

# **Citrix CloudPlatform (powered by Apache CloudStack) Version 4.3.0.3 Release Notes**

Revised September 30, 2015 01:30 PM IST



## **Citrix CloudPlatform (powered by Apache CloudStack) Version 4.3.0.3 Release Notes**

**Revised September 30, 2015 01:30 PM IST**

© 2015 Citrix Systems, Inc. All rights reserved. Specifications are subject to change without notice. Citrix Systems, Inc., the Citrix logo, Citrix XenServer, Citrix XenCenter, and CloudPlatform are trademarks or registered trademarks of Citrix Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders.

Release notes for Citrix CloudPlatform version 4.3.0.3.

---

<b>1. Submitting Feedback and Getting Help</b>	<b>1</b>
<b>2. Support Matrix</b>	<b>3</b>
2.1. Supported OS Versions for Management Server .....	3
2.2. Supported Hypervisor Versions .....	3
2.3. Supported External Devices .....	4
2.4. System VM Templates .....	4
2.5. Supported Browsers .....	7
<b>3. What's New in 4.3.0.3</b>	<b>9</b>
3.1. Fixed Issues .....	9
3.2. Known Issues .....	27
3.3. CloudPlatform 4.3.0.3 Release Notes Addendum .....	28
3.3.1. Limitation on Dynamically Assigning vCPUs on Windows VMs .....	28
3.3.2. Providing a Display Name to the Guest Virtual Machines .....	28
3.3.3. Preventing the Logging of Duplicate VM.START Usage Events .....	29
<b>4. Upgrade Instructions</b>	<b>31</b>
4.1. Upgrade from 4.3.x.x to 4.3.0.3 .....	31
4.2. Upgrade from 4.2.x to 4.3.0.3 .....	41
4.3. Upgrade from 3.0.x to 4.3.0.3 .....	51
4.4. Upgrade CloudPlatform Baremetal Agent on PXE and DHCP Servers .....	63
4.5. Updating SystemVM.ISO .....	64
4.6. Upgrading and Hotfixing XenServer Hypervisor Hosts .....	64
4.6.1. Upgrading to a New XenServer Version .....	65
4.6.2. Applying Hotfixes to a XenServer Cluster .....	66
4.6.3. Install CloudPlatform XenServer Support Package (CSP) .....	68
4.6.4. Upgrading to XenServer 6.2 SP1 Hotfix XS62ESP1005 .....	69



# Submitting Feedback and Getting Help

The support team is available to help customers plan and execute their installations. To contact the support team, log in to [the Support Portal](#)<sup>1</sup> by using the account credentials you received when you purchased your support contract.

---

<sup>1</sup> <http://support.citrix.com/cms/kc/cloud-home/>



# Support Matrix

This section describes the operating systems, browsers, and hypervisors that have been newly tested and certified compatible with CloudPlatform 4.3.0.3. Most earlier OS and hypervisor versions are also still supported for use with 4.3.0.3. For a complete list, see the System Requirements section of the CloudPlatform 4.3 Installation Guide.

## 2.1. Supported OS Versions for Management Server

- RHEL versions 5.10, 6.2, 6.3, 6.4, and 6.5
- CentOS versions 5.10, 6.2, 6.3, 6.4 and 6.5

## 2.2. Supported Hypervisor Versions

The following new hypervisor support has been added:

- Windows Server 2012 R2 (with Hyper-V Role enabled)
- Hyper-V Server 2012 R2
- XenServer version 6.2 SP1 Hotfix XS62ESP1005
- XenServer version 6.2 SP1 Hotfix XS62ESP1004
- XenServer version 6.2 SP1 Hotfix XS62ESP1003
- XenServer version 6.2 Hotfix ESP1015
- VMware vCenter 5.5 Update 1b
- VMware vCenter 5.1 Update 2a
- VMware vCenter 5.0 Update 3a

Other supported hypervisors for CloudPlatform:

- XenServer versions 5.6 SP2 with latest hotfixes.
- XenServer versions 6.0.2 with latest hotfixes (for CloudPlatform 3.0.2 and greater)
- XenServer versions 6.0 with latest hotfixes (for CloudPlatform 3.0.0 and greater)
- XenServer versions 6.1 with latest hotfixes.
- KVM versions 6.2 and 6.3
- Bare metal hosts are supported, which have no hypervisor. These hosts can run the following operating systems:
  - RHEL or CentOS, v6.2 or 6.3



**Note**

Use libvirt version 0.9.10 for CentOS 6.3

- Fedora 17
- Ubuntu 12.04

For more information, see the Hypervisor Compatibility Matrix in the CloudPlatform Installation Guide.

### 2.3. Supported External Devices

- NetScaler MPX versions 9.3, 10.0.e, 10.1.e, and 10.5
- Netscaler VPX versions 9.3, 10.0.e, 10.1.e, and 10.5
- Netscaler SDX version 9.3
- SRX (Model srx100b) versions 10.3 to 10.4 R7.5
- F5 11.X

### 2.4. System VM Templates

CloudPlatform 4.3.0.3 supports 64-bit System VM templates. This release does not provide 32-bit support for System VM templates. For the latest System VM fixes, follow the procedure given in [Upgrading System VM Template without Upgrading Management Server in CloudPlatform](#)<sup>1</sup>.

Hypervisor	Description
XenServer	Name: systemvm-xenserver-4.3  Description: systemvm-xenserver-4.3  URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-xen.vhd.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-xen.vhd.bz2</a>  Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, select each zone and individually register the template to make the template available in all the XenServer zones.  (Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform

<sup>1</sup> <http://support.citrix.com/article/CTX200024>

Hypervisor	Description
	<p>deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: XenServer</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
KVM	<p>Name: systemvm-kvm-4.3</p> <p>Description: systemvm-kvm-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-kvm.qcow2.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-kvm.qcow2.bz2</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, select each zone and individually register the template to make the template available in all the zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: KVM</p> <p>Format: QCOW2</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>

Hypervisor	Description
VMware	<p>Name: systemvm-vmware-4.3</p> <p>Description: systemvm-vmware-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-vmware.ova">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-vmware.ova</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, select each zone and individually register the template to make the template available in all the zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: VMware</p> <p>Format: OVA</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
Hyper-V (Applicable only for 4.3)	<p>Name: systemvm-hyperv-4.3</p> <p>Description: systemvm-hyperv-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2</a><sup>2</sup></p> <p>Hypervisor: Hyper-V</p> <p>Format: VHD</p>

<sup>2</sup> <http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2>

Hypervisor	Description
	OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)  Extractable: no  Password Enabled: no  Public: no  Featured: no

## 2.5. Supported Browsers

- Internet Explorer versions 10 and 11
- Firefox versions 33.0.2
- Google Chrome versions 38.0.2125.111 m
- Safari 7.1 (Mac)



## What's New in 4.3.0.3

CloudPlatform 4.3.0.3 release focusses on resolving critical defects to enhance the quality of the product. This release includes no new features or API changes.

### 3.1. Fixed Issues

Issue ID	Description
CS-44022	<p><b>Problem:</b> VM usage is stopped after out of band migration. The VMSync job does not provide the start event.</p> <p><b>Root cause:</b> If the old state of a VM is Started and the new state is Running, the VM is in the Started state. In case of migration through Vmotion, the VM state goes from Stopped to Running.</p> <p><b>Solution:</b> Add another conditional check for VmStarted with the old state as either Starting or Stopped and the new state as Running.</p>
CS-44018	<p><b>Problem:</b> Usage job generates usage even if there is an exception while processing usage events. This corrupts the usage data.</p> <p><b>Root cause:</b> The exception that is raised while processing usage events gets suppressed inside the catch block.</p> <p><b>Solution:</b> Removed the catch block that suppresses the exception.</p>
CS-43788	<p><b>Problem:</b> A template that is copied to a new zone stopped charging after the original template is removed.</p> <p><b>Root cause:</b> This is due to incorrect query, which does not have the zone filter.</p> <p><b>Solution:</b> Changes has been made to add zone filter while processing the TEMPLATE.DELETE event.</p>
CS-43738	<p><b>Problem:</b> Usage for an uploaded volume is not collected.</p> <p><b>Root cause:</b> Volume create event is not getting generated when an uploaded volume is attached to a VM.</p> <p><b>Solution:</b> Added the volume create event when a volume is moved from secondary to primary storage.</p>
CS-43468	<p><b>Problem:</b> Incorrect calculation of the allocated primary storage.</p>

Issue ID	Description
	<p>Root cause: While calculating the primary storage allocated, CloudPlatform does not account the size of the template. Because of this primary storage allocated is wrongly calculated. This may result in storage allocator assigning the storage, which is already full.</p> <p>Solution: Added populating template size while template registration and allocated storage calculations.</p>
CS-43461	<p>Problem: Static NAT configured on VPC VR on additional public subnet IP address, the traffic from the static NAT VM goes out through default gateway instead of additional public subnet interface. This issue is seen only in VPC case.</p> <p>Root cause: When static NAT is configured on the VPC VR, no proper IP table rules or IP routes are configured for the marked packets. Because of this, traffic originated from the VM is not routed via the correct interfaces.</p> <p>Solution: Fixed the IP table rules or IP routes for the marked packets to address this problem.</p>
CS-43432	<p>Problem: Calling disabling static NAT API within a small interval (one to three seconds) on multiple IP addresses leaves the disabled static NAT IP addresses on the VR public interface.</p> <p>Root cause: In static NAT implementation, when disable API is called on one IP address, all the IP addresses of the subnet are send to VR to delete this IP address and add the remaining IP addresses. There is no synchronisation for these delete operation due to this last disable operation delete one IP address and add remaining IP addresses. This leads to the IP address on the VR interface, but in the Released state in the CloudPlatform database.</p> <p>Solution: Added a lock for this concurrent disabled static NAT.</p>
CS-43256	<p>Problem: After restarting Management Server, the maintenance mode of the local storage gets disabled.</p> <p>Root cause: On starting Management Server, while performing the host attach to the storage pool, there was no check for pool status with maintenance.</p>

Issue ID	Description
	<p>Solution: If in maintenance mode, added a conditional check while performing the host attach to the storage pool.</p>
CS-42970	<p>Problem: Datastore file search operations are taking a longer time to complete due to large number of empty folders, which are created to store only data volumes.</p> <p>Root cause: During VM lifecycle, CloudPlatform deletes VM folder by destroying VM. But this destroy API call does not delete VM folders on datastores other than the location where the VM resides. A folder gets created when a volume gets live migrated to a datastore or when a volume was created on datastore, through storage tag, which is different from where VM is present.</p> <p>Solution: Move out all the files associated with the volume before deleting the VM folder on datastore that contains the volume as part of detaching volume. Delete would be attempted only if folder is empty. This ensures avoiding accumulation of folders over time.</p>
CS-42832	<p>Problem: When a VM snapshot is created with memory, its size is calculated incorrectly.</p> <p>Root cause: Memory snapshot is not fetched while computing the VM snapshot size.</p> <p>Solution: Modifications have been made to fetch and add memory snapshots while computing the VM snapshot size.</p>
CS-42641	<p>Problem: The size of detached volumes are not considered while computing resource count.</p> <p>Root cause: This is due to incorrect query that does not fetch the detached volumes.</p> <p>Solution: Modifications have been made to the resource count query to consider detached volumes.</p>
CS-42524	<p>Problem: Due to an issue in CloudPlatform version 3.0.7, even though the <code>vmware.root.disk.controller</code> parameter set to "ide", CloudPlatform creates user instances with root volume on the SCSI controller. This causes root volume to move from SCSI to IDE controller after the upgrade results in boot failure.</p>

Issue ID	Description
	<p>Root cause: After upgrading from CloudPlatform version 3.0.7, as this issue has already been fixed in 4.x release, a restore VM operation should result in moving the root volume from the SCSI controller's virtual node 0 to the IDE controller. Guest OS looks at SCSI controller's virtual node 0, which has a data volume instead of root volume before the restore VM operation. CloudPlatform would have retrieved controller information from the root volume, however, restore VM deletes existing root volume and creates a new root volume, which does not have any chain_info. That is, relying on the parameter that has been set to "ide" results in moving root volume from SCSI to IDE.</p> <p>Solution: Before deleting root volume as part of the restore VM operation, persist device bus information for root volume if it does not exist for the VM. This can occur if CloudPlatform has been upgraded from versions before 4.2.1-3.</p>
CS-42523	<p>Problem: CloudPlatform fails to resize the Windows XP data volume.</p> <p>Root cause: Resize for the IDE virtual disks is not supported in VMware. As Windows XP uses IDE controller for volumes, the resize volume fails.</p> <p>Solution: Resize operation for the IDE virtual disks has been denied instead of failing the operation. The users can view an error message that suggests them to detach the volume, perform the resize operation, and, then, attach the volume.</p>
CS-42179	<p>Problem: Snapshot creation events missing for start and completion states. Only scheduled was reported.</p> <p>Root cause: All the events were not being generated for snapshot creation.</p> <p>Solution: Report the snapshot create event for Start and completion states.</p>
CS-42079	<p>Problem: Instances fail to start after unsuccessful compute offering upgrade.</p> <p>Root cause: In case of unsuccessful compute offering upgrade, the rollback does not take care of restoring the old service offering details.</p> <p>Solution: Code has been modified to restore the old service offering details.</p>

