



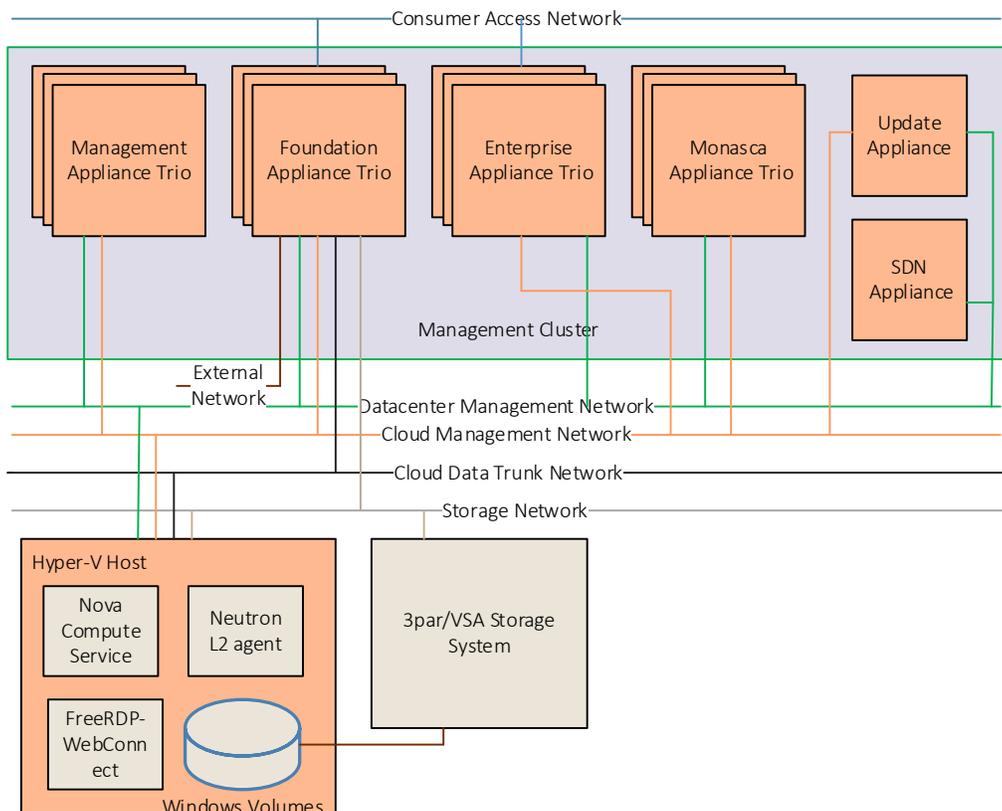
# HP Helion CloudSystem 9.0: Using Windows Hyper-V Server 2012 R2 as Compute Nodes

This white paper provides an overview of Hyper-V support in CloudSystem, including best practices and recommendations for using Hyper-V as compute nodes in an ESXi or KVM management hypervisor.

## Overview

HP CloudSystem is based on the HP Helion distribution of OpenStack cloud software. It integrates hardware and software to deliver core Infrastructure as a Service (IaaS) provisioning and lifecycle management for compute, network, and storage resources.

CloudSystem supports ESXi clusters, KVM and Hyper-V hosts as compute nodes. The following diagram shows the key components that enable CloudSystem to use Hyper-V.



The following is a summary of the networks shown in the diagram.

