the Name column, click an interface number.
This displays the properties of the interface.
3. For the **State** setting, verify that the interface is set to **Enabled**.
4. For the **LLDP** setting, verify that the property is set to **Transmit Only**.
5. For the **LLDP Attributes** setting, verify that the list of attributes in the **Send** field includes all Time Length Values (TLVs) that you want the BIG-IP system interface to send to neighbor devices.
6. Click the **Update** button.

After you perform this task, the interface is configured to send the specified LLDP information to neighbor devices.

Implementation result

This implementation results in this LLDP configuration:

- Support for the LLDP protocol is enabled on the BIG-IP system.
- For all BIG-IP system interfaces, the BIG-IP system attempts to transmit LLDPDUs to neighbor devices every 30 seconds, with a minimum delay between transmissions of 2 seconds.
- The maximum number of neighbors to which each BIG-IP system interface can send LLDPDUs is 10.
- Every BIG-IP system interface can send LLDPDUs to its neighbors.
- No BIG-IP system interface can receive LLDPDUs from its neighbors.

In addition, the content of the LLDPDUs that each BIG-IP system interface sends to its neighbors contains this information:

- Chassis ID
- Port ID
- Time-to-Live value
- Port description
- System name
- System description
- System capabilities
- Port VLAN ID
- Port and protocol VLAN ID
- VLAN name
- Protocol identity
- MAC/PHY config status
- Link aggregation
- Max frame size
- Product model

Configuring Remote High-Speed Logging

Overview: Configuring high-speed remote logging

You can configure the BIG-IP® system to log information about BIG-IP system processes and send the log messages to remote high-speed log servers. You can filter the data that the system logs based on alert-level and source.

This illustration shows the association of the configuration objects for remote high-speed logging of BIG-IP system processes.




Figure 18: Association of remote high-speed logging configuration objects

Task summary

Perform these tasks to configure BIG-IP® system logging.

Note: Enabling remote high-speed logging impacts BIG-IP system performance.

Creating a pool of remote logging servers

Creating a remote high-speed log destination

Creating a formatted remote high-speed log destination

Creating a publisher

Creating a logging filter

Disabling system logging

Troubleshooting logs that contain unexpected messages

About the configuration objects of high-speed remote logging

When configuring remote high-speed logging of BIG-IP system processes, it is helpful to understand the objects you need to create and why, as described here:

Object	Reason	Applies to
Pool of remote log servers	Create a pool of remote log servers to which the BIG-IP® system can send log messages.	Creating a pool of remote logging servers.
Destination (unformatted)	Create a log destination of Remote High-Speed Log type that specifies a pool of remote log servers.	Creating a remote high-speed log destination.
Destination (formatted)	If your remote log servers are the ArcSight, Splunk, IPFIX, or Remote Syslog type, create an additional log destination to format the logs in the required format and forward the logs to a remote high-speed log destination.	Creating a formatted remote high-speed log destination.
Publisher	Create a log publisher to send logs to a set of specified log destinations.	Creating a publisher.
Filter	Create a log filter to define the messages to be included in the BIG-IP system logs and associate a log publisher with the filter.	Creating a logging filter.

Creating a pool of remote logging servers

Before creating a pool of log servers, gather the IP addresses of the servers that you want to include in the pool. Ensure that the remote log servers are configured to listen to and receive log messages from the BIG-IP® system.

Create a pool of remote log servers to which the BIG-IP system can send log messages.

1. On the Main tab, click the applicable path.

- **DNS > Delivery > Load Balancing > Pools**
- **Local Traffic > Pools**

The Pool List screen opens.

2. Click **Create**.

The New Pool screen opens.

3. In the **Name** field, type a unique name for the pool.

4. Using the **New Members** setting, add the IP address for each remote logging server that you want to include in the pool:

- a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
- b) Type a service number in the **Service Port** field, or select a service name from the list.

***Note:** Typical remote logging servers require port 514.*

c) Click **Add**.

5. Click **Finished**.

Creating a remote high-speed log destination

Before creating a remote high-speed log destination, ensure that at least one pool of remote log servers exists on the BIG-IP® system.

Create a log destination of the **Remote High-Speed Log** type to specify that log messages are sent to a pool of remote log servers.

1. On the Main tab, click **System > Logs > Configuration > Log Destinations**.
The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **Remote High-Speed Log**.

***Important:** If you use log servers such as Remote Syslog, Splunk, or ArcSight, which require data be sent to the servers in a specific format, you must create an additional log destination of the required type, and associate it with a log destination of the **Remote High-Speed Log** type. With this configuration, the BIG-IP system can send data to the servers in the required format.*

The BIG-IP system is configured to send an unformatted string of text to the log servers.

5. From the **Pool Name** list, select the pool of remote log servers to which you want the BIG-IP system to send log messages.
6. From the **Protocol** list, select the protocol used by the high-speed logging pool members.
7. Click **Finished**.

Creating a formatted remote high-speed log destination

Ensure that at least one remote high-speed log destination exists on the BIG-IP® system.

Create a formatted logging destination to specify that log messages are sent to a pool of remote log servers, such as Remote Syslog, Splunk, or ArcSight servers.

1. On the Main tab, click **System > Logs > Configuration > Log Destinations**.
The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select a formatted logging destination, such as **IPFIX**, **Remote Syslog**, **Splunk**, or **ArcSight**.

***Important:** ArcSight formatting is only available for logs coming from Advanced Firewall Manager™ (AFM™), Application Security Manager™ (ASM™), and the Secure Web Gateway component of Access Policy Manager® (APM®). IPFIX is not available for Secure Web Gateway. Remote Syslog formatting*

is the only type supported for logs coming from APM. The Splunk format is a predefined format of key value pairs.

The BIG-IP system is configured to send a formatted string of text to the log servers.

5. If you selected **Remote Syslog**, from the **Syslog Format** list, select a format for the logs, and then from the **High-Speed Log Destination** list, select the destination that points to a pool of remote Syslog servers to which you want the BIG-IP system to send log messages.

Important: For logs coming from Access Policy Manager® (APM®), only the BSD Syslog format is supported.

6. If you selected **Splunk** or **IPFIX**, from the **Forward To** list, select the destination that points to a pool of high-speed log servers to which you want the BIG-IP system to send log messages.
7. Click **Finished**.

Creating a publisher

Ensure that at least one destination associated with a pool of remote log servers exists on the BIG-IP® system.

Create a publisher to specify where the BIG-IP system sends log messages for specific resources.

1. On the Main tab, click **System > Logs > Configuration > Log Publishers**.
The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. For the **Destinations** setting, select a destination from the **Available** list, and click << to move the destination to the **Selected** list.

Note: If you are using a formatted destination, select the destination that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.

5. Click **Finished**.

Creating a logging filter

Ensure that at least one log publisher is configured on the BIG-IP® system.

Create a custom log filter to specify the system log messages that you want to publish to a particular log.

1. On the Main tab, click **System > Logs > Configuration > Log Filters**.
The Log Filters screen opens.
2. In the **Name** field, type a unique, identifiable name for this filter.
3. From the **Severity** list, select the level of alerts that you want the system to use for this filter.

Note: The severity level that you select includes all of the severity levels that display above your selection in the list. For example, if you select **Emergency**, the system publishes only emergency messages to the log. If you select **Critical**, the system publishes critical, alert, and emergency-level messages in the log.

4. From the **Source** list, select the system processes from which messages will be sent to the log.

5. In the **Message ID** field, type the first eight hex-digits of the specific message ID that you want the system to include in the log. Use this field when you want a log to contain only each instance of one specific log message.

***Note:** BIG-IP system log messages contain message ID strings in the format: xxxxxxxx:x:. For example, in this log message: Oct 31 11:06:27 olgavmmgmt notice mcpd[5641]: 01070410:5: Removed subscription with subscriber id lind , the message ID string is: 01070410:5:. You enter only the first eight hex-digits: 01070410.*

6. From the **Log Publisher** list, select the publisher that includes the destinations to which you want to send log messages.
7. Click **Finished**.

Disabling system logging

When you no longer want the BIG-IP® system to log information about its internal systems, you can delete the log filter that you created. For example, when mitigating a DoS attack, if you created a log filter that includes only one specific message in the log, you can delete that log filter once you handle the attack.

1. On the Main tab, click **System > Logs > Configuration > Log Filters**.
The Log Filters screen opens.
2. Select the check box next to the name of the log filter that you want to delete. Click **Delete**, and then click **Delete** again.

Troubleshooting logs that contain unexpected messages

If you configured a filter to send all instances of a specific message ID to your remote logging servers and this message ID is still displaying in the local log in the BIG-IP system, you can disable legacy log message processing in order to display instances of this message ID only on the remote logging servers.

***Important:** When you create a filter that disables legacy log message processing, the legacy logs are completely disabled. Therefore, you must also create a filter for every source from which you want log messages to be sent to the pool of remote log servers.*

1. On the Main tab, click **System > Logs > Configuration > Log Filters**.
The Log Filters screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this filter.
4. From the **Severity** list, select **Debug**.
5. From the **Source** list, select **All**.
6. From the **Log Publisher** list, select **None**.
7. Click **Finished**.

Deploying Route Domains within a vCMP Guest

Overview: Deploying Route Domains within a vCMP Guest

With a vCMP® system, you typically create guests as a way to segment different types of application traffic. An alternative way to segment application traffic is to configure a feature known as route domains, within a single guest.

A *route domain* is a configuration object that isolates network traffic for a particular application on the network. Using route domains, you can assign the same IP address or subnet to multiple nodes on a network, provided that each instance of the IP address resides in a separate route domain.

The configuration described here manages traffic for three separate customers, where each customer has its own route domain to process and ensure isolation for a different type of application traffic. By using route domains within a guest, you can minimize the total number of guests you must create to manage customer traffic.

This illustration shows a redundant system configuration in which a single guest uses route domains for three separate customers.

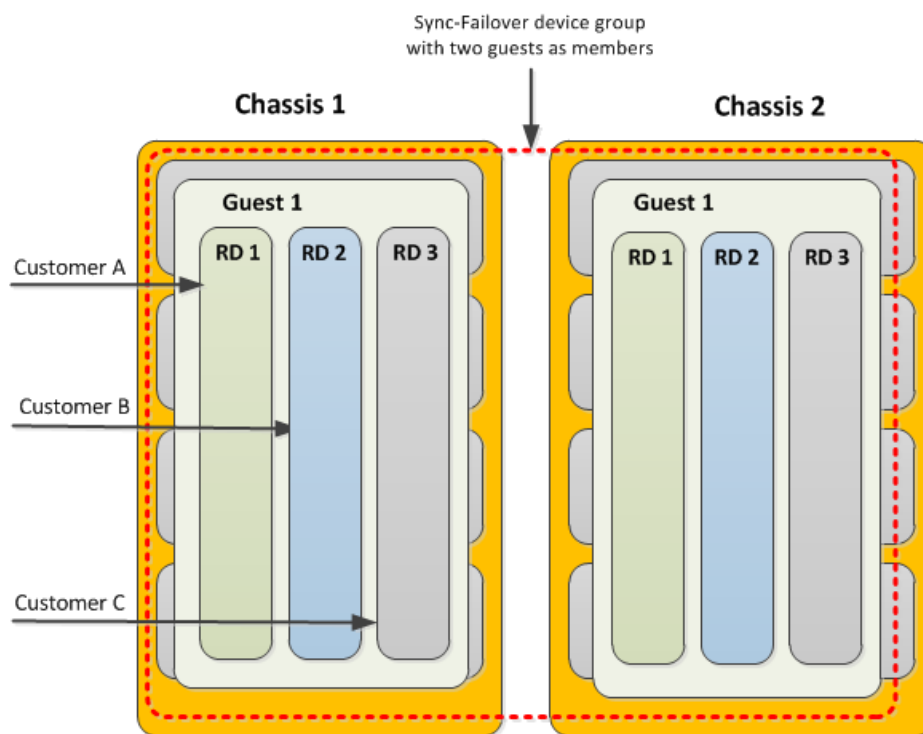


Figure 19: Route domains within a guest

Each route domain contains all of the network objects necessary for processing a specific type of traffic and ensuring failover to the other guest in the event that the system becomes unavailable. These network objects consist of floating self IP addresses associated with host-based VLANs, floating virtual IP addresses, and pool members defined on the guest. The floating addresses are further associated with an active traffic group on one instance of the guest and a standby traffic group on the other instance of the guest.

Prerequisite configuration tasks

Before you begin deploying route domains within a vCMP guest, ensure that you have configured the following on each chassis:

- The initial setup of the BIG-IP® base network on the VIPRION® chassis, prior to provisioning the system for vCMP®. This setup typically includes VLANs for the external and internal networks, as well as an additional internal VLAN for failover communications between device group members.
- The initial setup of the vCMP host. This includes provisioning the system for vCMP and creating guests, with the host VLANs published to the guest.
- Non-floating self IP addresses on the guest. These addresses are associated with the host-based external, internal, and high availability VLANs.
- A Sync-Failover device group consisting of two guests as its members (one guest per chassis). The guests on the two chassis should be identical with respect to memory, CPU, and slot allocation.

About VLAN and BIG-IP address configuration

When you initially configured the BIG-IP® base network on the VIPRION® system, you created three VLANs: two for the internal and external networks, and one for high availability communications, and you created their associated non-floating self IP addresses. Now you are ready to create additional VLANs and self IP addresses for processing each customer's application traffic. On a system provisioned for vCMP®, all VLANs reside on the vCMP host, while all self IP addresses (floating and non-floating) reside on the guest.

Illustration of VLAN and BIG-IP address configuration

This illustration shows the relationship of the VLANs on the host to the IP addresses within each route domain on the guest. Note that in our example, all three customers use the same self IP and virtual IP addresses but with unique route domain IDs. Also note that except for the non-floating self IP addresses in partition `Common`, the entire configuration is duplicated on the peer guest (not shown).

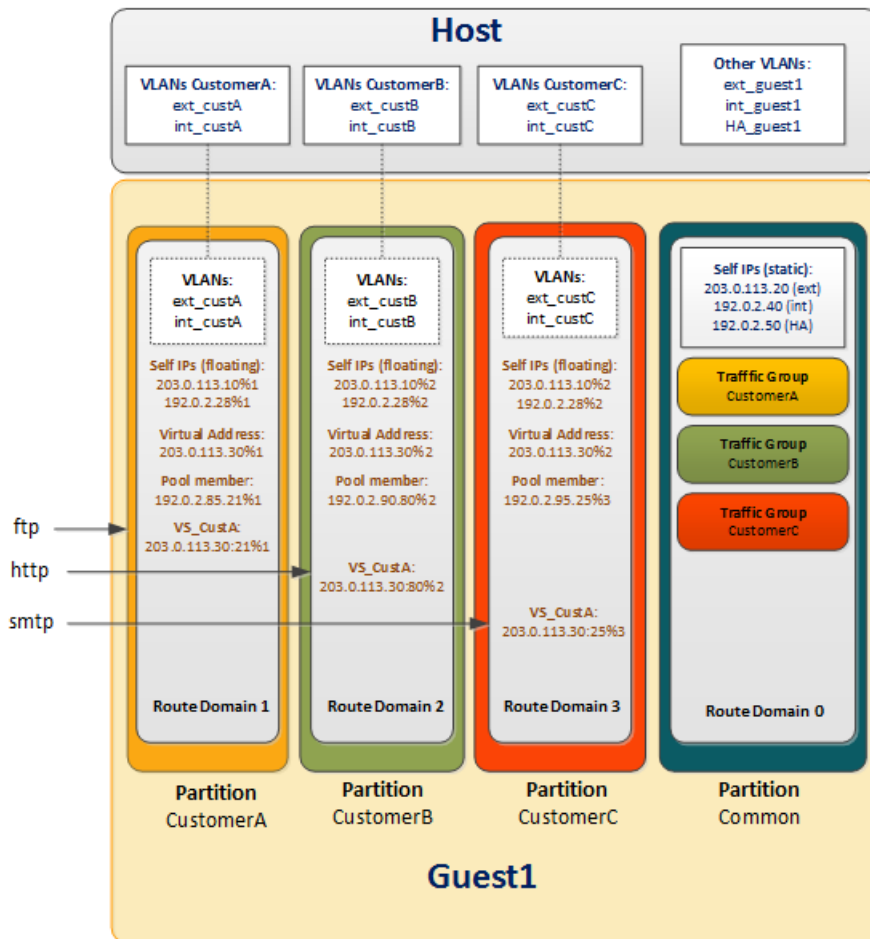


Figure 20: VLANs and BIG-IP addresses in a vCMP route domain configuration

In this illustration:

Blue text

Objects created by host administrator.

Black text

Objects created by guest administrator.

Brown text

Objects created by customer administrator.

Tasks for the host administrator

To set up a route domain configuration, the vCMP® host administrator needs to create VLANs for use by each customer.

