VIPRION® Systems: Configuration

Version 11.2.1

IT agility. Your way.

Table of Contents

Legal Notices

Publication Date

This document was published on September 28, 2012.

Publication Number

MAN-0312-03

Copyright

Trademarks

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.

Acknowledgments

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

This product includes software developed by the Politecnico di Torino, and its contributors.

This product includes software developed by the Swedish Institute of Computer Science and its contributors.

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the Computer Systems Engineering Group at the Lawrence Berkeley Laboratory.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

This product includes software developed by Jochen Pohl for the NetBSD Project.

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

This product includes software developed by SigmaSoft, Th. Lockert.

This product includes software developed for the NetBSD Project by Jason R. Thorpe.

This product includes software developed by Jason R. Thorpe for And Communications, http://www.and.com.

This product includes software developed for the NetBSD Project by Frank Van der Linden.

This product includes software developed for the NetBSD Project by John M. Vinopal.

This product includes software developed by Christos Zoulas.

This product includes software developed by the University of Vermont and State Agricultural College and Garrett A. Wollman.

This product includes software developed by Balazs Scheidler (bazsi@balabit.hu), which is protected under the GNU Public License.

This product includes software developed by Niels Mueller (nisse@lysator.liu.se), which is protected under the GNU Public License.

In the following statement, *This software* refers to the Mitsumi CD-ROM driver: This software was developed by Holger Veit and Brian Moore for use with 386BSD and similar operating systems. *Similar operating systems* includes mainly non-profit oriented systems for research and education, including but not restricted to NetBSD, FreeBSD, Mach (by CMU).

This product includes software developed by the Apache Group for use in the Apache HTTP server project (http://www.apache.org/).

This product includes software licensed from Richard H. Porter under the GNU Library General Public License (© 1998, Red Hat Software), www.gnu.org/copyleft/lgpl.html.

This product includes the standard version of Perl software licensed under the Perl Artistic License (© 1997, 1998 Tom Christiansen and Nathan Torkington). All rights reserved. You may find the most current standard version of Perl at http://www.perl.com.

This product includes software developed by Jared Minch.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product contains software based on oprofile, which is protected under the GNU Public License.

This product includes RRDtool software developed by Tobi Oetiker (http://www.rrdtool.com/index.html) and licensed under the GNU General Public License.

This product contains software licensed from Dr. Brian Gladman under the GNU General Public License (GPL).

This product includes software developed by the Apache Software Foundation (http://www.apache.org/).

This product includes Hypersonic SQL.

This product contains software developed by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and others.

This product includes software developed by the Internet Software Consortium.

This product includes software developed by Nominum, Inc. (http://www.nominum.com).

This product contains software developed by Broadcom Corporation, which is protected under the GNU Public License.

This product contains software developed by MaxMind LLC, and is protected under the GNU Lesser General Public License, as published by the Free Software Foundation.

Chapter 1

VIPRION Overview

Topics:

VIPRION overview

The VIPRION® system is a complete traffic management solution that offers high performance, reliability, scalability, and ease of management. Based on chassis and blade technology, this system is designed to meet the needs of large, enterprise networking environments that normally require multiple BIG-IP® systems to process large volumes of application traffic.

The VIPRION system consists of a chassis with a four-blade capacity. The four blades work together as a powerful system to process application traffic. Traffic comes into a single virtual server, and the system distributes that traffic over multiple blades, using the full multi-processing capacity of each blade. Moreover, if a blade unexpectedly becomes unavailable, another blade can complete the processing of the request.

VIPRION features

This table describes the VIPRION® system features.

