Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Release Notes

Revised on July 31, 2015 11:00 am IST

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Release Notes
Revised on July 31, 2015 11:00 am IST

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. Configuring 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 more information on this feature, refer to *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Getting Started Guide* and *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Hypervisor Configuration Guide*.

1.1.2. Uploading Templates from Your Computer to CloudPlatform

Along with using a URL to upload templates to CloudPlatform, you can upload the templates directly using a web browser and launch VMs using this template. You can upload the templates that you have saved to your computer. Also, this functionality eliminates the need of generating a URL using an external HTTP server to upload a template that you have saved on your computer.

For more information, refer to *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Administration Guide*.

1.1.3. Uploading Volumes from a Local File System to CloudPlatform

Along with using a URL to upload volumes, you can directly upload volumes from a local file system to CloudPlatform. For example, you can upload the volumes that you have saved to your computer. Also,

this functionality eliminates the need of generating a URL using an external HTTP server to upload a volume that you have saved on your computer. CloudPlatform allows you to upload multiple volumes at the same time.

For more information, refer to *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Administration Guide*.

1.1.4. Using DHCP and DNS Services Configured External to CloudPlatform

In a normal CloudPlatform deployment, the virtual router (VR) provides DHCP and DNS services. From the version 4.5.1, CloudPlatform configured with advanced zone and shared network supports the external DHCP and DNS services. You can leverage this feature to use your infrastructure to provide DHCP and DNS services. Also, this feature helps you provide user data that contains sensitive information.

This feature is supported on the following hypervisors:

- XenServer
- VMWare vCenter
- KVM

For more information, refer to *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Administration Guide*.

1.1.5. Disabling Storage Pools in CloudPlatform

When you place a primary storage pool in maintenance mode, the user VMs with volumes on that storage pool are stopped. The router and the system VMs on that storage pool stop first and restart with the volumes on another storage pool. This may have an impact on the end users' experience with the service that you are providing using your cloud infrastructure.

With the ability to disable the primary storage pool, the administrators can prevent further provisioning of storage on that storage pool without stopping the VMs with the volumes on that storage pool. This way, disabling the storage pool helps you manage your storage without affecting your end users.

You can disable all types of storage pools. These types include local, cluster, or zone level storage pools. The status of a disabled storage pool is displayed as Disabled. The status of an enabled storage pool is displayed as Up.

For more information, refer to *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Administration Guide*.

1.1.6. Enabling Local Storage for System VMs at Zone and Cluster Levels

From CloudPlatform 4.5.1, you can configure multiple zones in CloudPlatform that use the storage types configured at the zone level. For this, the scope of the configuration parameter system.vm.use.local.storage has been changed from global to zone level.

With this enhanced capability, CloudPlatform uses two default system VM offerings – one using the shared storage and the other using the local storage – for the two types of System VMs. These offerings are named uniquely. These unique names help the users select these offerings when they deploy System VMs.

When the Management server starts, these System VM offerings are created in the CloudPlatform database. If you are upgrading to version 4.5.1, CloudPlatform checks for the existing system VM offering, provides a unique name to it, and creates the offering that is not available in the CloudPlatform database.

For more information, refer to *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Administration Guide*.

1.1.7. Supporting Management and Storage Traffic on VMWare Distributed Virtual Switch

From version 4.5.1, VMware Distributed Virtual Switch (VDS) configured with CloudPlatform supports management and storage traffic along with guest and public traffic. This enables administrators to share the physical NICs among all traffic types. This provides administrators the flexibility in planning traffic types with less number of physical NICs.

For more information on this feature, refer to *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Hypervisor Configuration Guide*.

1.1.8. Supporting Live Migration for VMFS

From version 4.5.1, CloudPlatform supports the live migration of storage for the VMFS storage in a vSphere deployment. This is in addition to the support for the NFS storage.

For more information, refer to *Citrix CloudPlatform (powered by Apache CloudStack) Version 4.5.1 Administration Guide*.