On the host, for our sample configuration with three customers, you create a separate set of uniquely-tagged internal and external VLANs for each customer. You will therefore create at least six VLANs on the host (two per customer) that, when combined with the three existing VLANs, bring the total number of VLANs

on the host to nine. At this point, all VLANs reside in partition `Common`. Then you assign all nine host-based VLANs to the guest. This allows the guest to use those VLANs to process customer traffic.

To summarize, the objects that a host administrator creates are:

- VLANs created during base VIPRION® configuration
- Customer-specific VLANs for use by guest route domains

Creating customer VLANs on the vCMP host

You create additional VLANs on the vCMP® host that you then assign to the guest. Then, when logged in to the guest, you can selectively distribute the VLANs to different route domains within the guest. Each route domain corresponds to a different customer.

Note: You must create this same set of VLANs on the host of each vCMP system in the configuration.

Important: Ensure that the tags for all VLANs that you create are unique.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. Click **Create**.
The New VLAN screen opens.
3. In the **Name** field, type the name of the first VLAN.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.
The VLAN tag identifies the application traffic for the associated VLAN.

Important: Each VLAN tag that you specify in this field must be unique on the vCMP system.

5. If you want to use Q-in-Q (double) tagging, use the **Customer Tag** setting to perform the following two steps. If you do not see the **Customer Tag** setting, your hardware platform does not support Q-in-Q tagging and you can skip this step.
 - a) From the **Customer Tag** list, select **Specify**.
 - b) Type a numeric tag, from 1-4094, for the VLAN.

The customer tag specifies the inner tag of any frame passing through the VLAN.

6. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) If you specified a numeric value for the **Customer Tag** setting and from the **Tagging** list you selected **Tagged**, then from the **Tag Mode** list, select a value.
 - d) Click **Add**.
7. Click **Repeat** and repeat these steps to create additional VLANs.

After you complete this task on the vCMP host, VLAN objects exist on the system that you can assign to the guest.

Assigning VLANs to the vCMP guest

Before you perform this task, verify that you have created a vCMP® guest on the system. The guest should have an external, an internal, and a high availability VLAN assigned to the guest. Also verify that the guest is in the Configured or Provisioned state.

You assign host-based VLANs to a guest so that the guest can use those VLANs to process customer traffic. For the sample configuration, you assign all six customer-specific VLANs to the guest.

Important: *You must be logged in to the vCMP host to perform this task.*

1. On the Main tab, click **vCMP > Guest List**.
This displays a list of guests on the system.
2. In the Name column, click the name of the guest that you want to modify.
This displays the configured properties of the guest.
3. For the **VLAN List** setting, select all customer-specific VLANs from the **Available** list, and use the Move button to move the VLAN names to the **Selected** list.
4. Click **Update**.

After you perform this task, the guest can use the selected VLANs to process customer traffic.

Tasks for the guest administrator

You perform the remainder of the configuration on the vCMP® guest. First, you create an administrative partition for each customer. Then from within each customer's partition, you move the relevant customer-specific VLANs from **Common** to that partition.

Once each customer's VLANs have been moved to the relevant partition, you can create a route domain and a traffic group for each customer.

To summarize, the objects that a guest-wide administrator creates are:

- Administrative partitions
- Instances of host-based customer VLANs
- Route domains
- Traffic groups for failover

Creating an administrative partition for each customer

You perform this task to create administrative partitions within a vCMP® guest. An *administrative partition* creates an access control boundary for users and applications. Using this task, you create a separate administrative partition for each customer associated with the guest. Each administrative partition will contain a route domain that contains the Layer 3 objects associated with the relevant customer.

Important: *Before performing this task, log in to the guest using the guest IP address.*

1. On the Main tab, expand **System** and click **Users**.
The Users List screen opens.
2. On the menu bar, click **Partition List**.

3. Click **Create**.
The New Partition screen opens.
4. In the **Partition Name** field, type a unique name for the partition.
An example of a partition name is `CustomerA_partition`.
5. Type a description of the partition in the **Description** field.
This field is optional.
6. For the **Device Group** setting, ensure that the Sync-Failover device group containing this vCMP guest is selected.
7. For the **Traffic Group** setting, retain the default value, which is the floating traffic group `traffic-group-1`.

Note: You will change this value later in the route domain implementation process.

8. Click **Finished**.
9. Repeat these steps to create additional administrative partitions.

After you perform this task, the new partitions appear in the list of partitions on the guest, as well as in the **Partition** list in the upper right corner of every BIG-IP® Configuration utility screen.

About moving host-based VLANs to a customer partition

As guest administrator, you must switch to a specific customer administrative partition and move a customer-related VLAN from `Common` to that partition. You effectively move each VLAN by deleting the VLAN from `Common` and re-creating the VLAN in the relevant customer's partition.

For example, if you create route domain 1 in partition A for Customer A's traffic, you will then move VLANs `ext_custA` and `int_custA` from `Common` to partition A. This associates the VLAN with the new partition instead of partition `Common`, without changing the host's control of the VLAN's underlying Layer 2 (and lower) network resources.

Note: Although you are logged in to the guest and you move the VLANs from `Common` to the relevant partition, the VLANs continue to reside on the host.

Deleting VLANs in partition Common from within the guest

Before you perform this task, ensure that, on the vCMP® host, you have created all customer-relevant VLANs for this implementation and assigned all of them to the vCMP guest. Also, ensure that you are logged in to the guest, using the guest IP address.

You use this task to delete a VLAN in partition `Common` on a guest so that you can re-create the VLAN in a customer partition.

Note: You must be logged in to the guest to perform this task.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and ensure that partition `Common` is selected.
3. In the Name column, locate the relevant VLAN name.
4. In the Tag column, note the numeric ID.

You will specify this ID when you re-create this VLAN in a customer partition.

An example of a VLAN ID in the Tag column is **4094**.

5. If the VLAN has a customer tag (optional), then in the Customer Tag column, note the numeric ID.
You will specify this ID when you re-create this VLAN in a customer partition.
6. To the left of the VLAN name, select the check box and click **Delete**.
The system prompts you to confirm the delete action.
7. Click **Delete**.

After you perform this task, the VLAN in partition `Common` on the guest is deleted.

Re-creating VLANs in each administrative partition

Before you perform this task, ensure that you are logged in to the guest, using the guest IP address.

You perform this task to re-create a VLAN in a specific customer partition. You re-create a VLAN in a customer partition when you want to set up a route domain configuration within the guest. The VLAN you are re-creating is one that you previously created on the host in partition `Common` and then deleted from partition `Common` when you later logged in to the guest. Each route domain that you create in a partition requires you to assign one or more VLANs to that route domain, and those VLANs must reside in the same partition as the route domain.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and select the customer-specific administrative partition.

If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.

An example of a selected partition is `CustomerA_partition`.
Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.
3. Click **Create**.
The New VLAN screen opens.
4. Type a name for the VLAN.
You can specify the same name as the VLAN that you deleted from partition `Common` or you can type a unique name.
5. For the **Tag** field and the optional **Customer Tag** field, type the same ID that was previously assigned to the VLAN that you deleted from partition `Common`.

Important: For example, if VLAN `external_cust_A` on the host in partition `Common` has a VLAN tag of `4094`, then the VLAN that you re-create within the guest in partition `CustomerA_partition` must also have the tag `4094`.

6. Retain the values for all other settings as configured.
7. Click **Finished**.
This prompts you with the question: The VLAN has no interface, do you want to continue?
8. Click **OK**.

After you perform this task, the VLAN is associated with the customer's administrative partition.

Creating a route domain for each administrative partition

With this task, you can create a route domain and associate it with the administrative partition pertaining to a particular customer.

Important: Before performing this task, ensure that you are logged in to the guest, using the guest IP address.

1. On the Main tab, click **Network > Route Domains**.
The Route Domain List screen opens.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and select the customer-specific administrative partition.
If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.
An example of a selected partition is `CustomerA_partition`.
Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.
3. Click **Create**.
The New Route Domain screen opens.
4. In the **ID** field, type an ID number for the route domain.
This ID must be unique on the BIG-IP system; that is, no other route domain on the system can have this ID.
An example of a route domain ID is 1.
5. In the **Description** field, type a description of the route domain.
For example: This route domain applies to application traffic for Customer A.
6. For the **Strict Isolation** setting, select the **Enabled** check box to restrict traffic in this route domain from crossing into another route domain.
7. For the **Parent Name** setting, retain the default value.
8. For the **VLANs** setting, from the **Available** list, select a VLAN name and move it to the **Members** list.
The VLANs you select should be those pertaining to the customer for which you are creating this route domain.
For example, you can select VLANs `ext_custA` and `int_custA`.
9. For the **Dynamic Routing Protocols** setting, from the **Available** list, select one or more protocol names and move them to the **Enabled** list.
You can enable any number of listed protocols for this route domain.

Important: You must enable the BGP protocol.

10. From the **Bandwidth Controller** list, select a static bandwidth control policy to enforce a throughput limit on traffic for this route domain.
11. From the **Partition Default Route Domain** list, select **Make this route domain the Partition Default Route Domain**.
This value designates this route domain to be the default route domain for the current administrative partition.

Note: The **Partition Default Route Domain** setting appears only when the current partition is set to a partition other than `Common`.

After choosing this value, you are not required to append the route domain ID to any self IP or virtual IP address that you create later for this route domain. Instead, the BIG-IP system automatically associates an IP address with the default route domain in the partition, as long as you set this partition to be the current partition when you create the address.

12. Click **Finished.**

The system displays a list of route domains on the BIG-IP system, including the new route domain.

13. Repeat the process of creating a route domain for another customer for which you want to segment traffic, associating the relevant VLANs in the process.

After you perform this task repeatedly, you should have three separate route domains with unique route domain IDs, and each route domain should be associated with unique internal and external VLANs that pertain to a specific customer. Also, each route domain should be designated as the default route domain for its associated administrative partition.

Creating an empty traffic group for each customer

Before you perform this task, confirm that the current partition is set to `Common`.

Perform this task when you want to create a separate floating traffic group for each customer's traffic. You should perform this task on the guest on which you want the traffic groups to be active.

Important: *This procedure creates a traffic group but does not automatically associate the traffic group with failover objects such as self IP and virtual IP addresses. You associate a traffic group with specific failover objects when you create or modify each object.*

Note: *All traffic groups on the system must reside in partition `Common`.*

1. On the Main tab, click **Device Management > Traffic Groups**.
2. On the Traffic Groups screen, click **Create**.
3. In the **Name** field, type a name for the traffic group.
For example, you can name the traffic group `tg-customerA`.
4. In the **Description** field, type a description for the new traffic group.
For example, you can type `This traffic group manages failover for Customer B traffic`.
5. In the **MAC Masquerade Address** field, type a MAC masquerade address.
When you specify a MAC masquerade address, you reduce the risk of dropped connections when failover occurs. This setting is optional.
6. From the **Failover Method** list, select **HA Order**.
7. For the **Failover Order** setting, in the **Available** box, select the peer guest name, and using the Move button, move the name to the **Enabled** box.
This setting is optional. Only devices that are members of the relevant Sync-Failover device group are available for inclusion in the ordered list.
8. Click **Finished**.
9. Repeat these steps to create a traffic group for each additional customer.

You now have floating traffic groups with no members.

After you perform this task, you can associate each customer's traffic group with the relevant failover objects (self IP addresses, virtual servers, and so on).

Assigning a traffic group to each administrative partition

Before you perform this task, verify that you have created a unique administration partition for each customer.

You assign an individual traffic group to each customer partition to ensure that when failover occurs, the floating IP addresses defined in the named traffic group fail over to the peer guest and remain associated with the correct administrative partition.

1. On the Main tab, expand **System** and click **Users**.
The Users List screen opens.
2. On the menu bar, click **Partition List**.
3. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and ensure that partition `Common` is selected.
4. In the Name column, click a customer partition name.
5. For the **Traffic Group** setting, clear the check box labeled **Inherit traffic group from root folder** and from the list, select the name of a traffic group.
6. Click **Update**.
7. Repeat these steps to assign a traffic group to each of the other customer partitions.

After performing this task, each customer's floating IP addresses will remain associated with the correct administrative partition when failover occurs.

Tasks for each customer administrator

After the vCMP[®] host and guest administrators have set up the VLANs, partitions, route domains, and traffic groups, the customer administrator logging into the guest creates the necessary IP addresses for the application: internal and external floating self IP addresses, server pool member addresses, and a destination virtual server address. The customer administrator also modifies the floating virtual IP address (associated with the virtual server) to assign the relevant traffic group.

Creating floating self IP addresses

As a customer administrator, you create two floating self IP addresses for each customer route domain, one address for the internal network and one address for the external network.

For example, for customer A's internal and external networks, you create two self IP addresses to which you assign VLANs `int_custA` and `ext_custA` respectively, which have both been previously assigned to route domain **1**. Similarly, for customer B, you create self IP addresses and assign VLANs `int_custB` and `ext_custB` respectively, which have both been previously assigned to route domain **2**, and so on.

You also add the self IP addresses as members of a customer-related floating traffic group. This causes the self IP addresses to become floating addresses.

Important: Before performing this task, ensure that you are logged in to the guest, using the guest IP address.

1. On the Main tab, click **Network > Self IPs**.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and select the customer-specific administrative partition.

If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.

An example of a selected partition is `CustomerA_partition`.

Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.

3. Click **Create.**

The New Self IP screen opens.

4. In the **IP Address field, type an IP address.**

This IP address should represent the address space of a specific VLAN. Because the route domain for the VLAN that you will associate with this self IP address is the default route domain for the current administrative partition, you are not required to append the relevant route domain ID to this IP address.

The system accepts IP addresses in both the IPv4 and IPv6 formats.

5. In the **Netmask field, type the full network mask for the specified IP address.**

6. From the **VLAN/Tunnel list, select the VLANs that you want to associate with this self IP address.**

The VLANs you select are those that you moved from partition `Common` to the current administrative partition.

7. From the **Port Lockdown list, select a value.**

8. From the **Traffic Group list, select the floating traffic group for which you want this self IP address to be a member.**

Selecting a floating traffic group automatically causes the self IP address to be a floating address.

For example, you can select a traffic group named `tg-CustomerA`.

9. Click **Finished.**

The screen refreshes, and displays the new self IP address.

10. Repeat this task for each floating self IP address that you need to create.

After performing this task repeatedly, each floating traffic group on the guest should contain self IP addresses that are associated with the internal and external VLANs for each customer.

Creating a pool

You can create a pool of servers that you can group together to receive and process traffic. Once the pool is created, you can associate the pool with a virtual server.

Important: Before performing this task, ensure that you are logged in to the guest, using the guest IP address.

1. On the Main tab, click **Local Traffic > Pools.**

The Pool List screen opens.

2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition list and select the customer-specific administrative partition.**

If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.

An example of a selected partition is `CustomerA_partition`.

Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.

3. Click **Create.**

The New Pool screen opens.

4. In the **Name field, type a unique name for the pool.**

5. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.

***Note:** Because the route domain for this pool is the default route domain for the current administrative partition, you are not required to append the relevant route domain ID to this IP address.*

- c) In the **Service Port** field, type a port number, or select a service name from the list.
 - d) (Optional) In the **Priority** field, type a priority number.
 - e) Click **Add**.
6. Click **Finished**.
7. Repeat these steps to create each customer's pool.

After performing this task, the new pool appears in the Pools list.

Creating a virtual server

The purpose of this task is to create virtual servers that represent destination IP addresses for different types of application traffic.

***Important:** Before performing this task, ensure that you are logged in to the guest, using the guest IP address.*

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and select the customer-specific administrative partition.
If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.
An example of a selected partition is `CustomerA_partition`.
Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.
3. Click the **Create** button.
The New Virtual Server screen opens.
4. In the **Name** field, type a unique name for the virtual server.
5. In the **Destination Address** field, type the IP address in CIDR format.
The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `10.0.0.1` or `10.0.0.0/24`, and an IPv6 address/prefix is `ffe1::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

***Note:** The IP address you type must be available and not in the loopback network.*

6. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
 7. Configure all other settings as needed.
 8. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
 9. Click **Finished**.

Modifying a virtual IP address

The purpose of this task is to convert a non-floating virtual IP address to a floating address, by adding the address as a member of a traffic group.

Note: The BIG-IP® system automatically creates a virtual address when you create a virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers > Virtual Address List**.
The Virtual Address List screen opens.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and select the customer-specific administrative partition.
If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.
An example of a selected partition is `CustomerA_partition`.
Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.
3. In the Name column, click the virtual address that you want to assign to the traffic group.
This displays the properties of that virtual address.
4. From the **Traffic Group** list, select the traffic group for which you want this virtual address to be a member.
Selecting a floating traffic group automatically causes the virtual IP address to be a floating address.
For example, you can select a floating traffic group named `tg-CustomerA`.
5. Click **Update**.
6. Repeat these steps for each customer's virtual address.

