

BIG-IP[®] Device Service Clustering: Administration

Version 13.0



Table of Contents

Introducing BIG-IP Device Service Clustering.....	7
What is BIG-IP device service clustering?.....	7
DSC components.....	7
 Working with DSC Devices.....	 9
About IP addresses for config sync, failover, and mirroring.....	9
About device properties.....	10
Viewing device properties.....	10
Specifying values for device properties.....	10
Device properties.....	10
About device status.....	11
Viewing possible status types for a device.....	11
Viewing the status of a device.....	11
Device status.....	11
 Managing Device Trust.....	 13
What is device trust?.....	13
Types of trust authority.....	13
About certificate signing authorities.....	13
About peer authorities.....	13
About subordinate non-authorities.....	13
Device identity.....	14
Device discovery in a local trust domain.....	14
Establishing device trust.....	14
Adding a device to the local trust domain.....	15
Troubleshooting tips for establishing trust.....	16
Managing trust authority for a device.....	16
Viewing status for device trust.....	17
 Working with Device Groups.....	 19
About Sync-Failover device groups.....	19
Sample Sync-Failover configuration.....	19
Sync-Failover device group considerations.....	20
Creating a Sync-Failover device group.....	20
Viewing a list of device groups.....	22
Viewing the members of a device group.....	22
Adding a device to a device group.....	22
A note about folders and overlapping device groups.....	23
 Managing Configuration Synchronization.....	 25
About configuration synchronization.....	25
Specifying an IP address for config sync.....	25
Viewing config sync status for the local device.....	26
Viewing config sync status for all device groups and members.....	26
Manually synchronizing the BIG-IP configuration.....	27
About automatic vs. manual sync.....	28
Enabling and disabling automatic sync.....	28

About full vs. incremental sync.....	29
Enabling and disabling full sync.....	30
Troubleshooting the config sync process.....	30
Sync status for device groups.....	31
Sync status for device group members.....	32
Advanced config sync properties for a device.....	34
Managing Failover.....	35
Introduction to failover.....	35
What triggers failover?.....	35
About IP addresses for failover.....	35
Specifying IP addresses for failover communication.....	36
About traffic groups.....	37
About pre-configured traffic groups.....	37
Failover objects and traffic group association.....	38
Before you configure a traffic group.....	38
Creating a traffic group.....	38
Adding members to a traffic group.....	40
Viewing a list of traffic groups for a device.....	40
Viewing the members of a traffic group.....	40
Traffic group properties.....	40
Active and standby states.....	41
Managing failover using HA groups.....	43
Creating an HA group.....	43
Enabling an HA group for an existing traffic group.....	44
Example of an HA group deployment.....	45
About next-active device selection.....	46
About using HA scores to pick the next-active device.....	47
About using a preferred device order list to pick the next-active device.....	50
About using traffic load to pick the next-active device.....	51
About MAC masquerade addresses.....	53
Managing Connection Mirroring.....	55
About connection mirroring.....	55
About connection mirroring for VIPRION systems.....	55
Connection mirroring and traffic groups.....	56
Configuration task summary.....	56
Specifying an IP address for connection mirroring.....	57
Configuring connection mirroring between VIPRION clusters.....	58
Enabling connection mirroring for TCP and UDP connections.....	58
Enabling connection mirroring for SNAT connections.....	58
Enabling mirroring of persistence records.....	59
Working with Folders.....	61
About folders on the BIG-IP system.....	61
About folder attributes for redundancy.....	61
About the root folder.....	62
Viewing redundancy attributes for the root folder.....	62
Configuring the traffic group attribute for the root folder.....	62
About using HA scores to pick the next-active device.....	63
Creating an Active-Standby Configuration Using the Setup Utility.....	65
Overview: Creating a basic active-standby configuration.....	65

Task summary.....	66
Licensing and provisioning the BIG-IP system.....	66
Configuring a device certificate.....	67
Configuring the management port and administrative user accounts.....	67
Enabling ConfigSync and high availability.....	68
Configuring the internal network.....	68
Configuring the external network.....	68
Configuring the network for high availability.....	69
Configuring a ConfigSync address.....	69
Configuring failover and mirroring addresses.....	70
Discovering a peer device.....	70
Implementation result.....	70
 Creating an Active-Active Configuration Using the Setup Utility.....	73
Overview: Creating a basic active-active configuration.....	73
Task summary.....	74
Licensing and provisioning the BIG-IP system.....	74
Configuring a device certificate.....	75
Configuring the management port and administrative user accounts.....	75
Enabling ConfigSync and high availability.....	75
Configuring the internal network.....	76
Configuring the external network.....	76
Configuring the network for high availability.....	77
Configuring a ConfigSync address.....	77
Configuring failover and mirroring addresses.....	77
Establishing device trust.....	78
Creating a Sync-Failover device group.....	79
Creating an iApp application for the local device.....	79
Creating a traffic group for a remote device.....	80
Creating an iApp application for a remote device.....	80
Forcing a traffic group to a standby state.....	81
Syncing the BIG-IP configuration to the device group.....	81
Implementation Results.....	82
 Creating an Active-Standby Configuration using the Configuration Utility.....	83
Overview: Creating an active-standby DSC configuration.....	83
About DSC configuration on a VIPRION system.....	83
DSC prerequisite worksheet.....	85
Task summary.....	86
Specifying an IP address for config sync.....	86
Specifying an IP address for connection mirroring.....	87
Establishing device trust.....	87
Creating a Sync-Failover device group.....	88
Syncing the BIG-IP configuration to the device group.....	89
Specifying IP addresses for failover communication.....	90
Syncing the BIG-IP configuration to the device group.....	91
Implementation result.....	91
 Creating an Active-Active Configuration using the Configuration Utility.....	93
Overview: Creating an active-active DSC configuration.....	93
About DSC configuration on a VIPRION system.....	93
DSC prerequisite worksheet.....	95
Configurations using Sync-Failover device groups.....	96
Task summary.....	96

Specifying an IP address for config sync.....	97
Specifying an IP address for connection mirroring.....	97
Establishing device trust.....	98
Creating a Sync-Failover device group.....	99
Syncing the BIG-IP configuration to the device group.....	99
Specifying IP addresses for failover communication.....	100
Creating a second traffic group for the device group.....	101
Assigning traffic-group-2 to a floating virtual IP address.....	101
Assigning traffic-group-2 to floating self IP addresses.....	102
Syncing the BIG-IP configuration to the device group.....	102
Forcing a traffic group to a standby state.....	102
Implementation result.....	103
Useful troubleshooting tools.....	105
Useful command-line troubleshooting tools.....	105
Legal Notices.....	107
Legal notices.....	107

Introducing BIG-IP Device Service Clustering

What is BIG-IP device service clustering?

Device service clustering, or DSC[®], is an underlying architecture within BIG-IP[®] Traffic Management Operation System[®] (TMOS[®]). DSC provides synchronization and failover of BIG-IP configuration data at user-defined levels of granularity, among multiple BIG-IP devices on a network. More specifically, you can configure a BIG-IP device on a network to:

- Synchronize some or all of its configuration data among several BIG-IP devices
- Fail over to one of many available devices
- Mirror connections to a peer device to prevent interruption in service during failover

If you have two BIG-IP devices only, you can create either an active-standby or an active-active configuration. With more than two devices, you can create a configuration in which multiple devices are active and can fail over to one of many, if necessary.

By setting up DSC, you ensure that BIG-IP configuration objects are synchronized and can fail over at useful levels of granularity to the most-available BIG-IP devices on the network. You also ensure that failover from one device to another, when enabled, occurs seamlessly, with minimal to no interruption in application delivery.

Note: *If you use the Setup utility to create a DSC 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 Config Sync/HA Utility**.*

DSC components

Device service clustering (DSC[®]) is based on a few key components.

Devices

A *device* is a physical BIG-IP[®] system or a virtual BIG-IP system (BIG-IP Virtual Edition or vCMP[®] guest). Each device member has a set of unique identification properties that the BIG-IP system generates. For device groups configured for failover, it is important that the device with the smallest capacity has the capacity to process all traffic groups. This ensures application availability in the event that all but one device in the device group become unavailable for any reason.

Device groups

A *device group* is a collection of BIG-IP devices that trust each other and can synchronize, and sometimes fail over, their BIG-IP configuration data. A *Sync-Failover* device group contains devices that synchronize configuration data and support traffic groups for failover purposes when a device becomes unavailable. The BIG-IP system supports either homogeneous or heterogeneous hardware platforms within a device group.

Important: *BIG-IP module provisioning must be equivalent on all devices within a device group. For example, module provisioning is equivalent when all device group members are provisioned to run BIG-IP[®] Local Traffic Manager[™] (LTM[®]) and BIG-IP[®] Application Security Manager[™] (ASM[™]) only. Maintaining equivalent module provisioning on all devices ensures that any device in the device group can process module-specific application traffic in the event of failover from another device.*

Traffic groups

A *traffic group* is a collection of related configuration objects (such as a virtual IP address and a self IP address) that run on a BIG-IP device and process a particular type of application traffic. When a BIG-IP device becomes unavailable, a traffic group can float to another device in a device group to ensure that application traffic continues to be processed with little to no interruption in service.

Device trust and trust domains

Underlying the success of device groups and traffic groups is a feature known as device trust. *Device trust* establishes trust relationships between BIG-IP devices on the network, through mutual certificate-based authentication. A *trust domain* is a collection of BIG-IP devices that trust one another and is a prerequisite for creating a device group for config sync and failover operations. The trust domain is represented by a special system-generated and system-managed device group named `device_trust_group`, which is used to synchronize trust domain information across all devices.

Folders

Folders are containers for the configuration objects on a BIG-IP device. For every administrative partition on the BIG-IP system, there is a high-level folder. At the highest level of the folder hierarchy is a folder named `root`. The BIG-IP system uses folders to affect the level of granularity to which it synchronizes configuration data to other devices in the device group.

Working with DSC Devices

About IP addresses for config sync, failover, and mirroring

Each device in a device group must contain *device connectivity* information, that is, the IP addresses that you define on the device for configuration synchronization (config sync), failover, and connection mirroring.

Config sync IP address

This is the IP address that you want other devices to use when synchronizing configuration objects to the local device.

By default, the system uses the self IP address of VLAN `internal`. This is the recommended IP address to use for config sync. You can, however, use a different self IP address for config sync.

Important: *A self IP address is the only type of BIG-IP system address that encrypts the data during synchronization. For this reason, you cannot use a management IP address for config sync.*

Failover IP addresses

These are the IP addresses that you want another device to use when failing over to the local device. You can specify two types of addresses: unicast and multicast.

For appliance platforms, specifying two unicast addresses should suffice. For VIPRION[®] platforms, you should also retain the default multicast address that the BIG-IP system provides.

The recommended unicast addresses for failover are:

- The self IP address that you configured for either VLAN `HA` or VLAN `internal`. If you created VLAN `HA` when you initially ran the Setup utility on the local device, F5 recommends that you use the self IP address for that VLAN. Otherwise, use the self IP address for VLAN `internal`.
- The IP address for the local management port.

Mirroring IP addresses

These are the IP addresses that you want another device to use for mirroring connections to the local device. You specify both a primary address, as well as a secondary address for the system to use if the primary address is unavailable. If you configured VLAN `HA`, the system uses the associated self IP address as the default address for mirroring. If you did not configure VLAN `HA`, the system uses the self IP address of VLAN `internal`.

Note: *You can only mirror connections between identical hardware or virtual platforms. On a VIPRION[®] system, you can mirror connections between blades within the cluster or between two separate clusters in a device group. If you are mirroring between clusters in a device group, then for VIPRION systems that are not provisioned for vCMP, each chassis must have the same number of blades in the same slot numbers. For vCMP systems, each guest must be assigned to the same number of blades in the same slot numbers, with the same number of cores allocated per slot.*

About device properties

From the local BIG-IP® device, you can view or configure the properties of any device within the local trust domain, including the local device.

Viewing device properties

On each member of the local trust domain, the BIG-IP® system generates a set of information. This information consists of properties such as the device name, serial number, and management IP address. By default, every BIG-IP device in the local trust domain has a set of device properties. You can use the BIG-IP Configuration utility to view these properties.

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.

Specifying values for device properties

Using the BIG-IP® Configuration utility, you can specify values for a few of the properties for a device. Device properties provide information about the device that you can refer to when needed.

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 specify properties.
This displays a table of properties for the device.
3. In the **Description** field, type a description of the device.
4. In the **Location** field, type a location for the device.
5. In the **Contact** field, type contact information for the device.
6. In the **Comment** field, type a comment about the device.
7. Click **Update**.

Device properties

The following table lists and describes the properties of a device.

Property	Description
Name	The name of the device, such as <code>siterequest</code> .
Description	A user-created description of the device.
Location	The location of the device, such as <code>Seattle, Bldg. 1</code>
Contact	The name of the person responsible for this device.
Comment	Any user-specified remarks about the device.
Hostname	The host name of the device, such as <code>www.siterequest.com</code>
IP address	The IP address for the management port.
Serial Number	The serial number of the device.
MAC Address	The MAC address for the management port.
Time Zone	The time zone in which the device resides.

Property	Description
Platform ID	An identification for the platform.
Platform Name	The platform name, such as BIG-IP 8900.
Software Version	The BIG-IP version number, such as BIG-IP 11.0.0.
Status	The status of the device, such as <code>Device is active</code>
Active Modules	The complete list of active modules, that is, the modules for which the device is licensed.

About device status

A BIG-IP® device can have any status shown in the following table.

Status	Description
Active	A minimum of one floating traffic group is currently active on the device. This status applies to Sync-Failover device groups only.
Forced offline	An administrator has intentionally made the device unavailable for processing traffic.
Offline	The device is unavailable for processing traffic.
Standby	The device is available for processing traffic, but all traffic groups on the device are in a standby state. This status applies to Sync-Failover device groups only.
Unknown	The status of the device is unknown.

Viewing possible status types for a device

You can view a list of possible status types for a device.

1. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
2. In the status column, click **Status**.
This displays a list of all possible status types for a device.

Viewing the status of a device

You can view the status of a device in a device group. Viewing the status of a device can help with troubleshooting or to verify that the devices in the device group are working properly.

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, locate the name of the device for which you want to view status.
3. In the Status column, view the status of the device.

Device status

At all times, the BIG-IP® system displays a specific status for each device in a device group.

Table 1: Possible statuses of a DSC® device

Device status	Description
Active	The device is available and is processing traffic on the network. If the device is a member of a Sync-Failover device group, this status indicates that at least one traffic group is active on the device.
Forced Offline	An authorized user has intentionally taken the device offline, usually for maintenance purposes.
Offline	The device is offline for a reason other than being forced offline by an authorized user.
Standby	The device is available but is not processing traffic on the network. This applies to devices in a Sync-Failover device group only, and all traffic groups on the device are Standby traffic groups only.
Unknown/Not Watched	The BIG-IP system cannot determine the status of the device. This status usually occurs when the device has not yet joined a device group.

Managing Device Trust

What is device trust?

Before any BIG-IP® devices on a local network can be members of a Sync-Failover device group to synchronize configuration data or fail over to one another, they must establish a trust relationship known as device trust. *Device trust* between any two BIG-IP devices on the network is based on mutual authentication through the signing and exchange of x509 certificates.

Devices on a local network that trust one another constitute a trust domain. A *trust domain* is a collection of BIG-IP devices that trust one another.

The trust domain is represented by a system-generated device group named `device_trust_group`, which the system uses internally to synchronize trust domain information across all devices. You cannot delete this special device group from the system.

Note: You can add devices to a local trust domain from a single device on the network. You can also view the identities of all devices in the local trust domain from a single device in the domain. However, to maintain or change the authority of each trust domain member, you must log in locally to each device.

Types of trust authority

Within a local trust domain, in order to establish device trust, you designate each BIG-IP® device as either a peer authority or a subordinate non-authority.

About certificate signing authorities

A *certificate signing authority* can sign x509 certificates for another BIG-IP device that is in the local trust domain. For each authority device, you specify another device as a peer authority device that can also sign certificates. In a standard redundant system configuration of two BIG-IP devices, both devices are typically certificate signing authority devices.

Important: For security reasons, F5 Networks recommends you limit the number of authority devices in a local trust domain to as few as possible.

About peer authorities

A *peer authority* is another device in the local trust domain that can sign certificates if the certificate signing authority is not available. In a standard redundant system configuration of two BIG-IP devices, each device is typically a peer authority for the other.

About subordinate non-authorities

A *subordinate non-authority device* is a device for which a certificate signing authority device signs its certificate. A subordinate device cannot sign a certificate for another device. Subordinate devices provide an additional level of security because in the case where the security of an authority device in a trust domain is compromised, the risk of compromise is minimized for any subordinate device. Designating

devices as subordinate devices is recommended for device groups with a large number of member devices, where the risk of compromise is high.

Device identity

The devices in a BIG-IP® device group use x509 certificates for mutual authentication. Each device in a device group has an x509 certificate installed on it that the device uses to authenticate itself to the other devices in the group.

