

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

Version 11.5



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
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.....	13
 Additional Network Considerations.....	 15
Guest access to the management network.....	15
Network separation of Layer 2 and Layer 3 objects.....	15
Management IP addresses for bridged guests.....	16
About the VLAN publishing strategy.....	16
Overview of VLAN subscription.....	16
About VLANs with identical tags and different names.....	17
About VLANs with identical names and different tags.....	18
Solution for tag discrepancy between host and guest VLANs.....	19
Interface assignment for vCMP 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
Formula for host memory allocation to a guest.....	24
About slot assignment for a guest.....	25
About single-core guests.....	25
Resource allocation results.....	26
Scalability considerations.....	27
About initial slot assignment.....	27
About changing slot assignments.....	28

Effect of blade removal on a guest.....	28
Effect of blade re-insertion on a guest.....	29
About SSL and compression hardware.....	29
Guest states and resource allocation.....	29
Deployment Examples.....	31
A single-slot LTM guest on a standalone system.....	31
Dual-slot LTM guest within a device group.....	31
Multiple guests on multiple slots in device groups.....	32
Device Service Clustering for vCMP Systems.....	35
Overview: Device service clustering for vCMP systems.....	35
Required IP addresses for DSC configuration.....	36
Failover methods for vCMP guests.....	36
About HA groups for vCMP systems.....	37
About connection mirroring for vCMP systems.....	37
Initial vCMP Configuration Tasks.....	39
Overview: vCMP application volume management.....	39
Viewing disk space allocation for a vCMP application volume.....	39
Modifying disk space allocation for a vCMP application volume.....	39
vCMP host administrator tasks.....	40
Accessing the vCMP host.....	40
Provisioning the vCMP feature.....	41
Creating a vCMP guest.....	41
Setting a vCMP guest to the Deployed state.....	43
vCMP guest administrator tasks.....	44
Provisioning BIG-IP modules within a guest.....	44
Specifying cluster member IP addresses for a guest.....	45
Creating a self IP address for application traffic.....	45
Next steps.....	46
Configuration results.....	46
vCMP Core Upgrade for Deployed Guests.....	47
About increasing CPU cores for a deployed guest.....	47
Increasing CPU cores for a deployed guest.....	47
Managing vCMP Virtual Disks.....	49
Overview: Managing virtual disks.....	49
About virtual disk allocation.....	49
About virtual disk images.....	49
About virtual disk detachment and re-attachment.....	49
About virtual disk migration.....	51

Deleting a virtual disk from the BIG-IP system.....	51
Understanding Clusters.....	53
Overview: Managing a vCMP cluster.....	53
Viewing cluster properties.....	53
Cluster properties.....	53
Viewing cluster member properties.....	54
Cluster member properties.....	54
Enabling and disabling cluster members.....	54
Changing a cluster-related management IP address.....	55
Cluster-related IP addresses.....	55
Viewing vCMP Statistics.....	57
Overview: Viewing vCMP statistics.....	57
Viewing virtual disk statistics.....	57
Viewing vCMP guest statistics with the BIG-IP Configuration utility.....	57
Viewing statistics for physical disk usage	58
Best Practices.....	59
vCMP best practices.....	59
Calculation for Maximum Core Allocation.....	61
Calculation for determining maximum core allocation.....	61
Additional Tasks for Isolated Guests in Appliance Mode.....	63
Additional tasks for isolated guests in Appliance mode.....	63
Preparing an isolated guest for Appliance mode.....	63
Enabling Appliance mode on an isolated guest.....	64
Legal Notices.....	65
Legal Notices.....	65
Acknowledgments.....	66

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.

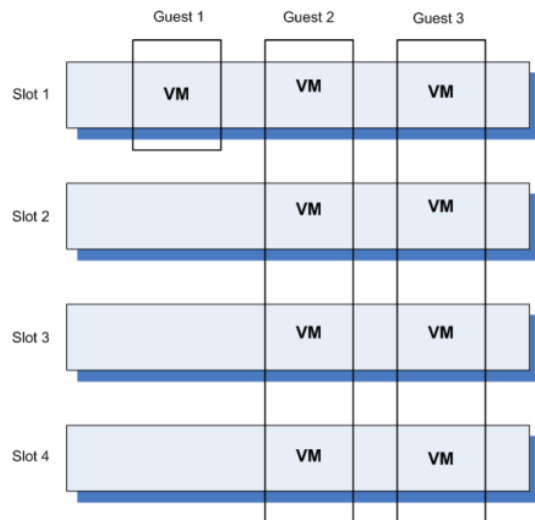


Figure 2: Guest VMs as cluster members

Virtual disk

A *virtual disk* is the portion of disk space on a slot that the system allocates to a guest VM. 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 any combination of BIG-IP® 11.x software. For example, in a three-guest configuration, the host can run version 11.2.1, while guests run 11.2, 11.3, and 11.4. With this type of version support, you can run multiple versions of the BIG-IP software simultaneously 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. To see the relevant matrix showing these version combinations, see 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.