Each floating virtual IP address for a route domain is now a member of the relevant traffic group.

Implementation results

After you have completed all tasks in this implementation, you have a Device Service Clustering (DSC®) configuration in which one of the guests on each vCMP® system contains three administrative partitions, each of which contains a default route domain with Layer 3 IP addresses pertaining to a specific type of traffic.

With this configuration, the BIG-IP® system can process network traffic for three separate customers. Because each set of addresses for a traffic type is contained in a route domain, all three sets of customer IP addresses can be identical except for the unique route domain ID that is implicitly part of each address.

Furthermore, each route domain is associated with a unique floating traffic group that can fail over to the other guest if the vCMP® system becomes unavailable for any reason.

Using Link Aggregation with Tagged VLANs for a One-network Topology

Overview: Configuring link aggregation using tagged VLANs on one network

You can use the BIG-IP[®] system in an aggregated two-interface load balancing topology. *Link aggregation* is the process of combining multiple links so that the links function as a single link with higher bandwidth. Aggregating multiple interfaces into a trunk to create a link has the following advantages:

- Link aggregation increases the bandwidth of the individual network interface cards (NICs) in an additive manner.
- If one link goes down, the other link can handle the traffic by itself.

Link aggregation occurs when you create a trunk. A *trunk* is a combination of two or more interfaces and cables configured as one link.

The examples in this implementation show a trunk that includes two tagged interfaces aggregated together. A *tagged interface* is an interface that is configured to process traffic for multiple VLANs. A VLAN tag identifies the specific VLAN and enables traffic to pass through that specific VLAN. To cause traffic for multiple VLANs to be passed through a single trunk, you must assign the same trunk to each VLAN.

In the example, we create a trunk (**trunk1**) that includes two interfaces, **1.1** and **1.2**, and then assign **trunk1** as a tagged interface to both VLAN **external** and VLAN **internal**. Both VLANs (**external** and **internal**) reside on the same network, and are combined to form a VLAN group.

With this configuration, inbound and outbound traffic passing between the BIG-IP system and the vendor switch can use either interface. For example, traffic destined for VLAN **external** can pass through either interface, **1.1** or **1.2**.

Illustration of link aggregation for a one-network topology

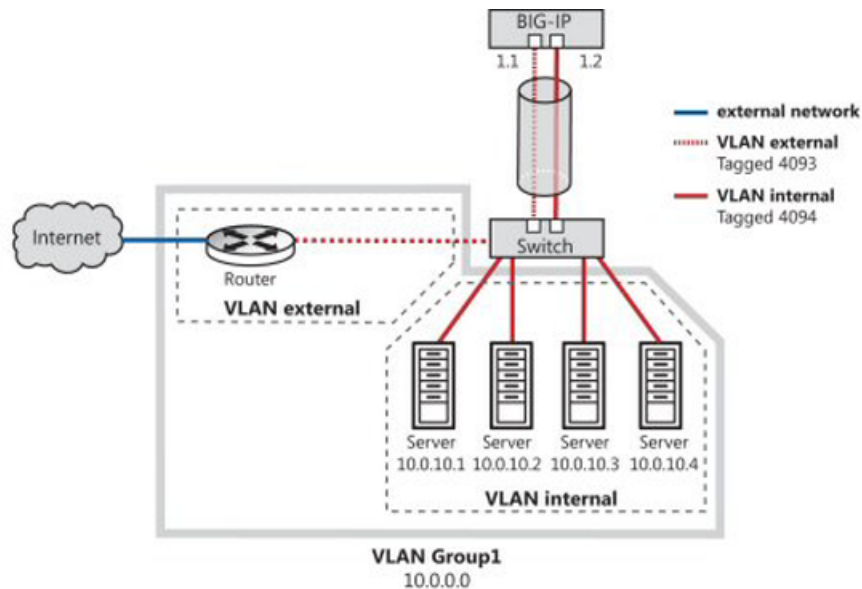


Figure 21: Link aggregation for a one-network topology

Task summary

Perform the following tasks to configure two interfaces (tagged VLANs) to function as a single link with higher bandwidth. In this implementation, you combine the two tagged VLANs into one VLAN group, where the two VLANs are on the same IP network.

Task list

Creating a trunk

Adding a tagged interface to a VLAN

Creating a load balancing pool

Creating a virtual server with source address affinity persistence

Removing the self IP addresses from the default VLANs

Creating a VLAN group

Creating a self IP for a VLAN group

Creating a trunk

You create a trunk on the BIG-IP® system so that the system can then aggregate the links to enhance bandwidth and ensure link availability.

1. On the Main tab, click **Network > Trunks**.
The Trunk List screen opens.
2. Click **Create**.

3. Name the trunk.
4. For the **Interfaces** setting, in the **Available** field, select an interface, and using the Move button, move the interface to the **Members** field. Repeat this action for each interface that you want to include in the trunk.
Trunk members must be untagged interfaces and cannot belong to another trunk. Therefore, only untagged interfaces that do not belong to another trunk appear in the **Available** list.
5. Select the **LACP** check box.
6. Click **Finished**.

After you create a trunk, the BIG-IP system aggregates the links to enhance bandwidth and prevent interruption in service.

Adding a tagged interface to a VLAN

After you aggregate the links, you assign the trunk to the VLAN as a tagged interface.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the Name column, click the relevant VLAN name.
This displays the properties of the VLAN.
3. For the **Interfaces** setting:
 - a) From the **Interface** list, select the trunk name.
 - b) From the **Tagging** list, select **Tagged**.
 - c) Click **Add**.

The trunk is assigned to the **external** and **internal** VLAN as a tagged interface.

Creating a load balancing pool

You can create a *load balancing pool* (a logical set of devices such as web servers that you group together to receive and process traffic) to efficiently distribute the load on your server resources.

Note: You must create the pool before you create the corresponding virtual server.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click << to move the monitor to the **Active** list.

Tip: Hold the Shift or Ctrl key to select more than one monitor at a time.

5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.
The default is **Round Robin**.

6. For the **Priority Group Activation** setting, specify how to handle priority groups:
 - Select **Disabled** to disable priority groups. This is the default option.
 - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.
 - c) In the **Service Port** field, type a port number, or select a service name from the list.
 - d) (Optional) In the **Priority** field, type a priority number.
 - e) Click **Add**.
8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server with source address affinity persistence

A virtual server represents a destination IP address for application traffic.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, in the **Address** field, type the IP address you want to use for the virtual server.
The IP address you type must be available and not in the loopback network.
5. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
6. Locate the relevant profile type for the traffic being managed, and either retain the default value or select a custom profile name.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. For the **Default Persistence Profile** setting, select **source_addr**.
This implements simple persistence, using the default source address affinity profile.

A client system now has a destination IP address on the BIG-IP system.

Removing the self IP addresses from the default VLANs

Remove the self IP addresses from the individual VLANs. After you create the VLAN group, you will create another self IP address for the VLAN group for routing purposes. The individual VLANs no longer need their own self IP addresses.

1. On the Main tab, click **Network > Self IPs**.
2. Select the check box for each IP address and VLAN that you want to delete.
3. Click **Delete**.
4. Click **Delete**.

The self IP address is removed from the Self IP list.

Creating a VLAN group

Create a VLAN group that includes the internal and external VLANs. Packets received by a VLAN in the VLAN group are copied onto the other VLAN. This allows traffic to pass through the BIG-IP® system on the same IP network.

1. On the Main tab, click **Network > VLANs > VLAN Groups**.
The VLAN Groups list screen opens.
2. Click **Create**.
The New VLAN Group screen opens.
3. In the **Name** field, type the name `myvlangroup`.
4. For the **VLANs** setting, select `internal` and `external` VLAN names in the **Available** list and move them to the **Members** list.
5. Click **Finished**.

Creating a self IP for a VLAN group

Before you create a self IP address, ensure that you have created at least one VLAN group.

You perform this task to create a self IP address for a VLAN group. The self IP address for the VLAN group provides a route for packets destined for the network. With the BIG-IP® system, the path to an IP network is a VLAN. However, with the VLAN group feature used in this procedure, the path to the IP network `10.0.0.0` is actually through more than one VLAN. As IP routers are designed to have only one physical route to a network, a routing conflict can occur. With a self IP address on the BIG-IP system, you can resolve the routing conflict by associating a self IP address with the VLAN group.

1. On the Main tab, click **Network > Self IPs**.
2. Click **Create**.
The New Self IP screen opens.
3. In the **IP Address** field, type an IPv4 or IPv6 address.
This IP address should represent the address space of the VLAN group that you specify with the **VLAN/Tunnel** setting.
4. In the **Netmask** field, type the network mask for the specified IP address.
For example, you can type `255.255.255.0`.
5. From the **VLAN/Tunnel** list, select the VLAN group with which to associate this self IP address.
6. From the **Port Lockdown** list, select **Allow Default**.
7. For the **Traffic Group** setting, choose one of the following actions:

Action	Result
Retain the default setting, traffic-group-local-only (non-floating).	The system creates a non-floating self IP address that becomes a member of <code>traffic-group-local-only</code> .
Select the check box labeled Inherit traffic group from current partition / path.	The system creates a floating self IP address that becomes a member of <code>traffic-group-1</code> .
Select a traffic group from the Traffic Group list.	The system creates a floating self IP address that becomes a member of the selected traffic group.

8. Click **Finished.**

The screen refreshes, and displays the new self IP address.

The BIG-IP system can send and receive traffic through the specified VLAN or VLAN group.

Using Link Aggregation with Tagged VLANs for a Two-network Topology

Overview: Configuring link aggregation of two interfaces using tagged VLANs on two networks

You can use the BIG-IP® system in an aggregated two-interface load balancing topology. *Link aggregation* is the process of combining multiple links so that the links function as a single link with higher bandwidth. Aggregating multiple interfaces into a trunk to create a link has the following advantages:

- Link aggregation increases the bandwidth of the individual network interface cards (NICs) in an additive manner.
- If one link goes down, the other link can handle the traffic by itself.

Link aggregation occurs when you create a trunk. A *trunk* is a combination of two or more interfaces and cables configured as one link.

The examples in this implementation show a trunk that includes two tagged interfaces aggregated together. A *tagged interface* is an interface that is configured to process traffic for multiple VLANs. A VLAN tag identifies the specific VLAN and allows traffic to be passed through that specific VLAN. To cause traffic for multiple VLANs to be passed through a single trunk, you must assign the same trunk to each VLAN.

In the examples, we create a trunk (**trunk1**) that includes two interfaces, **1.1** and **1.2**, and then assign **trunk1** as a tagged interface to both VLAN **external** and VLAN **internal**. One network is connected to VLAN **external**, and a separate network is connected to VLAN **internal**. Consequently, inbound and outbound traffic passing between the BIG-IP system and the vendor switch can use either interface. For example, traffic destined for VLAN **external** can pass through either interface, **1.1** or **1.2**.

Illustration of link aggregation for a two-network topology

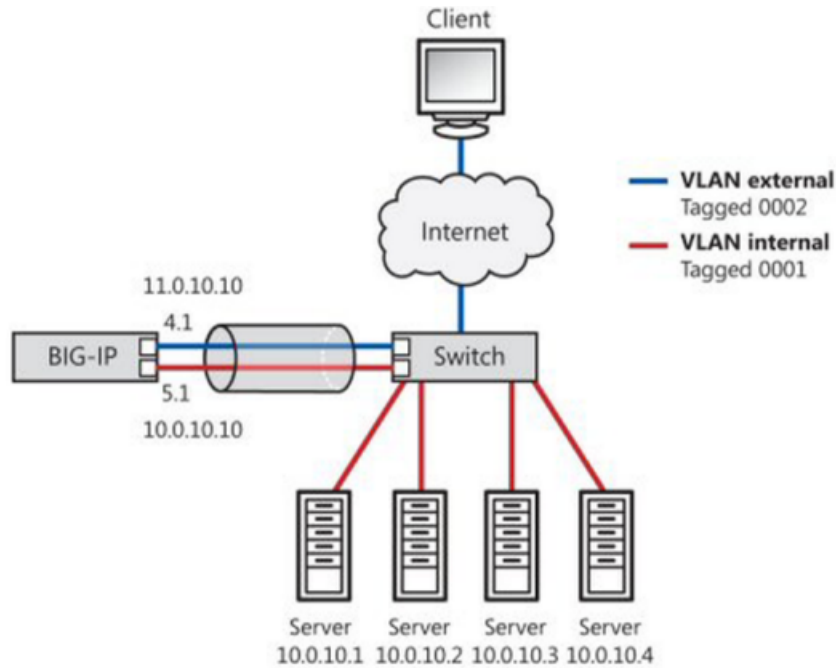


Figure 22: Link aggregation for a two-network topology

Task summary

Perform the following tasks to configure two interfaces (tagged VLANs) to function as a single link with higher bandwidth. In this implementation, each tagged VLAN is on a separate network.

Task list

Creating a trunk

Adding a tagged interface to a VLAN

Creating a load balancing pool

Creating a virtual server with source address affinity persistence

Creating a trunk

You create a trunk on the BIG-IP® system so that the system can then aggregate the links to enhance bandwidth and ensure link availability.

1. On the Main tab, click **Network > Trunks**.
The Trunk List screen opens.
2. Click **Create**.
3. Name the trunk.

4. For the **Interfaces** setting, in the **Available** field, select an interface, and using the Move button, move the interface to the **Members** field. Repeat this action for each interface that you want to include in the trunk.
Trunk members must be untagged interfaces and cannot belong to another trunk. Therefore, only untagged interfaces that do not belong to another trunk appear in the **Available** list.
5. Select the **LACP** check box.
6. Click **Finished**.

After you create a trunk, the BIG-IP system aggregates the links to enhance bandwidth and prevent interruption in service.

Adding a tagged interface to a VLAN

After you aggregate the links, you assign the trunk to the VLAN as a tagged interface.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the Name column, click the relevant VLAN name.
This displays the properties of the VLAN.
3. For the **Interfaces** setting:
 - a) From the **Interface** list, select the trunk name.
 - b) From the **Tagging** list, select **Tagged**.
 - c) Click **Add**.

The trunk is assigned to the **external** and **internal** VLAN as a tagged interface.

Creating a load balancing pool

You can create a *load balancing pool* (a logical set of devices such as web servers that you group together to receive and process traffic) to efficiently distribute the load on your server resources.

Note: You must create the pool before you create the corresponding virtual server.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click << to move the monitor to the **Active** list.

Tip: Hold the Shift or Ctrl key to select more than one monitor at a time.

5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.
The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:

- Select **Disabled** to disable priority groups. This is the default option.
 - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.
 - c) In the **Service Port** field, type a port number, or select a service name from the list.
 - d) (Optional) In the **Priority** field, type a priority number.
 - e) Click **Add**.
 8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server with source address affinity persistence

A virtual server represents a destination IP address for application traffic.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, in the **Address** field, type the IP address you want to use for the virtual server.
The IP address you type must be available and not in the loopback network.
5. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
6. Locate the relevant profile type for the traffic being managed, and either retain the default value or select a custom profile name.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. For the **Default Persistence Profile** setting, select **source_addr**.
This implements simple persistence, using the default source address affinity profile.

A client system now has a destination IP address on the BIG-IP system.

Configuring Packet Filtering

Overview: Setting up packet filtering

Packet filters enhance network security by specifying whether a BIG-IP® system interface should accept or reject certain packets based on criteria that you specify. Packet filters enforce an access policy on incoming traffic. They apply to incoming traffic only.

You implement packet filtering by creating packet filter rules. The primary purpose of a packet filter rule is to define the criteria that you want the BIG-IP system to use when filtering packets. Examples of criteria that you can specify in a packet filter rule are:

- The source IP address of a packet
- The destination IP address of a packet
- The destination port of a packet

You specify the criteria for applying packet filter rules within an expression. When creating a packet filter rule, you can instruct the Configuration utility to build an expression for you, in which case you need only choose the criteria from predefined lists, or you can write your own expression text, using the syntax of the `tcpdump` utility.

Note: *Packet filter rules are unrelated to iRules®.*

You can also configure global packet filtering that applies to all packet filter rules that you create.

Task summary

By setting up some basic IP routing and configuring packet filtering, specific hosts on the internal VLAN can connect to the internal VLAN's self IP address. These hosts can also use common Internet services such as HTTP, HTTPS, DNS, FTP, and SSH. Traffic from all other hosts in the internal VLAN is rejected.

Task list

Enabling SNAT automap for internal and external VLANs

Creating a default gateway pool

Creating a forwarding virtual server

Enabling packet filtering

Creating a packet filter rule

Enabling SNAT automap for internal and external VLANs

You can configure SNAT automapping on the BIG-IP system for internal and external VLANs.

1. On the Main tab, click **Local Traffic** > **Address Translation**.
The **SNAT List** screen displays a list of existing SNATs.