Device identity is a set of information that uniquely identifies that device in the device group, for the purpose of authentication. Device identity consists of the x509 certificate, plus this information:

- Device name
- Host name
- Platform serial number
- Platform MAC address
- Certificate name
- Subjects
- Expiration
- Certificate serial number
- Signature status

Tip: From the *Device Trust: Identity* screen in the *BIG-IP Configuration* utility, you can view the x509 certificate installed on the local device.

Device discovery in a local trust domain

When a BIG-IP® device joins the local trust domain and establishes a trust relationship with peer devices, the device and its peers exchange their device properties and device connectivity information. This exchange of device properties and IP addresses is known as *device discovery*.

For example, if a device joins a trust domain that already contains three trust domain members, the device exchanges device properties with the three other domain members. The device then has a total of four sets of device properties defined on it: its own device properties, plus the device properties of each peer. In this exchange, the device also learns the relevant device connectivity information for each of the other devices.

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 > Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. 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.
5. Click **Retrieve Device Information**.
6. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
7. In the **Name** field, verify that the name of the remote device is correct.
8. Click **Add Device**.

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.

Adding a device to the local trust domain

Verify that each BIG-IP® device that is to be part of a local trust domain has a device certificate installed on it.

Follow these steps to log in to any BIG-IP® device on the network and add one or more devices to the local system's local trust domain.

***Note:** Any BIG-IP devices that you intend to add to a device group at a later point must be members of the same local trust domain.*

1. On the Main tab, click **Device Management > Device Trust > Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. 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.

5. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
6. In the **Name** field, verify that the name of the remote device is correct.
7. Click **Add Device**.

After you perform this task, the local device and the device that you specified in this procedure have a trust relationship and, therefore, are qualified to join a device group.

Troubleshooting tips for establishing trust

This table lists possible problems that might occur when you are attempting to add a BIG-IP® device to a local trust domain. Each problem shows a recommended action.

Problem	Recommended action
Another device with the same name already exists in the trust domain.	Change the name of the device that you are adding to the trust domain.
The version of BIG-IP software on the device does not match the version of the devices in the trust domain.	Make sure that the BIG-IP version on the device you are adding exactly matches the version on the devices in the trust domain, including the hotfix version (if any).
The exact time reported on the device you are adding is out of sync with the time on the other devices in the trust domain.	Make sure that you have a Network Time Protocol (NTP) server configured for the device.
There is no config sync address configured on the device.	On the device you are adding, configure a config sync IP address. We recommend that you specify the self IP address associated with the device's internal VLAN.

Managing trust authority for a device

You can use a Reset Device Trust wizard in the BIG-IP® Configuration utility to manage the certificate authority of a BIG-IP device in a local trust domain. Specifically, you can:

- Retain the current authority (for certificate signing authorities only).
- Regenerate the self-signed certificate for a device.
- Import a user-defined certificate authority. In this case, a typical scenario is to generate another signing certificate and key through another certificate authority (such as OpenSSL) and then import the certificate to the BIG-IP system. The BIG-IP system then uses the certificate and key to sign the certificate signing request (CSR) that the BIG-IP generates. The resulting certificate is used to establish trust with other devices in the trust domain.

Warning: If you reset trust authority on a certificate signing authority by retaining the authority of the device, you must subsequently recreate the local trust domain and the device group. If you reset trust authority on a subordinate non-authority, the BIG system removes the non-authority device from the local trust domain. You can then re-add the device as an authority or non-authority device.

1. On the Main tab, click **Device Management > Device Trust > Local Domain**.
2. In the Trust Information area of the screen, click **Reset Device Trust**.
3. Choose a certificate signing authority option, and then click **Update**.
The system prompts you to confirm your choice.

When you confirm your choice, the system changes the **Authority Type**.

Viewing status for device trust

For any BIG-IP® devices that have a trust relationship, the BIG-IP® system automatically puts these devices into a special Sync-Only device group for device trust and syncs the trust information. (BIG-IP devices in this device group can also be members of other device groups that you create.) If any trust issue occurs between devices in the trust device group, this indicates that you need to re-establish trust between two or more devices. In this case, the BIG-IP system displays a config sync status of `Changes Pending`. You can use the BIG-IP® Configuration utility to view the config sync status of this trust device group, which has a system-supplied device group name.

Note: *A device in a trust domain that is not a member of a Sync-Failover device group normally shows a status in the BIG-IP Configuration utility of `NOT WATCHED`. If you add the device to a Sync-Failover device group, the status shows as either `In Sync` or `Changes Pending`.*

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, click the arrow next to 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.

Working with Device Groups

About Sync-Failover device groups

A *Sync-Failover* device group contains devices that synchronize their configuration data and fail over to one another when a device becomes unavailable. A Sync-Failover device group supports a maximum of eight devices.

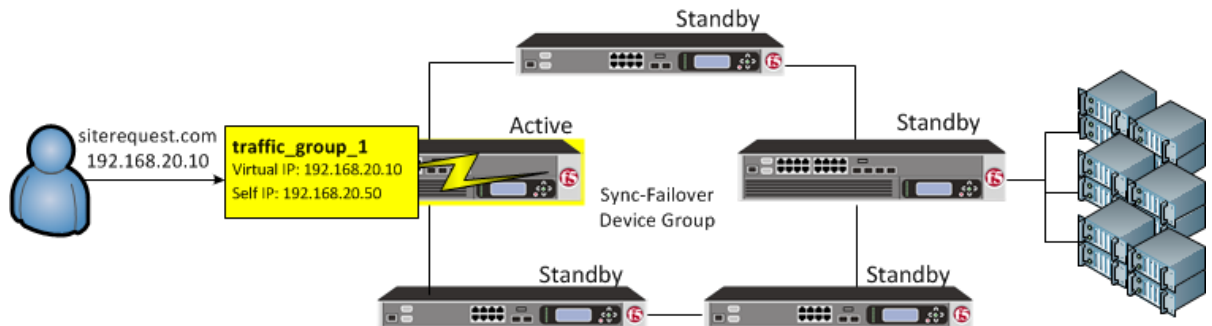


Figure 1: traffic_group_1 is active on a device in a Sync-Failover device group

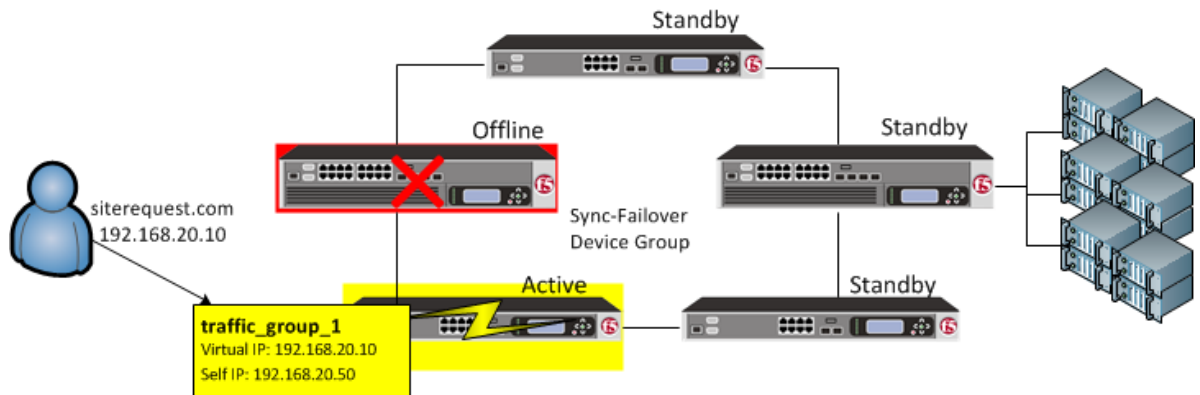


Figure 2: On failover, traffic_group_1 becomes active on another device in the Sync-Failover device group

For devices in a Sync-Failover group, the BIG-IP system uses both the device group and the traffic group attributes of a folder to make decisions about which devices to target for synchronizing the contents of the folder, and which application-related configuration objects to include in failover.

You can control the way that the BIG-IP chooses a target failover device. This control is especially useful when a device group contains heterogeneous hardware platforms that differ in load capacity, because you can ensure that when failover occurs, the system will choose the device with the most available resource to process the application traffic.

Sample Sync-Failover configuration

You can use a Sync-Failover device group in a variety of ways. This sample configuration shows two separate Sync-Failover device groups in the local trust domain. Device group A is a standard active-standby configuration. Prior to failover, only Bigip1 processes traffic for application A. This means that

Bigip1 and Bigip2 synchronize their configurations, and Bigip1 fails over to Bigip2 if Bigip1 becomes unavailable. Bigip1 cannot fail over to Bigip3 or Bigip4 because those devices are in a separate device group.

Device group B is also a standard active-standby configuration, in which Bigip3 normally processes traffic for application B. This means that Bigip3 and Bigip4 synchronize their configurations, and Bigip3 fails over to Bigip4 if Bigip3 becomes unavailable. Bigip3 cannot fail over to Bigip1 or Bigip2 because those devices are in a separate device group.

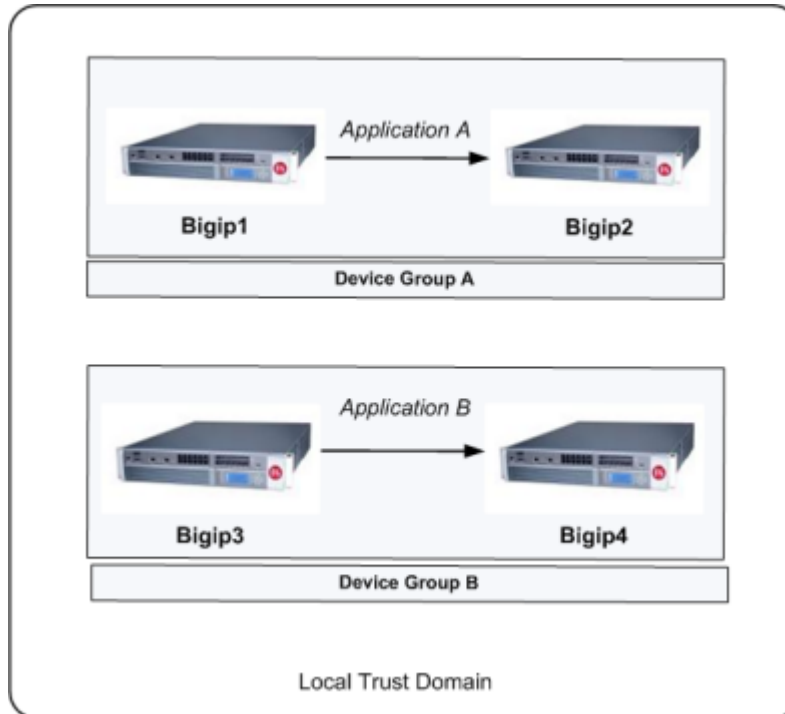


Figure 3: Sample Sync-Failover device groups in a trust domain

Sync-Failover device group considerations

The following configuration restrictions apply to Sync-Failover device groups:

- A specific BIG-IP device in a trust domain can belong to one Sync-Failover device group only.
- On each device in a Sync-Failover device group, the BIG-IP[®] system automatically assigns the device group name to the `root` and `/Common` folders. This ensures that the system synchronizes any traffic groups for that device to the correct devices in the local trust domain.
- The BIG-IP system creates all device groups and traffic-groups in the `/Common` folder, regardless of the partition to which the system is currently set.
- If no Sync-Failover device group is defined on a device, then the system sets the device group value that is assigned to the `root` and `/Common` folders to `None`.
- By default, on each device, the BIG-IP system assigns a Sync-Failover device group to any sub-folders of the `root` or `/Common` folders that inherit the `device group` attribute.
- You can configure a maximum of 127 floating traffic groups for a Sync-Failover device group.

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. In the **Name** field, type a name for the device group.
4. From the **Group Type** list, select **Sync-Failover**.
5. In the **Description** field, type a description of the device group.
This setting is optional.
6. From the **Configuration** list, select **Advanced**.
7. For the **Members** setting, 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.

8. From the **Sync Type** list:
 - Select **Automatic with Incremental Sync** when you want the BIG-IP system to automatically sync the most recent BIG-IP configuration changes from a device to the other members of the device group. In this case, the BIG-IP system syncs the configuration data whenever the data changes on any device in the device group.
 - Select **Manual with Incremental Sync** when you want to manually initiate a config sync operation. In this case, the BIG-IP system syncs the latest BIG-IP configuration changes from the device you choose to the other members of the device group. We strongly recommend that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.
 - Select **Manual with Full Sync** when you want to manually initiate a config sync operation. In this case, the BIG-IP system syncs the full set of BIG-IP configuration data from the device you choose to the other members of the device group. We strongly recommend that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.
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. 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 is the default value and 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.

11. In the **Link Down Time on Failover** field, use the default value of 0.0, or specify a new value.
This setting specifies the amount of time, in seconds, that interfaces for any external VLANs are down when a traffic group fails over and goes to the standby state. Specifying a value other than 0.0

for this setting causes other vendor switches to use the specified time to learn the MAC address of the newly-active device.

Important: This setting is a system-wide setting, and does not apply to this device group only. Specifying a value in this field causes the BIG-IP system to assign this value to the global bigdb variable `failover.standby.linkdowntime`.

12. Click **Finished**.

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

Viewing a list of device groups

You can perform this task when you want to display a list of the device groups of which the local device is a member. This list also displays other information such as the sync status of each device group and whether Auto Sync is enabled.

Note: Among this list of device groups is a special Sync-Only device group corresponding to the local trust domain. The BIG-IP system automatically creates this device group to internally sync trust information among the devices in the local trust domain, on an ongoing basis. You cannot delete this special device group.

On the Main tab, click **Device Management > Overview**.

After you perform this task, the list shows all device groups that include the local device as a member.

Viewing the members of a device group

You can list the members of a device group and view information about them, such as their management IP addresses and host names.

1. On the Main tab, click **Device Management > Device Groups**.
2. In the Group Name column, click the name of the relevant device group.

The screen shows a list of the device group members.

Adding a device to a device group

You must ensure that the device you are adding is a member of the local trust domain.

You can use this procedure to add a member to an existing device group.

1. On the Main tab, click **Device Management > Device Groups**.
2. In the Group Name column, click the name of the relevant device group.
3. In the Members 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. Use the Move button to move the host name to the **Selected** 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. If you are attempting to add a member to a Sync-Failover group and you do not see the member name in the list, it is possible that the device is already a member of another Sync-Failover device group. A device can be a member of one Sync-Failover group only.

4. Click **Update**.

5. On the Main tab, click **Device Management > Overview**.
6. In the Devices area of the screen, make sure that the device you are logged into is selected.
7. In the Sync Options area of the screen, click **Push the selected device configuration to the group**.

A note about folders and overlapping device groups

Sometimes when one BIG-IP® object references another, one of the objects gets synchronized to a particular device, but the other object does not. This can result in an invalid device group configuration.

For example, suppose you create two device groups that share some devices but not all. In the following illustration, Device A is a member of both Device Group 1 and Device Group 2.

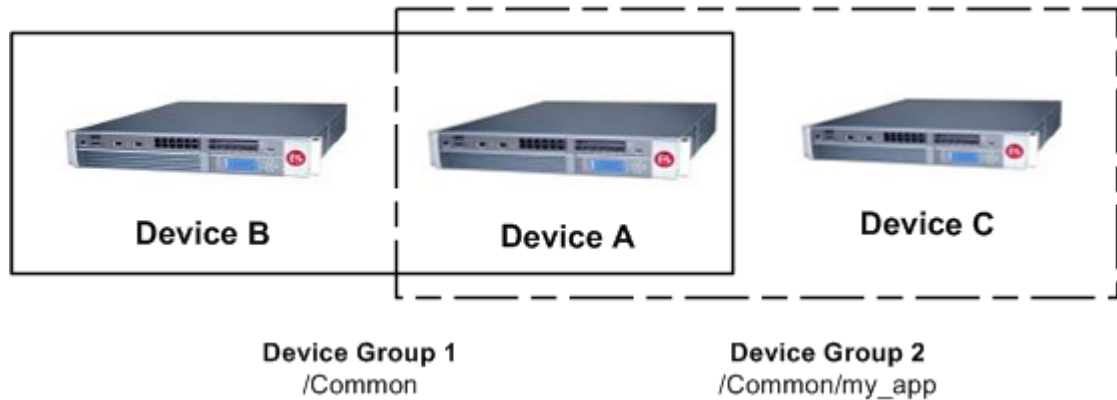


Figure 4: One device with membership in two device groups

Device Group 1 is associated with folder `/Common`, and Device Group 2 is associated with the folder `/Common/my_app`. This configuration causes Device A to synchronize all of the data in folder `/Common` to Device B in Device Group 1. The only data that Device A can synchronize to Device C in Device Group 2 is the data in the folder `/Common/my_app`, because this folder is associated with Device Group 2 instead of Device Group 1.

Now suppose that you create a pool in the `/Common/my_app` folder, which is associated with Device Group 2. When you create the pool members in that folder, the BIG-IP system automatically creates the associated node addresses and puts them in folder `/Common`. This results in an invalid configuration, because the node objects in folder `/Common` do not get synchronized to the device on which the nodes' pool members reside, Device C. When an object is not synchronized to the device on which its referenced objects reside, an invalid configuration results.