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.

Note: The reserved disk space protects against any possible resizing of the file system.

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.

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, or to be isolated from the management network but remain accessible by way of the host-only interface.

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`) instead of and 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

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.

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, or to be isolated from the management network but remain accessible by way of the host-only interface.

Network separation of Layer 2 and Layer 3 objects

On a vCMP system, you must 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. Typically, to ensure that each guest's data plane traffic is securely isolated from other guests, the host administrator 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:

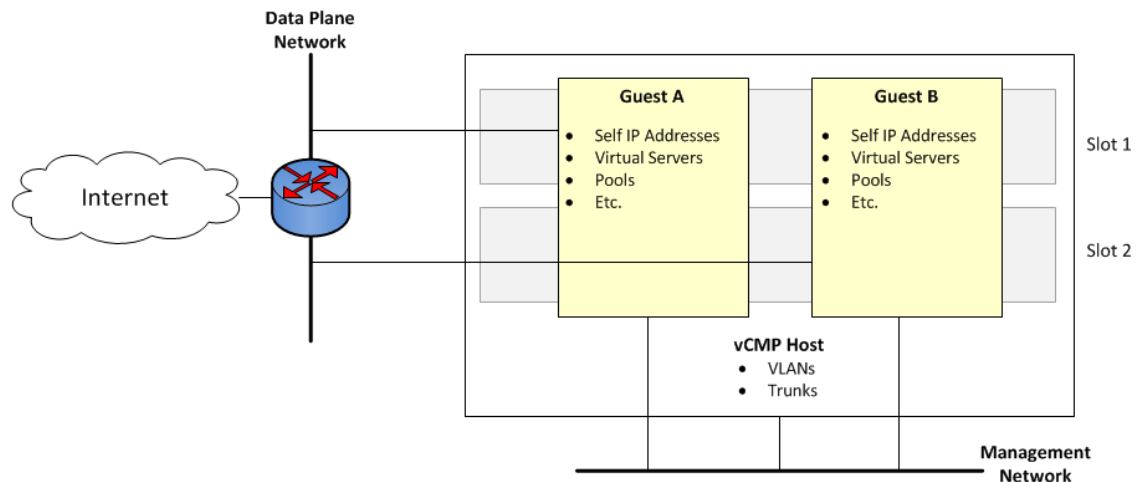


Figure 4: Isolation of network objects on the vCMP system

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.*

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.