Network	Number	Purpose	Connected to...
Data Center Management Network	1	This network connects virtual appliances to HP 3PAR, Hyper-V compute nodes, and enclosures for the administrator to manage	Management appliance trio Cloud controller trio Enterprise appliance trio Monitoring appliance trio Update appliance SDN appliance Hyper-V compute nodes
Cloud Management Network	1	This network connects all CloudSystem virtual appliances and Hyper-V compute nodes. This is a private network.	Management appliance trio Cloud controller trio Enterprise appliance trio Monitoring appliance trio Update appliance SDN appliance Hyper-V compute nodes
Consumer Access Network	1	This network provides cloud users with access to the CloudSystem OpenStack portal in the Cloud controller and the Marketplace Portal in Enterprise. Cloud users on this network are not able to access other CloudSystem virtual appliances managed by the cloud administrator.	Cloud controller trio Enterprise appliance trio
External Network	1	This network allows cloud end users to attach public IP addresses to their provisioned virtual machine instances.	Cloud controller trio
Cloud Data Trunk Network (Provider Networks and/or Private Networks)	At least 1	A Provider Network is a data center network routed through the existing data center infrastructure. A Private Network is created from a pool of VLANs. Both networks support instance communication.	Cloud controller trio Hyper-V compute nodes
Storage Network	1	This network connects virtual appliances and compute nodes to HP 3PAR	Management appliance trio Hyper-V compute nodes HP 3PAR ISCSI storage system

## Compute node requirements and notes

CloudSystem uses the Windows platform with the Hyper-V role enabled to run the Nova compute agent. The nova-compute service as well as necessary Python components are installed directly onto the Hyper-V server and run as the “openstack-compute” service.

See the *HP Helion CloudSystem Support Matrix* at [www.hp.com/go/cloudsystem/docs](http://www.hp.com/go/cloudsystem/docs) for Windows platforms supported as compute nodes.

CloudSystem differentiates Hyper-V compute nodes based on “Microsoft failover Clustering feature” installed on the Hyper-V servers.

- Standalone Hyper-V compute node has “Microsoft Failover Clustering feature” disabled
- Clustered Hyper-V compute node has “Microsoft Failover Clustering feature” enabled

### Planning for setup of Hyper-V compute nodes (standalone)

Before using Hyper-V compute with CloudSystem, you should take the following actions

- Allocate ample storage for virtual machine images running on the compute nodes. CloudSystem stores virtual machine images and hard disk files to the directory location specified under “**Virtual Hard Disks**” configuration under **Hyper-V Settings** in **Hyper-V manager**. Be sure to specify the directory location with sufficient storage under Hyper-V settings on the Hyper-V host for CloudSystem to use it for storing virtual machine image/hard disk files.

Do not change the directory location specified in “**Virtual Hard Disks**” under **Hyper-V Settings** in **Hyper-V Manager** after activating the compute node. This location is not updated in CloudSystem.

- CloudSystem requires Hyper-V compute nodes to have a minimum of three NICs: one connected to the **Data Center Management Network**, one connected to the **Cloud Management Network** and one connected to the **Cloud data trunk**. Optionally, one NIC is required for the **Storage Network** if Hyper-V compute nodes need to attach volumes provided by Cinder.
- The Hyper-V compute nodes must be synchronized with the NTP server running on the Management appliance to ensure proper operation of the compute nodes. If your environment is configured for Active Directory, you may need to synchronize time only for the Active Directory Domain Controller. If an external NTP server is used for Hyper-V compute nodes, be sure to have same NTP server used for CloudSystem.
- Hyper-V compute host names should not exceed 15 characters. This Microsoft limitation is described in [Windows Technet Library](#).
- For CloudSystem to prepare the Hyper-V compute node to be able to attach to volumes provided by Cinder, the default Initiator name under iSCSI Initiator settings should not be changed.
- CloudSystem needs “Administrator” account credentials to be provided for activating a Hyper-V host to manage the deployed virtual machines. If new “Local User” account credentials are used, ensure that the “Local User” is a member of the “**Administrators**” and “**Hyper-V Administrators**” groups.

**Important:** CloudSystem activation does not support using domain account credentials for the Hyper-V compute node.

### Planning for setup of Hyper-V compute nodes (clustered)

CloudSystem supports the “**Microsoft Failover clustering**” feature in Windows that enables high-availability for the virtual machines deployed on Hyper-V cluster compute nodes. Unlike ESXi in CloudSystem, each Hyper-V host in the “Microsoft Failover Cluster” is treated as an individual Hyper-V compute node and must be activated in CloudSystem individually.

