

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Release Notes

Revised July 13, 2015 04:00 PM IST



Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5 Release Notes

Revised July 13, 2015 04:00 PM IST

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

Experimental Features

CloudPlatform 4.5.1 includes experimental features for customers to test and experiment with in non-production environments, and share any feedback with Citrix. For any issues with these experimental features, customers can open a support ticket but Citrix cannot commit to debugging or providing fixes for them.

The following experimental features are included in this release:

- Linux Containers
- Supported Management Server OS and Supported Hypervisors: RHEL 7/CentOS 7 is only applicable for experimental use with Linux Containers.

Release notes for Citrix CloudPlatform version 4.5

1. What's New in 4.5	1
1.1. New Features	1
1.1.1. vGPU Support	1
1.1.2. KVM Qemu Enhancements	2
1.1.3. Rolling Upgrades Enhancements	3
1.1.4. (Experimental Feature) Advanced Networking in Baremetal	3
1.1.5. (Experimental Feature) Linux Containers	3
1.2. API Changes	4
1.2.1. (Experimental Feature) Advanced Networking in Baremetal	4
1.2.2. vGPU Support	4
1.2.3. (Experimental Feature) LXC Support	5
1.2.4. KVM/Qemu Support	5
1.3. System VM Templates	5
1.4. Fixed Issues	6
2. Support Matrix	11
2.1. Supported OS Versions for Management Server	11
2.2. Supported Hypervisor Versions	11
2.3. Supported External Devices	12
2.4. Supported Browsers	12
2.5. Software Requirements	12
3. Known Issues	13
4. Feature Parity Between CloudPlatform and Apache CloudStack	25
5. Submitting Feedback and Getting Help	27

What's New in 4.5

- [Section 1.1, “New Features”](#)
- [Section 1.2, “API Changes”](#)
- [Section 1.3, “System VM Templates”](#)
- [Section 1.4, “Fixed Issues”](#)

1.1. New Features

- [Section 1.1.1, “vGPU Support”](#)
- [Section 1.1.2, “KVM Qemu Enhancements”](#)
- [Section 1.1.3, “Rolling Upgrades Enhancements”](#)
- [Section 1.1.4, “\(Experimental Feature\) Advanced Networking in Baremetal ”](#)
- [Section 1.1.5, “\(Experimental Feature\) Linux Containers ”](#)

1.1.1. vGPU Support

CloudPlatform supports deploying guest VMs with Graphics Processing Unit (GPU) or Virtual Graphics Processing Unit (vGPU) capabilities on XenServer hosts. At the time of VM deployment or at a later stage, you can assign a physical GPU, known as GPU-passthrough, or a portion of a physical GPU card (vGPU) to a guest VM by changing the Service Offering. With this capability, the VMs running on CloudPlatform meet the intensive graphical processing requirement by means of the high computation power of GPU/vGPU, and CloudPlatform users can run multimedia rich applications, such as Auto-CAD, that they otherwise enjoy at their desk on a virtualized environment.

CloudPlatform leverages the XenServer support for NVIDIA GRID Kepler 1 and 2 series to run GPU/vGPU enabled VMs. NVIDIA GRID cards allows sharing a single GPU cards among multiple VMs by creating vGPUs for each VM. With vGPU technology, the graphics commands from each VM are passed directly to the underlying dedicated GPU, without the intervention of the hypervisor. This allows the GPU hardware to be time-sliced and shared across multiple VMs. XenServer hosts use the GPU cards in following ways: GPU passthrough and GRID vGPU.

CloudPlatform provides you with the following capabilities:

- Adding XenServer hosts with GPU/vGPU capability provisioned by the administrator.
- Creating a Compute Offering with GPU/vGPU capability.
- Deploying a VM with GPU/vGPU capability.
- Destroying a VM with GPU/vGPU capability.
- Allowing an user to add GPU/vGPU support to a VM without GPU/vGPU support by changing the Service Offering and vice-versa.
- Migrating VMs (cold migration) with GPU/vGPU capability.
- Managing GPU cards capacity.
- Querying hosts to obtain information about the GPU cards, supported vGPU types in case of GRID cards, and capacity of the cards.

- Dynamic Scaling of CPU and memory is supported with vGPU-enabled VMs on XenServer 6.5.

Consider the following non-functionalities:

- Deploying VMs with GPU/vGPU capability is not supported if hosts are not available with enough GPU capacity.
- Dynamic Scaling of CPU and memory is not supported with vGPU-enabled VMs on XenServer 6.2 Service Pack 1.
- A Service Offering cannot be created with the GPU values that are not supported by the CloudPlatform UI. However, you can make an API call to achieve this.
- Dynamic scaling is not supported. However, you can choose to deploy a VM without GPU support, and at a later point, you can change the System Offering to upgrade to the one with vGPU. You can achieve this by offline upgrade: stop the VM, upgrade the Service Offering to the one with GPU/vGPU, then start the VM.
- Live migration of GPU/vGPU enabled VM is not supported.
- Limiting GPU/vGPU resources per Account/Domain is not supported.
- Disabling GPU/vGPU at Cluster level is not supported.
- Notification thresholds for GPU/vGPU resource is not supported.
- Selecting hosts with no GPU cards over the hosts configured with GPU cards (GPU-enabled) to deploy VMs without GPU capabilities is not supported.

You can deploy a VM with Service Offering not enabled for GPU on GPU-enabled hosts. The VMs are randomly deployed on hosts irrespective of the availability of the GPU cards. To deploy the VMs without GPU functionality, currently, CloudPlatform does not consider the hosts with no GPU cards over GPU-enabled hosts. No mechanism available now to intelligently deploy these VMs on hosts without GPU cards, and consider the GPU-enabled hosts only when the hosts without GPU cards do not have enough capacity.

1.1.2. KVM Qemu Enhancements

In CloudPlatform 4.5, now packages a new set of patched binaries, `ccp-qemu-img`, to back up Snapshot from Primary to Secondary Storage, which does not overwrite the existing KVM binaries. In earlier versions of CloudPlatform, KVM qemu files have been overwritten by the binaries provided by CloudPlatform to achieve faster snapshotting, which broke the supportability from RedHat. In this release, snapshotting behaviour has been defaulted to that of the hypervisor, thereby enabling you to avail yourself of RedHat support. Depending on the size of the VM memory, you encounter a pause during backing up the VM snapshots. It implies that higher the memory size, longer the pause time.