1.2. API Changes

1.2.1. Configuring Advanced Networking in Bare Metal

API	Description
`addHost`	Added the following class-level parameters: • **`cpunumber`** - Number of CPUs in the host. • **`cpuspeed`** - CPU speed of the host. • **`memory`** - Memory of the host. • **`hostmac`** - MAC address of the host. You can use these parameters only for adding baremetal hosts.

1.2.2. Uploading Templates from Your Computer to CloudPlatform

API	Description
`getUploadParamsForTemplate`	Enable users to upload template from a local or network file share.

1.2.3. Uploading Volumes from a Local File System to CloudPlatform

API	Description
`getUploadParamsForVolume`	Enable users to upload volume from a local or network file share

1.2.4. Using DHCP and DNS Services Configured External to CloudPlatform

No API changes.

1.2.5. Disabling Storage Pools in CloudPlatform

API	Description
`updateStoragePool`	Added an optional parameter - **`enabled`** - to enable or disable a storage pool. If you do not specify this parameter, you cannot change the state of the storage pool.

1.2.6. Enabling Local Storage for System VMs at Zone and Cluster Levels

No API changes.

1.2.7. Supporting Management and Storage Traffic on VMWare Distributed Virtual Switch

No API changes.

1.2.8. Supporting Live Migration for VMFS

No API changes.

1.3. System VM Templates

CloudPlatform 4.5.1 supports 64-bit System VM templates. This release does not provide 32-bit support for System VM templates.

Hypervisor	Description Please provide the latest System VM info
XenServer	*http://download.cloud.com/templates/4.5.1/systemvm64template-2015-05-14-4.5.1-xen.vhd.bz2*
Hyper-V	*http://download.cloud.com/templates/4.5.1/systemvm64template-2015-05-14-4.5.1-hyperv.vhd.bz2*
KVM	*http://download.cloud.com/templates/4.5.1/systemvm64template-2015-05-14-4.5.1-kvm.qcow2.bz2*

Hypervisor	Description Please provide the latest System VM info
VMware	*http://download.cloud.com/templates/4.5.1/systemvm64template-2015-05-14-4.5.1-vmware.ova*

For more information, see **Appendix A: Latest System VM Templates** in the CloudPlatform 4.5.1 Installation Guide.

1.4. Fixed Issues

Issue ID	Description
CS-42434	Problem: CloudPlatform does not add a USB controller to Apple Mac OS X VMs created using an ESXi hypervisor. However, vSphere Client adds a USB Controller to the Mac OS VMs by default. Mac OS X machines require USB Controller for USB mouse and keyboard access. Root Cause: The Guest OS details are specified in the Virtual Machine Configuration Specification for creating the VM (using the SDK API) in the ESXi hypervisor. No USB Controller is added to the Virtual Machine Configuration Specification. As the guest OS identification details are specified in the VM Configuration Specification, it is assumed that the Create VM SDK API would create the defaults in the VM same as vSphere Client. However, as per the observation, USB Controller is not added to the Guest OS - Mac OS VM created through the SDK API. Solution: When the Guest OS is Apple Mac OS, add the USB Controller (EHCI+UHCI - Mac supported) to the Virtual Machine Configuration Specification before creating or starting the VM. For any existing Mac OS VMs, stop and start the VM to add the USB Controller. For new VMs with Mac OS, USB Controller is added automatically.
CS-42324	Problem: Storage XenMotion fails after applying CloudPlatform 4.5.0.0 HF 5. Root Cause: Storage motion of VM across clusters/XenServer-pools fails in a clustered Management Server configuration. In XenServer storage motion, the **`migrate_receive`** command is being sent to the destination host followed by the **`migrate_send`** command to the source host. The storage and network details of the destination host will be passed to the source through the **`migrate_send`** command. While migrating across clusters, the source and the destination resources are separate objects. To pass this information across resources, need to send separate migrate with storage receive and send commands to the resource. In a clustered Management Server setup these commands may have to be forwarded to another Management Server as the resource may be owned by it. In such a case, the serialization of the command and the answer objects fails as it does not understand the xapi storage and network objects. Solution: Ensure that the xapi objects are serialized/de-serialized in the resource layer as it is aware of the object types.
CS-42032	Problem: Stale NFS secondary storage on XenServer leads to volume creation failure from snapshot.