Issue ID	Description
CS-41525	<p>Problem: The <code>lsf</code> command inside SSVM reports "Cannot identify protocol".</p> <p>Root cause: Cloud service does not close the socket in case the SSVM cannot reach the Management Server.</p> <p>Solution: Code has been modified to close the socket properly in case of any exception.</p>
CS-41245	<p>Problem: Windows VMs does not unlock the screen when using Ctrl+Alt+Del button in the console window in IE version 9,10,and 11. This occurs only on Windows server 2012.</p> <p>Root cause: This issue occurs because the console loses focus when clicking Ctrl+Alt+Delete button combination in Internet Explorer. Again, the Ctrl+Alt+Delete button combination is not visible when focussing on the console by clicking anywhere in console window. This issue can be seen when the whole console does not fit into browser window. That is VM resolution is greater than 800x600 and scrollbar gets created to fit the console in browser window.</p> <p>Ctrl+Alt+Delete button combination works fine if the windows VM resolution is 800x600 or less. One more way to overcome this issue is by allowing user to resize window such that the whole console fits in the browser window without any scrollbar.</p> <p>Solution: Make the console window in IE resizable so that user can increase its size to fit the console without the need of scrollbar.</p>
CS-40944	<p>Problem: Several hosts get disconnected because the <code>DirectAgent</code> thread pool gets exhausted in the Management Server.</p> <p>Root cause: CloudPlatform uses the <code>DirectAgent</code> threads to perform the ping task along with running several agent commands and cron jobs. The Management Server pings the hosts that are connected to it at regular intervals (<code>ping.interval</code>). This is to monitor the status of the hosts. If the Management Server does not have free <code>DirectAgent</code> threads to perform the ping task, it cannot ping the hosts and eventually the hosts will be marked as disconnected.</p> <p>From the customer's logs, it is understood that the <code>Checks2SVpnConnectionsCommand</code>, which runs regularly against a virtual router</p>

Issue ID	Description
	<p>(<code>router.check.interval</code>) to check the validity of VPN connections, did not return the result to the Management Server. Because of this, the <code>DirectAgent</code> thread that executed the <code>Checks2SvpnConnectionsCommand</code> is blocked indefinitely and could not pick up any other request. As this command is designed to execute in sequence on a host, every execution of the command thereafter on that particular host ends up picking up a <code>DirectAgent</code> thread and waiting for the previous execution to complete. As a result, the host ends up using and blocking all the <code>DirectAgent</code> threads indefinitely.</p> <p>Solution: Ensure that CloudPlatform does not wait indefinitely for a script to execute on the virtual router by handling channel condition of EOF.</p>
CS-40942	<p>Problem: Volume snapshot operation fails on a particular volume.</p> <p>Root cause: In VMWare, creating a volume snapshot involves the following:</p> <ol style="list-style-type: none"> <li>1. Create a VM snapshot for the VM that the volume is attached to.</li> <li>2. Create a VM (clone) from the snapshot disk chain of the volume fro which the users want to create a snapshot.</li> <li>3. Export the VM created in step 2 to secondary storage.</li> <li>4. Destroy the cloned VM.</li> <li>5. Delete the VM snapshot taken in step 1 (This step involves VM consolidation in vCenter).</li> </ol> <p>The session between CloudPlatform and vCenter times out in 20 minutes. If any of the above steps (except VM export) takes more than 20 minutes to complete, snapshot backup operation fails and the snapshot is moved to the 'Error' state.</p> <p>Solution: Split the existing vCenter session timeout configuration into two.</p> <ul style="list-style-type: none"> <li>• To tune timeout value for CloudPlatform's session vCenter during snapshot backup operation - <code>vmware.snapshot.backup.session.timeout</code> (new; default: 20 minutes)</li> </ul>

Issue ID	Description
	<ul style="list-style-type: none"> <li>For all other operations - <code>vmware.vcenter.session.timeout</code> (existing; default: 20 minutes)</li> </ul> <p>Setting a high timeout value for the snapshot backup operation does not apply to vCenter connections made for any other operations.</p>
CS-40215	<p>Problem: CloudPlatform async call registers the failure of the <code>RevertToVMSnapshot</code> command as successful.</p> <p>Root cause: If the <code>RevertToVMSnapshot</code> command fails, CloudPlatform async call registers it as successful. This is because CloudPlatform never returns the call as unsuccessful even if it is a failure.</p> <p>Solution: If the <code>RevertToVMSnapshot</code> command fails, CloudPlatform returns async call as failure.</p>
CS-40164	<p>Problem: The CloudPlatform UI did not allow Domain administrators creating public templates from snapshots, because of the missing pre-filter in JavaScript.</p> <p>Solution: To resolve this issue, added the pre-filter in the JavaScript.</p>
CS-39957	<p>Problem: Unable to create snapshot of a volume.</p> <p>Root cause: Snapshot operation on Datadisk volumes DATA-20 and DATA-21 fails at the Management Server layer. This failure is not related to NPEs being thrown by SSVM.</p> <p>Each of the above mentioned DATA volumes has a snapshot that has inconsistent entries in the snapshots and snapshot_store_ref tables.</p> <p>There is a snapshot for each of these volumes that is marked as removed in snapshots table, but in snapshot_store_ref table it appears as Ready for Primary storage and Creating for the Secondary storage. Because of this mismatch, CloudPlatform throws an NPE while trying to retrieve parent snapshot details during the snapshot operation.</p> <p>Solution: For both disks, remove the snapshots entries in the snapshot_store_ref table that are inconsistent. For example,</p> <pre data-bbox="858 1937 1439 2027">delete from snapshot_store_ref where snapshot_id in (34, 35);</pre>

Issue ID	Description
CS-39478	<p>Problem: CloudPlatform UI does not show all the domains.</p> <p>Root cause: Due to an incorrect CSS property, if there are more than 25 domains, the last domain gets cut off from the UI.</p> <p>Solution: Fixed the CSS property such that even the last domain is shown properly on the UI.</p>
CS-39349	<p>Problem: The size of the volume differs from the value specified by the user.</p> <p>Root cause: Volume gets created with the size mentioned in the custom disk offering instead of the size specified by the user. For customized offering, do not expect the size of the volume to be available with the offering. If it is available with the offering, use it instead of the size of the volume specified by the user.</p> <p>Solution: Do not allow the custom disk size disk offering creation with disk size.</p>
CS-38920	<p>Problem: Management Server does not provide proper error message while deleting a default ACL.</p> <p>Root cause: ACL type should have been verified before evaluating access.</p> <p>Solution: Code changes has been done to verify ACL type and provide appropriate exception.</p>
CS-38850	<p>Problem: In CloudPlatform, the Domain Administrators can create public template from a snapshot, however, they cannot register a public template.</p> <p>Root cause: CloudPlatform allows domain administrators to create a public template from a snapshot, however, they cannot register the public template. This is a conflicting behaviour.</p> <p>Solution: Restrict the domain administrators from creating public template from a snapshot. Now, domain administrators will not be able to create public templates from snapshots. The behaviour is now consistent.</p>
CS-38823	<p>Problem: In CloudPlatform version 4.3.0.2, hosts without tag are not listed while listing the hosts for migration for instance with tag.</p> <p>Root cause: This is because of an issue with the code. While preparing the suitable hosts, CloudPlatform accidentally removes</p>

Issue ID	Description
	<p>the incompatible (host without tag) hosts from <code>otherhost</code> list (incorrect use of <code>List.retainAll()</code>. <code>hosts.retainAll()</code> retains all hosts with tags. So, if the <code>listByHostTag</code> returns an empty list, hosts list remains empty. This, in turn, makes the <code>otherHosts</code> list empty.</p> <p>Solution: Do not modify the hosts list using <code>retainAll</code>. Refer to the following <code>allocateTo</code> method:</p> <pre data-bbox="858 645 1439 815">public List&lt;Host&gt; allocateTo(VirtualMachineProfile vmProfile, DeploymentPlan plan, Type type, ExcludeList avoid, int returnUpTo,boolean considerReservedCapacity)</pre>
CS-38778	<p>Problem: When migrating to another storage, the resource tags on a disk are lost.</p> <p>Root cause: During cold volume migration, CloudPlatform duplicates volume entry in volumes table. When migration is complete, CloudPlatform updates the uuid of new entry and expunge the older entry. This results in removal of resource tags on volume as its resource id still pointing to older volume.</p> <p>Solution: While updating uuid for volume, CloudPlatform should also update <code>resource_id</code> for the tags.</p>
CS-38551	<p>Problem: Unable to delete IP tag.</p> <p>Root cause: Due to wrong search string, CloudPlatform fetches a resource that does not belong to the current account and domain. This results in permission exception.</p> <p>Solution: Code has been modified to fix the incorrect search string.</p>
CS-38341	<p>Problem: Do not allow VM snapshots and Volume snapshots related operations simultaneously.</p> <p>Root cause: Allowing operations documented in the Restriction on VM with VM snapshots section at <a href="https://cwiki.apache.org/confluence/display/CLOUDSTACK/VM+Snapshots">https://cwiki.apache.org/confluence/display/CLOUDSTACK/VM+Snapshots</a> could result in unpredictable code path execution. This could result in disk corruption.</p> <p>Solution: Do not allow the operations on VM with VM snapshots that would result in unpredictable</p>

Issue ID	Description
	<p>results. These changes have been ported from Goleta to 4.3.0.2.</p>
<p>CS-38219</p>	<p>Problem: Creating shared network fails when the network CIDR matches the existing guest CIDR in a zone.</p> <p>Root cause: The issue occurs because CloudPlatform on creating a shared network checks if there is any guest network already implemented with same CIDR and throws the error without checking if they have the same VLANs. Creating same CIDR shared network with different VLAN should be allowed.</p> <p>Solution: When creating a shared network, if CloudPlatform identifies an existing guest network with the same CIDR, CloudPlatform checks if they have the same VLAN. If they have same VLAN, CloudPlatform do not allow creating it. If they are same CIDR with different VLANs, CloudPlatform allows creating the network.</p>
<p>CS-38169</p>	<p>Problem: The ListUsageRecords API returns null usage_id for migrated volumes.</p> <p>Root cause: When a volume is migrated, a new entry is created in the volumes table and the old volume entry is marked as removed. Usage events for new volume creation are not logged.</p> <p>Solution: Moved volume usage events to be in sync with volume state machine. Correct volume usage events are logged during migration.</p>
<p>CS-34553</p>	<p>Problem: CloudPlatform does not accept the value set to the global configuration parameter: <code>system.vm.default.hypervisor</code>, which indicates the hypervisor type for System VMs.</p> <p>Root cause: CloudPlatform does not accept the value set to the global configuration parameter: <code>system.vm.default.hypervisor</code>. Instead, CloudPlatform accepts the hypervisor type of User VM's destination host.</p> <p>Solution: Add the hypervisor type of System VMs that is provided as the value for the global configuration parameter:<code>system.vm.default.hypervisor</code> to the list of supported hypervisors at zeroth index. Then, add the hypervisor type of the user VM's destination host. This will enable CloudPlatform to first deploy systemVM's</p>

Issue ID	Description
	hypervisor type before it deploys the hypervisor type of the userVM's destination host.
CS-34531	<p>Problem: In the Project view, site-to-site VPN gateways are not listed within a VPC.</p> <p>Root cause: The projectId is not considered during ACL check while listing S2S VPN gateways. This returns empty response.</p> <p>Solution: Included projectId during ACL check while listing S2S VPN gateways.</p>
CS-33144	<p>Problem: SSL v3 does not get disabled on Console Proxy VM (CPVM).</p> <p>Solution: Initializing <code>SSLContext.getInstance("TLS")</code> is not enough. You must disable the SSLv3.</p>
CS-33010	<p>Problem: Taking volume snapshots works as expected on VMware.</p> <p>Root cause: Inconsistencies between the snapshots table and the snapshot_store_ref table cause null pointer exception on the Management Server while taking snapshots, resulting in failure.</p> <p>Solution: This situation is now handled better to accommodate all scenarios.</p>
CS-33008	<p>Problem: Usage collection is not stopped for some of the deleted VMs.</p> <p>Root cause: When VM stopped directly from the Hypervisor, the VM synchronization operation tried to synchronize the state to Stopped in CloudPlatform. In this process the VM synchronization operation does not publish the usage events for the stopped VMs.</p> <p>Solution: Fixed this problem by publishing the usage even when the VM is stopped via VM synchronization.</p>
CS-32840	<p>Problem: The listVolumes API fails for a specific domain with NPE.</p> <p>Root cause: This occurs when the volume associated VM instance has NULL or invalid state.</p> <p>Solution: Modify the code to guard this situation because it should not block volume listing.</p>

