

vCMP[®] for VIPRION[®] Systems: Administration

Version 13.1



Table of Contents

Introduction to the vCMP System.....	7
What is vCMP?.....	7
Other vCMP system components.....	8
Supported BIG-IP system versions.....	9
BIG-IP license considerations for vCMP.....	9
About vCMP provisioning.....	9
Network isolation.....	10
System administration overview.....	10
Guest access to the management network.....	11
About bridged guests.....	11
About isolated guests.....	11
About Appliance mode.....	12
User access restrictions with Appliance mode.....	12
BIG-IP version restrictions with Appliance mode.....	12
 Additional Network Considerations.....	 15
Network separation of Layer 2 and Layer 3 objects.....	15
About the VLAN publishing strategy.....	15
Overview of VLAN subscription.....	15
About VLANs with identical tags and different names.....	16
About VLANs with identical names and different tags.....	17
Solution for tag discrepancy between host and guest VLANs.....	18
About the VLAN MTU setting.....	20
Interface assignment for vCMP guests.....	20
Management IP addresses for bridged guests.....	21
 Flexible Resource Allocation.....	 23
What is flexible resource allocation?.....	23
Resource allocation planning.....	23
Prerequisite hardware considerations.....	23
Understanding guest resource requirements.....	23
About core allocation for a guest.....	24
About core fragmentation when deploying a guest.....	24
Formula for host memory allocation to a guest.....	25
About slot assignment for a guest.....	25
About single-core guests.....	26
Scalability considerations.....	26
About initial slot assignment.....	26
About changing slot assignments.....	27
Effect of blade removal on a guest.....	27
Effect of blade re-insertion on a guest.....	29
Network throughput for guests.....	29
SSL resource allocation.....	30
About compression resource allocation.....	30
Guest states and resource allocation.....	30
 Deployment Examples.....	 33
Example: A single-slot LTM guest on a standalone system.....	33

Example: Dual-slot LTM guests within a device group.....	33
Example: Multi-slot guests within device groups.....	34
Device Service Clustering for vCMP Systems.....	37
Overview: Device service clustering for vCMP systems.....	37
Required IP addresses for DSC configuration.....	38
Failover methods for vCMP guests.....	39
About HA groups for vCMP systems.....	39
About connection mirroring for vCMP systems.....	40
About switchboard failsafe and vCMP guests.....	40
Initial vCMP Configuration Tasks.....	41
About vCMP application volume management.....	41
vCMP host administrator tasks.....	41
Accessing the vCMP host.....	41
Creating a vCMP guest.....	42
Setting a vCMP guest to the Deployed state.....	44
vCMP guest administrator tasks.....	45
Provisioning BIG-IP modules within a guest.....	45
Specifying cluster member IP addresses for a guest.....	46
Creating a self IP address for application traffic.....	46
Changing the MTU value on a VLAN (optional).....	46
Next steps.....	47
Configuration results.....	47
CPU Core Upgrade for Deployed Guests.....	49
About increasing CPU cores for a deployed guest.....	49
Increasing CPU cores for a deployed guest.....	49
Managing vCMP Virtual Disks.....	51
Overview: Managing virtual disks.....	51
About virtual disk allocation.....	51
About virtual disk images.....	51
About virtual disk templates.....	51
Viewing the list of virtual disk templates.....	52
Deleting virtual disk templates.....	52
Enabling and disabling the virtual disk template feature.....	53
Viewing the virtual disk templates db variable.....	53
About virtual disk detachment and re-attachment.....	54
Detaching virtual disks from a vCMP guest.....	54
Viewing virtual disks not attached to a vCMP guest.....	54
Attaching a detached virtual disk to a vCMP guest.....	54
About virtual disk migration.....	55
Deleting a virtual disk from the BIG-IP system.....	55
Deleting a vCMP application volume.....	55
Installing ISO images within vCMP guests.....	57
About ISO images.....	57
Viewing a list of host ISO images from within a guest.....	57
Installing a host ISO image from within a guest.....	58
Installing a host ISO image from within a guest using tmsh.....	58

Viewing vCMP Guest Status	59
About guest status.....	59
Viewing summary status for all guests.....	59
Viewing software status for a guest.....	60
Viewing resource provisioning for a guest.....	60
Viewing HA failure status.....	61
Viewing vCMP Statistics	63
Overview: Viewing vCMP statistics.....	63
Viewing virtual disk statistics.....	63
Viewing vCMP guest information.....	63
Viewing current vCMP guest statistics.....	64
Viewing srTCM policier statistics for vCMP guests.....	64
Viewing statistics for physical disk usage	64
Viewing historical statistics about vCMP.....	65
Sample vCMP Statistics reports.....	66
Understanding Clusters	69
Overview: Managing a vCMP cluster.....	69
Viewing cluster properties.....	69
Cluster properties.....	69
Viewing cluster member properties.....	69
Cluster member properties.....	70
Enabling and disabling cluster members.....	70
Best Practices	71
vCMP best practices.....	71
Calculation for Maximum Core Allocation	73
Calculation for determining maximum core allocation.....	73
Additional Tasks for Isolated Guests in Appliance Mode	75
Additional tasks for isolated guests in Appliance mode.....	75
Preparing an isolated guest for Appliance mode.....	75
Enabling Appliance mode on an isolated guest.....	75
Deploying Route Domains within a vCMP Guest	77
Overview: Deploying Route Domains within a vCMP Guest.....	77
Prerequisite configuration tasks.....	78
About VLAN and BIG-IP address configuration.....	78
Illustration of VLAN and BIG-IP address configuration.....	78
Tasks for the host administrator.....	79
Creating customer VLANs on the vCMP host.....	80
Assigning VLANs to the vCMP guest.....	81
Tasks for the guest administrator.....	81
Creating an administrative partition for each customer.....	81
About moving host-based VLANs to a customer partition.....	82
Creating a route domain for each administrative partition.....	83
Creating an empty traffic group for each customer.....	85
Assigning a traffic group to each administrative partition.....	85

Table of Contents

Tasks for each customer administrator.....	86
Creating floating self IP addresses.....	86
Creating a pool.....	87
Creating a virtual server.....	87
Modifying a virtual IP address.....	88
Implementation results.....	89
Legal Notices.....	91
Legal notices.....	91

Introduction to the vCMP System

What is vCMP?

Virtual Clustered Multiprocessing[™] (vCMP[®]) is a feature of the BIG-IP[®] system that allows you to provision and manage multiple, hosted instances of the BIG-IP software on a single hardware platform. A vCMP hypervisor allocates a dedicated amount of CPU, memory, and storage to each BIG-IP instance. As a vCMP system administrator, you can create BIG-IP instances and then delegate the management of the BIG-IP software within each instance to individual administrators.

A key part of the vCMP system is its built-in flexible resource allocation feature. With *flexible resource allocation*, you can instruct the hypervisor to allocate a different amount of resource, in the form of *cores*, to each BIG-IP instance, according to the particular needs of that instance. Each *core* that the hypervisor allocates contains a fixed portion of system CPU and memory.

Furthermore, whenever you add blades to the VIPRION[®] cluster, properly-configured BIG-IP instances can take advantage of those additional CPU and memory resources without traffic interruption.

At a high level, the vCMP system includes two main components:

vCMP host

The *vCMP host* is the system-wide hypervisor that makes it possible for you to create and view BIG-IP instances, known as *guests*. Through the vCMP host, you can also perform tasks such as creating trunks and VLANs, and managing guest properties. For each guest, the vCMP host allocates system resources, such as CPU and memory, according to the particular resource needs of the guest.

vCMP guests

A *vCMP guest* is an instance of the BIG-IP software that you create on the vCMP system for the purpose of provisioning one or more BIG-IP[®] modules to process application traffic. A guest consists of a TMOS[®] instance, plus one or more BIG-IP modules. Each guest has its own share of hardware resources that the vCMP host allocates to the guest, as well as its own management IP addresses, self IP addresses, virtual servers, and so on. In this way, each guest effectively functions as its own multi-blade VIPRION[®] cluster, configured to receive and process application traffic with no knowledge of other guests on the system. Furthermore, each guest can use TMOS[®] features such as route domains and administrative partitions to create its own multi-tenant configuration. Each guest requires its own guest administrator to provision, configure, and manage BIG-IP modules within the guest. The maximum number of guests that a fully-populated chassis can support varies by chassis and blade platform.

This illustration shows a basic vCMP system with a host and four guests. Note that each guest has a different set of modules provisioned, depending on the guest's particular traffic requirements.




Figure 1: Example of a four-guest vCMP system

Other vCMP system components

In addition to the host and guests, the vCMP® system includes these components:

Virtual machine

A *virtual machine (VM)* is an instance of a guest that resides on a slot and functions as a member of the guest's virtual cluster. This illustration shows a system with guests, each with one or more VMs.




Figure 2: Guest VMs as cluster members

Virtual disk

A *virtual disk* is the portion of disk space on a slot that the system allocates to a guest VM. A virtual disk image is typically a 100 gigabyte sparse file. For example, if a guest spans three slots, the system creates three virtual disks for that guest, one for each blade on which the guest is provisioned. Each virtual disk is implemented as an image file with an `.img` extension, such as `guest_A.img`.

Core

A *core* is a portion of a blade's CPU and memory that the vCMP host allocates to a guest. The amount of CPU and memory that a core provides varies by blade platform.

Supported BIG-IP system versions

On a vCMP[®] system, the host and guests can generally run different combinations of BIG-IP[®] software versions simultaneously. With this type of version support, you can run multiple versions of BIG-IP software for testing, migration staging, or environment consolidation.

The exact combination of host and guest BIG-IP versions that F5 Networks[®] supports varies by blade platform. For details, see the vCMP host and supported guest version matrix on the AskF5 Knowledge Base at <http://support.f5.com>.

BIG-IP license considerations for vCMP

The BIG-IP[®] system license authorizes you to provision the vCMP[®] feature and create guests with one or more BIG-IP system modules provisioned. Note the following considerations:

- Each guest inherits the license of the vCMP host.
- The host license must include all BIG-IP modules that are to be provisioned across all guest instances. Examples of BIG-IP modules are BIG-IP Local Traffic Manager[™] and BIG-IP Global Traffic Manager[™].
- The license allows you to deploy the maximum number of guests that the specific blade platform allows.
- If the license includes the Appliance mode feature, you cannot enable Appliance mode for individual guests; when licensed, Appliance mode applies to all guests and cannot be disabled.

You activate the BIG-IP system license when you initially set up the vCMP host.

About vCMP provisioning

To enable the vCMP[®] feature, you perform two levels of provisioning.

First, you provision the vCMP feature as a whole. When you do this, the BIG-IP[®] system, by default, dedicates most of the disk space to running the vCMP feature, and in the process, creates the host portion of the vCMP system.

Second, once you have configured the host to create the guests, each guest administrator logs in to the relevant guest and provisions the required BIG-IP modules. In this way, each guest can run a different combination of modules. For example, one guest can run BIG-IP[®] Local Traffic Manager[™] (LTM[®]) only, while a second guest can run LTM[®] and BIG-IP ASM[™].

Important: *Once you provision the vCMP feature, you cannot provision any BIG-IP modules, such as BIG-IP LTM, on the vCMP host. Moreover, if any BIG-IP modules are already provisioned on the system before you provision the vCMP feature, those modules are de-provisioned when you provision the vCMP feature. This, in turn, interrupts any application traffic currently being processed.*

Network isolation

The vCMP[®] system separates the data plane network from the management network. That is, the host operates with the hardware switch fabric to control the guest data plane traffic. Each slot in the chassis has its own network interface for data plane traffic that is separate from the management network. This separation of the data plane network from the management network provides true multi-tenancy by ensuring that traffic for a guest remains separate from all other guest traffic on the system.

The following illustration shows the separation of the data plane network from the management network.




Figure 3: Isolation of the data plane network from the management network

System administration overview

Administering a vCMP[®] system requires two distinct types of administrators: a vCMP host administrator who creates guests and allocates resources to those guests, and a vCMP guest administrator who provisions and configures BIG-IP modules within a specific guest.

At a minimum, these tasks must be performed on the vCMP host, by a host administrator:

- Provision the vCMP feature
- Create vCMP guests, including allocating system resources to each guest
- Create and manage VLANs
- Create and manage trunks
- Manage interfaces
- Configure access control to the host by other host administrators, through user accounts and roles, partition access, and so on

These tasks are performed on a vCMP guest by a guest administrator:

- Provision BIG-IP modules
- Create self IP addresses and associate them with host VLANs
- Create and manage features within BIG-IP modules, such as virtual servers, pools, policies, and so on
- Configure device service clustering (DSC)
- Configure access control to the guest by other guest administrators, through user accounts and roles, partition access, and so on

Important: vCMP host administration tasks can only be performed on the vCMP host and not from within a guest. This prevents a guest administrator from accessing either the host or other guests on the system, thereby ensuring the separation of administrative tasks across the vCMP deployment.

After you initially set up the vCMP host, you will have a standalone, multi-tenant vCMP system with some number of guests defined. A guest administrator will then be ready to provision and configure the BIG-IP modules within a guest to process application traffic. Optionally, if the host administrator has set up a second chassis with equivalent guests, a guest administrator can configure high availability for any two equivalent guests.

Guest access to the management network

As a vCMP host administrator, you can configure each vCMP® guest to be either bridged to or isolated from the management network, or to be isolated from the management network but remain accessible by way of the host-only interface.

Important: F5 Networks recommends that you configure all vCMP guests to be bridged to the management network, unless you have a specific business or security requirement that requires guests to be isolated from the management network.

About bridged guests

When you create a vCMP® guest, you can specify that the guest is a bridged guest. A *bridged* guest is one that is connected to the management network. This is the default network state for a vCMP guest. This network state bridges the guest's virtual management interface to the physical management interface of the blade on which the guest virtual machine (VM) is running.

You typically log in to a bridged guest using its cluster management IP address, and by default, guest administrators with the relevant permissions on their user accounts have access to the `bash` shell, the BIG-IP® Configuration utility, and the Traffic Management Shell (`tmsh`). However, if per-guest Appliance mode is enabled on the guest, administrators have access to the BIG-IP Configuration utility and `tmsh` only.

Although the guest and the host share the host's Ethernet interface, the guest appears as a separate device on the local network, with its own MAC address and IP address.

Note that changing the network state of a guest from isolated to bridged causes the vCMP host to dynamically add the guest's management interface to the bridged management network. This immediately connects all of the guest's VMs to the physical management network.

Important: If you want to easily make TCP connections (for SSH, HTTP, and so on) from either the host or the external network to the guest, or from the guest to the host or external network, you can configure a guest's management port to be on the same IP network as the host's management port, with a gateway identical to the host's management gateway. However, you should carefully consider the security implications of doing so.

About isolated guests

When you create a vCMP® guest, you can specify that the guest is an isolated guest. Unlike a bridged guest, an *isolated* guest is disconnected from the management network. As such, the guest cannot communicate with other guests on the system. Also, because an isolated guest has no management IP address for administrators to use to access the guest, the host administrator, after creating the guest, must use the `vconsole` utility to log in to the guest and create a self IP address that guest administrators can then use to access the guest.

About Appliance mode

Appliance mode is a BIG-IP system feature that adds a layer of security in two ways:

- By preventing administrators from using the `root` user account.
- By granting administrators access to the Traffic Management Shell (`tmsh`) only, instead of the advanced (`bash`) shell.

You can implement Appliance mode in one of two ways:

System-wide through the BIG-IP license

You can implement Appliance mode on a system-wide basis through the BIG-IP® system license. However, this solution might not be ideal for a vCMP® system. When a vCMP system is licensed for Appliance mode, administrators for all guests on the system are subject to Appliance mode restrictions. Also, you cannot disable the Appliance mode feature when it is included in the BIG-IP system license.

On a per-guest basis

Instead of licensing the system for Appliance mode, you can enable or disable the appliance mode feature for each guest individually. By default, per-guest Appliance mode is disabled when you create the guest. After Appliance mode is enabled, you can disable or re-enable this feature on a guest at any time.