Issue ID	Description
	Root Cause: This issue comes in case of parallel operation on snapshots or any thing else that involves mounting of secondary storage on host. CloudPlatform mounts secondary storage by creating SR on XS host where SR name-label is derived from the XenServer uuid. In case of XenServer pool (with two or more hosts), if two hosts are mounting the secondary storage with same mount point by creating two SRs, second SR-scan will run into error. This is because XenServer hosts in a pool share the same XAPI database and if XenServer scans two SR with same files - that is, VDIs - XAPI throws the Db_exn.Uniqueness_constraint_violation error. Solution: Create a single SR with same mount point across the pool by changing the SR name-label derivation method and handle the remove SR part as well in the same manner. That is, if the SR is being used by some other operation, you do not delete it.

CS-42015	Problem: **`prepareTemplate`** API call does not work properly with XenServer and local SR. The following error message is displayed:
	
	`Db_exn.Uniqueness_constraint_violation`
	
	Root Cause: CloudPlatform creates a SR on each host, which points to the template location on the secondary storage (secondary_Storage/template/tmpl/<account_id>/<template_id>). This causes the database unique constraint violation when each XenServer tries to scan the SR created on each host. The host that scans the SR last, throws the exception because VDI was recognized already from the SR scan of the first host.
	
	Solution: SR label is modified so that it can be queried by all the hosts in the pool and reuse the same SR if available. The prepareTemplate API now works fine with the local storage pools. Additionally, you must perform the following steps to perform SR clean-up:
	
	1. Logon to your XenServer host and run the following command:
	
	`xe sr-list`
	
	This lists all the SRs. If the name-label has a path such as /var/cloud_mount/*, remove it.
	
	2. List the pbds for that SR.
	
	`xe sr-param-list uuid=<uuid of sr you want to remove>`
	
	3. Unplug the pbd.
	
	`xe pbd-unplug uuid=<uuid of the pbd obtained in previous step>`
	
	4. Remove/forget the SR.
	
	`xe sr-forget uuid=<uuid of the sr>`

CS-41997	Problem: Unable to delete an IP tag.

Issue ID	Description
	Root Cause: Due to wrong search string, CloudPlatform fetches a resource, which does not belong to current account and domain. This results in permission exception. Solution: Code has been modified to fix the incorrect search string.
CS-41964	Problem: In an RVR network with the router in the Stopped state, the network is restarted in the backend. However, UI/API shows error. Root Cause: While implementing network in **`VirtualRouterElement`**, called the **`deployVirtualRouterInGuestNetwork`** method to start the router. On implementing, two routers are expected from the **`deployVirtualRouterInGuestNetwork`** method. But, it returns one router, which had been stopped previously. Due to this, the implementation method throws the following exception: **`Can't find all necessary running routers!`** Solution: In **`virtualnetworkappliancemanagerimpl`** updated **`findOrDeployVirtualRouterInGuestNetwork`** to return two routers - the router that is in the Running state and the router that is previously stopped - so that implementation method can get two routers.
CS-41913	Problem: Usage server reports indicate incorrect usage of VM snapshots. Root Cause: Delete snapshot event is not processed properly and this results in duplicate usage for a VM snapshot. Solution: Code has been modified to process delete event appropriately.
CS-41901	Problem: When a root administrator tries to create a port forwarding rule on an IP address that belongs to a different domain, the UI calls listNetowrks API by passing the id parameter. Because the network does not belong to the root administrator and CloudPlatform does not pass **`listall=true`** to the API, the network does not get listed. This causes the next **`listVirtualMachines`** API to fail, causing the error. Solution: Passing the domain id and the account id along with the id of network to make sure that the network is getting listed.
CS-41834	Problem: In Windows VMs, default gateway is configured for non-default network. This occurs even after setting the value of the **`network.dhcp.nondefaultnetwork.setgateway.guestos`** configuration parameter to NULL. The default value of this parameter is **`Windows`**. Root Cause: While creating the DhcpEntry command, the **`network.dhcp.nondefaultnetwork.setgateway.guestos`** configuration parameter is not referred. Solution: Code has been modified to use the **`network.dhcp.nondefaultnetwork.setgateway.guestos`** configuration parameter while creating the **`DhcpEntry`** command.
CS-41826	Problem: Cannot attach more than 6 disks to a Windows 2012 R2 VM. Root Cause: While trying to obtain the SCSI id to attach a disk, CloudPlatform should ignore the reserved SCSI id 7. However, this