Feature	Description
A chassis with blades	The multi-slot chassis significantly reduces the amount of rack space required for the BIG-IP® systems by housing blades instead of traditional switch systems. Hardware resources such as cooling and power systems, normally required for individual BIG-IP systems, are now part of the chassis instead.
Cluster technology	The VIPRION system's SuperVIP™ cluster technology is the core feature that coordinates all of the blades into a single high-performance system. A SuperVIP cluster is the group of slots in the VIPRION system chassis. Each slot in the cluster represents a cluster member, and any blades that you insert into the slots of a cluster work together to process application traffic. Cluster technology provides

Feature	Description
	the processing power of multiple blades, but you manage the entire cluster as a single system.
Live installation	When you upgrade the BIG-IP software on a running system, the system automatically upgrades the BIG-IP software on all blades in the cluster.
Cluster synchronization	The primary blade automatically propagates the system configuration to all secondary blades, even when a new blade is introduced into the cluster.
Connection mirroring	Connection mirroring ensures that if a blade, or a cluster within a device service clustering (redundant system) configuration, becomes unavailable, the system can still process any existing connections.

Related documents

You may find it useful to have an understanding of certain background concepts before performing VIPRION® configuration tasks.

- For more information about configuring required BIG-IP® network objects (trunks, VLANs, and self IP addresses), refer to the *BIG-IP® TMOS®: Concepts Guide*.
- For more information about configuring the BIG-IP system (or vCMP® guests) to manage local area network traffic (concepts pertaining to virtual servers, various types of traffic profiles, load balancing pools and pool members, and so on) refer to the *BIG-IP® Local Traffic Manager: Concepts Guide*.

These product guides are available from the AskF5 Knowledge Base web site, `http://support.f5.com`.

Chapter 2

Initial VIPRION Setup

Topics:

Overview: Initial VIPRION setup

To set up a newly installed VIPRION® chassis, you configure a number of BIG-IP® system objects. First trunks (external and internal), then VLANs (external, internal, and high availability), and finally, self IP addresses.

This illustration depicts a VIPRION chassis configured with a single active blade.

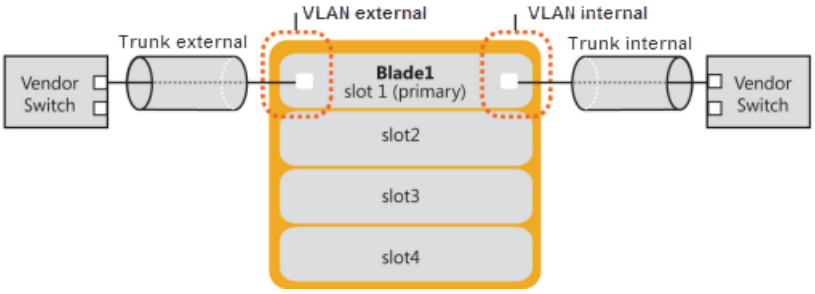

Task summary

Activating the BIG-IP license for VIPRION
Creating trunks
Creating VLANs
Creating self IP addresses for VLANs

VIPRION deployment worksheet

There are a number of points during the VIPRION® deployment process at which you will need to make decisions or provide values. Use this table as a prompt for gathering the answers and values you will need, so that you can provide them when performing the initial setup.

Configuration component	Considerations
External gateway address	What is the gateway address (next hop) for external traffic?
FQDN	What is the fully-qualified domain name (FQDN) for your BIG-IP® system?
Link aggregation control protocol	Do your trunks require LACP mode?
Network mask	What is the network mask?
Primary cluster IP address	What is the primary cluster IP address? The management IP address assigned to the chassis' primary cluster during chassis installation is used to access the VIPRION.
User role	Do you have a user role of Administrator? You need to have a user role of Administrator to perform the tasks in this process.

Activating the BIG-IP license for VIPRION

To activate the BIG-IP® license, you need access to a browser and the base registration key. The *base registration key* is a character string that the license server uses to verify the type and number of F5 Networks products that you are entitled to license. If you do not have a base registration key, contact the F5 Networks sales group (`http://www.f5.com`).

You activate the BIG-IP license from the Setup Utility License screen.

1. From a workstation attached to the network on which you configured the management interface, use a browser and type the following URL syntax where `<management_IP_address>` is the address you configured for device management:
   `https://<management_IP_address>`
2. At the prompts, type the user name `admin` and the password `admin`.
3. Click **Log in**.
   The Setup Utility screen opens.
4. Click **Activate**.
   The License screen opens.
5. In the **Base Registration Key** field, paste your base registration key.
6. Click **Next**.
   The End User License Agreement (EULA) displays.
7. Review the EULA.
   When you click **Accept**, the Platform screen opens.