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 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
}
```

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.

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

You use this task when you want to delete a VLAN from within a vCMP® guest. One reason for deleting a VLAN from within a guest is to help resolve a tag discrepancy between a guest VLAN and a host VLAN.

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 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 on a particular guest. One reason for deleting a VLAN from within a guest is to help resolve a tag discrepancy between a guest VLAN and a host VLAN.

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 vCMP 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

You 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 vCMP 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

You 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 which the host published to the guest has appeared on the guest, with the correct tag.

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.

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 CPU and a fixed amount of memory. You should therefore 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.

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, 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, each guest in the device group must run on the same blade and chassis models.

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 never allocates any other odd number of cores per slot for a guest (such as three, five, or seven cores).

The illustration shows a possible configuration where the host has allocated a single core to one of the guests.




Figure 5: A vCMP configuration with a single-core guest

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® Global Traffic Manager™ (GTM™) only
- BIG-IP® Global Traffic Manager™ (GTM™) standalone only

Resource allocation results

This illustration shows an example of a system with three guests that the vCMP® host has distributed across slots in varying ways. The way that the host distributes the guests depends on the way you configured guest properties when you initially created each guest.




Figure 6: Three guests with varying amounts of core allocation and slot assignment

This illustration shows three guests configured on the system. For blade model B2100, which provides approximately 2 GB of memory per core:

- The blue guest is configured with two cores per slot, providing a total memory allocation of four GB per slot for the guest (2 cores x 2 GB). The guest spans four slots.
- The red guest is configured with two cores per slot, providing a total memory allocation of four GB per slot for the guest (2 cores x 2 GB). The guest spans two slots.
- The green guest is configured with four cores per slot, providing a total memory allocation of eight GB per slot for the guest (4 cores x 2 GB). The guest spans two slots.

Important: For an individual guest, the number of cores allocated to one (not all) of the guest's slots determines the guest's total memory allocation.

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.

Fortunately, there is no reduction in memory allocation, given that when you create a guest, you instruct the host to allocate the full amount of required memory for that guest to every slot in the guest's cluster (through the guest's **Cores per Slot** property). However, each connection causes some amount of memory use, which means that when a blade becomes unavailable and the host redistributes its connections to other blades, the percentage of memory use on these remaining blades increases. In some cases, the increased memory use could exceed the amount of memory allocated to each of those slots.

For example, if a guest spans three slots which process 1,000,000 connections combined, each slot is processing a third of the connections to the guest. If one of the blades becomes unavailable, reducing the guest's cluster to two slots, then the two remaining blades must each process half of the guest's connections (500,000), resulting in a memory use per slot that could be higher than what is allocated for that slot. Assigning as many slots as possible to each guest reduces this risk.

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.

About SSL and compression hardware

On systems that include SSL and compression hardware processors, the vCMP® feature shares these hardware resources among all guests on the system, in a round robin fashion.

When sharing SSL hardware, if all guests are using similar-sized keys, then each guest receives an equal share of the SSL resource. Also, if any guests are not using SSL keys, then other guests can take advantage of the extra SSL resource.

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

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.

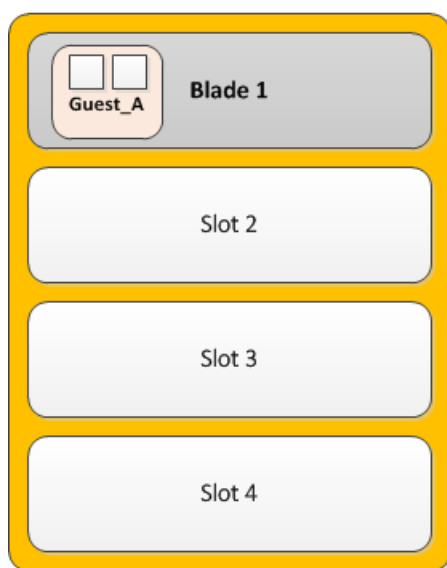


Figure 7: Single slot guest on a standalone VIPRION system

Dual-slot LTM guest 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 BIG-IP Local Traffic Manager (LTM) module than a configuration where the BIG-IP LTM module 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 guests in a device group

Multiple guests on multiple slots in 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 BIG-IP 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 three guests on each chassis. Each guest and its equivalent guest on the other chassis form a separate Sync-Failover device group.

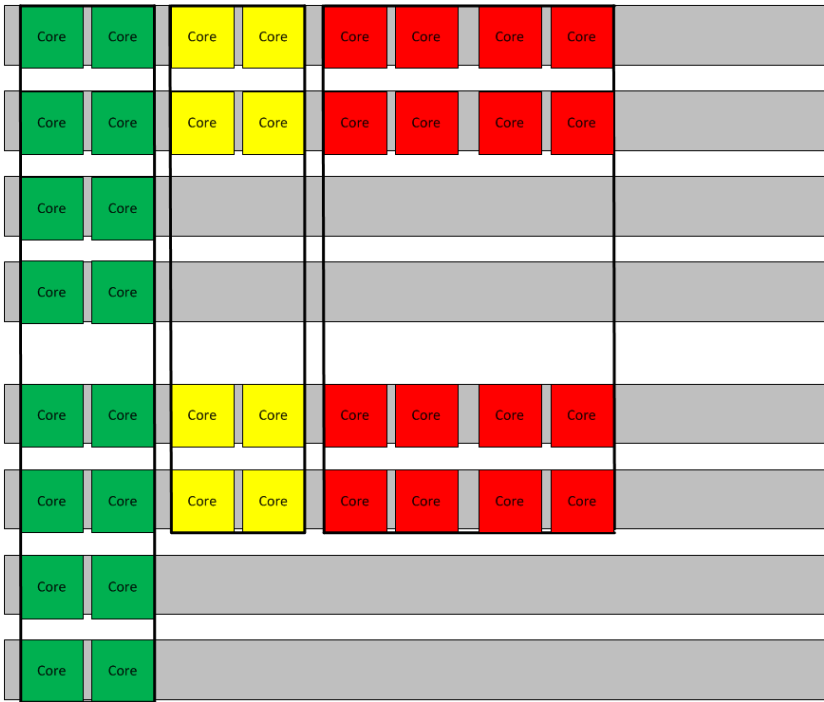


Figure 10: vCMP guests forming three device groups across two chassis

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.
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.
- 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.

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).