Issue ID	Description
	is not being honored in case of VMs with SCSI controller of type **`VirtualLsiLogicSASController`**. So, in case of Windows 2012 R2 VMs, CloudPlatform chooses to attach the 7th disk on the reserved SCSI id and this fails on vCenter. Solution: During disk attach, while trying to obtain the controller key for SCSI controller, look for device with the generic SCSI controller type, that is, **`VirtualSCSIController`**.
CS-41822	Problem: The **`lsof`** command inside SSVM reports cannot identify the protocol. Root Cause: Cloud service does not close the socket in case the Management Server is not reachable from SSVM. Solution: Code has been modified to close the socket properly in case of any exception.
CS-41720	Problem: After upgrading from CloudPlatform version 3.0.7 to version 4.5.0, snapshot creation fails for the volumes that contain previous snapshots. Root Cause: While populating snapshot_store_ref table as part of upgrade, snapshots that are removed are not considered. Solution: Provided the following script with this hotfix to consider the removed snapshots with the dependent snapshots: **`upgradesnasphot.py`** This script has been integrated in the hotfix installation utility. So, the users do not need to perform any additional task to run this script.
CS-41715	Problem: Upgrade of CloudPlatform version 4.5 to version 4.5.1 fails with a foreign key constraint. Root Cause: The upgrade failed while adding the following constraint: **`ALTER TABLE `cloud`.`guest_os_hypervisor` ADD FOREIGN KEY (`guest_os_id`) REFERENCES `cloud`.`guest_os`(`id`);`** The DB is in an inconsistent state. For upgrade to succeed the below query should have returned an empty list: **`mysql> select * from cloud.guest_os_hypervisor where guest_os_id not in (select id from cloud.guest_os);`** Solution: All entries returned by the above select query needs to be fixed manually or to be cleaned up.
CS-41711	Problem: Increased latency for the list* APIs at scale. Root Cause: ACS commit b0c6d4734724358df97b6fa4d8c5beb0f447745e (one of the earlier modifications of the code) caused performance regression, because the code scanned the entire response string to remove sensitive data before logging it. At scale, it caused performance issues. Solution: To improve performance of the list* APIs, removed the part of the code that scans the sensitive data.

Issue ID	Description
	 **Note** Also, this causes the display of sensitive information from API responses in the API logs.
CS-41693	Problem: A windows VM on XenServer with three or more disks attached to it fails to start. Root Cause: A windows VM or any HVM guest has only four (0, 1, 2, 3) vbds available for attaching disks/iso when the VM is in powered down state. CloudPlatform uses device id 0 for root disk and device id 3 for iso. Data disks were connected on device id 1 and device id 2. When checked if another vbd is available for checking, it returned false. When tried autodetecting vbd, it failed on a powered down VM. The vbd available check is not mandatory and if a valid device id 4 or higher number is used, it should work. Solution: Fixed the issue by using a valid device id.
CS-41692	Problem: Cannot start Windows VMs with more than three disks after upgrading to CloudPlatform 4.5. Root Cause: In case of delta snapshot, CloudPlatform needs the latest snapshot of the volume in the primary storage as well as in the image store. However, CloudPlatform is not checking for the latest snapshot of that volume in primary store. This leads to delta snapshot failure. Solution: In case of delta snapshot, check for the latest snapshot of that volume in the primary storage and the image store.
CS-41677	Problem: In a scaled up environment (with more than 5000 VMs), CloudPlatform performance is impacted because the database is getting slow due to a continuous stream of certain un-optimised queries generated as a result of VM sync (runs every ping.interval for each host). Root Cause: Queries are slow as certain fields are not indexed in the vm_instance table. Solution: Added an index based on the following query: **`ALTER TABLE cloud.vm_instance ADD INDEX i_vm_instance__instance_name(instance_name);`**
CS-41127	Problem: CloudPlatform UI lists IP addresses from the first NIC of the VM while adding port-forwarding or load balancer rule, regardless whether the NIC is default or not. Solution: Instead of passing the id of the first NIC, loop over NICs to identify the default NIC and pass the id of that NIC to the API.
CS-41079	Problem: Instances fail to start after unsuccessful compute offering upgrade.

