

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.3.0.2 Release Notes

Revised December 15, 2014 3:00 PM PST



Citrix CloudPlatform (powered by Apache CloudStack) Version 4.3.0.2 Release Notes

Revised December 15, 2014 3:00 PM PST

© 2014 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 4.3.0.2.

1. Submitting Feedback and Getting Help	1
2. Support Matrix	3
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	6
3. What's New in 4.3.0.2	9
3.1. Fixed Issues	9
3.2. Known Issues	40
4. Upgrade Instructions	53
4.1. Upgrade from 4.3.x.x to 4.3.0.2	53
4.2. Upgrade from 4.2.x to 4.3.0.2	63
4.3. Upgrade from 3.0.x to 4.3.0.2	73
4.4. Upgrade CloudPlatform Baremetal Agent on PXE and DHCP Servers	85
4.5. Updating SystemVM.ISO	86
4.6. Upgrading and Hotfixing XenServer Hypervisor Hosts	86
4.6.1. Upgrading to a New XenServer Version	87
4.6.2. Applying Hotfixes to a XenServer Cluster	88
4.6.3. Install CloudPlatform XenServer Support Package (CSP)	90
4.6.4. Upgrading to XenServer 6.2 SP1 Hotfix XS62ESP1005	91

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](#)¹ by using the account credentials you received when you purchased your support contract.

¹ <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 4.3.0.2. Most earlier OS and hypervisor versions are also still supported for use with 4.3.0.2. 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 4.3.0.2 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 CloudPlatform¹](#).

Hypervisor	Description
XenServer	Name: systemvm-xenserver-4.3 Description: systemvm-xenserver-4.3 URL (64-bit system VM template): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-xen.vhd.bz2 Zone: 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 XenServer zones. Hypervisor: XenServer Format: VHD

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

Hypervisor	Description
	<p>OS Type: Debian GNU/Linux 7.0 (32-bit and 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	<p>Name: systemvm-hyperv-4.3</p> <p>Description: systemvm-hyperv-4.3</p> <p>URL (64-bit system VM template): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-hyperv.vhd.bz2²</p> <p>Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running Hyper-V, choose All Zones to make the template available in all the Hyper-V zones.</p> <p>Hypervisor: Hyper-V</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (32-bit and 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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-kvm.qcow2.bz2</p> <p>Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes</p>

² <http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-hyperv.vhd.bz2>

Hypervisor	Description
	<p>multiple zones running KVM, choose All Zones to make the template available in all the KVM zones.</p> <p>Hypervisor: KVM</p> <p>Format: QCOW2</p> <p>OS Type: Debian GNU/Linux 7.0 (32-bit and 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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-vmware.ova</p> <p>Zone: 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 VMware zones.</p> <p>Hypervisor: VMware</p> <p>Format: OVA</p> <p>OS Type: Debian GNU/Linux 7.0 (32-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>

2.5. 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 4.3.0.2

CloudPlatform 4.3.0.2 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-26155	<p>Problem: Every time when a VM worker job is placed into sync_queue_item table in the Management Server log, a warning, "Was unable to find lock for the key vm_instance286 and thread id 1581663601", is displayed indicating that lock is not properly acquired in sync vm worker job into the queue.</p> <p>Root cause: The jobs in a queue use Integer.MAX_VALUE as the timeout period for acquiring a lock, which internally will be convert to milliseconds and will be out of integer range, so all the jobs are executed without acquiring any lock. Therefore, when lock is released it fails to find the lock and thus throw a warning.</p> <p>Solution: Refactor all the VM job queue code to place the lock acquire in a central location. Additionally, configured timeout is used instead of Integer.MAX_VALUE which is too long.</p>
CS-26144	<p>Problem: Usage job fails due to duplicate IP assign events occurs simultaneously for the same IP.</p> <p>Root cause: This is a synchronization issue in the Management Server while marking an IP as allocated. When multiple IPs are acquired for the same network, the same IP is marked as allocated multiple times, resulting in multiple event generation both in usage_event table and in event bus.</p> <p>Solution: Avoid marking IPs that are already in Allocated state as Allocated again. Use the row lock to ensure that previous state is either Allocating or Free. This will in turn avoid logging duplicate events.</p>
CS-25973	<p>Problem: The global setting for vSphere, VM.instancename, uses just the instancename without the prefix.</p> <p>Root cause: In version 4.2 MR1, by default VM name in vCenter is "InternalName-VM", for example: i-2-3-VM. However, if the CloudPlatform global configuration,</p>

Issue ID	Description
	<p>vm.instancename.flag, is set to true, the VM name in vCenter will be given as the VM display name.</p> <p>Solution: VM name in vCenter should be VM's display name.</p>
CS-25905	<p>Problem: Attaching a volume to a running VM does not work as expected on XenServer 6.2 SP1 with Xentools 6.2.</p> <p>Root cause: CloudPlatform doesn't allow attaching a volume to a windows VM if the PV driver is not updated.</p> <p>Solution: The check to know the risk of running windows VM without latest PV driver has been removed. This workaround is not officially supported by XenServer.</p>
CS-25847	<p>Problem: The usage records of deleted volumes are not removed from cloud_usage tables.</p> <p>Root cause: This issue occurs when some volumes are deleted from the database; however all these volumes are using the query: select distinct usage_id from cloud_usage.cloud_usage where usage_type = 6 and usage_id not in (select id from cloud.volumes); .</p> <p>Solution: To stop usage for deleted volumes:</p> <ol style="list-style-type: none"> 1. Run the following: <div data-bbox="810 1319 1350 1435" style="border: 1px solid gray; padding: 5px; margin: 5px 0;"> <pre># update usage_volume set deleted = <vol_deleted_ts> where id = <vol_id></pre> </div> 2. Run the following: <div data-bbox="810 1527 1350 1671" style="border: 1px solid gray; padding: 5px; margin: 5px 0;"> <pre># delete from cloud_usage where usage_type = 6 and usage_id = <vol_id> and start_date > <vol_deleted_ts></pre> </div>
CS-25817	<p>Problem: Unable to resize or storage live migrate VM or volume if volume is attached to Windows VMs with guest OS of Windows 7 or later versions, for example, Windows 2008.</p> <p>Root cause: The adapter type in the VMDK file associated with volumes attached are using LsiLogic SAS controllers, which is not supported by the DISKLIB module of ESXi while controller type used was LSI logic SAS.</p>

Issue ID	Description
	<p>Solution: While attaching a volume, ensured that the VMDK file has supported VMDK adapter type. Additionally, during resize or volume migration, CloudPlatform checks whether a VMDK file has already supported the VMDK adapter type, and update adapter type otherwise. Further, safe guard checks has been added before doing real VM and data migration.</p>
CS-25687	<p>Problem: Templates source_template_id is null when it is created from Snapshot with its corresponding volume removed.</p> <p>Root cause: Volume was removed and finding volume only by the findbyId method.</p> <p>Solution: When setting the source template through the volume, search for the volume including the removed one because the source template might be removed.</p>
CS-25452	<p>Problem: Destroyed VMs aren't being expunged.</p> <p>Root Cause: The VM worker job expunge thread that handles VM cleanup runs periodically and removes the completed jobs, which are an hour old, from the job queue in the async_job table. The defect is that all the completed jobs regardless of how old they were, were picked up and deleted due to a wrong criteria. Because of which the job wasn't properly completed and a Null Pointer exception was raised. This caused the job to get stuck in the queue, not allowing any further jobs to execute on the VM. The UserVM Scavenger thread picks up destroyed VMs for expunge every 'expunge.interval' time, and this thread waits for each VM expunge work to complete before proceeding to the next. However, if one of the expunge jobs gets blocked, in this particular case all the jobs on the VM are blocked because of the previous thread, the UserVM scavenger thread will be blocked and no further VM expunge will happen in CloudPlatform.</p> <p>Solution: Two fixes are made: Ensured that VM worker expunge thread correctly expunges the VM worker jobs. Ensured that if a job is already expunged from async job table, CloudPlatform doesn't make updates to the job in the database and only purges the job from the job queue. This will take care of the NPE that resulted in the pending job.</p>

Issue ID	Description
CS-25158	<p>Problem: IP allocation fails due to subnet overlap.</p> <p>Root cause: In version 3.0.7, VLAN ID is saved as 3000 in the database, whereas in version 4.x it is saved as vlan://3000. CloudPlatform considers these two as separate VLANs.</p> <p>Solution: Run the following MySQL statement when upgrading to 4.3.0.2</p> <pre data-bbox="767 607 1348 748">mysql> update vlan set vlan_id=concat('vlan://', vlan_id) where vlan_type = "VirtualNetwork" and vlan_id not like "vlan://%";</pre>
CS-25153	<p>Problem: VMs do not restart with more than 7 NICs on VMware hosts.</p> <p>Root cause: The vSphere API, ConfigureVM, behaves abnormally if NIC devices are numbered from 0. This is an API limitation elaborated in the KB article 2088490.</p> <p>Solution: While creating a VM, NIC numbering in the configuration is started from 7 instead of 0. This is due to an undocumented behaviour of vSphere API elaborated in KB article 2088490</p>
CS-25012	<p>Problem: A PF rule cannot be set when more than 20 VMs exist in a network.</p> <p>Root cause: UI issue. The listVirtualMachines API call that is issued has page size set to 20. It is observed that this issue affects all VMs that do not appear on page 1.</p> <p>Solution: Fixed the UI.</p>
CS-24971	<p>Problem: Creating snapshot from volume is failed when the task is performed repeatedly in a zone-wide primary storage.</p> <p>Root cause: While taking snapshot of a volume, CloudPlatform determines which storage pool contains the volume. CloudPlatform then chooses the endpoint to perform snapshot backup operation by selecting any host that has the storage pool mounted on it. In case of zone-wide primary storage, since every host present in the zone will have the storage mounted, a random host will be chosen. If the host chosen is not in the cluster that contains the VM that the volume is attached to, the snapshot operation will fail.</p>