Consider the following changes for KVM:

- The supported host OS are RHEL 6.3, 6.4, and 6.5.

RHEL 6.2 and below versions are no longer supported due to Java 7 requirements in CloudPlatform.

- For RHEL 6.3 and 6.4 versions, no patched KVM binaries are provided because the existing binaries can perform the backup operation.
- For RHEL 6.5, CloudPlatform provides a patched `ccp-qemu-img` binaries to back up Snapshot from Primary to Secondary Storage, which will not overwrite existing binaries. The binaries are installed

under `/usr/bin/` directory, which are used by the CloudPlatform KVM Agent during the snapshot backup.

CloudPlatform installer detects the version of KVM host, and if it's RHEL 6.3 or 6.4 version, the `ccp-gemu-img` binaries are not installed, otherwise, it is installed automatically.

1.1.3. Rolling Upgrades Enhancements

VR scalability has been improved as compared to previous versions. Earlier versions of CloudPlatform executed VR operations, such as creating or restarting VR, as individual commands, instead of performing a batch operation. This led to hours of pricey delay in larger deployments. With CloudPlatform 4.5, a new approach of VR reprogramming has been introduced to eliminate time consumed activities for VR operations. The new mechanism is to execute the commands in one batch inside VR, and restart the related services, if necessary, after the whole batch is completed.

1.1.4. (Experimental Feature) Advanced Networking in Baremetal

CloudPlatform adds advanced network capabilities for Baremetal. A new plug-in has been introduced in CloudPlatform which enables automatic VLAN programming on a physical switch to which baremetal instances are connected. In an Advanced zone, Baremetal instances gain VLAN isolation provided by CloudPlatform which is particularly useful if you want to provide Baremetal As a Service for public clouds. Baremetal As a Service cannot function standalone; it works in conjunction with a physical switch either from vendor's SDK or from an in-switch agent for white box switch.

For detailed instructions, see CloudPlatform 4.5 Hypervisor Configuration Guide.

1.1.5. (Experimental Feature) Linux Containers

CloudPlatform 4.5 rolls out experimental support for LXC hosts on RHEL 7. In addition to NFS and local storage, RBD (Ceph) is supported for Primary storage. If you want to run guest VMs (containers) on LXC hosts, install CloudPlatform Agents on each LXC hosts. For detailed instruction, see the Installing LXC for CloudPlatform chapter in the CloudPlatform 4.5 Installation Guide.

Before you use LXC, review the following list of unsupported features.

- Console access
- Live migration
- Migration across clusters
- Storage migration
- Snapshots
- Uploading and downloading volumes
- Template creation from ROOT volume
- ISO support for creating VM

1.2. API Changes

1.2.1. (Experimental Feature) Advanced Networking in Baremetal

API	Description
AddBaremetalRCT	A new field name, rctUrl, has been added. rctUrl specifies an HTTP link pointing to RCT on an accessible HTTP server.

1.2.2. vGPU Support

API	Description
createServiceOffering	<p>The following keys have been added for the serviceofferingdetails parameter. Specify them in key-value pairs.</p> <ul style="list-style-type: none"> pciDevice: Corresponding values are: <ul style="list-style-type: none"> Group of NVIDIA Corporation GK107GL [GRID K1] GPUs Group of NVIDIA Corporation GK104GL [GRID K2] GPUs Any other GPU group name vgpuType: Corresponding values are: <ul style="list-style-type: none"> GRID K100 GRID K120Q GRID K140Q GRID K200 GRID K220Q GRID K240Q GRID K260Q
listServiceOffering	<p>If service offering is GPU/vGPU enabled, the GPU card and type of vGPU are displayed in key-value pairs.</p> <ul style="list-style-type: none"> GPU vGPUType
listHosts	<p>If the host is GPU enabled, the following GPU details are listed:</p> <ul style="list-style-type: none"> The GPU cards (GPU groups) managed by the host.

API	Description
	<ul style="list-style-type: none"> Enabled vGPU Types: <ul style="list-style-type: none"> Video RAM Maximum Heads Maximum X resolution Maximum Y resolution Maximum vGPU resources per GPU card (pgpu) Remaining Capacity <p>The remaining number of vGPU that can be created with this type of vGPU in the given GPU group.</p> Maximum Capacity
listVirtualMachines	If the VM has GPU support, the following VM statistics is listed: vGPU type

1.2.3. (Experimental Feature) LXC Support

No new APIs.

1.2.4. KVM/Qemu Support

No new APIs.

1.3. System VM Templates

CloudPlatform 4.5 supports 64-bit System VM templates. This release does not provide 32-bit support for System VM templates.

Hypervisor	Description
XenServer	http://download.cloud.com/templates/4.5/systemvm64template-2014-12-18-4.5.0.0-xen.vhd.bz2
Hyper-V	http://download.cloud.com/templates/4.5/systemvm64template-2014-12-18-4.5.0.0-hyperv.vhd.bz2
KVM	http://download.cloud.com/templates/4.5/systemvm64template-2014-12-18-4.5.0.0-kvm.qcow2.bz2
VMware	http://download.cloud.com/templates/4.5/systemvm64template-2014-12-18-4.5.0.0-vmware.ova

For more information, see the Prepare the System VM Template section in the CloudPlatform 4.5 Installation Guide.