Managing Configuration Synchronization

About configuration synchronization

Configuration synchronization (also known as *config sync*) is the operation that the BIG-IP® system performs to propagate BIG-IP configuration changes, including device trust information, to all devices in a device group. BIG-IP devices that contain the same configuration data can work in tandem to more efficiently process application traffic on the network.

If you want to exclude certain devices from config sync, you simply exclude them from membership in that particular device group.

You can sync some types of data on a global level across all BIG-IP devices, while syncing other data in a more granular way, on an individual application level to a subset of devices. For example, you can set up a large device group to sync resource and policy data (such as iRules® and profiles) among all BIG-IP devices in a data center, while setting up a smaller device group for syncing application-specific data (such as virtual IP addresses) between the specific devices that are delivering those applications.

Whenever synchronization occurs, either automatically or manually, the BIG-IP system attempts to optimize performance by syncing only the data that changed since the last config sync operation.

Important: To synchronize configuration data among device group members, all members must be running the same version of the BIG-IP system software.

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. Near the top of the screen, click **ConfigSync**.
5. From the **Local Address** list, 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.

Viewing config sync status for the local device

You can use the BIG-IP® Configuration utility to view the config sync status of the local device relative to the other members of the device group. If you have configured the device group for manual synchronization, you can use the config sync status information to determine whether you need to perform a manual sync operation.

1. Display any BIG-IP Configuration utility screen.
2. In the upper left corner of the screen, view the status of the device group:
 - If the sync status is green (`In Sync`), the local device is synchronized with all device group members, and you do not need to perform a config sync operation.
 - If the sync status is yellow (`Changes Pending`), the BIG-IP configuration on the local device is out of sync with one or more device group members, or device trust is not fully established. You must therefore ensure that a config sync operation occurs for the relevant device group. If the **Automatic Sync** setting is enabled for the device group, the BIG-IP system synchronizes the configuration automatically, and no user action is required.

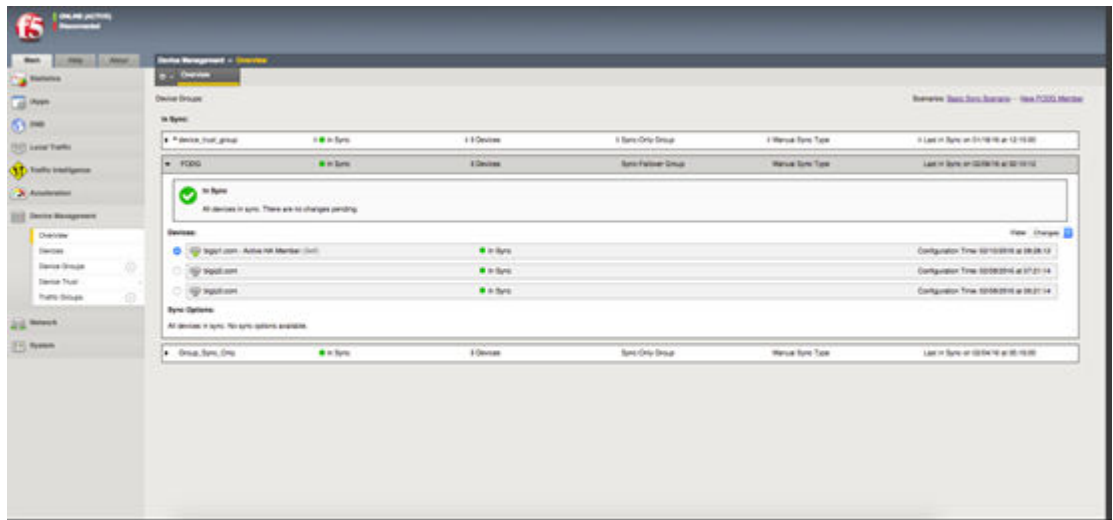
For more details, you can click the status, which displays the Overview screen. Using this screen, you can view a detailed message about the status, as well as the status of each device group member.

Viewing config sync status for all device groups and members

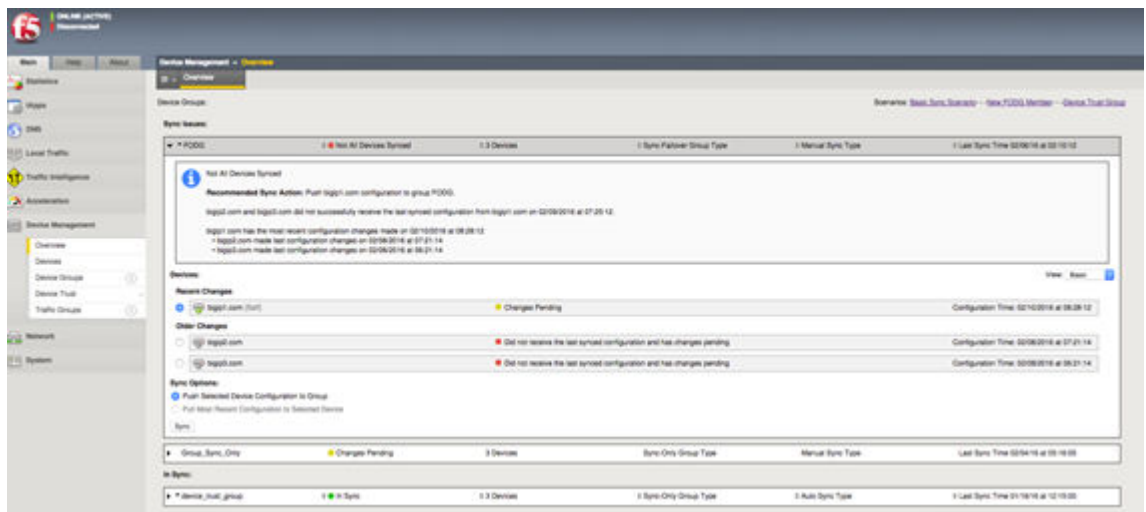
You can use the BIG-IP® Configuration utility to view the config sync status of any device group and each of its members, including the special Sync-Only device group for device trust. If the **Automatic Sync** setting is disabled for a device group, you can use the config sync status information to determine whether you need to do a manual sync operation.

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, click the arrow next to 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, view the sync status of each device:
 - If all devices show a sync status of green, the configurations of all device members are synchronized, and you do not need to perform a config sync operation. Here is a sample Overview screen showing a status of `In Sync`:



- If any device shows a sync status of **Changes Pending**, you must synchronize the configuration on that device to the other members of the device group. Here is a sample Overview screen showing a status of **Changes Pending**:



A status of **Changes Pending** for a device indicates that the device contains recent configuration changes that have not yet been synchronized to the other members of the device group.

Manually synchronizing the BIG-IP configuration

Before you perform this task, verify that device trust has been established and that all devices that you want to synchronize are members of a device group.

You perform this task when the automatic sync feature is disabled and you want to manually synchronize BIG-IP® configuration data among the devices in the device group. This synchronization ensures that any device in the device group can process application traffic successfully. You can determine the need to perform this task by viewing sync status in the upper left corner of any BIG-IP Configuration utility screen. A status of **Changes Pending** indicates that you need to perform a config sync within the device group.

Important: You can log into any device in the device group to perform this task.

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, click the arrow next to 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, choose a device.
4. In the Sync Options area of the screen, choose an option:

Option	Description
Push the selected device configuration to the group	Select this option when you want to synchronize the configuration of the selected device to the other device group members.
Pull the most recent configuration to the selected device	Select this option when you want to synchronize the most recent configurations of one or more device group members to the selected device.

5. Click **Sync**.

After you initiate a manual config sync, the BIG-IP® system compares the configuration data on the local device with the data on each device in the device group, and synchronizes the most recently-changed configuration data from one or more source devices to one or more target devices. Note that the system does not synchronize non-floating self IP addresses.

About automatic vs. manual sync

You can configure the BIG-IP® system to synchronization configuration data automatically, or you can manually initiate synchronization:

Automatic

Automatic synchronization (also known as *auto sync*) ensures that the BIG-IP system automatically synchronizes the configuration among device group members whenever you make a change to any one of those devices.

Manual

If you do not enable auto sync, you must manually synchronize the BIG-IP configuration among device group members to ensure that the devices remain in sync. With manual synchronization, the BIG-IP system notifies you whenever configuration data within the group has changed and therefore needs to be synchronized.

Enabling and disabling automatic sync

You can use the BIG-IP® Configuration utility to enable or disable automatic synchronization for device groups. When you enable automatic synchronization, a BIG-IP device in the device group automatically synchronizes its configuration data to the other members of the device group whenever its configuration data changes.

By default, the BIG-IP system syncs only the data that changed since the previous sync, rather than the entire set of configuration data.

1. On the Main tab, click **Device Management > Device Groups**.
2. In the Group Name column, click the name of the relevant device group.
3. For the **Automatic Sync** setting, select or clear the check box:

Action	Result
--------	--------

Select (Enable)	<p>Select the check box when you want the BIG-IP system to automatically sync configuration data to device group members whenever a change occurs. When you enable this setting, the BIG-IP system automatically syncs, but does not save, the configuration change on each device (this is the default behavior). To save the updated configuration on each device, you can log in to each device and, at the <code>tmsh</code> prompt, type <code>save sys config</code>. Alternatively, you can change the default behavior so that the system automatically saves configuration changes on target devices after an automatic config sync. You make this change by logging in to one of the devices in the device group and, at the <code>tmsh</code> prompt, typing <code>modify cm device-group name save-on-auto-sync true</code>.</p> <hr/> <p>Warning: Enabling the <code>save-on-auto-sync</code> option can unexpectedly impact system performance when the BIG-IP system automatically saves a large configuration change to each device.</p> <hr/> <p>Automatically saving configuration changes on target devices can provide a best practice for synchronizing configuration changes throughout a device group; however, in some instances, there is a potential to lose changes made on a local device while a remote peer device in the device group is rebooting. To prevent the possibility of an older configuration on a remote peer device from overwriting the latest changed configuration on a local device, complete the following steps.</p> <ol style="list-style-type: none"> 1. Disable automatic sync on all device groups that include the local device with the latest changed configuration. 2. Reboot the remote peer device. The device group indicates changes pending. 3. Change an object, such as the device description, on the local device if it appears in all device groups, or on a local device in each device group. 4. Manually sync the device group to each local device. 5. Enable automatic sync on all device groups.
Clear (Disable)	<p>Clear the check box when you want to disable automatic sync. When this setting is disabled, you must manually initiate each config sync operation. F5 Networks® recommends that you perform a config sync whenever configuration data changes on one of the devices in the device group. After you perform a manual config sync, the BIG-IP system automatically saves the configuration change on each device group member.</p>

4. Click Update.

After you enable automatic sync, the BIG-IP system automatically syncs future configuration changes to each device group member.

Whenever an automatic sync operation occurs, you must log in to each device group member and use the Traffic Management shell to save the configuration on the device. An alternative is to configure the `tmsh save-on-auto-sync` option for the device group.

About full vs. incremental sync

You can configure the BIG-IP® system to perform either full or incremental synchronization operations whenever a config sync is required:

Full

When you enable *full sync*, the BIG-IP system syncs the entire set of BIG-IP configuration data whenever a config sync operation occurs. You can only do a full sync operation if you have enabled manual sync; Full sync operations are not available when automatic sync is enabled.

Incremental

When you enable *incremental sync*, the BIG-IP system syncs only the changes that are more recent than those on the target device. The BIG-IP system accomplishes this by comparing 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. F5 networks recommends that you use incremental sync, for optimal performance. The incremental sync feature is a performance improvement feature and is the default value. You can use incremental sync with either automatic or manual sync operations.

You can also configure the cache size for any configuration changes slated for incremental sync. (This applies to incremental sync only.) For example, using the default cache size value of 1024, if you make more than 1024 KB worth of incremental changes, the system performs a full synchronization operation. Using incremental synchronization operations can reduce the per-device sync/load time for configuration changes.

Enabling and disabling full sync

You can enable or disable full synchronization for device groups. When you enable *full sync*, the BIG-IP® system syncs the entire set of configuration data whenever a sync operation occurs. When you disable full synchronization, the BIG-IP system performs *incremental synchronization*, which causes the system to sync only the changes that are more recent than the changes on the target device. The incremental sync feature is a performance improvement feature.

1. On the Main tab, click **Device Management > Device Groups**.
2. In the Group Name column, click the name of the relevant device group.
3. 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.

4. Click **Update**.

After you configure this feature, the BIG-IP system performs either a full or an incremental sync whenever a sync operation occurs.

Troubleshooting the config sync process

The BIG-IP® Configuration utility displays a number of different statuses and messages to help you diagnose and correct a config sync problem. These statuses and messages pertain to both device groups and individual device group members.

Sync status for device groups

At all times, the BIG-IP® system displays a specific sync status for each device group.

Table 2: Possible sync status for device groups

Sync Status	Summary Message	Explanation and Recommended Action
In Sync	All devices in the device group are in sync	All devices in the device group contain the current configuration. Recommended action: None.
Standalone	None.	The local trust domain contains one member only, which is the local device. Recommended action: None. You can optionally add other devices to the local trust domain.
Awaiting Initial Sync	The device group is awaiting the initial config sync.	All devices have been recently added to the device group and are awaiting an initial config sync. Recommended action: Sync any one of the devices to the device group.
Awaiting Initial Sync	<i>Device_name1, device_name2, etc.</i> awaiting the initial config sync	One or more of the devices in the device group has either not yet synchronized its data to the device group members or has not yet received a sync from another member. Recommended action: View the individual sync status of each device group member, and then sync the device with the most current configuration to the other devices.
Syncing	None.	A sync operation is in progress. Recommended action: None.
Changes Pending	Changes Pending	One or more devices in the device group has recent configuration changes that have not yet been synchronized to the other members of the device group. Recommended action: View the individual sync status of each device group member, and then sync the device with the most current configuration to the device group.
Changes Pending	There is a possible change conflict between <i>device_name1, device_name2, etc.</i>	There is a possible conflict among two or more devices because more than one device contains changes that have not been synchronized to the device group. Recommended action: View the individual sync status of each device group member, and then sync the device with the most current configuration to the device group.

Sync Status	Summary Message	Explanation and Recommended Action
Not All Devices Synced	<code>Device_name1, device_name2, etc.</code> did not receive last sync successfully.	<p>One or more of the devices in the device group does not contain the most current configuration.</p> <p>Recommended action: View the individual sync status of each device group member, and then sync the device with the most current configuration to the device group.</p>
Sync Failure	A validation error occurred while syncing to a remote device	<p>Because of a validation error, the named device was unable to accept a sync successfully.</p> <p>Recommended action: Review the <code>/var/log/ltn</code> file on the affected device.</p>
Unknown	The local device is not a member of the selected device group	<p>The device that you are logged into is not a member of the selected device group.</p> <p>Recommended action: Add the local device to the device group to view sync status for the device group.</p>
Unknown	Not logged into the primary cluster member	<p>The system cannot determine the sync status of the device group because you are logged in to a secondary cluster member instead of the primary cluster member. Pertains to VIPRION[®] systems only.</p> <p>Recommended action: Log out and then log in to the primary cluster member, using the primary cluster IP address.</p>
Unknown	Error in trust domain	<p>The trust relationships among devices in the device group are not properly established.</p> <p>Recommended action: On the local device, reset device trust and then re-add all relevant devices to the local trust domain.</p>
None.	<code>X</code> devices with <code>Y</code> different configurations	<p>The configuration time for two or more devices in the device group differs from the configuration time of the other device group members. This condition causes one of these status messages to appear for each relevant device:</p> <ul style="list-style-type: none"> <code>Device_name</code> awaiting initial config sync <code>Device_name</code> made last configuration change on <code>date_time</code> <p>Recommended action: Identify a device with the most current configuration and sync the device to the device group.</p>

Sync status for device group members

At all times, the BIG-IP[®] system displays a specific sync status for each device within a device group.

Table 3: Possible sync status for individual devices

Sync Status	Explanation and Recommended Action
Awaiting Initial Sync	<p>The local device is waiting for the initial ConfigSync. The device has not received a sync from another device and has no configuration changes to be synced to other members of the device group.</p> <p>Recommended action: Determine the device that has the latest or most desired configuration and sync the configuration from that device.</p>
Changes Pending	<p>The device has recent configuration changes that have not been synced to other device group members.</p> <p>Recommended action: Sync the device with the most recent configuration to the other members of the device group.</p>
Awaiting Initial Sync with Changes Pending	<p>This status indicates one of the following conditions:</p> <ul style="list-style-type: none"> The configuration on the device has changed since the device joined the device group. <p>Recommended action: Sync the device to the device group.</p> <ul style="list-style-type: none"> The device has not yet received a sync from another device but has configuration changes to be synced to other members of the device group. <p>Recommended action: Sync the device with the most recent configuration to this device.</p>
Does not have the last synced configuration, and has changes pending	<p>The device received at least one sync previously but did not receive the last synced configuration, and the configuration on the device has changed since the last sync.</p> <p>Recommended action: Determine the device that has the latest or most desired configuration and sync the configuration from that device.</p>
Disconnected	<p>The iQuery communication channel between the devices was terminated or disrupted. This may be a result of one of the following:</p> <ul style="list-style-type: none"> The disconnected device is not a member of the local trust domain. The disconnected device does not have network access to one or more device group members. <p>Recommended actions:</p> <ul style="list-style-type: none"> View the screens at Device Management > Device Trust to see if the disconnected device is a member of

Sync Status	Explanation and Recommended Action
	<p>the local trust domain, and if not, add the device to the domain.</p> <ul style="list-style-type: none"> Use the screen at Device Management Devices to view the specified config sync address of the disconnected device and determine whether the local device has a route to that address.