Issue ID	Description
	<p>If a host that belongs to a different cluster is chosen as the endpoint, the following happens in CloudPlatform VMware resource layer: During snapshot creation of a volume, first snapshot of the VM that the volume is attached to is taken. In case a host that doesn't contain the volume is chosen as the endpoint then the look-up for the VM in the cluster is negative. Because the VM cannot be found, a worker VM is created to which the disk is attached, whose snapshot has to be taken (data path for detached volumes). This operation fails with an error in vCenter.</p> <p>Solution: While taking a snapshot of a volume, CloudPlatform chooses the endpoint to perform backup snapshot operation by selecting any host that has the storage containing the volume mounted on it. Instead, if the volume is attached to a VM, the endpoint chosen by CloudPlatform should be the host that contains the VM.</p>
CS-24473	<p>Problem: Restarting VR on an isolate network with egress policy set to allow does not honor the default rule.</p> <p>Root cause: After VR reboot in a network with no egress rule created, the IPtables rules are not created to allow egress traffic.</p> <p>Solution: After VR reboot, rules are configured on VR to add egress default rules to allow traffic.</p>
CS-24057	<p>Problem: Message Bus SSL and Virtual Host configuration results in Management Service failure.</p>
CS-24056	<p>Problem: Attaching a local data disk to a VM does not work as expected.</p> <p>Root cause: Storage migration of a VM with volume on local was failing. When a plan with hostid included was passed to the local storage pool allocator, it returns all the local storage pools in the cluster, instead of just the local pool on the given host in the plan. This occurs because the search at the host level was taking place only for data disk. Additionally, the query to list the storage pools on a host is failing if the pool has tags. If no tags are used for storage pool, 0 pools are returned to LocalStoragePoolAllocator by the method PrimaryDataStoreDaoImpl::findLocalStoragePoolsByHostAndTags(). This occurs due to the Inner join used between storage_pool and storage_pool_details.</p>

Issue ID	Description
CS-23097	<p>Solution: Handle case for no tags.</p> <p>Problem: VM cannot be created in shared and isolated networks.</p> <p>Root cause: The network reached the maximum public IP limit.</p> <p>Solution: Public IP resources are not checked when deploying a VM on a shared network.</p>
CS-23065	<p>Problem: Volume sync does not block agent connection.</p> <p>Root cause: During SSVM agent connect on restarting Management Server, or starting or stopping SSVM, template and volume are synced. However, volume sync should happen between the secondary storage volumes and the database records of the volumes during Management Server and SSVM agent handshake. Sync mismatch, if occurs, is not resolved and a NPE is thrown preventing SSVM agent to connect.</p> <p>Solution: The volume sync operation is disabled by commenting the call to it to fix the defects on volume sync.</p>
CS-23064	<p>Problem: CHMOD with recursive option is timed out on large CloudPlatform 4.3.x installations.</p> <p>Root cause: The 'chmod -r' may take very long time on a slow network connection. It probably happens when management server and secondary storage are geographically separated.</p> <p>Solution: Run 'chmod' without '-r' option.</p>
CS-22856	<p>Problem: During the host connect VM sync reports that the VM state is stopped on the host and then proceeds with stopping the VM.</p> <p>Root cause: The earlier version of the VM sync operation was not able to handle the out-of-band VM changes, such as by the DRS operation in VMware, which resulted in running VMs getting stopped incorrectly.</p> <p>Solution: New VM sync changes has been added to handle all such scenarios.</p>
CS-22760	<p>Problem: VMs are not expunged in CloudPlatform 4.3.0.1, resulting extra storage usage.</p> <p>Root cause: The thread local context which is used to handle connections to a vCenter is</p>

Issue ID	Description
	<p>being re-used intermittently. When it is reused, if the existing context corresponds to the other vCenter, then the operation being executed will fail. The reuse is happening because after the PingTask completion the ThreadLocal context is not cleared.</p> <p>Solution: Clear the ThreadLocal context upon completion of the PingTask. Then, add validation to ensure that if ThreadLocal context is being reused, validate that the context corresponding to the right vCenter API session.</p>
CS-22756	<p>Problem: Data disk attach failed for VMs in zone-wide primary storage if there are mixed storage pools in zone.</p> <p>Root cause: When you copy or create Volume, storage allocators are called to get the appropriate pool. Firstly the ClusterScopeStoragePoolAllocator is called to check for podId, and if pod is null then ZoneWideStoragePoolAllocator is used. In this case pod passed to ClusterScopeStoragePoolAllocator is not null; therefore, ClusterScopeStoragePoolAllocator is allocating the pool for Data disk which is resulting in failure later. The pod which is getting passed to ClusterScopeStoragePoolAllocator is not null because VM's pod id is passed which is obviously not null and resulting in pool getting allocated from the cluster scope.</p> <p>Solution: Use pool's pod ID instead of VM's pod ID.</p>
CS-22660, CS-22378	<p>Problem: FTP inbound service is not working in isolated networks. Creating ingress rules for FTP on public IP and accessing FTP service in VM from the public network is failed.</p> <p>Root cause: FTP connection tracking modules are missing in VR, due to which the FTP data connections are failing.</p> <p>Solution: Load nf_nat_ftp module in VR by using the modprobe command. The nf_nat_ftp and nf_conntrack_ftp modules are loaded.</p>
CS-22533	<p>Problem: ClassCastException in VirtualMachineManagerImpl in handling various Agent command answer.</p> <p>Root cause: Exception is thrown when host is down or disconnected, where agent Command cannot even reach ServerResource, and thus</p>

Issue ID	Description
	<p>a plain Answer object is returned instead of specific Answer instance.</p>
CS-22451	<p>Problem: User is not logged in to the UI, even though Management Server log indicates user being logged in.</p> <p>Root cause: The java RabbitMQ AMQP client library used in CloudPlatform is an older version.</p> <p>Solution: The java RabbitMQ AMQP client library is updated to the latest version.</p>
CS-22450,CS-22449	<p>Problem: When an exception occurs in a publishing event to message bus, the thread will hold up the message bus lock without releasing it, this will lock the message bus and block all other jobs from acquiring the lock and publishing event when they are complete, so many jobs will show as in progress forever even though the action involved in the job is already finished.</p> <p>Root cause: MessageBus routine does not use try-catch-finally pattern to always release the lock acquired before try block, thus holding up the lock and causing other threads waiting to acquire the lock forever.</p> <p>Solution: Fix MessageBus class to always use try-catch-finally pattern and always release the lock in finally block.</p>
CS-22410	<p>Problem: Resizing detached volumes fails with errors.</p> <p>Root cause: In case of VMware, for disk re-size we use the VM associated with the disk to perform the re-size operation. For a detached volume since there is no associated VM the re-size failed with 'VM does not exist' error.</p> <p>Solution: For a detached volume, create a worker VM, attach the volume to be resized to this VM and perform the resize operation. Once the re-size is complete, detach the volume from the worker VM and destroy it.</p>
CS-22396, CS-22395	<p>Problem: HA takes long time to trigger.</p> <p>Root cause: Non-hypervisor specific investigators, such as PingInvestigator takes a considerable time, whereas KVMInvestigator could determine that host is down.</p>

Issue ID	Description
	<p>Solution: Placing hypervisor-specific investigators ahead, host status can be determined faster. This triggers HA faster.</p>
CS-21901	<p>Problem: If multiple SSVMs exist in a zone, downloading template fails because the URL which is generated goes to the wrong SSVM.</p> <p>Root cause: Wrong download URL is generated when using multiple SSVMs in a zone. The public IP of the URL would sometime point to the wrong SSVM when the URL was created on another one.</p> <p>Solution: Fix the bug by removing the command <code>CreateEntityDownloadURLCommand</code> from the host delegation. This results in same SSVM for creating the symlink on SSVM and the same public IP being used for generating the URL on Management Server.</p>
CS-21451	<p>Problem: Provide virtual host support in CloudPlatform.</p> <p>Root cause: Earlier AMQP integration in CloudPlatform assumed that the entire AMQP server is accessible, which might not be valid in a production environment. Therefore, CloudPlatform should support virtual host concept supported by AMQP servers.</p> <p>Solution: Added code to support virtual host concepts of RabbitMQ. To enable this feature, set 'virtualHost' property in the <code>spring-event-bus-context.xml</code> file at <code>/usr/share/cloudstack-management/webapps/client/WEB-INF/classes/META-INF/cloudstack/core</code>. with an appropriate virtual host name. Username and password specified should be that of the virtual host. For example:</p> <pre data-bbox="853 1579 1439 2033"> <beans xmlns="http:// www.springframework.org/schema/beans" xmlns:xsi="http://www.w3.org/2001/ XMLSchema-instance" xmlns:context="http:// www.springframework.org/schema/context" xmlns:aop="http://www.springframework.org/ schema/aop" xmlns:util="http://www.springframework.org/ schema/util" xsi:schemaLocation="http:// www.springframework.org/schema/beans http://www.springframework.org/schema/ beans/spring-beans-3.0.xsd http://www.springframework.org/schema/aop </pre>

Issue ID	Description
	<pre> http://www.springframework.org/schema/aop/ spring-aop-3.0.xsd http://www.springframework.org/schema/ context http://www.springframework.org/schema/ context/spring-context-3.0.xsd http://www.springframework.org/schema/util http://www.springframework.org/schema/util/ spring-util-3.0.xsd" > <bean id="eventNotificationBus" class="org.apache.cloudstack. mom.rabbitmq.RabbitMQEventBus"> <property name="name" value="eventNotificationBus"/> <property name="server" value="127.0.0.1"/> <property name="port" value="5672"/> <property name="username" value="temp"/> <property name="password" value="password"/ > <property name="virtualHost" value="temp"/> <property name="exchange" value="cloudstack-events"/> </bean> </beans> </pre>
CS-21373	<p>Problem: Templates with the same name have the same unique_name in the database.</p> <p>Root cause: Creating templates with the same names from volumes is deleted.</p> <p>Solution: Change the unique name generation to the standard one used by registertemplate.</p>
CS-21224	<p>Problem: If a VMware datacenter has a space in it's name, adding host to an existing VMware cluster (of that datacenter) fails if it is already being managed by CloudPlatform.</p> <p>Root cause: While adding host CloudPlatform attempts to insert the validated inventory URL path in the database. However, this inserts encoded URL into database implies whitespace is stored as '+' symbols. The URL from the API parameter string is being converted to URI object as part of validation, where the URL path is getting encoded.</p> <p>Solution: Skip updating the cluster URL in the cluster_details table while adding a host to the existing cluster.</p>
CS-21220,CS-20844	<p>Problem: SourceNAT,StaticNAT and Portforwarding is not working with VMware DVS.</p> <p>Root cause: Change in vCenter 5.5 API from prior versions forced code change in CloudPlatform.</p>