***Note:** If the license for the BIG-IP system includes Appliance mode, the system ignores the per-guest Appliance mode feature and permanently enforces Appliance mode for the vCMP host and all guests on the system.*

User access restrictions with Appliance mode

When you enable Appliance mode on a guest, the system enhances security by preventing administrators from accessing the root-level advanced shell (`bash`).

For bridged guests

For a bridged guest with Appliance mode enabled, administrators can access the guest through the guest's management IP address. Administrators for a bridged guest can manage the guest using the BIG-IP® Configuration utility and `tmsh`.

For isolated guests

For an isolated guest with Appliance mode enabled, administrators must access a guest through one of the guest's self IP addresses, configured with appropriate port lockdown values. Administrators for an isolated guest can manage the guest using the BIG-IP Configuration utility and `tmsh`.

***Important:** When you enable Appliance mode on a guest, any accounts with advanced shell access automatically lose that permission and the permission reverts to `tmsh`. If you disable Appliance mode later, you can re-assign advanced shell access to those accounts.*

BIG-IP version restrictions with Appliance mode

If you want to use the BIG-IP® version 11.5 Appliance mode feature on a guest, both the host and the guest must run BIG-IP version 11.5 or later.

Warning: *If you enable Appliance mode on a guest, and a previous version of the BIG-IP software is installed in another boot location, a guest administrator with an Administrator user role can boot to the previous version and obtain advanced shell access.*

Additional Network Considerations

Network separation of Layer 2 and Layer 3 objects

On a vCMP system, you typically configure BIG-IP® Layer 2 objects, such as trunks and VLANs, on the vCMP host and then selectively decide which of these objects you want each guest to inherit. To ensure that each guest's data plane traffic is securely isolated from other guests, the host administrator usually creates a separate VLAN for each guest to use. Other objects such as self IP addresses, virtual servers, pools, and profiles are configured on the guest by each guest administrator. With this separation of Layer 2 from Layer 3 objects, application traffic is targeted directly to the relevant guest, further allowing each guest to function as a fully-independent BIG-IP® device.

The following illustration shows the separation of Layer 2 objects from higher-layer objects on the vCMP system:




Figure 4: Isolation of network objects on the vCMP system

About the VLAN publishing strategy

For both host and guest administrators, it is important to understand certain concepts about VLAN configuration on a vCMP system:

- VLAN subscription from host to guest
- System behavior when a host and a guest VLAN have duplicate names or tags

Overview of VLAN subscription

As a vCMP® host administrator, when you create or modify a guest, you typically *publish* one or more host-based VLANs to the guest. When you publish a host-based VLAN to a guest, you are granting a subscription to the guest for use of that VLAN configuration, with the VLAN's underlying Layer 2 resources.

Additional Network Considerations

When you publish a VLAN to a guest, if there is no existing VLAN within the guest with the same name or tag as the host-based VLAN, the vCMP system automatically creates, on the guest, a configuration for the published VLAN.

If you modify a guest's properties to remove a VLAN publication from a guest, you are removing the guest's subscription to that host-based VLAN. However, the actual VLAN configuration that the host created within the guest during initial VLAN publication to the guest remains there for the guest to use. In this case, any changes that a host administrator might make to that VLAN are not propagated to the guest.

In general, VLANs that appear within a guest can be either host-based VLANs currently published to the guest, host-based VLANs that were once but are no longer published to the guest, or VLANs that the guest administrator manually created within the guest.

This example shows the effect of publishing a host-based VLAN to, and then deleting the VLAN publication from, a guest that initially had no VLANs.

```
# Within guest G1, show that the guest has no VLANs configured:
[root@G1:/S1-green-P:Active:Standalone] config # tmsh list net vlan

# From the host, publish VLAN v1024 to guest G1:
[root@host_210:/S1-green-P:Active:Standalone] config # tmsh modify vcmp guest G1 vlans add
{ v1024 }

# Within guest G1, list all VLANs:
[root@G1:/S1-green-P:Active:Standalone] config # tmsh list net vlan

net vlan v1024 {
if-index 96
tag 1024
}

# On the host, delete the host-based VLAN publication from guest G1:
[root@host_210:/S1-green-P:Active:Standalone] config # tmsh modify vcmp guest G1 vlans del
{ v1024 }

# Notice that the host-based VLAN still exists within the guest:
[root@G1:/S1-green-P:Active:Standalone] config # tmsh list net vlan

vlan v1024 {
if-index 96
tag 1024
}
```

Note: For any host-based VLAN published to a guest, all properties of the VLAN, such as name, VLAN ID, and so on, must be managed on the host. An exception is the maximum transmission unit (MTU) property. If a guest administrator wants to change the default VLAN MTU value for a guest, the administrator can (and should) log into the guest and modify the MTU value for the host-based VLAN from within the guest. Changing the MTU size when logged into the host has no effect on the guest's ability to process traffic or manage routing.

About VLANs with identical tags and different names

Sometimes a host administrator might publish a VLAN to a guest, but the guest administrator has already created, or later creates, a VLAN with a different name but the same VLAN tag. In this case, the guest VLAN always overrides the host VLAN. The VLAN can still exist on the host (for other guests to subscribe to), but it is the guest VLAN that is used.

Whenever host and guest VLANs have different names but the same tags, traffic flows successfully across the host from the guest because the VLAN tag alignment is correct. That is, when the tags match, the underlying Layer 2 infrastructure of the VLANs matches, thereby enabling the host to reach the guest.

The example here shows the `tmsh` command sequence for creating two separate VLANs with different names and the same tag, and the resulting successful traffic flow.

```
# On the host, create a VLAN with a unique name but with a tag matching that of a guest
VLAN VLAN_A:

[root@host_210:/S1-green-P:Active:Standalone] config # tmsh create net vlan VLAN_B tag 1000

# On the host, publish the host VLAN to the guest:

[root@host_210:/S1-green-P:Active:Standalone] config # tmsh modify vcmp guest guest1 vlans
add { VLAN_B }

# Within the guest, show that the guest still has its own VLAN only, and not the VLAN
published from the host:

[root@G1:/S1-green-P:Active:Standalone] config # tmsh list net vlan all

net vlan VLAN_A {
    if-index 192
    tag 1000
}

# On the guest, create a self IP address for VLAN_A:

[root@G1:/S1-green-P:Active:Standalone] config # tmsh create net self 10.1.1.1/24 vlan
VLAN_A

# On the host, delete the self IP address on VLAN_A (this VLAN also exists on the guest)
and re-create the self IP address on VLAN_B (this VLAN has the same tag as VLAN_A):

[root@host_210:/S1-green-P:Active:Standalone] config # tmsh delete net self 10.1.1.2/24
[root@host_210:/S1-green-P:Active:Standalone] config # tmsh create net self 10.1.1.2/24
vlan VLAN_B

# From the host, open a connection to the guest, and notice that because the two VLANs have
the same tags, the connection succeeds:

[root@host_210:/S1-green-P:Active:Standalone] config # ping -c2 10.1.1.1

PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=255 time=3.35 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=255 time=0.989 ms

--- 10.1.1.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 0.989/2.170/3.352/1.182 ms
```

About VLANs with identical names and different tags

Sometimes a host administrator might publish a VLAN to a guest, but the guest administrator has already created, or later creates, a VLAN with the same name but with a different VLAN tag. In this case, the guest VLAN always overrides the host VLAN. The VLAN can still exist on the host (for other guests to subscribe to), but it is the guest VLAN that is used.

Whenever host and guest VLANs have the same names but different tags, traffic cannot flow between the identically-named VLANs at Layer 2. That is, when the tags do not match, the underlying Layer 2 infrastructure of the VLANs does not match, thereby preventing the host from reaching the guest.

Additional Network Considerations

The example here shows the `tmsh` command sequence for creating two separate VLANs with the same names and different tags, and the resulting traffic flow issue.

```
# While logged into the guest, create a VLAN:
[root@G1:/S1-green-P:Active:Standalone] config # tmsh create net vlan VLAN_A tag 1000

# Show that no VLANs exist on the host:
[root@host_210:/S1-green-P:Active:Standalone] config # tmsh list net vlan all
[root@host_210:/S1-green-P:Active:Standalone] config #

# On the host, create a VLAN with the same name as the guest VLAN but with a unique tag on
the host:
[root@host_210:/S1-green-P:Active:Standalone] config # tmsh create net vlan VLAN_A tag 1001

# Publish the host VLAN to the guest:
[root@host_210:/S1-green-P:Active:Standalone] config # tmsh modify vcmp guest guest1 vlans
add { VLAN_A }

# Within the guest, show that the guest still has its own VLAN only, and not the VLAN
published from the host:
[root@G1:/S1-green-P:Active:Standalone] config # tmsh list net vlan all

net vlan VLAN_A {
  if-index 192
  tag 1000
}

# Within the guest, create a self IP address for the VLAN:
[root@G1:/S1-green-P:Active:Standalone] config # tmsh create net self 10.1.1.1/24 vlan
VLAN_A

# On the host, create a self IP address for the identically-named VLAN:
[root@host_210:/S1-green-P:Active:Standalone] config # tmsh create net self 10.1.1.2/24
vlan VLAN_A

# From the host, open a connection to the guest, and notice that because the two VLANs have
different tags, the connection fails:
[root@host_210:/S1-green-P:Active:Standalone] config # ping -c2 10.1.1.1

PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
From 10.1.1.2 icmp_seq=1 Destination Host Unreachable
From 10.1.1.2 icmp_seq=2 Destination Host Unreachable

--- 10.1.1.1 ping statistics ---
2 packets transmitted, 0 received, +2 errors, 100% packet loss, time 3000ms
pipe 2
```

Solution for tag discrepancy between host and guest VLANs

When a host-based VLAN and a guest-created VLAN have identical names but different VLAN tags, traffic flow at Layer 2 is impeded between host and guest. Fortunately, you can resolve this issue by performing these tasks, in the sequence shown:

- Within the guest, delete the relevant VLAN from within the guest.
- On the host, remove the VLAN publication from the guest.
- On the host, modify the tag of the host-based VLAN.
- On the host, publish the VLAN to the guest.

- Within the guest, view the VLAN from within the guest.

Deleting the VLAN within the guest

Perform this task when you want to delete a VLAN from within a vCMP guest.

Important: To perform this task, you must be logged in to the relevant vCMP guest.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the Name column, locate the name of the VLAN for which you want to change the partition, and to the left of the name, select the check box and click **Delete**.
The system prompts you to confirm the delete action.
3. Click **Delete**.

After performing this task, you will no longer see the VLAN name in the list of VLANs on the guest.

Removing the VLAN publication on the guest

You perform this task when you want to remove a VLAN subscription for a particular guest.

Important: To perform this task, you must be logged in to the vCMP host.

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

Modifying the tag of the host-based VLAN

Perform this task to change a VLAN tag on a vCMP host to ensure that the tag matches that of a VLAN on a guest.

Important: To perform this task, you must be logged in to the vCMP host.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the Name column, click the relevant VLAN name.
This displays the properties of the VLAN.
3. In the **Tag** field, type the same tag that was assigned to the VLAN you previously deleted.
4. If the host and guest VLANs have an optional customer tag, type the same customer tag that was assigned to the VLAN you previously deleted.
5. Click **Update**.

Publishing the VLAN to the guest

You perform this task when you want to publish a host-based VLAN to a particular guest.

Important: To perform this task, you must be logged in to the vCMP host.

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

After performing this task, the guest can use the selected host-based VLAN.

Viewing the new VLAN within the guest

Perform this task to verify that the VLAN that the host published to a guest appears on the guest, with the correct tag.

Important: To perform this task, you must be logged in to the relevant vCMP guest.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the Name column, click the name of the VLAN that the host published to the guest.
3. In the **Tag** field, verify that the correct tag is shown.
4. Click **Cancel**.

After you perform this task, you can see that the VLAN that the host published to the guest has appeared on the guest, with the correct tag.

About the VLAN MTU setting

One of the properties of a VLAN on a vCMP system is the maximum transmission unit (MTU). When a vCMP guest subscribes to a host-based VLAN, the MTU value on the vCMP host is independent of the MTU value on the guest.

If you need to change the VLAN MTU value for a specific guest, you should always log into the guest to make the change. Changing the MTU size on the host has no effect on the guest's ability to process traffic and manage routing.

Note that for any host-based VLAN, the MTU property is the only VLAN property that you can change when you're logged in to a guest. All other VLAN properties, such as name, VLAN ID, and so on, must be managed on the host.

Note: To avoid any traffic interruptions, make sure the neighboring network devices support the guest's VLAN MTU size.

Interface assignment for vCMP guests

The virtualized nature of vCMP® guests abstracts many underlying hardware dependencies, which means that there is no direct relationship between guest interfaces and the physical interfaces assigned to VLANs on the vCMP host.

Rather than configuring any interfaces on a guest, a guest administrator simply creates a self IP address within the guest, specifying one of the VLANs that the host administrator previously configured on the host and assigned to the guest during guest creation.

As host administrator, if you want to limit the guest to using specific physical interfaces, you simply change the physical interface assignments on the VLANs that you assign to that guest.

Management IP addresses for bridged guests

When a system administrator initially configured the VIPRION system, the administrator specified a primary cluster management IP address for the system as a whole, as well as a separate management IP address for each slot in the VIPRION cluster. On a vCMP system, because each guest functions like an independent VIPRION cluster, a vCMP host or guest administrator assigns a similar set of IP addresses for each guest:

A cluster IP address

This is the unique IP address that a host administrator assigns to a guest during guest creation. The cluster IP address is the management IP address that a guest administrator uses to log in to a guest to provision, configure, and manage BIG-IP® modules. This IP address is required for each guest.

One or more cluster member IP addresses

These are unique IP addresses that a guest administrator assigns to the virtual machines (VMs) in the guest's cluster, for high-availability purposes. For example, if a guest on a four-slot system is configured to run on four slots, then the guest administrator must create an IP address for each of those four slots. These addresses are management addresses, and although optional for a standalone system, these addresses are required for a device service clustering (DSC®) configuration. In this case, a second set of unique cluster member IP addresses must be configured on the peer system.

These IP addresses are the addresses that the guest administrator will specify when configuring failover for each guest that is a member of a Sync-Failover device group.

As an example, suppose you have a pair of VIPRION 2400 chassis, where the two guests on one chassis also reside on the other chassis to form a redundant configuration. In this case, as host administrator, you must assign a total of four cluster IP addresses (one per guest for four guests).

If each guest spans four slots, then each guest administrator must then assign four cluster member IP addresses per guest per chassis, for a total of eight. The result is a total of 20 unique vCMP-related management IP addresses for the full redundant pair of chassis containing two guests per chassis (four cluster IP addresses and 16 cluster member IP addresses).

Important: *F5 Networks recommends that you assign a cluster member IP address to every slot in the guest's cluster, even for slots not assigned to the guest. This simplifies the task of assigning slots to a guest later if you need to do so.*

Flexible Resource Allocation

What is flexible resource allocation?

Flexible resource allocation is a built-in vCMP® feature that allows vCMP host administrators to optimize the use of available system resources. Flexible resource allocation gives you the ability to configure the vCMP host to allocate a different amount of CPU and memory to each guest through core allocation, based on the needs of the specific BIG-IP® modules provisioned within a guest.

When you create each guest, you specify the number of logical cores that you want the host to allocate to the guest, and you identify the specific slots that you want the host to assign to the guest. Configuring these settings determines the total amount of CPU and memory that the host allocates to the guest. With flexible allocation, you can customize CPU and memory allocation in granular ways that meet the specific resource needs of each individual guest.

Resource allocation planning

When you create a vCMP® guest, you must decide the amount of dedicated resource, in the form of CPU and memory, that you want the vCMP host to allocate to the guest. You can allocate a different amount of resources to each guest on the system.

