Veritas Cluster Server Application Note: Support for HP-UX Integrity Virtual Machines

HP-UX 11i v2, HP-UX 11i v3

Augut 2009

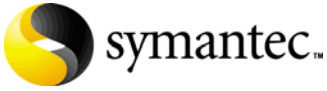

Application Note: Veritas Cluster Server Support for HP-UX Integrity Virtual Machines

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec Corporation product. Such third-party software is licensed separately by its copyright holder.

Technical support

Visit http://www.symantec.com/business/support/assistance_care.jsp for product assistance. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service. If you encounter an error when using a product, include the error number preceding the message when contacting Technical Services. You can also use the error number to search for information in TechNotes or documents on the Web site.

Veritas Cluster Server Support for HP-UX Integrity Virtual Machines

Introduction

This document provides information about support for HP Integrity Virtual Machines (HP Integrity VM) with Veritas Cluster Server.

Review this entire document before installing Veritas Cluster Server on Integrity VM. For information about installing Veritas Cluster Server, refer to the *Veritas Cluster Server Installation Guide for HP-UX*.

About HP Integrity Virtual Machines

Integrity Virtual Machines is a virtualization technique which supports virtualization on HP integrity servers. It lets you manage several virtual operating systems running on the same physical machine. You can consolidate and centrally manage your workloads on one system.

HP Integrity Virtual Machines (Integrity VM) is a soft partitioning and virtualization technology, within the HP Virtual Server Environment, which enables you to create multiple virtual servers or machines with shared resourcing within a single HP Integrity server or nPartition.

Figure 2-1 illustrates the architecture of HP Integrity Virtual Machines.

Figure 2-1 HP Integrity Virtual Machines architecture

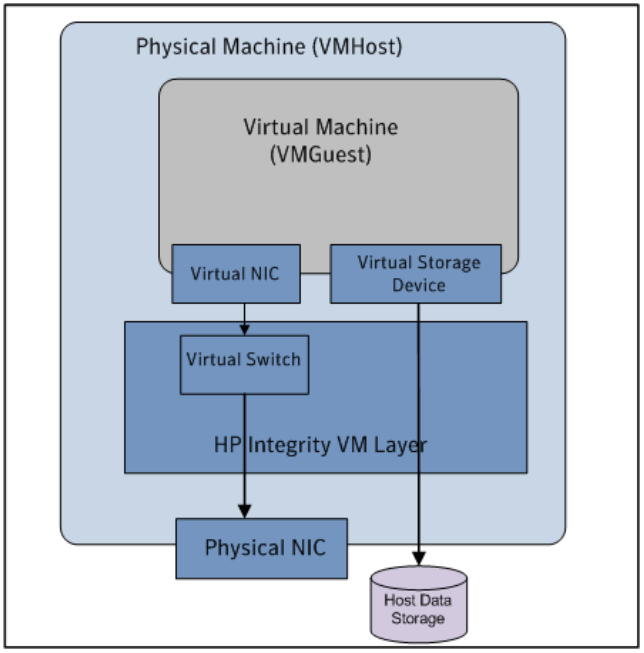

As shown in the figure, a virtual machine (VMGuest) runs within the physical machine (VMHost). The virtual machine is connected to a virtual storage device and a virtual switch (VSwitch).

Virtual switches are virtual entities created by the Integrity virtual machine that behave like a normal switch to the virtual machines.

The virtual switch is mapped to the physical NIC on the physical machine. The virtual switch acts as a bridge between the physical NIC and the virtual NIC.

A virtual storage device connects the physical storage device on the VMHost system to the virtual machine. The VMHost system must have sufficient physical storage for the VMHost and for all of the virtual machines.

For more information about virtual switches, virtual storage devices, and Integrity Virtual Machines, refer to the HP documentation.

Terminology for HP Integrity VM

The following terminology is helpful in understanding VCS for HP Integrity VM.

Table 2-1 HP Integrity VM terminology

Term	Definition
Integrity VM	Integrity Virtual Machines. It is a virtualization technology provided by HP on Integrity Servers.
Attached I/O	A device wholly given to a virtual machine without being virtualized by the VMHost. The meta-data can be seen through the VMGuest.
Shared I/O	A physical disk on VMHost is virtualized and different VMGuests share the disk. Partitions on the disk are presented as a disk or a DVD to Virtual Machines.
VMGuest	A virtual machine with its own operating system, resources, and identity within a physical host.
VMHost	HP Integrity Server which has virtual machines running on top of it. It hosts the Integrity VM package.
VM-PM	VCS supported configuration in which a cluster is formed between VMGuests and physical machines.
PM-PM	VCS supported configuration in which a cluster is formed between VMHosts and is mainly used to manage VMGuests running inside them.
VM-VM	VCS supported configuration in which a cluster is formed between the VMGuests running on top of VMHosts.
VSwitch	A virtual entity that Integrity VM creates. It is mapped to a physical NIC which acts as a switch to VMGuests.
Backing store	The physical device on the VMHost that is allocated to guests, such as a network adapter, disk, or file.
Online VM guest migration	Enables a running guest and its applications to be moved from one VMHost to another without service interruption.