Issue ID	Description
	<p>Solution: Update property value of the property, VirtualE1000.deviceInfo.summary, is accommodated.</p>
CS-21107	<p>Problem: Adding a host to existing VMware cluster with Nexus 1000v as a network backend fails.</p> <p>Root cause: While adding a host or cluster, CloudPlatform attempts to validate session authentication of associated Nexus VSM. This failed in case of adding host whereas it works fine for case of adding cluster. Reason is unavailability of Nexus VSM metadata in the code path that gets executed when a host is added to existing cluster. In case of addCluster, the API call provides Nexus VSM metadata including credentials that would be useful for session validation, hence that works fine.</p> <p>Solution: Check if a Nexus VSM associated with this cluster already is validated and exists in the database, otherwise perform the validation which would be the case of addCluster API call. It is not required to perform session authentication validation of Nexus VSM of the cluster while adding host to an existing cluster.</p>
CS-21077	<p>Problem: Host remains in Alert after restarting vCenter.</p> <p>Root cause: When a vCenter goes down, Management Server loses connection to the vCenter. At this point, since Management Server can't determine the state of any of the hosts, it puts all hosts into Alert state. This is because when Management Server can't determine the host state and it has to do an investigation to determine the real state of the host, the host is put into Alert state.</p> <p>Solution: add appropriate synchronization to ensure that the PingMap running in the Management Server doesn't add back an agent into the PingMap once the agent has been disconnected from the Management Server.</p>
CS-21073	<p>Problem: CloudPlatform enters host name thrice in the addHost string.</p> <p>Root cause: The AddHost API command fails if either the VMware datacenter name or the VMware cluster name contains spaces or special characters, implies, if the URL encoded value</p>

Issue ID	Description
	<p>for the VMware datacenter or VMware cluster is different from the non-decoded value.</p> <p>Solution: The URL decoding of vCenter path during cluster discovery before trying to find a match.</p>
CS-20970	<p>Problem: In a VMware setup with more than 200 ESX hosts, after running Management Server for several hours, system stops responding. The log gives the "OutOfMemory: cannot create native thread" error.</p> <p>Root cause: The ping task in CloudPlatform frequently runs to get the current status of the host, where VmwareContext has to be initialized to connect to vCenter. Since CloudPlatform 4.2 onwards, code has been modified to use VmwareContextPool to manage VmwareContext. However, in this ping task, after using the VMware context, it has not been released back to the pool to be reused by other operations, causing VmwareContext leak. Since each VmwareContext connecting to vCenter will use system thread handle, which leads to thread handle shortage, causing error.</p> <p>Solution: Fixed VmwareContext leak in the ping task. However, modify the ulimit settings to increase the max user threads for large VMware deployments:</p> <pre>stack size (kbytes, -s) 10240 max user processes (-u) 12288</pre>
CS-20962	<p>Problem: Netscaler VPX cannot be added to CloudPlatform.</p> <p>Root cause: The Netscaler VPX forces an idle timeout of 100,000 seconds (~1 day and 4 hours) during login. The request is to remove the default timeout value. An API change also caused this issue.</p> <p>Solution: The "setting timeout option" has been removed from the code.</p>
CS-20942	<p>Problem: VMFS local storage does not support over provisioning. With a XenApp workload VM is deployed with large volumes. On the backend there is plenty of storage available as VMware is thin provisioning; however, CloudPlatform does not allocate any more volumes as it does not provision past actual capacity.</p>

Issue ID	Description
	<p>Root cause: CloudPlatform applies the overprovisioning factor when a storage gets added and accordingly the capacity entry for that storage reflects the overprovisioned capacity. Changes are made to apply this factor to VMFS storage pools along with NFS. However, the local VMFS storage gets added to CloudPlatform under the poolType 'LVM' which is used for XenServer. This causes CloudPlatform to not consider the storage for overprovisioning.</p> <p>Solution: When a local VMFS storage is added to CloudPlatform, set the poolType correctly to VMFS. This will make the entry in op_host_capacity table reflect over provisioned capacity.</p>
CS-20940	<p>Problem: A network is being incorrectly shut down by the garbage collector in spite of having running VMs in the network.</p> <p>Root cause: CloudPlatform runs a "network garbage calculator or network GC" daemon at periodic intervals. Its responsibility is to find the isolated networks which has no user VMs running in the network. For each such network, it will destroy the VR and shuts down the network. To figure out the networks that are ready for network GC, Management Server keeps track of running user VMs in a network using the NIC count, stored in op_networks table. The count is incremented when a VM is launched and decremented when VM is stopped. When DRS is enabled in ESXi clusters, VMs are migrated across the hosts that results in vCenter giving power state of the VM to be 'off' temporarily. CloudPlatform has a component VMSync which track the state of VMs. When a state change occurs outside CloudPlatform, Management Server take a corrective actions internally. So when a VM is reported as 'powered off' by vCenter, Management Server decrements the NIC count. When vCenter reports VM as 'powered on' once migrated, there is no corresponding NIC count increment. Which leads to NIC count ending up as 0, even though there are VMs running and eventually network GC destroys the VR and network is shut down.</p> <p>The fix is not to depend on the NIC count stored in the op_network table to decide if network needs to be garbage collected. But find the actual number of running VMs in a network, and only when it is zero, the network GC decides to</p>

Issue ID	Description
	<p>destroy the VR and shuts down the network. As a result even though vCenter reports temporary states as VM is powered off, which result in NIC count going wrong, because CloudPlatform does not use NIC count any more there will not be network GC if the network has at least one VM running.</p> <p>Solution: The dependency on NIC count has been removed in deciding which network to be garbage collected. All the isolated non-persistent networks are checked if there are any active NICs (NICs responsible for running or starting VMs) by running a database query, instead of using NIC count in the op_network table.</p>
CS-20936	<p>Problem: Template/snapshot download is extremely slow at 2 Mbps from Secondary Storage VM.</p> <p>Root cause: The template download java code syncs each write operation during the download process, thus the download speed is very slow.</p> <p>Solution: Don't sync each write operation, but sync only after the download is finished. This would speed up download speed on par with wget command.</p>
CS-20881	<p>Problem: On VMware and KVM, if a snapshot creation is failed, then all the snapshot creation operations on the same volume are failed, consequently.</p> <p>Root cause: Snapshot creation has two stages: First, creating snapshot on primary storage. If it's succeed, snapshot state is changed to "CreatedOnPrimary", otherwise, the state will be "Error". Secondly, backing up snapshot from primary storage to secondary storage. If succeeds, the state is "BackedUp", otherwise, the state stuck at "CreatedOnPrimary". When other snapshot creation operations coming in on the same volume, Management Server mistakenly consider that another ongoing snapshot operation is running, because, it can only take snapshot no snapshot status is in "CreatedOnPrimary/BackingUp" state on this volume. Subsequently snapshot operation on the volume fails. This error occurs only on VMware and KVM because they check the status of snapshot before taking a new snapshot.</p>

Issue ID	Description
	<p>Solution: In case of backing up snapshot failure, snapshot state is changed to "Error", instead of "createdonPrimary".</p>
CS-20872	<p>Problem: Storage motion feature was broken when vmwork job queue changes were introduced. Gson serialization was working for volume to pool mapping. This caused storage motion of a VM with volume to fail.</p> <p>Root cause: Migration of a virtual machine with its volume was broken because management server wasn't able to serialize the volume to pool mapping. The mapping is used to identify which volume should be migrated to which storage pool. This issue wasn't specific to volumes on local storage pool but also affected volumes on shared storage pools.</p> <p>Solution: The volume to pool mapping object was updated so that it can be serialized and migration of a virtual machine with its volume can take place.</p>
CS-20840	<p>Problem: For VMware, creating a template from a large snapshot is failed.</p> <p>Root cause: VMware resource code doesn't honor the timeout, the <code>create.private.template.from.snapshot.wait</code> parameter in the global configuration, send by the Management Server while creating template from snapshot. By default, VMware resource times out in 1440 seconds; therefore, if the snapshot is large the issue is triggered.</p> <p>Solution: The cloudstack java agent running inside SSVM, needs to honor the timeout information send from Management server, during creating template from snapshot.</p>
CS-20826	<p>Problem: LDAP connection timeout is hardcoded to 500 minutes.</p> <p>Solution: A new global parameter, <code>ldap.read.timeout</code>, has been added to configure this value.</p>
CS-20788	<p>Problem: VMs using local storage cannot be live migrated. When you attempt to migrate a VM with its volume on local storage pool, it fails. CloudPlatform picks up the wrong destination storage pool for the volume as the pool was not available on the destination host.</p>

Issue ID	Description
	<p>Root cause: A defect in the LocalStoragePoolAllocator causes to return all the local storage pools in the cluster, instead of returning just the local pool on the given host in the deployment plan, when the allocator was queried for local storage pools on a host. When CloudPlatform initiates migrating storage with the root volume on local, system picks the wrong local storage pool that which wasn't available on the given host. This caused Storage Resource not available exception when hypervisor tries storage migration.</p> <p>Solution: Fixed the local storage pool allocator to right the correct storage pool which was available on the destination host.</p>
CS-20770	<p>Problem: LDAP Group imports does not work for Microsoft ActiveDirectory.</p> <p>Root cause: If the group has only one user, import works as expected, but if group has other groups as members then importing users fails with Nullpointer exception.</p> <p>Solution: Handled the case when the member of a group is not an user.</p>
CS-20747	<p>Problem: VOLUME.DELETE usage event is missing for VMs in ERROR state.</p> <p>Root cause: When deploying a VM is failed, database entry for this VM is updated to ERROR and the volume entry to Destroyed. However, publish VOLUME.EVENT event is not published in invoking internal destroyVolume command. Instead, it scatters VOLUME.EVENT publish code around in the service layer code. Because publish VOLUME.EVENT is not called in this particular case, no VOLUME.EVENT is published when deploying a VM is failed.</p> <p>Solution: Refactor the code to always publish VOLUME.EVENT event in invoking internal destroyVolume routine to mark a volume as Destroyed.</p>
CS-20718, CS-20716, CS-20715	<p>Problem: Japanese keys are incorrectly mapped on CentOS6.5 x64</p> <p>Root cause: Japanese keys mapping in the javascript key mapping file has been incorrectly mapped. The mapping of keys to the output characters in the VM console is incorrect.</p>