2. Click **Create**.
3. Name the new SNAT.
4. From the **Translation** list, select **Automap**.
5. For the **VLAN / Tunnel List** setting, in the **Available** list, select **external** and **internal**, and using the Move button, transfer the VLANs to the **Selected** list.
6. Click the **Finished** button.

SNAT automapping on the BIG-IP system is configured for internal and external VLANs.

Creating a default gateway pool

Create a default gateway pool for the system to use to forward traffic.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, from the **Available** list, select the **gateway_icmp** monitor and move the monitor to the **Active** list.
5. Using the **New Members** setting, add each router that you want to include in the default gateway pool:
 - a) Type the IP address of a router in the **Address** field.
 - b) Type an asterisk (*) in the **Service Port** field, or select ***All Services** from the list.
 - c) Click **Add**.
6. Click **Finished**.

Creating a forwarding virtual server

A virtual server represents a destination IP address for application traffic.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For a network, in the **Destination Address** field, type an IPv4 or IPv6 address in CIDR format to allow all traffic to be translated.

The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is 0.0.0.0/0, and an IPv6 address/prefix is ::/0.
5. From the **Service Port** list, select ***All Ports**.
6. In the Configuration area of the screen, from the **Type** list, select **Forwarding (IP)**.
7. From the **Protocol** list, select ***All Protocols**.
8. From the **VLAN/Tunnel Traffic** list, select **Enabled On**.
9. For the **VLAN List** setting, from the **Available** box, select **internal**, and click the Move button to move the VLAN name to the **Selected** box.
- 10.

11. In the Resources area of the screen, locate the **Default Pool** setting and select the pool you created previously.
12. Click **Finished**.

You now have a destination IP address on the BIG-IP system for application traffic.

Enabling packet filtering

Before creating a packet filtering rule, you must enable packet filtering. When you enable packet filtering, you can specify the MAC addresses, IP addresses, and VLANs that you want to be exempted from packet filter evaluation.

1. On the Main tab, click **Network > Packet Filters**.
The Packet Filters screen opens.
2. From the **Packet Filtering** list, select **Enabled**.
3. From the **Unhandled Packet Action** list, select **Accept**.
4. For the **Options** setting, retain the default value or select the check boxes as needed.
5. For the **Protocols** setting, retain the default value or clear the check boxes as needed.
6. From the **MAC Addresses** list, specify a value:

Value	Description
None	When you select this value, all MAC addresses are exempt from packet filter evaluation.
Always Accept	When you select this value, you can specify the MAC addresses that are exempt from packet filter evaluation, and the BIG-IP Configuration utility displays additional settings.

7. If you directed the **MAC Addresses** setting to always accept specific MAC addresses, provide the details:
 - a) In the **Add** field, type a MAC address and click **Add**.
The MAC address appears in the **MAC Address List** field.
 - b) Repeat this step for each MAC address that you want the system to exempt from packet filter evaluation.

8. From the **IP Addresses** list, specify a value:

Value	Description
None	When you select this value, all IP addresses are exempt from packet filter evaluation.
Always Accept	When you select this value, you can specify the IP addresses that are exempt from packet filter evaluation. The BIG-IP Configuration utility displays additional settings.

9. If you directed the **IP Addresses** setting to always accept specific IP addresses, provide the details:
 - a) In the **Add** field, type an IP address and click **Add**.
The IP address appears in the **IP Address List** field.
 - b) Repeat this step for each IP address that you want the system to exempt from packet filter evaluation.

10. From the **VLANs** list, specify a value:

Value	Description
None	When you select this value, all VLANs are exempt from packet filter evaluation.
Always Accept	When you select this value, you can specify the VLANs that are exempt from packet filter evaluation. The BIG-IP Configuration utility displays additional settings.

11. If you configured the **VLANs** setting to always accept specific VLANs, then use the **Move** button to move one or more VLAN names from the **Available** list to the **Selected** list.
12. Click **Update**.

After you enable packet filtering, the BIG-IP® system filters packets according to the criteria in the packet filter rule and the values you configured when enabling the packet filter.

Creating a packet filter rule

When implementing packet filtering, you need to create a packet filter rule.

1. On the Main tab, click **Network > Packet Filters**.
The Packet Filters screen opens.
2. Click **Rules**.
3. Click **Create**.
4. Name the rule.
5. From the **Order** list, select **First**.
6. From the **Action** list, select **Reject**.
7. From the **VLAN / Tunnel** list, select **internal**.
8. From the **Logging** list, select **Enabled**.
9. From the **Filter Expression Method** list, select **Enter Expression Text**.
10. In the **Filter Expression** field, type an expression.
For example: `not dst port 80 and not dst port 443 and not dst port 53 and not dst port 22 and not dst port 20 and not dst port 21 and not dst host <internal self IP address>`

Note: Replace `<internal self IP address>` with the actual self IP address of VLAN internal.

11. Click **Finished**.

The packet filter rule is available.

Referencing an External File from within an iRule

Overview: Referencing an external file from an iRule

Using the BIG-IP® Configuration utility or **tmsh**, you can import a file or URL from another system to the BIG-IP system, with content that you want an iRule to return to a client, based on some iRule event. Possible uses for this feature are:

- To send a web page other than the page that the client requested. For example, you might want the system to send a maintenance page instead of the requested page.
- To send an image.
- To use a file as a template and modify the file in the iRule before sending the file.
- To download policy information from an external server and merge that data with a locally-stored policy.

The file that an iRule accesses is known as an *iFile*, and can be any type of file, such as a binary file or a text file. These files are read-only files.

This example shows an iRule that references an iFile named `ifileURL`, in partition `Common`:

```
ltm rule ifile_rule {
  when HTTP_RESPONSE {
    # return a list of iFiles in all partitions
    set listifiles [ifile listall]
    log local0. "list of ifiles: $listifiles"

    # return the attributes of an iFile specified
    array set array_attributes [ifile attributes "/Common/ifileURL"]
    foreach {array attr} [array get array_attributes ] {
      log local0. "$array : $attr"
    }

    # serve an iFile when http status is 404.
    set file [ifile get "/Common/ifileURL"]
    log local0. "file: $file"
    if { [HTTP::status] equals "404" } {
      HTTP::respond 200 ifile "/Common/ifileURL"
    }
  }
}
```

iRule commands for iFiles

This list shows the commands available for referencing an iFile within an iRule. All of these commands return a string, except for the command `[ifile attributes IFILENAME]`, which returns an array.

Available iRule commands for referencing an iFile

```
[ifile get IFILENAME]
```

```
[ifile listall]
[ifile attributes IFILENAME]
[ifile size IFILENAME]
[ifile last_updated_by IFILENAME]
[ifile last_update_time IFILENAME]
[ifile revision IFILENAME]
[ifile checksum IFILENAME]
[ifile attributes IFILENAME]
```

Task summary

You can import an existing file to the BIG-IP® system, create an iFile that is based on the imported file, and then write an iRule that returns the content of that file to a client system, based on an iRule event.

Task list

Importing a file for an iRule

Creating an iFile

Writing an iRule that references an iFile

Importing a file for an iRule

Before you perform this task, the file you want to import must reside on the system you specify.

You can import a file from another system onto the BIG-IP® system, as the first step in writing an iRule that references that file.

1. On the Main tab, click **System > File Management > iFile List > Import**.
2. For the **File Name** setting, click **Browse**.
The system opens a browse window so that you can locate the file that you want to import to the BIG-IP system.
3. Browse for the file and click **Open**.
The name of the file you select appears in the **File Name** setting.
4. In the **Name** field, type a new name for the file, such as `lk.html`.
The new file name appears in the list of imported files.
5. Click the **Import** button.

After you perform this task, the file that you imported resides on the BIG-IP system.

Creating an iFile

As a prerequisite, ensure that the current administrative partition is set to the partition in which you want the iFile to reside. Also ensure that the file has been imported to the BIG-IP® system.

You perform this task to create an iFile that you can then reference in an iRule.

1. On the Main tab, click **Local Traffic > iRules > iFile List**.
2. Click **Create**.
3. In the **Name** field, type a new name for the iFile, such as `ifileURL`.

4. From the **File Name** list, select the name of the imported file object, such as `lk.html`.
5. Click **Finished**.
The new iFile appears in the list of iFiles.

The result of this task is that you now have a file that an iRule can reference.

Writing an iRule that references an iFile

You perform this task to create an iRule that references an iFile.

***Note:** If the iFile resides in partition `/Common`, then specifying the partition when referencing the iFile is optional. If the iFile resides in a partition other than `/Common`, such as `/Partition_A`, you must include the partition name in the iFile path name within the iRule.*

1. On the Main tab, click **Local Traffic > iRules**.
The iRule List screen opens, displaying any existing iRules.
2. Click **Create**.
The New iRule screen opens.
3. In the **Name** field, type a name, such as `my_irule`.
The full path name of the iRule cannot exceed 255 characters.
4. In the **Definition** field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.
For complete and detailed information iRules syntax, see the F5 Networks DevCentral web site (<http://devcentral.f5.com>).
5. Click **Finished**.
The new iRule appears in the list of iRules on the system.

Implementation result

You now have an iRule that accesses a file on the BIG-IP® system, based on a particular iRule event.

Configuring Remote User Authentication and Authorization

Overview: Remote authentication and authorization of BIG-IP user accounts

The BIG-IP® system includes a comprehensive solution for managing BIG-IP administrative accounts on your network. With this solution, you can:

Use a remote server to store BIG-IP system user accounts.

The BIG-IP system includes support for using a remote authentication server to store BIG-IP system user accounts. After creating BIG-IP system accounts on the remote server (using the server vendor's instructions), you can configure the BIG-IP system to use remote user authentication and authorization (access control) for that server type.

Assign group-based access.

The BIG-IP system includes an optional feature known as *remote role groups*. With the *remote role groups* feature, you can use existing group definitions on the remote server to define the access control properties for users in a group. This feature not only provides more granularity in assigning user privileges, but also removes any need to duplicate remote user accounts on the BIG-IP system for the purpose of assigning those privileges.

Propagate a set of authorization data to multiple BIG-IP systems.

The BIG-IP system includes a tool for propagating BIG-IP system configuration data to multiple BIG-IP devices on the network. This tool is known as the Single Configuration File (SCF) feature.

Task summary

You can configure the BIG-IP® system to authorize user accounts that are stored on a remote authentication server.

Important: *If you configure access control settings for group-based accounts (using the remote role groups feature), the BIG-IP system always applies those settings, rather than the default access control settings, to group-based accounts.*

The BIG-IP® system supports several types of authentication servers for storing BIG-IP system administrative user accounts. The actual procedure you use to specify the type of remote server differs, depending on the server type.

Task list

Specifying LDAP or Active Directory server information

Specifying client certificate LDAP server information

Specifying RADIUS server information

Specifying TACACS+ server information

Configuring access control for remote user groups

Saving access control settings to a file

Importing BIG-IP configuration data onto other BIG-IP systems

Specifying LDAP or Active Directory server information

Before you begin:

- Verify that the BIG-IP® system user accounts have been created on the remote authentication server.
- Verify that the appropriate user groups, if any, are defined on the remote authentication server.
- If you want to verify the certificate of the authentication server, import one or more SSL certificates.

You can configure the BIG-IP system to use an LDAP or Microsoft® Windows® Active Directory® server for authenticating BIG-IP system user accounts, that is, traffic that passes through the management interface (MGMT).

Important: The values you specify in this procedure for the **Role**, **Partition Access**, and **Terminal Access** settings do not apply to group-based access control. These values represent the default values that the BIG-IP system applies to any user account that is not part of a remotely-stored user group. Also, for the *Other External Users* user account, you can modify the **Role**, **Partition Access**, and **Terminal Access** settings only when your current partition on the BIG-IP system is set to *Common*. If you attempt to modify these settings when your current partition is other than *Common*, the system displays an error message.

1. On the Main tab, click **System > Users > Authentication**.
2. On the menu bar, click **Authentication**.
3. Click **Change**.
4. From the **User Directory** list, select **Remote - LDAP** or **Remote - Active Directory**.
5. In the **Host** field, type the IP address of the remote server.
The route domain to which this address pertains must be route domain 0.
6. For the **Port** setting, retain the default port number (389) or type a new port number.
This number represents the port number that the BIG-IP system uses to access the remote server.
7. In the **Remote Directory Tree** field, type the file location (tree) of the user authentication database on the LDAP or Active Directory server.
At minimum, you must specify a domain component (that is, `dc=[value]`).
8. For the **Scope** setting, retain the default value (`Sub`) or select a new value.
This setting specifies the level of the remote server database that the BIG-IP system should search for user authentication.
9. For the **Bind** setting, specify a user ID login for the remote server:
 - a) In the **DN** field, type the distinguished name for the remote user ID.
 - b) In the **Password** field, type the password for the remote user ID.
 - c) In the **Confirm** field, re-type the password that you typed in the **Password** field.
10. In the **User Template** field, type a string that contains a variable representing the distinguished name of the user, in the format `%s`.
This field can contain only one `%s` and cannot contain any other format specifiers.
For example, you can specify a user template such as `%s@siterequest.com` or `uxml:id=%s,ou=people,dc=siterequest,dc=com`.
The result is that when a user attempts to log on, the system replaces `%s` with the user name specified in the Basic Authentication dialog box, and passes that name as the distinguished name for the bind operation. The system also passes the associated password as the password for the bind operation.
11. For the **Check Member Attribute in Group** setting, select the check box if you want the system to check the user's member attribute in the remote LDAP or AD group.

12. To enable SSL-based authentication, from the **SSL** list select **Enabled** and, if necessary, configure these settings:
 - a) From the **SSL CA Certificate** list, select the name of a chain certificate, that is, the third-party CA or self-signed certificate that normally resides on the remote authentication server.
 - b) From the **SSL Client Key** list, select the name of the client SSL key.
Use this setting only when the remote server requires that the client present a certificate.
 - c) From the **SSL Client Certificate** list, select the name of the client SSL certificate.
Use this setting only if the remote server requires that the client present a certificate.
13. In the **Login LDAP Attribute** field, type the account name for the LDAP server.
The value for this option is normally the user ID. However, if the server is a Microsoft® Windows® Active Directory® server, the value must be the account name `sAMAccountName` (case-sensitive). The default value is none.
14. From the **Client Certificate Name Field** list:
 - a) Select either a subject alternate name or the subject name (**Common Name**).
 - b) If you select the subject alternate name **Other Name**, then in the **OID** field, type an object identifier (OID).
The OID indicates the format and semantics of the subject alternate name.
15. From the **Role** list, select the user role that you want the BIG-IP system to assign by default to all BIG-IP system user accounts authenticated on the remote server.
16. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
17. From the **Terminal Access** list, select either of these as the default terminal access option for remotely-authenticated user accounts:

Option	Description
Disabled	Choose this option when you do not want the remotely-stored user accounts to have terminal access to the BIG-IP system.
tmsh	Choose this option when you want the remotely-stored user accounts to have only <code>tmsh</code> access to the BIG-IP system.
18. Click **Finished**.

You can now authenticate administrative user accounts that are stored on a remote LDAP or Active Directory server. If you have no need to configure access control for remotely-stored user groups, your configuration tasks are complete.

Specifying client certificate LDAP server information

Verify that the required user accounts for the BIG-IP® system exist on the remote authentication server.

For authenticating BIG-IP system user accounts (that is, traffic that passes through the management interface [MGMT]), you can configure the BIG-IP system to authenticate certificates issued by a certificate authority's Online Certificate Status Protocol (OCSP) responder.

Important: The values you specify in this procedure for the **Role**, **Partition Access**, and **Terminal Access** settings do not apply to group-based authorization. These values represent the default values or locally configured user accounts (which override the default role) that the BIG-IP system applies to any user account that is not part of a remote role group.

1. On the Main tab, click **System > File Management > Apache Certificate List > Import**, browse for the certificate file to import, type a name, and click **Import**.
The certificate will be added to the Apache Certificate list.
2. On the Main tab, click **System > Users > Authentication**.
3. On the menu bar, click **Authentication**.
4. Click **Change**.
5. From the **User Directory** list, select **Remote - ClientCert LDAP**.
6. In the **Host** field, type the IP address of the remote server.
The route domain to which this address pertains must be route domain 0.
7. For the **Port** setting, retain the default port number (389) or type a new port number.
This number represents the port number that the BIG-IP system uses to access the remote server.
8. In the **Remote Directory Tree** field, type the file location (tree) of the user authentication database on the client certificate server.
At minimum, you must specify a domain component (that is, `dc=[value]`).
9. For the **Scope** setting, retain the default value (Sub) or select a new value.
This setting specifies the level of the remote server database that the BIG-IP system should search for user authentication.
10. For the **Bind** setting, specify a user ID login for the remote server:
 - a) In the **DN** field, type the distinguished name for the remote user ID.
 - b) In the **Password** field, type the password for the remote user ID.
 - c) In the **Confirm** field, re-type the password that you typed in the **Password** field.
11. To enable SSL-based authentication, from the **SSL** list select **Enabled** and, if necessary, configure these settings:
 - a) From the **SSL CA Certificate** list, select the name of a chain certificate; that is, the third-party CA or self-signed certificate that normally resides on the remote authentication server.
 - b) From the **SSL Client Key** list, select the name of the client SSL key.
Use this setting only when the remote server requires that the client present a certificate.
 - c) From the **SSL Client Certificate** list, select the name of the client SSL certificate.
Use this setting only if the remote server requires that the client present a certificate.
12. In the **CA Certificate** field, type the absolute folder path of `apache-ssl-cert fileobject` for the CA signing authority.
The absolute folder path is `/Common/<folder path>/<certificate name>`. To determine the absolute folder path of the `apache-ssl-cert fileobject`, click **System > File Management > Apache Certificate List** and note the target certificate's partition and path.