Creating trunks

To configure trunks for the VIPRION® system, the four external interfaces must be cabled to your Internet gateway, external bridge, or vendor switch.

The first objects you configure are trunks that tie the internal and external vendor switch interfaces to the corresponding VIPRION interfaces.

1. Use a browser to log in to the VIPRION® chassis's management IP address.
   This logs you in to the floating IP address for the cluster.
2. On the peer (vendor) switch on the external network, create a trunk that includes the four external interfaces to which you have physically connected the external interfaces of the four blades.

   If the peer switch is configured to use Link Aggregation Control Protocol (LACP), you must enable LACP.
3. Create a trunk, and if the peer switch is configured to use LACP, enable LACP on the new trunk:
   a) On the Main tab, expand **Network**, and click **Trunks**.
      The Trunks screen opens.
   b) At the upper right corner of the screen, click **Create**.
      The New Trunk screen opens.
   c) Assign the name `trunk_ext`, and assign an external interface of `blade 1` to the trunk.
   d) Enable LACP mode, if required.
   e) Click **Finished**.
4. Repeat the previous step, but this time, configure a trunk that ties the vendor switch internal interface to the VIPRION internal interface. Assign the name **trunk_int**.

Creating VLANs

VLANs associate with your trunks.

1. Use a browser to log in to the VIPRION® chassis's management IP address.
   This logs you in to the floating IP address for the cluster.
2. On the Main tab, expand **Network**, and click **VLANs**.
   The VLANs screen opens.
3. Click **Create**.
   The New VLAN screen opens.
4. Configure a VLAN named `external`, and assign it to the trunk named `trunk_ex` as an untagged interface.
5. Click **Finished**.
6. Repeat the last three steps, but this time, configure a VLAN named `internal`, and assign it to the trunk named `trunk_int`.
7. Repeat steps 3 through 5 one more time, but this time, configure a VLAN named `HA`, assign it to the trunk named `trunk_int` as a tagged interface.

Creating self IP addresses for VLANs

You need at least one VLAN or VLAN group configured before you create a self IP address.

Self IP addresses enable the BIG-IP® system, and other devices on the network, to route application traffic through the associated VLAN or VLAN group. Repeat the steps in this task for each VLAN.

1. On the Main tab, click **Network** > **Self IPs**.
   The Self IPs screen opens.
2. Click **Create**.
   The New Self IP screen opens.
3. In the **Name** field, type a unique name that readily identifies the VLAN to which it will associate for the self IP.

   Name the self IP for the internal VLAN `Internal`, name the external VLAN `External`, and name the HA VLAN `HA`.
4. In the **IP Address** field, type an IP address.

   This IP address must be within the address space that corresponds to the VLAN for which it is created (Internal, External or HA).

   The system accepts IP addresses in both the IPv4 and IPv6 formats.
5. In the **Netmask** field, type the network mask for the specified IP address.
6. From the **VLAN/Tunnel** list, select the VLAN to associate with this self IP address:
   - For the internal network, select the VLAN that is associated with an internal interface or trunk.
   - For the external network, select the VLAN that is associated with an external interface or trunk.
   - For the HA network, select the VLAN that is associated with an internal interface or trunk.
7. From the **Port Lockdown** list, select **Allow Default**.
8. Repeat the last 4 steps, but this time specify an address from your external network in step 4 and select the VLAN named `external` in step 6.

9. Repeat steps 3 through 7 one more time, but this time specify an address on your internal network in step 4 and select the VLAN named `HA` in step 6.
10. Click **Finished**.
    The screen refreshes, and displays the new self IP address in the list.

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

Overview: Verifying initial VIPRION configuration

Verifying your VIPRION configuration confirms that the setup performed up to this point is functioning properly. Once you establish that the VIPRION® configuration is correct, you will likely need to create a profile, pools, and virtual server that are tailored to your network topology before you can begin processing LTM® traffic.