Issue ID	Description
	<p>Solution: The JP keyboard mapping issue has been fixed for Windows or CentOS VMs on VMware hypervisor. To work with Japanese keyboard on VM console proxy window, ensure that VMs are deployed with "keyboard=jp" parameter and the keyboard language selected in the console proxy keyboard dropdown menu is Japanese.</p>
CS-20670, CS-20711	<p>Problem: Consider disk-chain renaming behavior in vCenter while locating disk-chain by name.</p> <p>Root cause: When vCenter performs snapshot operation, it does so by stacking a delta disk on top of the chain; the name tracking in various orchestration flow may run into out-of-sync situations. Consider this disk-chain renaming behavior in vCenter while locating disk-chain by name.</p> <p>Solution: While searching for a disk in vCenter by matching against its name, trim the postfix appended to the disk name by vCenter after snapshot operation.</p>
CS-20701	<p>Problem: Reset VM fails if a VM snapshot exists in the system.</p> <p>Root cause: When a Reset VM operation called on a VM that has VM snapshots associated with it, the operation fails with an exception.</p> <p>Solution: If Reset VM operation is invoked for VMs that have VM snapshots associated with them, ensure that CloudPlatform gracefully fails the operation with an appropriate error message.</p>
CS-20629	<p>Problem: Upgrade to version 4.3.0.1 fails if a VMware setup has multiple zones managing the same VMware datacenter.</p> <p>Root cause: As part of the new mapping model for CloudStack zone and VMware datacenter, CLOUDSTACK-1963¹, each of the VMware datacenter is populated into either cloud.vmware_datacenter or cloud.legacy_zones table during upgrade. However, the upgrade script doesn't handle a VMware Datacenter that is managed by two different zones and hence the upgrade fails.</p>

¹ <https://issues.apache.org/jira/browse/CLOUDSTACK-1963>

Issue ID	Description
	<p>Solution: Updated the upgrade script to handle this deployment model.</p>
CS-20612	<p>Problem: System VMs are failed to in a zone which has Nexus 1000v as backend for public/guest traffic.</p> <p>Root cause: During deployment of system VMs CloudPlatform needs to talk to vCenter by establishing a session. This session object was stored in VmwareContext object. The object is per each host. Nexus 1000v credentials are stored in the session context that was created for this specific host. This works until there is a 1:1 association between host and session. Because in version 4.2, pooling has been introduced for VmwareContexts, which implies CloudPlatform performs recycle and reuse. Each time a session object is received from pool, there is no guarantee that the same object is received where Nexus 1000v credentials are stored previously. Therefore, VSM credentials stored in session context cannot be retrieved always correctly.</p> <p>Solution: Fix is to register the VSM credentials after fetching context and the context is recycled after use. This implies that a session context is fetched it might not be the same one that you received from the previous fetch attempt.</p>
CS-20606	<p>Problem: Default value of XenServer "Max guest limit" is not honoured by CloudPlatform, implies that the value set in default field of Max guest limit for XenServer has no effect on XenServer which are added to CloudPlatform before setting the value. Only those XenServer hosts which are added after setting the value listen the new value.</p> <p>Root cause: The logic to compare the number of running VMs to the max limit was not correct.</p> <p>Solution: The logic has been corrected to compare the limit.</p>
CS-20600	<p>Problem: While creating a snapshot, only volume ID is mentioned in the events. For example, "Scheduled async job for creating snapshot for volume Id:270". On looking into the notification, volume cannot be identified. Provide the volume name or UUID in the events.</p> <p>Root cause: Earlier an effort has been made to replace all the occurrences of internal entity,</p>

Issue ID	Description
	<p>ID, with UUID. Events has been left out at that time; therefore, ID is continued to be used in the events.</p> <p>Solution: UUID is recorded instead of internal IDs in the event messages.</p>
CS-20534	<p>Problem: Users are confused because a quickview would stay open after hovering the mouse off the row, and assuming they are on the previous row, execute actions from the wrong quickview.</p> <p>Root cause: This was caused by the extra margin padding around the quickview.</p> <p>Solution: Remove the extra padding around the quickview, so that once the mouse is out of the box, it is hidden immediately.</p>
CS-20526	<p>Problem: Static NAT does not work after fail over in RVR. This issue exists only for the additional public subnets case.</p> <p>Root cause: For additional public subnet case in RVR, when fail over occurs there is no mechanism to add routes for the additional subnets in enable_pubip.sh When backup is switched to master, enable_pubip.sh is called, which brings up public interfaces and add routes for the interface. Due to this the ingress traffic coming in eth3 is going out via eth2 to add routes gateway and device information which is not available in the router dynamically.</p> <p>Solution: Once the gateway and device information is available in VR, you can add routes for the additional subnets. Therefore, the gateway and device information are maintained in /var/cache/cloud/ifaceGwIp in VR. Using this information routes for additional public subnet interfaces are added in enable_pubip.sh when VR switches to master.</p>
CS-20525	<p>Problem: On an additional public subnet, removing public IP does not delete SNAT rules. The issue is observed when the first IP is added, but while removing the first IP also is removed.</p> <p>Root cause: For additional public subnets (non-sourceNAT network) the first one is selected as the first IP from the list, which is retrieved from the database. When you have few IPs the delete/add operations on them changes the order of the IPs. A static NAT rule is configured on the last IP first, and later few static NAT rules are</p>

Issue ID	Description
	<p>configured on other IPs. While removing, the last IP is removed first so that it is not selected as the first IP. Therefore, the SNAT rules configured are untouched on the VR. While adding, source NAT rules are added for the first IP.</p> <p>Solution: To delete SNAT rules, while disabling static NAT on IP from the non source NAT network, set the source NAT flag to true. This is to make sure that the SNAT rules are got removed.</p>
CS-20500	<p>Problem: A VMware cluster cannot be added in the legacy Zone in 4.3.0.1.</p> <p>Root cause: While adding a VMware cluster, CloudPlatform check if the zone is legacy or not. To check this legacy_zones table is checked with the key, id, instead of zone_id, which is yielding a wrong row resulting in a legacy zone to show up as a normal zone.</p> <p>Solution: The field zone_id is queried in table legacy_zones to check if a zone is legacy zone or not.</p>
CS-20487	<p>Problem: While attempting to destroy a VM, it goes to Stopped state.</p> <p>Root cause: This was caused when DestroyVM command is issued while VMsync operation is in progress. The VMsync operation discovers that the database state is changing to stopped state and hypervisor state is running. Therefore, it attempts to stop the VM. The DestroyVM thread, after getting a response from the stopvmcommand, attempts to put the VM in stopped state, but does not succeed because the transition has already been made.</p> <p>Solution: With the new design for VM sync is in place, CloudPlatform does not act on the VM state change with VM sync if VM state change operations are in progress ; therefore this issue would not occur.</p>
CS-20482	<p>Problem: The snapshot storage usage seems too high, consuming around 27 TB.</p> <p>Root cause: Snapshot storage usage is not being computed correctly. The vm_snapshot_chain_size is calculated for each volume that belongs to a VM by using the getVMSnapshotChainSize() method. The issue while computing the snapshot chain size is that the file associated with a volume is searched</p>

Issue ID	Description
	<p>in the entire datastore instead of looking for it inside the VM folder in the datastore. Because no match is found on the entire VMDK path, in case of a ROOT volume, the sizes of all the ROOT volumes present in the datastore are added up, instead of just the ROOT volume that belongs to the VM.</p>
CS-20479	<p>Problem: VMs created prior to CloudPlatform 4.2.1 does not connect to the guest network if the NIC adapter used by the VM is other than the default one, which is E1000.</p> <p>Root cause: From 4.2.1 onwards only while creating guest VMs, template-level setting of nicAdapter is used and nicAdapter property persists in user_vm_details table such that all subsequent start operation would use it rather than depending upon template-level setting of nicAdapter. This implies that the instances created before 4.2.1 would not have any information about nicAdapter present in the user_vm_details table, which results in switching to default nicAdapter, E1000, upon a stop and start operation post upgrade. Prior to 4.2.1 guest VMs used tonicAdapter.</p> <p>Solution: This needs data migration during upgrade. The workaround is to update the user_vm_details table so that nicAdapter property, which is from template_details table, whose value is not 'E1000' be stored.</p> <pre data-bbox="858 1346 1439 1480"># insert into cloud.user_vm_details (vm_id,name,value,display_detail)VALUES (<VM_ID>,'nicAdapter', <nicAdapterPrptyfrotemplatedetailstable>,1);</pre> <p>For example:</p> <pre data-bbox="858 1570 1439 1682"># insert into cloud.user_vm_details (vm_id,name,value,display_detail) VALUES (1111,'nicAdapter','Vmxnet3',1);</pre>
CS-20465	<p>Problem: When time is synced by Network Time Protocol daemon(NTPD), Redundant VR may result in FAULT state.</p> <p>Root cause: NTPD in VR would move time backwards to keep sync with the NTP server, which results in false alarm of keepalived monitoring process.</p>

Issue ID	Description
	<p>Solution: Three strikes rule are added for Redundant VR freezing detection.</p>
CS-20458	<p>Problem: Resource limits are not honored upon upgrade from CloudPlatform 3.0.7 to version 4.3.0.1.</p> <p>Root cause: If the resource limits is set 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 configuration parameter, which is set to 20.</p> <p>Solution: To make these resource limits unlimited, set the limits to -1 for these newly added resource types in the upgraded setup.</p>
CS-20452	<p>Problem: When a legacy zone is selected in the UI, listvmwaredcs API is called in and for a legacy zone it throws the following error: 'Invalid zone id error while listing VMware zones error'.</p> <p>Root cause: The listVmwareDcs API called by the UI throws the error.</p> <p>Solution: When listVmwareDcs API is called remove the validation to check if the zone is legacy or not.</p>
CS-20352	<p>Problem: The default conntrack max set on the router VM is quite low. Due to this a lot of packet drops occurs. When more number of connections are made VR is not able to handle the connections, and therefore connections are dropped.</p> <p>Root cause: Setting the <i>ip_conntrack_max</i> value in the <i>sysctl.conf</i> file is failed. Value for the <i>ip_conntrack_max</i> is limited to 32000; however, the system would be capable of handling additional connections.</p> <p>Solution: The value of <i>proc ip_conntrack_max</i> has been increased.</p>
CS-20273	<p>Problem: Usage service fails to start after updating to java 1.7</p> <p>Root cause: The java path on that server is <code>/usr/lib/jvm/java-1.7.0-openjdk-1.7.0.55.x86_64/jre/</code></p>