Important: *Apache certificates can only be stored within /Common.*

13. In the **Login Name** field, type an LDAP search prefix that will contain the distinguished name (DN) from the user certificate, such as `CN`.
This specifies the LDAP attribute to be used as a login name. The default is disabled.
14. In the **Login LDAP Attribute** field, type the account name for the LDAP server.
The value for this option is normally the user ID. However, if the server is a Microsoft® Windows® Active Directory® server, the value must be the account name `sAMAccountName` (case-sensitive). The default value is none.
15. In the **Login Filter** field, type the LDAP attribute that contains the short name of the user.

This specifies the filter to be applied on the common name (CN) of the client certificate and usually this is the user ID or `sAMAccountName`. The filter is a regular expression used to extract required information from the CN of the client certificate that is matched against the LDAP search results. The default is disabled.

16. For the **Depth** setting, retain the default value (10) or type a new value for verification depth.
17. From the **Role** list, select the user role that you want the BIG-IP system to assign by default to all BIG-IP system user accounts authenticated on the remote server.
18. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
19. From the **Terminal Access** list, select either of these as the default terminal access option for remotely-authenticated user accounts:

Option	Description
Disabled	Choose this option when you do not want the remotely-stored user accounts to have terminal access to the BIG-IP system.
tmsh	Choose this option when you want the remotely-stored user accounts to have only <code>tmsh</code> access to the BIG-IP system.

20. Click **Finished**.

You can now authenticate administrative traffic for user accounts that are stored on a remote client certificate server. If you have no need to configure group-based user authorization, your configuration tasks are complete.

Specifying RADIUS server information

Before you begin:

- Verify that the BIG-IP® system user accounts have been created on the remote authentication server.
- Verify that the appropriate user groups, if any, are defined on the remote authentication server.

You can configure the BIG-IP system to use a RADIUS server for authenticating BIG-IP system user accounts, that is, traffic that passes through the management interface (MGMT).

Important: The values you specify in this procedure for the **Role**, **Partition Access**, and **Terminal Access** settings do not apply to group-based authorization. These values represent the default values that the BIG-IP system applies to any user account that is not part of a role group that is defined on the remote authentication server. Also, for the *Other External Users* user account, you can modify the **Role**, **Partition Access**, and **Terminal Access** settings only when your current partition on the BIG-IP system is set to *Common*. If you attempt to modify these settings when your current partition is other than *Common*, the system displays an error message.

1. On the Main tab, click **System > Users > Authentication**.
2. On the menu bar, click **Authentication**.
3. Click **Change**.
4. From the **User Directory** list, select **Remote - RADIUS**.
5. For the **Primary** setting:
 - a) In the **Host** field, type the name of the primary RADIUS server.
The route domain with which this host is associated must be route domain 0.
 - b) In the **Secret** field, type the password for access to the primary RADIUS server.

- c) In the **Confirm** field, re-type the RADIUS secret.
6. If you set the **Server Configuration** setting to **Primary and Secondary**, then for the **Secondary** setting:
 - a) In the **Host** field, type the name of the secondary RADIUS server.
The route domain with which this host is associated must be route domain 0.
 - b) In the **Secret** field, type the password for access to the secondary RADIUS server.
 - c) In the **Confirm** field, re-type the RADIUS secret.
7. From the **Role** list, select the user role that you want the BIG-IP system to assign by default to all BIG-IP system user accounts authenticated on the remote server.
8. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
9. From the **Terminal Access** list, select either of these as the default terminal access option for remotely-authenticated user accounts:

Option	Description
Disabled	Choose this option when you do not want the remotely-stored user accounts to have terminal access to the BIG-IP system.
tmsh	Choose this option when you want the remotely-stored user accounts to have only <code>tmsh</code> access to the BIG-IP system.
10. Click **Finished**.

You can now authenticate administrative traffic for BIG-IP system user accounts that are stored on a remote RADIUS server. If you have no need to configure access control for remotely-stored user groups, your configuration tasks are complete.

Specifying TACACS+ server information

Before you begin:

- Verify that the BIG-IP® system user accounts have been created on the remote authentication server.
- Verify that the appropriate user groups, if any, are defined on the remote authentication server.

You can configure the BIG-IP system to use a TACACS+ server for authenticating BIG-IP system user accounts, that is, traffic that passes through the management interface (MGMT).

Important: The values you specify in this procedure for the **Role**, **Partition Access**, and **Terminal Access** settings do not apply to group-based authorization. These values represent the default values that the BIG-IP system applies to any user account that is not part of a remote role group. Also, for the *Other External Users* user account, you can modify the **Role**, **Partition Access**, and **Terminal Access** settings only when your current partition on the BIG-IP system is set to *Common*. If you attempt to modify these settings when your current partition is other than *Common*, the system displays an error message.

1. On the Main tab, click **System > Users > Authentication**.
2. On the menu bar, click **Authentication**.
3. Click **Change**.
4. From the **User Directory** list, select **Remote - TACACS+**.
5. For the **Servers** setting, type an IP address for the remote TACACS+ server.
The route domain to which this address pertains must be route domain 0.

6. Click **Add**.
The IP address for the remote TACACS+ server appears in the **Servers** list.
7. In the **Secret** field, type the password for access to the TACACS+ server.

Warning: Do not include the symbol # in the secret. Doing so causes authentication of local user accounts (such as `root` and `admin`) to fail.

8. In the **Confirm Secret** field, re-type the TACACS+ secret.
9. From the **Encryption** list, select an encryption option:

Option	Description
Enabled	Specifies that the system encrypts the TACACS+ packets.
Disabled	Specifies that the system sends unencrypted TACACS+ packets.
10. In the **Service Name** field, type the name of the service that the user is requesting to be authenticated to use (usually `ppp`).
Specifying the service causes the TACACS+ server to behave differently for different types of authentication requests. Examples of service names that you can specify are: `ppp`, `slip`, `arap`, `shell`, `tty-daemon`, `connection`, `system`, and `firewall`.
11. In the **Protocol Name** field, type the name of the protocol associated with the value specified in the **Service Name** field.
This value is usually `ip`. Examples of protocol names that you can specify are: `ip`, `lcp`, `ipx`, `atalk`, `vines`, `lat`, `xremote`, `tn3270`, `telnet`, `rlogin`, `pad`, `vpdn`, `ftp`, `http`, `deccp`, `osicp`, and `unknown`.
12. From the **Role** list, select the user role that you want the BIG-IP system to assign by default to all BIG-IP system user accounts authenticated on the remote server.
13. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
14. From the **Terminal Access** list, select either of these as the default terminal access option for remotely-authenticated user accounts:

Option	Description
Disabled	Choose this option when you do not want the remotely-stored user accounts to have terminal access to the BIG-IP system.
tmsh	Choose this option when you want the remotely-stored user accounts to have only <code>tmsh</code> access to the BIG-IP system.

15. Click **Finished**.

You can now authenticate administrative traffic for BIG-IP system user accounts that are stored on a remote TACACS+ server. If you have no need to configure access control for remotely-stored user groups, your configuration tasks are complete.

Configuring access control for remote user groups

You perform this task to assign a user role, a corresponding administrative partition, and a type of terminal access to a remotely-stored group of user accounts. For a given user group, you can assign as many role-partition combinations as you need, as long as each role is associated with a different partition. If the partition you associate with a role is `All`, this entry might or might not take effect, depending on whether the `All` designation conflicts with other role-partition combinations for that user group. For any conflicts,

line order in the configuration is a consideration. To assign multiple role-partition combinations for a user group, you repeat this task for each combination, specifying the same attribute string for each task.

1. On the Main tab, click **System > Users**.
2. On the menu bar, click **Remote Role Groups**.
3. Click **Create**.
4. In the **Group Name** field, type the group name that is defined on the remote authentication server.
An example of a group name is **BigIPOperatorsGroup**.
5. In the **Line Order** field, type a number.

This value specifies the order of this access control configuration in the file `/config/bigip/auth/remoterole` for the named group. The LDAP and Active Directory servers read this file line by line. The order of the information is important; therefore, F5 Networks recommends that you specify a value of 1000 for the first line number. This allows you, in the future, to insert lines before the first line.

6. In the **Attribute String** field, type an attribute.

An example of an attribute string is

```
memberOf=cn=BigIPOperatorsGroup,cn=users,dc=dev,dc=net.
```

The BIG-IP system attempts to match this attribute with an attribute on the remote authentication server. On finding a match, the BIG-IP system applies the access control settings defined here to the users in that group. If a match is not found, the system applies the default access control settings to all remotely-stored user accounts (excluding any user account for which access control settings are individually configured).

7. From the **Remote Access** list, select a value.

Option	Description
Enabled	Choose this value if you want to enable remote access for the defined user group.
Disabled	Choose this value if you want to disable remote access for the defined user group. Note that if you configure multiple instances of this remote role group (one instance for each role-partition pair for the attribute string), then choosing a value of Disabled disables remote access for all user group members, regardless of the remote role group instance.

8. From the **Assigned Role** list, select a user role for the remote user group.

9. From the **Partition Access** list, select an administrative partition value.

Option	Description
All	Choose this value to give users in the defined group access to their authorized objects in all partitions on the BIG-IP system.
<i>partition_name</i>	Choose a specific partition name to give users in the defined group access to that partition only.
Common	Choose this value to give users in the defined group access to partition Common only.

10. From the **Terminal Access** list, select the type of command-line access you want to grant users in the group, if any.

11. Click **Finished** or **Repeat**.

After you perform this task, the user group that you specified has the assigned role, partition access, and terminal access properties assigned to it.

Saving access control settings to a file

You can save the running configuration of the system, including all settings for remote user authentication and authorization, in a flat, text file with a specified name and the extension `.scf`.

1. On the BIG-IP® system, access a command-line prompt.
2. At the prompt, open the Traffic Management Shell by typing the command `tmsh`.
3. Type `sys save filename`.
`sys save myConfiguration053107` creates the file `myConfiguration053107.scf` in the `var/local/scf` directory.
`sys save /config/myConfiguration` creates the file `myConfiguration.scf` in the `/config` directory.

You can now import this file onto other BIG-IP devices on the network.

Importing BIG-IP configuration data onto other BIG-IP systems

You can use the `tmsh sys load` command to import a single configuration file (SCF), including access control data, onto other BIG-IP® devices on the network.

Note: This task is optional.

1. On the BIG-IP system on which you created the SCF, access a command-line prompt.
 2. Copy the SCF that you previously created to a location on your network that you can access from the system that you want to configure.
 3. Edit the SCF to reflect the management routing and special passwords of the BIG-IP system that you want to configure:
 - a) Open the SCF in an editor.
 - b) Where necessary, change the values of the management IP address, network mask, management default route, self IP addresses, virtual server IP addresses, routes, default routes, and host name fields to the values for the new system.
 - c) If necessary, change the passwords for the `root` and `admin` accounts using the command `user name password none newpassword password`.
-

Important: When configuring a unit that is part of a redundant system configuration and that is using the SCF from the peer unit, do not modify the `root` and `admin` accounts. These accounts must be identical on both units of the redundant system.

- d) Save the edited SCF.
4. On the BIG-IP system that you want to configure, open the Traffic Management Shell by typing the command `tmsh`.
 5. Type `sys load scf_filename`.
`sys load myConfiguration053107.scf` saves a backup of the running configuration in the `/var/local/scf` directory, and then resets the running configuration with the configuration contained in the SCF you are loading.

Configuring Administrative Partitions to Control User Access

Overview: Administrative partitions for user access control

The BIG-IP® system includes a powerful authorization feature known as administrative partitions. Using the *administrative partitions* feature, you ensure that BIG-IP system grants administrative users exactly the right type and amount of access to BIG-IP system resources. As a result, you can tailor user access to resources to exactly fit the needs of your organization.

Task summary

There are two main tasks for controlling user access to BIG-IP® system objects.

Task list

- Creating an administrative partition*
- Assigning roles to a user account*

Creating an administrative partition

You perform this task to create an administrative partition. An *administrative partition* creates an access control boundary for users and applications.

1. On the Main tab, expand **System** and click **Users**.
The Users List screen opens.
2. On the menu bar, click **Partition List**.
3. Click **Create**.
The New Partition screen opens.
4. In the **Partition Name** field, type a unique name for the partition.
An example of a partition name is `Spanned_VIP`.
5. Type a description of the partition in the **Description** field.
This field is optional.

6. For the **Device Group** setting, choose an action:

Action	Result
Retain the default value.	Choose this option if you want the folder corresponding to this partition to inherit the value of the device group attribute from folder <code>root</code> .
Clear the check box and select the name of a device group.	Choose this option if you do not want the folder corresponding to this partition to inherit the value of the device group attribute from folder <code>root</code> .

7. For the **Traffic Group** setting, choose an action:

Action	Result
Retain the default value.	Choose this option if you want the folder corresponding to this partition to inherit the value of the traffic group attribute from folder <code>root</code> .
Clear the check box and select the name of a traffic group.	Choose this option if you do not want the folder corresponding to this partition to inherit the value of the traffic group attribute from folder <code>root</code> .

8. Click **Finished**.

The new partition appears in the partition list.

Assigning roles to a user account

Before performing this task, ensure that you have a user role of Administrator or that you have a role of User Manager for the relevant partition.

You perform this task to change the user roles that are assigned to a user account. You can assign a different role for each partition to which the user has access. By default, the user role that the BIG-IP® system assigns to a user account on each partition is No Access.

Important: *If you are performing this task while the user is logged into the system through `tmsh`, the BIG-IP system terminates the user's `tmsh` session when the user subsequently issues another `tmsh` command. This behavior ensures that the user is notified of the change in permissions and that data integrity is maintained.*

1. Access the BIG-IP® Configuration utility.
2. In the upper-left corner of the screen, confirm that the **Partition** list is set to the partition in which the user account that you want to modify resides.
3. On the Main tab, click **System > Users**.
The BIG-IP system displays the list of user accounts that reside in the current partition and in partition `Common`. Note that all users except those with a user role of No Access have at least read access to partition `Common`.
4. In the User Name column, click the user account name.
5. For the **Partition Access** setting:
 - a) From the **Role** list to select a user role.
 - b) From the **Partition** list, select a partition name.
 - c) Click the **Add** button.
A user role pertaining to a partition now appears in the box.
 - d) Repeat these steps for each partition to which you want to assign a role for this user.

Partition Access

Role:

Certificate Manager

Partition:

Common

Add

Role	Partition
User Manager	USERS
Certificate Manager	Common

Edit

Delete

Figure 23: Granting partition access to a BIG-IP user account

After you configure this setting, one or more role-partition combinations are specified for assignment to this user account.

- 6. Click the **Update** button.

Working with Single Configuration Files

Overview: Working with single configuration files

A *single configuration file (SCF)* is a flat, text file that contains a series of `tmsh` commands, and the attributes and values of those commands, that reflect the configuration of the BIG-IP® system. Specifically, the SCF contains the local traffic management and TMOS® configuration of the BIG-IP system. This figure shows a small part of a sample SCF.

```
vlan external {
  tag 4093
  interfaces 1.3
}
vlan internal {
  tag 4094
  interfaces 1.10
}
pool dev_https3 {
  members {
    10.60.10.105:https{}
    10.60.10.106:https{}
  }
}
```

The single configuration file feature allows you to save the configuration of a BIG-IP system in a text file. You can then use the text file to easily replicate the configuration across multiple BIG-IP systems. This not only saves you time, but also allows you to create a consistent, secure, comprehensive local traffic management environment on your network.

tmsh commands for single configuration files (SCFs)

You use `tmsh` to manage a single configuration file (SCF). This table lists an overview of `tmsh` commands used to manage SCF files.

tmsh command	Description
<code>save sys config file [filename]</code>	<p>Saves a copy of the currently running configuration to an SCF.</p> <hr/> <p>Important: Saving a configuration to an SCF does not affect the running or stored configuration of the BIG-IP® system on which you run the command.</p> <hr/>
<code>load sys config file [filename]</code>	<p>Replaces or restores an SCF with a saved configuration. When you use this command, the system saves any previously running configuration to the <code>/var/local/scf/</code> directory, by default.</p>

tmsh command	Description
<code>load sys config default</code>	Restores the factory default settings of the configuration file, while retaining the management IP address and the administrator user name and password.