Creating a pool to manage HTTP traffic
Creating a virtual server to manage HTTP traffic

Creating a pool to manage HTTP traffic

You can create a pool to manage HTTP connections.

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, from the **Available** list, select the **http** monitor, and click << to move the monitor to the **Active** list.
5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.

   The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
   - Select **Disabled** to disable priority groups. This is the default option.
   - Select **Less than**, and in the **Available Members** field, type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:
   a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
   b) Type `80` in the **Service Port** field, or select **HTTP** from the list.
   c) (Optional) Type a priority number in the **Priority** field.
   d) Click **Add**.
8. Click **Finished**.

The new pool appears in the Pools list.

Creating a virtual server to manage HTTP traffic

You can create a virtual server to manage HTTP traffic as either a host virtual server or a network virtual server.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen displays a list of existing virtual servers.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, in the **Address** field, type the IP address you want to use for the virtual server.

   The IP address you type must be available and not in the loopback network.
5. In the **Service Port** field, type `80`, or select **HTTP** from the list.
6. From the **HTTP Profile** list, select **http**.
7. In the Resources area of the screen, from the **Default Pool** list, select a pool name.
8. Click **Finished**.

The HTTP virtual server appears in the list of existing virtual servers on the Virtual Server List screen.

Chapter 3

Create an Active-Standby Configuration

Topics:

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 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 the peer device, and application processing is handled by that device.