Issue ID	Description
	<p>bin/java. The woraround is to add the java_home and JDK_DIRS in /etc/init.d/cloudstack-usage.</p> <p>Solution: The hard coded jdk directory setting and java home are removed by using readlink and dirname.</p>
CS-20259	<p>Problem: HA does not occur when the uploaded volume is attached to the VM.</p> <p>Root cause: For the uploaded volume, the spoolType property was not getting set, causing an NPE during the HA process.</p> <p>Solution: Set the poolType and added a check in the code to avoid the NPE.</p>
CS-20252	<p>Problem: A network offering cannot be created.</p> <p>Root cause: The createNetworkOffering API call exceeds limit of API call size because too many parameters need to be passed when creating a network offering.</p> <p>Solution: No parameter are passed whose value entered on UI happens to be the same as its default value at the server-side.</p>
CS-20243, CS-19859	<p>Problem: Re-copying templates to other zones doesn't work.</p> <p>Root cause: There are two factors to this issue: first, when a template is recopied the same entry is updated in the template_zone_ref and the removed column is set to null. In order to set this to null DAO layer is used. DAO layer does not permit setting the removed column to null, and therefore this is always set to a time stamp. Second, when copying a template is failed the state in the template_store_ref is not properly updated. The entry created while downloading the template is left as it is without cleaning it up. Additionally, in version 4.2.1, the list templates API does not consider the destroyed flag in the templates store ref. From versions 4.3 onwards, the destroyed flag is not used, instead, two new states active and inactive namely were introduced, so the list API fix was not needed in the later version.</p> <p>Solution: Updated the removed column of the template_zone_ref table. The stale entries are cleaned up when retrying the template copy. The list templates API has been updated to consider the destroyed flag.</p>

Issue ID	Description
CS-20175	<p>Problem: A network cannot be deleted. A VM is assigned to a NIC but later it has been removed. You cannot delete the network later, complained VM is still running, though there is no NIC in the VM belong to the network.</p> <p>Root cause: When checking for NICs to be removed, CloudPlatform counted removed NICs as well before checking the VM's state for removed NIC, which is incorrect.</p> <p>Solution: Available NICs are searched only while checking if it's safe to delete a network.</p>
CS-20168, CS-20461	<p>Problem: The value of vCPUs should not be higher than the XenServer specified limit.</p> <p>Root cause: When dynamic scaling is enabled and a non-Windows VM is created on XenServer, <i>vcpu-max</i> parameter is set to 32, which is above the supported XenServer limits. The supported limit for vCPUs is 16 for Linux VMs.</p> <p>When dynamic scaling is disabled and a non-Windows VM is created on XenServer, <i>vcpu-max</i> value is set to 32, which is not required. The value set for <i>vcpu-max</i> should be equal to the vCPU value set in the service offering when dynamic scaling is disabled.</p> <p>Solution: A new configuration parameter, <i>xen.vm.vcpu.max</i>, has been introduced with the default value of 16. When dynamic scaling is enabled and non-Windows VMs are created on XenServer, <i>vpcu-max</i> value is set to the value specified in the <i>xen.vm.vcpu.max</i> parameter. When dynamic scaling is disabled and a VM is created, values of <i>vpcu-max</i> and <i>vcpu-startup</i> are set to the value equal to the vCPU value specified in the service offering.</p>
CS-20164	<p>Problem: The listCapacity API has missing types for certain zones.</p> <p>Root cause: The SQL query which gathers this information is incorrect. The query groups on the <i>capacity_type</i>. Therefore, all the capacities of a particular type are added to and returned as the capacity for one particular zone.</p> <p>Solution: The SQL query and the associated method have been modified.</p>

Issue ID	Description
CS-20162	<p>Problem: Attaching a datadisk for VMs that have VM snapshots results in 'Unexpected exception'.</p> <p>Root cause: When an attempt is made to attach datadisks to VMs with snapshots an 'Unexpected exception' is thrown, instead of reflecting the correct error message. While attaching a disk to a VM, the input parameter validation, which includes verifying if the VM has any snapshots associated with it, occurs inside the Job Queue framework. The error that is thrown inside the framework is not being handled correctly.</p> <p>Solution: As with other API commands, ensured that the input parameter validation for AttachVolume happens outside the VM job queue.</p>
CS-20104	<p>Problem: Creating storage pool operation is failing if path of VMFS datastore contains space.</p> <p>Root cause: While creating a storage pool, path to the storage pool is being stored in encoded form. This results in any white space in storage pool path is getting stored in the storage_pool table in the cloud database, as '+'. Operations, such as validating storage pool which uses the path, which read from the database are failing with invalid path or no such object exists.</p> <p>Solution: The decoded path of storage pool is stored.</p>
CS-20103	<p>Problem: Creating multiple-core VMs.</p> <p>Root cause: Enhancement request.</p> <p>Solution: Allow creation of an instance on VMware with multiple cores per socket. If a template has been registered and "cpu.corespersocket=X" template details has been added for it, then any instance deployed from that template should have X cores per socket.</p>
CS-20101	<p>Problem: Usage is generated for volumes even after the volume is destroyed and expunged.</p> <p>Root cause: When a VM is reset, old ROOT volume is deleted and a new volume is created for the VM. However, usage events for old volume deletion and new volume creation are missing, because of which the obsolete volume usage does not stop.</p>

Issue ID	Description
	<p>Solution: Volume usage events are added during VM reset.</p>
<p>CS-20097</p>	<p>Problem: Quickly attaching multiple data disks to a new VM fails.</p> <p>Root cause: When trying to attach multiple data disks to a VM in quick succession, AttachVolumeCmds consistently fails. This is only reproducible when attaching volumes to a VM in fast succession. In VMware the following steps are followed during disk attach to a VM: preparing the disk that needs to be attached. Re-configuring the VM to attach the prepared disk device.</p> <p>The first step involves figuring out what should be the device number on the controller key that the disk is connected to. These two steps are not synchronized in the code, yet. Therefore, when attaching two disks in quick succession, while preparing the second disk if the VM is still being re-configured with the first disk, that is second step is in progress for the first disk, then device number of the first disk will be chosen for the second disk too. Therefore, it results in Invalid configuration for device 'x' error where x is the device number on the controller key that the first disk is connected to.</p> <p>Solution: Synchronize the tasks of disk preparation and reconfiguration of a VM with the disk. This will ensure CloudPlatform doesn't try to attach a disk to a VM with a wrong device number on the controller key.</p>
<p>CS-20021</p>	<p>Problem: The action events should have project ID, when the current logged in user or account is part of a project.</p> <p>Root cause: The UI has not been passing projectid when creating and deleting a project.</p> <p>Solution: The events published on the rabbitMQ server has the project ID in the details if the account is part of a project.</p>
<p>CS-20010</p>	<p>Problem: Creating a VM by using jclouds fails with an exception.</p> <p>Root cause: In listProjectsInternal function, if domainId parameter and the caller domain ID matches, exception is thrown.</p>

Issue ID	Description
	<p>Solution: Set the correct condition for domainId verification and do not throw exception in case of equality of domains.</p>
CS-19973	<p>Problem: some action event messages have project ID and none of them have UUID and Event Types.</p> <p>Root cause: By design.</p> <p>Solution: The action events published on the rabbitMQ server has project ID in the message if the account is part of project.</p>
CS-19927	<p>Problem: When a network is restarted with cleanup=true in basic or advanced shared network, DNS of VMs in the network may not work.</p> <p>Root cause: The IP of the VR changes during recreating a network because it's not limited to the first IP in the network as is the case for Isolated or VPC networks. When this occurs, the DNS entries of old VMs are pointed to the previous VR's IP, resulting DNS failure.</p> <p>Solution: A placeholder NIC is added an IP is assigned to the VR for the first time in a basic or shared network. Use the same IP later whenever VR is recreated, thus keeping the same IP for the VR for the entire network life cycle.</p>
CS-19899	<p>Problem: When deploying a VM from a template that is available across zones, the Management Server randomly attempt to mount wrong secondary storage on XenServer, which in turn causes VM deployment failure.</p> <p>Root cause: Management Server attempts to mount wrong secondary storage on XenServer to copy template because no filter in the code that uses the zone ID; therefore, templates are randomly copied from the right secondary storage and sometimes from the wrong one leading to VM deployment failure.</p> <p>Solution: Use the filtering with zone ID to find the right template and secondary storage location for creating a VM.</p>
CS-19892	<p>Problem: Typing a pipe ' ' character in VM console returns a question mark and pipe '? '.</p> <p>Root cause: One key was mapped to two characters in the javascript key mapping file.</p>

Issue ID	Description
	<p>Solution: Fixed the issue so the keypress of ' ' key is printing the ' ' (pipe) character in the VM console.</p>
CS-19858	<p>Problem: Secondary storage limit is reaching maximum.</p> <p>Root cause: The issue is observed in an upgraded setup in which size of the snapshots taken before upgrade are still stored as zero in the database. The full snapshot size was getting updated in snapshot_store_ref table even if it was a delta snapshot and for used_secondary_storage_space calculation, CloudPlatform was summing up the sizes of all the delta + full snapshots from the previous mentioned table. Even though delta snapshots was using the space of full snapshot in sec storage, CloudPlatform was considering it as full snapshot in terms of size. Because of this sec storage was reaching the maximum limits.</p> <p>Solution: As a fix, now CloudPlatform stores the physical size of delta snapshots and using this physical size in resource utilization calculation.</p> <p>In an upgraded setup, perform the following:</p> <ul style="list-style-type: none"> • Update resource count for the ROOT domain by using the action buttons in the Domain detailView in the UI or by using the updateResourceCount API. • Size of the snapshots taken before upgrading to 4.3.0.2 are not updated and remains stored as zero in the database. This will cause inconsistency in secondary storage resource count and the actual secondary storage capacity. To avoid that, update the size and physical_size column in the snapshot_store_ref table, manually or by using a script with the physical size of snapshots specified in the secondary storage.
CS-19850	<p>Problem: There are not sufficient logs to diagnose LDAP integration failures. Add additional logs when encountering failures in integration with Active Directory.</p> <p>Root cause: Enhancement request.</p> <p>Solution: Exception logs and stacktrace have been added whenever an LDAP bind error occurs.</p>