Task summary

You can perform three main tasks with respect to single configuration files.

Task list

Creating and saving an SCF

Loading an SCF onto a target BIG-IP system

Using an SCF to restore a BIG-IP system configuration

Creating and saving an SCF

You can use `tmsh` to create and save a single configuration file (SCF).

Important: The system configuration data contained in the text file includes any local device certificate and keys used to establish device trust between this system and the other devices in a BIG-IP device group. These certificates and keys are unencrypted in the text file and are not included in the `.tar` file.

Note: If you create an SCF file twice (on two different occasions), you can compare the contents of the two files.

1. Open the Traffic Management Shell (`tmsh`).

```
tmsh
```

2. Create and save an SCF.

```
save sys config file [filename]
```

Note: If you include the `.scf` extension in the file name, the system does not add an additional file extension.

The system gathers all of the commands that make up the running configuration, and then saves the configuration to a `.scf` file with the name you specify. The system also creates a corresponding `.tar` file. By default, the system stores these files in the `/var/local/scf` directory, but you can specify a different path if you prefer.

Loading an SCF onto a target BIG-IP system

You can use `tmsh` to load a single configuration file (SCF) on one BIG-IP® system that you created on another BIG-IP system (hereafter referred to as the target BIG-IP system). This saves you from having to recreate the configuration multiple times. Loading an SCF resets the running configuration with the values contained in the stored configuration.

Important: If you run a `load` command or restart the system before you save your changes to the stored configuration, you will lose any changes.

Note: To successfully load a configuration that you have replicated, make sure that no line of the configuration is longer than 4096 characters. If there are more than 4096 characters in a single line, the system reverts to the previous running configuration.

1. Open the Traffic Management Shell (tmsh).
`tmsh`
2. On the target BIG-IP system, load the saved SCF file.
`tmsh load sys config file [filename]`
 The system saves the stored configuration to a backup file named `/var/local/scf/backup.scf`, and then uses the configuration stored in the SCF that you are loading.
3. Use a text editor to open the SCF and edit any data that is unique to the target BIG-IP system, such as the management IP address.
4. Save the SCF to the target BIG-IP system.
`sys save config file [filename]`
 If a backup SCF already exists, the system appends a number to the name of the existing backup file, and then creates a new backup file. In the case of multiple load and save operations:
 - The first time the system backs up the running configuration during a load operation, the system names the backup file `/var/local/scf/backup.scf`.
 - The next time the system backs up the running configuration, the system renames the file from `/var/local/scf/backup.scf` to `/var/local/scf/backup-1.scf` and creates a new file named `/var/local/scf/backup.scf`.
 - If you run the `load` command a third time, the system renames the file from `/var/local/scf/backup-1.scf` to `/var/local/scf/backup-2.scf`, renames the `/var/local/scf/backup.scf` file to `/var/local/scf/backup-1.scf`, and again creates a new file named `/var/local/scf/backup.scf`.

Using an SCF to restore a BIG-IP system configuration

You can use `tmsh` to restore a BIG-IP® system configuration using either a specific single configuration file (SCF) or the factory default configuration.

1. Open the Traffic Management Shell (tmsh).
`tmsh`
2. Restore the system configuration using one of these options:
 - Restore a system to the factory default configuration by using `tmsh load sys config default`. This command retains the management IP and the assigned root and administrator passwords. When you use this command, the system first saves the running configuration in the `backup.scf` file, and then resets the local traffic management and the operating system configuration to the factory default settings by loading the factory default SCF (`/defaults/defaults.scf`).
 - Restore a system with values defined in the specified SCF by using `tmsh load sys config file [filename]`. When you use this command, the system first saves the running configuration in the `backup.scf` file, and then resets the running configuration to the values contained in the specified SCF.

Note: *You must run the `save sys config partitions all` command to save the running configuration in the stored configuration files.*

Configuring a One-Arm Deployment Using WCCPv2

Overview: Configuring a one-arm deployment using WCCPv2

In certain cases, it is not advantageous or even possible to deploy the BIG-IP® system inline. For example, in the case of a collapsed backbone where the WAN router and the LAN switch are in one physical device, you might not be able to deploy the BIG-IP system inline.

If you choose not to deploy the BIG-IP system inline, you can use a one-arm deployment. In a *one-arm deployment*, the BIG-IP system has a single (hence, one-arm) connection to the WAN router or LAN switch. The WAN router (or switch) redirects all relevant traffic to the BIG-IP system. In this configuration, the WAN router typically uses Web Cache Communication Protocol version 2 (WCCPv2) to redirect traffic to the BIG-IP system.




Figure 24: Network topology for a one-arm connection

The traffic flow sequence in this illustration is as follows:

1. The client initiates a session.
2. A WAN router redirects traffic to the BIG-IP system.
3. The BIG-IP1 processes traffic and sends it back to the WAN router.
4. The WAN router forwards traffic across the WAN.

About WCCPv2 redirection on the BIG-IP system

The BIG-IP® system includes support for Web Cache Communication Protocol version 2 (WCCPv2). WCCPv2 is a content-routing protocol developed by Cisco® Systems. It provides a mechanism to redirect traffic flows in real time. The primary purpose of the interaction between WCCPv2-enabled routers and a BIG-IP® system is to establish and maintain the transparent redirection of selected types of traffic flowing through those routers.

To use WCCPv2, you must enable WCCPv2 on one or more routers connected to the BIG-IP® system, and configure a service group on the BIG-IP system that includes the router information. The BIG-IP system

then receives all the network traffic from each router in the associated service group, and determines both the traffic to optimize and the traffic to which to apply a service.

In configuring WCCPv2 on a network, you define a *service group* on the BIG-IP system, which is a collection of WCCPv2 services configured on the BIG-IP system. A WCCPv2 *service* in this context is a set of redirection criteria and processing instructions that the BIG-IP system applies to any traffic that a router in the service group redirects to the BIG-IP system. Each service matches a service identifier on the router.

The following illustration shows a one-arm configuration on one side of the WAN and an inline (bridge) configuration on the other side.




Figure 25: Example of a one-arm configuration

Before you begin configuring an iSession connection

Before you configure an iSession™ connection on the BIG-IP® system, make sure that you have completed the following general prerequisites.

- You must have an existing routed IP network between the two locations where the BIG-IP devices will be installed.
- One BIG-IP system is located on each side of the WAN network you are using.
- The BIG-IP hardware is installed with an initial network configuration applied.
- F5® recommends that both units be running the same BIG-IP software version.
- The Application Acceleration Manager™ license is enabled.
- Application Acceleration Manager (AAM) is provisioned at the level **Nominal**.
- The management IP address is configured on the BIG-IP system.
- You must have administrative access to both the Web management and SSH command line interfaces on the BIG-IP system.
- If there are firewalls, you must have TCP port 443 open in both directions. Optionally, you can allow TCP port 22 for SSH access to the command line interface for configuration verification, but not for actual BIG-IP iSession traffic. After you configure the BIG-IP system, you can perform this verification from the Configuration utility (**Acceleration > Symmetric Optimization > Diagnostics**).

Task summary

To use WCCPv2 for traffic redirection, you configure a service group on the BIG-IP® system that includes at least one service. You also configure this service on the WCCPv2-enabled router connected to the BIG-IP system.

For optimization, you also need to configure the BIG-IP system on the other side of the WAN to complete the connection. The BIG-IP system on the other side of the WAN can be set up in either a one-arm or inline configuration.

Note: The example described in this implementation applies to the Cisco 3750 and Cat 6500 routers.

Prerequisites

Before you begin configuring WCCPv2 for traffic redirection, ensure that you have performed the following actions on the other devices in your network.

- The interface and associated VLAN have been configured on the router or switch. For instructions, refer to the Cisco documentation for your device.
- IP addresses have been assigned on the Cisco router or switch interface. Note the router identification address, which you will use when configuring WCCPv2 on the BIG-IP system.

Task list

Creating a VLAN for a one-arm deployment

Creating a self IP address for a one-arm deployment

Defining a route

Configuring WCCPv2

Verifying connectivity

Verifying WCCPv2 configuration for one-arm deployment

Creating an iSession connection

Validating iSession configuration in a one-arm deployment

Configuring the Cisco router for a one-arm deployment using WCCPv2

Viewing pertinent configuration details from the command line

Creating a VLAN for a one-arm deployment

For a one-arm deployment, you create only one VLAN on the BIG-IP® system, because the system has only a single connection to the WAN router or switch.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. Click **Create**.
The New VLAN screen opens.
3. In the **Name** field, type `wan`.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.
The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. From the **Customer Tag** list:

- a) Retain the default value of **None** or select **Specify**.
- b) If you chose **Specify** in the previous step, type a numeric tag, between 1-4094, for the VLAN.

The customer tag specifies the inner tag of any frame passing through the VLAN.

6. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number or trunk name.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) If you specified a numeric value for the **Customer Tag** setting and from the **Tagging** list you selected **Tagged**, then from the **Tag Mode** list, select a value.
 - d) Click **Add**.
 - e) Repeat these steps for each interface or trunk that you want to assign to the VLAN.
7. If you want the system to verify that the return route to an initial packet is the same VLAN from which the packet originated, select the **Source Check** check box.
8. In the **MTU** field, retain the default number of bytes (**1500**).
9. Configure the sFlow settings or retain the default values.
10. Click **Finished**.
The screen refreshes, and displays the new VLAN in the list.

Creating a self IP address for a one-arm deployment

A VLAN must be configured before you create a self IP address.

This self IP address is the local endpoint for the iSession™ connection.

1. On the Main tab, click **Network > Self IPs**.
2. Click **Create**.
The New Self IP screen opens.
3. In the **Name** field, type a descriptive name for the self IP address, for example `onearm`.
4. In the **IP Address** field, type an IP address that is not in use and resides on the `wan` VLAN you created.
In the example shown, this is `10.150.3.1`.
5. In the **Netmask** field, type the network mask for the specified IP address.
For example, you can type `255.255.255.0`.
6. From the **VLAN/Tunnel** list, select **wan**.
7. From the **Port Lockdown** list, select **Allow None**.
This selection avoids potential conflicts (for management and other control functions) with other TCP applications. However, to access any of the services typically available on a self IP address, select **Allow Custom**, so that you can open the ports that those services need.
8. In the **Traffic Group** field, clear the check box, and select **traffic-group-local-only (non-floating)** from the drop-down menu.
9. Click **Finished**.
The screen refreshes, and displays the new self IP address.

The self IP address is assigned to the external (WAN) VLAN.

The screenshot shows the 'Properties' screen for a self IP address. The breadcrumb navigation at the top reads 'Network >> Self IPs >> clientside'. Below the breadcrumb is a 'Properties' tab. The main area is titled 'Configuration' and contains a table with the following fields and values:

Name	clientside
Partition / Path	Common
IP Address	10.150.3.1
Netmask	255.255.255.0
VLAN / Tunnel	wan
Port Lockdown	Allow None
Traffic Group	<input type="checkbox"/> Inherit traffic group from current partition / path traffic-group-local-only (non-floating)

At the bottom of the configuration area are three buttons: 'Update', 'Cancel', and 'Delete'.

Figure 26: Example of the Properties screen for the self IP address you created

Use this self IP address on the WAN Optimization Quick Start screen for the **WAN Self IP Address**, which is the local endpoint for the iSession connection.

Defining a route

You must define a route on the local BIG-IP® system for sending traffic to its destination. In the example shown, the route defined uses the default gateway to send traffic to the router.

1. On the Main tab, click **Network > Routes**.
2. Click **Add**.
The New Route screen opens.
3. In the **Name** field, type `default-gateway`.
4. In the **Destination** field, type the IP address `0.0.0.0`.
An IP address of `0.0.0.0` in this field indicates that the destination is a default route.
5. In the **Netmask** field, type `0.0.0.0`, the network mask for the default route.
6. From the **Resource** list, select **Use Gateway**.
The gateway represents a next-hop or last-hop address in the route.
7. For the **Gateway Address** setting, select **IP Address** and type an IP address. In the example shown, this is `10.150.3.254`.

Configuring WCCPv2

To configure traffic redirection using WCCPv2 for a one-arm deployment, follow these steps on the BIG-IP® system. This implementation specifies the Layer 2 (L2) method of traffic forwarding and mask assignment as the load-balancing method for a WCCPv2 service.

Note: The values you select for **Redirection Method**, **Return Method**, and **Traffic Assign** are automatically selected by the Cisco router or switch, provided that the Cisco device supports these settings.

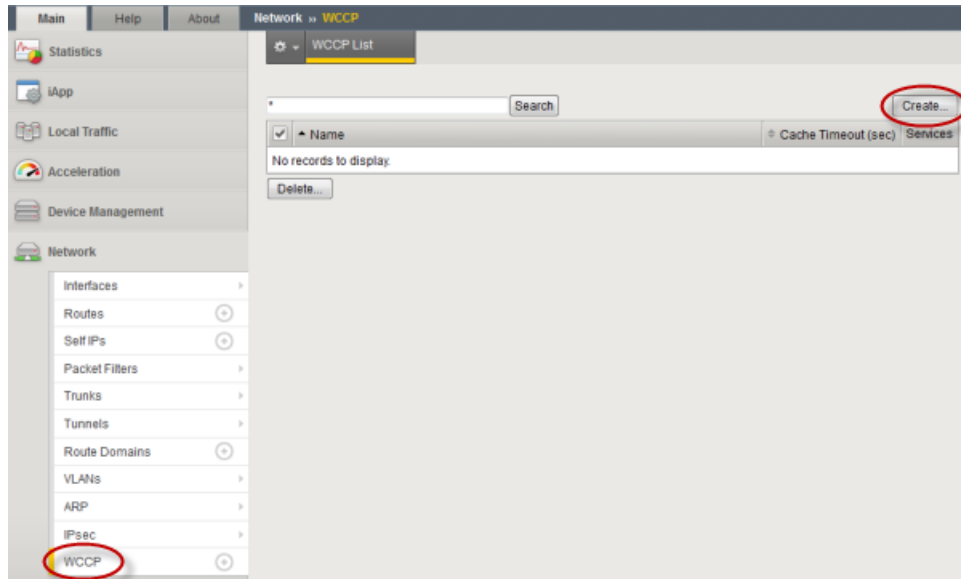


Figure 27: Example showing browser interface for configuring WCCP

1. On the Main tab of the BIG-IP® system user interface, click **Network > WCCP**.
2. Click the **Create** button.
The New WCCP List screen opens.
3. In the **Service Group** field, type a name for the service group, for example, `service-wccp`.
4. In the **Service** field, type a service group identifier, which is a number between 51 and 255.
This number must match the service ID you configure on the Cisco router. In the illustration shown, this number is 75.
5. From the **Port Type** list, select **Destination**.
If you specify a port in the **Port List**, this setting specifies the port on which the server listens for incoming traffic that has been redirected by WCCP. For best results, select **Destination**, even if you do not specify a port.
6. From the **Redirection Method** list, select **L2**.
This setting specifies the method the router uses to redirect traffic to the BIG-IP system. Typically, L2 has a faster throughput rate than GRE, but GRE traffic has the advantage that it can be forwarded by a Layer-3 router. This example uses **L2**.

Note: The router or switch uses the same redirection method, if supported.

7. From the **Return Method** list, select **L2**.
This setting specifies the method the BIG-IP system uses to return pass-through traffic to the router. Typically, L2 has a faster throughput rate than GRE, but GRE traffic has the advantage that it can be forwarded by a Layer-3 router. This example uses **L2**.

Note: The router or switch uses the same return method, if supported.

8. From the **Traffic Assign** list, select **Mask**.

This setting specifies whether load balancing is achieved by a hash algorithm or a mask. This example uses a mask.

Note: *The router or switch uses the same setting, if supported.*

9. In the **Routers** field, type the IP address of the Cisco router, and click **Add**.

In the illustration shown, this is 10.150.3.254.

Important: *Do not use a secondary IP address for the Cisco router or switch.*

10. In the **Port List** field, select an application, or leave it blank to indicate all ports.

11. For the **Router Identifier** setting, type the Router Identifier IP address of the router.

If you do not know the Router Identifier IP address, consult the Cisco documentation that applies to the router or switch you are using.

12. In the **Client ID** field, type the IP address of the VLAN that connects to the Cisco router.

In the illustration shown, this is 10.150.3.1.

13. Click **Finished**.

The BIG-IP is configured for WCCPv2 traffic redirection in a one-arm deployment. The completed screen looks similar to the following example.