1.4. Fixed Issues

Issue ID	Description
CS-31703	ExternalNetworkResourceUsageCommand no longer throws SAXException with F5 as Load balancer.
CS-29526	Invalid zone ID error no longer generated while listing VMware zones error.
CS-28524	No inconsistency in usage dates.
CS-26148	The NET.IPRELEASE events are successfully added to usage_event on IP range deletion from physical networks.
CS-25965	Cluster.cpu.allocated.capacity.disablethreshold is no longer ignored when host.capacityType.to.order.clusters is set to RAM.
CS-25961	Disabled host no longer remains in up status after unmanaging a cluster.
CS-25884/CS-25413	VM state is now in sync between cloud database and KVM host.
CS-25815	Resizing volumes that are attached to Windows 2008 and beyond works as expected.
CS-25795	The listSnapshots API command now returns zone ID.
CS-25688	vm_template.source_template_id of template created from snapshot is null.
CS-25155	VMs with more than 7 NICs on VMware successfully restarts.
CS-25148	If VM deployment fails, the Event tab no longer shows the "Successfully completed, Starting VM" status.
CS-24972	Snapshot can now be created from a volume when the task is performed repeatedly in zone-wide primary storage.
CS-24484	Sync job failures are no longer reported as success on Event Bus.
CS-24474	Chain FW_EGRESS_RULES and Chain FW_OUTBOUND no longer disappear when VR is rebooted by using CloudPlatform UI.
CS-23830	Port forwarding rules are successfully configured on the VR after VMware is crashed.
CS-22952	Multiple threads are no longer used to collect the stats from the same VR.
CS-22858	Attaching data disk no longer fails for clusters with only zone-wide primary storage.
CS-22857	Correct version of StopCommand is displayed with or without cleanup flag.
CS-22853	VM Sync operation no longer reports running VMs as stopped on VMware hosts.
CS-22842	Attaching a volume works as expected.
CS-22761	VM is successfully expunged in 4.3.0.1.
CS-22665	The new VM state machine changes now capture the usage events.
CS-22661	FTP modules are loaded in VR.
CS-22637	Performance of cleanString method is improved in StringUtils.
CS-22606	CloudPlatform no longer collects usage for VMs which do not exist anymore.
CS-22596	After a volume is migrated, the usage table no longer shows the old volume ID.
CS-22562	Taking snapshot of a volume works as expected.
CS-22377	A NIC can now be removed from a VM without any error.

Issue ID	Description
CS-22321	VMs are no longer deleted from hypervisor after deleting from UI when using zone-wide primary storage.
CS-21374	Templates with the same name no longer have the same unique_name in the database.
CS-21348	CloudPlatform Management Server connects to hosts soon after rebooting.
CS-21332	No BSOD occurs for Windows 2008 R2 and 7 or earlier VMs on KVM.
CS-21319	Network offering with VPC now allows selecting system offering for VR.
CS-21289	Events can now be archived or deleted for a deleted account.
CS-21108	VMware host can now be added to an existing cluster.
CS-21083	HA no longer reboots several VMs while they were still running on a disconnected host.
CS-21078	Host no longer remains in Alert state after vCenter restart.
CS-21074	CloudPlatform no longer attempts to enter host name three times in addHost string.
CS-20938	Connection time out in vSphere server no longer causes hosts to be disconnected in CloudPlatform.
CS-20894	Applying load balancer rules no longer breaks existing connections and causes short outage.
CS-20893	SSVM no longer gets disconnected due to the channel is closed error.
CS-20891	The listnetworks API no longer fails with the exception, "Incorrect number".
CS-20853	Running_VM and ALLOCATED_VM usage records no longer contains old service offering ID after service offering has been changed for the VM by using the API changeServiceForVirtualMachine.
CS-20845	Source NAT, Static NAT and Port forwarding now works as expected with VMware DVS.
CS-20789	Live migrating VMs with local storage no longer fails from CloudPlatform and continues to work as expected by using Storage XenMotion.
CS-20787	System VMs can be started after upgrading 4.2.1 to 4.3.0.1. No errors are observed with VMware network labels.
CS-20785	passwd_server attempts no longer terminates with the exit code 137 while starting.
CS-20754	Attaching datadisk for VMs that have VM snapshot no longer results in screenful JAVA errors.
CS-20752	Latest OS X VPN client now works as expected.
CS-20748	VOLUME.DELETE usage event now available for VMs in ERROR state.
CS-20699	VM password length and complexity has been increased.
CS-20686	Template deletion is not prevented when template is in use.
CS-20628	The queryAsyncJobResult API now returns jobinstanceid.
CS-20611	VMs can now be deleted successfully.
CS-20609	irqbalance can now be run on multi-core CPU of VRs which is created from the new 64-bit system template.

Issue ID	Description
CS-20608	Default value of XenServer, max guest limit, is now honoured by CloudPlatform.
CS-20601	CloudPlatform no longer records ID in events tables instead of UUID.
CS-20533	VMware worker VMs are now cleaned up as expected.
CS-20451	Upgrading to 4.3.0.1 no longer fails when a VMware setup has multiple zones managing the same VMware datacenter.
CS-20258	HA works as expected when the uploaded volume is attached to the VM.
CS-20131	Volume migration fails when using local disk and XenServer 6.2.
CS-20124	Connection to host no longer fails when the VM Snapshot is stuck in Creating state.
CS-20102	The listUsageRecords API no longer returns usage of volumes with Expunged status.
CS-20098	Attaching multiple data disks to a new VM works as expected.
CS-20086	CloudPlatform management catalina.out log now rotates as expected.
CS-20048	CloudPlatform no longer allows creating VMs with the same display name when vm.instanceName.flag is set to true.
CS-19929/CS-19927	IP address of the VR does not change if its destroyed and re-created in Basic zone.
CS-19916	Storage migration between cluster-wide and zone-wide storage works as expected.
CS-19908	Storage can be chosen when attaching a uploaded data volume to VM.
CS-19836	Volume migration between pools works as expected in both CloudPlatform and XenServer.
CS-19765	VR can be immediately destroyed with RestartNetworkCmd and create new VR instead of releasing the PF/FW rules on VR.
CS-19530	Template ordering in UI works as expected.
CS-19303	The updateNetwork API passes empty string when network domain is null.
CS-19135	Downloading template works as expected when having three SSVMs in a single zone.
CS-18752	For Basic zone only setup clicking the Add guest network option no longer shows API error.
CS-25952	After cold migration of a VM, starting VM no longer results in warning messages in CloudPlatform logs. After a cold migration, VMs on VMware start on the first attempt itself.
CS-19675	[VMware] In clusters with multiple primary storages configured VMs no longer fail to restart when either Reset VM operation is performed or the compute offering has the Volatile option enabled.
CLOUDSTACK-5815	[Hyper-V] Two SNAT rules for one isolated network is no longer created if the acquired IP is from a different VLAN.
CLOUDSTACK-5501	Creating more than one VPN connection per customer gateway is now supported.
CLOUDSTACK-5485	[VMware] When 10 hourly snapshots are scheduled in parallel,all of them are being simultaneously processed.

Issue ID	Description
CLOUDSTACK-5463	[Hyper-V] Stopped VMs are now reported when out of band state changes occurred on VMs or hosts are not being reconciled by CloudPlatform.
CLOUDSTACK-4475	If cluster-wide and zone-wide primary storage are mixed together, the data disk by default will no longer be created on cluster-wide primary storage.
CLOUDSTACK-4364	Restore VM needs to log usage event for volume so that it is correctly charged for usage.
CLOUDSTACK-3212	Default guest network can have multiple subnets per VLAN, and the IP range list page displays the netmask and gateway for each subnet.