Issue ID	Description
CS-32557	<p>Problem: Agent LB fails to rebalance some of the hosts. These hosts get stuck in the op_host_transfer queue forever.</p> <p>Root cause: Hosts that are not being rebalanced are not routing hosts. Non-routing hosts such as SSVM and CPVM are picked up as part of the rebalance process. The transfer of these hosts get stuck in the op_host_transfer queue.</p> <p>Solution: Ensure that only the routing hosts are picked up as part of agent rebalance.</p>
CS-31790	<p>Problem: VOLUME.DELETE usage event missing for the VMs in the ERROR state.</p> <p>Root cause: If VM fails while deploying it, CloudPlatform updates the database entry for this VM as ERROR and marks the volume entry for this VM as Destroyed. However, CloudPlatform does not publish the VOLUME.EVENT event while invoking internal destroyVolume command. Instead, it scatters the VOLUME.EVENT publish code around in the service layer code. Because CloudPlatform fails in publishing VOLUME.EVENT, no VOLUME.EVENT is published when a VM deployment fails</p> <p>Solution: VOLUME.DELETE usage event now available for VMs in ERROR state.</p>
CS-30961	<p>Problem: Planner allocates new storage pools for Ready data disks within the same cluster. This triggers data migration that can be avoided.</p> <p>Root cause: StoragePool reuse logic starts only if ROOT volume is Ready. However, for restoreVM, ROOT is re-created. So this logic does not get invoked for restoreVM.</p> <p>Solution: Reuse the storage pools for any Ready disk if the pool fits the deployment plan (dc/pod/cluster).</p>
CS-30247	<p>Problem: When trying to launch a VM on a network that uses the network offering with the Loadbalancer that F5 provides, errors are observed in the logs and the VM creation ultimately fails.</p> <p>Root cause: The listVLAN response of F5 BigIP has been modified in 11.x.All VLANs are prefixed with the partition/path information.</p>

Issue ID	Description
	<p>Solution: The <code>listVlan</code> response handler has been modified to remove path information.</p>
CS-30246	<p>Problem: When trying to launch a VM on a network that uses the network offering with LB provided by F5, errors are observed in the logs and the VM creation operation ultimately fails.</p> <p>Root cause: The <code>listVLAN</code> response of F5 BigIP has been modified in 11.x. All VLANs are prefixed with the partition/path information.</p> <p>Solution: The <code>listVLAN</code> response handler has been modified to remove the path information.</p>
CS-30096	<p>Problem: VM are stopped minutes after they are started in CloudPlatform.</p> <p>Root cause: HA work, which is "Done", is scheduled for some HA jobs. This retries infinitely.</p> <p>Solution: You do not schedule HA work, which is "Done. By default, it will retry five times. You can change the global configuration parameter:<code>max.retries</code> to set a different retry time.</p>
CS-29627	<p>Problem: Deployment planner goes in infinite loop while creating a VM with a tagged local data disk and the pool is not tagged.</p> <p>Root cause: Creating a VM with two disks - one on shared storage and the other one on local storage - and one storage is full (cannot be allocated) and other storage is not full does not place the cluster in the avoid list. So, CloudPlatform checks for deployment in that cluster again and again and this results in an infinite loop.</p> <p>Solution: Ensure that the cluster is placed in the avoid list.</p>
CS-28607	<p>Problem: From the <b>Domains</b> tab, unable to add the Admin type user.</p> <p>Root cause: Incorrect declaration and initialization of variable in the JavaScript.</p> <p>Solution: Modifications have been made to declare and initialize the JavaScript variable.</p>
CS-27149	<p>Problem: From remote access VPN client, not able to resolve the DNS name of the VMs in the guest network.</p>

Issue ID	Description
	<p>Root cause: Dnsmasq is not listening on the ppp + interface. Due to this, the DNS queries are not answered.</p> <p>Solution: Updated the Dnsmasq to listen on all the interfaces except ethic and ether.</p>
CS-26200	<p>Problem: Disk space is getting full because the <code>daemon.log</code> file gets huge.</p> <p>Root cause: The problem is with the <code>contrackd</code> package. This daemon gets started in the VR. By default, the stats logging gets enabled. The stats logging by this daemon fills up the VR space.</p> <p>Solution: Do the following:</p> <ol style="list-style-type: none"> <li>1. Disable the option at <code>/etc/contrackd.conf</code> file to switch off the <code>contrackd</code> stats logging.</li> <li>2. Do not start the <code>contrackd</code> daemon on Non-RVR.</li> </ol>
CS-26147	<p>Problem: The metering for IP range, which is tied to an account, was not stopped even after range was deleted.</p> <p>Root cause: This occurs in the following scenario:</p> <ol style="list-style-type: none"> <li>1. Add an IP range and add account to it. This will add usage events for the metering with event <code>NET.IPASSIGN</code></li> <li>2. Do not allocate any IP addresses from this range to that account.</li> <li>3. Delete the IP range. This will not publish any usage events for <code>IP RELEASE</code>.</li> </ol> <p>Solution: In the code earlier we were publishing usage events only when there are allocated IPs in that range. (NEED TO REWORD.)</p>
CS-25880	<p>Problem: After issuing a stop command on a Windows guest VM running on KVM host, the operation in CloudPlatform failed, however, VM eventually was stopped by the agent leaving the state in cloud database as running and on the host as stopped (does not show in <code>virsh list -all</code> output).</p> <p>Root cause: Logs show that the stop command has been failed due to time out. Because of this, CloudPlatform reverts its state from Stopping to Running with the <code>host_id</code> set to 1. CloudPlatform was successfully stopped on the agent. After a</p>

Issue ID	Description
	<p>while, when VMsync runs the missing VM report, it detects that the VM's power report from agent is missing and marks the VM's status as Stopped in the database.</p> <p>Solution: Updated the VM Synchronization process to take this scenario into account.</p>
CS-25423	<p>Problem: After a volume is migrated, the usage table shows the old volume id.</p> <p>Root cause: After migrating a volume (does not apply for Live Migration), CloudPlatform destroys old volume and creates the new volume with the same details. As part of this process, it should publish Usage Event for the old volume removal and should publish Volume Create event for creating the new volume. However, this was not happening and the usage table was displaying the old volume id.</p> <p>Solution: Fixed this by using Volume State Machine to publish the events instead of publishing at operation level.</p>
CS-25256	<p>Problem: The Event tab shows the failed APIs as successful.</p> <p>Root cause: Set the default value of success to <code>true</code> instead of <code>false</code> in the <code>ActionEventInterceptor.invoke()</code> method. This causes ActionEventInterceptor to log event as successful even in case of failure.</p> <p>Solution: Place the default value of success as <code>false</code>.</p>
CS-22841	<p>Problem: CloudPlatform fails to attach volume to a VM.</p> <p>Root cause: To attach a volume to a VM, CloudPlatform creates the volume in primary storage if it has not already been created. During this, CloudPlatform tries to select a storage pool to deploy the volume.</p> <p>This storage pool selection fails in case of the following conditions:</p> <ul style="list-style-type: none"> <li>• ROOT volume of the VM to which CloudPlatform wants to attach a data disk is in a zone-wide primary storage.</li> <li>• The <code>vm.allocation.algorithm</code> global configuration parameter has been set to <code>userdispersing</code>.</li> </ul>

Issue ID	Description
	<p>Failure occurs because CloudPlatform's cluster-wide storage pool allocator sends the VM's pod ID instead of sending the storage pool's pod ID (NULL), which further results in an NPE.</p> <p>Solution: Ensure that CloudPlatform sends the storage pool's pod ID instead of sending the VM's pod ID to the storage pool allocators.</p>
CS-22502	<p>Problem: Not able to remove NIC from VM.</p> <p>Root cause: While removing NIC, the query to check for PF rules only considers destination IP. It is possible to have same destination IPs in different networks.</p> <p>Solution: Include network id also as a part of the filter in the query. This will remove the IPs outside NIC's network to be considered.</p>
CS-21347	<p>Problem: After rebooting CloudPlatform, a delay is noticed in connecting hosts.</p> <p>Root cause: The following points explain this issue:</p> <ol style="list-style-type: none"> <li>1. The overall ping timeout is computed using the formula "total ping timeout = ping.interval* ping.timeout".</li> <li>2. During restart, the Management Server picks up the hosts in the "Disconnected" state, whose 'last ping' timestamp is less than this cutoff timeout, for reconnecting (that is, "last ping &lt; (current time - total ping timeout)")</li> <li>3. At the time of restart, the Management Server resets the last ping timestamp of hosts by setting it to "current time - 10mins" based on an already existing logic.</li> <li>4. From 2 and 3, it can be inferred that time taken for hosts to get reconnected is effectively "total ping timeout - 10mins". This causes the delay in connecting the hosts.</li> <li>5. Also, it can be inferred that for "total ping timeout &lt; 10 mins", the reconnect will happen immediately.</li> </ol> <p>Solution: Remove the hardcoding of 10 mins from the code. Instead, use the configurations <code>ping.timeout</code> and <code>ping.interval</code>.</p>
CS-21287	<p>Problem: Unable to delete or archive the events that belong to a deleted user.</p>

Issue ID	Description
	<p>Solution: In the <code>getAccountIdsForDomains</code> method, call <code>customSearchIncludingremoved</code> instead of calling <code>customsearch</code> to include the deleted accounts.</p>
CS-20758	<p>Problem: Random selection of secondary storage by CloudPlatform when there are more than one secondary storage available.</p> <p>Root cause: If there are more than one secondary storage available, CloudPlatform randomly selects one of them. CloudPlatform does not check whether the secondary storage is full before it selects the secondary storage.</p> <p>Solution: CloudPlatform now checks the size of the secondary storage. If there are more than one secondary storage, CloudPlatform does not select the secondary storage if its used percentage exceeds 90%. If the used percentages of all the available secondary storages exceed 90%, CloudPlatform selects one of them.</p>
CS-20627	<p>Problem: The <code>queryAsyncJobResult</code> API function does not return the <code>jobinstanceid</code> response.</p> <p>Root cause: This occurs because the <code>asyncjob</code> framework resets the <code>instance_id</code> and the <code>instance_type</code> fields in the <code>async_job</code> table after completing an <code>async</code> job.</p> <p>Solution: No need to reset the <code>instance_id</code> and the <code>instance_type</code> fields after completing an <code>asyncjob</code>. So, removed that part of code.</p>
CS-20123	<p>Problem: When a VM snapshot is stuck in the Creating state, observed connection issues (host does not connect to the Management Server) with the host to which the VM is associated.</p> <p>Root cause: A failed VM snapshot (failed because of reasons such as timeout) does not move the VM snapshot into the 'Error' state, which results in the host connection issues.</p> <p>Solution: Handling timeout exceptions and other snapshot error scenarios so that unless success operation result, the VM snapshot would be marked as in the "Error" state.</p>
CS-20085	<p>Problem: The <code>catalina.out</code> log file (available at <code>/var/log/cloudstack/management/</code></p>

Issue ID	Description
	<p><b>catalina.out</b>) keeps on growing without a proper daily rotation and compression.</p> <p>Root cause: There is no log rotate configuration file for the catalina.out file created by CloudPlatform. So, this file keeps on growing to a large size and causes problems to the Management Server.</p> <p>Solution: Fixed this issue by creating tomcatcloudstack logrotate configuration file at <code>/etc/logrotate.d/tomcat-cloudstack</code>. CloudPlatform management <b>catalina.out</b> log rotates as expected.</p>
CS-20047	<p>Problem: CloudPlatform allows creation of VMs with the same Display name when the <code>vm.instance.name.flag</code> is set to 'true'.</p> <p>Root cause: This occurs because CloudPlatform restricts two VMs from having the same name only if they are in the same network. But, if <code>vm.instance.name.flag</code> is set to 'true', VM name in vCenter will be its display name. And vCenter does not allow two VMs to have the same name under the same DC, which results in the failure of VM start in CloudPlatform and is not handled correctly.</p> <p>Solution: During VM creation, if <code>vm.instance.name.flag</code> is set to 'true' and hypervisor type is VMware, check if VM with the same hostname already exists in the zone.</p>
CS-19914	<p>Problem: Volume migration times out on CloudPlatform, even though the actual migration on XenServer completes successfully. This leaves the volume on CloudPlatform with incorrect data for the pool location and path.</p> <p>Root cause: Live migration of a volume that is attached to a running VM from one storage pool to another is a long running task. By default, the agent waits for (30*2) minutes for any command to finish on the resource. If volume migration takes longer than that, the agent assumes that the command has not finished and moves ahead. Later, the command completes on the resource, but the answer is not processed by the agent.</p> <p>Solution: Set the timeout value for the <code>MigrateVolumeCommand</code>. This makes sure that the agent waits for the resource to complete migration of volume on the hypervisor.</p>

Issue ID	Description
CS-19816	<p>Problem: Log messages and alerts are created for SSVM disconnection for the running SSVMs.</p> <p>Root cause: Sometimes, Management Server will not respond to agent's ping command on time. In KVM's case, libvirt will take time to respond, which will cause the agent reconnecting to Management Server and eventually failing all the on-going tasks on the host.</p> <p>Solution: Do not try reconnecting in such cases.</p>