Supported software

Table 2-2 describes the supported VCS and Integrity VM versions on the host and guest systems.

Table 2-2 Supported software for HP Integrity VM

VCS version	Integrity VM version	Host OS version	Guest OS version	Supported configuration	Online VM guest migration support
5.0, 5.0 MP2	3.5, 4.0	HP-UX 11i v2	HP-UX 11i v2, HP-UX 11i v3	■ VM-VM ■ VM-PM	No
5.0 MP2 RP1	3.5, 4.0	HP-UX 11i v2	HP-UX 11i v2, HP-UX 11i v3	■ VM-VM ■ VM-PM ■ PM-PM	
5.0 RP3 on HP-UX 11i v3	4.0	HP-UX 11i v3 September 2008 and later	HP-UX 11i v2 HP-UX 11i v3	■ VM-VM ■ VM-PM ■ PM-PM	No
	4.1	HP-UX 11i v3 March 2009	HP-UX 11i v2 HP-UX 11i v3	■ VM-VM ■ VM-PM ■ PM-PM	Yes* *Only for VM-VM configuration

VCS supported configurations using Integrity VM

Clusters can be formed using various combinations of physical machines (VMHost) and virtual machines (VMGuest) running within the physical machines.

VCS is installed either on the physical machine or on the virtual machine. The following cluster configurations are supported:

- **Cluster among VMGuests (VM-VM)**
  A cluster is formed among virtual machines (VMGuests). The virtual machines in the cluster can be on the same physical machine or on different physical machines. VCS is installed on the virtual machines in the cluster.

VCS manages and controls the applications and resources which are running within the virtual machines. If an application or service group fails, it is failed over to another VMGuest in the cluster.
This configuration does not take care of VMGuest failovers because VCS is running within the VMGuest.

- **Cluster among VMGuests and Physical machines (VM-PM)**
  A cluster is formed among the virtual machines and physical machines. VCS is installed on the virtual machines and on different physical machines in the cluster. The virtual machines are connected to physical machines through the network of their VMHosts.
  This VCS cluster manages services and applications running on the cluster nodes that can either be virtual machines or physical machines. Any faulted application on one node is failed over to other node that can either be a virtual machine or a physical machine.
- **Cluster among VMHosts (PM-PM) [without resource monitoring within guests]**
  The physical machines hosting the virtual machines form a cluster. VCS is installed on the physical machines in the cluster. VCS does not monitor resources within VMGuests. VCS only manages the virtual machines (VMGuests). VCS includes two new agents - HPVirtualMachine agent and HPVSwitch agent to manage the VMGuests and virtual switch for VMGuests respectively.

Cluster among VMGuests (VM-VM)

The following different clusters are possible:

- Cluster among VMGuests on two different VMHosts
- Cluster among VMGuests on the same VMHost

Cluster among VMGuests on two different VMHosts

Figure 2-2 shows a configuration in which a cluster is formed using VMGuests as nodes in a cluster. Faulted applications can fail over to other VMGuest nodes in the cluster.

Figure 2-2 VCS cluster among VMGuests on two different VMHosts

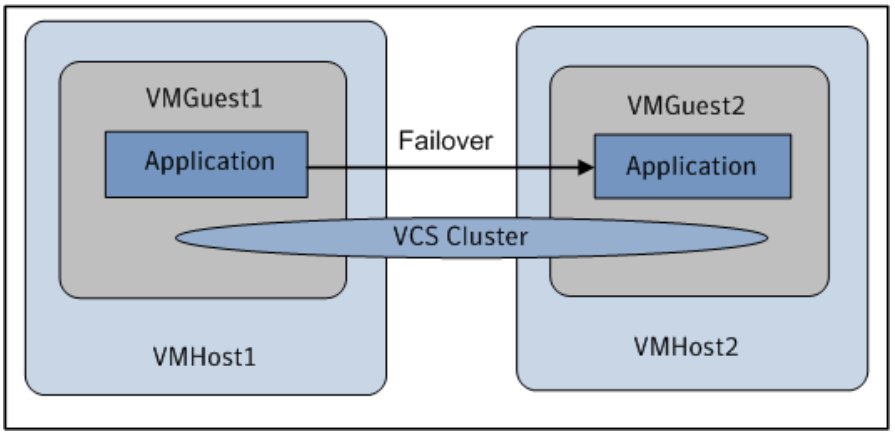

Cluster among VMGuests on the same VMHost

Figure 2-3 shows a configuration in which a cluster is formed between two virtual machines on the same node.