Before using Hyper-V cluster compute nodes with CloudSystem, take the following actions.

- Ensure all the standalone Hyper-V compute requirements (applicable and stated above) are met.

- Hyper-V compute nodes must be part of “Failover Cluster Manager” and need to have at least one “Clustered Shared Volume” enabled. Refer to [Microsoft Technet Library](#) for information about installing a Failover cluster.
- CloudSystem uses **Clustered Shared Volumes (CSV)** created on the Hyper-V compute node to store virtual machine configuration and disk files. If multiple CSVs are available, CloudSystem picks the CSV with more disk space to deploy instances. Best practice is to use/extend single a CSV with ample space instead of multiple CSVs for optimal disk usage on Hyper-V compute nodes.
- CloudSystem requires that you activate all Hyper-V compute nodes that are part of the Failover Cluster before deploying instances. Otherwise, the instances will be unmanaged in CloudSystem if the instances are moved to a non-activated Hyper-V compute node in the Failover cluster.

## Storage requirements and notes

### Boot image storage and caching

Hyper-V compute nodes support VHD and VHDX file formats for Virtual Machine image files. An image cache is created on each Windows volume when a Glance image is used for the first time to boot instances on it.

**VHD:** A copy of the Glance image is made for each flavor used (according to the boot disk size specified in the flavor). The boot disks of provisioned VM instances are created as linked clones and the images in cache are used as the base disks.

**VHDX:** A copy of the Glance image is made and **dynamically extended** to the flavor size for each **provisioned** VM instance. The boot disks of provisioned VM instances are created as linked clones and the images in the cache are used as the base disks.

All boot disks are thick provisioned. The VM instances are linked clones, and they may not use as much space as that allocated by the flavor. The OpenStack Nova scheduler evaluates compute nodes’ available disk space based on disk allocation and not actual utilization. You may see more free space available than what CloudSystem shows as allocated.

The image cache is not cleared automatically from Hyper-V compute nodes when all instances are deleted. The image cache requires manual cleanup. Following is the location where the image caches will be stored on Hyper-V compute nodes.

- **Standalone Hyper-V compute nodes:** The image cache is stored in “\_base” directory under **Virtual Hard Disks** directory location specified under **Hyper-V Settings** in **Hyper-V manager**.
- **Clustered Hyper-V compute nodes:** The image cache is stored in “\_base” directory under **Clustered Shared Volumes** created on the compute nodes.

### Prepare images for use with Hyper-V Compute

Hyper-V compute nodes support the VHD and VHDX file format for virtual machine instances. Detailed instructions for installing virtual machines on Hyper-V can be found in [Microsoft Technet Library](#).

After you have successfully created a virtual machine, you can then upload the image to Glance.

### Snapshots

VM snapshots result in a new Glance image that is created from a VM instance. The image is created from the VM’s boot disk and the size is determined by the flavor. It is important to keep in mind that VM boot disks are thick provisioned. This size could be much larger than the original boot image that was used to provision the VM. The time required to upload the snapshot into Glance is proportionally higher.

## Cinder Block Storage

CloudSystem supports HP 3PAR and HP VSA StoreVirtual block storage iSCSI volumes. Hyper-V compute nodes use the Microsoft iSCSI Initiator running on the host to connect to the 3PAR/VSA iSCSI storage system through the Storage Network. Use the CloudSystem Operations Console to register HP 3PAR/VSA block storage drivers and volume types.

### Supported operations

- Data Disks – Attach or detach a Cinder volume
- Boot Disk – Create a volume from a bootable Glance VMDK image

## Networking requirements and notes

Hyper-V compute nodes run the neutron-agent layer 2 service that works with the OpenStack Neutron server. These work in tandem to configure Neutron networks on the VM instance provisioned on Hyper-V compute nodes.