Issue ID	Description
	Root Cause: In case of unsuccessful compute offering upgrade, the rollback does not take care of restoring the old service offering details. Solution: Code has been modified to restore the old service offering details.
CS-40989	Problem: Logging on to Linux Guest VM from Firefox 36 or Firefox 37 results in an abnormal reboot of the Guest VM. Root Cause: While accessing the VM console from the CloudPlatform UI, the Ctrl+Alt+Del keyboard combination gets activated on pressing some keys (for example, the '/' key). Pressing ENTER followed by this results in the VM reboot. Solution: Disable quick find keyboard shortcut "/" on VM console window.
CS-40945	Problem: Several hosts get disconnected because the **`DirectAgent`** thread pool get exhausted in the Management Server. Root Cause: CloudPlatform Management Server uses the DirectAgent threads to execute the ping task at regular intervals to monitor the host status for every host that it owns. Lack of free DirectAgent threads results in the non-execution of ping task for the hosts and eventually the hosts are marked as disconnected. In this particular case, it is noticed that the **`DirectAgent`** thread used for running the **`CheckS2SVpnConnectionsCommand`** script causes this indefinite delay. Solution: Ensure that CloudPlatform does not wait indefinitely for a script to execute in the virtual router by handling channel condition of End of File (EOF).
CS-40941	Problem: Volume snapshot operation fails on a particular volume. Root Cause:If any of the steps involved in taking the snapshot of a volume takes more than 20 minutes (the session between CloudPlatform and vCenter times out in 20 minutes) to complete, snapshot backup operation fails and the snapshot moves to the 'Error' state. If your environment contains a volume that is larger in size, snapshot operation on the volume attached to this VM takes a longer time to complete. This is because vCenter takes a longer period to consolidate VMs. To address this problem, you may need to set the duration of CloudPlatform's session to vCenter during snapshot operation to a value higher than the time taken for the snapshot operation. However, setting the existing global configuration to such a higher value could affect all operations in the cloud. So, this is not considered as an ideal workaround. Solution: Split the existing vCenter session timeout configuration into the following two configurations: • Use the **`vmware.snapshot.backup.session.timeout`** parameter to tune timeout value for CloudPlatform's session to vCenter during the snapshot backup operation. Default value of this parameter is 20 minutes. • Use the **`vmware.vcenter.session.timeout`** parameter for all other operations. Default va;lue of this parameter is 20 minutes.