### 3.2. Known Issues

Issue ID	Description
CS-44712	<p>Problem: The volume resize operation fails after cold migration from a zone-wide primary storage to a cluster-wide primary storage.</p> <p>Workaround: There is no workaround for this problem. You may need to perform the volume resize operation before you migrate the volume.</p>
CS-44708	<p>Problem: Volume resize fails after the cold migration of a volume to a different storage pool in the same cluster.</p> <p>Root cause: Due to migration to a different storage pool, the volume path at vCenter gets corrupted (goes out of sync) with the actual location of volume files on the datastore.</p> <p>Workaround: If migration is already performed, just start the VM and try to resize volume. If the migration has not been performed, resize the volume before performing volume migration.</p>
CS-44693	<p>Problem: Management server generates multiple 'VOLUME.DELETE' events during the VM expunge operation.</p> <p>Workaround: Duplicate VOLUME.DELETE events does not have any effect on usage generation. The first event will be considered for calculating the usage.</p>
CS-44667	<p>Problem: The resource_count table shows irregular decrease in the resource count for an account with the values are going into negative.</p> <p>Workaround: To correct the resource count, do one of the following:</p> <ol style="list-style-type: none"> <li>1. Use the <code>updateResourceCount</code> API to update the resource count.</li> </ol>

Issue ID	Description
	<p>(For more information on the <code>updateResourceCount</code> API, see: <a href="https://cloudstack.apache.org/api/apidocs-4.3/root_admin/updateResourceCount.html">https://cloudstack.apache.org/api/apidocs-4.3/root_admin/updateResourceCount.html</a>).</p> <ol style="list-style-type: none"> <li data-bbox="762 412 1355 481">2. In the CloudPlatform UI, in the <b>Domains</b> tab, use the <b>Update Resource Count</b> button.</li> </ol>

## 3.3. CloudPlatform 4.3.0.3 Release Notes Addendum

### 3.3.1. Limitation on Dynamically Assigning vCPUs on Windows VMs



#### Note

Read the following limitation as part of the **11.12.1. Prerequisites and Guidelines** section in the *CloudPlatform (powered by Apache CloudStack) 4.3.0.2 Administrator's Guide*:

- Dynamically assigning vCPUs on Windows VMs does not work in case of XenServer. This is because the Windows VMs need to be restarted to make the changes take effect.

### 3.3.2. Providing a Display Name to the Guest Virtual Machines



#### Note

This section has been updated in the *CloudPlatform (powered by Apache CloudStack) 4.3.0.2 Administrator's Guide* as section: **11.6. Providing a Display Name to the Guest Virtual Machines..**

Every guest VM has an internal name. The host uses the internal name to identify the guest VMs. CloudPlatform gives you an option to provide a display name to guest VMs. You can add this display name to the internal name so that it is displayed in vCenter. This feature is intended to make the correlation between instance names and internal names easier in large datacenter deployments.

To provide display names to a VM, you need to set the global configuration parameter `vm.instance.name.flag` to `true`. The default value of this parameter is `false`.

The default format of the internal name is `i-<user_id>-<vm_id>-<instance.name>`, where `instance.name` is a global parameter. After you set the `vm.instance.name.flag` parameter to `true` and provide a display name during the creation of a guest VM, the display name will be displayed in vCenter for the guest.

**Note**

The VMs that are deployed for CloudPlatform version 3.0.7 will continue displaying the VM names in vCenter in the **InternalName-DisplayName** format. The VMs that are deployed for the higher versions of CloudPlatform will display their display name in vCenter.

The following table explains how a VM name is displayed in different scenarios.

In the following table, Display Name represents the user-supplied display name.

User-Provided Display Name	vm.instancename.flag	Host name on the VM	Name on vCenter
Yes	True	Display name	Display Name
No	True	<instance.name>-<UUID>	<instance.name>-<UUID>
Yes	False	Display name	i-<user_id>-<vm_id>-<instance.name>
No	False	<instance.name>-<UUID>	i-<user_id>-<vm_id>-<instance.name>

### 3.3.3. Preventing the Logging of Duplicate VM.START Usage Events

**Note**

Read the following content as part of the **11.4. Creating VMs** section in the *CloudPlatform (powered by Apache CloudStack) 4.3.0.2 Administrator's Guide*:

CloudPlatform creates a worker job when a VM deployment operation starts. The VM synchronization operation thread cannot update the VM state as long as this worker job is on queue. This worker job gets deleted after the completion of the VM deployment operation, which, ideally, must not exceed one hour.

If the VM deployment operation exceeds one hour, the worker job clean up thread that is configured using the `vm.op.cleanup.wait` parameter removes this worker jobs from the database. This allows the VM synchronization thread to update the state of the VM. This results in the logging of duplicate VM.START usage event. As a result of this, the copy template operation may consume more than the expected time to complete.

You can address this issue by modifying the value of the `vm.op.cleanup.wait` parameter. If the VM deployment operation exceeds one hour, you must modify the value of the `vm.op.cleanup.wait` parameter to two hours.



# Upgrade Instructions

## 4.1. Upgrade from 4.3.x.x to 4.3.0.3

Perform the following to upgrade from version 4.3.x.x to version 4.3.0.3.

1. Download the latest System VM templates:

For the latest System VM fixes, follow the procedure given in [Upgrading System VM Template without Upgrading Management Server in CloudPlatform<sup>1</sup>](#).

Hypervisor	Description
XenServer	<p>Name: systemvm-xenserver-4.3</p> <p>Description: systemvm-xenserver-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-xen.vhd.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-xen.vhd.bz2</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, select each zone and individually register the template to make the template available in all the XenServer zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: XenServer</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
KVM	Name: systemvm-kvm-4.3

<sup>1</sup> <http://support.citrix.com/article/CTX200024>

Hypervisor	Description
	<p>Description: systemvm-kvm-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-kvm.qcow2.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-kvm.qcow2.bz2</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, select each zone and individually register the template to make the template available in all the zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: KVM</p> <p>Format: QCOW2</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
VMware	<p>Name: systemvm-vmware-4.3</p> <p>Description: systemvm-vmware-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-vmware.ova">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-vmware.ova</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, select each zone and individually register the template to make the template available in all the zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running</p>

Hypervisor	Description
	<p>VMware, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: VMware</p> <p>Format: OVA</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
<p>Hyper-V</p> <p>(Applicable only for 4.3)</p>	<p>Name: systemvm-hyperv-4.3</p> <p>Description: systemvm-hyperv-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2</a><sup>2</sup></p> <p>Hypervisor: Hyper-V</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>

2. Ensure that the latest System VM are copied to all the primary storages.
3. (KVM on RHEL 6.0/6.1 only) If your existing CloudPlatform deployment includes one or more clusters of KVM hosts running RHEL 6.0 or RHEL 6.1, you must first upgrade the operating system version on those hosts before upgrading CloudPlatform itself.

Run the following commands on every KVM host.

- a. Download the CloudPlatform 4.3.0.3 RHEL 6.3 binaries from <https://www.citrix.com/downloads/cloudplatform.html>.

<sup>2</sup> <http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2>

- b. Extract the binaries:

```
# cd /root
# tar xvf CloudPlatform-4.3.0.3-1-rhel6.3.tar.gz
```

- c. Create a CloudPlatform 4.3 qemu repo:

```
# cd CloudPlatform-4.3.0.3-1-rhel6.3/6.3
# createrepo
```

- d. Prepare the yum repo for upgrade. Edit the file `/etc/yum.repos.d/rhel63.repo`. For example:

```
[upgrade]
name=rhel63
baseurl=url-of-your-rhel6.3-repo
enabled=1
gpgcheck=0
[cloudstack]
name=cloudstack
baseurl=file:///root/CloudPlatform-4.3.0.3-1-rhel6.3/6.3
enabled=1
gpgcheck=0
```

- e. Upgrade the host operating system from RHEL 6.0 to 6.3:

```
yum upgrade
```

4. Stop all Usage Servers if running. Run this on all Usage Server hosts.

```
# service cloudstack-usage stop
```

5. Stop the Management Servers. Run this on all Management Server hosts.

```
# service cloudstack-management stop
```

6. On the MySQL master, take a backup of the MySQL databases. We recommend performing this step even in test upgrades. If there is an issue, this will assist with debugging.

In the following commands, it is assumed that you have set the root password on the database, which is a CloudPlatform recommended best practice. Substitute your own MySQL root password.

```
# mysqldump -u root -p<mysql_password> cloud >> cloud-backup.dmp
# mysqldump -u root -p<mysql_password> cloud_usage > cloud-usage-backup.dmp
```

7. (RHEL/CentOS 5.x) If you are currently running CloudPlatform on RHEL/CentOS 5.x, use the following command to set up an Extra Packages for Enterprise Linux (EPEL) repo:

```
rpm -Uvh http://mirror.pnl.gov/epel/5/i386/epel-release-5-4.noarch.rpm
```

8. Download CloudPlatform 4.3.0.3 onto the management server host where it will run. Get the software from the following link:

<https://www.citrix.com/English/ss/downloads/>.

You need a [My Citrix Account](#)<sup>3</sup>.

9. Upgrade the CloudPlatform packages. You should have a file in the form of "CloudPlatform-4.3.0-N-OSVERSION.tar.gz". Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-4.3.0-N-OSVERSION.tar.gz
# cd CloudPlatform-4.3.0-N-OSVERSION
# ./install.sh
```

You should see a few messages as the installer prepares, followed by a list of choices.

10. Choose "U" to upgrade the package

```
>U
```

You should see some output as the upgrade proceeds, ending with a message like "Complete! Done."

11. If you have made changes to your existing copy of the configuration files db.properties or server.xml in your previous-version CloudPlatform installation, the changes will be preserved in the upgrade. However, you need to do the following steps to place these changes in a new version of the file which is compatible with version 4.3.



### Note

How will you know whether you need to do this? If the upgrade output in the previous step included a message like the following, then some custom content was found in your old file, and you need to merge the two files:

```
warning: /etc/cloud.rpmsave/management/server.xml created as /etc/cloudstack/management/
server.xml.rpmnew
```

- a. Make a backup copy of your previous version file. For example: (substitute the file name in these commands as needed)

```
# mv /etc/cloudstack/management/server.xml /etc/cloudstack/management/server.xml-
backup
```

- b. Copy the \*.rpmnew file to create a new file. For example:

```
# cp -ap /etc/cloudstack/management/server.xml.rpmnew /etc/cloudstack/management/
server.xml
```

<sup>3</sup> <http://www.citrix.com/lang/English/publicindex.asp?destURL=%2FEnglish%2FmyCitrix%2Findex.asp%3F#>

- c. Merge your changes from the backup file into the new file. For example:

```
# vi /etc/cloudstack/management/server.xml
```

12. Repeat steps 7 - 11 on each management server node.

13. Start the first Management Server. Do not start any other Management Server nodes yet.

```
# service cloudstack-management start
```

Wait until the databases are upgraded. Ensure that the database upgrade is complete. After confirmation, start the other Management Servers one at a time by running the same command on each node.



### Note

Failing to restart the Management Server indicates a problem in the upgrade. Restarting the Management Server without any issues indicates that the upgrade is successfully completed.

14. Start all Usage Servers (if they were running on your previous version). Perform this on each Usage Server host.

```
# service cloudstack-usage start
```

15. (VMware only) If you have existing clusters created in CloudPlatform 3.0.6, additional steps are required to update the existing vCenter password for each VMware cluster.

These steps will not affect running guests in the cloud. These steps are required only for clouds using VMware clusters:

- a. Stop the Management Server:

```
service cloudstack-management stop
```

- b. Perform the following on each VMware cluster:

- i. Encrypt the vCenter password:

```
java -classpath /usr/share/cloudstack-common/lib/jasypt-1.9.0.jar  
org.jasypt.intf.cli.JasyptPBEStrEncryptionCLI encrypt.sh  
input=<_your_vCenter_password_> password=`cat /etc/cloudstack/management/key`  
verbose=false
```

Save the output from this step for later use. You need to add this in the `cluster_details` and `vmware_data_center` tables in place of the existing password.

- ii. Find the ID of the cluster from the `cluster_details` table:

```
mysql -u <username> -p<password>
```

```
select * from cloud.cluster_details;
```

- iii. Update the existing password with the encrypted one:

```
update cloud.cluster_details set value = <_ciphertext_from_step_i_> where id = <_id_from_step_ii_>;
```

- iv. Confirm that the table is updated:

```
select * from cloud.cluster_details;
```

- v. Find the ID of the VMware data center that you want to work with:

```
select * from cloud.vmware_data_center;
```

- vi. Change the existing password to the encrypted one:

```
update cloud.vmware_data_center set password = <_ciphertext_from_step_i_> where id = <_id_from_step_v_>;
```

- vii. Confirm that the table is updated:

```
select * from cloud.vmware_data_center;
```

- c. Start the CloudPlatform Management server

```
service cloudstack-management start
```

16. (KVM only) Additional steps are required for each KVM host. These steps will not affect running guests in the cloud. These steps are required only for clouds using KVM as hosts and only on the KVM hosts.



### Note

After the software upgrade on a KVM machine, the Ctrl+Alt+Del button on the console view of a VM doesn't work. Use Ctrl+Alt+Insert to log in to the console of the VM.

- a. Copy the CloudPlatform 4.3.0.3.tgz download to the host, untar it, and change to the resulting directory.
- b. Stop the running agent.