Figure 2-3 VCS cluster within the same VMHost

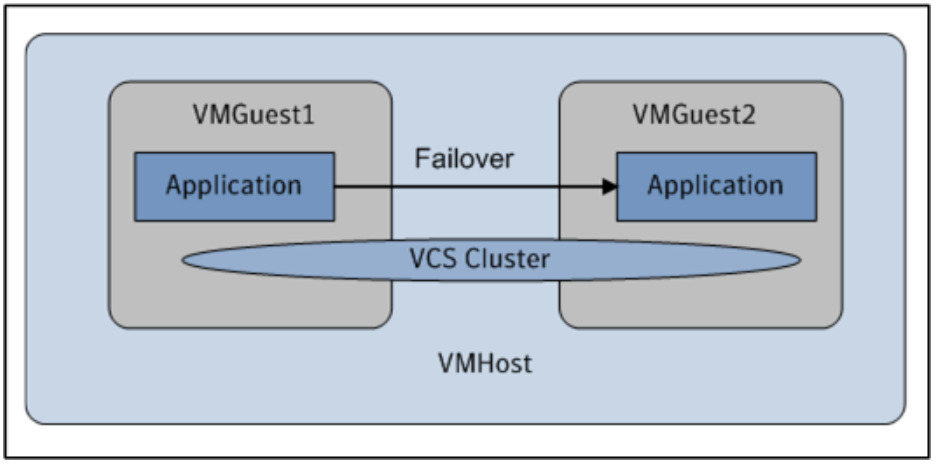

Support for online VM guest migration within clusters

Online VM guest migration is supported in the VM-VM configuration. You can migrate the virtual machine (VMGuest) with VCS and applications running inside it from one physical host to another without service interruption with the help of the online migration command:

```
/opt/hpvm/bin/hpvmmigrate -o -P  <VMGuest> -h  <host>
```

For more information about online VM guest migration, refer to the HP documentation.

Drawbacks of the VM-VM configuration

- This configuration cannot handle virtual machine failovers as VCS is running within the virtual machines.
- As HP Integrity VM does not support Host Bus Adapter (HBA) virtualization, the two virtual machines cannot share the same HBA. Hence, I/O fencing is not supported for virtual machines running on the same VMHost.
  I/O fencing can be supported only using iSCSI LUNS for data and coordinator disks.

Network configuration for VM-VM cluster

To manage the VCS cluster between the virtual machines, you must configure the network and the storage domains for the cluster. The setup details for network and storage configurations are explained in the sections that follow.

Figure 2-4 shows a cluster setup between two virtual machines VMGuest1 and VMGuest2 running on nodes VMHost1 and VMHost2 respectively.

Figure 2-4 Network configuration for virtual machines on different nodes

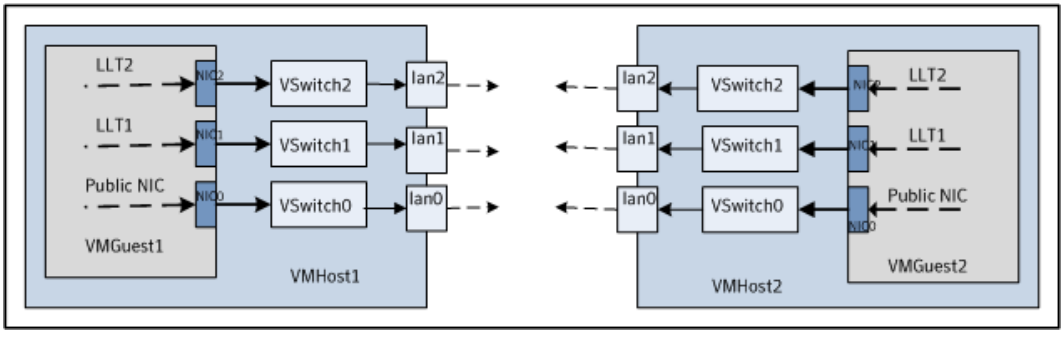

Both nodes VMHost1 and VMHost2 have three physical NICs out of which lan0 is a public NIC, and lan1 and lan2 are private NICs. The private NICs of VMHost1 and VMHost2 are connected to each other through private heartbeat links.

The connections for virtual machine VMGuest1 are as follows:

- Virtual switch VSwitch0 is mapped to the public physical NIC lan0 on VMHost1. A virtual NIC, NIC0 on VMGuest1, is connected to VSwitch0.
- Virtual switch VSwitch1 is mapped to the private physical NIC, lan1 on VMHost1. A virtual NIC named NIC1 on VMGuest1 is connected to VSwitch1.
- Virtual switch VSwitch2 is mapped to another private physical NIC lan2 on VMHost1. A virtual NIC named NIC2 is connected to VSwitch2.

Set up public and private heartbeat network connections for VMGuest2 on the other node VMHost2 in a similar manner. The vswitch names can be different on both the nodes in the cluster.

Figure 2-5 shows the network configuration of a cluster between four VMGuests running on two physical machines. Each of them has two private heartbeat connections which are connected to a vswitch.