Prerequisite hardware considerations

Blade models vary in terms of how many cores the blade provides and how much memory each core contains. Also variable is the maximum number of guests that each blade model supports. For example, a single B2100 blade provides eight cores and approximately 3 gigabytes (GB) of memory per core, and supports a maximum of four guests.

Before you can determine the number of cores to allocate to a guest and the number of slots to assign to a guest, you should understand:

- The total number of cores that the blade model provides
- The amount of memory that each blade model provides
- The maximum number of guests that the blade model supports

By understanding these metrics, you ensure that the total amount of resource you allocate to guests is aligned with the amount of resource that your blade model supports.

For specific information on the resources that each blade model provides, see the vCMP® guest memory/CPU core allocation matrix on the AskF5™ Knowledge Base at <http://support.f5.com>.

Understanding guest resource requirements

Before you create vCMP® guests and allocate system resources to them, you need to determine the specific CPU and memory needs of each guest. You can then decide how many cores to allocate and slots to assign to a guest, factoring in the resource capacity of your blade model.

To determine the CPU and memory resource needs, you must know:

- The number of guests you need to create
- The specific BIG-IP® modules you need to provision within each guest
- The combined memory requirement of all BIG-IP modules within each guest

About core allocation for a guest

When you create a guest on the vCMP® system, you must specify the total number of cores that you want the host to allocate to the guest based on the guest's total resource needs. Each core provides some amount of virtual CPU (vCPU) and a fixed amount of memory.

Memory allocation

You should specify enough cores to satisfy the combined memory requirements of all BIG-IP® modules that you provision within the guest. When you deploy the guest, the host allocates this number of cores to every slot on which the guest runs, regardless of the number of slots you have assigned to the guest.

It is important to understand that the total amount of memory available to a guest is only as much as the host has allocated to each slot. If you instruct the host to allocate a total of two cores per slot for the guest (for example, 6 GB of memory depending on blade model) and you configure the guest to run on four slots, the host does not aggregate the 6 GB of memory on each slot to provide 24 GB of memory for the guest. Instead, the guest still has a total of 6 GB of memory available. This is because blades in a chassis operate as a cluster of independent devices, which ensures that if the number of blades for the guest is reduced for any reason, the remaining blades still have the required memory available to process the guest traffic.

If a blade suddenly becomes unavailable, the total traffic processing resource for guests on that blade is reduced and the host must redistribute the load on that slot to the remaining assigned slots. This increases the number of connections that each remaining blade must process and therefore the amount of memory actually used per blade. The increased percentage of memory used per blade compared to the amount of memory allocated could cause swapping and degraded performance. You can prevent this result by making sure you allocate enough cores to the guest, per slot, when you create the guest.

vCPU allocation

Unlike memory allocation, the total allocation of vCPUs for a guest is the sum total of vCPUs allocated to the guest on all blades in the chassis. This means that if you remove a blade from the chassis, total vCPU allocation for a guest decreases by the number of vCPUs allocated to the guest on that blade.

Important: Sometimes when you attempt to set a guest to the Deployed state, the action fails and the system displays this message:

```
Could not allocate vCMP guest guest-name because slot 1: fragmented resources
```

This message is caused by core fragmentation. Fragmentation occurs when one or more single-CPU guests consume only one of the two hyperthreads that make up a core. Guests that consume a single hyperthread from a core leave the core's other hyperthread available to other guests, but not in a contiguous way that most other guests can use. That is, most guests are multi-CPU guests and therefore require that each core allocated be a full core consisting of two hyperthreads. The vCMP system cannot allocate two hyperthreads by combining available non-contiguous (that is, fragmented) hyperthreads to a guest.

About core fragmentation when deploying a guest

Sometimes when you attempt to set a guest to the Deployed state, the action fails and the system displays a message similar to this:

```
Could not allocate vCMP guest guest-name because slot 1: fragmented resources
```

This message is caused by core fragmentation. *Fragmentation* occurs when a single-CPU guest consumes only one of the two hyperthreads that make up a core, and you have deployed more than one single-CPU

guest. Guests that consume a single hyperthread of a core each leave the remaining hyperthread available to other guests, but not in a contiguous way that other guests can use. That is, most guests are multi-CPU guests and therefore require that each core allocated be a full core consisting of two hyperthreads. The vCMP system cannot allocate a core to a guest by combining two non-contiguous (that is, fragmented) hyperthreads.

The best way to fix the problem is to set one of the single-CPU guests back to the Configured state and then re-deploy it. In this case, the vCMP system will most likely allocate a hyperthread from a different core to that guest, freeing up a full, two-hyperthread core for the guest that triggered the error message.

Formula for host memory allocation to a guest

You can use a formula to confirm that the cores you plan to allocate to a specific guest are sufficient, given the guest's total memory requirements:

```
(total_GB_memory_per_blade - 3 GB) x (cores_per_slot_per_guest / total_cores_per_blade) =
amount of guest memory allocation from host
```

Important: For metrics on memory and CPU support per blade model, refer to the vCMP® guest memory/CPU allocation matrix at <http://support.f5.com>.

The variables in this formula are defined as follows:

total_GB_memory_per_blade

The total amount of memory in gigabytes that your specific blade model provides (for all guests combined).

cores_per_slot_per_guest

The estimated number of cores needed to provide the total amount of memory that the guest requires.

total_cores_per_blade

The total number of cores that your specific blade model provides (for all guests combined).

For example, if you have a VIPRION® 2150 blade, which provides approximately 32 GB memory through a maximum of eight cores, and you estimate that the guest will need two cores to satisfy the guest's total memory requirement of 8 GB, the formula looks as follows:

```
(32 GB - 3 GB) x (2 cores / 8 cores) = 7.25 GB memory that the host will allocate to the
guest per slot
```

In this case, the formula shows that two cores will not provide sufficient memory for the guest. If you specify four cores per slot instead of two, the formula shows that the guest will have sufficient memory:

```
(32 GB - 3 GB) x (4 cores / 8 cores) = 14.5 GB memory that the host will allocate to the
guest per slot
```

Note that except for single-core guests, the host always allocates cores in increments of two. For example, for B2150 blade models, the host allocates cores in increments of 2, 4, and 8.

Once you use this formula for each of the guests you plan to create on a slot, you can create your guests so that the combined memory allocation for all guests on a slot does not exceed the total amount of memory that the blade model provides.

About slot assignment for a guest

On the vCMP® system, the host assigns some number of slots to each guest based on information you provide when you initially create the guest. The key information that you provide for slot assignment is the maximum and minimum number of slots that a host can allocate to the guest, as well as the specific

slots on which the guest is allowed to run. With this information, the host determines the number of slots and the specific slot numbers to assign to each guest.

As a best practice, you should configure every guest so that the guest can span all slots in the cluster whenever possible. The more slots that the host can assign to a guest, the lighter the load is on each blade (that is, the fewer the number of connections that each blade must process for that guest).

***Note:** In device service clustering (DSC[®]) configurations where mirroring is enabled, all guests in the device group must have the same core allocation and module provisioning, and the guests must match with respect to number of slots and the exact slot numbers. Also, when mirroring is enabled, all guests in the device group must run on the same blade model and chassis model.*

About single-core guests

On platforms with hard drives, the vCMP[®] host always allocates cores on a slot for a guest in increments of two cores. In the case of blades with solid-state drives, however, the host can allocate a single core to a guest, but only for a guest that requires one core only; the host does not allocate any other odd number of cores per slot for a guest (such as three, five, or seven cores).

Because a single-core guest has a relatively small amount of CPU and memory allocated to it, F5 Networks supports only these products or product combinations for a single-core guest:

- BIG-IP[®] Local Traffic Manager[™] (LTM[®]) only
- BIG-IP[®] Local Traffic Manager[™] (LTM[®]) and BIG-IP[®] DNS (previously Global Traffic Manager) only
- BIG-IP[®] DNS (previously Global Traffic Manager) standalone only

Scalability considerations

When managing a guest's slot assignment, or when removing a blade from a slot assigned to a guest, there are a few key concepts to consider.

About initial slot assignment

When you create a vCMP[®] guest, the number of slots that you initially allow the guest to run on determines the maximum total resource allocation possible for that guest, even if you add blades later. For example, in a four-slot VIPRION[®] chassis that contains two blades, if you allow a guest to run on two slots only and you later add a third blade, the guest continues to run on two slots and does not automatically expand to acquire additional resource from the third blade. However, if you initially allow the guest to run on all slots in the cluster, the guest will initially run on the two existing blades but will expand to run on the third slot, acquiring additional traffic processing capacity, if you add another blade.

Because each connection causes some amount of memory use, the fewer the connections that the blade is processing, the lower the percentage of memory that is used on the blade compared to the total amount of memory allocated on that slot for the guest. Configuring each guest to span as many slots as possible reduces the chance that memory use will exceed the available memory on a blade when that blade must suddenly process additional connections.

If you do not follow the best practice of instructing the host to assign as many slots as possible for a guest, you should at least allow the guest to run on enough slots to account for an increase in load per blade if the number of blades is reduced for any reason.

In general, F5 Networks strongly recommends that when you create a guest, you assign the maximum number of available slots to the guest to ensure that as few additional connections as possible are redistributed to each blade, therefore resulting in as little increase in memory use on each blade as possible.

About changing slot assignments

At any time, you can intentionally increase or decrease the number of slots a guest runs on explicitly by re-configuring the number of slots that you initially assigned to the guest. Note that you can do this while a guest is processing traffic, to either increase the guest's resource allocation or to reclaim host resources.

When you increase the number of slots that a guest is assigned to, the host attempts to assign the guest to those additional slots. The host first chooses those slots with the greatest number of available cores. The change is accepted as long as the guest is still assigned to at least as many slots as dictated by its **Minimum Number of Slots** value. If the additional number of slots specified is not currently available, the host waits until those additional slots become available and then assigns the guest to these slots until the guest is assigned to the desired total number of slots. If the guest is currently in a deployed state, VMs are automatically created on the additional slots.

When you decrease the number of slots that a guest is assigned to, the host removes the guest from the most populated slots until the guest is assigned to the correct number of slots. The guest's VMs on the removed slots are deleted, although the virtual disks remain on those slots for reassignment later to another guest. Note that the number of slots that you assign to a guest can never be less than the minimum number of slots configured for that guest.

Effect of blade removal on a guest

If a blade suddenly becomes unavailable, the total traffic processing resource for guests on that blade is reduced and the host must redistribute the load on that slot to the remaining assigned slots. This increases the number of connections that each remaining blade must process and therefore the amount of memory used per blade. Fortunately, when you instruct the host to allocate some amount of memory to the guest, the host allocates that amount of memory to every slot in the guest's cluster.

Be aware, however, that if a blade goes offline so that the number of connections per blade increases, the increased percentage of memory used per blade compared to the amount of memory allocated could cause swapping and degraded performance. You can prevent this result by making sure you allocate enough cores to the guest, per slot, when you create the guest.

Example of blade removal and memory use

A blade going offline increases the amount of memory being used on the remaining blades. The following example helps to explain this concept.

Important: *The memory use values shown in these illustrations are for example purposes only and are not meant to represent typical values.*

Suppose you have a guest spanning four slots that process 1,000,000 connections combined, where each slot is processing a quarter of the connections to the guest. Notice that the host administrator has allocated 4 GB of memory to the guest, and there is a current memory use of 3 GB for every 250,000 connections.




Figure 5: All blades are functional with normal memory use per blade

Now suppose a blade goes offline. In this case, each remaining blade must now process a third of the connections (333,333), which might increase memory use per blade to 4.5 GB (for example).

The following illustration shows that when a blade goes offline, memory use can exceed the 4 GB available on each blade due to the increase in number of connections per blade:




Figure 6: An offline blade causes memory use to approach memory allocation per blade

As you can see in the example, the increase in percentage of memory use per blade could be avoided by allocating four cores per slot instead of two, so that 8 GB of memory is available per blade. This removes any risk of performance degradation when blade loss occurs.

Effect of blade re-insertion on a guest

When you remove a blade from the chassis, the host remembers which guests were allocated to that slot. If you then re-insert a blade into that slot, the host automatically allocates cores from that blade to the guests that were previously assigned to that slot.

Whenever the host assigns guests to a newly-inserted blade, those guests that are below their **Minimum Number of Slots** threshold are given priority; that is, the host assigns those guests to the slot before guests that are already assigned to at least as many slots as their **Minimum Number of Slots** value. Note that this is the only time when a guest is allowed to be assigned to fewer slots than specified by its **Minimum Number of Slots** value.

Network throughput for guests

To manage network throughput for a vCMP® guest, you should understand the throughput capacity of your blade type, as well as the throughput limit you want to apply to each guest:

Throughput capacity per blade

Each blade type on a VIPRION® system has a total throughput capacity, which defines the combined upper limit on throughput for guests on a blade. For example, on a B2100 blade with one single-slot guest, the guest can process up to 40Gbps (with ePVA enabled). If the single-slot guest needs to process more than 40Gbps, you can expand the guest to run on more slots.

Throughput limits per guest

Throughput requirements for a guest are typically lower than the throughput capacity of the blades on which the guest runs. Consequently, you can define a specific network throughput limit for each guest. When vCMP is provisioned on the system, you define a guest's throughput limit by logging in to the vCMP host and creating a rate shaping object known as a *Single Rate Three Color Marker (srTCM) Policer*. You then assign the policer to one or more guests when you create or modify the guests. It is important that the srTCM values that you assign to a guest do not exceed the combined throughput capacity of the blades pertaining to that guest.

SSL resource allocation

On certain models, you can control the way that the vCMP system allocates high-performance SSL hardware resources to vCMP guests. Specifically, you can configure one of three SSL modes for each guest:

Shared mode

This mode causes the guest to share its consumption of SSL hardware resource with other guests that are also in Shared mode. In this case, guests with the most need for SSL resources consume more of the total resource available. This is the default SSL mode.

Dedicated mode

This mode dedicates a fixed amount of SSL hardware resource to a guest. When you configure this option for a guest, the amount of resource that the system allocates to the guest is based on the guest's core allocation. In Dedicated mode, the guest is guaranteed a fixed amount of resource and this amount is not affected by the amount of resource that other guests consume.

None

This option prevents a guest from consuming any SSL hardware resources. This option also prevents the guest from consuming compression hardware resources.

Note: Regardless of the current guest state (Deployed, Provisioned, or Configured), you can change the SSL mode for a guest from Shared to Dedicated, or the reverse, at any time. However, if you want to change a guest to or from None mode, ensure that the guest is in the Configured state.

About compression resource allocation

On blade models that include compression hardware processors, the vCMP® host allocates an equal share of the hardware compression resource among all guests on the system, in a round robin fashion.

Additionally, on certain blade models, the vCMP host automatically disables the allocation of compression hardware resources to a guest whenever you also disable the allocation of SSL hardware resources to that guest.

Guest states and resource allocation

As a vCMP® host administrator, you can control when the system allocates or de-allocates system resources to a guest. You can do this at any time, by setting a guest to one of three states: Configured, Provisioned, or Deployed. These states affect resource allocation in these ways:

Configured

This is the initial (and default) state for a newly-created guest. In this state, the guest is not running, and no resources are allocated. If you change a guest from another state to the Configured state, the

vCMP host does not delete any virtual disks that were previously attached to that guest; instead, the guest's virtual disks persist on the system. The host does, however, automatically de-allocate other resources such as CPU and memory. When the guest is in the Configured state, you cannot configure the BIG-IP® modules that are licensed to run within the guest; instead, you must set the guest to the Deployed state to provision and configure the BIG-IP modules within the guest.

Provisioned

When you change a guest state from Configured to Provisioned, the vCMP host allocates system resources to the guest (CPU, memory, and any unallocated virtual disks). If the guest is new, the host creates new virtual disks for the guest and installs the selected ISO image on them. A guest does not run while in the Provisioned state. When you change a guest state from Deployed to Provisioned, the host shuts down the guest but retains its current resource allocation.