The Neutron L2 agent is automatically deployed and configured during Hyper-V compute activation. The L2 agent on the Hyper-V compute uses Hyper-V Native Virtual switch capabilities for provisioning networks on the launched VM instances. The Native Hyper-V Virtual Switch created on the Cloud Data trunk network on Hyper-V compute nodes will be used for network connectivity for virtual machines.

CloudSystem also supports NIC Teaming of Cloud Data trunk interfaces. CloudSystem provides the choice of specifying an existing Hyper-V virtual switch created on the Cloud Data trunk interface or teamed interfaces, or Cloud Data trunk physical interface names on which Hyper-V Virtual switch will be automatically created.

Following is the default NIC teaming configuration that CloudSystem configures on activation:

- **Teaming Mode:** Switch Independent
- **Load Balancing Algorithm:** Dynamic

### Network provisioning in Hyper-V compute nodes

**Important:** CloudSystem supports only VLAN networks with Hyper-V compute nodes.

The following is a high-level summary of network connection configuration for provisioned VMs. This process ensures VM instances are booted up with proper network connectivity.

- When a VM is provisioned on a Hyper-V compute node, Nova-compute provisions the VM with the required NICs as defined in the provisioning request. It then waits for the port-binding to happen on the VM NICs with appropriate Network VLANs from Neutron.
- The L2 Neutron agent on a Hyper-V compute node detects the VM spawned and binds VM NIC ports with the appropriate Network VLAN IDs. It then notifies the port-binding events to Nova on successful binding of ports with VM NICs
- When Nova-compute receives the VM port-binding events, it powers on the VM instance so that the VM has proper network connectivity.

## Security group support

Hyper-V compute nodes use “Microsoft Hyper-V Virtual Switch Extended ACL feature” to support OpenStack security groups. Hyper-V supports stateful firewall rules so you can define OpenStack security group rules in the same way as you do for KVM.

### Important:

- If the number of security group rules per port is more than 100, you will see degraded performance of the Hyper-V host, and VM instance boot time will increase.
- Security group rules created with ICMP Protocol and Remote Group ID also need to have a reverse direction rule for ICMP Protocol to work correctly.

## Activating Hyper-V compute nodes

Before activating Hyper-V compute with CloudSystem, you should take the following actions.

- Bring up the Cloud Management Network on the compute node. A new IP address is acquired for the compute node from the Management appliance DHCP server.

```
ipconfig /release <Hyper-V server network adapter connected to the CLM>
```

```
ipconfig /renew <Hyper-V server network adapter connected to the CLM>
```

- Hyper-V compute node must meet certain prerequisites before it can be activated in CloudSystem. Follow the below procedure to download and execute the prerequisite script.

Download the script CSHyperVPreReq.ps1 from <http://< Management appliance VIP>:8085/scripts/> on the Hyper-V compute and run the script in Powershell as follows.

```
PS C:\Users\Administrator.CLOUD\Desktop> powershell -ExecutionPolicy Bypass .\CSHyperVPreReq.ps1 <Hyper-V Compute CLM-IPAddress>
```

- CloudSystem uses the CloudBase Solutions FreeRDP-Webconnect solution to enable RDP access to Hyper-V VM instances using the Launch Console feature in the OpenStack user portal. The administrator is responsible for security and other updates to FreeRDP-Webconnect. Follow the below procedure to download and execute pre-requisite script

Download FreeRDP-Webconnect from [CloudBase Solutions](#) and copy it to the Management appliance (ma1) under `/var/csm/www/msi` directory and then perform the following actions.