Support Matrix

This section describes the operating systems, browsers, and hypervisors that have been newly tested and certified compatible with CloudPlatform 4.5. Most earlier OS and hypervisor versions are also still supported for use with 4.5. For a complete list, see the System Requirements section of the CloudPlatform 4.5 Installation Guide.

2.1. Supported OS Versions for Management Server

- RHEL versions 6.3 and 6.5

RHEL 6.4 is not supported.

- CentOS versions 6.3, 6.4, and 6.5
- RHEL/CentOS 7

RHEL 7 and CentOS 7 is supported only for experimental use with Linux Containers.

RHEL 7 is supported only for the fresh installation of CloudPlatform 4.5.

2.2. Supported Hypervisor Versions

Support for the following hypervisors has been added:

- LXC Host Containers on RHEL 7

Only for experimental use.

- KVM 6.5

The libvirt versions supported for KVM on RHEL 6.5 are libvirt-0.10.2-41 version and above.

- VMware vCenter versions 5.0 Update 3a
- VMware vCenter versions 5.1 Update 2a
- VMware vCenter versions 5.5 Update up to 2

Other supported hypervisors for CloudPlatform:

- Windows Server 2012 R2 (with Hyper-V Role enabled)
- Hyper-V Server 2012 R2
- VMware vCenter versions 5.0 up to Update 3a
- VMware vCenter versions 5.1 unto Update 2a
- VMware vCenter versions 5.5 unto Update 1b
- KVM 6.3

The libvirt versions supported for KVM on RHEL 6.x are libvirt-0.10.2-41 and beyond.

- XenServer version 6.5
- XenServer version 6.5 SP1

- XenServer version 6.2 SPI with Hotfix XS62ESP1004 and beyond
- 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, 10.1.e, and 10.5
 - Netscaler VPX versions 9.3, 10.1.e, and 10.5
- Supported only on XenServer, KVM, and VMware

- NetScaler SDX versions 9.3, 10.1.e, and 10.5
- SRX (Model srx100b) versions 10.3 to 10.4 R7.5
- F5 11.X
- Force 10 Switch version S4810 for Baremetal Advanced Networks

2.4. Supported Browsers

- Mozilla Firefox versions 26 to 33
- Google Chrome 38.x
- Apple Safari 7.1
- Microsoft Internet Explorer versions 9, 10 and 11

2.5. Software Requirements

- Java 1.7
- MySQL 5.6 (RHEL 7)
- MySQL 5.1 (RHEL 6.x)

Known Issues

Issue ID	Description
CS-32825	<p>Problem: If a guest network is deleted and expunged, the VLAN tag is not deleted from the hypervisors guest interface.</p> <p>Root cause: CloudPlatform doesn't clean the networks on XenServer when the guest networks are deleted.</p> <p>Solution: Manually delete directly from the hypervisor host.</p>
CS-32756	<p>Problem: SystemVM agents will not start if the <i>instance.name</i> parameter value starts with "_" .</p> <p>Root cause: The character "-" is not an allowed character in the name of a host with *NIX operating systems. Hostnames, including the domain components, are only allowed to contain alphabetic characters, numeric characters, hyphen and period. This was originally specified in RFC 952 and updated in RFC 1123 though Windows allow the use of hostnames which violate this syntax. However, it is left to administrator to decide to use such syntax based on existence of non-Microsoft hosts/DNS servers.</p> <p>Solution: Do not use the special character, "_", in the value you specify for the <i>instance.name</i> parameter.</p>
CS-32750	<p>Problem: Dynamic scaling is not restricted when destination offering has changes in the vGPU type.</p> <p>Root cause: During scaling up of a running VM, CloudPlatform does not check if this VM is vGPU-enabled or not. If the new service offering has different vGPU type as compared to the current service offering, this problem occurs.</p>
CS-32584	<p>Problem: Deploying new VMs fails on KVM zone with traffic labels.</p> <p>Root cause: This is a pre-upgrade issue on versions 4.3 and below.</p> <p>Solution: Update the networks with the untagged VLAN:</p> <pre data-bbox="544 1503 1453 1585"># mysql> update networks set broadcast_uri = 'vlan://untagged' where id = <network_id> ;</pre>
CS-32411	<p>Problem: After VM cold migration, VM snapshot creation will fail.</p> <p>Root cause: After a VM or its storage is cold migrated, VM will not be in the right state because vCenter is not aware of the new location of the disks. Because of this performing VM operations, for example VM snapshot creation, after a cold migration and before starting the VM results in failure.</p> <p>Solution: Upon cold migration of a VM or VM's disk start the VM before performing any further operations.</p>
CS-32089	<p>Problem: Admin is prompted for password change even after password has been changed on session expiry on deployments where no zones are configured. This behaviour is observed in the following scenarios:</p>

Issue ID	Description
	<ul style="list-style-type: none"> • Clear recent history, and attempt to log in to the Management Server as admin. • Log in to the Management Server as admin by using another browser. • Log in to the Management Server as admin by using another machine. • Log in to the Management Server after a day or whenever the UI session expires. <p>In all these cases, user will be allowed to change the password to the existing password.</p>
CS-32111	<p>Problem: Storage on separate NIC is not recognized on KVM and SSVM fails to launch.</p> <p>Solution: To use multiple NICs on KVM host, manually create multiple bridges on the host, then add that in CloudPlatform.</p>
CS-32042	<p>Problem: Database HA fails when cloudstack-setup-database is executed without cloud user password.</p> <p>Solution: This issue is caused when EMPTY password is provided to the database user while creating a database user. You should not allow empty passwords while creating users.</p>
CS-32079	<p>Problem: Migrating a downloaded volume does not function as expected.</p> <p>Root cause: When you download a volume, an entry is created in volume_store_ref table. The volume is marked in ready state. When you migrate that volume, another entry is created with same volume id. The state is marked as allocated. Later you try to list only one dataobject in datastore for state transition during volume migration. If the listed volume's state is allocated, it passes otherwise it fails.</p> <p>Solution: Download URL expires after 4 hours. After 4 hours migration should pass as there will be unique dataobject in datastore for the volume.</p>
CS-31878	<p>Problem: Attach volume operation fails if Management Server is restarted while the operation is in progress.</p> <p>Root cause: After restarting Management Server, the volume is attached to a VM in XenServer side, but CloudPlatform assumes that the volume is not attached to the VM; therefore, when CloudPlatform attempts to attach another volume to the VM, it complains that the device id has been used by previous volume.</p> <p>Solution: Perform the following:</p> <ol style="list-style-type: none"> 1. Destroy the unattached VBD in volume attach. 2. Use autodetect as device ID in XenServer; XenServer picks up unused device ID for this VBD. 3. If the previous volume, which is unattached in CloudPlatform and attached in XenServer, can't be attached to any VM, detach the volume in XenServer to make XenServer and CloudPlatform in sync.