Deployed

When you change a guest to the Deployed state, the vCMP host activates the guest virtual machines (VMs), and the guest administrator can then provision and configure the BIG-IP modules within the guest. For a guest in this state, the vCMP host starts and maintains a VM on each slot for which the guest has resources allocated. If you are a host administrator and you reconfigure the properties of a guest after its initial deployment, the host immediately propagates the changes to all of the guest VMs and also propagates the list of allowed VLANs.

Deployment Examples

Example: A single-slot LTM guest on a standalone system

The simplest example of the deployment of a vCMP® system is a standalone system configured with one guest that is provisioned to run BIG-IP® Local Traffic Manager™ (LTM®) on a single slot in the VIPRION® cluster.

The following illustration depicts a single-slot, two-core LTM guest on a standalone VIPRION chassis.




Figure 7: Single-slot, two-core guest on a standalone system

Example: Dual-slot LTM guests within a device group

If you have a redundant system consisting of two VIPRION® chassis, you can deploy a vCMP® guest on each chassis, where each guest is provisioned to run BIG-IP® Local Traffic Manager™ (LTM®) on two slots in the VIPRION cluster.

With this configuration, the host has allocated twice the amount of CPU and memory to the guest than a configuration where the guest is assigned to a single slot only. By putting both guests in a BIG-IP Sync-Failover device group, you are assured that when failover occurs, the LTM guest can continue processing application traffic.

Note: For best results, particularly when connection mirroring is enabled, configure the two guests so that the slot numbers and amount of core allocation for the two guests match.

The following illustration depicts the deployment of LTM within a two-slot, four-core guest on each VIPRION chassis in a two-member device group.




Figure 8: Dual-slot, four-core guests in a device group

Example: Multi-slot guests within device groups

A common use of a vCMP[®] system is to create a redundant system configuration with multiple guests, where each guest contains a different set of BIG-IP[®] modules, with varying amounts of system resource allocated to each guest. In this case, the system is in a redundant configuration consisting of two separate VIPRION[®] systems. For each guest, you can create an equivalent peer guest on the other VIPRION system and create a Sync-Failover device group with the two equivalent guests as members. If failover occurs, the equivalent guest on the peer system can assume the processing of the guest's application traffic.

The following illustration depicts the deployment of BIG-IP guests on multiple populated slots, on two VIPRION chassis. The illustration shows that each guest has an equivalent guest on a peer chassis and that each pair of equivalent guests comprises a separate device group, resulting in a total of four device groups.

Each guest in the first three device groups has either eight, four, or six cores, and spans either four two, or three slots, respectively. The guests in the fourth device group are single-core, single-slot guests.




Figure 9: Multiple guests in device groups

Device Service Clustering for vCMP Systems

Overview: Device service clustering for vCMP systems

One of the tasks of a vCMP[®] guest administrator is to configure device service clustering (DSC[®]). Using *DSC*, a guest administrator can implement config sync, failover, and mirroring across two or more chassis. Configuring DSC is the same on a vCMP system as on non-virtualized systems, except that the members of a device group are virtual devices (guests) rather than physical devices.

When configuring DSC, a guest administrator creates a device group that consists of vCMP guests as members, where each member is deployed on a separate chassis.

For example, a Sync-Failover device group in an active-standby configuration can consist of:

- `guest_A on chassis_1 and guest_A on chassis_2`
- `guest_B on chassis_1 and guest_B on chassis_2`
- `guest_C on chassis_1 and guest_C on chassis_2`

Creating a device group that consists of guests on separate chassis ensures that if a chassis goes out of service, any active traffic groups on a guest can fail over to a device group member on another chassis.

This illustration shows this DSC configuration. The illustration shows two four-slot chassis, with four guests on each chassis. Each guest and its equivalent guest on the other chassis are homogeneous (same slot numbers and same number of cores) and form a separate Sync-Failover device group. Note that homogeneous guests in a device group are only required when connection mirroring is enabled.




Figure 10: vCMP guests forming four device groups across two chassis

Required IP addresses for DSC configuration

This table describes the types of IP addresses that a guest administrator specifies when configuring device service clustering (DSC[®]) on a vCMP[®] system.

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

Configuration feature	IP addresses required
Device trust	The cluster IP address that the vCMP host administrator assigned to the guest during guest creation.
Config sync	Any non-floating self IP address on the guest that is associated with an internal VLAN on the host.
Failover	<ul style="list-style-type: none"> Recommended: A unicast non-floating self IP address on the guest that is associated with an internal VLAN on the host (preferably VLAN_{HA}), as well as a multicast address. Alternate to a multicast address: The guest-unique cluster member IP addresses assigned to all slots in the guest's cluster.

Configuration feature	IP addresses required
Connection mirroring	For both the primary and the secondary IP addresses, a non-floating self IP address on the guest that is associated with an internal VLAN on the host. The secondary address is optional.

Failover methods for vCMP guests

Each traffic group in a device service clustering (DSC[®]) device group has a property known as a failover method. The *failover method* dictates the way that the system chooses a target device for failover. Available failover methods that the user can choose from are: load-aware failover, an ordered list, and an HA group.

The specific core allocation and slot assignments for a guest in a Sync-Failover device group determine the particular failover method that is appropriate for a DSC traffic group within the guest:

- Guests in a device group that are identical in terms of core allocation and slot assignment are considered to be *homogeneous* guests. In this case, an ordered list would be an appropriate failover method, since relative capacity is equal among all guests. Note that guests in a device group are required to be homogeneous when connection mirroring is enabled.
- Guests in a device group that differ from one another in terms of core allocation and slot assignments are considered to be *heterogeneous* guests. In this case, load-aware failover is an appropriate failover method because the guest administrator can define a relative capacity and relative traffic load for each guest. For example, an eight-core, four-slot guest has a relative capacity that is twice that of a four-core, two-slot guest. Note that you cannot enable connection mirroring for heterogeneous guests in a device group.

Important: For guests in a device group, you can enable connection mirroring for homogeneous guests only, that is, guests that are assigned to the same slot numbers with the same number of cores.

An additional type of failover method is an HA group, which applies to both homogeneous and heterogeneous guests.

About HA groups for vCMP systems

For failover configuration, an alternative to using load-aware failover or an ordered list is to use HA groups. An *HA group* is a specification of certain pools or host trunks (or any combination of these) that a guest administrator associates with a traffic group instance. The most common reason to configure an HA group is to ensure that failover is triggered when some number of trunk members become unavailable.

The BIG-IP[®] system uses an HA group to calculate an overall health score for the instance of a traffic group on a guest. The instance of a traffic group that has the best overall score at any given time becomes or remains the active traffic group instance. With an HA group, the system triggers failover of a traffic group based on changes to trunk or pool health instead of on system, gateway, or VLAN failure.

Because trunks and HA groups are never synchronized among guests as part of a config sync operation, you must assign a separate HA group to each traffic group instance. For example, you could create `ha_group_A` to reference the host trunk `my_trunk` and assign the HA group to `traffic-group-1` on `guest_A`. You could then create another HA group, `ha_group_B`, to also reference `my_trunk` and assign the HA group to the same traffic group (`traffic-group-1`) on `guest_B`.

About connection mirroring for vCMP systems

Connection mirroring is a device service clustering (DSC[®]) feature that allows a device to mirror its connection and persistence information to another device. Connection mirroring prevents interruption in service during failover. On a vCMP[®] system, the devices that mirror their connections to each other are virtual devices (vCMP guests).

Important: *When you enable connection mirroring within a device group, a guest can only mirror its connections to one other guest. In this case, the two guests must be homogenous. That is, as mirrored peers, the guests must each reside on a separate chassis where the two chassis and the guests' blades are the same model. Also, the guests must have the same number of slots assigned, on the same slot numbers, and with the same number of cores per slot.*

About switchboard failsafe and vCMP guests

If a vCMP[®] guest is a member of a device group, make sure the guest's switchboard failsafe setting is set to the default value. If you need to change the default switchboard failsafe configuration, always do this on the vCMP host, and not the guest.

Initial vCMP Configuration Tasks

About vCMP application volume management

When you provisioned the vCMP[®] feature as part of VIPRION[®] system setup, the BIG-IP[®] system allocated most of the total disk space to the vCMP application volume (by default, all but 30 gigabytes). Known as the *reserve disk space*, this 30 gigabytes of disk space is left available for other uses, such as for installing additional versions of the BIG-IP system in the future.

Important: Do not attempt to change the amount of reserved disk space after you have provisioned the vCMP feature. Changing the reserved disk space after provisioning produces unwanted results.

vCMP host administrator tasks

Important: Before you configure the vCMP[®] host, make sure you followed the VIPRION setup tasks described in the guide *VIPRION Systems: Configuration*. After using that guide, you should now have a VIPRION[®] system that is provisioned for vCMP, with the standard external, internal, and high-availability VLANs configured.

As a vCMP[®] host administrator, you have the important task of initially planning the amount of total system CPU and memory that you want the vCMP host to allocate to each guest. This decision is based on the resource needs of the particular BIG-IP[®] modules that guest administrators intend to provision within each guest, as well as the maximum system resource limits for the relevant hardware platform. Thoughtful resource allocation planning prior to creating the guests ensures optimal performance of each guest. Once you have determined the resource allocation requirements for the guests, you are ready to configure the host.

Overall, your primary duties are to create and manage guests, ensuring that the proper system resources are allocated to those guests.

Task summary

Accessing the vCMP host

Creating a vCMP guest

Setting a vCMP guest to the Deployed state

Accessing the vCMP host

Before accessing the vCMP[®] host, verify that you have created a primary cluster management IP address. For information on creating this address, see the guide titled *VIPRION[®] Systems: Configuration*.

Performing this task allows you to access the vCMP host. Primary reasons to access the host are to create and manage vCMP[®] guests, manage virtual disks, and view or manage host and guest properties. You can also view host and guest statistics.

1. From a system on the external network, display a browser window.
2. In the URL field, type the primary cluster management IP address for the chassis, as follows:

```
https://<ip_address>
```

The browser displays the login screen for the BIG-IP[®] Configuration utility.

Creating a vCMP guest

Before creating a guest on the system, verify that you have configured the base network on the system to create any necessary trunks, as well as VLANs for guests to use when processing application traffic.

You create a guest when you want to create an instance of the BIG-IP software for the purpose of running one or more BIG-IP® modules to process application traffic. For example, you can create a guest that runs BIG-IP® Local Traffic Manager™ and BIG-IP® DNS. When creating a guest, you specify the number of logical cores per slot that you want the vCMP host to allocate to each guest, as well as the specific slots that you want the host to assign to the guest.

Note: When creating a guest, if you see an error message such as *Insufficient disk space on /shared/vmdisks. Need 24354M additional space.*, you must delete existing unattached virtual disks until you have freed up that amount of disk space.

Important: If you are planning to add this guest to a Sync-Failover device group and enable connection mirroring with a guest on another chassis, you must ensure that the two guests are configured identically with respect to slot assignment and core allocation. That is, the number of cores, the number of slots, and even the slot numbers on which the guests reside must be the same. Therefore, you must ensure that on each guest of the mirrored pair, the values match for the **Cores per Slot**, **Number of Slots**, **Minimum Number of Slots**, and **Allowed Slots** settings.

1. Use a browser to log in to the VIPRION® chassis, using the primary cluster management IP address. If you provisioned the system for vCMP®, this step logs you in to the vCMP host.
2. On the Main tab, click **vCMP > Guest List**. This displays a list of guests on the system.
3. Click **Create**.
4. From the **Properties** list, select **Advanced**.
5. In the **Name** field, type a name for the guest.
6. In the **Host Name** field, type a fully-qualified domain name (FQDN) name for the guest. If you leave this field blank, the system assigns the name `localhost.localdomain`.
7. From the **Cores Per Slot** list, select the total number of logical cores that the guest needs, based on the guest's memory requirements.

The value you select causes the host to assign that number of cores to each slot on which the guest is deployed. The host normally allocates cores per slot in increments of two (two, four, six, and so on).

Important: Cores for a multi-slot guest do not aggregate to provide a total amount of memory for the guest. Therefore, you must choose a **Cores per Slot** value that satisfies the full memory requirement of the guest. After you finish creating the guest, the host allocates this amount of memory to each slot to which you assigned the guest. This ensures that the memory is sufficient for each guest if any blade becomes unavailable. For blade platforms with solid-state drives, you can allocate a minimum of one core per guest instead of two. For metrics on memory and CPU support per blade model, see the vCMP® guest memory/CPU allocation matrix at <http://support.f5.com>.

8. From the **Number of Slots** list, select the maximum number of slots that you want the host to allocate to the guest.
9. From the **Minimum Number of Slots** list, select the minimum number of chassis slots that must be available for this guest to deploy.

Important: The minimum number of slots you specify must not exceed the maximum number of slots you specified.

10. From the **Allowed Slots** list, select the specific slots that you want the host to assign to the guest and then use the Move button to move the slot number to the **Selected** field.

Important: If you want to allow the guest to run on any of the slots in the chassis, select all slot numbers. For example, if you configure the **Number of Slots** value to be 2, and you configure the **Allowed Slots** values to be 1, 2, 3, and 4, then the host can assign any two of these four slots to the guest. Note that the number of slots in the **Allowed Slots** list must equal or exceed the number specified in the **Minimum Number of Slots** list.

11. From the **Management Network** list, select a value:

Value	Result
Bridged (Recommended)	Connects the guest to the management network. Selecting this value causes the IP Address setting to appear.
Isolated	Prevents the guest from being connected to the management network and disables the host-only interface.

Important: If you select **Isolated**, do not enable the **Appliance Mode** setting when you initially create the guest. For more information, see the step for enabling the **Appliance Mode** setting.

Host-Only	Prevents the guest from being connected to the management network but ensures that the host-only interface is enabled.
------------------	--

12. If the **IP Address** setting is displayed, specify the required information:

- In the **IP Address** field, type a unique management IP address that you want to assign to the guest.
You use this IP address to access the guest when you want to manage the BIG-IP modules running within the guest.
- In the **Network Mask** field, type the network mask for the management IP address.
- In the **Management Route** field, type a gateway address for the management IP address.

Important: Assigning an IP address that is on the same network as the host management port has security implications that you should carefully consider.

13. From the **Initial Image** list, select an ISO image file for installing TMOS[®] software onto the guest's virtual disk.

14. In the **Virtual Disk** list, retain the default value of **None**.

Note that if an unattached virtual disk file with that default name already exists, the system displays a message, and you must manually attach the virtual disk. You can do this using the `tmsh` command line interface, or use the Configuration utility to view and select from a list of available unattached virtual disks.

The BIG-IP system creates a virtual disk with a default name (the guest name plus the string `.img`, such as `guestA.img`).

15. For the **VLAN List** setting, subscribe to host-based VLANs:

- Select the external and internal VLANs from the **Available** list.
- Use the Move button to move the VLANs to the **Selected** list.

After you create the guest, the guest will use the selected VLANs to process application traffic. As an option, the guest administrator can create additional VLANs later from within the guest.

16. From the **Requested State** list, select **Provisioned**.

Once the guest is created, the vCMP host allocates all necessary resources to the guest, such as cores and virtual disk.

17. If you want to enable Appliance mode for the guest, select the **Appliance Mode** check box.

Warning: Before enabling this feature on an isolated guest, you must perform some prerequisite tasks, such as creating a self IP address on the guest. Failure to perform these prerequisite tasks will make the guest unreachable by all host and guest administrators. Therefore, you must create the isolated guest with Appliance mode disabled, perform the prerequisite tasks, and then modify the guest to enable this setting. For more information, see the relevant appendix of this guide.

When you enable **Appliance Mode** for a guest, the system enhances security by denying access to the `root` account and the `Bash` shell for all administrators.