Within any Sync-Failover device group on a vCMP system, connection mirroring can only be implemented between exactly two guests, and the two guests must reside on separate chassis and be homogeneous. The term *homogeneous* normally refers to platforms that are the same model of hardware, but in the context of vCMP, the term refers to guests that are equivalent in terms of resource allocation; that is, the guests are deployed on the same number of slots with the same number of cores allocated to them.

Initial vCMP Configuration Tasks

Overview: vCMP application volume management

The BIG-IP® system allocates all but 30 gigabytes of the total disk space to the vCMP® application volume. Known as the *reserve disk space*, the remaining 30 gigabytes of disk space are left available for other uses, such as for installing additional versions of the BIG-IP system in the future. The vCMP disk space allocation, as well as the creation of the reserve disk space, occurs when you initially provision the vCMP feature as part of vCMP host configuration.

If you want the system to reserve more than the standard 30 gigabytes of disk space for non-vCMP uses, you must do this prior to provisioning the vCMP feature. Adjusting the reserved disk space after you have provisioned the vCMP feature can produce unwanted results.

Important: *When increasing the reserve disk space for additional BIG-IP installations, the recommended amount of space to reserve is 8 gigabytes per installation.*

Viewing disk space allocation for a vCMP application volume

Using this procedure, you can view the amount of disk space, in megabytes, that the system has allocated to a vCMP application volume.

1. In the URL field, type the management IP address that you previously assigned to the system.
`https://<ip_address>`
The browser displays the login screen for 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, in the Volume column, locate the vCMP application volume.
7. In the Size (MB) column, view the size of the application volume, in megabytes.

Modifying disk space allocation for a vCMP application volume

When you provision the BIG-IP system for vCMP, the BIG-IP system dedicates all but 30 gigabytes of disk space to running the vCMP feature. (The 30 gigabytes of reserved disk space protects against any possible resizing of the file system.) Before provisioning the vCMP feature, you can reserve additional space for a

logical disk. Use this procedure if you decide that you need to change the amount of disk space (in megabytes) that the system allocates to a vCMP application volume.

1. In the URL field, type the management IP address that you previously assigned to the system.
`https://<ip_address>`
The browser displays the login screen for 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 **Reserved (MB)** field, increase the amount of disk space that you want to reserve for the logical disk.
The more space you reserve, the less disk space is available for the vCMP application volume.
7. Click **Update**.

vCMP host administrator tasks

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 provision the vCMP feature and to create and manage guests, ensuring that the proper system resources are allocated to those guests.

Task summary

Accessing the vCMP host

Provisioning the vCMP feature

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.

Provisioning the vCMP feature

Before performing this task, ensure that the amount of reserve disk space that the provisioning process creates is sufficient. Attempting to adjust the reserve disk space after you have provisioned the vCMP® feature produces unwanted results.

Performing this task creates the vCMP host (the hypervisor) and dedicates most of the system resources to running vCMP.

Warning: *If the system currently contains any BIG-IP® module configuration data, this data will be deleted when you provision the vCMP feature.*

1. On the Main tab, click **System > Resource Provisioning**.
2. Verify that all BIG-IP modules are set to **None**.
3. From the **vCMP** list, select **Dedicated**.
4. Click **Update**.

After provisioning the vCMP feature, the system reboots TMOS® and prompts you to log in again. This action logs you in to the vCMP host, thereby allowing you to create guests and perform other host configuration tasks.

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® Global Traffic Manager™. 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.*

1. Use a browser to log in to the vCMP® host, using the primary cluster management IP address.

Note: *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 vCMP 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.

Important: If you are planning to add this guest as a mirroring peer for a guest on another chassis, you must ensure that the two guests are configured identically. That is, the number of cores, the number of slots and even the slot numbers on which the guests reside must be the same. Additionally, on each guest individually, the values of the **Minimum Number of Slots** and **Number of Slots** settings must match.

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:

- If you want the guest to be connected to the management network, select the recommended value, **Bridged**.
- If you do not want the guest to be connected to the management network, select **Isolated**.

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.

Selecting **Bridged** causes the **IP Address** setting to appear.

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

- a) 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.
- b) In the **Network Mask** field, type the network mask for the management IP address.
- c) 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, select both an internal and an external VLAN name from the **Available** list, and use the Move button to move the VLAN names to the **Selected** list.
 The VLANs in the **Available** list are part of the vCMP host configuration.
 After you create the guest, the guest can use the selected VLANs to process application traffic.
16. From the **Requested State** list, select **Provisioned**.
 After 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. 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 image installed..