Issue ID	Description
CS-30980	<p>Problem: If a guest network is deleted and expunged in CloudPlatform, guest VLAN tags are not deleted on XenServer.</p> <p>Root cause: If a guest network is deleted in CloudPlatform, the network is not deleted on the hypervisor.</p> <p>Solution: To remove the network, the administrator can manually delete them on XenServer using the following commands.</p> <ol style="list-style-type: none"> 1. Run the following to list all the networks on a host and the associated UUID: <pre data-bbox="592 633 1458 696"># xe network-list</pre> <ol style="list-style-type: none"> 2. Select the UUID of the network you want to delete. 3. Run the following to destroy the network on the host. <pre data-bbox="592 846 1458 909"># xe network-destroy uuid=<network-uuid></pre>
CS-30610	<p>Problem: When a VM is in running state, custom service offering is not displayed available for the upgrade.</p> <p>Solution: Stop the VM before upgrading the custom service offering.</p>
CS-30608/CS-30609	<p>Problem:(VMware) After upgrading from 3.0.x to 4.5 and higher versions, restoring/ cold migrating the existing VM which has an additional disk fails to boot.</p> <p>Root cause: Due to a defect 3.0.x, even though global configuration setting <i>vmware.root.disk.controller</i> is set to IDE, root disk/volume is created by using SCSI controller itself. However, after upgrading to 4.5, during restore/ cold migration operation root volume is created by using IDE. This implies that the boot volume has moved off from original device node and results in a boot failure, because data disk is searched for OS instead of root disk.</p> <p>Solution: Perform either of the following on user VMs:</p> <ul style="list-style-type: none"> • This issue is specific to user VMs created in version 3.0.7 before upgrade. If a restore/cold migration operation need to be performed over such VMs, detach data volumes before the operation. Once restore/cold migration is completed attach the data volumes in the same order. • Update database: <pre data-bbox="571 1753 1458 1843"># insert into user_vm_details(vm_id,name,value,display) values(100,'rootDiskController','scsi',1);</pre> <p>Where 100 is the ID of the user instance.</p>
CS-30331	<p>Problem: CloudPlatform allows creation of VMs with the same display name when <i>vm.instance.name.flag</i> is set to true.</p>

Issue ID	Description
	<p>Root cause: This happens because CloudPlatform restricts two VMs from having the same name only if they are in the same network. But in case <i>vm.instanceName.flag</i> is set to true, VM name in vCenter will be its CloudPlatform display name. And vCenter doesn't allow two VMs to have the same name under the same datacenter. Which results in the failure of VM start in CloudPlatform and is not handled correctly.</p> <p>Solution: During VM creation, if <i>vm.instanceName.flag</i> is set to true and hypervisor type is VMware, CloudPlatform checks if VM with the same host name already exists in the zone. The verification is being made under the zone because:</p> <ul style="list-style-type: none"> • In a non-legacy zone, since a zone maps to a VMware datacenter, by checking under the CloudPlatform zone, it's ensured that no two VMs with same name are created under the same datacenter. • In a legacy zone, no 1-1 mapping is available between CloudPlatform zones and VMware datacenters. However, entire zone must be checked to avoid any potential failures for other CloudPlatform operations. For example, in case of cold migration, CloudPlatform allows a volume to be moved across between VMware datacenters if they belong to the same CloudPlatform zone. Now if a VM with an existing VM name is allowed to be created under a different datacenter, but same CloudPlatform zone, the after a cold migration, when a VM is attempted to be started it fails with conflicting name issue.
CS-30165	<p>Problem: When all the hosts in a cluster go down, no HA available on user and system VMs.</p> <p>Root cause: Currently, CloudPlatform cannot reliably determine if a host is up or down if all the hosts in a cluster go down or have connectivity issues. In such cases, host is marked in "Alert" and there is no HA triggered for user VMs / System VMs. These will continue to be marked as being in "Running" state. The agent state of SSVM and CPVM will be shown as being "Alert". However, no HA is triggered for both user and system VMs in this case.</p>
CS-30101	<p>Problem: Isolated guest network state remains in implemented state when the last VM created in this network is stopped outside CloudPlatform.</p> <p>Root cause: When a single VM in a network is stopped outside of CloudPlatform, the VMsync operation is run to sync the state to stopped. If last VM on this network is stopped outside of CloudPlatform, the network remains in implemented state. In this scenario, CloudPlatform is not designed to clean the network.</p>
CS-30089	<p>Problem: Live storage migration is not support on VMFS based storage pools.</p>
CS-29850	<p>Problem: Data disk is not detected with Windows 2012 R2 VM ISO on VMware deployments when root controller global configuration is set to SCSI. You cannot install a VM.</p> <p>Root cause: Latest Windows no longer support LSI Logic. Windows default to SAS controller.</p>