18. From the **SSL-Mode** list:

- Select **Dedicated** to assign dedicated SSL hardware resources, in the form of SSL cores, to the guest. A guest in **Dedicated** mode has a fixed amount of SSL hardware resource available and does not share that resource with other guests on the system. Consequently, SSL performance for a guest in **Dedicated** mode is not impacted by other guests' use of SSL hardware resources. The number of SSL cores that the system assigns to the guest is based on the number of vCMP cores allocated to the guest.
- Select **Shared** to give the guest access to all available SSL hardware resources, that is, resources not used by guests in **Dedicated** mode. In **Shared** mode, the guest shares SSL hardware resources with all guests that are also in **Shared** mode. This option can impact SSL performance for the guest, depending on use of SSL resources by other guests. Guests in **Shared** mode do not impact the SSL performance of guests in **Dedicated** mode.
- Select **None** to prevent the guest from accessing SSL hardware resources. When you select **None**, the guest has no access to SSL hardware resources, but can access SSL software resources.

Important: If you do not see the **SSL-Mode** setting, your hardware platform does not support this feature.

19. From the **Single Rate TCM Policer** list:

- Select **None** if you do not want to meter network traffic using a Single Rate Three Color Marker (srTCM) policer.
- Select the name of an existing srTCM policer if you want the BIG-IP system to classify network traffic as green, yellow, or red using the srTCM standard.

20. Click **Finish**.

The system installs the selected ISO image onto the guest's virtual disk and displays a status bar to show the progress of the resource allocation.

You now have a new vCMP guest on the system in the Provisioned state with an ISO imaged installed.

After you create the guest, if an administrator needs to change the maximum transmission unit (MTU) size on a host-based VLAN to optimize the guest's application traffic, the administrator can (and must) change the MTU value from within the guest. An administrator for a specific guest should never try to change the MTU value of a host-based VLAN when logged into the vCMP host.

Setting a vCMP guest to the Deployed state

Setting a guest to the Deployed state enables a guest administrator to then provision and configure the BIG-IP® modules within the guest.

Warning: For any isolated guest with Appliance mode enabled, you must first perform some additional tasks before deploying the guest. For more information, see the relevant appendix of this guide.

1. Ensure that you are logged in to the vCMP host.
2. On the Main tab, click **vCMP > Guest List**.
This displays a list of guests on the system.
3. In the Name column, click the name of the vCMP guest that you want to deploy.

4. From the **Requested State** list, select **Deployed**.
5. Click **Update**.

After moving a vCMP® guest to the Deployed state, a guest administrator can provision and configure the BIG-IP modules within the guest so that the guest can begin processing application traffic.

vCMP guest administrator tasks

The primary duties of a vCMP® guest administrator are to provision BIG-IP® modules within the guest, configure the correct management IP addresses for the slots pertaining to the guest, and configure any self IP addresses that the guest needs for processing application traffic. The guest administrator must also configure all BIG-IP modules, such as creating virtual servers and load balancing pools within BIG-IP Local Traffic Manager™ (LTM®).

Optionally, a guest administrator who wants a redundant system configuration can create a device group with the peer guests as members.

Task list

Provisioning BIG-IP modules within a guest

Specifying cluster member IP addresses for a guest

Creating a self IP address for application traffic

Changing the MTU value on a VLAN (optional)

Provisioning BIG-IP modules within a guest

Before a guest administrator can access a guest to provision licensed BIG-IP® modules, the vCMP® guest must be in the Deployed state.

To run BIG-IP modules within a guest, the guest administrator must first provision them. For example, a guest administrator for `guestA` who wants to run LTM® and DNS must log into `guestA` and provision the LTM and BIG-IP DNS modules.

Note: For guests that are isolated from the management network, you must access them using a self IP address instead of a management IP address.

1. Open a browser, and in the URL field, specify the management IP address that the host administrator assigned to the guest.
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**.
This displays the Resource Provisioning screen.
4. For each licensed BIG-IP module in the list, select the check box and select **Minimal**, **Nominal**, or **Dedicated**.
5. Click **Next**.
This displays the Certificate Properties screen.
6. Click **Next**.
This displays some general properties of the guest.
7. Click **Next**.
This displays the screen for specifying the guest's cluster member IP addresses.
8. Click **Next**.
9. Click **Finished**.

Specifying cluster member IP addresses for a guest

For each vCMP® guest, the guest administrator needs to create a unique set of management IP addresses that correspond to the slots of the VIPRION® cluster. Creating these addresses ensures that if a blade becomes unavailable, the administrator can log in to another blade to access the guest.

1. On the Setup utility screen for resource provisioning, in the Cluster Member IP Address area, type a management IP address for each slot in the VIPRION chassis, regardless of how many blades are installed or how many slots are assigned to the guest.
Each IP address must be on the same subnet as the management IP address that the host administrator assigned to the guest (displayed).
2. Click **Next**.
3. Click **Finished**.

After performing this task, a guest administrator can log in to a specific slot for a guest if blade availability becomes compromised.

Creating a self IP address for application traffic

A vCMP® guest administrator creates a self IP address within a guest, assigning a VLAN to the address in the process. The self IP address serves as a hop for application traffic destined for a virtual server configured within the guest. On a standalone system, the self IP address that a guest administrator creates is a static (non-floating) IP address. Note that the administrator does not need to create VLANs within the guest; instead, the VLANs available for assigning to a self IP address are VLANs that a host administrator previously created on the vCMP host.

1. On the Main tab of the BIG-IP Configuration utility, click **Network > Self IPs**.
2. Click **Create**.
The New Self IP screen opens.
3. In the **Name** field, type a unique name for the self IP address.
4. In the **IP Address** field, type an IPv4 or IPv6 address.
This IP address should represent the address space of the VLAN that you specify with the **VLAN/Tunnel** setting.
5. In the **Netmask** field, type the network mask for the specified IP address.
For example, you can type `255.255.255.0`.
6. From the **VLAN/Tunnel** list, select the VLAN to associate with this self IP address.
 - On the internal network, select the internal or high availability VLAN that is associated with an internal interface or trunk.
 - On the external network, select the external VLAN that is associated with an external interface or trunk.
7. From the **Port Lockdown** list, select **Allow Default**.
8. Click **Finished**.
The screen refreshes, and displays the new self IP address.

After creating a self IP address, the BIG-IP system can send and receive traffic destined for a virtual server that allows traffic through the specified VLAN.

Changing the MTU value on a VLAN (optional)

Do this task when you need to adjust the maximum transmission unit (MTU) size on a VLAN for the vCMP® guest that you are logged into. Changing a VLAN's MTU size can help to optimize application

traffic for the guest. You can do this task for either a host-based VLAN or a VLAN that you created from within the guest.

Important: Always do this task when you're logged into the guest and not the host.

1. Log into the guest using the guest's management IP address.
The BIG-IP® Configuration utility opens.
2. On the Main tab of the BIG-IP Configuration utility, click **Network > VLAN**.
A list of VLANs appears.
3. In the Name column, double-click the name of the VLAN you want to modify.
This displays the properties of the VLAN.
4. In the MTU field, change the value to whatever is appropriate for the guest.
5. Click Update.

vCMP guest administrator tasks

Creating a self IP address for application traffic

Next steps

After all guests are in the Deployed state, each individual guest administrator can configure the appropriate BIG-IP modules for processing application traffic. For example, a guest administrator can use BIG-IP® Local Traffic Manager™ (LTM®) to create a standard virtual server and a load-balancing pool. Optionally, if guest redundancy is required, a guest administrator can set up device service clustering (DSC®).

Another important task for a guest administrator is to create other guest administrator accounts as needed.

Important: If the guest has an isolated (rather than bridged) management network, you must grant access to the Traffic Management Shell (`tmssh`) to all guest administrator accounts. Otherwise, guest administrators have no means of logging in to the guest, due to the lack of access to the management network.

Configuration results

After you and all guest administrators have completed the initial configuration tasks, you should have a VIPRION® system provisioned for vCMP, with one or more guests ready to process application traffic.

When logged in to the vCMP® host, you can see the VLANs and trunks configured on the VIPRION system, as well as all of the guests that you created, along with their virtual disks. When using the BIG-IP Configuration utility, you can also display a graphical view of the number of cores that the host allocated to each guest and on which slots.

You can also view the current load on a specific guest in terms of throughput, as well as CPU, memory, and disk usage.

When logged in to a guest, the guest administrator can see one or more BIG-IP® modules provisioned and configured within the guest to process application traffic. If the guest administrator configured device service clustering (DSC®), the guest is a member of a device group.

CPU Core Upgrade for Deployed Guests

About increasing CPU cores for a deployed guest

Overview

If you have a guest that's actively processing application traffic and you decide that it needs additional CPU cores than its original core allocation, you can increase the number of cores per slot for the guest. To do this, you'll need to temporarily disable the guest while you add cores, and then re-deploy the guest.

About CPU core availability

When you increase the **Cores per Slot** value on a guest, the new value must be within the total number of unused CPU cores on the slot; the system never allocates more cores per slot to a guest than the number of cores still available on the slot.

For example, suppose you have three guests on a 12-core slot, with this core allocation:

Guest Name	Current Cores per Slot
Guest A	6
Guest B	2
Guest C	2

With a total allocation of 10 out of 12 cores allocated per slot, this means that only two cores remain unused and available to add to a guest. Therefore, if you try to add four additional cores to **Guest C** (from **2** to **6**), the system displays a message that you are exceeding the number of cores available on the slot. In this case, the maximum number of cores per slot that you can allocate to **Guest C** is **4**.

A potential workaround for this would be to first decrease the number of cores per slot allocated to **Guest A**, to free up additional cores for **Guest C**.

You can determine the number of cores still available on each slot by logging in to the BIG-IP Configuration utility and displaying the vCMP Guest List screen.

Increasing CPU cores for a deployed guest

Before doing this task, confirm that your BIG-IP[®] user account has an Administrator role assigned to it.

If you have a deployed vCMP guest, and you decide that you need to allocate more CPU cores per slot to the guest than you had originally allocated, you can increase the number of cores per slot for the guest.

To do this, you'll first need to set the guest state from **Deployed** to **Configured**. Once you've increased the number of cores per slot for the guest, you can set the guest back to the **Deployed** state.

Note that when switching between guest states, the process can take several minutes for the guest to shut down cleanly and restart again.

Important: When you set a guest to the **Configured** state, the guest is removed from service until you set the guest back to the **Deployed** state.

1. Log in to the BIG-IP Configuration utility on the vCMP host, using the VIPRION[®] system's primary management IP address.
2. On the Main tab, click **vCMP > Guest List**.
3. View the graphic to determine the number of unused cores available on the slot.
This tells you how many cores are available to add to the guest.
4. In the Name column, find the name of the guest that needs additional CPU cores and confirm that the Requested State is **Deployed**.
5. To the left of the guest name, select the check box.
6. Click the **Disable** button.
The guest state changes from **Deployed** to **Configured**.

***Important:** It might take a few minutes for the guest state to change to **Configured**. Although the Guest List screen does not display a progress indicator during this state change, you can log into the console and view the status of the guest at any time, using the command `tmsl show vcmp guest`.*

7. Once the screen shows that the guest is now in the **Configured** state, click the guest name.
The guest properties screen opens.
8. From the **Cores per Slot** list, select a higher number of CPU cores.

***Note:** Do not select a value that exceeds the total number of unused cores available for use on the slot.*

For example, if you currently have two cores per slot allocated to the guest and you want to add two cores, then from the **Cores per Slot** list, select **4**.

9. From the **Requested State** list, select **Deployed**.
10. Click **Update**.
The Guest List screen opens again and you will see that the guest state is changing.

***Important:** It might take a few minutes for the guest state to change back to the **Deployed** state.*

After you complete this task, the guest has additional CPU cores per slot and is capable of processing application traffic.

Managing vCMP Virtual Disks

Overview: Managing virtual disks

A *virtual disk* is the portion of disk space on a slot that the system has allocated to a guest. For example, if a guest spans three slots, the system creates three virtual disks for that guest, one per slot. Each virtual disk is implemented as an image file with an `.img` extension, such as `guest_A.img`.

You do not explicitly create virtual disks. The vCMP® system automatically creates a virtual disk when you set a guest to the Provisioned or Deployed state.

Using the BIG-IP® Configuration utility or the Traffic Management Shell (`tmsh`), you can delete virtual disks on the system as a way to optimize disk space.

About virtual disk allocation

For each slot that you assign to a vCMP® guest, the host automatically creates a sparse file to be used as a virtual disk. This amount of disk space can grow to 100 GB, and is not dependent on the number of cores per slot that you configure for that guest. For example, a slot with two cores allocated to `guest_A` could provide the same amount of available disk space for the guest as a slot with four cores allocated to that guest.

Note that you cannot explicitly create virtual disks; instead, the BIG-IP® system creates virtual disks when the guest changes to a Provisioned or Deployed state. You can create a guest that remains in the Configured state, but in this case, the guest has no virtual disk allocated to it.

About virtual disk images

A virtual disk is in the form of an image that resides in the `/shared/vmdisks` directory on each physical blade. The default file name that the BIG-IP® system initially assigns to a virtual disk is the guest name plus an `.img` extension (for example, `guestA.img`). Using the BIG-IP Configuration utility or the Traffic Management Shell (`tmsh`), you identify and manage virtual disks on the system using these file names.

A virtual disk image for a guest resides on each slot assigned to that guest.

About virtual disk templates

If you need to create multiple guests, you most likely want to minimize the time that the vCMP® system needs to create all of the virtual disks. The vCMP system automatically accomplishes this through a feature known as virtual disk templates. A *virtual disk template* is a virtual disk image that contains a fresh installation of an initial ISO image. Its purpose is to minimize the time that the system uses to create virtual disks on the system.

When you provision a guest on the system, with a specific version of BIG-IP software installed to the relevant blades, the system automatically creates a virtual disk template locally on each blade, pertaining to that ISO image. For example, when you provision a guest on four slots of the cluster, the system creates a template locally on each of the four associated blades. Later, when you create other guests that use the same ISO image, the system instantiates a copy of the virtual disk template to more rapidly create

the virtual disks for those guests. The vCMP system creates a separate virtual disk template for each initial image that you initially configure for a guest.

No user intervention is required to use this feature. On the vCMP system, you can view a list of the system-created templates, or you can delete a template, but you cannot explicitly create or modify a template.

Important: By default, the virtual disk template feature is enabled on hardware platforms with solid state drives and disabled on platforms with spinning hard drives. If you want to use virtual disk templates on platforms with spinning drives, you must explicitly enable the feature, using the `db` variable `vcmp.installer.use_vdisk_templates`.

Viewing the list of virtual disk templates

Before performing this task, confirm that you have created and provisioned at least one vCMP® guest after upgrading the host to the latest version.

You perform this task when you want to view the virtual disk templates that the vCMP system has created.

Note: The virtual disk template list shows a separate virtual disk template for each initial image that you initially configured for a guest.

1. On the Main tab, click **vCMP > Template List**.
2. View all information displayed.
For example, the following shows a sample list of virtual disk templates on the vCMP host.




Figure 11: List of virtual disk templates

After performing this task, you can see the virtual disk templates that the vCMP system can use when installing the initial image.

Deleting virtual disk templates

You perform this task when you want to delete a virtual disk template on the vCMP host. On the host, there is a separate virtual disk template corresponding to each initial image that you previously installed

on a guest. The reason for deleting virtual disk templates is to conserve disk space. You should delete any virtual disk templates that the host will no longer use when creating vCMP guests.

1. On the Main tab, click **vCMP > Template List**.
2. In the Name column, locate the name of the virtual disk template that you want to delete.
3. To the left of the virtual disk template name, select the check box.
4. Click **Delete**.
The system prompts you to confirm the delete action.
5. Click **Delete**.

After performing this task, the deleted virtual disk template is no longer available for the vCMP system to use. Note, however, that the system can recreate the template if another guest is provisioned using that same software version.

Enabling and disabling the virtual disk template feature

You can perform this task to enable or disable the virtual templates feature on any vCMP-enabled system. The virtual templates feature is useful for minimizing the time that the system uses to create virtual disks on the system. By default, the feature is enabled on platforms with solid-state drives. On platforms with spinning drives, the virtual disk templates feature is automatically disabled due to potential stress and latency on spinning drives during guest provisioning. For this reason, F5 Networks recommends that for platforms with spinning drives, you enable virtual disk templates in a test environment only, whenever you need to create multiple guests running the same BIG-IP software version.