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 vCMP 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

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 GTM™ must log into `guestA` and provision the LTM and GTM 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**.
The Self IPs screen opens.
2. Click **Create**.
The New Self IP screen opens.
3. In the **Name** field, type a unique name for the self IP.
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.
6. From the **VLAN/Tunnel** list, select the VLAN to associate with this self IP address.
 - On the internal network, select the VLAN that is associated with an internal interface or trunk.
 - On the external network, select the 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.

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 (tmsh) 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.

vCMP 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**.

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 for each guest. Each virtual disk is implemented as an image file with an `.img` extension, such as `guest_A.img`.

The vCMP® system automatically creates a virtual disk when you create a vCMP guest. However, after you have created and deployed all guests, you can manage the virtual disks on the system as a way to optimize performance. You must be logged into the vCMP host to manage virtual disks on the system.

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 guest, the host allocates a fixed amount of disk space (typically 100 gigabytes) for that guest virtual machine (VM). This amount of disk space remains fixed regardless of the number of cores per slot that you configure for that guest. For example, a slot with two cores allocated to `guest_A` provides 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 a `.img` extension (for example, `guestA.img`). You identify and manage virtual disks on the system using these file names.

A copy of the virtual disk image for a guest resides on each slot pertaining to that guest (either a single slot or all slots, depending on the guest configuration).

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 vCMP 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 vCMP guests on the system.
2. In the Name column, click the name of the vCMP 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

In some cases, the virtual disk for a guest migrates automatically from one slot to other slots, or from multiple slots to one slot. You do not need to explicitly manage this migration.

Virtual disk migration occurs when you change the state of a single-slot guest from Provisioned to Configured, and then back to Provisioned. In this case, the system moves the guest to a new slot, copying the virtual disk also.

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**.

Understanding Clusters

Overview: Managing a vCMP cluster

One of the tasks that a guest administrator performs is managing the VIPRION[®] 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	Displays the IP address assigned to the cluster. Click this IP address to change it.
Network Mask	Displays the network mask for the cluster IP address.
Primary Member	Displays the number of the slot that holds the primary blade in the cluster.
Software Version	Displays the version number of the BIG-IP [®] software that is running on the cluster.
Software Build	Displays the build number of the BIG-IP software that is running on the cluster.
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.

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

Changing a cluster-related management IP address

A guest administrator can perform this task to change one or more management IP addresses for a vCMP® guest cluster.

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

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

Cluster-related IP addresses

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

Setting Type	Setting	Description
Cluster IP address	IP Address	Specifies the management IP address that you want to assign to the guest cluster. This IP address is used to access the guest, as well as to function as an identifier for the peer guest in a device service clustering (DSC®) configuration.
Cluster IP address	Network Mask	Specifies the network mask for the cluster IP address.
Cluster IP address	Management Route	Specifies the gateway for the cluster IP address. Typically, this is the default route.
Cluster Member IP Address	Slot 1 IP Address	Specifies the management IP address associated with slot 1 of the cluster. You can also set this value to None .
Cluster Member IP Address	Slot 2 IP Address	Specifies the management IP address associated with slot 2 of the cluster. You can also set this value to None .
Cluster Member IP Address	Slot 3 IP Address	Specifies the management IP address associated with slot 3 of the cluster. You can also set this value to None .
Cluster Member IP Address	Slot 4 IP Address	Specifies the management IP address associated with slot 4 of the cluster. You can also set this value to None .

Viewing vCMP Statistics

Overview: Viewing vCMP statistics

After creating vCMP[®] guests to process application traffic, you can display vCMP guest 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 statistics with the BIG-IP Configuration utility

Before viewing guest statistics, 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**.

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

- The state configured for the guest
- The slot numbers on which the guest is running or slated to run
- The guest's management IP address and netmask

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

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.
Slot assignment to guests	Whenever possible, configure a 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.

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, 6, 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 11: 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:

```
Number of physical cores per physical processor * Number of cores per physical core = 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. 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 system prompt, type `vconsole guest_name any_guest_slot_number`.
In this syntax, the variable `any_guest_slot_number` refers to any slot on which the guest is running. Note that for single-slot guests, the slot number is not required.
For example, you can type `vconsole guest_A 1`, where 1 represents slot 1 of the guest.
The system prompts you to enter a user name and password.
2. Log in to the guest using the `root` account and the password `default`.
A system prompt is displayed.
3. At the prompt, determine the primary slot number by typing `tmsh show sys cluster` and locating the Primary Slot ID.
4. If the system output indicates that you are not currently logged into the primary slot of the cluster, type `ssh primary_slot_and_number`.
For example, if the primary slot is slot 2, you can type `ssh slot2`.
Typing this command logs you into the primary slot of the cluster.
5. Type the command `tmsh modify auth user admin shell tmsh`.
This command grants `tmsh` access to the `admin` user account.
6. 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.

7. At the prompt, exit the guest by typing `exit`.
8. Exit the vConsole utility by typing the key sequence `ctrl-]`.
This displays the prompt `telnet>`.
9. 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.

Note: 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 vCMP 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 `tmsh` to manage the guest.

Legal Notices

Legal Notices

Publication Date

This document was published on October 27, 2017.

Publication Number

MAN-0376-07

Copyright

Copyright © 2013-2016, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, APM, Application Acceleration Manager, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, F5 SalesXchange [DESIGN], F5 Synthesis, f5 Synthesis, F5 Synthesis [DESIGN], F5 TechXchange [DESIGN], Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LTM, LineRate, LineRate Systems [DESIGN], LROS, LTM, Message Security Manager, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, Signalling Delivery Controller, SDC, SSL Acceleration, software designed applications services, SDAC (except in Japan), StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WebSafe, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at:

<http://www.f5.com/about/guidelines-policies/patents>

Acknowledgments

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

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

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

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

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

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

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

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

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

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

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

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

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

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

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

This product includes software developed by Christos Zoulas.

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

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

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

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

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

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

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

This product includes software developed by Jared Minch.

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

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

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

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

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

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

This product includes Hypersonic SQL.

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

This product includes software developed by the Internet Software Consortium.

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

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

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

This product includes Intel QuickAssist kernel module, library, and headers software licensed under the GNU General Public License (GPL).

This product includes software licensed from Gerald Combs (gerald@wireshark.org) under the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or any later version. Copyright ©1998 Gerald Combs.

This product includes software developed by Thomas Williams and Colin Kelley. Copyright ©1986 - 1993, 1998, 2004, 2007

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,

2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions. This software is provided "as is" without express or implied warranty to the extent permitted by applicable law.

This product contains software developed by Google, Inc. Copyright ©2011 Google, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Jeremy Ashkenas and DocumentCloud, and distributed under the MIT license. Copyright © 2010-2013 Jeremy Ashkenas, DocumentCloud.

This product includes gson software, distributed under the Apache License version 2.0. Copyright © 2008-2011 Google Inc.

Index

A

- access control
 - administrative 10
- administrator tasks 10, 40, 44
- admin user account
 - modifying 63
- allocation
 - for vCMP application volume 39
- Appliance mode
 - additional tasks 63
 - and boot locations 13
 - and user access restrictions 12
 - enabling on existing guest 64
- Appliance mode types
 - described 12
- audience 10

B

- BIG-IP instances 7
- BIG-IP modules
 - and guest states 29
 - and memory allocation 24
 - provisioning 9–10
 - provisioning within guests 44
- BIG-IP software
 - versions supported 9
- BIG-IP version requirements 13
- blade insertion
 - and slot allocation 29
- blade platforms
 - physical vs. logical cores 61
- blade reduction
 - effects of 25
- blade removal
 - effect on resource allocation 27
 - effects of 28
- bridged guests
 - described 11
- bridged network
 - setting 41
- bridges
 - and management interfaces 59

C

- cluster availability
 - and vCMP guests 37
- cluster IP address
 - modifying 55
- cluster management IP addresses
 - configuring 40
- cluster member IP addresses
 - specifying 45
- cluster member properties
 - described 54
 - viewing 54