Figure 2-5 Cluster between multiple nodes

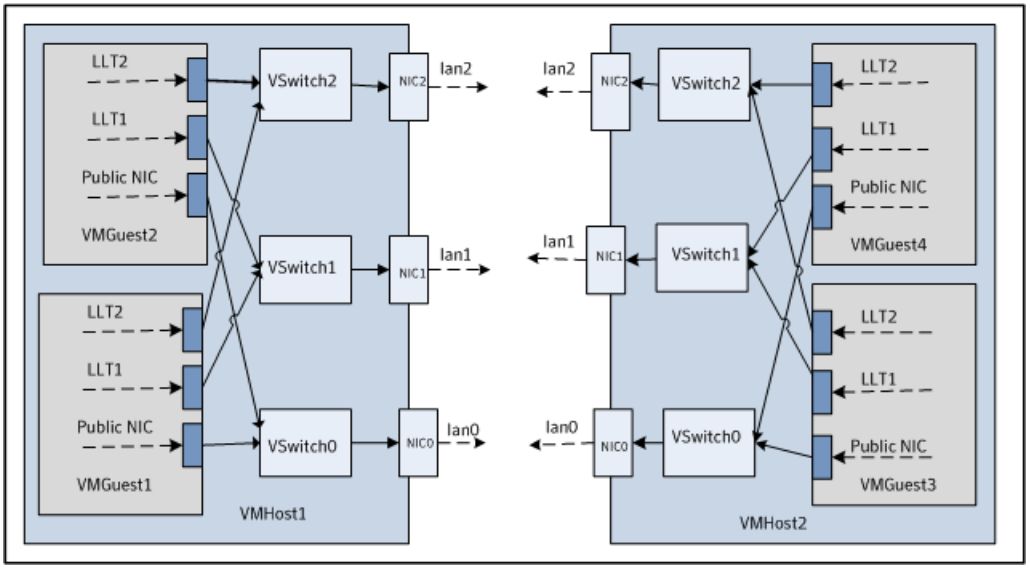

For a VCS cluster having multiple cluster nodes on the same VMHost, an extra VSwitch need not be created for every VMGuest private heartbeat connection.

After a VSwitch is created, it can be used by several VMGuests on the same VMHost.

Storage configuration for VM-VM cluster

Raw disks are supported for cluster nodes in VM-VM configuration.

A whole disk is provided to a virtual machine without any disk virtualization. The disk can be a physical disk connected to the system or it can be a shared disk.

Use of raw disks with no disk groups or volume groups configured is not recommended in VM-VM configuration, as it can lead to data corruption.

Setting up virtual machines

Following is a high-level overview of the steps for setting up virtual machines. For detailed instructions, refer to HP documentation.

To set up virtual machines

- Before creating virtual machines, ensure that CPU and memory resources are available to create VMGuests on all nodes in the cluster.
- Install the HP Integrity VM package on the VMHosts.
- Create VMGuests.
  Before creating VMGuests, you must have the following:
  - Virtual switches to enable networking for guests.
  - Storage for guests (optional, can be configured while creating VMGuests).
- Install the operating system in the VMGuests.
- Repeat the above steps for all VMGuests that you want to be a part of the cluster.
- Install VCS on all the VMGuests. For information about installing VCS, refer to the *Veritas Cluster Server Installation Guide*.
- Configure the VCS resources that you want VCS to manage.
- If you intend to use the online VM guest migration feature and move the running VMGuest on a different physical machine, Symantec recommends that you set the VCS_GAB_TIMEOUT value in the /opt/VRTSvcs/bin/vcsenv file on all the VMGuests, to avoid the VCS engine missing heartbeats with GAB on a loaded system during online migration.

  ```
  VCS_GAB_TIMEOUT=30000
  export VCS_GAB_TIMEOUT
  ```

Cluster among virtual machines and physical machines (VM-PM)

A cluster is formed among a VMGuest on one VMHost and other physical machines in the cluster. VCS is installed on the virtual machines and on different physical machines in the cluster. The virtual machines are connected to physical machines through the network of their VMHosts.

This VCS cluster manages and monitors services and applications running on the cluster nodes that can either be virtual machines or physical machines. Any faulted application on the virtual machine can fail over to the other physical machine in cluster and vice-versa.

Figure 2-6 shows a cluster between a physical machine and a virtual machine running on another physical machine. VCS is installed on the virtual machine and the other physical machine.

Figure 2-6 VM-PM cluster

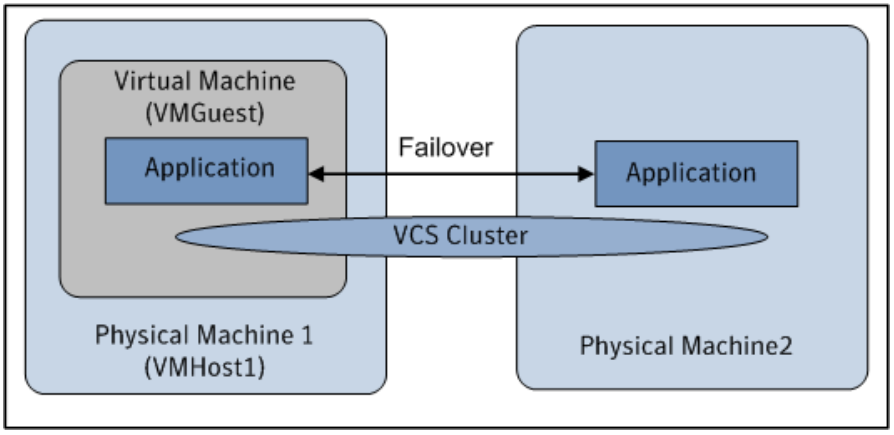

Note: Symantec does not recommend a cluster between a VMGuest and its own VMHost for VCS clustering solutions.