Network » WCCP » New WCCP Service...

WCCP Group

Service Group: serv-wccp Select...

Cache Timeout (sec): 10

Configuration

Service: 75

Priority: 100

IP Protocol: TCP

Port Type: Destination

Weight: 50

Redirection Method: L2

Return Method: L2

MD5 Password:

Traffic Assign: Mask

Resources

Routers

Address: 10.150.3.254

Add

10.150.3.254

Remove Edit Up Down

Port List

Service Port: Select...

Add

Remove Edit Up Down

Router Identifier

Address: 192.168.3.161

Add

192.168.3.161

Remove Edit Up Down

Client ID

10.150.3.1

Cancel Repeat Finished

Figure 28: Example of completed configuration screen

Verifying connectivity

Important: Use this task as a checkpoint before proceeding with the one-arm setup.

You can verify connectivity from the command-line interface.

1. Ping the router interface using the command-line access to the BIG-IP® system.
2. Use TCPdump on TCP traffic between the servers at both sites to verify that TCP packets are redirected when you initiate TCP traffic.
3. Review the log `/var/log/wccpd.log` and look for the `SESSION up` message.

The following example is an excerpt from the log of a one-arm configuration.

```
Aug  2 17:26:18 clientside3600 notice router_ip 10.150.3.254
Aug  2 17:26:18 clientside3600 notice ports: 0,0,0,0,0,0,0,0,
Aug  2 17:26:18 clientside3600 notice tunnel_remote_addr: 192.31.3.161
Aug  2 17:26:18 clientside3600 notice
Aug  2 17:26:18 clientside3600 notice wccpd-1[1db1:f73f46d0]
WccpMcpInterface.cpp:113 :
Aug  2 17:26:18 clientside3600 notice wccpd-1[1db1:f73f46d0] WccpApp.cpp:208
: Failover status active 0
Aug  2 17:26:18 clientside3600 notice wccpd-1[1db1:f73f46d0] WccpApp.cpp:208
: Failover status active 1
Aug  2 17:26:18 clientside3600 notice wccpd-1[1db1:f73f46d0]
ServiceGroup.cpp:194 : Sending Wccp Capabilities Service group 75, Forwarding
Type: L2, Return Type: L2, Assignment Type: MASK
Aug  2 17:26:18 clientside3600 notice wccpd-1[1db1:f73f46d0]
ServiceGroup.cpp:468 : Final Wccp Capabilities Service group 75, Redirection:
L2, Return: L2, Traffic Assign: MASK
Aug  2 17:26:18 clientside3600 notice wccpd-1[1db1:f73f46d0]
ServiceGroup.cpp:615 : SESSION up
```

Verifying WCCPv2 configuration for one-arm deployment

You can use the command line interface to verify the WCCPv2 configuration on the BIG-IP® system.

1. Log on to the command-line interface using the root account.
2. At the command prompt, type `tmsh list net wccp`, and verify the WCCP values you configured. A listing similar to the following appears.

```
net wccp server-wccp
services
  75
    port-type dest
    redirection-method l2
    return-method l2
    routers { 10.150.1.254 }
    traffic-assign mask
    tunnel-local-address 10.150.3.1
    tunnel-remote-addresses { 10.150.2.1 }
```

Creating an iSession connection

You cannot view the Quick Start screen until you have defined at least one VLAN and at least one self IP on a configured BIG-IP® system that is provisioned for symmetric optimization.

Use the Quick Start screen to set up symmetric optimization for a one-arm deployment.

1. Log in to the BIG-IP system that you want to configure.
The default login value for both user name and password is `admin`.

2. On the Main tab, click **Acceleration > Quick Start > Symmetric Properties**.
3. In the **WAN Self IP Address** field, type the local endpoint IP address.
In the example shown, this is 10.150.3.1.
4. Verify that the **Discovery** setting is set to **Enabled**.
If you disable the **Discovery** setting, or discovery fails, you must manually configure any remote endpoints and advertised routes.
5. In the **Select VLANs** field, select the wan VLAN for both the **LAN VLANs** and **WAN VLANs** settings.
You select only one VLAN, because the system has only a single connection to the WAN router or switch.
6. Click **Apply**.

This example shows a completed Quick Start screen.

The screenshot displays the 'Quick Start: Symmetric Properties' configuration interface. At the top, there are tabs for 'Quick Start' and 'Deploy Applications'. The 'Local Endpoint' section includes a 'WAN Self IP Address' field with the value '10.150.3.1' and a 'Discovery' dropdown set to 'Enabled'. Below this, the 'Select VLANs' section has two rows: 'LAN VLANs' and 'WAN VLANs'. Each row has 'Selected' and 'Available' lists. For 'LAN VLANs', '/Common/lan' is selected. For 'WAN VLANs', '/Common/wan' is selected. The 'Authentication' section shows 'Outbound iSession to WAN' set to 'serverssl' and 'Inbound iSession from WAN' set to 'wom-default-clientssl'. The 'IP Encapsulation' section shows 'IP Encapsulation Type' set to 'None'. An 'Apply' button is located at the bottom left.

Figure 29: Example of completed Quick Start screen

After you configure the iSession™ endpoints, use an iApp template to select the application traffic for optimization. Click **Acceleration > Quick Start > Deploy Applications**. Click **Create**, from the **Template** list select **f5.replication**, and follow the online instructions.

Validating iSession configuration in a one-arm deployment

At this point, you have finished configuring BIG-IP® systems at opposite sides of the WAN, and the systems have discovered their remote iSession™ endpoints.

Important: Use this task as a checkpoint to allow for troubleshooting before you complete the setup.

You can validate the configuration using the browser and command-line interfaces.

1. Run diagnostics to verify the configuration.
 - a) On the Main tab, click **Acceleration > Symmetric Optimization > Diagnostics**.
 - b) Next to **Diagnose WOM Configuration**, click **Run**.
 - c) Correct any configuration errors as indicated on the screen.
2. Transfer data between the servers at the two sites, and verify that the transfer was successful.
3. Using the command-line interface, enter `tmsh show wom remote-endpoint all`, and verify the remote endpoint IP address and the `STATE: Ready` message.
The following listing is an example of the results for this command.

```
-----
Remote endpoint: 10.150.2.1          □-----
-----
Status
  HOSTNAME: server_bridge3600.example.net
  MGMT ADDR: 192.X.X.X  VERSION: 11.4.0
  UUID: 195f:74a0:d242:eab6:57fe:c3a:c1d2:6e22
  enabled                                STATE: ready □-----
  BEHIND NAT: no
  CONFIG STATUS: none
  DEDUP CACHE: 43.5G
  REFRESH count: 0                      REFRESH timestamp: 12/31/12 16:00:00
  ALLOW ROUTING: enabled

-----
Endpoint Isession Statistic: _tunnel_data_10.150.2.1
-----
Connections                                Current  Maximum  Total
Connections OUT IDLE:                      0         0         0
Connections OUT ACTIVE:                    1         1         1
Connections IN ACTIVE:                     0         0         0
Direction                                Action    Raw      Opt
  Out (to WAN) bits      Deduplication    880     1.2K
  Out (to WAN) bits      Compression     1.2K     1.2K
Direction                                Action    Opt      Raw
  In (from WAN) bits     Decompression   273.9M   273.8M
  In (from WAN) bits     Deduplication   272.6M   272.5M
```

4. Using the browser interface, view the green status indicator on the Remote Endpoints screen.
5. On the Main tab, click **WAN Optimization > Dashboard**, and view the traffic optimization data.

Configuring the Cisco router for a one-arm deployment using WCCPv2

To configure traffic redirection using Web Cache Communication Protocol version 2 (WCCPv2) for a one-arm deployment, follow these steps on the Cisco router.

1. Configure the service ID that you configured on the BIG-IP® device.
 - a) Enable WCCP globally.
 - b) In Command mode, configure the service ID; for example, 75.In the example shown, the command line might look like the following.

```
(config)#ip wccp 75
```

2. Using the router interface that is connected to the client from which you want to redirect traffic, associate the VLAN with the service ID you configured.
- In the example shown, the command-line interface might look like the following.

```
(config)#interface vlan 254  
(config)#ip wccp 75 redirect in
```

The following listing is an example of the information displayed for a Cisco router configured to redirect traffic to the BIG-IP system using WCCPv2.

```
Clientside_Top_switch#sh run  
Building configuration...  
Current configuration : 4848 bytes  
version 12.2  
no service pad  
hostname Clientside_Top_switch  
!  
no aaa new-model  
switch 1 provision ws-c3750g-48ts  
system mtu routing 1500  
vtp mode transparent  
ip subnet-zero  
ip routing  
ip wccp 75  
!  
interface GigabitEthernet1/0/4  
  switchport access vlan 200  
  switchport mode access  
!  
interface GigabitEthernet1/0/5  
  switchport access vlan 100  
  switchport mode access  
!  
interface GigabitEthernet1/0/6  
!  
interface GigabitEthernet1/0/7  
  switchport access vlan 254  
  switchport mode access  
!  
interface Vlan1  
  ip address 192.31.3.161 255.255.255.0  
!  
interface Vlan100  
  ip address 10.15.3.254 255.255.255.0  
!  
interface Vlan200  
  ip address 10.15.2.254 255.255.255.0  
!  
interface Vlan254  
  ip address 10.15.1.254 255.255.255.0
```

```
ip wccp 75 redirect in
!
```

Viewing pertinent configuration details from the command line

You can view details of the BIG-IP® iSession™ configuration from the command line.

1. Log on to the command-line interface of the BIG-IP system using the root account.
2. At the command prompt, type `tmssh`.
3. At the command prompt, type `list all-properties`.

The following listing is an example of the pertinent information displayed for a one-arm configuration.

```
ltm profile tcp wom-tcp-lan-optimized {
  abc enabled
  ack-on-push enabled
  app-service none
  close-wait-timeout 5
  cmetrics-cache disabled
  congestion-control high-speed
  defaults-from tcp-lan-optimized
  deferred-accept disabled
  delay-window-control disabled
  delayed-acks disabled
  description none
  dsack disabled
  ecn disabled
  fin-wait-timeout 5
  idle-timeout 600
  init-cwnd 0
  init-rwnd 0
  ip-tos-to-client 0
  keep-alive-interval 1800
  limited-transmit enabled
  link-qos-to-client 0
  max-retrans 8
  md5-signature disabled
  md5-signature-passphrase none
  nagle enabled
  partition Common
  pkt-loss-ignore-burst 0
  pkt-loss-ignore-rate 0
  proxy-buffer-high 1228800
  proxy-buffer-low 98304
  proxy-mss disabled
  proxy-options disabled
  receive-window-size 65535
  reset-on-timeout enabled
  rfc1323 enabled
  selective-acks enabled
  selective-nack disabled
  send-buffer-size 65535
  slow-start disabled
  syn-max-retrans 3
  syn-rto-base 0
  tcp-options none
  time-wait-recycle enabled
  time-wait-timeout 2000
  verified-accept disabled
  zero-window-timeout 20000
}
```

```

ltm profile tcp wom-tcp-wan-optimized {
    abc enabled
    ack-on-push disabled
    app-service none
    close-wait-timeout 5
    cmetrics-cache enabled
    congestion-control high-speed
    defaults-from tcp-wan-optimized
    deferred-accept disabled
    delay-window-control disabled
    delayed-acks disabled
    description none
    dsack disabled
    ecn disabled
    fin-wait-timeout 5
    idle-timeout 600
    init-cwnd 0
    init-rwnd 0
    ip-tos-to-client 0
    keep-alive-interval 1800
    limited-transmit enabled
    link-qos-to-client 0
    max-retrans 8
    md5-signature disabled
    md5-signature-passphrase none
    nagle enabled
    partition Common
    pkt-loss-ignore-burst 8
    pkt-loss-ignore-rate 10000
    proxy-buffer-high 196608
    proxy-buffer-low 131072
    proxy-mss disabled
    proxy-options disabled
    receive-window-size 2048000
    reset-on-timeout enabled
    rfc1323 enabled
    selective-acks enabled
    selective-nack enabled
    send-buffer-size 2048000
    slow-start disabled
    syn-max-retrans 3
    syn-rto-base 0
    tcp-options none
    time-wait-recycle enabled
    time-wait-timeout 2000
    verified-accept disabled
    zero-window-timeout 300000
}
ltm virtual isession-virtual {
    app-service none
    auth none
    auto-lasthop default
    clone-pools none
    cmp-enabled yes
    connection-limit 0
    description none
    destination 10.150.3.1:any
    enabled
    fallback-persistence none
    gtm-score 0
    http-class none
    ip-protocol tcp
    last-hop-pool none
    mask 255.255.255.255
    mirror disabled
    nat64 disabled
    partition Common
    persist none
    pool none

```

```

profiles {
    isession {
        context clientside
    }
    wom-default-clientssl {
        context clientside
    }
    wom-tcp-lan-optimized {
        context serverside
    }
    wom-tcp-wan-optimized {
        context clientside
    }
}
rate-class none
rules none
snat none
source-port preserve
traffic-classes none
translate-address enabled
translate-port disabled
vlans none
vlans-disabled
}
net interface 1.1 {
    app-service none
    description none
    enabled
    flow-control tx-rx
    force-gigabit-fiber disabled
    mac-address 0:1:d7:79:9a:84
    media none
    media-active 1000T-FD
    media-fixed auto
    media-max 1000T-FD
    media-sfp auto
    mtu 1500
    prefer-port sfp
    stp enabled
    stp-auto-edge-port enabled
    stp-edge-port true
    stp-link-type auto
    vendor none
}
net route def {
    description none
    gw 10.150.3.254
    mtu 0
    network default
    partition Common
}
net self "clientside Self" {
    address 10.150.3.1/24
    allow-service none
    app-service none
    description none
    floating disabled
    inherited-traffic-group false
    partition Common
    traffic-group traffic-group-local-only
    unit 0
    vlan wan
}
net vlan wan {
    app-service none
    auto-lasthop default
    description none
    failsafe disabled
    failsafe-action failover-restart-tm

```

```

    failsafe-timeout 90
    interfaces {
        1.1 {
            app-service none
            untagged
        }
    }
    learning enable-forward
    mtu 1500
    partition Common
    source-checking disabled
    tag 4094
}
sys datastor {
    cache-size 1066
    description none
    disk enabled
    high-water-mark 90
    low-water-mark 80
    store-size 97152
}
sys disk application-volume datastor {
    logical-disk HD1
    owner datastor
    preservability discardable
    resizeable false
    size 97152
    volume-set-visibility-restraint none
}
sys management-route default {
    app-service none
    description none
    gateway 192.31.3.129
    mtu 1500
    network default
}
sys provision wom {
    app-service none
    cpu-ratio 0
    description none
    disk-ratio 0
    level nominal
    memory-ratio 0
}
sys provision woml {
    app-service none
    cpu-ratio 0
    description none
    disk-ratio 0
    level none
    memory-ratio 0
}
wom deduplication {
    description none
    dictionary-size 256
    disk-cache-size 97152
    enabled
    max-endpoint-count 1
}
wom endpoint-discovery {
    auto-save enabled
    description none
    discoverable enabled
    discovered-endpoint enabled
    icmp-max-requests 1024
    icmp-min-backoff 5
    icmp-num-retries 10
    max-endpoint-count 0
    mode enable-all
}

```

```

}
wom local-endpoint {
  addresses { 10.150.3.1 }
  allow-nat enabled
  description none
  endpoint enabled
  ip-encap-mtu 0
  ip-encap-profile { /Common/default-ipsec-policy-isession }
  ip-encap-type ipsec
  no-route passthru
  server-ssl serverssl
  snat none
  tunnel-port https
}
wom profile isession isession-http {
  adaptive-compression enabled
  app-service none
  compression enabled
  compression-codecs { deflate lzo bzip2 }
  data-encryption disabled
  deduplication enabled
  defaults-from isession
  deflate-compression-level 1
  description none
  mode enabled
  partition Common
  port-transparency enabled
  reuse-connection enabled
  target-virtual virtual-match-all
}
wom remote-endpoint 10.150.2.1 {
  address 10.150.2.1
  allow-routing enabled
  app-service none
  description none
  endpoint enabled
  ip-encap-mtu 0
  ip-encap-profile none
  ip-encap-type default
  origin manually-saved
  server-ssl none
  snat default
  tunnel-encrypt enabled
  tunnel-port https
}
wom server-discovery {
  auto-save enabled
  description none
  filter-mode exclude
  idle-time-limit 0
  ip-ttl-limit 5
  max-server-count 50
  min-idle-time 0
  min-prefix-length-ipv4 32
  min-prefix-length-ipv6 128
  mode enabled
  rtt-threshold 10
  subnet-filter none
  time-unit days
}

```

Implementation result

After you complete the tasks in this implementation, the BIG-IP® system is configured in a one-arm deployment. For symmetric optimization, you must also configure the other side of the WAN. The other BIG-IP deployment can be in bridge, routed, or one-arm mode.