Issue ID	Description
CS-19803	<p>Problem: KVM VMs don't start after upgrading to version 4.2.1-4.</p> <p>Root cause: The path prefix for a volume is added twice to the volume path, then CloudPlatform can't find the volume in primary storage.</p> <p>Solution: Add the path prefix only once for a volume, and make sure the volume path is right.</p>
CS-19793	<p>Problem: LDAP global configuration defaults to openldap values.</p> <p>Solution: LDAP global configuration defaults to Microsoft Active Directory values.</p>
CS-19726	<p>Problem: Multiple network with the same VLAN is created in a cluster.</p> <p>Root cause: This issue occurs because of a typo in the logic that handles the race condition for creating network in a XenServer pool, which causes multiple multiple XenServer networks to be created for the same network.</p> <p>Solution: The race condition for creating network has been correctly handled in XenServer pool.</p>
CS-19720	<p>Provide SSL support for EventBus so that the communication between the client in the Management Server and AMQP server is secure.</p>
CS-19716	<p>Problem: Timed out connection to vSphere server causes hosts to be disconnected. For a clustered environment using HAProxy, in case of network glitch, SSVM/CPVM agents will reconnect back to the Management Server, but may connect back to a Management Server node different from the one before network outage. The previous Management Server still keeps sending messages to those agents, causing "Channel is closed" error.</p> <p>Root cause: The issue is caused by the AgentManager code, which has an internal cache for agents managed by the Management Server. In case of disconnect caused by network outage and reconnecting back to a different Management Server node, CloudPlatform cannot robustly invalidate the old cache and the previous Management Server node still sends messages to the agents that are already not managed by it.</p>

Issue ID	Description
	<p>Solution: Before sending message to the agents, Management Server will check database to see which node is currently connected by the agents. If the node ID is changed, it will forward the message to the new Management Server.</p>
CS-19701	<p>Problem: VR is started even if any PF rules fail to apply during VR startup.</p> <p>Solution: The VR is stopped during PF rule failures.</p>
CS-19532	<p>Problem: Multiple threads are being used to collect the stats from the same VR.</p> <p>Root cause: The same threads are being instantiated by the multiple managers(VirtualNetworkApplianceManagerImpl and VpcVirtualNetworkApplianceManagerImpl) in the Management Server. This causes multiple threads to run for the network usage tasks.</p> <p>Solution: The duplicate task scheduled by VpcVirtualNetworkApplianceManagerImpl has been removed.</p>
CS-19251	<p>Problem: The IPtables rules occasionally fail to program on a XenServer host.</p> <p>Root cause: XenServer uses VM name, port, and chain name to compose IPtable chain name; the name often exceeds maximum character limit of 28, and cause IPtables programming failure.</p> <p>Solution: Checking IPtables chain name each time before applying the chain to IPtables; if the name length is found 28 or more, truncate it to less than 28.</p>
CS-19137	<p>Problem: When more than one SSVM exist in a zone for a VMware setup, downloading template gives wrong URL that is pointing to the wrong SSVM IP and therefore cannot be downloaded.</p> <p>Root cause: This is only an issue existing for VMware setup. CloudPlatform orchestration layer will pick a SSVM to send download template command and cache that SSVM IP in orchestration layer which is used to generate the download URL. However, before sending download template to SSVM, VmwareGuru will randomly pick a SSVM (in the case of multiple SSVMs) which may be different from the one picked by upper orchestration layer, thus the URL generated by orchestration layer will point to the wrong SSVM.</p>

Issue ID	Description
	<p>Solution: The cached SSVM IP is updated in the endpoint selector after VmwareGuru delegates the command to a different SSVM.</p>
CS-19067	<p>Problem: If remote VPN access is enabled on a particular IP address, CloudPlatform reports the following error: "A Remote Access VPN already exists for this public IP address".</p> <p>Root cause: A remote access VPN configured on public IP is failing in the backend because it was unreachable. The state entry, <code>remote_access_vpn</code>, in the database is marked as 'Running', which is the previous state. Due to this, the public IP is displayed as enabled in the UI.</p> <p>Solution: Returns success when no remote access VPN exists on a public IP.</p>
CS-18930	<p>Problem: When a RHEL 5.5 VM is deployed the OS type is marked as "Other" in the vCenter.</p> <p>Root cause: VmwareGuestOsMapper missed to correct the entries for RHEL 5.5 to 6.4 OS, and CloudPlatform uses that mapping table to pass correct information to invoke vCenter API.</p> <p>Solution: Added correct mappings for RHEL 5.5 up to RHEL 6.4 OS in VmwareGuestOsMapper.</p>
CS-18788	<p>Problem: Creating volume from custom disk offering does not work as expected. The volume size is incorrectly displayed.</p> <p>Root cause: The Management Server code takes a cache of the size of custom disk offering, so the subsequent custom disk offering will use the size of first created custom disk offering.</p> <p>Solution: Fixed the code.</p>
CS-18405	<p>Problem: The <code>job.cancel.threshold.minutes</code> parameter does not apply when the other wait time, such as <code>backup.snapshot.wait</code>, <code>copy.volume.wait</code>, <code>create.volume.from.snapshot.wait</code>, and <code>migrate.wait</code> is longer, because it would override others and force stop the job.</p> <p>Root cause: By design.</p> <p>Solution: The <code>job.cancel.threshold.minutes</code> is set on the API layer. Other wait time parameters, such as <code>backup.snapshot.wait</code>, <code>copy.volume.wait</code>, <code>migratawait</code>, and <code>create.volume.from.snapshot.wait</code> are set on</p>

Issue ID	Description
	<p>the AgentManager. They have no correlation and don't make call backs to each other. This is by design. However, when an API job is cancelled by some reason after it was sent to the AgentManager, this event never gets propagated to the agent side, and therefore the command on the agent doesn't get cancelled, and might still succeed on the backend. Because CloudPlatform does not prevent one parameter to be greater than the other; they act independently.</p> <p>Ensure that the <code>job.cancel.threshold.minutes</code> parameter should always be greater than the any waiting period set up on the agent side. This parameter applies only for the jobs that are being synced on a certain object on the API side. For snapshot, it can happen only when you perform per-host synching.</p>

3.2. Known Issues

Issue ID	Description
CS-16008	<p>In a clustered management server deployment, hosts are not load balanced across management servers in cluster. This is by design.</p> <p>Workaround: All Management server in cluster must be synced by running:</p> <pre data-bbox="762 1346 1350 1406"># ntpdate 0.xenserver.pool.ntp.org</pre> <pre data-bbox="762 1435 1350 1496"># service ntpd start</pre>
CS-16373	<p>[KVM] When a KVM cluster is taken to the Unmanaged state, then returned to the Managed state, the hosts do not come into the UP state.</p> <p>Workaround: Manually restart cloud-agent on the KVM hosts to bring up the hosts.</p>
CS-18561	<p>[VMware] After upgrading from 3.0.x to 4.2 and higher versions, restoring the existing VM which has an additional disk fails to boot.</p> <p>Workaround:</p> <p>If the <code>vmware.root.disk.controller</code> global parameter is set to <code>ide</code> in 3.0.x setup, after upgrade perform following:</p>

Issue ID	Description
	<ul style="list-style-type: none"> • Before performing any VM operations, such as start and restore, set <code>vmware.root.disk.controller</code> to <code>scsi</code>. • Restart the Management Server. <p>If <code>vmware.root.disk.controller</code> is set to <code>scsi</code> in 3.0.x setup, you need not change anything, because the controller setting is consistent across upgrade operations.</p>
CS-18616	Event messages should provide VM name along with VM ID when deleting VMs.
CS-18605	Order of templates and ISOs not honored by UI or API.
CS-18558	Version 4.2 does not show account information on UI for dedicated host.
CS-18743	[XenServer] VM state is incorrectly reflected in CloudPlatform if VM is deleted outside of CloudPlatform. In this case, the VM state is marked as Stopped in CloudPlatform. Depending on whether or not the on-disk information is still maintained, you may or may not be able to start it again in CloudPlatform.
CS-18834	[Hyper-V] More than 13 disks cannot be attached to a guest VM.
CS-18973	<p>[VMware] Volumes cannot be downloaded after SSVM is HAed. Download fails with the "Failed to copy the volume from the source primary storage pool to secondary storage" error.</p> <p>Workaround: Either remove the host that experiences the issue from the vCenter or bring it back.</p> <p>This issue is caused by limitations from vCenter when one of host is at disconnect or down state. If the host is at disconnect or down state in vCenter, vCenter will encounter an internal server error when it serves the URL request for file downloading and uploading operations to its datastores.</p>
CS-18991	HA rebooted several VMs while they were still running on a disconnected host.
CS-19109	Problem: Async response from <code>addAccountToProject</code> doesn't contain useful information. Generally async responses for commands that modify resources include the resource ID and description in the async response. This

Issue ID	Description
	<p>is not true for <code>addAccountToProject</code> or <code>deleteAccountFromProject</code>.</p> <p>Root cause: All the delete commands always return true/false value for success because at the end the object is removed (<code>deleteVolume</code>) or released (<code>disassociateIpAddress</code>), and there is nothing to return in the response. Therefore, the current output would suffice for the <code>deleteAccountFromProject</code>. For <code>addAccountToProject</code>, the response cannot be fixed at this point as it would break the API compatibility. Additionally, <code>resourceId</code> in this case will be ambiguous, as project and account resources are linked together.</p> <p>Solution: This is a known anomaly and there is no solution.</p>
CS-19177	CloudPlatform does not support external LB's private interface on a different network segment than the guest network.
CS-19259	Templates created from snapshots are not replicated to multiple secondary storage.
CS-19253	[XenServer] Discrepancy in the CloudPlatform and XenServer view of available memory.
CS-19285	<p>[VMware] When changes to a VM state is performed out-of-band, VR goes out of sync and is eventually shuts down.</p> <p>Workaround: Stop and restart the VM by using the CloudPlatform Management Server.</p>
CS-19405	[XenServer] <code>vm.instanceName.flag = true</code> has no effect when creating a VM.
CS-19530	Template ordering in the UI does not work as expected.
CS-19659	[Hyper-V] VRs might be force stopped when guest VMs are deployed across more than 20 isolated networks in parallel.
CS-19675	[VMware] In clusters with multiple primary storages configured VMs fail to restart when either Reset VM operation is performed or the compute offering has the Volatile option enabled.
CS-19685	VMware Distributed vSwitch is only supported for public and guest networks, but not for management and storage networks.
CS-19707	[VMware] Legacy Windows VMs cannot be restarted after attaching a DATA volume. This issue is observed only when the value for <code>vmware.root.disk.controller</code> is changed