1. Log in to the BIG-IP system and access `tmsh`.
2. At the `tmsh` command prompt, type `modify sys db vcmp.installer.use_vdisk_templates value default|enabled|disabled`

Value	Description
default	When set to default , the db variable <code>vcmp.installer.use_vdisk_templates</code> enables the virtual disk templates feature on any vCMP-enabled platforms with solid-state drives and disables virtual disk templates on any vCMP-enabled platforms with spinning drives. The default value is default .
enabled	When set to enabled , the db variable <code>vcmp.installer.use_vdisk_templates</code> enables the virtual disk templates feature on all vCMP-enabled hardware platforms, regardless of drive type.
disabled	When set to disabled , the db variable <code>vcmp.installer.use_vdisk_templates</code> disables the virtual disk templates feature on all vCMP-enabled hardware platforms, regardless of drive type.

Note: The virtual disk template feature is not supported on the B4200 platform.

Viewing the virtual disk templates db variable

You can perform this task to view the current value of the db variable `vcmp.installer.use_vdisk_templates`.

1. Log in to the BIG-IP system and access `tmsh`.
2. At the `tmsh` command prompt, type `list sys db vcmp.installer.use_vdisk_templates`
The BIG-IP system displays the current value of the db variable `vcmp.installer.use_vdisk_templates`.

About virtual disk detachment and re-attachment

When a vCMP® guest has no virtual disk and moves from the Configured state to the Provisioned state, the system creates a virtual disk and attaches the disk to the guest. This attachment ensures that only that guest can use the virtual disk. A guest can have only one virtual disk attached to it at any one time.

A virtual disk can become unattached from a guest when you perform one of these actions:

- Delete a guest.
- Change the **Virtual Disk** property of the guest to **None**. Note that to perform this action, you must first change the guest state to Configured.

With either of these actions, the system retains the virtual disks on the system for future use.

You can attach an existing, unattached virtual disk to a new guest that you create. Attaching an existing virtual disk to a newly-created guest saves the BIG-IP® system from having to create a new virtual disk for the guest.

Detaching virtual disks from a vCMP guest

Before you can detach a virtual disk from a guest, you must be logged into the vCMP host. Also, you must change the **Requested State** property on the guest to **Configured**.

You can detach a virtual disk from the guest, but retain the virtual disk on the BIG-IP® system so that you can attach it to another guest later.

Important: *Unattached virtual disks consume disk space on the system. To prevent unattached virtual disks from depleting available disk space, routinely monitor the number of unattached virtual disks that exist on the system.*

1. On the Main tab, click **vCMP > Guest List**.
This displays a list of guests on the system.
2. In the Name column, locate the relevant guest name, and to the left of the name, select the check box.
3. Click the **Configured** button.
4. In the Name column, click the guest name.
5. From the **Virtual Disk** list, select the default value, **None**.
6. Click **Update**.

The vCMP guest no longer has any virtual disk attached to it.

Viewing virtual disks not attached to a vCMP guest

Before you can view unattached virtual disks, you must be logged into the vCMP host.

You can view virtual disks that are not attached to a vCMP® guest so that you can monitor virtual disks that might be unused but still consuming disk space.

1. On the Main tab, click **vCMP > Virtual Disk List**.
2. Locate the Virtual Disk List area of the screen.
3. To the right of the list of virtual disk names, note any disks that do not have any guest names associated with them. These disks are unattached.

Attaching a detached virtual disk to a vCMP guest

Before you begin this task, ensure that:

- You are logged into the vCMP® host.
- The guest to which you are attaching the virtual disk is in the Configured state.
- The virtual disk is not currently be attached to another guest.

It is possible for a virtual disk to become detached from a vCMP guest. A disk that is no longer attached to a guest is known as an *unattached virtual disk*.

You can attach an unattached virtual disk to another guest either when you create the guest or when you modify the **Virtual Disk** property of a guest.

1. On the Main tab, click **vCMP > Guest List**.
This displays a list of guests on the system.
2. In the Name column, click the name of the guest that you want to view.
3. From the **Properties** list, select **Advanced**.
4. From the **Virtual Disk** list, select a file name.
The guest uses the newly-selected virtual disk when being deployed.
5. Click **Update**.

About virtual disk migration

Whenever the vCMP® system re-assigns a guest to other slots, the system must also migrate the guest's virtual disks to the new slots. This virtual disk migration occurs automatically; you do not need to explicitly manage this migration.

Deleting a virtual disk from the BIG-IP system

Before deleting a virtual disk, ensure that you are logged into the vCMP® host.

Using the BIG-IP® Configuration utility, you can delete a virtual disk from the system.

Important: *This is the only way to delete a virtual disk from the system. If you delete the associated guest instead, the system retains the virtual disk for re-use by another guest later.*

1. On the Main tab, click **vCMP > Virtual Disk List**.
2. Locate the Virtual Disk List area of the screen.
3. In the Name column, locate the name of the virtual disk that you want to delete.
4. To the left of the virtual disk name, select the check box.
5. Click **Delete**.
The system prompts you to confirm the delete action.
6. Click **Delete**.

Deleting a vCMP application volume

Whenever you de-provision the vCMP® feature, you must also delete its vCMP application volumes (named **vmdisks**) from the relevant software volume (boot location). There is one **vmdisks** volume for each blade that is assigned to one or more guests, for a specific software volume. De-provisioning the vCMP feature and deleting its application volumes allows you to perform certain disk management tasks such as increasing the amount of disk space that the BIG-IP® system reserves for non-vCMP uses.

Warning: *Deleting vCMP application volumes deletes all guest configuration data. Therefore, prior to deleting vCMP application volumes, F5 Networks® strongly recommends that you create a UCS file for*

each guest configuration. This allows you to easily re-create the guests if you decide to provision the vCMP feature again later.

Important: When the BIG-IP system initially created a vCMP application volume for each assigned blade, the system also created a set of 2-GB, MySQL volumes in the same software volume as the vCMP application volumes. If you decide to de-provision vCMP and delete its application volumes, you should also delete the MySQL volumes in that software volume. Retaining these MySQL volumes consumes disk space that could negatively impact your ability to successfully provision other BIG-IP modules later. Be careful, however, not to delete MySQL volumes that reside in other software volumes.

1. Use a browser and the 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 > Disk Management**.
The display shows the logical disks and application volumes from the perspective of the vCMP host.
 3. Click the logical disk for which you want to reserve disk space.
An example of a logical disk is HD1.
 4. On the menu bar, click **Image List** if displayed.
The screen displays a list of the installed images on the system.
 5. If a list of images appears, locate the relevant image, and in the Disk column, click the logical disk name.
 6. In the Contained Application Volumes area of the screen, to the left of the list of application volume names, select the boxes for the per-blade vCMP application volumes (named `vmdisks`), as well as any associated MySQL volumes in that same software volume.
-

Important: Be careful not to delete MySQL application volumes pertaining to other software volumes.

7. Click **Delete**.

After you perform this task, the BIG-IP system should have enough disk space to accommodate the provisioning of other BIG-IP modules.

Installing ISO images within vCMP guests

About ISO images

BIG-IP® software images that are stored and managed on the vCMP® host are available for vCMP guests to install. The vCMP host presents a list of those images within each guest for guest administrators to use as needed.

Installing updates and hotfixes on the host for guests to use offers these benefits:


- You save time because you remove the need to repeatedly copy the same ISO image into each guest's `/shared/images` folder.
- You reduce the impact on the management network.
- You conserve storage space on the vCMP system.

Viewing a list of host ISO images from within a guest

vCMP® guest administrators perform this task to view any ISO images that reside on the vCMP host and are available for installation on the guest. All ISO images that the host administrator has imported into the host's `/shared/images` folder automatically appear on each guest as available for installation.

1. On the Main tab, click **System > Software Management > Image List**.
The Image List screen displays a list of existing image files.
2. In the **Available Images** area of the screen, in the Image Source column, view the ISO images that show a value of **Host**.

For example, the following shows a sample list of ISO images available on the vCMP host for installation on the guest.



The screenshot shows the 'Image List' interface with two main sections: 'Installed Images' and 'Available Images'. The 'Installed Images' section contains a table with columns for Product, Version, Build, Media, and Install Status. The 'Available Images' section contains a table with columns for Status, Software Image, Version, Last Modified, Image Size, MD5 Verified, Image Source, and Available. Three ISO images are listed as available, all with 'Host' as the source.

Installed Images				
Product	Version	Build	Media	Install Status
BIG-IP	11.5.0	0.0.25	td	complete
None	None	None	td	empty
None	None	None	td	empty

Available Images								Import...
✓ Status	Software Image	Version	Last Modified	Image Size	MD5 Verified	Image Source	Available	
<input type="checkbox"/>	BIGIP-11.5.0.0.0.264.iso	11.5.0	Thu Jun 26 03:50:51 2014	1903 MB	Yes	Host	Yes	
<input type="checkbox"/>	BIGIP-11.5.0.0.0.264-11.5.0.0.264-11.5.0.0.264.iso	11.5.0	Thu Jul 17 16:13:53 2014	1908 MB	Yes	Host	Yes	
<input type="checkbox"/>	BIGIP-11.5.0.0.0.264-11.5.0.0.264-11.5.0.0.264.iso	11.5.0	Tue Jul 1 13:19:37 2014	1903 MB	Yes	Host	Yes	

Figure 12: List of ISO images shared from host

After you perform this task, you can see the images that reside on the vCMP host and are available for installation on the guest.

Installing a host ISO image from within a guest

vCMP® guest administrators perform this task to install an ISO image that resides on the vCMP host. All ISO images that the host administrator has imported into the host's `/shared/images` folder automatically appear on each guest as available for installation.

1. On the Main tab, click **System > Software Management > Image List**.
The Image List screen displays a list of existing image files.
2. In the **Available Images** area of the screen, in the check box column, select an ISO image that shows **Host** in the corresponding Image Source column.
The Install Software Image screen opens.
3. For the **Select Disk** setting, select the disk on which to install the software (for example, MD1 or HD1).

Note: You can install software only on inactive volumes. To install software to the active volume, you must boot to a different volume.

4. For the **Volume set name** setting, select the volume on which to install the software.
5. Click **Install**.

A progress indicator displays as the BIG-IP system installs the software image.

After you perform this task, an ISO image shared by the vCMP host is installed on the guest.

Installing a host ISO image from within a guest using tmsh

vCMP® guest administrators perform this task when using the Traffic Management Shell (`tmsh`) to install an ISO image that resides on the vCMP host. All ISO images that the host administrator has imported into the host's `/shared/images` folder automatically appear on each guest as available for installation.

1. On a vCMP guest, log in to the BIG-IP® system and access `tmsh`.
2. At the `tmsh` prompt, type `install sys software block-device-image image_name volume volume_name` and press Enter.
For example: `install sys software block-device-image BIGIP-11.3.0.2806.0.iso volume HD1.1`

After you perform this task, an ISO image shared by the vCMP host is installed on the guest.

Viewing vCMP Guest Status

About guest status

As a vCMP[®] host administrator, you can log into the vCMP host and view status information about each guest. Using the BIG-IP[®] Configuration utility or the Traffic Management Shell (tmsh), you can view this information in two forms:

- A summary of information for all guests on the vCMP system.
- Detailed information about a specific guest, such as software status, resource provisioning, and high availability (HA) status for specific services running on the guest.


Viewing summary status for all guests

vCMP[®] administrators can view guest summary information while logged into the vCMP host. The vCMP system displays this information on a single screen of the BIG-IP[®] Configuration utility for all guests on the vCMP system. The summary information consists of:

- Guest names.
- The product and version number of the currently-active software volume per guest.
- A list of the specific BIG-IP modules provisioned per guest.
- Per-slot command-line interface prompt status. This status consists of the slot numbers for clustered guests, status color, a slot designation of **P** (primary) or **S** (secondary), and high availability (HA) status.
- HA failure status. This status indicates an HA failure on the guest, and if applicable, a link to the HA Failure screen for the guest.

On the Main tab, click **vCMP > Guest Status**.

For example, the following shows a list of sample guests with summary information.



Name	Product	Version	Build	Active Volume	Resource Provisioning	Prompt Status	HA Failure
guest_1	BIG-IP	11.0.0	0.0.371	HD1.1	LTM	Active	No
guest_2	BIG-IP	11.0.0	0.0.371	HD1.1	LTM	Active	No

Figure 13: List of guests with summary information

Viewing software status for a guest

From the vCMP® host, you perform this task to view information about the software installed on a specific vCMP guest on the system.

1. On the Main tab, click **vCMP > Guest List**.
This displays a list of guests on the system.
2. In the Name column, click the name of the guest that you want to view.
3. On the menu bar, click **Software Status**.
The following shows an example of a guest's installation information.




Figure 14: Guest installation information

Viewing resource provisioning for a guest

From the vCMP® host, you perform this task to view detailed information about current core, memory, and disk allocation for a guest. You can also view a list of the BIG-IP® modules that a vCMP guest administrator has provisioned and the level of provisioning for each module (Dedicated, Nominal, or Minimal).

1. On the Main tab, click **vCMP > Guest List**.
This displays a list of guests on the system.
2. In the Name column, click the name of the vCMP guest for which you want to view status about resource provisioning.
This displays the properties of the guest.
3. On the menu bar, click **Resource Provisioned**.
The following shows an example of a guest's resource provisioning.

Current Guest Resource Allocation	
Number of Cores	2
Memory Use	6.9GB
Disk Use	4.3GB

Module	Provisioning
AFM	None
AM	None
APM	None
ASM	None
AVR	None
FPS	None
GTM	None
LC	None
LTM	Nominal
SWG	None

Figure 15: Resource provisioning information for a guest

Viewing HA failure status

From the vCMP® host, you perform this task to view any high availability (HA) failures pertaining to services running on the guest. For example, you can view whether the `cluster-time-sync` feature within the `CLUSTERED` service has failed. You can also view the specific action that the BIG-IP system took when the failure occurred, such as rebooting the system on the relevant slot.

1. On the Main tab, click **vCMP > Guest List**.
This displays a list of guests on the system.
2. In the Name column, click the name of the guest that you want to view.
3. On the menu bar, click **HA Failure**.
The following shows an example of a guest's HA failure status.

Slot	Feature	Service	Description	Action	Fail
0	Hypervisor Offline	CHMAND	common.ha.dynandvscription	Go Offline	No
0	Crypto Failafe	ON-CRYPTO-0	common.ha.on-origto-0description	No Action	No
0	NIC Failafe	TMM	Traffic Management	Reboot System	No
0	Reboot Request	SDD	Fallover	Reboot System	No
0	Ready For Work	TMM	Traffic Management	No Action	No
0	TMM Detect Failure	TMM	Traffic Management	Fallover	No
0	Forced Offline	SDD	Fallover	No Action	No
0	Software Update	LIND	common.ha.linddescription	Reboot System	No
0	Proc Run	TMM	Traffic Management	No Action	No
0	License Invalid	MCPD	Messaging & Configuration	No Action (HA Feature Disabled)	No
0	Proc Run	MCPD	Messaging & Configuration	No Action	No
0	Proc Run	NAMED	DNS Management	No Action	No
0	Overdog Control	WATCHDOG	common.ha.watchdogdescription	No Action	No
0	Proc Run	BDD	Health Monitors	No Action (HA Feature Disabled)	No
0	Daemon Heartbeat	VILAND	common.ha.vilandedescription	Restart Service (HA Feature Disabled)	No
0	Daemon Heartbeat	WCCPD	Traffic Redirection	Restart Service	No
0	Daemon Heartbeat	SDD	Fallover	Restart All	No
0	Daemon Heartbeat	TMM	Traffic Management	Go Offline & Down Links & restart	No

Figure 16: HA failure status for a guest

Viewing vCMP Statistics

Overview: Viewing vCMP statistics

After creating vCMP® guests to process application traffic, you can display vCMP statistics to better manage performance.