Advanced config sync properties for a device

A device in a device group has several advanced properties.

Property	Description
Current Commit Time	Indicates either the last time that a user updated the configuration locally, or, if the configuration on the device was synced from a remote device group member, the actual time that the synced configuration change was made on that remote device.
Current Commit Originator	Indicates the source of the most recent change to the configuration on the relevant device. More specifically, the CID originator is either: <ul style="list-style-type: none"> The relevant device itself (due to locally-made changes) Another device in the device group that synchronized a change to the relevant device
Previous Commit Time	Indicates the actual time that the synced configuration change was made on a remote device group member. Whenever a device in the device group syncs its configuration to the other device group members, the LSS time on each device is updated to reflect the Commit ID time of the configuration change on the device that initiated the sync operation.
Previous Commit Originator	Indicates the device that most recently performed a successful sync operation to the relevant device.
Device Last Sync Time	Indicates the last time that a sync was initiated or forced to or from the relevant device.
Device Last Sync Type	Indicates the type of sync. Possible values are: Manual Full Load, Manual Incremental, and Automatic.

Managing Failover

Introduction to failover

When you configure a Sync-Failover device group as part of device service clustering (DSC[®]), you ensure that a user-defined set of application-specific IP addresses, known as a *floating traffic group*, can fail over to another device in that device group if necessary. DSC failover gives you granular control of the specific configuration objects that you want to include in failover operations.

If you want to exclude certain devices on the network from participating in failover operations, you simply exclude them from membership in that particular Sync-Failover device group.

What triggers failover?

The BIG-IP system initiates failover of a traffic group according to any of several events that you define. These events fall into these categories:

System fail-safe

With *system fail-safe*, the BIG-IP system monitors various hardware components, as well as the heartbeat of various system services. You can configure the system to initiate failover whenever it detects a heartbeat failure.

Gateway fail-safe

With *gateway fail-safe*, the BIG-IP system monitors traffic between an active BIG-IP[®] system in a device group and a pool containing a gateway router. You can configure the system to initiate failover whenever some number of gateway routers in a pool of routers becomes unreachable.

VLAN fail-safe

With *VLAN fail-safe*, the BIG-IP system monitors network traffic going through a specified VLAN. You can configure the system to initiate failover whenever the system detects a loss of traffic on the VLAN and the fail-safe timeout period has elapsed.

HA groups

With an *HA group*, the BIG-IP system monitors the availability of resources for a specific traffic group. Examples of resources are trunk links, pool members, and VIPRION[®] cluster members. If resource levels fall below a user-defined level, the system triggers failover.

Auto-failback

When you enable *auto-failback*, a traffic group that has failed over to another device fails back to a preferred device when that device is available. If you do not enable auto-failback for a traffic group, and the traffic group fails over to another device, the traffic group remains active on that device until that device becomes unavailable.

About IP addresses for failover

Part of configuring a Sync-Failover device group is configuring failover. Configuring failover requires you to specify certain types of IP addresses on each device. Some of these IP addresses enable continual, high availability (HA) communication among devices in the device group, while other addresses ensure that application traffic processing continues when failover occurs.

The IP addresses that you need to specify as part of HA configuration are:

A local, static self IP address for VLAN HA

This unicast self IP address is the main address that other devices in the device group use to communicate continually with the local device to assess the health of that device. When a device in the device group fails to receive a response from the local device, the BIG-IP® system triggers failover.

A local management IP address

This unicast management IP address serves the same purpose as the static self IP address for VLAN HA, but is only used when the local device is unreachable through the HA static self IP address.

One or more floating IP addresses associated with a traffic group

These are the IP addresses that application traffic uses when passing through a BIG-IP system. Each traffic group on a device includes application-specific floating IP addresses as its members. Typical traffic group members are: floating self IP addresses, virtual addresses, NAT or SNAT translation addresses, and IP addresses associated with an iApp application service. When a device with active traffic groups becomes unavailable, the active traffic groups become active on other device in the device group. This ensures that application traffic processing continues with little to no interruption.

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. Near the top of the screen, click **Failover Network**.
5. Click **Add**.
6. From the **Address** list, select an IP address.

The unicast IP address you select depends on the type of device:

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

Platform**Action**

unicast failover addresses only (and not a multicast address), you must also select the existing, virtual static management IP addresses that you previously configured for all slots in the guest's virtual cluster. If you choose to select one or more unicast addresses and a multicast address, you do not need to select 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.

7. From the **Port** list, select a port number.

We recommend using port **1026** for failover communication.

8. 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.

9. 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**.

10. Click **Finished**.

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

About traffic groups

Traffic groups are the core component of failover. A *traffic group* is a collection of related configuration objects, such as a floating self IP address, a floating virtual IP address, and a SNAT translation address, that run on a BIG-IP[®] device. Together, these objects process a particular type of application traffic on that device.

When a BIG-IP[®] device goes offline, a traffic group floats (that is, fails over) to another device in the device group to make sure that application traffic continues to be processed with minimal interruption in service.

A traffic group is first active on the device you created it on. If you want an active traffic group to be active on a different device than the one you created it on, you can force the traffic group to switch to a standby state. This causes the traffic group to fail over to (and become active on) another device in the device group. The device it fails over to depends on how you configured the traffic group when you created it.

Note: A Sync-Failover device group can support a maximum of 127 floating traffic groups.

About pre-configured traffic groups

Each new BIG-IP[®] device comes with two pre-configured traffic groups:

traffic-group-1

A floating traffic group that initially contains any floating self IP addresses that you create on the device. If the device that this traffic group is active on goes down, the traffic group goes active on another device in the device group.

`traffic-group-local-only`

A non-floating traffic group that contains the static self IP addresses that you configure for VLANs `internal` and `external`. This traffic group never fails over to another device.

Failover objects and traffic group association

Any traffic group that you explicitly create on the BIG-IP® system is a floating traffic group. The types of configuration objects that you can associate with a floating traffic group are:

- Virtual IP addresses
- NATs
- SNAT translation addresses
- Self IP addresses
- Folders (such as an iApps® folder)

You can associate configuration objects with a traffic group in these ways:

- You can rely on the folders in which the objects reside to inherit the traffic group that you assign to the `root` folder.
- You can use the BIG-IP Configuration utility to directly associate a traffic group with an iApp application service, a virtual IP address, a NAT or SNAT translation address, or a floating self IP address.
- You can use the BIG-IP® Configuration utility to directly associate a traffic group with a folder.

Important: By default, floating objects that you create with the BIG-IP Configuration utility are associated with `traffic-group-1`. Non-floating objects are associated with `traffic-group-local-only`. You can change these associations by using the BIG-IP Configuration utility to change the traffic group that is associated with each floating IP address on the system.

Note: The only non-floating traffic group that resides on the system is the default non-floating traffic group named `traffic-group-local-only`.

Before you configure a traffic group

The following considerations apply to traffic groups:

- On each device in a Sync-Failover device group, the BIG-IP® system automatically assigns the default floating traffic group name to the `root` and `/Common` folders.
- The BIG-IP system creates all traffic groups in the `/Common` folder, regardless of the partition to which the system is currently set.
- Any traffic group named other than `traffic-group-local-only` is a floating traffic group.
- You can specify a floating traffic group on a folder only when the device group that is set on the folder is a Sync-Failover type of device-group.
- You can set a floating traffic group on only those objects that reside in a folder with a device group of type Sync-Failover.
- Setting the traffic group on a failover object to `traffic-group-local-only` prevents the system from synchronizing that object to other devices in the device group.

Creating a traffic group

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 task when you want to create a traffic group for a BIG-IP® device. You can perform this task on any BIG-IP device within the device group, and the traffic group becomes active on the local device.

Important: This procedure creates a traffic group with no members. After creating a traffic group, you must associate the traffic group with specific floating IP addresses such as a self IP address and a virtual address.

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 new traffic group.
 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. If you have created an HA group for monitoring trunk, pool, or VIPRION® cluster resources and for creating an HA health score, then from the **HA Group** list, select the HA group name.
This setting is optional.
 7. From the **Failover Method** list, select a failover method:
 - Choose **Failover to Device With Best HA Score** when you want the BIG-IP system to choose the next-active device based on an HA health score for the device. You can only choose this option when you have configured the **HA Group** setting to assign an existing HA group to this traffic group.
 - Choose **Failover using Preferred Device List and then Load Aware** when you want the BIG-IP system to choose the next-active device based on either an ordered list of devices or relative traffic group load.
 8. If you configured the **Failover Methods** setting with a value of **Failover using Preferred Device List and then Load Aware**, then configure the following settings. Otherwise, skip this step.
 - a) Select or clear the check box for the **Auto Failback** option.
 - b) If **Auto Failback** is selected, then in the **Auto Failback Timeout** field, type the number of seconds that you want the system to wait before failing back to the specified device. The range in seconds is from 0 to 300. The default is 60. A value of 40 to 60 seconds allows for state mirroring information to be re-mirrored for traffic groups.
 - c) For the **Failover Order** setting, in the **Load Aware** box, select a device name and, using the Move button, move the name to the **Preferred Order** 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 **Preferred 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 list is currently available or there are no devices in the list when failover occurs, BIG-IP system performs load-aware failover instead, using the **HA Load Factor** setting.
 - d) 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.
The BIG-IP system uses this value whenever it performs load-aware failover.
-

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

9. Click **Finished**.

You now have a floating traffic group with zero members

After creating the traffic group, you must add members to it. Possible members are floating IP addresses such as self IP addresses, virtual addresses, NAT or SNAT translation addresses, and iApp application services. Also, if you want the traffic group to become active on a device other than this local device, you can use the **Force to Standby** button. By forcing the traffic group into a standby state on the local device, you cause the traffic group to become active on another device.

Adding members to a traffic group

Before performing this task, verify that the traffic group exists on the BIG-IP system.

You perform this task to add members to a newly-created or existing traffic group. Traffic group members are the floating IP addresses associated with application traffic passing through the BIG-IP® system. Typical members of a traffic group are: a floating self IP address, a floating virtual address, and a floating SNAT translation address.

1. From the Main tab, display the properties page for an existing BIG-IP object, such as a self IP address or a virtual address.
For example, from the Main tab, click **Network > Self IPs**, and then from the Self IPs list, click a self IP address.
2. From the **Traffic Group** list, select the floating traffic group that you want the BIG-IP object to join.
3. Click **Update**.

After performing this task, the BIG-IP object belongs to the selected traffic group.

Repeat this task for each BIG-IP object that you want to be a member of the traffic group.

Viewing a list of traffic groups for a device

You can view a list of traffic groups for the device group. Using this list, you can add floating IP addresses to a traffic group, force a traffic group into a Standby state, and view information such as the current and next-active devices for a traffic group and its HA load factor.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, view the names of the traffic groups on the local device.

Viewing the members of a traffic group

You can use the BIG-IP® Configuration utility to view a list of all failover objects associated with a specific traffic group. For each failover object, the list shows the name of the object, the type of object, and the folder in which the object resides.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, click the name of the traffic group for which you want to view the associated objects.

This displays a list of all failover objects for the traffic group.

Traffic group properties

This table lists and describes the properties of a traffic group.

Property	Description
Name	The name of the traffic group, such as <code>traffic-group-1</code> .
Partition	The name of the folder or sub-folder in which the traffic group resides.
Description	A user-defined description of the traffic group.

Property	Description
MAC Masquerade Address	A user-created MAC address that floats on failover, to minimize ARP communications and dropped connections.
Current Device	The device on which a traffic group is currently running.
Next Active Device	The device currently most available to accept a traffic group if failover of that traffic group should occur.
HA Group	The HA group that you created and assigned to this traffic group. (This setting is optional.)
HA Group Status	Indicates whether an HA group is enabled for this traffic group.
Failover Method	The possible failover methods to configure: Failover to Device With Best HA Score and Failover using Preferred Device Order and then Load Aware . This property also shows whether auto-failback is enabled for this traffic group.
Failover Order	An ordered list of devices that the BIG-IP® system uses to determine the next-active device for the traffic group.
HA Load Factor	A numeric value pertaining to load-aware failover that represents the application traffic load of this traffic group relative to other active traffic groups on the same device.

Active and standby states

On each device, a particular floating traffic group is in either an active state or a standby state. In an *active* state, a traffic group on a device processes application traffic. In a *standby* state, a traffic group on a device is idle.

For example, on Device A, traffic-group-1 might be active, and on Device B, traffic-group-1 might be standby. Similarly, on Device B, traffic-group-2 might be active, and on Device A, traffic-group-2 might be standby.

When a device with an active traffic group becomes unavailable, the traffic group floats to (that is, becomes active on) another device. The BIG-IP® system chooses the target device for failover based on how you initially configured the traffic group when you created it. Note that the term *floats* means that on the target device, the traffic group switches from a standby state to an active state.

The following illustration shows a typical device group configuration with two devices and one traffic group (named `my_traffic_group`). In this illustration, the traffic group is active on Device A and standby on Device B prior to failover.

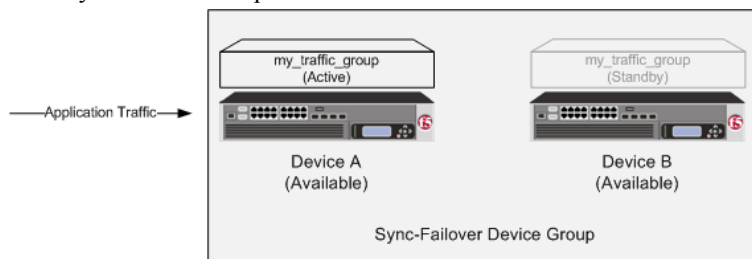


Figure 5: Traffic group states before failover

If failover occurs, the traffic group becomes active on the other device. In the following illustration, Device A has become unavailable, causing the traffic group to become active on Device B and process traffic on that device.

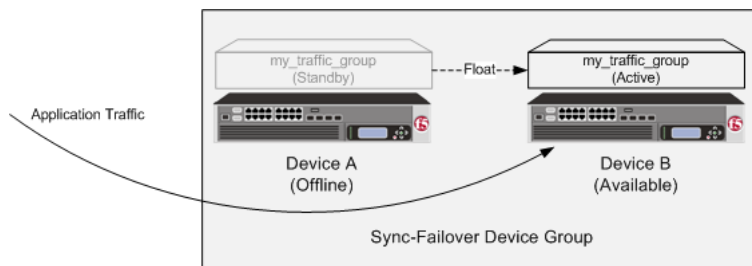


Figure 6: Traffic group states after failover

When `Device A` comes back online, the traffic group becomes standby on `Device A`.

About active-standby vs. active-active configurations

A device group that contains only one floating traffic group is known as an *active-standby* configuration.

A device group that contains two or more floating traffic groups is known as an *active-active* configuration. You can then choose to make all of the traffic groups active on one device in the device group, or you can balance the traffic group load by making some of the traffic groups active on other devices in the device group.

Viewing the failover state of a device

You can use the BIG-IP® Configuration utility to view the current failover state of a device in a device group. An *Active* failover state indicates that at least one traffic group is currently active on the device. A *Standby* failover state indicates that all traffic groups on the device are in a *Standby* state.

1. Display any screen of the BIG-IP Configuration utility.
2. In the upper left corner of the screen, view the failover state of the device.

Viewing the state of a traffic group

You can use the BIG-IP® Configuration utility to view the current state of all traffic groups on the device.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Failover Status area of the screen, view the state of all traffic groups on the device.

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.

Managing failover using HA groups

Sometimes a traffic group within a BIG-IP® Sync-Failover device group needs a certain number of resources to be up -- resources like pool members, trunk links, VIPRION® cluster members, or some combination of these.

With *HA groups*, you can define the minimum number of resources that a traffic group needs for it to stay active on its current device. If resources fall below that number, the traffic group fails over to a device with more resources. An HA group:

- Monitors resource availability on current and next-active devices for an active traffic group
- Calculates an HA "resource" score on each device for choosing the next-active device

Note: For an HA group to prevent a traffic group from failing over, all of the resource types that you specify in an HA group must meet the defined minimum thresholds for availability.

Creating an HA group

You use this task to create an HA group for a traffic group on a device in a BIG-IP® device group. An HA group is most useful when you want an active traffic group on a device to fail over to another device based on trunk and pool availability, and on VIPRION® systems, also cluster availability. You can create multiple HA groups on a single BIG-IP device, and you associate each HA group with the local instance of a traffic group.