Issue ID	Description
	<p>Solution: Configure the <code>vmware.root.disk.controller</code> global configuration to <code>osdefault</code>.</p>
CS-29660	<p>Problem: Success notification is generated for a failed Delete Snapshot operation.</p> <p>Root cause: This issue is caused when a async job is timed out while the snapshot backup is in progress, which giving user an illusion that create snapshot is failed. However, the Management Server is still backing up snapshot. It's a known issue, when async job timeout is less than the timeout set on each actual command.</p> <p>Solution: Set async job timeout to be the maximum of timeout value of each command. In this case, the value of async job, at least, to be ≥ 21600, which is the timeout value of <code>copy.snapshot</code>.</p>
CS-29411	<p>Problem: Live migrating a VM along with storage from a tagged host to an untagged host is allowed when using <code>migrateVirtualMachineWithVolume</code>.</p> <p>Root cause: This is by design. The <code>migrateVirtualMachineWithVolume</code> API can only be called by an admin. When the admin intends to migrate the VM on a specified host, CloudPlatform cannot block the admin's choice. When the admin lists hosts for migration, CloudPlatform shows suitable as well as non-suitable hosts in the results. A host that do not match the tag is an example for non-suitable host. However, even if CloudPlatform hints that a host is not suitable, admin is free to pick such a host for migration.</p> <p>Solution: A VM can be migrated from a tagged host to an untagged host. This is because migration of a VM is an admin operation and admin is free to choose where to place the VM.</p>
CS-29396	<p>Problem: Uploaded volume is in unusable state after it failed to migrate across cluster's primary storage. Migration of downloaded volume fails.</p> <p>Root cause: An entry is created in the <code>volume_store_ref</code> table when a volume is downloaded and is in the 'Ready' state. Now, when a volume migration across clusters is attempted, another entry is created in the same table for the volume, but in 'Allocated' state, during the copy of the volume from source primary storage to secondary storage. While copying from secondary to destination primary storage, when CloudPlatform tries to retrieve the migrated volume that is secondary it could potentially end up listing the volume that has been downloaded and marked as 'Ready'. Further, if it lists the wrong volume, migration fails with state transition failure exception. An attempt to attach after the failure to a VM in the same cluster succeeds.</p>
CS-29341	<p>Problem: Resizing Windows XP DATA volume on VMware fails. You see an invalid operation for device '0'.</p> <p>Root cause: Resize is not supported for IDE virtual disks in VMware. Resizing a volume is not supported if the volume is attached to a VM by using IDE disk controller. This is expected behaviour and identical to a resize operation initiated on such a volume from the vSphere client.</p>
CS-28540	<p>Problem: A zone is configured to use local storage but global configuration, <code>system.vm.use.local.storage</code>, is set to false. In this case, if no shared primary storage is configured system VMs does not start up.</p>

Issue ID	Description
	<p>Root cause: It is a limitation in the current code, and <code>system.vm.use.local.storage</code> needs to be made a zone-level configuration.</p> <p>Solution: Zone configuration UI shows an error message and possible workaround when local storage option is selected.</p>
CS-28479	<p>Problem: User VMs after HA may not boot up successfully when primary storage connectivity is lost and restored.</p> <p>Root cause: It is an expected behavior. Lost connection or forceful shutdown of hosts might lead to data lose in VMs.</p>
CS-28478	<p>Problem: When connectivity to primary storage is lost, KVM host continues to be in disconnected state. This occurs because the host gets stuck while trying to get restarted. No HA is triggered until the host is manually power off and power on. System VMs are also marked as being in disconnected state and no HA is triggered for the System VMs as well.</p> <p>Root cause: It's due to a behavior on KVM host, as given in Shutting down, rebooting and force-shutdown of a guest virtual machine¹ Normally, KVM host will try to suspend VMs on host during reboot, which might get stuck when primary storage is unavailable.</p> <p>Solution: Run <code>reboot -f</code></p>
CS-27425/CS-21217	<p>Problem: Windows guest VMs, when used with dynamic scaling, reports 4 times more the amount of RAM. For example, a VM with a 4GB compute offering shows 16GB or RAM within the OS, but only 4 of them are available for use.</p> <p>Root cause: CloudPlatform uses XenServer dynamic memory control (DMC) to enable dynamic scaling of VMs running on XenServer. To use the DMC CloudPlatform sets static max and dynamic max for VMs which are dynamically scalable. The static max parameter refers to the maximum memory to which the VM can scale. Dynamic memory is the actual memory or the current usable memory allocated to the VMs. In case of Windows VMs XenServer incorrectly shows the static max memory allocated to the VM in the Windows task manager.</p> <p>Problem: To get the actual memory allocated in case of Windows VMs, execute the following commands in the power shell as administrator.</p> <pre data-bbox="448 1608 1366 1697"> \$a = gwmi -n root\wmi -cl CitrixXenStoreSession \$a[0].getvalue("memory/target").value </pre>
CS-27564	<p>Problem: dnsmasq service does not start with IPv6 shared network.</p> <p>Root cause: Zone creation prompt for both IPv4 and IPv6 DNS details. If you don't provide IPv6 DNS server details during Zone creation, and if user</p>

¹ https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Virtualization_Deployment_and_Administration_Guide/sect-Managing_guest_virtual_machines_with_virsh-Shutting_down_rebooting_and_force_shutdown_of_a_guest_virtual_machine.html

Issue ID	Description
	<p>creates a shared network with IPv6 address and attempts to deploy a VM in that network, VM deployment would fail.</p> <p>Solution: Edit Zone configuration and provide IPv6 DNS server details.</p>
CS-27114	<p>Problem: When volume snapshot is initiated on KVM, errors are seen on the VM console.</p> <p>Root cause: This issue occurs because VM will be paused for several seconds during snapshot, which is longer time than what observed in previous release, and therefore kernel exception may be thrown after VM is back from pause state.</p> <p>In CloudPlatform 4.5, snapshotting behaviour has been defaulted to that of the hypervisor, thereby enabling you to avail yourself of RedHat support. Depending on the size of the VM memory, you encounter a pause during backing up the VM snapshots. It implies that higher the memory size, longer the pause time.</p>
CS-26006/ CS-27155	<p>Problem: The listVirtualMachines API returns the same keypair if two SSH keypairs are registered with the same public key.</p> <p>Root cause: CloudPlatform does not support registering two SSH keypairs with the same public key. In deploying a VM with keypair parameter passed, CloudPlatform stores its public key in the user_vm_details table with property name "SSH.PublicKey". In the case of deploying two VMs, the same SSH.PublicKey value will be stored in its detail entry. Therefore, the listVirtualMachine API query the ssh_keypairs table with the stored public key, which in turn, returns only the first entry.</p>
CS-25783	<p>Problem: A NIC cannot be removed from a VM on VMware. The guest OS does not respond to a hot-remove request for device ethernet1 in a timely manner.</p> <p>Root cause: On vCenter 5.5, hot unplug requires VMware tools to be installed on the guest OS. Therefore, to add or remove NIC to a guest VM in CloudPlatform 4.5 using vCenter 5.5, the tools should already be running on the guest VM. If the tools are not running, hot remove of NIC causes exception.</p> <p>For more information, see Knowledge Base Article 2081503².</p>
CS-25448	<p>Management Server does not create system VMs on VMware environment due to a template copy error. This is caused because '-R' is removed from the chmod command. This is by design.</p> <p>VMware uses a folder in the machine where Management Server is running to mount secondary storage. This is a bootstrap phase to start system VM, because unlike KVM, XenServer, Management Server cannot directly access VMware ESXI host to download systemVM template from secondary storage to primary storage. The secondary storage is usually managed by SSVM that uses root account to download templates; however, Management Server is using account 'cloud' to manipulate templates after secondary storage is</p>