```
# service cloud-agent stop
```

- c. Update the agent software.

```
# ./install.sh
```

- d. Choose "U" to update the packages.

- e. Upgrade all the existing bridge names to new bridge names by running this script:

```
# cloudstack-agent-upgrade
```

- f. Install a libvirt hook with the following commands:

```
# mkdir /etc/libvirt/hooks
# cp /usr/share/cloudstack-agent/lib/libvirtqemuhook /etc/libvirt/hooks/qemu
# chmod +x /etc/libvirt/hooks/qemu
```

- g. Restart libvirtd.

```
# service libvirtd restart
```

- h. Start the agent.

```
# service cloudstack-agent start
```

17. Log in to the CloudPlatform UI as administrator, and check the status of the hosts. All hosts should come to Up state (except those that you know to be offline). You may need to wait 20 or 30 minutes, depending on the number of hosts.



### Note

Troubleshooting: If login fails, clear your browser cache and reload the page.

Do not proceed to the next step until the hosts show in Up state. If the hosts do not come to the Up state, contact support.

18. Perform the following on all the System VMs including Secondary Storage VMs, Console Proxy VMs, and virtual routers.

- a. Upgrade Secondary Storage VMs and Console Proxy VMs either from the UI or by using the following script:

```
# cloudstack-sysvmadm -d <IP address> -u cloud -p <password> -s
```

<IP Address> is the IP address of the cloud database server. If you have not specified this, it will display as root. Also, you can specify any location to collect the logs. Default location is **cloud.log** under current directory.

- b. Selectively upgrade the virtual routers:

- i. Log in to the CloudPlatform UI as the root administrator.
- ii. In the left navigation, choose Infrastructure.
- iii. On Virtual Routers, click View More.

All the VRs are listed in the Virtual Routers page.

- iv. In Select View drop-down, select desired grouping based on your requirement:

You can use either of the following:

- Group by zone
- Group by pod
- Group by cluster
- Group by account

- v. Click the group which has the virtual routers to be upgraded.
- vi. Click the Upgrade button to upgrade all the virtual routers.

For example, if you have selected Group by zone, select the name of the desired zone .

- vii. Click OK to confirm.

19. (XenServer only) Upgrade all existing XenServer clusters to XenServer 6.2 SP1 Hotfix XS62ESP1005.

For more information, see [Section 4.6.4, “Upgrading to XenServer 6.2 SP1 Hotfix XS62ESP1005”](#).

For instructions for upgrading XenServer software and applying hotfixes, see [Section 4.6.2, “Applying Hotfixes to a XenServer Cluster”](#).

20. (VMware only) After upgrade, if you want to change a Standard vSwitch zone to a VMware dvSwitch Zone, perform the following:

- a. Ensure that the Public and Guest traffics are not on the same network as the Management and Storage traffic.
- b. Set `vmware.use.dvswitch` to true.
- c. Access the physical network for the Public and guest traffic, then change the traffic labels as given below:

```
<dvSwitch name>,<VLANID>,<Switch Type>
```

For example: `dvSwitch18,vmwaredvs`

VLANID is optional.

- d. Stop the Management server.
- e. Start the Management server.
- f. Add the new VMware dvSwitch-enabled cluster to this zone.

21. Manually update `systemvm.iso` as given in [Section 4.5, “Updating SystemVM.ISO”](#).

In the previous 4.x releases, the Management Server version stored in the database version table is in x.x.x format. For example, 4.3.0 and 4.3.0.3 are stored as 4.3.0 as only the first 3 digits are considered as release version. Therefore, because the Management Server version number is the same for both the releases, the latest `systemvm.iso` files are not pushed after upgrade. Therefore, you must manually push `systemvm.iso` after upgrade.

### Post-Upgrade Considerations

Consider the following:

- Restart the network with setting cleanup to true if DHCP services run concurrently on two VRs.

Service monitoring is enabled for redundant VR in 4.3, which causes DHCP services to run simultaneously on two VRs. Stopping service monitoring for the existing routers should resolve this issue.

- Troubleshooting tip: If passwords which you know to be valid appear not to work after upgrade, or other UI issues are seen, try clearing your browser cache and reloading the UI page.
- Prior to version 4.3, the VLAN ID in VLAN table is stored as a number, whereas in versions 4.3 and beyond, it is stored as `vlan://<vlanid>`. To accommodate this change, manually edit the database as follows:

```
# mysql> update vlan set vlan_id=concat('vlan://', vlan_id) where vlan_type =  
"VirtualNetwork" and vlan_id not like "vlan://%";
```

- (VMware only) After upgrade, whenever you add a new VMware cluster to a zone that was created with a previous version of CloudPlatform, the fields vCenter host, vCenter Username, vCenter Password, and vCenter Datacenter are required. The Add Cluster dialog in the CloudPlatform user interface incorrectly shows them as optional, and will allow you to proceed with adding the cluster even though these important fields are blank. If you do not provide the values, you will see an error message like "Your host and/or path is wrong. Make sure it's of the format, `http://hostname/path`."
- If you have set the resource limits in the pre-upgraded setup, you may experience resource limit issues with the newly added resource types, such as cpu, memory, primary storage, and secondary storage, in an upgraded setup. This is caused because the limits are not set for these resources, which are added as a part of the upgrade and CloudPlatform is taking the default limits value from the global configuration parameter, which is set to 20.

To make these resource limits unlimited, set the limits to -1 for these newly added resource types in the upgraded setup.

- Size of the snapshots taken before upgrade to 4.3.0.3 or beyond are not updated and remains stored as zero in the database. This leads to inconsistency in secondary storage resource count and the actual secondary storage capacity. To avoid this issue, update the `size` and `physical_size` columns in the `snapshot_store_ref` table, manually or by using a script, to the actual size of the snapshots specified in the secondary storage.

Update the resource count for the ROOT domain by using the action buttons in the Domain Details View in the UI or by using the `updateResourceCount` API.

- If you are using LDAP authentication, change the default values based on the LDAP server that you are using:

LDAP Attribute	OpenLDAP	Active Directory
ldap.user.object	inetOrgPerson	user
ldap.username.attribute	uid	sAMAccountName
ldap.group.object	groupOfUniqueNames	group
ldap.group.user.uniquemember	uniquemember	member
(optional) ldap.search.group.principal	customer-specified	customer-specified

## 4.2. Upgrade from 4.2.x to 4.3.0.3

Perform the following to upgrade from version 4.2.x to version 4.3.0.3.

1. Download the latest System VM templates:

Hypervisor	Description
XenServer	<p>Name: systemvm-xenserver-4.3</p> <p>Description: systemvm-xenserver-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-xen.vhd.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-xen.vhd.bz2</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, select each zone and individually register the template to make the template available in all the XenServer zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: XenServer</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>

Hypervisor	Description
KVM	<p>Name: systemvm-kvm-4.3</p> <p>Description: systemvm-kvm-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-kvm.qcow2.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-kvm.qcow2.bz2</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, select each zone and individually register the template to make the template available in all the zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: KVM</p> <p>Format: QCOW2</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
VMware	<p>Name: systemvm-vmware-4.3</p> <p>Description: systemvm-vmware-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-vmware.ova">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-vmware.ova</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, select each zone and individually register the template to make the template available in all the zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform</p>

Hypervisor	Description
	<p>deployment includes multiple zones running VMware, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: VMware</p> <p>Format: OVA</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
<p>Hyper-V</p> <p>(Applicable only for 4.3)</p>	<p>Name: systemvm-hyperv-4.3</p> <p>Description: systemvm-hyperv-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2</a><sup>4</sup></p> <p>Hypervisor: Hyper-V</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>

2. Ensure that the latest System VM are copied to all the primary storages.
3. (KVM on RHEL 6.0/6.1 only) If your existing CloudPlatform deployment includes one or more clusters of KVM hosts running RHEL 6.0 or RHEL 6.1, you must first upgrade the operating system version on those hosts before upgrading CloudPlatform itself.

Run the following commands on every KVM host.

<sup>4</sup> <http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2>

## Chapter 4. Upgrade Instructions

---

- a. Download the CloudPlatform 4.3.0.3 RHEL 6.3 binaries from <https://www.citrix.com/downloads/cloudplatform.html>.
- b. Extract the binaries:

```
# cd /root
# tar xvf CloudPlatform-4.3.0.3-1-rhel6.3.tar.gz
```

- c. Create a CloudPlatform 4.3 qemu repo:

```
# cd CloudPlatform-4.3.0.3-1-rhel6.3/6.3
# createrepo
```

- d. Prepare the yum repo for upgrade. Edit the file `/etc/yum.repos.d/rhel63.repo`. For example:

```
[upgrade]
name=rhel63
baseurl=url-of-your-rhel6.3-repo
enabled=1
gpgcheck=0
[cloudstack]
name=cloudstack
baseurl=file:///root/CloudPlatform-4.3.0.3-1-rhel6.3/6.3
enabled=1
gpgcheck=0
```

- e. Upgrade the host operating system from RHEL 6.0 to 6.3:

```
yum upgrade
```

4. Stop all Usage Servers if running. Run this on all Usage Server hosts.

```
# service cloudstack-usage stop
```

5. Stop the Management Servers. Run this on all Management Server hosts.

```
# service cloudstack-management stop
```

6. On the MySQL master, take a backup of the MySQL databases. We recommend performing this step even in test upgrades. If there is an issue, this will assist with debugging.

In the following commands, it is assumed that you have set the root password on the database, which is a CloudPlatform recommended best practice. Substitute your own MySQL root password.

```
# mysqldump -u root -p<mysql_password> cloud >> cloud-backup.dmp
# mysqldump -u root -p<mysql_password> cloud_usage > cloud-usage-backup.dmp
```

7. (RHEL/CentOS 5.x) If you are currently running CloudPlatform on RHEL/CentOS 5.x, use the following command to set up an Extra Packages for Enterprise Linux (EPEL) repo:

```
rpm -Uvh http://mirror.pnl.gov/epel/5/i386/epel-release-5-4.noarch.rpm
```

8. Download CloudPlatform 4.3.0.3 onto the management server host where it will run. Get the software from the following link:

<https://www.citrix.com/English/ss/downloads/>.

You need a [My Citrix Account](#)<sup>5</sup>.

9. Upgrade the CloudPlatform packages. You should have a file in the form of "CloudPlatform-4.3.0-N-OSVERSION.tar.gz". Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-4.3.0-N-OSVERSION.tar.gz
# cd CloudPlatform-4.3.0-N-OSVERSION
# ./install.sh
```

You should see a few messages as the installer prepares, followed by a list of choices.

10. Choose "U" to upgrade the package

```
>U
```

You should see some output as the upgrade proceeds, ending with a message like "Complete! Done."

11. If you have made changes to your existing copy of the configuration files db.properties or server.xml in your previous-version CloudPlatform installation, the changes will be preserved in the upgrade. However, you need to do the following steps to place these changes in a new version of the file which is compatible with version 4.3.



### Note

How will you know whether you need to do this? If the upgrade output in the previous step included a message like the following, then some custom content was found in your old file, and you need to merge the two files:

```
warning: /etc/cloud.rpmsave/management/server.xml created as /etc/cloudstack/management/
server.xml.rpmnew
```

- a. Make a backup copy of your previous version file. For example: (substitute the file name in these commands as needed)

```
# mv /etc/cloudstack/management/server.xml /etc/cloudstack/management/server.xml-
backup
```

- b. Copy the \*.rpmnew file to create a new file. For example:

<sup>5</sup> <http://www.citrix.com/lang/English/publicindex.asp?destURL=%2FEnglish%2FmyCitrix%2Findex.asp%3F#>

## Chapter 4. Upgrade Instructions

---

```
# cp -ap /etc/cloudstack/management/server.xml.rpmnew /etc/cloudstack/management/
server.xml
```

- c. Merge your changes from the backup file into the new file. For example:

```
# vi /etc/cloudstack/management/server.xml
```

12. Repeat steps 7 - 11 on each management server node.

13. Start the first Management Server. Do not start any other Management Server nodes yet.

```
# service cloudstack-management start
```

Wait until the databases are upgraded. Ensure that the database upgrade is complete. After confirmation, start the other Management Servers one at a time by running the same command on each node.



### Note

Failing to restart the Management Server indicates a problem in the upgrade. Restarting the Management Server without any issues indicates that the upgrade is successfully completed.

14. Start all Usage Servers (if they were running on your previous version). Perform this on each Usage Server host.

```
# service cloudstack-usage start
```

15. (VMware only) If you have existing clusters created in CloudPlatform 3.0.6, additional steps are required to update the existing vCenter password for each VMware cluster.