Issue ID	Description
	<p>from <i>ide</i> to <i>osdefault</i>, which in turn results in losing the previous controller information.</p> <p>Workaround: Update the <code>user_vm_details</code> table such that the information about the previous controller before changing to <code>osdefault</code> is persisted in database. Sample query:</p> <pre data-bbox="858 510 1439 622">#insert ignore user_vm_details (vm_id,name,value,display_detail) values(2,'rootDiskController','ide',1);</pre>
CS-19895	<p>[Realhostip] The certificate for a custom domain can't be reverted after it's uploaded. To use the <code>realhostip</code> domain, upload the <code>realhostip</code> certificate and key again.</p>
CS-19885	<p>[Realhostip] Certificates uploaded in 4.x versions without <code>realhostip</code>-related fixes fail after upgrading to a version with those defect fixes.</p> <p>Workaround: After the upgrade, upload the certificate again in the correct order with supplying all the parameters.</p>
CS-20150	<p>After upgrading from CloudPlatform 4.2.1 to 4.3, the VPN Customer Gateway functionality goes missing. The script to encrypt <code>ipsec_psk</code> during upgrade is missing in version 4.2.1, which causes this issue.</p> <p>Workaround: Run the following to encrypt the values of <code>ipsec_psk</code> in the <code>s2s_customer_gateway</code> table:</p> <pre data-bbox="858 1368 1439 1570"># java -classpath /usr/share/ cloudstack-common/lib/jasypt-1.9.0.jar org.jasypt.intf.cli. JasyptPBEStrEncryptionCLI encrypt.sh input=<clearText> password=<secretKey> verbose=false</pre> <p>Use the secret key for the database.</p>
CS-20181	<p>[Realhostip] The SSL certificates uploaded with invalid sequence numbers are not handled.</p> <p>Workaround: Upload certificates in the correct order. Use <code>id=1</code> for the first root certificate, then for the subsequent intermediate certificates use <code>id=2</code>, <code>id=3</code>, <code>id=4</code>, and so on.</p>
CS-20246	<p>[Realhostip] If invalid certificates are uploaded, the SSL error you get while viewing the console might refer to <code>realhostip.com</code>. Neglect this error for debugging. Uploading a valid certificate would resolve the issue.</p>

Issue ID	Description
CS-20319	<p>[Realhostip] The Console Proxy VMs need to be recreated after changing the communication mode from HTTP to HTTPS vice versa.</p> <p>Workaround: Stop all the running Console Proxy VMs so that they are directed to listen on the right port.</p>
CS-20120	<p>When a volume is being migrated from a storage pool to another, and if the Management Server is stopped; duplicate entries for the volume are left in the database. When the Management Server is started later, the volume is unusable and the database edits have to be made to make it work.</p> <p>Therefore, if a volume is being migrated to another storage, ensure that the Management Server is not be stopped or restarted. This is because volume migration is a long running task and if the Management Server is restarted updating database fails once the operation is completed on the hypervisor.</p>
CS-20632	<p>Problem: The native VPN client on OS X 10.9.3 cannot be connected to a CloudPlatform remote access VPN. However, native client on Windows 7 connects to the same VPN as expected.</p> <p>Root cause: OpenSwan version 1:2.6.37-3+deb7u1 fail to work with OSX/IOS after latest Debian security update. For more information, see Debian security update². Debian has not fixed the issue. For more information, see Debian Bug report logs - #744717³.</p> <p>To workaround, perform the following:</p> <ol style="list-style-type: none"> 1. Go to the VR that hosts the VPN connection and add the following line at the end of the <code>/etc/ipsec.d/12tp.conf</code> file. <pre data-bbox="810 1608 1350 1671">forceencaps=yes</pre> 2. Restart the ipsec service. <pre data-bbox="810 1760 1350 1823"># service ipsec restart</pre> <p>This workaround would fix the issue; however, Windows native VPN client might</p>

² <https://www.debian.org/security/2014/dsa-2893>

³ <https://bugs.debian.org/cgi-bin/bugreport.cgi?bug=744717>

Issue ID	Description
	<p>not be able to connect through the VPN. In order for the Windows VPN client to connect, remove the <code>forceencaps=yes</code> flag and restart the ipsec service.</p>
CS-20763	<p>System VMs cannot be upgraded after upgrade from CloudPlatform 4.2.1 to 4.3.0.1 due to an error with VMware network labels.</p>
CS-25952	<p>After cold migration of a VM, starting VM results in warning messages in CloudPlatform logs. After a cold migration, VMs on VMware fails to start on the first attempt. However, CloudPlatform retries to start the VM if the first attempt fails. When the attempt is made again, VM is successfully brought up.</p> <p>Root cause: Consider the following:</p> <ul style="list-style-type: none"> • After VM migration, VMDK name of the ROOT volume changes. In the volumes table in the database, the path is updated to reflect this change but the chain_info details remain the same. • VMware vCenter is not in sync with CloudPlatform when a VM is cold migrated and therefore is unaware of the change in the location of the VM's volumes. • When you start a VM, CloudPlatform retrieves the current disk information of the VM from vCenter, tear down the disks from the VM, reconfigure the VM with the disks, and power on the VM. During VM reconfiguration with the disks if disk information is returned by vCenter in the very first step, then that information is honored and used to reconfigure the VM. Otherwise CloudPlatform builds the disk information based on the volume's path. <p>Starting a VM after cold migration triggers the following chain of activities : When a VM is cold migrated, vCenter is unaware of the change in the ROOT volumes's location. Therefore, when you attempt to retrieve VM's disk information, vCenter returns the obsolete path, which is the path before the VM is migrated. CloudPlatform then reconfigures the VM with this wrong disk information and attempts to power on the VM. This power-on fails with a warning message. However, CloudPlatform retries when the VM fails to start on the first attempt. And on the second attempt, vCenter returns no disk information because CloudPlatform has already</p>

Issue ID	Description
	<p>torn down the disks during the first attempt to start the VM. Because vCenter returns no disk information, CloudPlatform builds the disk information based on the volume path, which has been rightly updated after the cold migration, reconfigure the VM with the right disk and successfully powers on the VM.</p>
CS-26167	<p>A large number of jobs remaining in the <code>async_jobs</code> table.</p> <p>The global configuration parameters, <code>job.expire.minutes</code> and <code>job.cancel.threshold.minutes</code> are incorrectly multiplied by a factor of 60. Therefore, to set a desired value, divide it by 60 to undo the incorrect multiply effect. For example, if you want to set <code>job.expire.minutes</code> to 120 minutes, set the value to 2 (120/60), instead. This is applicable to <code>job.cancel.threshold.minutes</code> as well.</p>
CS-26091	<p>Problem: Logs are filled up with error messages as given below:</p> <pre>com.mysql.jdbc.exceptions.jdbc4. MySQLIntegrityConstraintViolationException: Cannot delete or update a parent row: a foreign key constraint fails (`cloud`.`async_job_join_map`, CONSTRAINT `fk_async_job_join_map__join_job_id` FOREIGN KEY (`join_job_id`) REFERENCES `async_job` (`id`))"</pre> <p>Root cause: The problem occurs because there are duplicate constraints on the <code>async_job_join_map</code> table, and the second one doesn't have the ON DELETE CASCADE set, meaning that the parent row can't be removed until all dependencies are cleaned up.</p> <pre>CONSTRAINT `fk_async_job_join_map__job_id` FOREIGN KEY (`job_id`) REFERENCES `async_job` (`id`) ON DELETE CASCADE, CONSTRAINT `fk_async_job_join_map__join_job_id` FOREIGN KEY (`join_job_id`) REFERENCES `async_job` (`id`).</pre> <p>Solution: Remove the second constraint as follows as it duplicates the first one. Leave the constraint where ON DELETE CASCADE is specified.</p>

Issue ID	Description
	<pre data-bbox="874 264 1326 342"># ALTER TABLE async_job_join_map DROP FOREIGN KEY fk_async_job_join_map__join_job_id</pre>
CS-26519	<p data-bbox="850 376 1410 443">The VMs with Windows 8.1 guest OS does not work as expected on VMware.</p> <p data-bbox="850 472 1430 882">There is no unique Guest OS available for Microsoft Windows 8.1 (both 32-bit and 64-bit) on VMware vSphere. When you deploy a VM from an ISO that is registered as Windows 8.1, CloudPlatform defaults it to Other (32-bit)/ Other (64-bit) guest OS as principle no mapping for this version is found on vSphere. Therefore, the VM deployed on vSphere will have the Guest OS type as Other (32-bit)/ Other (64-bit). For more information on the issue, refer to Microsoft Windows 8.1 guest operating system option is not available (2067000)⁴.</p> <p data-bbox="850 911 1362 978">Workaround: Based on your environment, perform the following:</p> <ul data-bbox="850 1008 1406 1167" style="list-style-type: none"> <li data-bbox="850 1008 1342 1075">• Register Windows 8.1 template\ISO as Windows 8.0 <li data-bbox="850 1104 1406 1167">• Register Windows 2012 R2 template\ISO as Windows 2012
CS-27136	<p data-bbox="850 1189 1426 1285">You might get the following error message while doing the Live volume migration operation: "Message: No such disk device:"</p> <p data-bbox="850 1314 1370 1382">Workaround: Stop and start the VM before proceeding further.</p>
CS-27344	<p data-bbox="850 1402 1442 1565">[VMware] VM snapshot usage calculation for root disk is not proper when vmware.create.full.clone is set to false. The root disk size of the VM snapshot is way less than the actual template size from which the VM is created.</p>
CS-27642	<p data-bbox="850 1588 1394 1852">Operations involving ROOT volume, such as Volume migration, VM migration with storage, Root Volume Snapshot fails with the "Live VM Migration with volumes" error. This error occurs due to the ROOT volumes meta data inconsistency between CloudPlatform and vCenter during live migration. Suggested workaround resolves the inconsistency.</p>