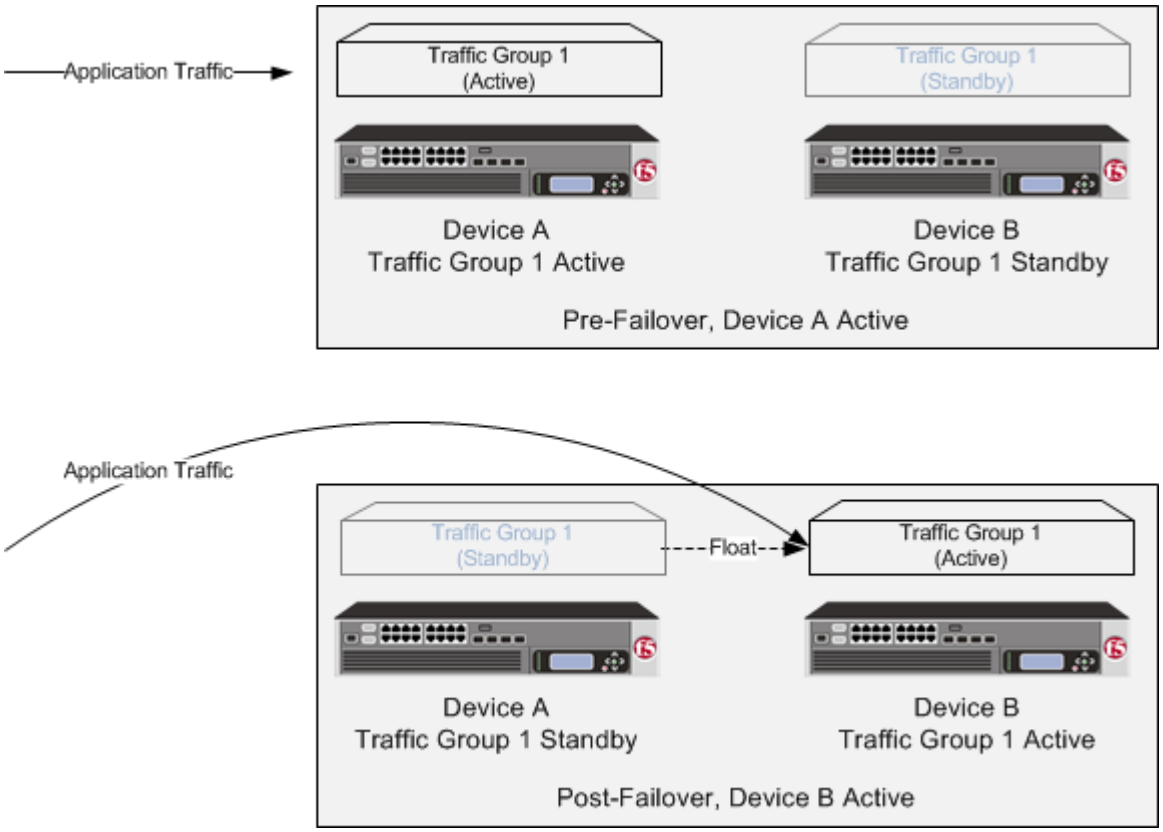

Figure 1: A Sync-Failover device group for an active-standby configuration

About DSC configuration on a VIPRION system

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

For non-vCMP systems

On a VIPRION system that is not provisioned for vCMP, the management IP address that you specify for establishing device trust and enabling failover should be the system's primary cluster IP address. This is a floating management IP address.

For vCMP systems

On a vCMP system, the devices in a device group are virtual devices, known as *vCMP guests*. You configure config sync and failover 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 1: Sample device groups for two VIPRION systems with vCMP

Device groups for vCMP	Device group members
`Device-Group-A`	• `Guest1` on chassis1 • `Guest1` on chassis2
`Device-Group-B`	• `Guest2` on chassis1 • `Guest2` on chassis2
`Device-Group-C`	• `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 self IP addresses that you specify per guest for config sync and failover should be the self IP addresses that you previously configured on the guest (not the host). Similarly, the management IP address that you specify per guest for device trust and failover should be the cluster IP address of the guest.

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 2: DSC deployment worksheet

Configuration component	Considerations
Hardware, licensing, and provisioning	Devices in a device group must match as closely as possible with respect to hardware platform, product licensing, and module provisioning. If you want to configure mirroring, ensure that the hardware platforms of the mirrored devices match.
BIG-IP software version	Each device must be running BIG-IP version 11.x. This ensures successful configuration synchronization.

Configuration component	Considerations
Management IP addresses	Each device must have a management IP address, a network mask, and a management route defined.
FQDN	Each device must have a fully-qualified domain name (FQDN) as its host name.
User name and password	Each device must have a user name and password defined on it that you will use when logging in to the BIG-IP Configuration utility.
`root` folder properties	The platform properties for the `root` folder must be set correctly (`Sync-Failover` and `traffic-group-1`).
VLANs	You must create these VLANs on each device, if you have not already done so: • A VLAN for the internal network, named `internal` • A VLAN for the external network, named `external` • A VLAN for failover communications, named `HA`
Self IP addresses	You must create these self IP addresses on each device, if you have not already done so: • Two self IP addresses (floating and non-floating) on the same subnet for VLAN `internal`. • Two self IP addresses (floating and non-floating) on the same subnet for VLAN `external`. • A non-floating self IP address on the internal subnet for VLAN `HA`.  ***Note:*** *When you create floating self IP addresses, the BIG-IP system automatically adds them to the default floating traffic group,* `traffic-group-1`*. To add a self IP address to a different traffic group, you must modify the value of the self IP address* ***Traffic Group*** *property.*
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**.
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, `traffic-group-1`.
Time synchronization	The times set by the NTP service on all devices must be synchronized. This is a requirement for configuration synchronization to operate successfully.
Device certificates	Verify that each device includes an x509 device certificate. Devices with device certificates can authenticate and therefore trust one another, which is a prerequisite for device-to-device communication and data exchange.

Task summary

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

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

Task list

Specifying an IP address for config sync
Specifying IP addresses for connection mirroring
Establishing device trust
Creating a Sync-Failover device group
Syncing the BIG-IP configuration to the device group
Specifying IP addresses for failover
Syncing the BIG-IP configuration to the device group

Specifying an IP address for config sync

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

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

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

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

After performing this task, the other devices in the device group can sync their configurations to the local device.

Specifying IP addresses for connection mirroring

Before configuring mirroring addresses, verify that the mirroring peers have the same hardware platform.