These steps will not affect running guests in the cloud. These steps are required only for clouds using VMware clusters:

- a. Stop the Management Server:

```
service cloudstack-management stop
```

- b. Perform the following on each VMware cluster:

- i. Encrypt the vCenter password:

```
java -classpath /usr/share/cloudstack-common/lib/jasypt-1.9.0.jar
org.jasypt.intf.cli.JasyptPBEStrEncryptionCLI encrypt.sh
input=<_your_vCenter_password_> password="`cat /etc/cloudstack/management/key`"
verbose=false
```

Save the output from this step for later use. You need to add this in the `cluster_details` and `vmware_data_center` tables in place of the existing password.

- ii. Find the ID of the cluster from the cluster\_details table:

```
mysql -u <username> -p<password>
```

```
select * from cloud.cluster_details;
```

- iii. Update the existing password with the encrypted one:

```
update cloud.cluster_details set value = <_ciphertext_from_step_i_> where id = <_id_from_step_ii_>;
```

- iv. Confirm that the table is updated:

```
select * from cloud.cluster_details;
```

- v. Find the ID of the VMware data center that you want to work with:

```
select * from cloud.vmware_data_center;
```

- vi. Change the existing password to the encrypted one:

```
update cloud.vmware_data_center set password = <_ciphertext_from_step_i_> where id = <_id_from_step_v_>;
```

- vii. Confirm that the table is updated:

```
select * from cloud.vmware_data_center;
```

- c. Start the CloudPlatform Management server

```
service cloudstack-management start
```

16. (KVM only) Additional steps are required for each KVM host. These steps will not affect running guests in the cloud. These steps are required only for clouds using KVM as hosts and only on the KVM hosts.



### Note

After the software upgrade on a KVM machine, the Ctrl+Alt+Del button on the console view of a VM doesn't work. Use Ctrl+Alt+Insert to log in to the console of the VM.

- a. Copy the CloudPlatform 4.3.0.3.tgz download to the host, untar it, and change to the resulting directory.
- b. Stop the running agent.

## Chapter 4. Upgrade Instructions

---

```
# service cloud-agent stop
```

- c. Update the agent software.

```
# ./install.sh
```

- d. Choose "U" to update the packages.

- e. Upgrade all the existing bridge names to new bridge names by running this script:

```
# cloudstack-agent-upgrade
```

- f. Install a libvirt hook with the following commands:

```
# mkdir /etc/libvirt/hooks  
# cp /usr/share/cloudstack-agent/lib/libvirtqemuhook /etc/libvirt/hooks/qemu  
# chmod +x /etc/libvirt/hooks/qemu
```

- g. Restart libvirtd.

```
# service libvirtd restart
```

- h. Start the agent.

```
# service cloudstack-agent start
```

17. Log in to the CloudPlatform UI as administrator, and check the status of the hosts. All hosts should come to Up state (except those that you know to be offline). You may need to wait 20 or 30 minutes, depending on the number of hosts.



### Note

Troubleshooting: If login fails, clear your browser cache and reload the page.

Do not proceed to the next step until the hosts show in Up state. If the hosts do not come to the Up state, contact support.

18. Perform the following on all the System VMs including Secondary Storage VMs, Console Proxy VMs, and virtual routers.

- a. Upgrade Secondary Storage VMs and Console Proxy VMs either from the UI or by using the following script:

```
# cloudstack-sysvmadm -d <IP address> -u cloud -p <password> -s
```

<IP Address> is the IP address of the cloud database server. If you have not specified this, it will display as root. Also, you can specify any location to collect the logs. Default location is **cloud.log** under current directory.

- b. Selectively upgrade the virtual routers:
  - i. Log in to the CloudPlatform UI as the root administrator.
  - ii. In the left navigation, choose Infrastructure.
  - iii. On Virtual Routers, click View More.  
All the VRs are listed in the Virtual Routers page.
  - iv. In Select View drop-down, select desired grouping based on your requirement:  
You can use either of the following:
    - Group by zone
    - Group by pod
    - Group by cluster
    - Group by account
  - v. Click the group which has the virtual routers to be upgraded.
  - vi. Click the Upgrade button to upgrade all the virtual routers.  
For example, if you have selected Group by zone, select the name of the desired zone .
  - vii. Click OK to confirm.

19. (XenServer only) Upgrade all existing XenServer clusters to XenServer 6.2 SP1 Hotfix XS62ESP1005.

For more information, see [Section 4.6.4, “Upgrading to XenServer 6.2 SP1 Hotfix XS62ESP1005”](#).

For instructions for upgrading XenServer software and applying hotfixes, see [Section 4.6.2, “Applying Hotfixes to a XenServer Cluster”](#).

20. (VMware only) After upgrade, if you want to change a Standard vSwitch zone to a VMware dvSwitch Zone, perform the following:

- a. Ensure that the Public and Guest traffics are not on the same network as the Management and Storage traffic.
- b. Set `vmware.use.dvswitch` to true.
- c. Access the physical network for the Public and guest traffic, then change the traffic labels as given below:

```
<dvSwitch name>,<VLANID>,<Switch Type>
```

For example: `dvSwitch18,vmwaredvs`

VLANID is optional.

- d. Stop the Management server.
- e. Start the Management server.
- f. Add the new VMware dvSwitch-enabled cluster to this zone.

21. Manually update `systemvm.iso` as given in [Section 4.5, “Updating SystemVM.ISO”](#).

In the previous 4.x releases, the Management Server version stored in the database version table is in x.x.x format. For example, 4.3.0 and 4.3.0.3 are stored as 4.3.0 as only the first 3 digits are considered as release version. Therefore, because the Management Server version number is the same for both the releases, the latest `systemvm.iso` files are not pushed after upgrade. Therefore, you must manually push `systemvm.iso` after upgrade.

### Post-Upgrade Considerations

Consider the following:

- Troubleshooting tip: If passwords which you know to be valid appear not to work after upgrade, or other UI issues are seen, try clearing your browser cache and reloading the UI page.
- If you have set the resource limits in the pre-upgraded setup, you may experience resource limit issues with the newly added resource types, such as `cpu`, `memory`, `primary storage`, `sec storage`, in the upgraded setup. This is caused because the limits are not set for these resources which are added as a part of the upgrade and CloudStack is taking the default limits value from the global config parameter, which is set to 20.

To make these resource limits unlimited, set the limits to -1 for these newly added resource types in the upgraded setup.

- Size of the snapshots taken before upgrade to 4.3.0.3 or beyond are not updated and remains stored as zero in the database. This leads to inconsistency in secondary storage resource count and the actual secondary storage capacity. To avoid this issue, update the `size` and `physical_size` columns in the `snapshot_store_ref` table, manually or by using a script, to the actual size of snapshots specified in the secondary storage.

Update the resource count for the ROOT domain by using the action buttons in the Domain Details View in the UI or by using the `updateResourceCount` API.

- Prior to version 4.3, the VLAN ID in VLAN table is stored as a number, whereas in versions 4.3 and beyond, it is stored as `vlan://<vlanid>`. To accommodate this change, manually edit the database as follows:

```
# mysql> update vlan set vlan_id=concat('vlan://', vlan_id) where vlan_type =  
"VirtualNetwork" and vlan_id not like "vlan://%";
```

- (VMware only) After upgrade, whenever you add a new VMware cluster to a zone that was created with a previous version of CloudPlatform, the fields `vCenter host`, `vCenter Username`, `vCenter Password`, and `vCenter Datacenter` are required. The Add Cluster dialog in the CloudPlatform user interface incorrectly shows them as optional, and will allow you to proceed with adding the cluster even though these important fields are blank. If you do not provide the values, you will see an error message like "Your host and/or path is wrong. Make sure it's of the format, `http://hostname/path`."
- If you are using LDAP authentication, change the default values based on the LDAP server that you are using:

LDAP Attribute	OpenLDAP	Active Directory
ldap.user.object	inetOrgPerson	user
ldap.username.attribute	uid	sAMAccountName
ldap.group.object	groupOfUniqueNames	group
ldap.group.user.uniquemember	uniquemember	member
(optional) ldap.search.group.principal	customer-specified	customer-specified

### 4.3. Upgrade from 3.0.x to 4.3.0.3

Perform the following to upgrade from version 3.0.0, 3.0.1, 3.0.2, 3.0.3, 3.0.4, 3.0.5, 3.0.6, or 3.0.7 to version 4.3.0.3.

1. If you are upgrading from 3.0.0 or 3.0.1, ensure that you query your IP address usage records and process them; for example, issue invoices for any usage that you have not yet billed users for.

Starting in 3.0.2, the usage record format for IP addresses is the same as the rest of the usage types. Instead of a single record with the assignment and release dates, separate records are generated per aggregation period with start and end dates. After upgrading, any existing IP address usage records in the old format will no longer be available.

2. While running the 3.0.x system, log in to the UI as root administrator.
3. Using the UI, add a new System VM template for each hypervisor type that is used in your cloud. In each zone, add a system VM template for each hypervisor used in that zone.



#### Note

You might notice that the size of the system VM template has increased compared to previous CloudPlatform versions. This is because the new version of the underlying Debian template has an increased disk size.

- a. In the left navigation bar, click Templates.
- b. In Select view, click Templates.
- c. Click Register template.

The Register template dialog box is displayed.

- d. In the Register template dialog box, specify the following values depending on the hypervisor type (do not change these):

Hypervisor	Description
XenServer	Name: systemvm-xenserver-4.3 Description: systemvm-xenserver-4.3

Hypervisor	Description
	<p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-xen.vhd.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-xen.vhd.bz2</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, select each zone and individually register the template to make the template available in all the XenServer zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: XenServer</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
KVM	<p>Name: systemvm-kvm-4.3</p> <p>Description: systemvm-kvm-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-kvm.qcow2.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-kvm.qcow2.bz2</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, select each zone and individually register the template to make the template available in all the zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple</p>

Hypervisor	Description
	<p>zones running KVM, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: KVM</p> <p>Format: QCOW2</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
VMware	<p>Name: systemvm-vmware-4.3</p> <p>Description: systemvm-vmware-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-vmware.ova">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-vmware.ova</a></p> <p>Zone: (4.3 and beyond) Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, select each zone and individually register the template to make the template available in all the zones.</p> <p>(Prior to version 4.3): Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, choose All Zones to make the template available in all the zones.</p> <p>Hypervisor: VMware</p> <p>Format: OVA</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>
Hyper-V	Name: systemvm-hyperv-4.3

Hypervisor	Description
(Applicable only for 4.3)	<p>Description: systemvm-hyperv-4.3</p> <p>URL (64-bit system VM template): <a href="http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2">http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2</a><sup>6</sup></p> <p>Hypervisor: Hyper-V</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (64-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p>

- e. Watch the screen to be sure that the template downloads successfully and enters the READY state. Do not proceed until this is successful
- f. If you use more than one type of hypervisor in your cloud, repeat these steps to download the system VM template for each hypervisor type.



**Warning**

If you do not repeat the steps for each hypervisor type, the upgrade will fail.

- 4. Ensure that the latest System VM are copied to all the primary storages.
- 5. (KVM on RHEL 6.0/6.1 only) If your existing CloudPlatform deployment includes one or more clusters of KVM hosts running RHEL 6.0 or RHEL 6.1, you must first upgrade the operating system version on those hosts before upgrading CloudPlatform itself.

Run the following commands on every KVM host.

- a. Download the CloudPlatform 4.3.0.3 RHEL 6.3 binaries from <https://www.citrix.com/downloads/cloudplatform.html>.
- b. Extract the binaries:

```
# cd /root
```

<sup>6</sup> <http://download.cloud.com/templates/4.3/systemvm64template-2015-08-20-4.3-hyperv.vhd.bz2>

```
# tar xvf CloudPlatform-4.3.0.3-1-rhel6.3.tar.gz
```

- c. Create a CloudPlatform 4.3.0.3 qemu repo:

```
# cd CloudPlatform-4.3.0.3-1-rhel6.3/6.3
# createrepo
```

- d. Prepare the yum repo for upgrade. Edit the file `/etc/yum.repos.d/rhel63.repo`. For example:

```
[upgrade]
name=rhel63
baseurl=url-of-your-rhel6.3-repo
enabled=1
gpgcheck=0
[cloudstack]
name=cloudstack
baseurl=file:///root/CloudPlatform-4.3.0.3-1-rhel6.3/6.3
enabled=1
gpgcheck=0
```

- e. Upgrade the host operating system from RHEL 6.0 to 6.3:

```
yum upgrade
```

6. Stop all Usage Servers if running. Run this on all Usage Server hosts.

```
# service cloud-usage stop
```

7. Stop the Management Servers. Run this on all Management Server hosts.

```
# service cloud-management stop
```

8. On the MySQL master, take a backup of the MySQL databases. We recommend performing this step even in test upgrades. If there is an issue, this will assist with debugging.