Important: Once you create an HA group on one device and associate the HA group with a traffic group, you must create an HA group on every other device in the device group and associate it with that same traffic group.

1. Log in to a device in the device group (such as BIG-IP A), using the device's management IP address.
The login screen of the BIG-IP® Configuration utility opens.
2. On the Main tab, click **System > High Availability > HA Groups**
3. Click **Create**.
4. In the **HA Group Name** field, type a name such as `ha_group_deviceA_tg1`.
5. In the **Active Bonus** field, keep the default value.
The purpose of the active bonus is to boost the HA score to prevent failover when minor or frequent changes occur to the availability of a pool, trunk, or cluster.
6. In the Pools area of the screen, click **Add**.
If the **Add** button is grayed out, there are no pools on the BIG-IP system.
The **Add Pool to HA Group** dialog box appears.
7. From the **Pool** list, select a pool name.
8. Using the drop-down list, select the minimum number of active pool members required for this device to process traffic.

This value is the minimum number of pool members that you want to be up in order for the active instance of a specific traffic group to remain on its current device. You will assign this HA group to the traffic group later.

9. In the weight field, retain the default value or type a value and for the number of active pool members that are sufficient to be up for calculating the weight, select a value.

For example, if the total number of pool members is **6**, but the value of the **Sufficient Threshold** setting is **4** and there are only two pool members currently available, the BIG-IP system calculates the score by multiplying the weight you configured for the pool by the percentage of pool members available as compared to the *sufficient value*, not to the total number of pool members. If the weight we configure for the pool is **50**, and 50% of the pool members are up (2 of 4), then the HA score calculation for the pool is $50 \times 50\% = 25$.

10. Click **Add**.

This displays the New HA Group screen and shows the pool member criteria that must be met to prevent the traffic group from failing over.

11. In the Trunks area of the screen, click **Add**.

If the **Add** button is grayed out, there are no trunks on the BIG-IP system.

The **Add Trunk to HA Group** dialog box appears.

12. From the **Trunk** list, select a trunk name.

13. Using the drop-down list, select a minimum number of active links required for this device to process traffic, which in our example, is **3**.

This value is the minimum number of trunk links that you want to be up in order for `traffic-group-1` to remain on its current device. You will assign this HA group to the traffic group later.

14. For the weight field, type a value such as **50**, and for the number of active trunk links that are sufficient to be up for calculating the weight, select a value such as **3**.

For example, if the total number of trunk links is **4**, but the value of the **Sufficient Threshold** setting is **3** and there are only two links currently available, the BIG-IP system calculates the score by multiplying the weight you configured for the trunk by the percentage of links available as compared to the *sufficient value*, not to the total number of links. If the weight we configure for the trunk is **50**, and 66% of the links are up (2 of 3), then the HA score calculation for the trunk is $50 \times 66\% = 33$.

15. Click **Add**.

This displays the New HA Group screen and shows the trunk member criteria that must be met to prevent the traffic group from failing over.

16. Click **Create HA Group**.

17. Log in to each of the remaining devices in the device group and repeat this task, giving each HA group a unique name.

You can use the same weights and resource criteria within each HA group that you specified for this HA group.

For example, on `Device_A`, if you create `HA_GroupA_TG1` and associate it with `traffic-group-1`, then on `Device_B` you can create `HA_GroupB_TG1` and also associate it with `traffic-group-1`.

You now have an HA group that the BIG-IP system can use to trigger failover for whatever traffic group instance you assign this HA group to. If you intend to configure the traffic group to select the next-active device based on an HA score, this HA group will calculate an HA score for this device.

After creating an HA group on the local device, you must assign it to a traffic group on the local device.

Enabling an HA group for an existing traffic group

You use this task to associate an HA group with an existing traffic group. You associate an HA group with a traffic group when you want the traffic group to fail over to another device in the device group due to issues with trunk, pool, and/or VIPRION® cluster availability. Once a BIG-IP® device determines through this association that an active traffic group should fail over, the system chooses the next-active

device, according to the failover method that you configure on the traffic group: An ordered list of devices, load-aware failover based on device capacity and traffic load, or the HA score derived from the HA group configuration.

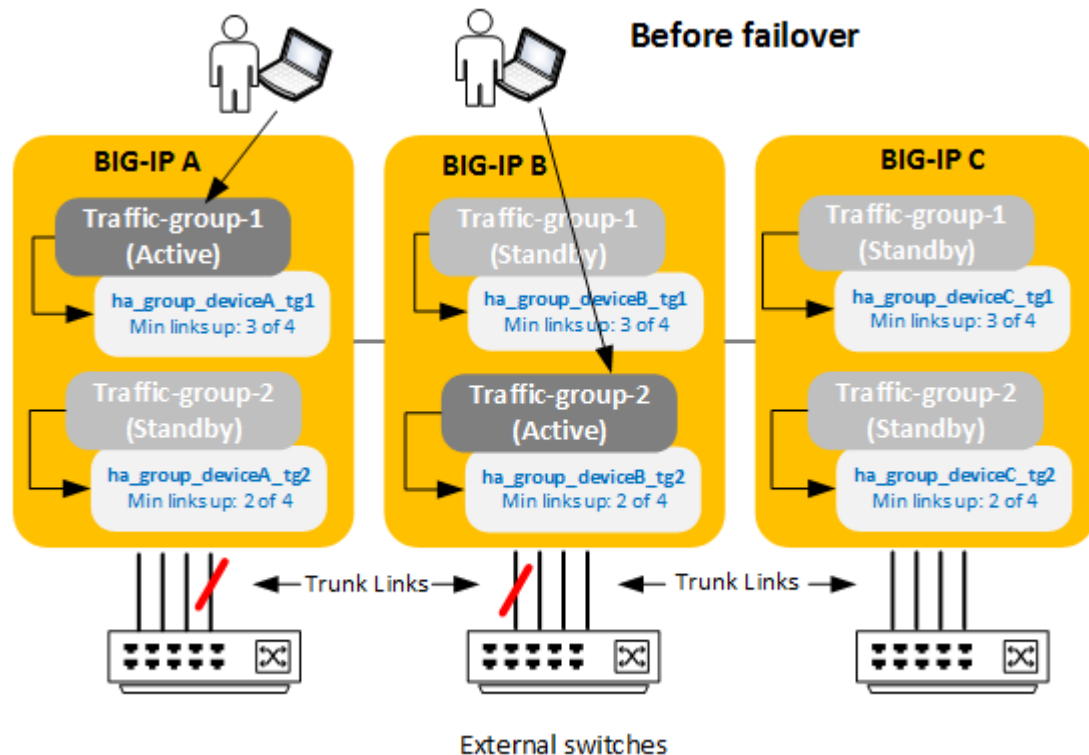
Important: HA groups are not included in config sync operations. For this reason, you must associate a different HA group on every device in the device group for this traffic group. For example, if the device group contains three devices and you want to create an HA group for `traffic-group-1`, you must associate a different HA group for `traffic-group-1` on each of the three devices separately. In a typical device group configuration, the values of the HA group settings on the traffic group will differ on each device.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, click the name of a traffic group on the local device.
This displays the properties of the traffic group.
3. From the **HA Group** list, select an HA group.
4. Click **Update**.

After you perform this task for the same traffic group on each device group member, the BIG-IP system ensures that the traffic group, when active, will fail over to another device when a configured number of trunk links, pool members, or VIPRION cluster members becomes unavailable.

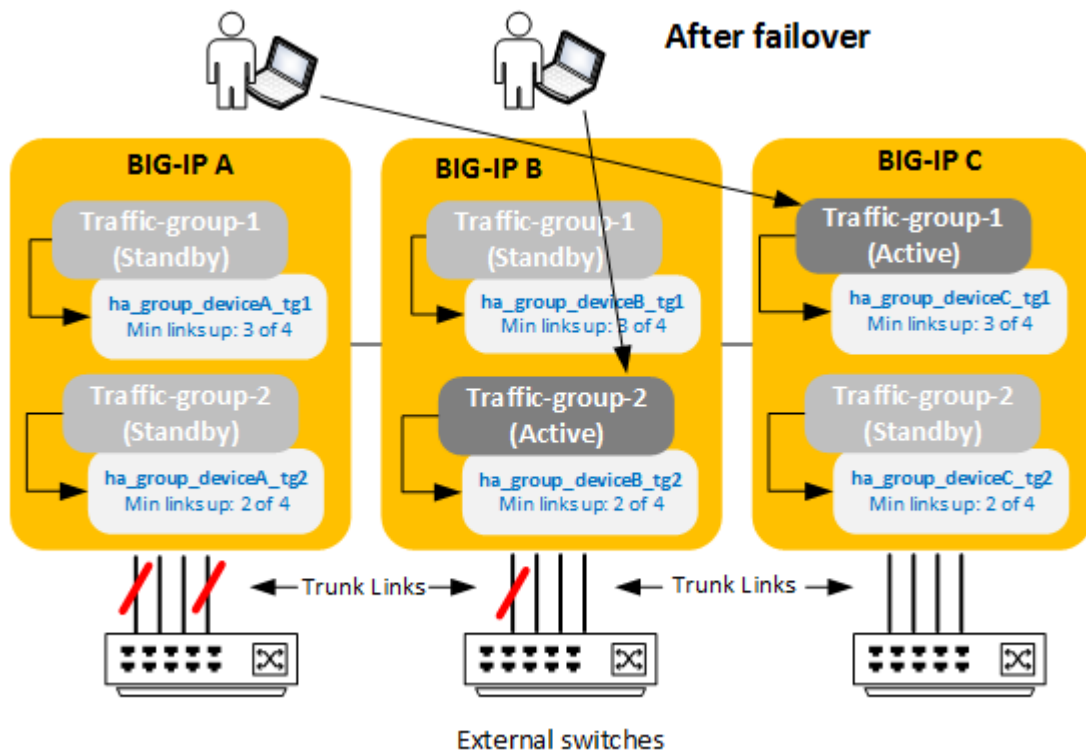
Example of an HA group deployment

This illustration shows three sample devices with two active traffic groups. We've configured both traffic groups to use HA groups to define acceptable criteria for trunk health. Although it's not shown here, we'll assume that `traffic-group-1` and `traffic-group-2` use the HA score and the Preferred Device Order failover methods, respectively, to pick their next-active devices.



In our example, we see that on both BIG-IP A and BIG-IP B, three of four trunk links are currently up, which meets the minimum criteria specified in the HA groups assigned to `traffic-group-1` and `traffic-group-2` on those devices. This allows each traffic group to stay active on its current device.

Now suppose that the trunk on BIG-IP A loses another link. We see that even though BIG-IP A is still up, `traffic-group-1` has failed over because BIG-IP A no longer meets the HA group criteria for hosting the traffic group: only two of four trunk links are now up on that device.



Because we've configured `traffic-group-1` to use HA scores to select the next-active device, the traffic group fails over to BIG-IP C, because this is the device with the most trunk links up and therefore has the highest HA score for hosting this traffic group.

As for `traffic-group-2`, it stays on its current device because BIG-IP B still meets the minimum criteria specified in its HA group.

About next-active device selection

For every active traffic group in your device group, the BIG-IP® Configuration utility displays the *current* device, meaning the device that a traffic group is currently active on.

The BIG-IP® system can also tell you the device that is to be the next-active device. The *next-active* device is the device that the traffic group will fail over to if the traffic group has to fail over for some reason.

The device labeled as next-active for a traffic group can change at any time, depending on:

- Which devices are currently available in the device group
- Which device is best able to take on extra traffic group load
- Which device has the most available trunk, pool, or VIPRION® cluster members (if you're using the HA groups feature)

You can tell the BIG-IP system how to choose a next-active device for a traffic group by configuring the traffic group's *failover method*. The available failover methods are **Failover to Device With Best HA Score** and **Failover using Preferred Device Order and then Load Aware**.

About using HA scores to pick the next-active device

An *HA score* is a numeric value that the BIG-IP® system calculates independently for each instance of a particular traffic group, when you have assigned an HA group to each traffic group instance. For each traffic group instance, the HA group's monitoring function determines the availability of certain resources such as trunk links, pool members, or VIPRION® cluster members.

The BIG-IP system uses these per-instance scores to decide which device has the most resources that the traffic group needs, such as trunk links or pool members. The higher the score for a traffic group instance, the higher the availability of needed resources.

You must have an HA group assigned to each instance of the same traffic group in order for the system to calculate an HA score. An HA score is calculated based on how the corresponding HA group is configured. Whenever the HA group for the active traffic group decides to trigger failover, the traffic group automatically fails over to the device with the highest score.

To get the BIG-IP to base the selection of a traffic group's next-active device on an HA score, you configure the **Failover to Device with Best HA Score Failover Method** setting on a floating traffic group.

Factors in HA score calculation

The BIG-IP® system calculates an HA health score per traffic group on a device, based on weight, minimum threshold, sufficient threshold, and active bonus values that you specify when you configure an HA group.

HA score weight value

A *weight* is a health value that you assign to each member of the HA group (that is, a pool, trunk, and/or VIPRION® cluster). The weight that you assign to each HA group member must be in the range of 10 through 100.

The maximum overall weight that the BIG-IP system can potentially calculate is the sum of the individual weights for the HA group members, plus the active bonus value. There is no limit to the sum of the member weights for the HA group as a whole.

HA score minimum threshold value (optional)

For each member of an HA group, you can specify a setting known as a minimum threshold. A *minimum threshold* is a value that specifies the number of object members that must be available to prevent failover. The system factors in a threshold value when it calculates the overall score for the traffic group or device.

The way that the BIG-IP system calculates the score depends on the number of object members that are actually available as compared to the configured minimum threshold value:

- If the number of available object members is less than the threshold, the BIG-IP system assigns a score of 0 to the HA group member so that the score of that HA group member no longer contributes to the overall score.

For example, if a trunk in the HA group has four trunk members and you specify a minimum threshold value of 3, and the number of available trunk members falls to 2, then the trunk contributes a score of 0 to the total score for the traffic group or device.

- If the number of available object members equals or exceeds the minimum threshold value, or you do not specify a minimum threshold, the BIG-IP system calculates the score as described previously, by multiplying the percentage of available object members by the weight for each HA group member and then adding the scores to determine the overall score for the traffic group or device.

The minimum threshold that you define for pools can be less than or equal to the number of members in the pool. For clusters, the threshold can be less than or equal to the number of possible blades in the chassis, and for trunks, the minimum threshold can be less than or equal to the number of possible members in a trunk for that platform.

Tip: Do not configure the `tms` attribute `min-up-members` on any pool that you intend to include in the HA group.

HA score sufficient threshold value (optional)

When you've configured the BIG-IP® system to use HA scores to pick the next-active device for a traffic group, the traffic group will fail over whenever another device has a higher score for that same traffic group. This means that an active traffic group could potentially fail over frequently because it will fail over even when its HA group's minimum threshold value is still met.

To mitigate this problem, you can define a sufficient threshold value. The *sufficient threshold* value specifies the amount of available resource (of a trunk, pool, or cluster) that is considered good enough to prevent the traffic group from failing over when another device has a higher score.

The default value for the **Sufficient Threshold** setting is **All**, which means that the system considers the amount of available resource to be sufficient when all of its component members are available. For example, if a trunk has a total of four links, and you specify the default sufficient threshold value, then all of the trunk links must be up to prevent failover when another device has a higher HA score. If you specify a sufficient threshold of 3, then only three of the four trunk links must be up to prevent failover when another device has a higher HA score.

HA score active bonus value

An *active bonus* is an amount that the BIG-IP system automatically adds to the overall HA score of an active traffic group or device. An active bonus ensures that the traffic group or device remains active when its score would otherwise temporarily fall below the score of the standby traffic group on another device. The active bonus that you configure can be in the range of 0 to 100.

A common reason to specify an active bonus is to prevent failover due to *flapping*, the condition where failover occurs frequently as a trunk member switches rapidly between availability and unavailability. In this case, you might want to prevent the HA scoring feature from triggering failover each time a trunk member is lost. You might also want to prevent the HA scoring feature from triggering failover when you make minor changes to the BIG-IP system configuration, such as adding or removing a trunk member.

For example, suppose that the HA group for a traffic group on each device contains a trunk with four members, and you assign a weight of 30 to each trunk. Without an active bonus defined, if the trunk on one device loses some number of members, failover occurs because the overall calculated score for that traffic group becomes lower than that of a peer device. You can prevent this failover from occurring by specifying an active bonus value.

The BIG-IP system uses an active bonus to contribute to the HA score of an active traffic group only; the BIG-IP system never uses an active bonus to contribute to the score of a standby traffic group.

Note: An exception to this behavior is when the active traffic group score is 0. In this case, the system does not add the active bonus to the active traffic group or active device score.

To decide on an active bonus value, calculate the trunk score for some number of failed members (such as one of four members), and then specify an active bonus that results in a trunk score that is greater than the weight that you assigned to the trunk.