This task configures connection mirroring between two devices to ensure that in-process connections are not dropped when failover occurs. You can mirror connections between a maximum of two devices in a device group.

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

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

   The recommended IP address is the self IP address for either VLAN `HA` or VLAN `internal`.
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**.

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 `A`, `B`, and `C` 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 `A` and add devices `B` and `C` to the local trust domain. Note that there is no need to repeat this process on devices `B` and `C`.

1. On the Main tab, click **Device Management** > **Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type an IP address, administrator user name, and administrator password for the remote BIG-IP® device.

   This IP address can be either a management IP address or a self IP address.
4. Click **Retrieve Device Information**.

5. Verify that the certificate of the remote device is correct.
6. Verify that the name of the remote device is correct.
7. Verify that the management IP address and name of the remote device are correct.
8. Click **Finished**.

The device you added is now a member of the local trust domain.

Repeat this task for 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 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 perform this task on any one of the authority devices within the local trust domain.

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

1. On the Main tab, click **Device Management** > **Device Groups**.
2. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **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. A device can be a member of one Sync-Failover group only.
5. For the **Network Failover** setting:
   - Select the **Enabled** check box if you want device group members to handle failover communications by way of network connectivity.
   - Clear the **Enabled** check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

   Serial failover is not available for device groups with more than two members.
6. 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 has been 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, in the Name column, select the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
   The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Specifying IP addresses for failover

This task specifies the local IP addresses that you want other devices in the device group to use for failover communications with the local device. You must perform this task locally on each device in the device group.

Note: *The failover addresses that you specify must belong to route domain* `0`.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management** > **Devices**.
   This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Failover.
5. For the Failover Unicast Configuration settings, retain the displayed IP addresses.

   You can also click **Add** to specify additional IP addresses that the system can use for failover communications. F5 Networks recommends that you use the self IP address assigned to the `HA` VLAN.
6. If the BIG-IP® system is running on a VIPRION® platform, then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
7. If you enable **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.

   If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

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

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust has been 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, in the Name column, select the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
   The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data 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.

Chapter 4

Understanding Clusters

Topics:

Cluster overview

The slots in a VIPRION® chassis work together as a single, powerful unit. This entity is called a *cluster*. The size of the cluster depends on the number of running blades installed in the chassis. Blades in the cluster share the overall workload, and can be configured to mirror each others' connections, so that if a blade is taken out of service or becomes unavailable for some reason, any in-process connections remain intact.

When a blade is installed in a slot and turned on, it automatically becomes a member of the cluster.

One of the first tasks performed as part of the platform installation is to insert blades and assign a unique cluster IP address to the primary blade in the cluster. The cluster IP address is a floating management IP address used to access the primary blade to configure the system. If the primary blade becomes unavailable for any reason, the primary designation moves to a different blade, and the cluster IP address floats to that blade. This ensures that you can always access the cluster using the cluster IP address, even when the primary blade changes.

When you log on to the system using the cluster IP address, you can configure features such as trunks, VLANs, administrative partitions, and virtual servers. If you have a redundant system configuration, you can configure failover IP addresses, as well as connection mirroring between clusters.

Viewing cluster properties

You can use the BIG-IP® Configuration utility to view the properties for the cluster.

1. Use a browser to log in to the VIPRION® chassis's management IP address.
   This logs you in to the floating IP address for the cluster.
2. On the Main tab, click **System** > **Clusters**.
   The Cluster screen opens, showing the properties of the cluster, and listing the cluster members.

Cluster properties

The Cluster screen displays the properties of the cluster.

Property	Description
Name	Displays the name of the cluster.
Cluster IP Address	Displays the IP address assigned to the cluster. Click this IP address to change it.
Network Mask	Displays the network mask for the cluster IP address.
Primary Member	Displays the number of the slot that holds the primary blade in the cluster.
Software Version	Displays the version number of the BIG-IP® software that is running on the cluster.
Software Build	Displays the build number of the BIG-IP software that is running on the cluster.