Issue ID	Description
	This way, setting a high timeout value for snapshot backup operation does not apply to vCenter connections made for any other operations.
CS-40743	Problem: CloudPlatform async call registers the failure of the RevertToVMSnapshot command as successful. Root Cause: If the RevertToVMSnapshot command fails, CloudPlatform async call registers it as successful. This is because CloudPlatform never returns the call as unsuccessful even if it is a failure. Solution: If the RevetToVMSnapshot command fails, CloudPlatform returns async call as failure.
CS-40584	Problem: The base API URL to interact with Dell Switch is hardcoded. Solution: A new configuration **`baremetal.switch.base.url`** is added. Using this parameter, the administrators can customize the base API URL for Dell Switch. The default value is: **`/api/running/ftos/interface/`**. This configuration is valid only in Advanced Zone for Bare Metal.
CS-40579	Problem: VMs experience an undue delay in completing their deployments in CloudPlatform over VMware. Root Cause: The CheckRouter command that runs periodically makes CloudPlatform fetch the systemvm key file and transfer it to VR using an SSH connection. In this process, CloudPlatform connects to vCenter, which results in delay and affects performance. Because CloudPlatform fetches the systemvm key file more than once in the life cycle, the VM deployments experience an undue delay. Solution: To address this issue, CloudPlatform avoids connecting to vCenter when it fetches the systemvm key file location/path. Also, the fetch operation would now run only once in life time of the service, because the systemvm key file location/path would be stored in the VmwareManagerImpl class variable for further references.
CS-40568	Problem: Migrating primary storage causes the name_label field to appear blank. Root Cause: After migrating root disk or data disk, CloudPlatform was not setting any name label to destination VDI. Solution: Set name-label of vdi same as the name given in CloudPlatform.
CS-40545	Problem: vStorage API for Array Integration (VAAI) for NAS deployment does not function properly in CloudPlatform. Root Cause: While setting up template in primary storage, CloudPlatform takes the base snapshot over template. This snapshot over template prevents VAAI to start during the clone operations over such template. This makes ESXi host to perform the storage operations. Solution: Introduced a new zone-level configuration parameter: vmware.create.base.snapshot that enables/disables snapshot creation over template while setting up template in primary storage. If you want to leverage VAAI integration, set the value of this parameter to 'false' to prevent CloudPlatform creating snapshot over template. Also, you must set the value

Issue ID	Description
	of the global configuration parameter: vmware.create.full.clone to ‘true’ for avoiding the failure of template set up in primary storage. These settings let VAAI kick in for offloaded storage operations.
CS-40506	Problem: The database server configured with CloudPlatform responds slowly after the users upgrade to CloudPlatform 3.0.7 Patch G. Root Cause: In CloudPlatform 3.0.7 Patch G, CloudPlatform synchronizes the time offset data of VMs periodically. As part of this synchronization, the resource layer sends the time offset data of the VMs to the Management Server. The response of the database that caters to a huge number of VMs decreases when it is queried very often for updating the time offset data of the VMs. Solution: To address this issue, you must ensure the following: • Do not send the time offset data from resource layer if the VM does not have this data. • Send the time offset data of the VM only if it has changed. For this, you must keep the time offset data of the VMs, which have changed, in memory. • Query the database only once for each host if there is a change in the time offset data of any VMs. Then, update the database accordingly.
CS-40430	Problem: CloudPlatform deployment over VMware fails to configure a tier in VPC VR. Root Cause: Failure to configure a tier in VPC VR is because of the delay in detecting NIC by the guest OS (Debian 7 x64) or VR. The type of Ethernet adapter (VMXNET3) in the environment causes the delay in detecting NIC by the guest OS. Solution: To address this issue, the timeout period for NIC to be detected by guest OS is increased from 15 seconds to 40 seconds. Use the global configuration parameter: vmware.nic.hotplug.wait.timeout to configure this time out period. Also, improved the logic to wait for the NIC and NIC’s index inside VPC VR, limiting the successful check over each NIC/MAC to 1.
CS-40244	Problem: Usage sanity checker should report errors if there is duplicate usage. Root Cause: New checks are being added to sanity. Solution: Added check to find duplicate usage entries with reported usage > aggregation range.
CS-40192	Problem: SSVM reboot/shutdown during snapshot operation results in orphan workervm, snapshot, and disk. Root Cause: When SSVM is destroyed while a snapshot operation is in progress, disks that are partially created in secondary storage are not cleaned up. These disks will be found under the snapshot's install path.