Viewing virtual disk statistics

Using the BIG-IP® Configuration utility, you can view information about the virtual disks that are currently allocated to vCMP® guests:

- The virtual disk names
- The slot number corresponding to each virtual disk image
- The size in gigabytes of each virtual disk
- The name of the guest to which each virtual disk is currently allocated

1. On the Main tab, click **vCMP > Virtual Disk List**.
2. Locate the Virtual Disk List area of the screen.

The following table shows sample statistics for three separate virtual disks.

Virtual Disk Name	Slot ID	Operating System	Status	Disk use
GuestA.img	1	TMOS	Ready	64.4G
GuestB.img	1	Unknown	Unknown	64.4G
GuestC.img	1	TMOS	Ready	64.4G

Viewing vCMP guest information

Before viewing a list of vCMP guests, you must be logged in to the vCMP host.

Using the BIG-IP® Configuration utility, you can list the names of, and information about, the vCMP® guests that are currently on the system.

1. On an external system, open a browser window and access the vCMP host, using the vCMP host's management IP address.
This displays the login window for the BIG-IP Configuration utility.
2. Using your user credentials, log in to the BIG-IP Configuration utility.
3. On the Main tab, click **vCMP > Guest List**.

The system displays a list of vCMP guest names, as well as this information:

- The state configured for each guest
- The number of cores allocated to each guest
- The slot numbers on which each guest is running or slated to run
- The management IP address and netmask for each guest
- The minimum number of slots allocated to each guest

- The slot numbers on which each guest is allowed to run

Viewing current vCMP guest statistics

Before viewing vCMP[®] statistics, you must be logged in to the vCMP host.

You can review current vCMP statistics for all guests on the BIG-IP[®] system. The information shown includes the guest name, bytes, packets, multicast packets, dropped packets, average CPU use, and slot information.

1. On the Main tab, click **VCMP > Statistics**.
The vCMP Guest screen opens and summarizes vCMP activity on the system.
2. You can adjust the display options to change the data format.

Viewing srTCM policier statistics for vCMP guests

Before performing this task, confirm that you have created a single rate three-color marker (srTCM) policier and assigned the policier to a vCMP guest.

You can use this task to view throughput statistics associated with an srTCM policier and its associated guests.

1. On the vCMP host, open the Traffic Management shell (TMSH).
2. At the `tms` command-line prompt, type: `show net rate-shaping sr-tcm-policier`.
This command sequence displays statistics for each srTCM policier on the system and its associated guest.
3. View the results.
The following shows sample srTCM statistics for a policier named `standardSLA` that is associated with guest `myGuest`:

```
-----  
Net::Rate srTCM: standardSLA  
-----  
Settings          CIR  CBS  EBS  
limits           50Mbps 10M  20M  
Statistics  
-----  
VCMP Guest      Green Yellow Red Dropped  
myGuest         53.9K 12.1K 5.1K   5.1K
```

Viewing statistics for physical disk usage

Using the BIG-IP[®] Configuration utility, you can view information about usage of the physical disk on a vCMP[®] system:

- Disk name
- The slot numbers corresponding to the disk name
- The number of virtual disks
- The total vCMP application volume size, in gigabytes
- The available vCMP application volume size, in gigabytes

1. On the Main tab, click **VCMP > Virtual Disk List**.
2. Locate the Disk Usage area of the screen.

The following table shows sample statistics.

Disk	Slot ID	Number of Virtual Disks	Total Volume Size (GB)	Available Volume Size (GB)
HD1	2	1	84	14

Viewing historical statistics about vCMP

To view vCMP® statistics, you must be logged in to the Virtual Clustered Multiprocessing™ (vCMP) host.

You can review detailed historical vCMP statistics in graphical form on the BIG-IP® system. The statistics provide an overview of vCMP performance, network throughput, CPU usage, and disk usage over time.

1. On the Main tab, click **Statistics > Analytics > vCMP**.
The vCMP Overview screen opens and summarizes vCMP activity on the system.
2. You can change the time period for which to examine statistics; adjust the time for each widget or for all widgets (using the override time range).
3. If you want to add new information to the Overview screen, click **Add Widget**.
The Add New Widget popup screen opens.
4. Specify the page, information, range, the details, and measurements to display, and click **Done**.
A new widget with your specifications is added to the vCMP Overview.
5. From the menu bar, select the type of vCMP statistics you want to view.

Select this option

To see these vCMP statistics

Overview

Top statistical information about vCMP traffic on your system, such as the top vCMP guests by average CPU usage. You can customize the information that is displayed by adding widgets that show the information you want from the other screens.

Network

Average throughput or bytes in or out per vCMP guest, interface, or chassis slot.

CPU Usage

Average CPU usage per vCMP guest or chassis slot.

Disk Usage

Average bytes or requests read or written per vCMP guest or chassis slot.

6. From the **View By** list, select the item for which to display statistics.

*Tip: You can also click **Expand Advanced Filters** to filter the information that displays.*

7. You can select a different time for which to view the statistics, and you can also customize the **Time Period** by marking the appropriate zone one line chart using the mouse (hold and draw to select the required period).
8. To focus in on the specific details you want more information about, click the chart, an item in the details list, or the pie chart on the right (for some entities).
For example, if you are displaying information about vCMP Guests, you can click one of the guests to display a chart that shows details about that guest.
As you drill down into the statistics, you can locate more details and view information about a specific item on the charts.
9. If you want to export the information in any of the charts, click **Export** and specify your options for how and where to send the data.

To send reports by email, the system requires an SMTP configuration.

Viewing vCMP Statistics

The statistics show an overview of vCMP performance: network throughput, CPU usage, and disk usage. The data can be displayed per guest, interface, or chassis slot depending on the selected statistics page. Review the vCMP statistics to understand how the guests and chassis are using resources on the system. As a result, you become more familiar with the system and its resource utilization, and you can troubleshoot the system as needed.

Sample vCMP Statistics reports

This figure shows a sample vCMP® Statistics report showing a system on which there are two guests. The chart shows the average CPU usage for the guests over the past day.




Figure 17: Sample vCMP Overview

You can view other statistics, such as network statistics, by clicking items on the menu bar. This figure shows network statistics for vCMP guests during the last hour, but you can also view statistics by vCMP interfaces or chassis slots. You can also change the time frame for which to view the statistics.




Figure 18: Sample vCMP Network statistics

By clicking `guest_1` in the table below the chart, you can drill down to see what is happening for that guest. For example, here you can see the throughput for each of the interfaces on `guest_1`.




Figure 19: Sample vCMP Network statistics after drill down

You can further drill down by clicking an interface to see additional details, or view CPU or disk usage by clicking the menu bar.

Understanding Clusters

Overview: Managing a vCMP cluster

One of the tasks that a guest administrator performs is managing the cluster for a guest.

Viewing cluster properties

A guest administrator can use this task to view the properties of the guest's cluster.

1. Open a browser, and in the URL field, specify the management IP address that the host administrator assigned to the guest.
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	Specifies 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.
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

A guest administrator can use this task to view the properties of the guest's cluster members.

Important: When logged into the guest, never change the vCMP management IP address. Doing so produces unexpected results. You can, however, change a cluster member IP address.

1. Open a browser, and in the URL field, specify the management IP address that the host administrator assigned to the guest.
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 a blade goes out of service, a guest administrator can mark that cluster member disabled. When the blade has been returned to service, the guest administrator must enable the blade again.

1. Use a browser and the cluster management IP address to log in to the system 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**.

Best Practices

vCMP best practices

F5 Networks makes the following recommendations for managing a vCMP® system.

Category	Recommendation
vCMP® disk management	Ensure that you allocate enough disk space for other installation slots for the vCMP host before you provision the vCMP feature.
Network setup	Before setting up a vCMP system, verify that each slot's management interface is physically wired to an external bridge.
Change of vCMP management IP address	You should only change the vCMP management IP address when logged into the vCMP host. Changing this address when logged into a guest could produce unexpected results. Note that changing individual cluster member IP addresses when logged into a guest is fully supported.
Slot assignment to guests	Whenever possible, configure a guest to allow the guest to run on more slots than are actually populated with blades. The result is an automatic expansion of the guest cluster when you insert an additional blade.
Virtual disk management	To prevent unattached virtual disks from consuming disk space over time, consider deleting unwanted virtual disks from the system. Otherwise, previously provisioned virtual disks remain on disk after their associated vCMP guest configurations have been deleted.
Protection from performance degradation	To protect a guest from performance degradation if a blade failure occurs, configure high availability if possible. You do this by setting up device service clustering (DSC®). For a standalone vCMP system, consider deploying guests with sufficient cores and slots to ensure that a single blade failure does not result in unacceptable service degradation.
Adjusting the MTU size for a guest	When a guest is subscribing to a host-based VLAN, you should only adjust the VLAN's MTU size when you are logged into the guest. Changing the MTU size when logged into the host has no effect on the guest's ability to process traffic or manage routing.

Calculation for Maximum Core Allocation

Calculation for determining maximum core allocation

When you are creating a vCMP® guest and allocating cores to that guest, the BIG-IP Configuration utility assists you by displaying only valid amounts of cores in the **Cores per Slot** setting. For example, for a chassis with B2100 blades, the BIG-IP Configuration utility displays **Cores per Slot** values of **2**, **4**, and **8**, because these are the only valid choices for that blade platform. Some users, however, might want more detailed information about these selections to enhance their own understanding of core allocation on the vCMP system.

The total number of cores that you can allocate to all vCMP® guests (combined) on a blade depends on the number of physical cores that a single physical processor contains on a particular blade platform. For example, on a blade platform with hyper-threading support, each physical core represents two logical cores. Consequently, a blade platform with two physical processors, each with six physical cores (that is, 12 cores), has a total of 24 logical cores that the host can allocate for that slot. This illustration shows an example of the relationship of physical processors to logical cores.




Figure 20: Relationship of physical processors to logical cores

In addition to the total number of logical cores available for allocation on that slot, there is a maximum number of logical cores that the host can allocate to an individual guest on that slot. This number is restricted to the number of physical cores per physical processor, which means that you cannot allocate additional logical cores to a guest VM from any other processor on the blade. Therefore, if you know the number of physical cores per physical processor on your blade platform, you can use this simple calculation to understand the maximum number of logical cores that you can allocate to a guest on a slot:

$$\text{Number of physical cores per physical processor} * \text{Number of cores per physical core} = \text{Maximum number of logical cores per guest}$$

For example, if a blade platform has six physical cores per physical processor, and the number of cores per physical core is 2, then the maximum number of logical cores per guest on that slot is 12 ($6 * 2 = 12$).

Additional Tasks for Isolated Guests in Appliance Mode

Additional tasks for isolated guests in Appliance mode

To ensure that guest administrators can access an isolated guest and manage the BIG-IP[®] software within the guest, you must create the isolated guest with Appliance mode disabled, perform some additional tasks, and then modify the guest to enable Appliance mode. These additional tasks are:

- Creating a self IP address for guest administrators to use to access the guest, and granting `tmsh` access to the guest's `admin` user account.
- Enabling Appliance mode on the guest.

After performing these tasks, administrators for an isolated guest are restricted to using either the BIG-IP[®] Configuration utility or `tmsh` to manage BIG-IP modules within the guest (when port lockdown settings on the self IP address allow such traffic).

Preparing an isolated guest for Appliance mode

You use this task to prepare an isolated guest to operate in Appliance mode. Specifically, you use this task to:

- Grant access to the Traffic Management Shell (`tmsh`) for the `admin` user account within a vCMP[®] guest. Because the `admin` user for an isolated guest in Appliance mode is restricted to using `tmsh`, you must first grant the `admin` account permission to use `tmsh`. By default, the `admin` account for a guest has no access to `tmsh`.
- Create a self IP address for guest administrators to use to access the guest. This is necessary because an isolated guest is not connected to the management network and therefore has no management IP address assigned to it.

You perform this task by accessing the guest from the vCMP[®] host.

1. From the vCMP host, access the Bash shell by typing `vconsole guest_name`.
For example, you can type `vconsole guest_A`
The system prompts you to enter a user name and password.
2. Type the `root` account and the password `default`.
The system logs you into the guest and displays the guest's system prompt.
3. Type the command `tmsh modify auth user admin shell tmsh`.
This command grants `tmsh` access to the `admin` user account.
4. Type the command `tmsh create net self address ip_address/netmask vlan vlan_name allow-service default`.
This creates the specified IP address on the guest and makes required adjustments to the port lockdown settings.
5. At the prompt, exit the guest by typing `exit`.
6. At the Bash prompt, log out of the Linux system by typing `exit`, if necessary.
7. Exit the vConsole utility by typing the key sequence `ctrl-]`.
This displays the prompt `telnet>`.
8. Type `q`.

Enabling Appliance mode on an isolated guest

You use this task to enable Appliance mode on an existing guest that is isolated from the management network.

***Note:** You can perform this task while the guest is in the Deployed or Provisioned state; there is no need to set the guest state to Configured prior to performing this task.*

1. Use a browser to log in to the vCMP® host, using the primary cluster management IP address.
2. On the Main tab, click **vCMP > Guest List**.
This displays a list of guests on the system.
3. In the Name column, click the name of the guest that you want to modify.
This displays the configured properties of the guest.
4. For the **Appliance Mode** setting, select the check box.
When you enable **Appliance Mode** for an isolated guest, the system enhances security by denying access to the `root` account and the `Bash` shell for all guest administrators.
5. Click **Update**.

The guest is now running in Appliance mode. All guest administrators are restricted to using the BIG-IP® Configuration utility and `tmssh` to manage the guest.

Deploying Route Domains within a vCMP Guest

Overview: Deploying Route Domains within a vCMP Guest

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

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

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

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




Figure 21: Route domains within a guest

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

active traffic group on one instance of the guest and a standby traffic group on the other instance of the guest.

Prerequisite configuration tasks

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

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

About VLAN and BIG-IP address configuration

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

Illustration of VLAN and BIG-IP address configuration

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




Figure 22: VLANs and BIG-IP addresses in a vCMP route domain configuration

In this illustration:

Blue text

Objects created by host administrator.

Black text

Objects created by guest administrator.

Brown text

Objects created by customer administrator.

Tasks for the host administrator

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

On the host, for our sample configuration with three customers, you create a separate set of uniquely-tagged internal and external VLANs for each customer. You will therefore create at least six VLANs on the host (two per customer) that, when combined with the three existing VLANs, bring the total number of VLANs on the host to nine. At this point, all VLANs reside in partition `Common`. Then you assign all nine host-based VLANs to the guest. This allows the guest to use those VLANs to process customer traffic.

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

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

Creating customer VLANs on the vCMP host

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

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

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

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

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

5. If you want to use Q-in-Q (double) tagging, use the **Customer Tag** setting to perform the following two steps. If you do not see the **Customer Tag** setting, your hardware platform does not support Q-in-Q tagging and you can skip this step.
 - a) From the **Customer Tag** list, select **Specify**.
 - b) Type a numeric tag, from 1-4094, for the VLAN.
The customer tag specifies the inner tag of any frame passing through the VLAN.
6. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) If you specified a numeric value for the **Customer Tag** setting and from the **Tagging** list you selected **Tagged**, then from the **Tag Mode** list, select a value.
 - d) Click **Add**.
7. Click **Repeat** and repeat these steps to create additional VLANs.

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

Assigning VLANs to the vCMP guest

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

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

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

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

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

Tasks for the guest administrator

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

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

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

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

Creating an administrative partition for each customer

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

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

1. On the Main tab, expand **System** and click **Users**.
The Users List screen opens.
2. On the menu bar, click **Partition List**.
3. Click **Create**.
The New Partition screen opens.
4. In the **Partition Name** field, type a unique name for the partition.

An example of a partition name is `CustomerA_partition`.