⁴ http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2067000

Issue ID	Description
	<p>Take a volume snapshot of the ROOT volume once the live VM migration with volumes is completed successfully. Use the following API to snapshot the root volume.</p> <pre data-bbox="767 409 1350 499">command=createSnapshot&volumeid= >uuid of root volume<</pre>
CLOUDSTACK-1717	<p>Local region entry that is added by default should not include "/api" for its end_point. Additionally, the endpoint should have the actual hostname instead of localhost.</p>
CLOUDSTACK-2112	<p>VM will go into stopped state after a failed live migration during a scale up VMs operation.</p> <p>Workaround: Manually restart the VM.</p>
CLOUDSTACK-2293	<p>DeletePhysicalNetworkCmd is not deleting the external devices.</p>
CLOUDSTACK-2646	<p>When firewall and LB service providers are different, CloudPlatform incorrectly allows both the rules on the same public IP. Workaround: Admin should not create network offering with different service providers for firewall and LB, while keeping conserve mode on.</p>
CLOUDSTACK-2910	<p>Ctrl combined with > is not working on SC IME.</p> <p>Workaround: Click the "Chinese/Western Punctuation(Ctrl+.)" in the IME tool bar to switch the punctuation between full-width and half-width.</p>
CLOUDSTACK-3111	<p>Volume listing screen shows Hypervisor column as empty if the volumes are attached to instances running in KVM Hypervisor.</p>
CLOUDSTACK-3212	<p>Default guest network can now have multiple subnets per VLAN, but the IP range list page does not display the netmask and gateway for each subnet.</p> <p>Workaround: Use the API listVlanIPRanges to get the complete details.</p>
CLOUDSTACK-3317	<p>Management and storage network traffic cannot be configured to use VMware Distributed vSwitch (DVS). Continue to use standard vSwitch.</p>
CLOUDSTACK-3895	<p>VM Migration across VMware clusters which are added with different switches (Standard Switch, VMware DVS, Cisco Nexus 1000v) is not supported.</p>
CLOUDSTACK-3680	<p>[KVM on CentOS 5.5, 5.6] While accessing console view of a guest virtual machine, the keystrokes tab, ctrl, \, tilde, single quote,</p>

Issue ID	Description
	double quote, and caret ^ do not work on CentOS 5.5\5.6 running on KVM. This is due to a known bug in CentOS (see http://www.centos.org/modules/newbb/viewtopic.php?topic_id=33233&forum=55 ⁵).
CLOUDSTACK-3968	Distributed port groups on DV Switch are not removed when the associated account from CloudPlatform is removed.
CLOUDSTACK-4016	The listPublicIpAddresses API lists the portable IP that was already transferred to a different Isolated network.
CLOUDSTACK-4139	<p>[VMware] The volumes created from snapshots on VMware deployments cannot be resized when attached to a running VM. The volume is created with IDE disk instead of SCSI disk which cannot be resized.</p> <p>Workaround: Detach the volume created from a snapshot and resize it, and then reattach it to the VM.</p>
CLOUDSTACK-4207	<p>The following exception is observed when the Management Server is started after upgrade from any older versions to CloudPlatform 4.2.</p> <pre> jsonParseException: The JsonSerializer com.cloud.agent.transport. ArrayTypeAdaptor@2426e26f failed to deserialize json object </pre> <p>Ignore this exception, this would stop after you upgrade the System VM. However, if you want to prevent this, stop system VM from the hypervisor before upgrade.</p>
CLOUDSTACK-4364	Restore VM needs to log usage event for volume so that it is correctly charged for usage.
CLOUDSTACK-4475	<p>If cluster-wide and zone-wide primary storage are mixed together, the data disk by default will be created on cluster wide primary storage.</p> <p>Workaround: If admin wants data disk to be created on zone-wide primary storage, then create a disk offering with the tag on zone-wide primary storage.</p>
CLOUDSTACK-4492	Uploaded volume state was not set to "Uploaded" in CloudPlatform 3.0.6. After upgrade to 4.x, volume attach fails because of volume being in incorrect state.

⁵ http://www.centos.org/modules/newbb/viewtopic.php?topic_id=33233&forum=55

Issue ID	Description
	<p>Workaround: Upload and attach volume after the upgrade.</p>
CLOUDSTACK-4517	<p>Deployment of VM using CentOS 6.2 template registered before upgrade is failing.</p>
CLOUDSTACK-4578	<p>[VMware] If the host where the SSVM is running goes down, the SSVM is not being recreated on another host in the cluster.</p> <p>Workaround: Forcefully stop the SSVM through the CloudPlatform API call stopSystemVm. Then the new SSVM will be created on a second host.</p>
CLOUDSTACK-4593	<p>Live Storage Migration and VM Snapshot features are not fully functional after upgrade.</p> <p>Workaround: Stop and then start the VM post upgrade.</p>
CLOUDSTACK-4622	<p>If a VM from a guest network is added to a network tier of a VPC, then IP reservation allows the CIDR to be the superset of Network CIDR for that VPC tier.</p>
CLOUDSTACK-5452	<p>[KVM] Agent is not able to connect back if the Management Server was restarted when pending tasks to the hosts are remaining.</p> <p>Workaround: Restart the agent.</p>
CLOUDSTACK-5463	<p>[Hyper-V] Stopped VMs are not reported, because out of band state changes occurred on VMs or hosts are not reconciled by CloudPlatform.</p>
CLOUDSTACK-5485	<p>[VMware] When 10 hourly snapshots are scheduled in parallel, only 5 of them are being simultaneously processed actively.</p> <p>To increase the number of simultaneous commands processed in SSVM (increase the count of worker threads), modify the agent properties file in SSVM to specify the number of workers.</p> <ul style="list-style-type: none"> • Stop the cloud service: <pre data-bbox="791 1697 1350 1765">service cloud stop</pre> • In SSVM, update the following file to add the number of line workers: <pre data-bbox="791 1888 1350 1977">/usr/local/cloud/systemvm/conf/agent.properties</pre> • Run the cloud service:

Issue ID	Description
	<pre>service cloud start</pre>
CLOUDSTACK-5501	Creating more than one VPN connection per customer gateway is not supported.
CLOUDSTACK-5753	[Hyper-V] ConsoleProxyLoadReportCommand does not honor the default value of consoleproxy.loadscan.interval, which is 10 second.
CLOUDSTACK-5815	[Hyper-V] Two SNAT rules for one isolated network is created if the acquired IP is from a different VLAN.

Upgrade Instructions

4.1. Upgrade from 4.3.x.x to 4.3.0.2

Perform the following to upgrade from version 4.3.x.x to version 4.3.0.2.

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*¹.

Hypervisor	Description
XenServer	<p>Name: systemvm-xenserver-4.3</p> <p>Description: systemvm-xenserver-4.3</p> <p>URL (64-bit system VM template): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-xen.vhd.bz2</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

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

Hypervisor	Description
	<p>Description: systemvm-kvm-4.3</p> <p>URL (64-bit system VM template): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-kvm.qcow2.bz2</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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-vmware.ova</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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-hyperv.vhd.bz2²</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 4.3.0.2 RHEL 6.3 binaries from <https://www.citrix.com/downloads/cloudplatform.html>.

² <http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-hyperv.vhd.bz2>

- b. Extract the binaries:

```
# cd /root
# tar xvf CloudPlatform-4.3.0.2-1-rhel6.3.tar.gz
```

- c. Create a CloudPlatform 4.3 qemu repo:

```
# cd CloudPlatform-4.3.0.2-1-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-4.3.0.2-1-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 4.3.0.2 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](#)³.

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
```

³ <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 4.3.0.2.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
```

Substitute your own IP address of Secondary Storage VMs and Console Proxy VMs.

- 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 4.3.0.2 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 4.3.0.2 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 4.3.0.2

Perform the following to upgrade from version 4.2.x to version 4.3.0.2.

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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-xen.vhd.bz2</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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-kvm.qcow2.bz2</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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-vmware.ova</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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-hyperv.vhd.bz2⁴</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.

⁴ <http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-hyperv.vhd.bz2>

Chapter 4. Upgrade Instructions

- a. Download the CloudPlatform 4.3.0.2 RHEL 6.3 binaries from <https://www.citrix.com/downloads/cloudplatform.html>.
- b. Extract the binaries:

```
# cd /root
# tar xvf CloudPlatform-4.3.0.2-1-rhel6.3.tar.gz
```

- c. Create a CloudPlatform 4.3 qemu repo:

```
# cd CloudPlatform-4.3.0.2-1-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-4.3.0.2-1-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 4.3.0.2 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](#)⁵.

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:

⁵ <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.

- Copy the CloudPlatform 4.3.0.2.tgz download to the host, untar it, and change to the resulting directory.
- 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
```

Substitute your own IP address of Secondary Storage VMs and Console Proxy VMs.

- 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 4.3.0.2 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 4.3.0.2 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
<code>ldap.user.object</code>	<code>inetOrgPerson</code>	<code>user</code>

LDAP Attribute	OpenLDAP	Active Directory
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 4.3.0.2

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 4.3.0.2.

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	<p>Name: systemvm-xenserver-4.3</p> <p>Description: systemvm-xenserver-4.3</p> <p>URL (64-bit system VM template): http://download.cloud.com/templates/4.3/</p>

Hypervisor	Description
	<p>systemvm64template-2014-09-30-4.3-xen.vhd.bz2</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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-kvm.qcow2.bz2</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>

Hypervisor	Description
	<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): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-vmware.ova</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>

Hypervisor	Description
	<p>URL (64-bit system VM template): http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-hyperv.vhd.bz2⁶</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 4.3.0.2 RHEL 6.3 binaries from <https://www.citrix.com/downloads/cloudplatform.html>.
- b. Extract the binaries:

```

# cd /root
# tar xvf CloudPlatform-4.3.0.2-1-rhel6.3.tar.gz
```

⁶ <http://download.cloud.com/templates/4.3/systemvm64template-2014-09-30-4.3-hyperv.vhd.bz2>

- c. Create a CloudPlatform 4.3.0.2 qemu repo:

```
# cd CloudPlatform-4.3.0.2-1-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-4.3.0.2-1-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 4.3.0.2 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](#)⁷.

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 4.3.0.2



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:

⁷ <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 4.3.0.2, 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 4.3.0.2.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
```

Substitute your own IP address of Secondary Storage VMs and Console Proxy VMs.

- 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 4.3.0.2 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 4.3.0.2 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:

- 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 4.3.0.2 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

4.4. 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 4.3.0.2 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](#)⁸.

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
```

⁸ <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 4.3.0.2:

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.