Set privileges on the .msi file:

```
chmod 644 FreeRDPWebConnect.msi
```

Change the owner of the .msi file:

```
chown root:root FreeRDPWebConnect.msi
```

CloudSystem auto imports Hyper-V compute nodes to be activated and can be seen from the Operations Console UI. You can use Operations Console to activate imported Hyper-V compute nodes. The activation process performs a set of validations that ensure that the Hyper-V compute node meets the prerequisites. It also automatically deploys the Nova-Compute Service and Neutron-agent Service on the Hyper-V host. After successful activation, the Nova scheduler can use the Hyper-V compute node for provisioning VM instances.

## Recommendations for using activated Hyper-V compute resources

After activating Hyper-V compute nodes in CloudSystem, follow these recommendations for optimal usage of Hyper-V compute resources.

- CloudSystem supports a maximum of 40 instances deployed on a single Hyper-V compute node. For better managing of Hyper-V compute resources with density of instances supported per Hyper-V compute:
  - a. Create multiple “Host aggregates” with Hyper-V compute nodes based on compute capacity.
  - b. Use the best rated “Availability zones” to launch instances based on the image size for effective usage of Hyper-V compute resources.
- If the Hyper-V compute node is part of cluster, the instances are deployed to Clustered Shared Volumes. Extend the existing Clustered Shared volume with more disk space instead of creating separate Clustered Shared Volumes. Creating separate CSVs consumes disk space for the virtual machine images deployed on it.
- If you have a mix of activated Hyper-V standalone and clustered compute nodes and all of your deployed instances need to be highly available, create a Host aggregate with all of the Hyper-V compute nodes in the cluster and use an availability zone to launch instances.

## Procedures to note with Hyper-V compute nodes in Cluster

### Procedure 1: Add a new Hyper-V compute node to an existing Microsoft Failover Cluster

CloudSystem requires that you activate any new Hyper-V compute nodes added to an existing Microsoft failover cluster, otherwise any VM instances migrated to the newly added Hyper-V host in the cluster will not be managed by CloudSystem.

**NOTE:** Existing instances created with volumes in the cluster cannot migrate to the newly added Hyper-V host in the cluster. The Microsoft Failover Cluster does not migrate instances to a new Hyper-V host as existing instance volumes are not presented to it.

### Procedure 2: Add an activated Hyper-V compute node to a Microsoft Failover Cluster

To add an activated Hyper-V compute node to a Microsoft failover Cluster:

- Deactivate the standalone Hyper-V compute node.
- Add the compute node to the cluster, then activate it from CloudSystem to deploy it as a cluster compute node.

Note: Existing instances on the compute node will continue to operate if you do not change the compute node host name. However, existing instances on the compute node will not be highly available.