Network configuration

The network connection for the physical machine is similar to any other node in a VCS cluster. The virtual machine is connected to the physical machine through VSwitches and physical NIC on its VMHost.

Figure 2-7 shows the network configuration for the VM-PM configuration.

Figure 2-7 Network configuration for VM-PM configuration

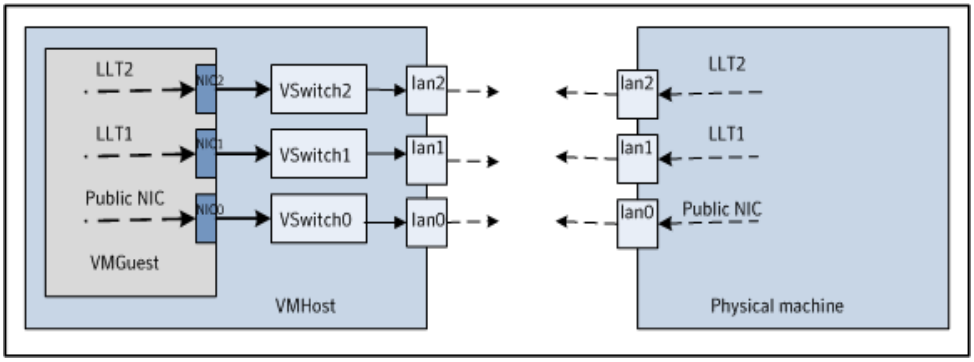

The network connections for the virtual machine VMGuest are as follows:

- A virtual switch named VSwitch0 is mapped to public physical NIC lan0 on VMHost. A virtual NIC, NIC0 on VMGuest is connected to VSwitch0.
- A virtual switch named Vswitch1 is mapped to the private NIC lan1 on VMHost. A virtual NIC, NIC1 on VMGuest is connected to VSwitch1.
- A virtual switch named Vswitch2 is mapped to the private NIC lan2 on VMHost. A virtual NIC, NIC2 on VMGuest is connected to VSwitch2.

A typical diagram of network configuration consisting of a physical machine and two VMGuests on same VMHost can be represented as shown in Figure 2-8.

Figure 2-8 Network configuration consisting of two VMGuests

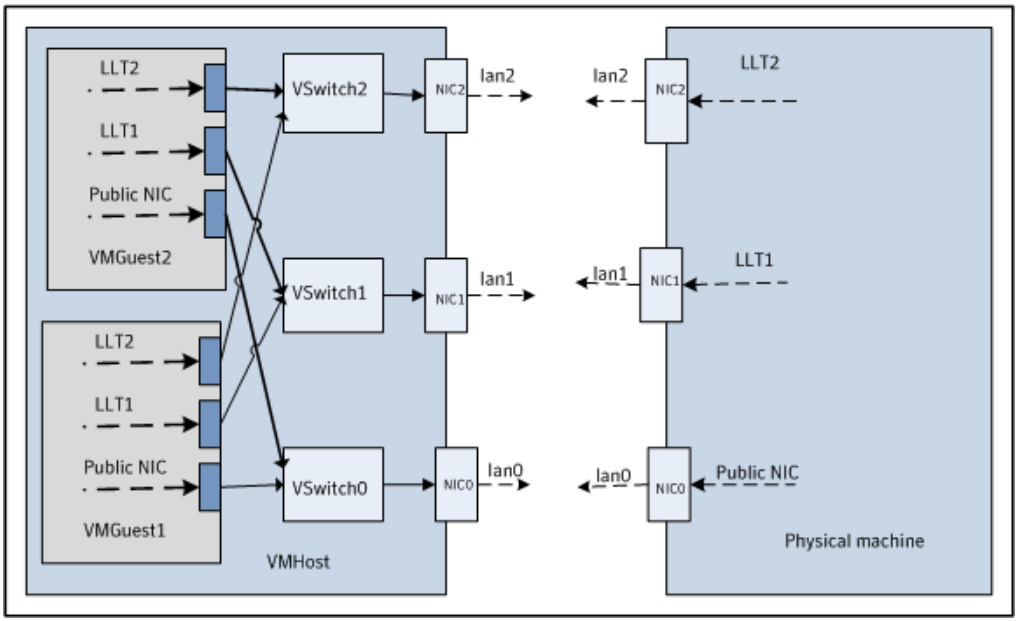

Storage configuration

The physical storage support for VM-PM configuration is similar to the VM-VM configuration. For a VCS cluster, the physical machine in the cluster can have storage support for only those devices that the VMGuests support. Review the list of devices the VMGuests support.

See "Storage configuration for VM-VM cluster" on page 13.

Setting up a VM-PM cluster

Following are the high-level steps for setting up a VM-PM cluster. For detailed instructions, refer to HP documentation.

- Ensure that CPU and memory resources are available to create VMGuests on all nodes in the cluster.
- Install the HP Integrity VM package on the VMHosts.
- Create VMGuests.
  Before creating VMGuests, you must have the following:
  - Virtual switches to enable networking for guests.