² http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2081503

Issue ID	Description
	<p>mounted. After the new systemVM template is registered in CloudPlatform as a normal upgrade procedure, the old SSVM downloads the template by using root account, but Management Server will create a new SSVM from the new template using account 'cloud', and a permission denied error is thrown.</p> <p>CloudPlatform handles this by running 'chm -R' to the folder to which secondary storage is mounted every time the Management Server mounts secondary storage. Unfortunately, this method causes low performance if the network connection between the Management Server and secondary storage is slow. Because of that, you must manually run 'Chm -R' to the folder 'templates' on secondary storage, after registering new systemVM template. For example:</p> <pre data-bbox="448 680 1362 741"># chmod -R 0777 /path_to_secondary_storage_mount_point/template/</pre>
CS-26519	<p>The VMs with Windows 8.1 guest OS does not work as expected on VMware.</p> <p>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, and 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)³.</p> <p>Workaround: Based on your environment, perform the following:</p> <ul data-bbox="448 1211 1187 1308" style="list-style-type: none"> • Register Windows 8.1 template\ISO as Windows 8.0 • Register Windows 2012 R2 template\ISO as Windows 2012 <p>For more information, see http://support.citrix.com/article/CTX200336.</p>
CS-23123	<p>Problem: A network is incorrectly being shutdown.</p> <p>Root cause: VMsync functionality is designed to handle out-of-band VM movement from one host to another. For example, one triggered by DRS in VMware. There is no need to decrement NIC count during PowerOff/ PowerMissing operations and increment it during PowerOn. As during out-of-band movement there is PowerOff report on old hosts and PowerOn report on new host. If there is only a out-of-band stop operation or start operation on VM, which is different from out-of-band movement where the VM is always started on a new host, the NIC count and other VM statistics would not get updated as part of VMsync. VMsync will only update the state of the VM. NIC and other statistics would get corrected when an explicit operation is done from CloudPlatform on that VM subsequently.</p>
CS-22985	<p>Baremetal planner may incorrectly deploy an instance with a compute offering having an host tag to a host with no host tag.</p>

³ http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2067000

Issue ID	Description
CS-21274/CS-24354	<p>Problem: Windows VMs with vGPU cards and PV drivers are going to repair state when deployed with Dynamic Scaling option enabled.</p> <p>Root cause: This is a XenServer issue. Dynamic Scaling is not supported in vGPU enabled VMs on XenServer 6.2.5; XenServer 6.5 onwards it works as expected.</p> <p>Solution: Dynamic scaling can be performed on a VM if either of the following is true:</p> <ul style="list-style-type: none"> • VM is deployed with template which has PV tools already installed. • VM is deployed with ISO, tools are installed, and Dynamic scaling option is enabled on this VM. <p>Do not enable Dynamic scaling option on a VM which does not have the VM tools installed.</p>
CS-21259	<p>The <code>job.cancel.threshold.minutes</code> overrides the related waiting time to finish if there is a related waiting time depending on async job's attribute.</p> <p>Therefore, the <code>job.cancel.threshold.minutes</code> should always be greater than the any <code>.wait</code> global parameter, such as <code>create.volume.from.snapshot.wait</code>, <code>migrate.wait</code>, <code>backup.snapshot.wait</code>, <code>copy.volume.wait</code> set on the agent side.</p>
CS-21238	<p>Problem: Host HA will not be triggered in the following cases:</p> <ul style="list-style-type: none"> • KVM host is gracefully powered down. • The admin manually stops cloudstack-agent on KVM host. <p>The host state is changed to "Disconnected" state in the above cases.</p>
CS-21202	XenServer 6.0.2 hotfix releases are not supported for CloudPlatform 4.5.
CS-20847	<p>Problem: Template download does not get initiated after upgrading XenServer host from 5.6 FP2 to 6.2.</p> <p>Root Cause: This is a XenServer issue, as given in XOP-607⁴. During live migration of system VM, it could cause high CPU usage which would cause the VM to be unresponsive.</p> <p>Solution: The workaround is to stop and start the VM.</p>
CS-20721	<p>Problem: Attaching a volume fails with a null pointer exception while attaching an uploaded volume to a VM.</p> <p>Root cause: Attaching an uploaded volume fails because the OVA that was uploaded was incompatible with the underlying ESXi hypervisor host in the infrastructure. In this particular case upon reading the OVF, it was found that the OVA hardware version is 2 is not compatible with ESXi host 4.1.</p>

⁴ <https://issues.citrite.net/browse/XOP-607>

Issue ID	Description
	<p>Solution: No fix from VMware. For more information, see ESXi/ESX hosts and compatible virtual machine hardware versions list (2007240)⁵.</p>
CS-20703	<p>Problem: Baremetal planner may incorrectly deploy an instance with a compute offering having an host tag to an host with a different host tag. For example, an instance with compute offering with host tag A is deployed on a host with host tag B.</p>
CS-20314	<p>Problem: (RealhostIP) CloudPlatform does not detect or throw any error when an expired SSL certificate is uploaded.</p> <p>Root cause: Currently, no validation has been implemented on SSL certification expiry time.</p>
CS-20721/CS-19938	<p>Problem: Attaching volume fails with NULLpointer exception on VMware VMs. This issue is observed when an OVA file is uploaded.</p> <p>Root cause: No mechanism to check for valid hardware versions in a OVF file.</p>
CS-20070	<p>Problem: After upgrading from CloudPlatform 4.2.1 to 4.5, the VPN Customer Gateway functionality goes missing.</p> <p>Root cause: The script to encrypt ipsec_psk during upgrade is missing in version 4.2.1.</p> <p>Solution: To workaround, run the following to encrypt the values of ipsec_psk in the s2s_customer_gateway table:</p> <pre data-bbox="448 1122 1366 1319"> # 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-19967	<p>Problem: CloudPlatform does not advertise remote subnets when configured with a client VPN. You must set up the subnet routing in guest VMs manually.</p>
CS-19993	<p>Problem: The error message for attach volume failure is not self explanatory to reflect the failed operation due to low global timeout value.</p> <p>Root cause: Attaching large data volume might take more than an hour, and this cause subsequent attach volume failure with timeout.</p> <p>Solution: Neglect the error message and increase the value of the <code>migrate.wait</code> global parameter.</p>
CS-19272	<p>Problem: In a XenServer cluster with a primary storage, when a new host is added after all the existing hosts in the XenServer pool is deleted, the existing primary storage is not getting added to the host.</p>

⁵ http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2007240.