Property	Description
Hotfix Build	Displays the build number of any BIG-IP software hotfix that is running on the cluster.
Chassis 400-level BOM	Displays the bill-of-materials (BOM) number for the chassis.
Status	Displays an icon and descriptive text that indicates whether there are sufficient available members of the cluster.

Viewing cluster member properties

You can use the BIG-IP® Configuration utility to view the properties for cluster members.

1. Use a browser to log in to the VIPRION® chassis's management IP address.
   This logs you in to the floating IP address for the cluster.
2. On the Main tab, click **System** > **Clusters**.
   The Cluster screen opens, showing the properties of the cluster, and listing the cluster members.
3. To display the properties for one cluster member, click the slot number of that member.
   The Cluster Member properties screen opens, showing the properties of that member.

Cluster member properties

In addition to displaying the properties of the cluster, the Cluster screen also lists information about members of the cluster. The table lists the information associated with each cluster member.

Property	Description
Status	The Status column indicates whether the cluster member is available or unavailable.
Slot	The Slot column indicates the number of the slot. Click this number to display the properties of that cluster member.
Blade serial number	The Blade Serial Number column displays the serial number for the blade currently in that slot.
Enabled	The Enabled column indicates whether that cluster member is currently enabled.
Primary	The Primary column indicates whether that cluster member is currently the primary slot.
HA State	The HA State column indicates whether the cluster member is used in a redundant system configuration for high availability.

Enabling and disabling cluster members

To gracefully drain the connections from a cluster member before you take that blade out of service, you can mark that cluster member disabled. Before you can return that member to service, you need to enable it.

Important: *Perform this task while logged in to the vCMP® host; not from a guest.*

1. Use a browser and the cluster management IP address of the vCMP® host to log in to the vCMP host (hypervisor) and access the BIG-IP® Configuration utility.
2. On the Main tab, click **System** > **Clusters**.
   The Cluster screen opens, showing the properties of the cluster, and listing the cluster members.
3. Locate the cluster member you want to enable or disable, and select the box to the left of the Status icon.
4. Click **Enable** or **Disable/Yield**.

Changing a cluster-related management IP address

You can use the BIG-IP® Configuration utility to view or change the properties for a vCMP® cluster.

Important: *Perform this task while logged in to the vCMP host; not from a guest.*

1. Use a browser and the cluster management IP address of the vCMP® host to log in to the vCMP host (hypervisor) and access the BIG-IP® Configuration utility.
2. On the Main tab, click **System** > **Clusters**.
   The Cluster screen opens, showing the properties of the cluster, and listing the cluster members.
3. On the menu bar, click **Management IP Address**.
   The Management IP Address screen opens.
4. Locate the specific management IP address or cluster member IP address that you would like to change, and type the new IP address.
5. Click **Update**.

The specific management IP address or cluster member IP address that you edited is changed. You can now use that new address to access the cluster.

Cluster-related IP addresses

The cluster-related addresses that you can modify are defined in the table.

Setting Type	Setting	Description
Cluster IP address	**IP Address**	Specifies the management IP address that you want to assign to the cluster. This IP address is used to access the Configuration utility, as well as to function as a cluster identifier for the peer cluster in a device service clustering configuration.

Setting Type	Setting	Description
Cluster IP address	**Network Mask**	Specifies the network mask for the cluster IP address.
Cluster IP address	**Management Route**	Specifies the gateway for the cluster IP address. Typically, this is the default route.
Cluster Member IP Address	**Slot 1 IP Address**	Specifies the management IP address associated with slot 1 of the cluster. You can also set this value to **None**.
Cluster Member IP Address	**Slot 2 IP Address**	Specifies the management IP address associated with slot 2 of the cluster. You can also set this value to **None**.
Cluster Member IP Address	**Slot 3 IP Address**	Specifies the management IP address associated with slot 3 of the cluster. You can also set this value to **None**.
Cluster Member IP Address	**Slot 4 IP Address**	Specifies the management IP address associated with slot 4 of the cluster. You can also set this value to **None**.

Index

X