In the following commands, it is assumed that you have set the root password on the database, which is a CloudPlatform recommended best practice. Substitute your own MySQL root password.

```
# mysqldump -u root -p<mysql_password> cloud >> cloud-backup.dmp
# mysqldump -u root -p<mysql_password> cloud_usage > cloud-usage-backup.dmp
```

9. (RHEL/CentOS 5.x) If you are currently running CloudPlatform on RHEL/CentOS 5.x, use the following command to set up an Extra Packages for Enterprise Linux (EPEL) repo:

```
rpm -Uvh http://mirror.pnl.gov/epel/5/i386/epel-release-5-4.noarch.rpm
```

10. Download CloudPlatform 4.3.0.3 onto the management server host where it will run. Get the software from the following link:

<https://www.citrix.com/English/ss/downloads/>

You need a [My Citrix Account](#)<sup>7</sup>.

11. Upgrade the CloudPlatform packages. You should have a file in the form of "CloudPlatform-4.3.0-N-OSVERSION.tar.gz". Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-4.3.0-N-OSVERSION.tar.gz
# cd CloudPlatform-4.3.0-N-OSVERSION
# ./install.sh
```

You should see a few messages as the installer prepares, followed by a list of choices.

12. Choose "U" to upgrade the package

```
>U
```

You should see some output as the upgrade proceeds, ending with a message like "Complete! Done."

13. If you have made changes to your existing copy of the configuration files components.xml, db.properties, or server.xml in your previous-version CloudPlatform installation, the changes will be preserved in the upgrade. However, you need to do the following steps to place these changes in a new version of the file which is compatible with version 4.3.0.3



### Note

How will you know whether you need to do this? If the upgrade output in the previous step included a message like the following, then some custom content was found in your old file, and you need to merge the two files:

```
warning: /etc/cloud.rpmsave/management/components.xml created as /etc/cloudstack/management/components.xml.rpmnew
```

- a. Make a backup copy of your previous version file. For example: (substitute the file name components.xml, db.properties, or server.xml in these commands as needed)

```
# mv /etc/cloudstack/management/components.xml /etc/cloudstack/management/
components.xml-backup
```

- b. Copy the \*.rpmnew file to create a new file. For example:

```
# cp -ap /etc/cloudstack/management/components.xml.rpmnew /etc/cloudstack/management/
components.xml
```

- c. Merge your changes from the backup file into the new file. For example:

---

<sup>7</sup> <http://www.citrix.com/lang/English/publicindex.asp?destURL=%2FEnglish%2FmyCitrix%2Findex.asp%3F#>

```
# vi /etc/cloudstack/management/components.xml
```

14. Repeat steps 9 - 13 on each management server node.

15. Start the first Management Server. Do not start any other Management Server nodes yet.

```
# service cloudstack-management start
```

Wait until the databases are upgraded. Ensure that the database upgrade is complete. After confirmation, start the other Management Servers one at a time by running the same command on each node.



### Note

Failing to restart the Management Server indicates a problem in the upgrade. Restarting the Management Server without any issues indicates that the upgrade is successfully completed.

16. Start all Usage Servers (if they were running on your previous version). Perform this on each Usage Server host.

```
# service cloudstack-usage start
```



### Note

After upgrade from 3.0.4 to 4.3.0.3, if the usage server fails to restart then copy db.properties from /etc/cloudstack/management to /etc/cloudstack/usage. Then start the Usage Server.

17. (VMware only) If you are upgrading from 3.0.6 or beyond and you have existing clusters created in 3.0.6, additional steps are required to update the existing vCenter password for each VMware cluster.

These steps will not affect running guests in the cloud. These steps are required only for clouds using VMware clusters:

a. Stop the Management Server:

```
service cloudstack-management stop
```

b. Perform the following on each VMware cluster:

i. Encrypt the vCenter password:

```
java -classpath /usr/share/cloudstack-common/lib/jasypt-1.9.0.jar  
org.jasypt.intf.cli.JasyptPBEStrEncryptionCLI encrypt.sh
```

```
input=<_your_vCenter_password_> password="`cat /etc/cloudstack/management/key`"  
verbose=false
```

Save the output from this step for later use. You need to add this in the `cluster_details` and `vmware_data_center` tables in place of the existing password.

- ii. Find the ID of the cluster from the `cluster_details` table:

```
mysql -u <username> -p<password>
```

```
select * from cloud.cluster_details;
```

- iii. Update the existing password with the encrypted one:

```
update cloud.cluster_details set value = <_ciphertext_from_step_i_> where id =  
<_id_from_step_ii_>;
```

- iv. Confirm that the table is updated:

```
select * from cloud.cluster_details;
```

- v. Find the ID of the VMware data center that you want to work with:

```
select * from cloud.vmware_data_center;
```

- vi. Change the existing password to the encrypted one:

```
update cloud.vmware_data_center set password = <_ciphertext_from_step_i_> where  
id = <_id_from_step_v_>;
```

- vii. Confirm that the table is updated:

```
select * from cloud.vmware_data_center;
```

- c. Start the CloudPlatform Management server

```
service cloudstack-management start
```

18. (KVM only) Additional steps are required for each KVM host. These steps will not affect running guests in the cloud. These steps are required only for clouds using KVM as hosts and only on the KVM hosts.



### Note

After the software upgrade on a KVM machine, the `Ctrl+Alt+Del` button on the console view of a VM doesn't work. Use `Ctrl+Alt+Insert` to log in to the console of the VM.

- a. Copy the CloudPlatform 4.3.0.3.tgz download to the host, untar it, and cd into the resulting directory.
- b. Stop the running agent.

```
# service cloud-agent stop
```

- c. Update the agent software.

```
# ./install.sh
```

- d. Choose "U" to update the packages.
- e. Edit `/etc/cloudstack/agent/agent.properties` to change the resource parameter from `com.cloud.agent.resource.computing.LibvirtComputingResource` to `com.cloud.hypervisor.kvm.resource.LibvirtComputingResource`.
- f. Upgrade all the existing bridge names to new bridge names by running this script:

```
# cloudstack-agent-upgrade
```

- g. Install a libvirt hook with the following commands:

```
# mkdir /etc/libvirt/hooks  
# cp /usr/share/cloudstack-agent/lib/libvirtqemuhook /etc/libvirt/hooks/qemu  
# chmod +x /etc/libvirt/hooks/qemu
```

- h. Restart libvirtd.

```
# service libvirtd restart
```

- i. Start the agent.

```
# service cloudstack-agent start
```

19. Log in to the CloudPlatform UI as administrator, and check the status of the hosts. All hosts should come to Up state (except those that you know to be offline). You may need to wait 20 or 30 minutes, depending on the number of hosts.



### Note

Troubleshooting: If login fails, clear your browser cache and reload the page.

Do not proceed to the next step until the hosts show in Up state. If the hosts do not come to the Up state, contact support.

20. If you are upgrading from 3.0.1 or 3.0.2, perform the following:

- a. Ensure that the admin port is set to 8096 by using the "integration.api.port" global parameter.

This port is used by the cloudstack-sysvadm script later in the upgrade procedure. For information about how to set this parameter, see "Setting Configuration Parameters" in the Installation Guide.

- b. Restart the Management Server.



### Note

If you don't want the admin port to remain open, you can set it to null after the upgrade is done and restart the Management Server.

21. Perform the following on all the System VMs including Secondary Storage VMs, Console Proxy VMs, and virtual routers.

- a. Upgrade Secondary Storage VMs and Console Proxy VMs either from the UI or by using the following script:

```
# cloudstack-sysvadm -d <IP address> -u cloud -p <password> -s
```

<IP Address> is the IP address of the cloud database server. If you have not specified this, it will display as root. Also, you can specify any location to collect the logs. Default location is **cloud.log** under current directory.

- b. Selectively upgrade the virtual routers:

- i. Log in to the CloudPlatform UI as the root administrator.
- ii. In the left navigation, choose Infrastructure.
- iii. On Virtual Routers, click View More.

All the VRs are listed in the Virtual Routers page.

- iv. In Select View drop-down, select desired grouping based on your requirement:

You can use either of the following:

- Group by zone
  - Group by pod
  - Group by cluster
  - Group by account
- v. Click the group which has the virtual routers to be upgraded.
  - vi. Click the Upgrade button to upgrade all the virtual routers.

For example, if you have selected Group by zone, select the name of the desired zone .

- vii. Click OK to confirm.
22. If you would like additional confirmation that the new system VM templates were correctly applied when these system VMs were rebooted, SSH into the System VM and check the version.

Use one of the following techniques, depending on the hypervisor.

### XenServer or KVM:

SSH in by using the link local IP address of the system VM. For example, in the command below, substitute your own path to the private key used to log in to the system VM and your own link local IP.

Run the following commands on the XenServer or KVM host on which the system VM is present:

```
# ssh -i /root/.ssh/id_rsa.cloud <link-local-ip> -p 3922
# cat /etc/cloudstack-release
```

The output should be like the following:

```
Cloudstack Release 4.3.0.3 Mon Oct 14 15:10:04 PST 2013
```

### ESXi

SSH in using the private IP address of the system VM. For example, in the command below, substitute your own path to the private key used to log in to the system VM and your own private IP.

Run the following commands on the Management Server:

```
# ssh -i /var/cloudstack/management/.ssh/id_rsa <private-ip> -p 3922
# cat /etc/cloudstack-release
```

The output should be like the following:

```
Cloudstack Release 4.3.0.3 Fri Oct 8 15:10:04 PST 2014
```

23. If you want to close the admin port again (recommended in production systems), set `integration.api.port` to null. Then restart the Management Server.

For information about how to set `integration.api.port`, see “Setting Configuration Parameters” in the Installation Guide.

24. (XenServer only) Upgrade all existing XenServer clusters to XenServer 6.2 SP1 Hotfix XS62ESP1005.

For more information, see [Section 4.6.4, “Upgrading to XenServer 6.2 SP1 Hotfix XS62ESP1005”](#).

For instructions for upgrading XenServer software and applying hotfixes, see [Section 4.6.2, “Applying Hotfixes to a XenServer Cluster”](#).

25. (VMware only) After upgrade, if you want to change a Standard vSwitch zone to a VMware dvSwitch Zone, perform the following:

## Chapter 4. Upgrade Instructions

---

- a. Ensure that the Public and Guest traffics are not on the same network as the Management and Storage traffic.
- b. Set `vmware.use.dvswitch` to `true`.
- c. Access the physical network for the Public and guest traffic, then change the traffic labels as given below:

```
<dvSwitch name>,<VLANID>,<Switch Type>
```

For example: `dvSwitch18,,vmwaredvs`

VLANID is optional.

- d. Stop the Management server.
- e. Start the Management server.
- f. Add the new VMware dvSwitch-enabled cluster to this zone.

### Post-Upgrade Considerations

Consider the following:

- Troubleshooting tip: If passwords which you know to be valid appear not to work after upgrade, or other UI issues are seen, try clearing your browser cache and reloading the UI page.
- If you have set the resource limits in the pre-upgraded setup, you may experience resource limit issues with the newly added resource types, such as `cpu`, `memory`, `primary storage`, `secondary storage`, in the upgraded setup. This is caused because the limits are not set for these resources which are added as a part of the upgrade and CloudPlatform is taking the default limits value from the global configuration parameter, which is set to 20.

To make these resource limits unlimited, set the limits to `-1` for these newly added resource types in the upgraded setup.

- Size of the snapshots taken before upgrading to 4.3.0.3 or beyond are not updated and remains stored as zero in the database. This leads to inconsistency in secondary storage resource count and the actual secondary storage capacity. To avoid this issue, update the `size` and `physical_size` columns in the `snapshot_store_ref` table, manually or by using a script, to the actual size of snapshots specified in the secondary storage.

Update the resource count for the ROOT domain by using the action buttons in the Domain Details View in the UI or by using the `updateResourceCount` API.

- Prior to version 4.3, the VLAN ID in VLAN table is stored as a number, whereas in versions 4.3 and beyond, it is stored as `vlan://<vlanid>`. To accomodate this change, annually edit the database as follows:

```
# mysql> update vlan set vlan_id=concat('vlan://', vlan_id) where vlan_type =  
"VirtualNetwork" and vlan_id not like "vlan://%";
```

- (VMware only) After upgrade, whenever you add a new VMware cluster to a zone that was created with a previous version of CloudPlatform, the fields `vCenter host`, `vCenter Username`, `vCenter Password`, and `vCenter Datacenter` are required. The Add Cluster dialog in the CloudPlatform user interface incorrectly shows them as optional, and will allow you to proceed with adding the cluster

even though these important fields are blank. If you do not provide the values, you will see an error message like "Your host and/or path is wrong. Make sure it's of the format, http://hostname/path.

- If you are using LDAP authentication, change the default values based on the LDAP server that you are using:

LDAP Attribute	OpenLDAP	Active Directory
ldap.user.object	inetOrgPerson	user
ldap.username.attribute	uid	sAMAccountName
ldap.group.object	groupOfUniqueNames	group
ldap.group.user.uniquemember	uniquemember	member
(optional) ldap.search.group.principal	customer-specified	customer-specified

## 4.4. Upgrade CloudPlatform Baremetal Agent on PXE and DHCP Servers

If you installed bare metal clusters using a previous version of CloudPlatform, use the following steps to upgrade the baremetal agent in order to get the latest bug fixes for 4.3.0.

1. Log in as root to the host or virtual machine running the Baremetal PXE server and DHCP server.
2. Download CloudPlatform 4.3.0.3 onto the PXE or DHCP server. Get the software from the following link:

<https://www.citrix.com/English/ss/downloads/>.

You need a [My Citrix Account](#)<sup>8</sup>.

3. Upgrade the CloudPlatform packages. You should have a file in the form of "CloudPlatform-4.3.0-N-OSVERSION.tar.gz". Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-4.3.0-N-OSVERSION.tar.gz
# cd CloudPlatform-4.3.0-N-OSVERSION
# ./install.sh
```

You should see a few messages as the installer prepares, followed by a list of choices.