- Storage for guests (optional, can be configured while creating VMGuests).

- Install the operating system in the VMGuests.
- Repeat the above steps for all VMGuests that you want to be a part of the cluster.
- Install VCS on all the physical machines and VMGuests which are to be part of the cluster. For information about installing VCS, refer to the *Veritas Cluster Server Installation Guide*.
- Configure the VCS resources that you want VCS to manage.

Cluster between physical machines (PM-PM)

The physical machines hosting the virtual machines form a cluster. VCS is installed on the physical machines in the cluster. VCS does not monitor resources within VMGuests. VCS only manages the virtual machines (VMGuests).

Figure 2-9, shows a VCS cluster between VMHost1 and VMHost2.

Figure 2-9 PM-PM configuration

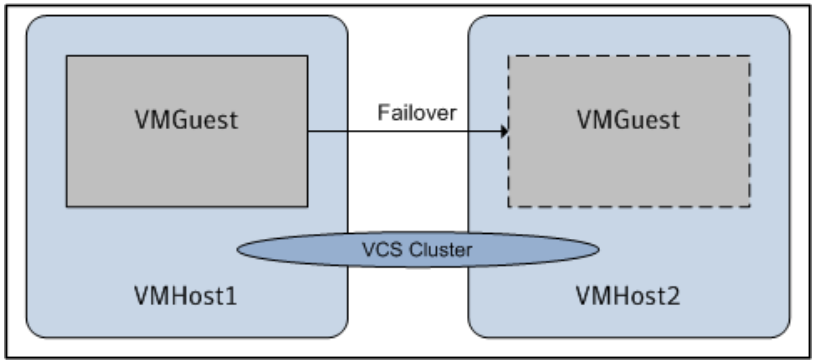

If the virtual machine on one of the VMHost faults, it is failed over to the other VMHost. VMGuests, which are used as resources in VCS, must have the same configuration across all VMHosts. The failover of VMGuest when it is running is not supported in this configuration.

The HPVirtualMachine agent manages the virtual machines and the HPVSwitch agent manages the virtual switches for the virtual machines. The agents use online, offline, monitor, and clean functions to manage the virtual machines and virtual switches.

The storage for the VMGuests must be accessible to all the VMHosts in the cluster. Refer to the following special considerations for details.

Special Considerations

- The HPVirtualMachine agent does not wait for the operating system to load fully. This agent returns online after the loading of boot manager only.
- The MAC addresses for the guests should be manually configured and should be the same across VMHosts.

Network and storage considerations

The network configuration and storage configuration for the physical machines (VMHosts) is the same as the nodes in VCS cluster configurations.

For information on configuring the physical machines, refer to the *Veritas Cluster Server Installation and Configuration Guide*.

However, virtual machines (VMGuests) which are used for failover between VMHosts need to have specific network and storage requirements as shown in Figure 2-10.

Figure 2-10 Network and storage considerations for PM-PM configuration

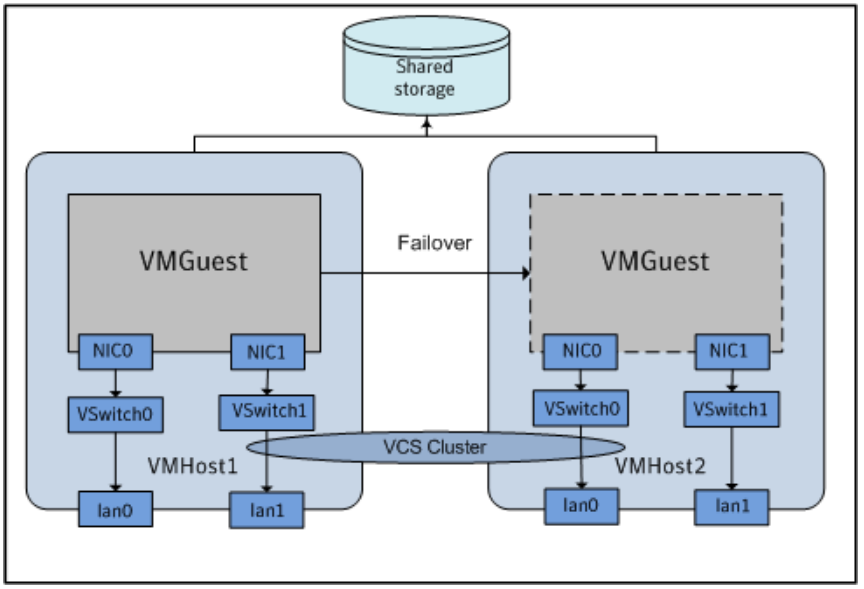

Setting up virtual machines

Following is a high-level overview of the steps for setting up the virtual machines. For detailed instructions, refer to HP documentation.

To set up virtual machines

- Before creating virtual machines, ensure that:
  - Physical machines form a cluster with VCS installed on them. For information about installing VCS, refer to the *Veritas Cluster Server Installation Guide*.
  - CPU and memory resources are available to create VMGuests on all nodes in the cluster.
- Install the HP Integrity VM package on all cluster nodes.
- Create VMGuests.
  Before creating virtual guests, you must have:
  - Virtual switches to enable networking for guests.
  - Storage for guests (optional, can be configured while creating VMGuests).
- Install the operating system of the VMGuests in a shared storage accessible to all the VCS cluster nodes.
- Configure the VMGuest as a resource in VCS.