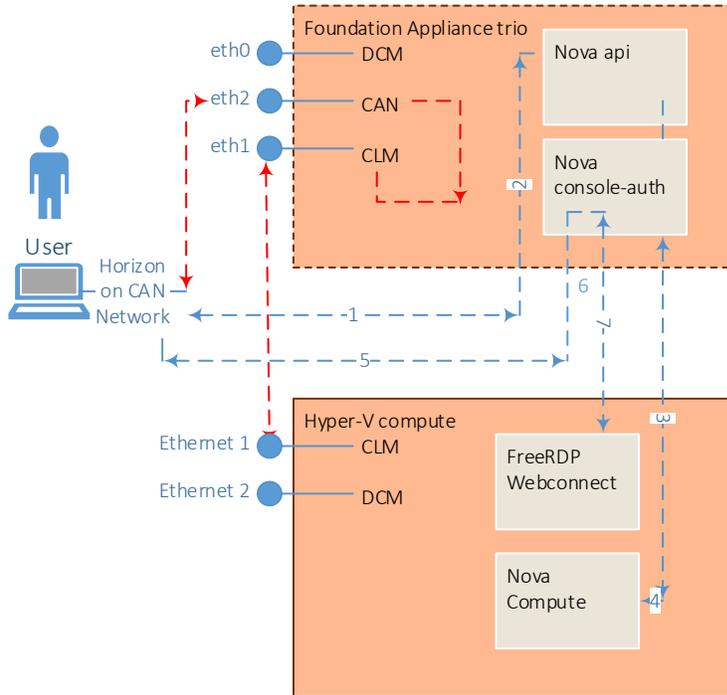
### IMPORTANT

- Hyper-V Cluster compute nodes return the total disk space as an aggregate of all Cluster shared volumes created on the Hyper-V compute node. The “total disk space” statistics will be the same across all Hyper-V compute nodes in the cluster.
- CloudSystem synchronizes instances migrated across Hyper-V cluster compute nodes by Microsoft Failover Cluster every two minutes, so that the migrated instances can be managed in CloudSystem.

## Deactivating Hyper-V compute nodes

Hyper-V compute nodes can be deactivated from the CloudSystem Operations Console when there are no instances provisioned by CloudSystem on the compute node. If you want to decommission the Hyper-V compute host, first deactivate the Hyper-V compute node in CloudSystem and then delete the host.

## Hyper-V Remote console support using FreeRDP-Webconnect



Hyper-V compute nodes use the FreeRDP-WebConnect solution from CloudBase Solutions to enable RDP access to VM instances launched from the OpenStack user portal.

In the diagram above, red lines show how the instance console accessed using the Consumer Access Network is routed to the Hyper-V host to access the RDP Console. Following is the process involved

- User accesses the OpenStack user portal using the Consumer Access Network.
- When the instance console is accessed for the Hyper-V VM, the request reaches the Consumer Access Network Interface on the Cloud controller.
- The request on the Consumer Access Network is NATed to the Cloud Management Network and is routed to the correct Hyper-V compute node on which the VM instance resides.

Blue lines show how the instance console is accessed using the FreeRDP-Webconnect service on the Hyper-V compute node.

- User accesses the OpenStack user portal using the Consumer Access Network
- The request reaches the Nova API, and then to nova-console auth services on the Cloud controller.

- Nova-console auth service creates authentication token for accessing the VM instance console and stores it in cache.
- The request reaches the Nova-compute running on the Hyper-V host which returns the URL on which FreeRDP-Webconnect can expose the VM console.
- The OpenStack user portal queries the nova-console auth service to authenticate the token for accessing the VM console and connects to Hyper-VM console.

## Monitoring

CloudSystem uses the HP open source Monasca solution for compute node monitoring ESXi and KVM compute nodes. In CloudSystem 9.0, monitoring of Hyper-V compute resources is not supported. The CloudSystem Operations Console does not show Hyper-V compute resource metrics.

## Summary and limitations of OpenStack operations

The following OpenStack operations differ from what is generally described in OpenStack documentation.

### Volume Attach/Detach

Hyper-V cluster compute nodes limitations in supporting HP 3PAR/VSA volumes attach/detach operations:

- CloudSystem does not support boot from volume operation for instances spawned on Hyper-V cluster compute nodes. This is limitation from Microsoft Failover Clustering feature in Windows.
- If one of the Hyper-V compute nodes in a cluster is powered down, CloudSystem recommends that you do not attach volumes to instances spawned on Hyper-V cluster compute nodes as the volumes attached will not be presented to the powered down host.
- Attach/detach volumes are not supported if instances are in a running state in Hyper-V compute cluster. Instances should be powered off before you can attach or detach a volume. This is a limitation in the Failover Clustering feature in Hyper-V.

### Migrate/Live Migrate

Live migrate using the OpenStack Nova CLI/API is not supported on Hyper-V compute nodes. However, VMs can be migrated out of band from a Microsoft Failover Cluster either manually or automatically via cluster.

### Interface Attach/Detach

Additional network interfaces cannot be attached to an instance after it is created.

### Learn more about HP Helion CloudSystem:

<http://www.hp.com/go/CloudSystem>

<http://www.hp.com/go/CloudSystem/docs>

---

© Copyright 2014–2015 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.