Issue ID	Description
	Solution: Storage garbage collector has been enhanced to handle this. While deleting 'ERROR' snapshots, CloudPlatform will now also clean up any disks that are left behind in the secondary storage.
CS-40166	Problem: VM migration with storage fails intermittently in a clustered management server. Root Cause: When a request is made to migrate a VM along with its storage, CloudPlatform spawns an internal VM work job (VmWorkMigrateWithStorage). If this job is being processed by a Management Server that is not the owner of the host that the VM belongs to, then the request will be forwarded to the Management Server that owns the VM's host. In case of any forwarded request, CloudPlatform attempts to de-serialize the JSON request. But in case of 'MigrateWithStorageCommand' this de-serialization fails. This issue will be intermittently observed in a clustered Management Server because the issue will only occur if the VM Migrate job is not being processed by the Management Server that owns the host that the VM being migrated belongs to. Solution: Update **`MigrateWithStorageCommand`** such that when forwarded to another Management server, JSON de-serialization is successful.
CS-40141	Problem: The order in which the templates are returned in the API response are inconsistent for a given page with a given page size. Solution: Added additional order-by clause on the temp_zone_pair column in the SQL query to fetch the templates.
CS-40079	Problem: CloudPlatform UI allows domain administrators to create public templates from snapshots. This occurs because of a missing pre-filter in javascript. Solution: Added the pre-filter in the javascript to address this issue.
CS-39934	Problem: Using the listTemplates API with the "templatefilter=all" parameter lists all the templates that are available with all domains in the system. Root Cause: The "templatefilter=all" filter is not implemented properly for domain administrators. Solution: Added additional filters to list all the templates within its own domain for a domain administrator
CS-39913	Problem: The Hyper-V agent fails to start on a Hyper-V host. Because of this problem, the Hyper-V agent installation rolls back. Root Cause: When the Hyper-V agent starts on a Hyper-V host, it queries the NIC details on the host. In certain cases, the Hyper-V agent fails to query the details and subsequently the agent fails to start. Solution: Fixed the agent code to not to query the details during service start-up. Agent initializes after the management server contacts it.
CS-39578	Problem: Template created within the project context is not visible in the project view.

Issue ID	Description
	Root Cause: This issue occurs because the new templates are created against caller's account id. Solution: A new parameter 'projectId' has been added to the createTemplate API and the ownership is decided based on the current user account and the projectId.
CS-39572	Problem: CloudPlatform deploys VR at the same cluster or at the same host. This undermines the purpose of redundancy because the redundant VR cannot function if a problem occurs to the cluster or the host where both the routers are deployed. Root Cause: This deployment occurs because the deployment planner does not consider the destination details configured for the master router. Solution: To address this issue, the deployment planner has been fixed in such a way that it deploys the backup router based on the destination details of the master router. Ideally, the deployment planner deploys the master router and the backup router on different cloud resources. If the planner is unable to find the suitable destination for the backup router, it may deploy the backup router on the same cloud resource where the master router has been deployed.
CS-39570	Problem: Windows VMs does not unlock the screen when using Ctrl+Alt+Del button in the console window in IE version 9,10,and 11. This occurs only on Windows server 2012. Root Cause: This issue occurs because the console loses focus when clicking Ctrl+Alt+Delete button combination in Internet Explorer. Again, the Ctrl+Alt+Delete button combination is not visible when focussing on the console by clicking anywhere in console window. This issue can be seen when the whole console does not fit into browser window. That is VM resolution is greater than 800x600 and scrollbar gets created to fit the console in browser window. Ctrl+Alt+Delete button combination works fine if the windows VM resolution is 800x600 or less. One more way to overcome this issue is by allowing user to resize window such that the whole console fits in the browser window without any scrollbar. Solution: Make the console window in IE resizeable so that user can increase its size to fit the console without the need of scrollbar.
CS-39510	Problem: After an event fails, its status is wrongly updated in the usage_event table as a create event. This results in the incorrect calculation of the total resource usage of the customer. Root Cause: When an event is created, its status is logged in the 'usage_event' table. However, CloudPlatform does not log usage events for certain state transitions of VMs. Solution: CloudPlatform has been enabled to log appropriate events for various VM state transitions in the usage_event table.
CS-39479	Problem: cloudPlatform UI does not show all the domains.