Note: Online VM guest migration within clusters is not supported in PM-PM configuration.

Bundled agents for HP Integrity VM

The following agents are used to manage the virtualization environment and provide high availability for application and process-related resources.

- HPVirtualMachine agent
- HPVSwitch agent

For information on the Mount, Volume, and DiskGroup agents, refer to the *Veritas Cluster Server Bundled Agents Reference Guide*.

HPVirtualMachine agent

The HPVirtualmachine agent brings online, takes offline, and monitors virtual machines (VMGuests) that are running on the physical host (VMHost).

Agent functions

The functions that are associated with this agent are:

Online	Uses the `hpvmstart` command to start the virtual machine (VMGuest).
Offline	Attempts a graceful shutdown of the virtual machine. Uses `hpvmstop` command to stop the VMGuest.
Monitor	Uses the `hpvmstatus` command to detect the virtual machine's state. Returns the following status: ■ If the virtual machine is missing, it returns the state as UNKNOWN. ■ If the virtual machine is running, then it returns the state as ONLINE. ■ If the virtual machine is not running, it returns OFFLINE.
Clean	Forcefully shuts down the virtual machine. It uses `hpvmstop` command with `-h` argument.

State definitions

The states that are associated with this agent are:

ONLINE	Indicates that the virtual machine (VMGuest) is up and has a heartbeat.
OFFLINE	Indicates that the virtual machine is turned off.
FAULTED	Indicates that the virtual machine has failed to start up using the online operation. This problem can occur due to an issue with the VMGuest configuration. It can also occur due to a sudden unexpected shutdown of the virtual machine.
UNKNOWN	Indicates the agent cannot determine the virtual machine's state. This state can occur if the virtual machine has not been yet created or the resource type definition of the HPVirtualMachine agent is not configured properly.

Agent attributes

Refer to the required and optional attributes while configuring the HPVirtualmachine agent.

Table 2-3 shows the required attributes for configuring the HPVirtualmachine agent.

Table 2-3 Agent required attributes

Required attribute	Description
VMName	The VMGuest name which the agent monitors. This attribute is unique and must be configured. Type and dimension: string-scalar Default: N/A Example: vcsivm

Table 2-4 shows the optional attributes for configuring the HPVirtualmachine agent.

Table 2-4 Agent optional attributes

Optional attribute	Description
DelayAfterOnline	Defines the maximum time that the VMGuest can take to reach the EFI (Extensible Firmware Interface) shell. If the VMGuest leaves the EFI shell before the time provided by DelayAfterOnline, the online function exits at that time only. This attribute is added as to ensure that the boot time of VMGuest up to EFI shell remains flexible and can be modified as per the requirement. Type and dimension: integer-scalar Default: 30 Example: 50

Resource type definition

```
type HPVirtualMachine (
    static str ArgList[] = { VMName, DelayAfterOnline }
    str VMName
    int DelayAfterOnline = 30
)
```

Agent limitation

The agent cannot detect if the operating system hangs. Even if the virtual machine reports the VMGuest state as ONLINE, it does not mean that the operating system running within the guest is functioning properly.

HPVSwitch agent

Use the HPVSwitch agent to manage and control the virtual switches that are associated with the network connection of the virtual machines (VMGuests). Virtual switches are virtual entities that resemble a normal switch to VMGuests. Virtual switches are mapped to the physical network interface card (NIC) on the physical machine (VMHost).

Requirements

For the HPVSwitch resource agents to work, the virtual switches must be present on the VMHost because virtual switches are controlled from the VMHost. VMGuest has no information about them.

Note: If the virtual switch is started upon system reboot, you may observe a concurrency violation for the HPVSwitch resource configured as a part of a failover service group.

Agent functions

The functions associated with this agent are:

Online	Starts the virtual switch using the native `hpvmnet` commands. The online function requires the VSwitchName attribute to operate. When a switch is started, all the virtual NICs mapped to that VSwitch activate.
Offline	Stops the virtual switch using the `hpvmnet` command. After the shutdown of VMGuest, no network is required. The VSwitch is halted using native `hpvmnet` command.
Monitor	This function checks the status of VSwitch and returns the result. ■ If VSwitch is running, it returns ONLINE. ■ If VSwitch is not running or is halted, it returns OFFLINE. ■ If VSwitch is not present on the system, then it returns the state UNKNOWN.
Clean	Stops the virtual switch. It is similar to the Offline function.

State definitions

The states that are associated with this agent are:

ONLINE	The VSwitch for HPVSwitch agent is running properly.
OFFLINE	The VSwitch is halted. It is currently switched off.

FAULTED	The VSwitch has unexpectedly turned off or it failed to start.
UNKNOWN	The VSwitch is not configured properly.

Agent attributes

Refer to the required attributes while configuring the HPVSwitch agent.