4. Choose "U" to upgrade the package

```
>U
```

You should see some output as the upgrade proceeds, ending with a message like "Complete! Done."

5. Run the bare metal setup script:

```
cloudstack-setup-baremetal
```

<sup>8</sup> <http://www.citrix.com/lang/English/publicindex.asp?destURL=%2FEnglish%2FmyCitrix%2Findex.asp%3F#>

### 4.5. Updating SystemVM.ISO

- On CloudPlatform versions 3.0.5.x and 3.0.7.x `systemvm.iso` will get propagated automatically; therefore, no separate procedure is required.
- On CloudPlatform versions 4.2.1.x and 4.3.x, perform the following based on the hypervisor that you use:
  - XenServer: No action is required.
  - KVM
    - a. On the KVM host, stop the CloudPlatform agent.
    - b. Upgrade the CloudPlatform agent.
    - c. Restart the CloudPlatform agent.
    - d. Stop and Start SystemVMs.
  - HyperV (for CloudPlatform versions 4.3 and above)
    - a. Stop all the Management Servers.
    - b. Remove `systemvm-4.3.x.x.iso` from the `systemvm` directory in the Secondary Storage directory, `\\<secondary_storage_path>\systemvm\`.
    - c. Remove `systemvm-4.3.x.x.iso` from each Hyper-V host.

The location of the file is `C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks`.
    - d. Start the Management Server.
    - e. Destroy SystemVMs.

New SystemVMs will be spawned and the new iso, `systemvm-4.3.x.x.iso`, is copied to the secondary storage and Hypervisor host.
  - VMware
    - a. Stop all the Management Servers.
    - b. Remove the old `systemvm<version>.iso` file from the `systemvm` directory, `\\<secondary_storage_path>\systemvm\`.

Where `<version>` denotes the Management Server version number.
    - c. Start the Management Server.

Verify if the new `systemvm.iso` is pushed to the `systemvm` folder in the Secondary Storage directory.
    - d. Stop and Start SystemVMs.

### 4.6. Upgrading and Hotfixing XenServer Hypervisor Hosts

In CloudPlatform 4.3.0, you can upgrade XenServer hypervisor host software without having to disconnect the XenServer cluster. You can upgrade XenServer 5.6 GA, 5.6 FP1, or 5.6 SP2 to any

newer version that is supported by CloudPlatform. The actual upgrade is described in XenServer documentation, but there are some additional steps you must perform before and after the upgrade.

### 4.6.1. Upgrading to a New XenServer Version

To upgrade XenServer hosts when running CloudPlatform 4.3.0.3:

1. Edit the file `/etc/cloudstack/management/environment.properties` and add the following line:

```
manage.xenserver.pool.master=false
```

2. Restart the Management Server to put the new setting into effect.

```
# service cloudstack-management restart
```

3. Find the hostname of the master host in your XenServer cluster (pool):

- a. Run the following command on any host in the pool, and make a note of the host-uuid of the master host:

```
# xe pool-list
```

- b. Now run the following command, and find the host that has a host-uuid that matches the master host from the previous step. Make a note of this host's hostname. You will need to input it in a later step.

```
# xe host-list
```

4. On CloudPlatform, put the master host into maintenance mode. Use the hostname you discovered in the previous step.



#### Note

In the latest XenServer upgrade procedure, even after putting the master host into maintenance mode, the master host continues to stay as master.

Any VMs running on this master will be automatically migrated to other hosts, unless there is only one UP host in the cluster. If there is only one UP host, putting the host into maintenance mode will stop any VMs running on the host.

5. Disconnect the XenServer cluster from CloudPlatform. It will remain disconnected only long enough to upgrade one host.
  - a. Log in to the CloudPlatform UI as root.
  - b. Navigate to the XenServer cluster, and click Actions – Unmanage.
  - c. Watch the cluster status until it shows Unmanaged.
6. Upgrade the XenServer software on the master host:

- a. Insert the XenServer CD.
  - b. Reboot the host.
  - c. Upgrade to the newer version of XenServer. Use the steps in XenServer documentation.
7. Cancel the maintenance mode on the master host.
  8. Reconnect the XenServer cluster to CloudPlatform.
    - a. Log in to the CloudPlatform UI as root.
    - b. Navigate to the XenServer cluster, and click Actions – Manage.
    - c. Watch the status to see that all the hosts come up.
  9. Upgrade the slave hosts in the cluster:
    - a. Put a slave host into maintenance mode.

Wait until all the VMs are migrated to other hosts.
    - b. Upgrade the XenServer software on the slave.
    - c. Cancel maintenance mode for the slave.
    - d. Repeat steps [a](#) through [c](#) for each slave host in the XenServer pool.
  10. You might need to change the OS type settings for VMs running on the upgraded hosts, if any of the following apply:
    - If you upgraded from XenServer 5.6 GA to XenServer 5.6 SP2, change any VMs that have the OS type CentOS 5.5 (32-bit), Oracle Enterprise Linux 5.5 (32-bit), or Red Hat Enterprise Linux 5.5 (32-bit) to Other Linux (32-bit). Change any VMs that have the 64-bit versions of these same OS types to Other Linux (64-bit).
    - If you upgraded from XenServer 5.6 SP2 to XenServer 6.0.2 or higher, change any VMs that have the OS type CentOS 5.6 (32-bit), CentOS 5.7 (32-bit), Oracle Enterprise Linux 5.6 (32-bit), Oracle Enterprise Linux 5.7 (32-bit), Red Hat Enterprise Linux 5.6 (32-bit) , or Red Hat Enterprise Linux 5.7 (32-bit) to Other Linux (32-bit). Change any VMs that have the 64-bit versions of these same OS types to Other Linux (64-bit).
    - If you upgraded from XenServer 5.6 to XenServer 6.0.2 or higher, do all of the above.

### 4.6.2. Applying Hotfixes to a XenServer Cluster

1. Edit the file `/etc/cloudstack/management/environment.properties` and add the following line:

```
manage.xenserver.pool.master=false
```

2. Restart the Management Server to put the new setting into effect.

```
# service cloudstack-management restart
```

3. Find the hostname of the master host in your XenServer cluster (pool):

- a. Run the following command on any host in the pool, and make a note of the host-uuid of the master host:

```
# xe pool-list
```

- b. Now run the following command, and find the host that has a host-uuid that matches the master host from the previous step. Make a note of this host's hostname. You will need to input it in a later step.

```
# xe host-list
```

4. On CloudPlatform, put the master host into maintenance mode. Use the hostname you discovered in the previous step.

Any VMs running on this master will be automatically migrated to other hosts, unless there is only one UP host in the cluster. If there is only one UP host, putting the host into maintenance mode will stop any VMs running on the host.

5. Disconnect the XenServer cluster from CloudPlatform. It will remain disconnected only long enough to hotfix one host.

- a. Log in to the CloudPlatform UI as root.
- b. Navigate to the XenServer cluster, and click Actions – Unmanage.
- c. Watch the cluster status until it shows Unmanaged.

6. Hotfix the master host:

- a. Add the XenServer hot fixes to the master host.

- i. Assign a UUID to the update file:

```
xe patch-upload file-name=XS602E015.xsupdate
```

The command displays the UUID of the update file:

```
33af688e-d18c-493d-922b-ec51ea23cfe9
```

- ii. Repeat the `xe patch-upload` command for all other XenServer updates: `XS62ESP1005.xsupdate`, `XS62ESP1003.xsupdate`.

Take a note of the UUIDs of the update files. The UUIDs are required in the next step.

- b. Apply XenServer hot fixes to master host:

```
xe patch-apply host-uuid=<master uuid> uuid=<hotfix uuid>
```

- c. Repeat `xe patch-apply` command for all the hot fixes.
- d. Install the required CSP files.

```
xe-install-supplemental-pack <csp-iso-file>
```

- e. Restart the master host.
7. Cancel the maintenance mode on the master host.
8. Reconnect the XenServer cluster to CloudPlatform.
  - a. Log in to the CloudPlatform UI as root.
  - b. Navigate to the XenServer cluster, and click Actions – Manage.
  - c. Watch the status to see that all the hosts come up.
9. Hotfix the slave hosts in the cluster:
  - a. Put a slave host into maintenance mode.

Wait until all the VMs are migrated to other hosts.

- b. Apply the XenServer hot fixes to the slave host:

```
xe patch-apply host-uuid=<slave uuid> uuid=<hotfix uuid>
```

- c. Repeat Step a through b for each slave host in the XenServer pool.
- d. Install the required CSP files.

```
xe-install-supplemental-pack <csp-iso-file>
```

- e. Restart the slave hosts.

Wait until all the slave hosts are up. It might take several minutes for the hosts to come up.

10. Cancel the maintenance mode on the slave hosts.
11. You might need to change the OS type settings for VMs running on the upgraded hosts, if any of the following apply:
  - If you upgraded from XenServer 5.6 SP2 to XenServer 6.0.2, change any VMs that have the OS type CentOS 5.6 (32-bit), CentOS 5.7 (32-bit), Oracle Enterprise Linux 5.6 (32-bit), Oracle Enterprise Linux 5.7 (32-bit), Red Hat Enterprise Linux 5.6 (32-bit) , or Red Hat Enterprise Linux 5.7 (32-bit) to Other Linux (32-bit). Change any VMs that have the 64-bit versions of these same OS types to Other Linux (64-bit).
  - If you upgraded from XenServer 5.6 GA or 5.6 FP1 to XenServer 6.0.2, change any VMs that have the OS type CentOS 5.5 (32-bit), CentOS 5.6 (32-bit), CentOS 5.7 (32-bit), Oracle Enterprise Linux 5.5 (32-bit), Oracle Enterprise Linux 5.6 (32-bit), Oracle Enterprise Linux 5.7 (32-bit), Red Hat Enterprise Linux 5.5 (32-bit), Red Hat Enterprise Linux 5.6 (32-bit) , or Red Hat Enterprise Linux 5.7 (32-bit) to Other Linux (32-bit). Change any VMs that have the 64-bit versions of these same OS types to Other Linux (64-bit).

### 4.6.3. Install CloudPlatform XenServer Support Package (CSP)

Ensure that you install CloudPlatform XenServer Support Package (CSP) to enable security groups, elastic load balancing, and elastic IP on XenServer.

For more information, see the Install CloudPlatform XenServer Support Package (CSP) in the Installation Guide.

If your hosts on versions prior to 6.2 operated on bridge mode with CSP packages installed, after upgrade, run only the following to restore the desired Security Groups configuration:

1. If the XenServer host is part of a zone that uses basic networking, disable Open vSwitch (OVS):

```
# xe-switch-network-backend bridge
```

2. Restart the host machine when prompted.
3. If you are using XenServer 6.1 or greater, perform the following:

- a. Run the following commands:

```
echo 1 > /proc/sys/net/bridge/bridge-nf-call-iptables  
echo 1 > /proc/sys/net/bridge/bridge-nf-call-arptables
```

- b. To persist the above changes across reboots, set the following values in the `/etc/sysctl.conf` file. Run the following command:

```
sysctl -p /etc/sysctl.conf
```

Set these to 1:

```
net.bridge.bridge-nf-call-iptables = 1  
net.bridge.bridge-nf-call-arptables = 1
```

#### 4.6.4. Upgrading to XenServer 6.2 SP1 Hotfix XS62ESP1005

It is highly recommended that all XenServer clusters are upgraded to XenServer 6.2 SP1 Hotfix XS62ESP1005. You can upgrade from any prior version of XenServer to the latest version, which might include multiple hops as part of a single upgrade process. For example, if you are upgrading from 6.0.2, upgrade the master host by using the upgrade path given below, followed by each slave host upgrading to XenServer 6.2 SP1 Hotfix XS62ESP1005 by using this same upgrade path:

1. XenServer 6.0.2 to XenServer 6.2
2. XenServer 6.2 to XenServer 6.2 SP1
3. XenServer 6.2 SP1 to XenServer 6.2 SP1 Hotfix XS62ESP1005

After upgrading, ensure that XenServer Pool HA is enabled.

For information on enabling Pool HA for HA support, see Enabling Pool HA section in the Citrix CloudPlatform Installation Guide.