For example, if you assigned a weight of 30 to the trunk, and one of the four trunk members fails, the trunk score becomes 23 (75% of 30), putting the traffic group at risk for failover. However, if you specified an active bonus of 8 or higher, failover would not actually occur, because a score of 8 or higher, when added to the score of 23, is greater than 30.

Example of HA health score calculation

This example illustrates the way that HA group configuration results in the calculation of an HA health score for a traffic group on a specific device. Suppose that you previously created an HA group for `traffic-group-1` on all device group members and that `traffic-group-1` is currently active on device `Bigip_A`. Also suppose that on device `Bigip_B`, the HA group for `traffic-group-1` consists of two pools and a trunk, with weights that you assign:

Table 4: Sample HA group configuration for `traffic-group-1` on `Bigip_B`

HA group object	Member count	User-specified weight
http_pool	8	50
ftp_pool	6	20
trunk1	4	30

Now suppose that on device `Bigip_B`, the current member availability of pool `http_pool`, pool `ftp_pool`, and trunk `trunk1` is 5, 6, and 3, respectively. The resulting HA score that the BIG-IP system calculates for `traffic-group-1` on `Bigip_B` is shown here:

Table 5: Sample health score calculation for `traffic-group-1` on `Bigip_B`

HA group object	Member count	Available member count	User-specified weight	Current HA score
http_pool	8	5 (62.5%)	50	31 (60% x 50)
ftp_pool	6	6 (100%)	20	20 (100% x 20)
trunk1	4	3 (75%)	30	23 (75% x 30)
				Total score: 74

In this example, the total HA score for `traffic-group-1` on `Bigip_B` is currently 74. If this score is currently the highest score in the device group for `traffic-group-1`, then `traffic-group-1` will automatically failover and become active on `Bigip_B`.

About matching HA health scores

In rare cases, the BIG-IP® system might calculate that two or more traffic groups have the same HA score. In this case, the BIG-IP system needs an additional method for choosing the next-active device for an active traffic group.

The way that the BIG-IP system chooses the next-active device when HA health scores match is by determining the management IP address of each matching device and then calculating a score based on the highest management IP address of those devices.

For example, if `Bigip_A` has an IP address of `192.168.20.11` and `Bigip_B` has an IP address of `192.168.20.12`, and their HA scores match, the BIG-IP system calculates a score based on the address `192.168.20.12`.

About using a preferred device order list to pick the next-active device

A *Preferred Device Order list* is a static list of devices that you can assign to a floating traffic group as a way for the BIG-IP® system to choose the next-active device. The list tells the BIG-IP system the order to use when deciding which device to designate as the next-active device for the traffic group.

You create a preferred device order list by configuring the traffic group's **Failover Method** setting and choosing **Failover using Preferred Device Order and then Load Aware**. For example, for `traffic-group-1`, if you create a list that contains devices BIG-IP A, BIG-IP B, and BIG-IP C, in that order, the system checks to see if BIG-IP A is up and if so, designates BIG-IP A as the target device for `traffic-group-1`. If the system sees that BIG-IP A is down, it checks BIG-IP B to see if it's up, and if so, designates BIG-IP B as the target failover device for the traffic group, and so on.

If you assigned an HA group to the traffic group, the BIG-IP system not only selects the next-active device by checking to see if a device in the list is up, but also whether the device's trunk, pool, or cluster resources meet the minimum criteria defined in the HA group. In this case, if a device's resources don't meet the minimum criteria (and therefore its HA score is zero), the system will not designate that device as the next-active device and will check the next device in the list.

If the preferred device order list is empty or if none of the devices in the list is available, the BIG-IP system switches to using the load-aware failover method to choose the next-active device.

Note: When you enable the auto-failback feature for a traffic group, the BIG-IP system tries to ensure that the traffic group is always active on the first device in the list. If the first device in the list is unavailable, no fail-back occurs.

About auto-failback

The failover feature includes an option known as auto-failback. When you enable *auto-failback*, a traffic group that has failed over to another device fails back to a preferred device when that device is available. If you do not enable auto-failback for a traffic group, and the traffic group fails over to another device, the traffic group remains active on the new device until that device becomes unavailable.

You can enable auto-failback on a traffic group only when you have configured an ordered list with at least one entry, for that traffic group. In this case, if auto-failback is enabled and the traffic group has failed over to another device, then the traffic group fails back to the first device in the traffic group's ordered list (the preferred device) when that device becomes available.

Important: If the first device in the ordered list is unavailable, no fail-back occurs. The traffic group does not fail back to the next available device in the list and instead remains on its current device.

If a traffic group fails over to another device, and the new device fails before the auto-failback timeout period has expired, the traffic group will still fail back, to the original device if available. The maximum allowed timeout value for auto-failback is 300 seconds.

Creating an HA ordered list

You perform this task to create a prioritized, ordered list for a floating traffic group. The BIG-IP® system uses this list to determine the next-active device for this traffic group. This configuration option is most useful for device groups with homogeneous hardware platforms and similar application traffic loads, or for applications that require a specific target failover device, such as those that use connection mirroring. When failover occurs, the traffic group will fail over to the first available device in the list.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, click the name of a traffic group on the local device.

This displays the properties of the traffic group.

3. For the **Failover Method** setting, choose **Failover using Preferred Device Order and then Load Aware**.
4. Select or clear the check box **Always Failback to First Device if it is Available**:
 - Select the check box to cause the traffic group, after failover, to fail back to the first device in the traffic group's ordered list when that device (and only that device) is available. If that device is unavailable, no failback occurs and the traffic group continues to run on the current device.
 - Clear the check box to cause the traffic group, after failover, to remain active on its current device until failover occurs again.
5. If auto-failback is enabled, in the **Auto Failback Timeout** field, type the number of seconds that you want the system to wait before failing back to the specified device. The range is from 0 to 300 seconds. The default is 60. A value of 40 to 60 allows for state mirroring information to be re-mirrored for traffic groups.
6. For the **Failover Order** setting, in the **Load-Aware** box, select a device name and using the Move button, move the device name to the **Preferred Order** 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, make sure 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 none of the devices in the **Preferred Order** list is currently available when failover occurs, the BIG-IP system uses load-aware failover instead.
7. Click **Update**.

After you perform this task, the BIG-IP system designates the first available device that is highest in the ordered list as the next-active device for the traffic group. If you've assigned an HA group to the traffic group, the traffic group will fail over to the first available device in the list that has a non-zero HA score (that is, a device whose trunk, pool, or VIPRION® cluster resources meet the minimum criteria specified in the HA group).

About using traffic load to pick the next-active device

If you want the BIG-IP® system to base the next-active selection for a traffic group on application traffic load, you can use load-aware failover. *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 the amount of application traffic that traffic groups process on a device. The load-aware configuration option is most useful for device groups with varying application traffic loads.

The BIG-IP system implements load-aware failover by calculating a utilization score for each device, based on numeric values that you specify for each traffic group relative to the other traffic groups in the device group. The system then uses this current 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.

Note: *If you have varying hardware platforms in your device group, you can use `tmsh` to specify the relative capacity of each device, and this value factors into the score calculation along with the traffic load value. The `tmsh` command to do this is: `modify /cm device device_name ha-capacity integer`.*

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 device group.

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

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 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.

About the HA load factor

For each traffic group on a BIG-IP® device, you can assign an high availability (HA) load factor. An *HA load factor* is a number that represents the relative application traffic load that an active traffic group processes compared to other active traffic groups in the device group.

For example, if the device group has two active traffic groups, and one traffic group processes twice the amount of application traffic as the other, then you can assign values of 4 and 2, respectively. You can assign any number for the HA load factor, as long as the number reflects the traffic group's relative load compared to the other active traffic groups.

About metrics for the HA load factor

User-specified values for the HA load factor can be based on different metrics. For example, suppose you have the three devices `Bigip_A`, `Bigip_B`, and `Bigip_C`, and each device has one active traffic group with an HA load factor of 2, 4, or 8 respectively. These values could indicate either of the following:

- If each traffic group contains one virtual address, then the sample factor values could indicate that the virtual server for `Bigip_B` processes twice the amount of traffic as that of `Bigip_A`, and the virtual server for `Bigip_C` processes twice the amount of traffic as that of `Bigip_B`.
- If the traffic group on `Bigip_A` contains one virtual address, the traffic group on `Bigip_B` contains two virtual addresses, and the traffic group on `Bigip_C` contains four virtual addresses, this could indicate that the virtual servers corresponding to those virtual addresses each process the same amount of traffic compared to the others.

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, choose **Failover using Preferred Device Order and then 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.

About MAC masquerade addresses

A *MAC masquerade address* is a unique, floating Media Access Control (MAC) address that you create and control. You can assign one MAC masquerade address to each traffic group on a BIG-IP® device. By assigning a MAC masquerade address to a traffic group, you indirectly associate that address with any floating IP addresses (services) associated with that traffic group. With a MAC masquerade address per traffic group, a single VLAN can potentially carry traffic and services for multiple traffic groups, with each service having its own MAC masquerade address.

A primary purpose of a MAC masquerade address is to minimize ARP communications or dropped packets as a result of a failover event. A MAC masquerade address ensures that any traffic destined for the relevant traffic group reaches an available device after failover has occurred, because the MAC masquerade address floats to the available device along with the traffic group. Without a MAC masquerade address, on failover the sending host must relearn the MAC address for the newly-active device, either by sending an ARP request for the IP address for the traffic or by relying on the gratuitous ARP from the newly-active device to refresh its stale ARP entry.

The assignment of a MAC masquerade address to a traffic group is optional. Also, there is no requirement for a MAC masquerade address to reside in the same MAC address space as that of the BIG-IP device.

Note: *When you assign a MAC masquerade address to a traffic group, the BIG-IP system sends a gratuitous ARP to notify other hosts on the network of the new address.*

Managing Connection Mirroring

About connection mirroring

Purpose

BIG-IP® system high availability includes the ability for a device to mirror connection and persistence information to another device in a device service clustering (DSC®) configuration, to prevent interruption in service during failover. The BIG-IP system mirrors connection and persistence data over TCP port 1028 with every packet or flow state update.

How to enable connection mirroring

You enable connection mirroring on the relevant virtual server, and then on each device in the device group, you specify the self IP addresses that you want other devices to use when mirroring connections to the local device. This enables mirroring between an active traffic group and a mirroring peer in the device group. You can enable connections such as FTP, Telnet, HTTP, UDP, and SSL connections.

Note: In addition to enabling connection mirroring on the virtual server, you must also assign the appropriate profiles to the virtual server. For example, if you want the BIG-IP system to mirror SSL connections, you must assign one or more SSL profiles to the virtual server.

When to enable connection mirroring

You should enable connection mirroring whenever failover would cause a user session to be lost or significantly disrupted. For example, long-term connections such as FTP and Telnet are good candidates for mirroring. For this type of traffic, if failover occurs, an entire session can be lost if the connections are not being mirrored to a peer device. Conversely, the mirroring of short-term connections such as HTTP and UDP is typically not recommended, because these protocols allow for failure of individual requests without loss of the entire session, and the mirroring of short-term connections can negatively impact system performance.

Platform caveats

Important: Connection mirroring only works between devices with identical hardware platforms. Note that for VIPRION® systems, you configure the BIG-IP system to mirror connections between two chassis or between two vCMP® guests that reside in separate chassis. If the VIPRION system is not provisioned for vCMP, each chassis must have the same number of blades in the same slot numbers. For vCMP systems, each guest must be assigned to the same number of blades in the same slot numbers, with the same number of cores allocated per slot. For more information, see the section About connection mirroring for VIPRION systems.

About connection mirroring for VIPRION systems

For VIPRION® systems, each device in a Sync-Failover device group can be either a physical cluster of slots within a chassis, or a virtual cluster for a vCMP® guest. In either case, you can configure a device to mirror an active traffic group's connections to its next-active device.

Important: For mirroring to work, both the active device and its next-active device must have identical chassis platform and blade models.

You enable connection mirroring on the relevant virtual server, and then you configure each VIPRION cluster or vCMP guest to mirror connections by choosing one of these options:

Within a cluster

You can configure the BIG-IP system to mirror connections between blades within a single VIPRION cluster on the same chassis. This option is not available on VIPRION systems provisioned to run vCMP.

Note: With this option, the BIG-IP system mirrors Fast L4 connections only.

Between clusters (recommended)

You can configure the BIG-IP system to mirror connections between two chassis or between two vCMP guests that reside in separate chassis. When you choose this option, the BIG-IP system mirrors a traffic group's connections to the traffic group's next-active device. For VIPRION systems that are not provisioned for vCMP, each chassis must have the same number of blades in the same slot numbers. For VIPRION systems provisioned for vCMP, each guest must be assigned to the same number of blades in the same slot numbers, with the same number of cores allocated per slot.

In addition to enabling connection mirroring on the virtual server, you must also assign the appropriate profiles to the virtual server. For example, if you want the BIG-IP system to mirror SSL connections, you must assign one or more SSL profiles to the virtual server.

Connection mirroring and traffic groups

Connection mirroring operates at the traffic group level. That is, for each virtual server that has connection mirroring enabled, the traffic group that the virtual server belongs to mirrors its connections to its next-active device in the device group.

For example, if `traffic-group-1` is active on `Bigip_A`, and the next-active device for that traffic group is `Bigip_C`, then the traffic group on the active device mirrors its in-process connections to `Bigip_C`.

If `Bigip_A` becomes unavailable and failover occurs, `traffic-group-1` goes active on `Bigip_C` and begins mirroring its connections to the next-active device for `Bigip_C`.

Important: Connection mirroring only works between devices with identical hardware platforms. Note that for VIPRION® systems, you configure the BIG-IP system to mirror connections between two chassis or between two vCMP® guests that reside in separate chassis. If the VIPRION system is not provisioned for vCMP®, each chassis must have the same number of blades in the same slot numbers. For vCMP systems, each guest must be assigned to the same number of blades in the same slot numbers, with the same number of cores allocated per slot.

Configuration task summary

Configuring connection mirroring requires you to perform these specific tasks:

Specifying a local self IP address for connection mirroring (required)

This local self IP address is the address that you want other devices in a device group to use when other traffic groups mirror their connections to a traffic group on this device.

Enabling connection mirroring on a virtual server

The BIG-IP® can mirror TCP or UDP connections for a virtual server. When you enable connection mirroring on a virtual server, and you then make the relevant virtual address a member of an active floating traffic group, the traffic group can mirror its connections to its corresponding standby traffic group on another device.

Enabling connection mirroring on a SNAT

The BIG-IP system can mirror TCP or UDP connections for a SNAT.

Enabling persistence mirroring on a persistence profile

The BIG-IP system can mirror persistence information between peers for the following persistence profiles:

- Destination address affinity
- Hash
- Microsoft Remote Desktop (MSRDP)
- Session Initiation Protocol (SIP)
- Source address affinity
- SSL
- Universal

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.

Important: Connection mirroring only functions between devices with identical hardware platforms.

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. Near the top of the screen, click **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.

Configuring connection mirroring between VIPRION clusters

Before doing this task, you must enable connection mirroring on the relevant virtual server.

Using the BIG-IP® Configuration utility, you can configure connection mirroring between two VIPRION® or vCMP® clusters as part of your high availability setup:

- When you configure mirroring on a VIPRION system where vCMP is not provisioned (a bare-metal configuration), an active traffic group on one chassis mirrors its connections to the next-active chassis in the device group.
- When you configure mirroring on a vCMP guest, an active traffic group mirrors its connections to its next-active guest in another chassis.

Important: Connection mirroring requires that both devices have identical hardware platforms (chassis and blades).

Important: You must perform this task locally on every device (chassis or vCMP guest) in the device group. For VIPRION systems with bare-metal configurations (no vCMP provisioned), each chassis must contain the same number of blades in the same slot numbers. For VIPRION systems provisioned for vCMP, each guest must reside on a separate chassis, be assigned to the same number of blades in the same slot numbers, and have the same number of cores allocated per slot.

1. From a browser window, log in to the BIG-IP Configuration utility, using the cluster IP address.
2. On the Main tab, click **Device Management > Devices**.
The Devices screen opens.
3. In the Device list, in the Name column, click the name of the device you want to configure.
4. From the Device Connectivity menu, choose Mirroring.
5. From the **Network Mirroring** list, select **Between Clusters**.
6. Click **Update**.

Enabling connection mirroring for TCP and UDP connections

Verify that you have specified primary and secondary mirroring IP addresses on this device. Other traffic groups in the device group use these addresses when mirroring connections to this device.

You can perform this task to enable TCP or UDP connections for a virtual server. *Connection mirroring* is an optional feature of the BIG-IP® system, designed to ensure that when failover occurs, in-process connections are not dropped. You enable mirroring for each virtual server that is associated with a floating virtual address.

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. For the **Connection Mirroring** setting, select the check box.

Note: This setting only appears when the BIG-IP device is a member of a device group.

5. Click **Update** to save the changes.

Enabling connection mirroring for SNAT connections

You can perform this task to enable connection mirroring for source network address translation (SNAT). *Connection mirroring* is an optional feature of the BIG-IP® system, designed to ensure that when failover

occurs, in-process SNAT connections are not dropped. You can enable mirroring on each SNAT that is associated with a floating virtual address.