- cluster members
 - enabling and disabling 54
- cluster properties
 - described 53
 - viewing 53
- cluster-related IP addresses
 - described 55
- clusters 53
- components 8
- config sync
 - for vCMP systems 35
- config sync IP addresses 36
- configuration data
 - and vCMP provisioning 41
- configuration results 46
- Configured state
 - and disk attachment 49
 - described 29
- connection mirroring
 - on vCMP systems 37
- connections
 - and blade removal 28
 - and memory use 25
- control plane 10
- core allocation
 - about 23
 - based on blade model 23
 - configuring 41
 - determining 24
 - explained 26
 - on solid-state platforms 25
- core availability
 - determining 47
 - increasing 47
- core distribution
 - across slots 26
- cores
 - as system resource 23
 - defined 8
 - increasing 47
- cores per platform
 - calculating 61
- CPU allocation
 - based on blade model 23
- CPU cores
 - and guest states 29
 - increasing 47
- CPU resources
 - allocating 23
 - increasing 47
- custom resource allocation
 - defined 23

D

- data plane
 - vs. management network 10

- Deployed guest state
 - purpose of [43](#)
- Deployed state
 - described [29](#)
 - next steps [46](#)
- device groups
 - for vCMP systems [35](#)
- device trust IP addresses [36](#)
- disk space
 - and vCMP application volume [39](#)
 - and vCMP provisioning [41](#)
 - consuming [59](#)
 - insufficient [41](#)
 - modifying [39](#)
 - reserving [7](#), [9](#), [59](#)
 - viewing [39](#)
- disk space allocation
 - about [49](#)
- disk usage [58](#)
- dual-slot guests
 - example of [31](#)

E

- Ethernet interface
 - of host [11](#)

F

- failover
 - for vCMP systems [35](#), [37](#)
 - on vCMP systems [37](#)
- failover IP addresses [36](#)
- failover methods
 - for vCMP systems [36](#)
- flexible resource allocation
 - defined [23](#)
- floating IP addresses
 - configuring [40](#)

G

- guest access
 - with vconsole utility [11](#)
- guest administrators
 - about [10](#)
 - duties of [10](#)
- guest administrator tasks [44](#)
- guest failover [37](#)
- guest interfaces
 - bridging to physical interface [11](#)
- guest IP addresses
 - configuring [41](#)
- guest-related tasks [44](#)
- guest resource allocation
 - determining [24](#)
- guests
 - about [7](#)
 - additional tasks [63](#)
 - and licensing [9](#)
 - and resource requirements [23](#)
 - and virtual disks [49](#)

- guests (*continued*)
 - configuring BIG-IP modules on [46](#)
 - configuring for LTM [57](#)
 - creating [41](#)
 - provisioning BIG-IP modules for [44](#)
 - setting to Deployed state [43](#)
- guests, multi-slot
 - example of [32](#)
- guest states
 - and virtual disk migration [51](#)
 - configuring [41](#)
 - described [29](#)
- guest traffic load
 - viewing [46](#)
- guest VLANs
 - removing subscription from host [20](#)

H

- HA groups
 - for vCMP systems [37](#)
- hardware resources
 - sharing for SSL and compression [29](#)
- high availability
 - about [10](#)
 - for vCMP systems [35](#)
- host
 - about [7](#)
 - accessing [40](#)
 - and licensing [9](#)
- host administrators
 - about [10](#)
- host administrator tasks [40](#)
- host VLANs
 - publishing to guests [21](#)
- hypervisor
 - See also host
 - accessing [40](#)
 - See also host

I

- instances, BIG-IP [7](#)
- IP addresses
 - for DSC [36](#)
- ISO images
 - and guest states [29](#)
 - installing [41](#)
- isolated guests
 - accessing [44](#)
 - additional tasks [63](#)
 - and Appliance mode [64](#)
 - described [11](#)
- isolated network
 - setting [41](#)

L

- Layer 2/Layer 3 configuration [15](#)
- licensing
 - and Appliance mode [12](#)
 - and guests [9](#)

M

- management interfaces
 - on guests [10](#)
 - wiring [59](#)
- management IP addresses
 - [53](#)
 - configuring [40](#)
 - for guests [16](#), [45](#)
- management network
 - and bridged guests [11](#)
 - and connection to guests [11](#), [15](#)
 - and isolated guests [11](#)
 - vs. data plane network [10](#)
- management network mode
 - setting [41](#)
- memory allocation
 - about [23](#)
 - about calculating [24](#)
 - and blade removal [28](#)
 - based on blade model [23](#)
 - configuring [41](#)
 - determining [24](#)
- memory use
 - and connections [25](#)
- mirroring
 - on vCMP systems [37](#)
- mirroring IP addresses [36](#)
- module configuration [46](#)
- multiple-slot guests
 - example of [32](#)
- multi-tenancy [7](#)