Table 2-5 Agent attributes

Attribute	Description
VSwitchName	Name of the virtual switch (VSwitch) that the HPVSwitch agent monitors. This attribute is unique and must be configured. Type and dimension: integer-scalar Default: N/A Example: pub0

Resource type definition

```
type HPVSwitch (
    static str ArgList[] = { VSwitchName }
    str VSwitchName
)
```

Agent limitation

The HPVSwitch agent does not manage the virtual NIC on the VMGuest.

Configuring HP Integrity VM service groups

Following is the high-level procedure to configure HP Integrity VM service groups:

- Create Integrity VM service group
- Create HP Virtual Switch resource
- Create VMGuest resource
- Create backing storage resources for the VMGuest resource
- Sample dependency diagrams
- Sample configurations

Create Integrity VM service group

- Configure a failover service group called *IVM*.
- The SystemList attribute should contain all the cluster nodes where the VMGuest can failover.

Create HP Virtual Switch resource

- Configure an HPVSwitch resource vswitch inside this service group *IVM*.
- Assign the virtual switch name that you want to monitor in the attribute VSwitchName.
- Create a NIC resource *nic* inside the service group *IVM*.
- The Device attribute should contain the backing physical NIC of the virtual switch.
- Create a dependency between vswitch (parent) and nic (child).

Create VMGuest resource

- Configure a HPVirtualMachine resource *vm* inside the service group *IVM*. Add the name of the VMGuest that you want to monitor in the VMName attribute.
- Create a dependency between vm (parent) and vswitch (child).

Create backing storage resources for the VMGuest resource

- If you intend to use a disk group as the backing store for the VMGuest, configure a DiskGroup resource *dg* with the diskgroup name in the DiskGroup attribute. In case, you are using LVMs, raw disks, or CVM, use corresponding storage resources.
- In case, the VMGuest is installed on a VxVM volume, create a Volume resource *vol*. The volume name should be the backing storage for the VMGuest. Refer to hpvm commands to know more about backing storage for VMGuests.
- Create a dependency between vm (parent) and vol (child).
- Ensure that all resources are enabled before bringing them online.

Sample dependency diagrams

VMGuest with diskgroup as backing store

Figure 2-11 Sample service group for a HPVirtualMachine resource

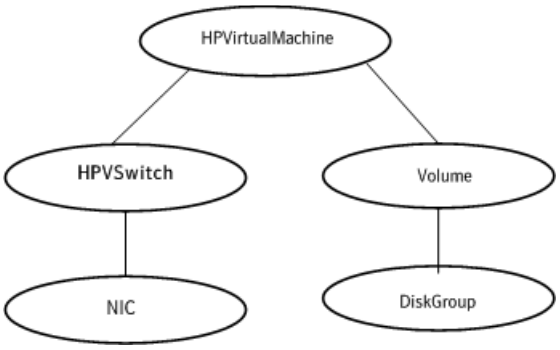

VMGuest with raw disk as backing store

Figure 2-12 Sample service group for a HPVSwitch resource

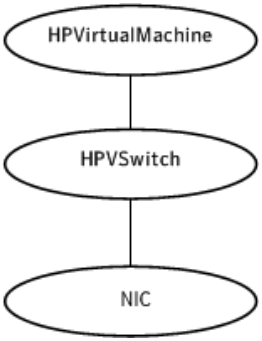

Sample configurations

Service group with disk group as the backing store

```
include "types.cf"
cluster ivmclus (
        )
system sysa (
```

```
        )
system sysb (
        )
group IVM (
        SystemList = { sysa = 0, sysb = 1 }
        )
        DiskGroup dg (
                DiskGroup = dg1
                )
        HPVSwitch pub0 (
                VswitchName = pub0
                )
        HPVirtualMachine vm (
                VMName = vcsivm1
                )
        NIC nic (
                Device = lan0
                )
        Volume vol (
                Volume = vol1
                DiskGroup = dg1
                )
        pub0 requires nic
        vm requires pub0
        vm requires vol
        vol requires dg

        // resource dependency tree
        //
        //      group IVM_OS
        //      {
        //      HPVirtualMachine vm
        //          {
        //          Volume vol
        //              {
        //              DiskGroup dg
        //              }
        //          HPVSwitch pub0
        //              {
        //              NIC nic
        //              }
        //          }
        //      }
```

Service group with raw disk as the backing store

```
include "types.cf"
cluster Test (
                )
system SysA (
        )
```

```
system SysB (
        )
group g1 (
        SystemList = { SysA = 0, SysB = 1 }
        )
        HPVSwitch vswitch (
            VSwitchName = public0
            )
        HPVirtualMachine vm (
            VMName = vmsharedhp
            )
        vm requires vswitch
```

Restrictions with VCS on VMGuests

- Controlling applications running within VMGuests is not supported in a PM-PM configuration.
- I/O Fencing is not supported in the HP Integrity VM Guest. I/O Fencing relies on SCSI-3 PGR support. HP Integrity VM does not support SCSI-3 PGR in the HP Integrity VM Guest for virtualized disks. I/O fencing can be supported only using iSCSI LUNs for data and coordinator disks.