Legal Notices

Legal notices

Publication Date

This document was published on May 9, 2016.

Publication Number

MAN-0379-10

Copyright

Copyright © 2012-2016, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

Trademarks

For a current list of F5 trademarks and service marks, see
<http://www.f5.com/about/guidelines-policies/trademarks/>.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at:
<http://www.f5.com/about/guidelines-policies/patents>

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual,

may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Index

A

- access control
 - configuring [152](#)
- access control properties
 - assigning to user groups [147](#)
- access control settings
 - saving [149](#)
- active-active configuration
 - described [19](#)
 - result of [29](#)
- Active Directory server information [142](#)
- active-standby configuration
 - creating [13](#)
 - described [13](#)
 - result of [18](#)
- address configuration for VLAN and BIG-IP
 - illustration [110](#)
- address exchange [17](#)
- administrative partitions
 - access to [152](#)
 - creating [91](#), [113](#), [151](#)
 - defined [151](#)
- administrative traffic
 - authenticating [142–143](#)
- administrative user accounts
 - configuring [14](#), [21](#)
- application load
 - and failover [55](#)
 - balancing [53](#)
- applications
 - creating [26–27](#)
- application traffic
 - isolating on network [91](#), [113](#), [151](#)
- ARP entries
 - populating manually for virtual network segments [69](#)
- AWS floating IP address [33](#), [43](#)

B

- bandwidth controller categories
 - adding [63](#)
- bandwidth controllers
 - compared with rate shaping [59](#)
- bandwidth control policies
 - about Max User Rate setting [61](#)
 - adding categories to [63](#)
 - and SNMP [59](#)
 - dynamic, about [61](#)
 - dynamic, adding to virtual server [64](#)
 - dynamic, classifying traffic [63](#)
 - dynamic, creating [62](#)
 - dynamic, prerequisites [61](#)
 - overview [59](#)
 - static, about [59](#)
 - static, adding to virtual server [60](#)
 - static, creating [60](#)
- base network components [13](#)

- BIG-IP main dashboard
 - customizing [11](#)
- BIG-IP system
 - provisioning [14](#), [21](#)
 - restoring SCF [157](#)
- BIG-IP system licenses [14](#), [21](#)

C

- CCLDAP, See remote server authentication
- Cert-LDAP, See remote server authentication
- Cisco router
 - configuring for one-arm deployment [169](#)
- cloud
 - about connectivity in [65](#)
- configsync
 - configuring for VIPRION systems [31](#), [41](#)
- config sync
 - See also configuration synchronization.
 - disabling for tunnels [77](#)
 - See also configuration synchronization.
 - config sync addresses [17](#), [24](#)
- See also configuration synchronization
 - specifying [34](#), [45](#)
 - See also configuration synchronization
- configuration data
 - creating [156](#)
 - importing [149](#)
 - loading [156](#)
 - restoring [157](#)
 - saving [156](#)
- configuration objects
 - and traffic groups [19](#)
- configuration synchronization
 - [17](#), [24](#)
 - and Setup utility [13](#)
 - syncing to group [28](#), [38](#), [40](#), [48](#), [51](#)
- connection mirroring
 - configuring [35](#), [45](#)
 - enabling [15](#), [22](#)
- connection mirroring addresses, See mirroring addresses
- connections
 - preserving on failover [35](#), [45](#)
- connectivity
 - checking [166](#)
- content
 - of LLDPDUs [100](#)
- control-plane logging, overview [103](#)
- customer administrator tasks
 - for deploying route domains within a vCMP guest [118](#)
- custom log filters
 - and disabling legacy system logging [107](#)
 - and disabling logging [107](#)
 - creating [106](#)

D

- DAG tunnel
 - specifying a port number [75](#)
- dashboard, BIG-IP main
 - customizing [11](#)
- dashboard windows
 - customizing [11](#)
- default gateway pools
 - creating [134](#)
- default traffic groups [13, 19](#)
- destinations
 - for logging [105](#)
 - for remote high-speed logging [105](#)
- device certificate configuration [14, 21](#)
- device discovery
 - for device trust [25, 36, 47](#)
 - of peer devices [18](#)
- device groups
 - configuring for VIPRION systems [31, 41](#)
 - creating [26, 37, 48](#)
- devices
 - and mirroring limit [35, 45](#)
- device trust
 - configuring for VIPRION systems [31, 41](#)
 - establishing [25, 36, 47](#)
- device utilization
 - about [53–54](#)
 - examples of [56](#)
- DSC deployment worksheet [33, 43](#)
- dynamic bandwidth control policies, *See* bandwidth control policies. [61, 63–64](#)

E

- external files
 - and iRules [137](#)
- external network
 - configuring [16, 23](#)
- external switches
 - incorporating into network [81](#)

F

- failover
 - and next-active device [53](#)
 - configuring for VIPRION systems [31, 41](#)
- failover addresses
 - configuring [17](#)
 - exchanging during discovery [18](#)
 - specifying [24](#)
- failover devices
 - targeting [36, 46, 54](#)
- failover IP addresses
 - specifying [39, 49](#)
- failover methods
 - specifying [15, 22](#)
- failover objects
 - associating with traffic groups [117](#)
- FDB entries
 - populating manually for virtual network segments [69](#)

- files
 - importing [137–138](#)
- floating IP address
 - for AWS [33, 43](#)
- floating IP addresses
 - configuring [15–16, 22–23](#)
- Force to Standby option [117](#)
- forwarding virtual servers
 - creating [134](#)
- FQDN (fully-qualified domain name) [14, 21](#)
- fully-qualified domain name (FQDN) [14, 21](#)

G

- global LLDP properties
 - configuring [100](#)
- guest-wide administrator tasks
 - for deploying route domains within a vCMP guest [113](#)

H

- HA load factor
 - about [55](#)
- HA load factors
 - examples of [56](#)
- hardware platforms
 - and failover [36, 46, 54](#)
 - heterogeneous [56](#)
- HA traffic load
 - about [55](#)
- health monitors
 - assigning to pools [82, 86, 94, 125, 131](#)
- high availability
 - and tunnels [77](#)
 - and VLANs [17, 24](#)
 - and VXLAN [77](#)
 - enabling [15, 22](#)
- high-speed logging
 - and server pools [104](#)
- high-speed remote logging
 - configuring [104](#)
- host administrator tasks
 - for deploying route domains within a vCMP guest [111](#)

I

- iApp applications
 - creating [26–27](#)
- ifile commands [137](#)
- iFiles
 - creating [138](#)
- illustration
 - of VLAN and BIG-IP address configuration [110](#)
- implementation results [121](#)
- imported files
 - listing [138](#)
- interfaces
 - and external VLAN configuration [16, 23](#)
 - and HA VLAN configuration [17, 24](#)
 - and internal VLAN configuration [15, 22](#)
 - tagging [82, 92](#)

- internal network
 - configuring 15, 22
- iRule commands
 - for iFiles 137
- iRule events 138–139
- iRules
 - and bandwidth control policies 63
 - and external files 137
 - and iFiles 138–139
- iSession
 - prerequisites for configuring 160
- iSession configuration
 - validating for one-arm 169

L

- L2 location records
 - populating manually 69
- LDAP server information
 - client certificate 143
 - specifying 142
- licenses
 - activating 14, 21
- link aggregation
 - creating 124, 130
 - described 123, 129
 - tasks for 124, 130
- Link Layer Discovery Protocol, *See* LLDP
- LLDPDU contents 101
- LLDP messages
 - sending and receiving 100
- LLDP properties
 - global 100
 - per interface 100
- LLDP protocol
 - overview 99
- LLDP tasks 100
- load-aware failover
 - about 36, 46, 53–54
 - task summary 54
- local trust domain
 - and device groups 26, 37, 48
 - defined 25, 36, 47
- log filters
 - and disabling system logging 107
 - creating 106
- logging
 - and destinations 105
 - and pools 104
 - and publishers 106
 - system alerts 106

M

- MAC addresses
 - adding to virtual network forwarding table 69
 - removing from virtual network forwarding table 69
- main BIG-IP dashboard
 - customizing 11
- management IP addresses
 - and ConfigSync 17, 24

- management port
 - configuring 14, 21
 - specifying for failover 24
- maximum rate of throughput, *See* bandwidth control policies
- message content
 - for LLDPDUs 100
- messages
 - transmitting and receiving 100
- mirroring
 - See also* connection mirroring
 - configuring for VIPRION systems 31, 41
 - See also* connection mirroring
- mirroring addresses
 - configuring 17
 - exchanging 18
 - specifying 24
- monitors
 - assigning to pools 82, 86, 94, 125, 131

N

- network
 - configuring one-arm deployment 159
- network failover
 - configuring 26, 37, 48
 - specifying 15, 22
- network traffic
 - about segmenting 109
- network virtualization
 - about tunneling types for 68
 - centralized vs. decentralized model 67
 - configuring BIG-IP system as gateway 65
 - creating tunnels for 66
- network virtualization tunnels
 - considerations for configuring 69
- next-active devices
 - controlling 36, 46, 54
- NVGRE
 - configuration example using tmsh 71
 - defined 67

O

- one-arm deployment
 - configuration result 176
 - configuring Cisco router 169
 - configuring WCCPv2 163
 - overview 159
 - using WCCPv2 161
 - verifying WCCPv2 configuration 167
 - viewing iSession configuration 171
- OSCP, *See* remote server authentication
- overlay networks
 - and VXLAN-GPE tunnels 78
 - and VXLAN tunnels 75
 - bridging traffic to physical network 76
 - using VXLAN 73
- OVSDB management component
 - about configuring VXLAN tunnels with 77
 - setting up 77

P

- packet filtering
 - enabling 135
- packet filter rules
 - creating 136
- partitions, *See* administrative partitions
- peer devices
 - and traffic groups 27
 - discovering 18
- performance monitors
 - assigning to pools 82, 86, 94, 125, 131
- persistence
 - and source address affinity 126, 132
- pools
 - creating 82, 86, 94, 119, 125, 131
 - for high-speed logging 104
- prerequisites
 - for configuring iSession 160
- prerequisite tasks
 - for deploying route domains within a vCMP guest 110
- publishers
 - creating for logging 106

Q

- Quick Start screen
 - configuring one-arm deployment 167

R

- RADIUS protocol
 - for remote server authentication 145
- rate shaping
 - compared with bandwidth controllers 59
- remote server authentication 141
- remote servers
 - and destinations for log messages 105
 - for high-speed logging 104
- route advertisement
 - configuring 96
- route domains
 - about 89
 - adding routes for 96
 - creating 93, 116
 - described 109
 - sample configuration 90
 - tasks for 91
- routes
 - and route domains 96
 - defining default 163
- routing
 - one-arm mode 159

S

- SCF file configuration
 - tasks for 156
- SCF files
 - and access control 149
- self IP addresses
 - and VLAN groups 127

- self IP addresses (*continued*)
 - and VLANs 110
 - assigning to traffic group 51
 - creating 127
 - creating for default route domains 118
 - creating for one-arm deployment 162
 - for default route domains 93
 - for external network 16, 23
 - for HA network 17, 24
 - for internal network 15, 22
 - removing from VLANs 126
- serial cable failover
 - specifying 15, 22
- servers
 - and destinations for log messages 105
 - and publishers for log messages 106
 - for high-speed logging 104
- session persistence 126, 132
- Setup utility
 - and base network 19
 - and base network configuration 15–17, 22–24
 - and device discovery 18
 - for active-standby configurations 13
 - using 20
- single configuration file (SCF)
 - copying 156
 - creating 156–157
 - loading 156
 - saving 156–157
 - tmsh commands 155
- SNAT translation addresses
 - and traffic groups 19
- SNMP
 - and bandwidth control policies 59
- source address affinity persistence 126, 132
- standby state
 - forcing to 28, 52
- static bandwidth control policies, *See* bandwidth control policies
- switch configuration
 - tasks for 81
- switches
 - incorporating into network 81
- Sync-Failover device groups
 - creating 26, 37, 48
 - illustrated 44
- synchronization, *See* configuration synchronization
- system log filters, customizing 106
- system logging
 - configuring 104
 - disabling 107
 - disabling legacy 107
 - overview 103
- system provisioning 14, 21

T

- TACACS+ protocol
 - for remote server authentication 146
- tagged interfaces
 - configuring 82, 92
 - described 123, 129
 - for web hosting 85

- TLVs [101](#)
- tmsh commands
 - for SCF files [155](#)
- traffic groups
 - activating [117](#)
 - and failover objects [117](#)
 - and iApp applications [26–27](#)
 - as defaults [19](#)
 - assigning to each administrative partition [118](#)
 - creating [31, 34, 41, 44, 50](#)
 - default name of [13](#)
 - forcing to standby state [28, 52, 117](#)
 - for remote devices [27–28, 52](#)
 - specifying load for [55](#)
- traffic load
 - balancing [53](#)
- traffic redirection
 - about WCCPv2 [159](#)
- transmission
 - of LLDPDUs [100](#)
- Transparent Ethernet Bridging
 - described [67](#)
- trunks
 - [123, 129](#)
 - creating [124, 130](#)
- trust domains
 - and local trust domain [25, 36, 47](#)
- trusted peers
 - and address exchange [17](#)
- trust relationships
 - establishing [13](#)
- tunnels
 - about configuring for VXLAN-GPE [78](#)
 - about static configuration for network virtualization [69](#)
 - about types used for network virtualization [68](#)
 - configuring for network virtualization [65–66](#)
 - creating for VXLAN [75](#)
 - creating for VXLAN-GPE [78](#)
 - viewing statistics for [73](#)

U

- unicast failover addresses, *See* unicast IP addresses
- unicast IP addresses
 - specifying for failover [24](#)
- untagged interfaces
 - for web hosting [86](#)
- user access control
 - tasks for [151](#)
- user accounts
 - authenticating [142–143](#)
- user authorization
 - granting [151](#)
- user groups
 - assigning access control properties to [147](#)
- user roles
 - for system access [152](#)

V

- vCMP guests
 - about using route domains in [109](#)

- vCMP host
 - creating VLANs on [112](#)
- virtual addresses
 - advertising routes for [96](#)
 - assigning to traffic group [50, 121](#)
- Virtual eXtended LAN, *See* VXLAN
- virtual IP addresses
 - and traffic groups [19](#)
- virtualized networks
 - about tunneling types for [68](#)
 - configuring BIG-IP system as gateway for [65](#)
 - configuring on BIG-IP system [66](#)
 - terminology defined [66](#)
- virtual servers
 - See also* forwarding virtual servers
 - adding dynamic bandwidth control policy [64](#)
 - adding static bandwidth control policies [60](#)
 - and source address persistence [126, 132](#)
 - creating [95, 120](#)
 - creating for web hosting [83, 87](#)
 - listening on VXLAN VNI command example [73](#)
 - See also* forwarding virtual servers
- VLAN
 - adding tagged interface [125, 131](#)
- VLAN and BIG-IP address configuration
 - illustration [110](#)
- VLAN groups
 - and self IP addresses [127](#)
 - creating [127](#)
 - creating for VXLAN [76](#)
- VLAN IDs
 - configuring [15–17, 22–24](#)
- VLANs
 - and self IP addresses [110](#)
 - assigning to guests [113](#)
 - creating [112](#)
 - creating for one-arm deployment [161](#)
 - creating with tagged interfaces [82, 92](#)
 - creating with untagged interfaces [86](#)
 - enabling SNAT automap [133](#)
 - modifying for DAG tunnel [75](#)
 - moving from partition Common [114](#)
 - removing self IP addresses [126](#)
- VLAN tags, *See* VLAN IDs
- VTEP entries
 - adding to virtual network forwarding table [69](#)
- VXLAN
 - about [73](#)
 - about configuring BIG-IP system as gateway [73](#)
 - adding virtual server command example [73](#)
 - and high availability [77](#)
 - bridging with L2 VLAN network [73](#)
 - configuration example using tmsh [72](#)
 - considerations for configuring [74](#)
 - creating tunnels for [75](#)
 - creating VLAN groups for [76](#)
 - multicast mode defined [67](#)
 - pre-requisites for configuring [74](#)
 - terminology defined [66](#)
 - unicast mode defined [67](#)
- VXLAN-GPE
 - creating tunnels for [78](#)

VXLAN-GPE tunnels
 about configuring [78](#)
VXLAN header
 about adding fields [78](#)

W

WCCPv2
 checking connectivity [166](#)
 configuring [163](#)
 configuring one-arm deployment [159](#), [161](#)
 description [159](#)
 verifying configuration [167](#)

web customers
 hosting [81](#)
web hosting
 tasks for [81](#), [85](#)
 with no external switch [85](#)
 with route domains [91](#)

X

x509 certificates
 for device trust [25](#), [36](#), [47](#)