N

- network configuration
 - host vs. guest [15](#)
- network isolation [10](#)
- network state
 - changing [11](#)

P

- performance degradation
 - preventing [59](#)
- pool availability
 - and vCMP guests [37](#)
- pools
 - for BIG-IP modules [44](#), [46](#)
- Provisioned state
 - described [29](#)
- provisioning
 - for vCMP feature [41](#)
- provisioning process [9](#)

R

- redundancy
 - for vCMP systems [35](#)
- resource allocation
 - and guest states [29](#)
 - based on blade model [23](#)

- resource allocation (*continued*)
 - defined [23](#)
 - determining [24](#)
 - explained [26](#)
 - forms of [23](#)
 - on solid-state platforms [25](#)
- resource requirements
 - understanding [23](#)
- resources
 - allocating [9](#)

S

- self IP address configuration [15](#)
- self IP addresses
 - and VLANs [45](#)
 - creating [45](#)
 - for BIG-IP modules [44](#)
- single-core guests
 - and solid-state platforms [25](#)
- single slot guests
 - example of [31](#)
- single-slot guests
 - considerations for [26](#)
- size
 - of vCMP application volume [39](#)
- slot assignment
 - [27](#)
 - about changing [28](#)
 - and scalability [27](#)
 - best practice for [59](#)
 - for guests [25](#)
- slots
 - and virtual disk migration [51](#)
 - assigning to guests [26](#), [41](#)
- software
 - installing for guests [41](#)
- software versions [9](#)
- solid-state drives
 - and core allocation [25](#)
- statistics
 - and disk usage [58](#)
 - viewing for guests [57](#)
 - viewing for virtual disks [57](#)
- system administrator tasks [40](#), [44](#)
- system components [8](#)
- system provisioning
 - for vCMP feature [41](#)
- system resources
 - allocating [9](#), [23](#)

T

- tag discrepancy, VLANs
 - between host and guest [19](#)
- tags, VLANs
 - changing [20](#)
 - identical [17](#)
 - non-matching [18](#)
 - resolving discrepancies [20](#)
- tenants, See guests

- TMOS software
 - installing [41](#)
- tmsh access
 - granting [63](#)
- traffic load
 - viewing for guest [46](#)
- trunk availability
 - and vCMP guests [37](#)
- trunk configuration [15](#)
- trunks
 - about [10](#)

U

- user access restrictions [12](#)
- user account permissions [11](#)

V

- vCMP application volume
 - and disk space [39](#)
 - and disk space allocation [39](#)
- vCMP configuration results [46](#)
- vCMP feature
 - provisioning [9–10](#)
- vCMP guests, *See* guests
- vCMP host
 - accessing [40](#)
- vCMP systems
 - described [7](#)
 - provisioning [41](#)
- vconsole utility [11](#)
- vCPU cores
 - increasing [47](#)
- vCPUs, *See* cores
- version requirements [13](#)
- versions, software [9](#)
- virtual disks
 - about [49](#)
 - and disk space consumption [59](#)
 - and guest states [29](#)
 - as system resource [23](#)

- virtual disks (*continued*)
 - attaching [50](#)
 - creating [41](#)
 - defined [8](#)
 - deleting [51](#)
 - detaching and re-attaching [49–50](#)
 - effect on disk space [41](#)
 - file names and location of [49](#)
 - viewing unattached [50](#)
- virtual disk statistics
 - viewing [57](#)
- virtual interfaces
 - bridging to physical interface [11](#)
- virtual machines [7](#)

See also guests

See also guest VMs

- defined [8](#)
- See also* guests
- See also* guest VMs

- virtual servers
 - for BIG-IP modules [44, 46](#)
- VLAN configuration
 - and vCMP host [15](#)
- VLANs
 - about [10](#)
 - about publishing strategy [16](#)
 - about tag discrepancy solution [19](#)
 - and self IP addresses [45](#)
 - configuring guest use of [41](#)
 - deleting within the guest [20](#)
 - overview of subscription [16](#)
 - publishing to guests [21](#)
 - removing publication on the guest [20](#)
 - viewing within guests [21](#)
 - with identical name but different tags [18](#)
 - with identical tags but different names [17](#)
- VLAN tags
 - changing [20](#)
- VMs
 - defined [8](#)
 - propagating changes to [29](#)
- volumes, *See* vCMP application volume