Issue ID	Description
	<p>Root cause: CloudPlatform cannot reintroduce the pre-setup primary storage, because no storage detail exists for pre-setup primary storage in the CloudPlatform database.</p> <p>Solution: When all hosts in a cluster are deleted, the existing primary storage can't be used unless you manually perform the following:.</p> <ol style="list-style-type: none"> Run the following: <div data-bbox="592 533 1458 678" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <pre># xe sr-introduce content-type=user name-label=<uuid-of-the- primarystore-to-be-introduced> shared=true uuid=<uuid-of-the-primarystore-to-be- introduced>type=nfs</pre> </div> <p>Where UUID and name label are the UUID of the primary store that need to be introduced.</p> Run the following: <div data-bbox="592 860 1458 976" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <pre># xe pbd-create type=nfs host-uuid=<host UUID> sr-uuid=<SR UUID> device-config:server= <IP address> device-config:serverpath=</path/to/primarystore></pre> </div> Run the following: <div data-bbox="592 1066 1458 1128" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <pre># xe pbd-plug uuid=<UUID of the PBD></pre> </div> <p>For more information on introducing storage repository in XenServer, see Introducing an SR⁶.</p> Force reconnect the host from CloudPlatform.
CS-18799/CS-22788	<p>Problem: If connectivity to the primary storage is down on a KVM/XenServer host, SystemVM after being HAed might remain in alert state.</p> <p>Root cause: This issue occurs because the guest OS on the SystemVM is corrupted when HA is triggered. Therefore, the agent inside CPVM can't connect back to the Management Server.</p> <p>Solution: Manually destroy the SystemVMs.</p>
CS-18216	<p>Problem: Allocated memory more than total memory on a host.</p> <p>Root cause: The KVM agent code doesn't check the availability of memory on KVM host; however, Management Server depends on KVM agent to check the memory, in case of concurrent starting of VMs on a host. If concurrent starting of VMs on the host occurs, KVM agent starts whatever VMs directed by the Management Server and memory overprovisioning might have occurred.</p>

⁶ http://docs.vmd.citrix.com/XenServer/6.2.0/1.0/en_gb/reference.html#id477731

Feature Parity Between CloudPlatform and Apache CloudStack

The following features are available in CloudPlatform as a result of community contributions via Apache CloudStack. Though these features are available in CloudPlatform, Citrix does not offer support.

Feature Category	Feature Details
Network	KVM QinQ VLAN support
Network	Juniper Contrail SDN Plug-in
Network	Palo Alto Firewall integration
Network	NS SSL termination
Network	Nuage VSP Network Plugin
Network	Integration with external DNS provider
Network	Tags for Security Group Rules
SDN	Stratosphere SDN work
VR	VR Extension
VR	VR cleanup
Storage	Clustered LVM Storage support
Storage	Ceph RBD support
Storage	IOPS for data volumes in disk offering (Hypervisor or Storage based) for XenServer and VMware
Storage	IOPS for data volumes in disk offering (Hypervisor or Storage based) for KVM
Storage	IOPS for root volumes in compute offering (Hypervisor-based only)
Storage	Root volume resize
Storage	Volume provisioning type option: thin vs fat, for KVM
Storage	IOPS for root volumes in compute offering, for XenServer and VMware
Storage	Create GUI to add primary storage based on plug-ins
Storage	New NFS storage adapter for KVM hypervisor
Storage	Global Setting to Provision Volumes
Storage	CloudByte storage plugin
Security	SELinux support
Automation/ Puppet integration	Puppet integration
Console Proxy	Console Proxy enhancements
OS	Debian support
Management	Sync Domain/Account/User information across Regions
Management	CloudStack event enhancements

The following are the unsupported UI options in CloudPlatform 4.5.1:

Unsupported UI Options	UI Wizard
Hypervisors: OVM	<ul style="list-style-type: none"> • Infrastructure > Zones > Add Zone • Infrastructure > Clusters > Add Cluster • Infrastructure > Sockets Templates > Register Templates • Other places where hypervisors are listed
Isolation methods: GRE, VNS, SSP	Infrastructure > Zones > Add Zone (Advanced) > Setup Network > Isolation Method
Network Service providers: BigSwitch, MidoNet	Infrastructure > Zones > Select a Zone > Physical Network (Tab) > Select a Physical Network > Network Service Providers > Configure
Swift Storage	Infrastructure > Secondary Storage > Add Secondary Storage > Provider (Swift)
Disk IO Throttling (QoS) added by Solidfire	<ul style="list-style-type: none"> • Service Offerings > Add Compute Offering > Remove the following options: Disk read rate (BPS), Disk write rate (BPS), Disk read rate (IOPS), Disk write rate (IOPS) • Service Offerings > Add Disk Offering > QoS Type Hypervisor > Remove the following options: Disk read rate (BPS), Disk write rate (BPS), Disk read rate (IOPS), Disk write rate (IOPS) • Service Offerings > Add Disk Offering > QoS Type Storage > Remove the following options: Custom IOPS, Min IOPS, Max IOPS
Global Setting to Provision Volumes	<p>“provisioningtype” option will be added to the following:</p> <ul style="list-style-type: none"> • Create disk offering dialog • Disk offering details view
Tags for Security Group Rules	<ul style="list-style-type: none"> • Home > Network - Security Groups > Ingress Rule > Add/Edit/Delete Key,Value tags for the Ingress Rule. • Home > Network - Security Groups > EgressRule > Add/Edit/Delete Key,Value tags for the Egress Rule.
Integration with external DNS Provider	Under Network Provider configuration for a Zone > Credentials and endpoint of API

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/>