1. On the Main tab, click **Local Traffic > Address Translation**.
The **SNAT List** screen displays a list of existing SNATs.
2. In the Name column, click the relevant SNAT name.
3. For the **Stateful Failover Mirror** setting, select the check box.
4. Click **Update**.

In addition to enabling connection mirroring on a SNAT, you must also specify a mirroring IP address on this device. Other traffic groups in the device group use this address when mirroring their connections to this device.

Enabling mirroring of persistence records

Verify that you have specified primary and secondary mirroring IP addresses on this device. Other traffic groups in the device group use these addresses when mirroring persistence records to this device.

You can perform this task to mirror persistence records to another device in a device group.

1. On the Main tab, click **Local Traffic > Profiles > Persistence**.
The Persistence profile list screen opens.
2. In the Name column, click the name of the relevant persistence profile.
3. For the **Mirror Persistence** setting, select the check box.
4. Click **Update**.

Working with Folders

About folders on the BIG-IP system

A *folder* is a container for BIG-IP® configuration objects and files on a BIG-IP device. Virtual servers, pools, and self IP addresses are examples of objects that reside in folders on the system.

In the context of the BIG-IP system, a folder is a container for BIG-IP system objects. Folders resemble standard UNIX directories, in that the system includes a hierarchy of folders and includes a `root` folder (represented by the `/` symbol) that is the parent for all other folders on the system.

You can create sub-folders within a high-level folder. For example, if you have a high-level folder (partition) within the `root` folder named `Customer1`, you can create a sub-folder, such as `App_B`, within `Customer1`.

A folder can contain other folders.

One of the important ways that you can use folders is to set up full or granular synchronization and failover of BIG-IP configuration data in a device group. You can synchronize and fail over all configuration data on a BIG-IP device, or you can synchronize and fail over objects within a specific folder only.

You manage BIG-IP folders and sub-folders using the Traffic Management Shell (`tmssh`) command line interface.

About folder attributes for redundancy

Folders have two specific redundancy attributes that enable granular synchronization and failover of BIG-IP® system data within a device group. These two attributes are a device group name and a traffic group name.

Device group name

This attribute determines the scope of the synchronization, that is, the specific devices to which the system synchronizes the contents of the associated folder. When you create a Sync-Failover device group on a BIG-IP device, the system assigns that device group name as an attribute of folder `root`. Any other folders that you subsequently create on a device group member then inherit that same device group name, by default.

The result is that when you enable config sync for the local device, the contents of the `root` folder and any sub-folders are synchronized across the members of the specified device group.

Note: The device group assigned to a folder must contain the local BIG-IP device. Also, you cannot remove the local BIG-IP device from the Sync-Failover device group assigned to a folder.

Traffic group name

This attribute determines the scope of a failover action, that is, the specific configuration objects that will fail over if the device becomes unavailable. If you enabled failover on a device (as part of running the Setup utility or upgrading from a previous BIG-IP version), the device contains the default traffic group named `traffic-group-1`. The system assigns this traffic group name by default as an attribute of folder `root`. Any other folders that you subsequently create on a device group member inherit that same

traffic group name, by default. The result is that when the local device is a member of a Sync-Failover device group, all failover objects within the `root` folder and its hierarchy fail over based on the definition of the specified traffic group.

You can assign a different traffic group to a specific sub folder. For example, you can create an iApps™ application in a sub folder and change the inherited traffic group value of `traffic-group-1` to a traffic group that you create, such as `traffic-group-2`. You can then manually cause `traffic-group-2` to fail over to another device so that the iApp application runs on a separate device from `traffic-group-1`.

About the root folder

At the highest-level, the BIG-IP® system includes a `root` folder. The `root` folder contains all BIG-IP configuration objects on the system, by way of a hierarchical folder and sub-folder structure within it.

By default, the BIG-IP system assigns a Sync-Failover device group and a traffic group to the `root` folder. All folders and sub-folders under the `root` folder inherit these default assignments.

Viewing redundancy attributes for the root folder

You can view the device group and traffic group attributes assigned to the `root` folder. All eligible configuration objects in the `root` folder hierarchy synchronize to the named device group, and all failover objects in the hierarchy fail over with the named traffic group.

Note: All folders and sub-folders in the `root` folder hierarchy inherit these attribute values, by default.

1. On the Main tab, click **System > Platform**.
The Platform screen opens.
2. For the **Redundant Device Configuration** setting, view the device group and the traffic group attributes.

Configuring the traffic group attribute for the root folder

If you have two or more traffic groups defined on the BIG-IP® system, you can configure the traffic group attribute assigned to the `root` folder. By default, this value is `traffic-group-1`.

Note: All folders and sub folders in the `root` folder hierarchy inherit this attribute value, by default.

1. On the Main tab, click **System > Platform**.
The Platform screen opens.
2. If the system includes two or more traffic groups, then for the **Default traffic group** setting, select a traffic group from the list.
3. Click **Update**.

By default, all failover objects in the `root` folder hierarchy fail over with the named traffic group, when failover occurs.

About using HA scores to pick the next-active device

An *HA score* is a numeric value that the BIG-IP® system calculates independently for each instance of a particular traffic group, when you have assigned an HA group to each traffic group instance. For each traffic group instance, the HA group's monitoring function determines the availability of certain resources such as trunk links, pool members, or VIPRION® cluster members.

The BIG-IP system uses these per-instance scores to decide which device has the most resources that the traffic group needs, such as trunk links or pool members. The higher the score for a traffic group instance, the higher the availability of needed resources.

You must have an HA group assigned to each instance of the same traffic group in order for the system to calculate an HA score. An HA score is calculated based on how the corresponding HA group is configured. Whenever the HA group for the active traffic group decides to trigger failover, the traffic group automatically fails over to the device with the highest score.

To get the BIG-IP to base the selection of a traffic group's next-active device on an HA score, you configure the **Failover to Device with Best HA Score Failover Method** setting on a floating traffic group.

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.

Example

In this configuration example, the device group is named `Device Group A`. This device group contains two BIG-IP devices, named `Device 1` and `Device 2`, and these two devices are peers of one another. The default traffic group, named `traffic-group-1`, resides on each device.

`Device 1` actively processes traffic because `traffic-group-1` is in an Active state on that device. `Device 2` remains idle until failover occurs because `traffic-group-1` is in a Standby state on that device.

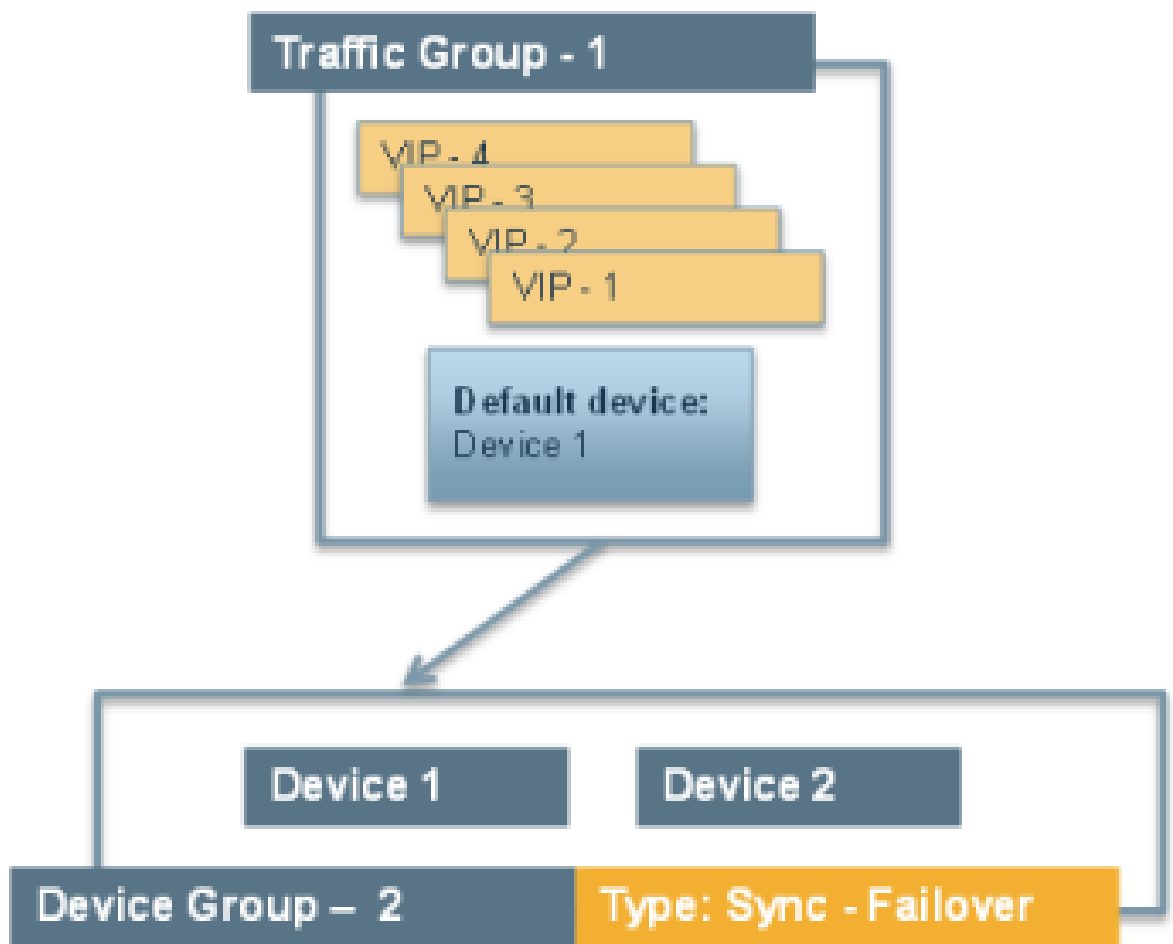


Figure 7: Example active-standby configuration

By implementing this configuration, you ensure that:

- Each device has base network components (such as self IPs and VLANs) configured.
- The two devices can synchronize their configuration to one another.
- Failover capability and connection mirroring are enabled on each device.

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

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. From the **SSH IP Allow** list, retain the default value of ***All Addresses**, or specify a range.
9. Click **Next**.
10. 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`.
- One end of an iSession™ connection for WAN traffic optimization.

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`.

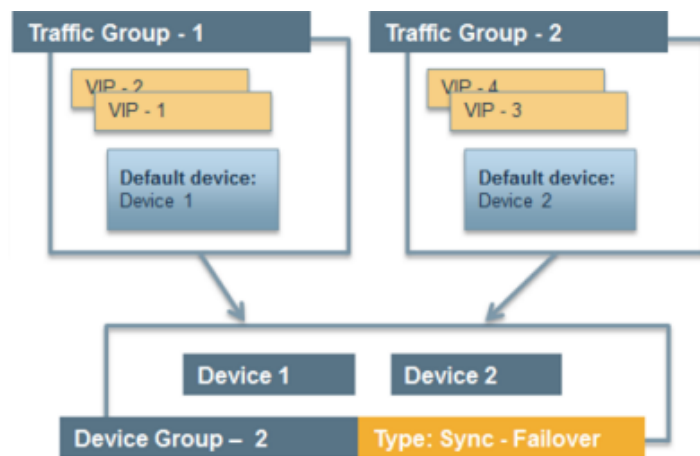


Figure 8: Device group with active-active configuration

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

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. From the **SSH IP Allow** list, retain the default value of ***All Addresses**, or specify a range.
9. Click **Next**.
10. 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 > Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. 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.
5. Click **Retrieve Device Information**.

6. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
7. In the **Name** field, verify that the name of the remote device is correct.
8. Click **Add Device**.

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. For the **Members** setting, 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.

4. 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 is the default value and 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.

5. 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 **Always Failback to First Device if it is Available**:
 - Select the check box to cause the traffic group, after failover, to fail back to the first device in the traffic group's ordered list when that device (and only that device) is available. If that device is unavailable, no failback occurs and the traffic group continues to run on the current device.
 - 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, click the arrow next to 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, choose the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Push the selected device configuration to the group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device 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.

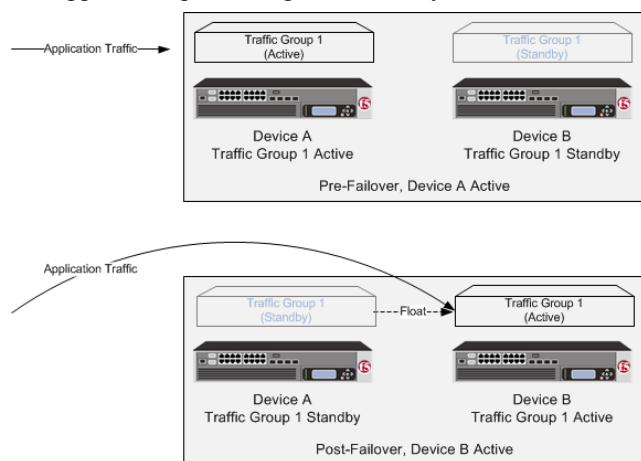


Figure 9: A two-member Sync-Failover device group for an active-standby configuration

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 6: 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 7: 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 8: 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 9: 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 should be running BIG-IP version 12.x or higher. 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 <code>root</code> folder must be set correctly (<code>Sync-Failover</code> and <code>traffic-group-1</code>).
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.
Switchboard failsafe	If your devices are provisioned for vCMP® and your guests are members of a device group, make sure the guests' switchboard failsafe setting is set to the default value. Any change from the default switchboard failsafe configuration must always be done on the vCMP host, and not on the guests.

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

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. Near the top of the screen, click **ConfigSync**.
5. From the **Local Address** list, 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.

Important: Connection mirroring only functions between devices with identical hardware platforms.

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. Near the top of the screen, click **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.

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 > Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. 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.
5. Click **Retrieve Device Information**.
6. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
7. In the **Name** field, verify that the name of the remote device is correct.
8. Click **Add Device**.

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. In the **Name** field, type a name for the device group.
4. From the **Group Type** list, select **Sync-Failover**.
5. In the **Description** field, type a description of the device group.
This setting is optional.
6. From the **Configuration** list, select **Advanced**.
7. For the **Members** setting, 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.

8. From the **Sync Type** list:

- Select **Automatic with Incremental Sync** when you want the BIG-IP system to automatically sync the most recent BIG-IP configuration changes from a device to the other members of the device group. In this case, the BIG-IP system syncs the configuration data whenever the data changes on any device in the device group.
- Select **Manual with Incremental Sync** when you want to manually initiate a config sync operation. In this case, the BIG-IP system syncs the latest BIG-IP configuration changes from the device you choose to the other members of the device group. We strongly recommend that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.
- Select **Manual with Full Sync** when you want to manually initiate a config sync operation. In this case, the BIG-IP system syncs the full set of BIG-IP configuration data from the device you choose to the other members of the device group. We strongly recommend that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.

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. 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 is the default value and 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.

11. In the **Link Down Time on Failover** field, use the default value of 0.0, or specify a new value.

This setting specifies the amount of time, in seconds, that interfaces for any external VLANs are down when a traffic group fails over and goes to the standby state. Specifying a value other than 0.0 for this setting causes other vendor switches to use the specified time to learn the MAC address of the newly-active device.

Important: This setting is a system-wide setting, and does not apply to this device group only. Specifying a value in this field causes the BIG-IP system to assign this value to the global bigdb variable `failover.standby.linkdowntime`.

12. 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, click the arrow next to 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, choose the device that shows a sync status of *Changes Pending*.
4. In the Sync Options area of the screen, select **Push the selected device configuration to the group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device 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. Near the top of the screen, click **Failover Network**.
5. Click **Add**.
6. From the **Address** list, select an IP address.

The unicast IP address you select depends on the type of device:

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

Platform**Action**

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.

7. From the **Port** list, select a port number.
We recommend using port **1026** for failover communication.
8. 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.
9. 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**.
10. Click **Finished**.

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, click the arrow next to 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, choose the device that shows a sync status of **Changes Pending**.
4. In the Sync Options area of the screen, select **Push the selected device configuration to the group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device 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.

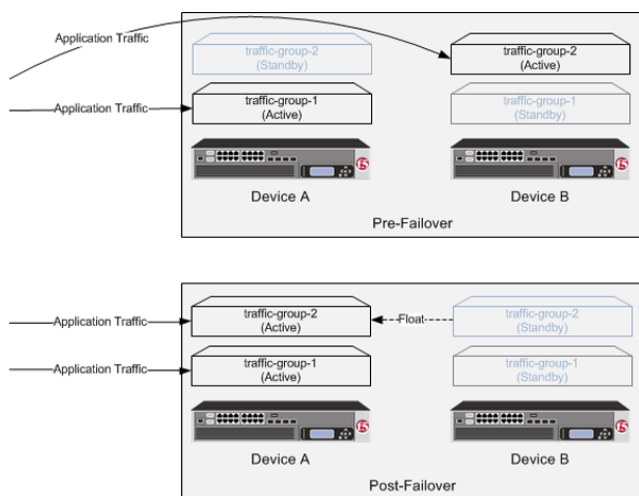


Figure 10: A two-member Sync-Failover group for an active-active configuration

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 10: 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 11: 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 12: 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 13: 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 should be running BIG-IP version 12.x or higher. 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 <code>root</code> folder must be set correctly (<code>Sync-Failover</code> and <code>traffic-group-1</code>).
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><i>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.</i></p> <hr/> <p><i>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.</i></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.
Switchboard failsafe	If your devices are provisioned for vCMP® and your guests are members of a device group, make sure the guests' switchboard failsafe setting is set to the default value. Any change from the default switchboard failsafe configuration must always be done on the vCMP host, and not on the guests.