5. Type a description of the partition in the **Description** field.
This field is optional.
6. For the **Device Group** setting, ensure that the Sync-Failover device group containing this vCMP guest is selected.
7. For the **Traffic Group** setting, retain the default value, which is the floating traffic group `traffic-group-1`.

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

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

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

About moving host-based VLANs to a customer partition

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

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

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

Deleting VLANs in partition `Common` from within the guest

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

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

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

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and ensure that partition `Common` is selected.
3. In the Name column, locate the relevant VLAN name.
4. In the Tag column, note the numeric ID.
You will specify this ID when you re-create this VLAN in a customer partition.
An example of a VLAN ID in the Tag column is **4094**.
5. If the VLAN has a customer tag (optional), then in the Customer Tag column, note the numeric ID.
You will specify this ID when you re-create this VLAN in a customer partition.
6. To the left of the VLAN name, select the check box and click **Delete**.
The system prompts you to confirm the delete action.

7. Click **Delete**.

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

Re-creating VLANs in each administrative partition

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

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

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

The VLAN List screen opens.

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

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

An example of a selected partition is `CustomerA_partition`.

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

3. Click **Create**.

The New VLAN screen opens.

4. Type a name for the VLAN.

You can specify the same name as the VLAN that you deleted from partition `Common` or you can type a unique name.

5. For the **Tag** field and the optional **Customer Tag** field, type the same ID that was previously assigned to the VLAN that you deleted from partition `Common`.

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

6. Retain the values for all other settings as configured.

7. Click **Finished**.

This prompts you with the question: The VLAN has no interface, do you want to continue?

8. Click **OK**.

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

Creating a route domain for each administrative partition

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

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

1. On the Main tab, click **Network > Route Domains**.

The Route Domain List screen opens.

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

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

After choosing this value, you are not required to append the route domain ID to any self IP or virtual IP address that you create later for this route domain. Instead, the BIG-IP system automatically associates an IP address with the default route domain in the partition, as long as you set this partition to be the current partition when you create the address.
 12. Click **Finished**.
The system displays a list of route domains on the BIG-IP system, including the new route domain.
 13. Repeat the process of creating a route domain for another customer for which you want to segment traffic, associating the relevant VLANs in the process.
- After you perform this task repeatedly, you should have three separate route domains with unique route domain IDs, and each route domain should be associated with unique internal and external VLANs that pertain to a specific customer. Also, each route domain should be designated as the default route domain for its associated administrative partition.

Creating an empty traffic group for each customer

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

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

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

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

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

You now have floating traffic groups with no members.

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

Assigning a traffic group to each administrative partition

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

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

1. On the Main tab, expand **System** and click **Users**.
The Users List screen opens.
2. On the menu bar, click **Partition List**.
3. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and ensure that partition `Common` is selected.
4. In the Name column, click a customer partition name.

5. For the **Traffic Group** setting, clear the check box labeled **Inherit traffic group from root folder** and from the list, select the name of a traffic group.
6. Click **Update**.
7. Repeat these steps to assign a traffic group to each of the other customer partitions.

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

Tasks for each customer administrator

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

Creating floating self IP addresses

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

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

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

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

1. On the Main tab, click **Network > Self IPs**.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and select the customer-specific administrative partition.
If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.
An example of a selected partition is `CustomerA_partition`.
Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.
3. Click **Create**.
The New Self IP screen opens.
4. In the **IP Address** field, type an IP address.
This IP address should represent the address space of a specific VLAN. Because the route domain for the VLAN that you will associate with this self IP address is the default route domain for the current administrative partition, you are not required to append the relevant route domain ID to this IP address.
The system accepts IP addresses in both the IPv4 and IPv6 formats.
5. In the **Netmask** field, type the full network mask for the specified IP address.
6. From the **VLAN/Tunnel** list, select the VLANs that you want to associate with this self IP address.
The VLANs you select are those that you moved from partition `Common` to the current administrative partition.

7. From the **Port Lockdown** list, select a value.
8. From the **Traffic Group** list, select the floating traffic group for which you want this self IP address to be a member.
Selecting a floating traffic group automatically causes the self IP address to be a floating address.
For example, you can select a traffic group named `tg-CustomerA`.
9. Click **Finished**.
The screen refreshes, and displays the new self IP address.
10. Repeat this task for each floating self IP address that you need to create.

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

Creating a pool

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

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

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and select the customer-specific administrative partition.
If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.
An example of a selected partition is `CustomerA_partition`.
Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.
3. Click **Create**.
The New Pool screen opens.
4. In the **Name** field, type a unique name for the pool.
5. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.

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

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

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

Creating a virtual server

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

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

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. In the upper-right corner of any the BIG-IP Configuration utility screen, locate the **Partition** list and select the customer-specific administrative partition.
If the partition selections are unavailable, you do not have a user role that allows you to change the current partition.
An example of a selected partition is `CustomerA_partition`.
Whenever you select a partition name from the list, the current administrative partition changes to the selected partition.
3. Click the **Create** button.
The New Virtual Server screen opens.
4. In the **Name** field, type a unique name for the virtual server.
5. In the **Destination Address** field, type the IP address in CIDR format.
The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `10.0.0.1` or `10.0.0.0/24`, and an IPv6 address/prefix is `fe1::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a `/32` prefix.

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

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

Modifying a virtual IP address

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

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

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

5. Click **Update**.
6. Repeat these steps for each customer's virtual address.

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

Implementation results

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

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

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

Legal Notices

Legal notices

Publication Date

This document was published on August 16, 2018.

Publication Number

MAN-0376-12

Copyright

Copyright © 2018, 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

- access control
 - administrative *10*
- address configuration for VLAN and BIG-IP
 - illustration *78*
- admin user account
 - modifying *75*
- administrative partitions
 - creating *81*
- administrator tasks *10, 41, 45*
- Appliance mode
 - additional tasks *75*
 - and boot locations *12*
 - and user access restrictions *12*
 - enabling on existing guest *75*
- Appliance mode types
 - described *12*
- application traffic
 - isolating on network *81*
- audience *10*

B

- BIG-IP instances *7*
- BIG-IP modules
 - and guest states *30*
 - and memory allocation *25*
 - and resource provisioning *60*
 - provisioning *9, 10*
 - provisioning within guests *45*
- BIG-IP software
 - versions supported *9*
- BIG-IP version requirements *12*
- blade insertion
 - and slot allocation *29*
- blade platforms
 - physical vs. logical cores *73*
- blade reduction
 - effects of *25*
- blade removal
 - effect on resource allocation *26*
 - effects of *27*
- bridged guests
 - described *11*
- bridged network
 - setting *42*
- bridges
 - and management interfaces *71*

C

- cluster availability
 - and vCMP guests *39*
- cluster management IP addresses
 - configuring *41*
- cluster member IP addresses
 - specifying *46*

- cluster member properties
 - described *70*
 - viewing *69*
- cluster members
 - enabling and disabling *70*
- cluster properties
 - described *69*
 - viewing *69*
- clusters *69*
- components *8*
- compression resource allocation
 - per guest *30*
- config sync
 - for vCMP systems *37*
- config sync IP addresses *38*
- configuration results *47*
- Configured state
 - and disk attachment *54*
 - described *30*
- connection mirroring
 - on vCMP systems *40*
- connections
 - and blade removal *27*
 - and memory use *25*
- control plane *10*
- core allocation
 - about *23*
 - based on blade model *23*
 - configuring *42*
 - determining *24*
 - on solid-state platforms *26*
- core availability
 - determining *49*
 - increasing *49*
- cores
 - as system resource *23*
 - defined *8*
 - increasing *49*
- cores per platform
 - calculating *73*
- CPU allocation
 - based on blade model *23*
- CPU cores
 - and guest states *30*
 - increasing *49*
- CPU resources
 - allocating *23*
 - increasing *49*
- custom resource allocation
 - defined *23*
- customer administrator tasks
 - for deploying route domains within a vCMP guest *86*

D

- daemon failures
 - on vCMP guests *61*
- data plane

Index

- data plane (*continued*)
 - vs. management network 10
- Deployed guest state
 - purpose of 44
- Deployed state
 - described 30
 - next steps 47
- device groups
 - example of 34
 - for vCMP systems 37
- device trust IP addresses 38
- disk creation time
 - minimizing 52
- disk space
 - and vCMP application volume 41
 - consuming 71
 - insufficient 42
 - reserving 7, 9, 71
- disk space allocation
 - about 51
- disk usage 64
- dual-slot guests
 - example of 33

E

- Ethernet interface
 - of host 11

F

- failover
 - for vCMP systems 37, 40
 - on vCMP systems 39
- failover IP addresses 38
- failover methods
 - for vCMP systems 39
- failover objects
 - associating with traffic groups 85
- failsafe
 - for vCMP systems 40
- flexible resource allocation
 - defined 23
- floating IP addresses
 - configuring 41
- Force to Standby option 85

G

- guest access
 - with vconsole utility 11
- guest administrator tasks 45
- guest administrators
 - about 10
 - duties of 10
- guest failover 40
- guest interfaces
 - bridging to physical interface 11
- guest IP addresses
 - configuring 42
- guest resource allocation
 - determining 24

- guest software
 - viewing 60
- guest states
 - and virtual disk migration 55
 - configuring 42
 - described 30
- guest statistics
 - viewing for vCMP 64
- guest status
 - about viewing from host 59
 - and resource provisioning 60
 - viewing summary of 59
- guest throughput limitations 29
- guest traffic load
 - viewing 47
- guest-related tasks 45
- guest-wide administrator tasks
 - for deploying route domains within a vCMP guest 81
- guests
 - about 7
 - additional tasks 75
 - and licensing 9
 - and resource requirements 23
 - and SSL resources 30
 - and virtual disks 54
 - configuring BIG-IP modules on 47
 - creating 42
 - provisioning BIG-IP modules for 45
 - setting to Deployed state 44

H

- HA failure
 - viewing status 61
- HA groups
 - for vCMP systems 39
- high availability
 - about 10
 - for vCMP systems 37
- host
 - about 7
 - accessing 41
 - and licensing 9
- host administrator tasks
 - for deploying route domains within a vCMP guest 79
- host administrators
 - about 10
- hotfixes
 - installing to guest 57, 58
- hypervisor
 - accessing 41

I

- illustration
 - of VLAN and BIG-IP address configuration 78
- implementation results 89
- instances, BIG-IP 7
- interfaces
 - tagging 46
- IP addresses
 - for DSC 38

- ISO images
 - and guest states 30
 - and virtual disk templates 53
 - installing 42
 - installing to guest 58
 - sharing with guests 57
 - viewing from guest 57

- isolated guests
 - accessing 45
 - additional tasks 75
 - and Appliance mode 75
 - described 11
- isolated network
 - setting 42

L

- Layer 2/Layer 3 configuration 15
- licensing
 - and Appliance mode 12
 - and guests 9

M

- management interfaces
 - on guests 10
 - wiring 71
- management IP addresses
 - configuring 41
 - for guests 21, 46
- management network
 - and bridged guests 11
 - and connection to guests 11
 - and isolated guests 11
 - vs. data plane network 10
- management network mode
 - setting 42
- memory allocation
 - about 23
 - about calculating 25
 - and blade removal 27
 - based on blade model 23
 - configuring 42
 - determining 24
- memory use
 - and connections 25
- mirroring
 - on vCMP systems 40
- mirroring IP addresses 38
- module configuration 47
- MTU setting
 - for VLANs 20
- multi-tenancy 7
- multiple-slot guests
 - example of 34

N

- network configuration
 - host vs. guest 15
- network isolation 10
- network state

- network state (*continued*)
 - changing 11
- network throughput statistics
 - for vCMP guests 64
- network traffic
 - about segmenting 77

P

- partitions, *See* administrative partitions
- performance degradation
 - preventing 71
- pool availability
 - and vCMP guests 39
- pools
 - creating 87
 - for BIG-IP modules 45, 47
- prerequisite tasks
 - for deploying route domains within a vCMP guest 78
- Provisioned state
 - described 30
- provisioning process 9

R

- rate shaping statistics
 - for vCMP guests 64
- redundancy
 - for vCMP systems 37
- reserve space
 - increasing 55
- resource allocation
 - and guest states 30
 - based on blade model 23
 - defined 23
 - determining 24
 - forms of 23
 - on solid-state platforms 26
- resource provisioning
 - viewing for guests 60
- resource requirements
 - understanding 23
- resources
 - allocating 9
- route domains
 - creating 83
 - described 77

S

- self IP address configuration 15
- self IP addresses
 - and VLANs 46, 78
 - creating 46
 - creating for default route domains 86
 - for BIG-IP modules 45
- single slot guests
 - example of 33
- single-core guests
 - and solid-state platforms 26
- slot assignment
 - about changing 27

- slot assignment (*continued*)
 - and scalability 26
 - best practice for 71
 - for guests 25
- slots
 - and virtual disk migration 55
 - assigning to guests 42
- software
 - installing for guests 42
- software images
 - sharing with guests 57
- software status
 - viewing 60
- software versions 9
- solid-state drives
 - and core allocation 26
- srTCM statistics
 - viewing 64
- SSL hardware cards
 - use of 30
- SSL modes
 - about 30
- statistics
 - and disk usage 64
 - sample vCMP reports 66
 - viewing for guests 63
 - viewing for vCMP 63
 - viewing for virtual disks 63
 - viewing historical charts 65
- status
 - of guests 59
 - viewing 59
- switchboard failsafe
 - for vCMP systems 40
- SYnc-Failover device groups
 - example of 34
- system administrator tasks 41, 45
- system components 8
- system resources
 - allocating 9, 23

T

- templates
 - viewing 52
- tenants, *See* guests
- throughput limitations 29
- TMOS software
 - installing 42
- tmsh access
 - granting 75
- traffic groups
 - activating 85
 - and failover objects 85
 - assigning to each administrative partition 85
 - forcing to standby state 85
- traffic load
 - viewing for guest 47
- trunk availability
 - and vCMP guests 39
- trunk configuration 15
- trunks

- trunks (*continued*)
 - about 10

U

- updates
 - installing to guest 57, 58
- user access restrictions 12
- user account permissions 11

V

- vCMP
 - sample statistics reports 66
 - viewing current statistics 64
 - viewing historical statistics 65
- vCMP application volume
 - and disk space 41
 - creating and deleting 55
- vCMP configuration results 47
- vCMP feature
 - provisioning 9, 10
- vCMP guests
 - about using route domains in 77
 - and SSL resources 30
 - and switchboard failsafe 40
 - See also* guests
- vCMP host
 - accessing 41
 - creating VLANs on 80
- vCMP systems
 - described 7
- vconsole utility 11
- vCPU cores
 - increasing 49
- vCPUs, *See* cores
- version requirements 12
- versions, software 9
- virtual addresses
 - assigning to traffic group 88
- virtual disk creation time
 - minimizing 52
- virtual disk space allocation
 - about 51
- virtual disk statistics
 - viewing 63
- virtual disk templates
 - about 51
 - enabling and disabling 53
 - viewing 52
- virtual disks
 - about 51
 - and disk space consumption 71
 - and guest states 30
 - as system resource 23
 - attaching 54
 - creating 42
 - defined 8
 - deleting 55
 - detaching and re-attaching 54
 - effect on disk space 42
 - file names and location of 51

- virtual disks (*continued*)
 - viewing unattached 54
- virtual interfaces
 - bridging to physical interface 11
- virtual machines
 - defined 8
- virtual servers
 - creating 87
 - for BIG-IP modules 45, 47
- VLAN
 - adding tagged interface 19
- VLAN and BIG-IP address configuration
 - illustration 78
- VLAN configuration
 - and vCMP host 15
- VLAN MTU setting
 - host vs. guest 20
- VLANs
 - about 10
 - and self IP addresses 46, 78
 - assigning to guests 81
 - configuring guest use of 42
 - creating 46, 80
 - moving from partition Common 82
 - tagged interfaces for 46
- VMs
 - defined 8
 - propagating changes to 30