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.

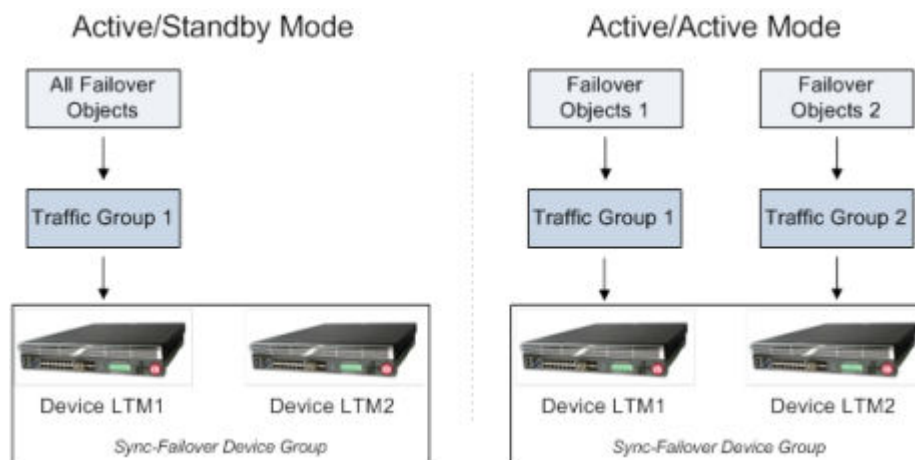


Figure 11: Comparison of Active-Standby and Active-Active device groups

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

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. Near the top of the screen, click **ConfigSync**.
5. From the **Local Address** list, 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.

Important: Connection mirroring only functions between devices with identical hardware platforms.

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. Near the top of the screen, click **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.

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 > Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. 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.
5. Click **Retrieve Device Information**.
6. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
7. In the **Name** field, verify that the name of the remote device is correct.
8. Click **Add Device**.

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. In the **Name** field, type a name for the device group.
4. From the **Group Type** list, select **Sync-Failover**.
5. In the **Description** field, type a short description that you can use to easily identify the device group.
6. For the **Members** setting, 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.
7. 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).
8. 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, click the arrow next to 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, choose the device that shows a sync status of *Changes Pending*.
4. In the Sync Options area of the screen, select **Push the selected device configuration to the group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device 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. Near the top of the screen, click **Failover Network**.
5. Click **Add**.
6. From the **Address** list, select an IP address.

The unicast IP address you select depends on the type of device:

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

7. From the **Port** list, select a port number.
We recommend using port **1026** for failover communication.
8. 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.
9. 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**.

10. Click **Finished.**

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 **Always Failback to First Device if it is Available**:
 - Select the check box to cause the traffic group, after failover, to fail back to the first device in the traffic group's ordered list when that device (and only that device) is available. If that device is unavailable, no failback occurs and the traffic group continues to run on the current device.
 - 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 **Load-Aware** box, select a device name and using the Move button, move the device name to the **Preferred Order** 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, make sure 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 none of the devices in the **Preferred Order** list is currently available when failover occurs, the BIG-IP system uses load-aware failover instead.
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 floating self IP addresses

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, click the arrow next to 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, choose the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Push the selected device configuration to the group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device 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.

Useful troubleshooting tools

Useful command-line troubleshooting tools

The `tmsh` and `tmctl` utilities include commands for troubleshooting device trust and device group operations. For detailed reference material on `tmsh` commands, see the F5 Networks® knowledge base at <http://support.f5.com>.

Table 14: Useful command-line troubleshooting tools

Command	Description
<code>tmsh run cm sniff-updates</code>	Displays the commit ID updates that occur over the configuration management communications channel.
<code>tmsh run cm watch-devicegroup-device</code>	Displays information about the devices in the device group to which the local device belongs.
<code>tmsh run cm watch-sys-device</code>	Displays information about the local device.
<code>tmsh run cm watch-trafficgroup-device</code>	Displays information about the traffic groups associated with devices in a device group.
<code>tmsh sys db configsync.timesyncthreshold</code>	Displays the time threshold for the time difference between devices in the trust domain. If the time difference between devices exceeds the configured threshold value, the BIG-IP system logs an error.
<code>tmsh show cm device</code>	Displays the time difference, in seconds, between the local device and each of the other devices in the trust domain.
<code>tmsh show cm traffic-group</code>	Displays status for all traffic groups on the local device, including the next-active device, the previously-active device, and the reason that an active traffic group is active on its current device. This information is also available with the <code>tmsh cm traffic-group all-properties</code> command.
<code>tmsh show cm sync-status</code>	Displays the current network connection status, either connected or disconnected.
<code>tmsh show sys ha-mirror</code>	Displays the current status of mirrored connections.
<code>tmsh show cm failover-status</code>	In addition to other information, displays log messages when: <ul style="list-style-type: none">• The local device first receives a SOD status message on its unicast addresses and a multicast address/interface (if any).• The local device stops receiving SOD status messages.• An interface on the local device begins receiving SOD status messages again.
<code>tmctl sod_tg_conn_stat</code>	Displays SOD messaging statistics for each type of message sent and received.

Command	Description
<code>tmctl sod_tg_msg_stat</code>	Displays the outgoing packets from the SOD daemon to the other devices in the device group.
<code>tmsh modify cm trust-domain Root ca-devices add { ip_address } name device_name username admin password admin</code>	<p>In addition to adding a device to a trust domain, returns error messages for these conditions:</p> <ul style="list-style-type: none">• A device with the specified device name already exists in the trust domain.• The BIG-IP software version (including hotfix version) on the specified device does not match the version on the local device.• The time on the specified device is out of sync with the current device by some number of seconds.• A config sync address is not configured on the specified device.

Legal Notices

Legal notices

Publication Date

This document was published on May 18, 2017.

Publication Number

MAN-0375-09

Copyright

Copyright © 2017, 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

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: <https://f5.com/about-us/policies/patents>.

Link Controller Availability

This product is not currently available in the U.S.

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

- active bonus values
 - for HA score calculations 48
- active state
 - defined 41
- active-active configuration
 - defined 42
 - described 73
 - result of 82
- active-standby configuration
 - creating 66
 - defined 42
 - described 65
 - result of 70
- address exchange 70
- administrative user accounts
 - configuring 67, 75
- application load
 - and failover 52
 - balancing 51
- applications
 - creating 79, 80
- ARP communications 53
- authentication
 - and device identity 14
 - and local trust domains 13
- authority
 - changing 13
- auto-failback feature
 - defined 50
- automatic synchronization
 - defined 28
 - enabling and disabling 28
- availability
 - during failover 37
- AWS floating IP address 85, 95

B

- base network components 65
- BIG-IP system
 - provisioning 66, 74
- BIG-IP system licenses 66, 74

C

- certificate authority
 - importing 16
 - managing and retaining 16
- certificate signing authorities
 - described 13
 - resetting trust on 16
- certificates
 - for device trust 15
- Changes Pending status
 - about 17, 26, 27
- config sync, *See* configuration synchronization.

- config sync address
 - described 9
- config sync addresses
 - specifying 25, 86, 97
 - See also* configuration synchronization
- config sync properties
 - advanced 34
- config sync status
 - determining 17, 26
 - displaying 30, 32
 - troubleshooting 31
 - viewing 22
- config sync types
 - defined 29
- configsync
 - configuring for VIPRION systems 83, 93
- configuration objects
 - and traffic group associations 38
 - and traffic groups 73
- configuration synchronization
 - about 25
 - and Setup utility 65
 - automating 28
 - preventing 38
 - scope of 61
 - syncing to group 27, 81, 89, 91, 99, 102
- connection mirroring
 - about 9
 - and SNATs 58
 - configuring 57, 87, 97
 - configuring for VIPRION 58
 - considerations for 55
 - enabling 68, 75
 - for TCP and UDP connections 58
- connection mirroring addresses, *See* mirroring addresses
- connections
 - preserving on failover 57, 87, 97

D

- default traffic groups
 - described 37
- device availability
 - defined 37
- device certificate configuration 67, 75
- device discovery
 - defined 14
 - for device trust 14, 15, 78, 87, 98
 - of peer devices 70
- device group assignments
 - to root and /Common folders 20
- device group attribute
 - described 61
 - viewing on root folder 62
- device group members
 - adding and viewing 22
- device group membership 20
- device group subset 61

- device groups
 - and root folder 62
 - configuration restrictions for 20
 - configuring for VIPRION systems 83, 93
 - creating 20, 79, 88, 99
 - defined 7
 - sync status for 31
 - viewing 22
- device identity
 - defined 14
- device objects
 - defined 7
- device properties 10
- device service clustering
 - about 7
- device status types
 - described 11
 - viewing 11
- device trust
 - about 13
 - adding domain members 15
 - configuring for VIPRION systems 83, 93
 - defined 7
 - establishing 14, 78, 87, 98
- device utilization
 - about 51
- devices
 - and mirroring limit 57, 87, 97
 - defined 7
 - discovering 14
 - excluding from config sync 25
 - selecting for failover 37
- dropped connections 56
- DSC deployment worksheet 85, 95

E

- errors
 - with device trust 16
- external network
 - configuring 68, 76

F

- failback
 - defined 50
- failover
 - about 35
 - and default traffic groups 37
 - and dropped packets 53
 - and failback 50
 - and HA groups 44
 - and HA scores 43, 47
 - and next-active device 51
 - and ordered lists 50
 - and traffic groups 37
 - configuring for VIPRION systems 83, 93
 - preventing 48
 - scope of 61
- failover addresses
 - configuring 70
 - exchanging during discovery 70

- failover addresses (*continued*)
 - specifying 77
- failover devices
 - about 46
- failover IP addresses
 - about 9
 - specifying 36, 90, 100
 - types of 35
- failover methods
 - specifying 68, 75
- failover objects
 - adding 40
 - associating with traffic groups 38
 - viewing 40
- failover states
 - viewing 42
- failover status
 - of traffic groups 40
- fast failover 43
- floating IP address
 - for AWS 85, 95
- floating IP addresses
 - and traffic groups 37
 - configuring 68, 76
- floating traffic groups
 - and traffic group states 41
 - defined 35
- folder attributes
 - described 61
- folder hierarchy 62
- folder inheritance 19
- folders
 - and traffic group associations 38
 - associating device groups with 23
 - defined 7, 61
- Force to Standby option 38
- FQDN (fully-qualified domain name) 67, 75
- FTP connections
 - and mirroring 55
- full sync
 - defined 29
- fully-qualified domain name (FQDN) 67, 75

G

- granular synchronization
 - about 25
 - with folders 61
- gratuitous ARPs 53

H

- HA groups
 - about 35, 43
 - and active bonus values 48
 - and matching health scores 49
 - and threshold values 47
 - and weight values 47
 - configuring 43, 44
 - example of 45
 - purpose of 43
- HA health scores

HA health scores (*continued*)

 matching 49

HA load factor

 about 52

 viewing 40

HA score failover method

 defined 47, 63

HA scores

 and active bonus values 48

 and threshold values 47

 and weight values 47

 calculating 43, 47

 purpose of 43

HA traffic load

 about 52

health scores 43

high availability

 and VLANs 69, 77

 enabling 68, 75

HTTP connections

 and mirroring 55

I

iApp applications

 creating 79, 80

iApps applications

 and traffic group associations 38

 and traffic groups 37

incremental sync

 defined 29

information exchange 14

interfaces

 and external VLAN configuration 68, 76

 and HA VLAN configuration 69, 77

 and internal VLAN configuration 68, 76

internal network

 configuring 68, 76

IP addresses

 as traffic group members 40

 for failover 35

 for redundancy 9

L

licenses

 activating 66, 74

load-aware failover

 about 35, 51

local trust domain

 and device group members 22

 and device groups 20, 79, 88, 99

 defined 13–15, 78, 87, 98

 joining 14

M

MAC masquerade addresses

 defined 53

management IP addresses

 and ConfigSync 69, 77

management port

management port (*continued*)

 configuring 67, 75

 specifying for failover 77

manual synchronization 28

mirroring

 configuring for VIPRION 58

 configuring for VIPRION systems 83, 93

 considerations for 55

 of connections 56

 See also connection mirroring

mirroring addresses

 configuring 70

 exchanging 70

 specifying 77

mirroring IP addresses 9

mirroring tasks 56

N

network failover

 configuring 20, 79, 88, 99

 specifying 68, 75

next-active devices

 about 46

 and failback 50

 and failover 44, 50

 choosing 35

 selection of 47, 63

O

object referencing 23

ordered failover lists

 about 50

 configuring 50

ordered lists

 about 35

P

peer authorities

 described 13

peer devices

 and traffic groups 80

 discovering 70

persistence mirroring

 about 56

persistence records

 mirroring 59

problems

 with device trust 16

profiles

 and persistence mirroring 59

R

redundancy attributes

 configuring 62

redundant system configuration

 described 7

relative load value

- relative load value (*continued*)
 - viewing 40
- relative traffic load values 52
- root folder 62
- root folder attributes
 - configuring 62
 - viewing 62

S

- self IP addresses
 - and traffic group associations 38
 - assigning to traffic group 102
 - for external network 68, 76
 - for HA network 69, 77
 - for internal network 68, 76
- self-signed certificates
 - regenerating 16
- serial cable failover
 - specifying 68, 75
- service interruptions 56
- Setup utility
 - and base network 73
 - and base network configuration 68, 69, 76, 77
 - and device discovery 70
 - for active-standby configurations 65, 66
 - using 74
- SNAT translation addresses
 - and traffic groups 73
- SNATs
 - and mirroring 58
 - and traffic group associations 38
- solutions
 - for device trust 16
- standby state
 - defined 41
 - forcing to 42, 81, 102
- static self IP addresses
 - and traffic groups 37
- status
 - for config sync 26
- status types
 - for devices 11
 - viewing 11
- subordinate non-authorities
 - defined 13
 - described 13
 - resetting trust on 16
- sufficient threshold value
 - described 48
- sync status
 - determining 17, 26
 - displaying 32
 - of configuration 32
 - troubleshooting 31
- sync types
 - defined 28, 29
- Sync-Failover configuration
 - example of 19
- Sync-Failover device groups
 - about 19
 - creating 20, 79, 88, 99

- Sync-Failover device groups (*continued*)
 - illustrated 96
- synchronization 25
- system provisioning 66, 74

T

- target failover devices
 - about 35, 46
- TCP connections
 - mirroring of 58
- Telnet connections
 - and mirroring 55
- threshold values
 - for HA score calculation 47
- traffic group attribute
 - described 61
 - viewing on root folder 62
- traffic group members
 - adding 40
- traffic group properties 40
- traffic group states
 - defined 41
 - viewing 42
- traffic groups
 - activating 38
 - and defaults 37
 - and failover 37
 - and failover objects 38, 40
 - and iApp applications 79, 80
 - and root folder 62
 - as defaults 73
 - assigning MAC masquerade addresses to 53
 - associating objects with 38
 - balancing load of 42
 - configuration restrictions for 38
 - creating 83, 86, 93, 96, 101
 - default name of 65
 - defined 7, 37
 - for remote devices 42, 80, 81, 102
 - forcing to standby state 38, 42, 81, 102
 - inheriting 38
 - maximum number supported 37
 - specifying load for 52
 - viewing list of 40
- traffic load
 - balancing 51
- traffic load values 52
- troubleshooting tips
 - for device trust 16
- troubleshooting tools 105
- trust authority
 - managing and resetting 16
- trust domains
 - and local trust domain 13–15, 78, 87, 98
- trust relationships
 - between devices 13
 - establishing 65
- trusted peers
 - and address exchange 70

U

- UDP connections
 - and mirroring [55](#)
 - mirroring of [58](#)
- unicast failover addresses, *See* unicast IP addresses
- unicast IP addresses
 - specifying for failover [77](#)

V

- vCMP systems
 - connection mirroring for [55](#)
- VIPRION systems
 - and connection mirroring [58](#)
 - connection mirroring for [55](#)
 - mirroring connections for [56](#)
 - mirroring connections on [9](#)
- virtual addresses
 - assigning to traffic group [101](#)
- virtual IP addresses
 - and traffic group associations [38](#)
 - and traffic groups [73](#)
- virtual servers
 - and mirroring [58](#)
- VLAN IDs
 - configuring [68](#), [69](#), [76](#), [77](#)
- VLAN tags, *See* VLAN IDs
- VLANs
 - and traffic groups [37](#)

W

- weight values
 - for HA score calculations [47](#)

X

- x509 certificates
 - and device identity [14](#)
 - and device trust [13](#)
 - for device trust [14](#), [78](#), [87](#), [98](#)