Issue ID	Description
	Root Cause: Due to an incorrect CSS property, if there are more than 25 domains, the last domain gets cut off from the UI. Solution: Fixed the CSS property such that even the last domain is shown properly on the UI.
CS-38921	Problem: Management Server does not provide proper error message while deleting a default ACL. Root Cause: ACL type should have been verified before evaluating the access. Solution: Code changes has been done to verify ACL type and raise appropriate exception.
CS-38869	Problem: Virtual routers (VR) do not start after upgrading to CloudPlatform 4.5.0.0. Root Cause: In XenServer, the VR configuration is first transferred to vmops plug-in as an argument of the command before it is transferred securely to the VR using SCP. The size of the argument in the command exceeds the argument size set by the ARG_MAX parameter. The large quantity of content in the command results in the failure of the command. Solution: To resolve this issue, first copy the file contents in bytes to the host using SCP. This helps in copying large quantity data to remote host. After the configuration file is created in the XenServer host, transfer it to the VR.
CS-38851	Problem: In CloudPlatform, the Domain Administrators can create public template from a snapshot, however, they cannot register a public template. Root Cause: CloudPlatform allows domain administrators to create a public template from a snapshot, however, they cannot register the public template. This is a conflicting behaviour. Solution: Restrict the domain administrators from creating public template from a snapshot. Now, domain administrators will not be able to create public templates from snapshots. The behaviour is now consistent.
CS-38821	Problem: CloudPlatform UI does not display source CIDR in VPC ACL configuration. Root Cause: The CIDR to network ACL mapping is re-factored into a new table, and the **`listNetworkAclItems`** API was not fetching the CIDR information during list call. Solution: Fixed the code such that the CIDR information is now populated in the network ACL items.
CS-38809	Problem: The **`listVMSnapshot`** API in CloudPlatform does not have the 'project' and the 'projectid' tags. Root Cause: CloudPlatform does not set the project ID and the project name in VM snapshot response object if the account is linked to project. Solution: Code has been modified to address this issue.
CS-38779	Problem: When migrating to another storage, the resource tags on a disk are lost.

Issue ID	Description
	Root Cause: During cold volume migration, CloudPlatform duplicates volume entry in volumes table. When migration is complete, CloudPlatform updates the uuid of new entry and expunge the older entry. This results in removal of resource tags on volume as its resource id still pointing to older volume. Solution: While updating uuid for volume, CloudPlatform should also update resource_id for the tags.
CS-38771	Problem: When VRs are moved out-of-band by VMware (due to native HA), they needs to be rebooted for all rules to be re-programmed. If there are multiple VRs are moved (for example, around 20 - 30 VRs are moved), CloudPlatform takes a long time to reboot all of them. Root Cause: Only a single thread is available to process VR reboots due to out-of-band movement. This slows down the entire process as only one VR is processed at a time. More over, the same thread is also used to run another task (query VR stats), which runs at fixed interval based on **`router.stats.interval`**. Solution: A separate pool of threads is used to handle VR reboots. Also, now, multiple VRs can be rebooted in parallel and this reduces the overall time taken for reboot.
CS-38573	Problem: When creating a VM with a custom compute offering, CloudPlatform allows the input of double-byte numbers. The VM is created, however, **`listVirtualMachines`** fails subsequently with an exception. Root Cause: When using the double byte numbers to specify values for custom compute offering, CloudPlatform persists the double byte values in the database. This causes the **`listVirtualMachines`** API failure. Solution: Modified the code such that the double byte values are not persisted in the database when using custom compute offering.
CS-38546	Problem: A VR was getting deployed on a host under a pod, which was disabled. Root Cause: A VR was getting deployed under System account. This resulted in the Management Server picking up the host, which may be available in a disabled pod. Solution: Deployed the VR in the context of the user. This makes sure that the disabled pods are not picked up for deploying VRs.
CS-38467	Problem: In XenServer, the VMs fail to start because of the change in the order of ROOT and DATA disk in the boot order. Root Cause: Because of the incorrect handling of disk device ID in CloudPlatform, the order of ROOT and DATA disks in the boot order changes. This occurs for the operations such as restoreVirtualMachine and stop/startVirtualMachine. Consequently, the VMs, especially Windows VMs, fail to start and become unusable. Solution: A fix has been introduced to ensure the proper handling of device ID, which will correct the boot order of ROOT and DATA disks. The ROOT disk will always boot first.

Issue ID	Description
CS-38465	Problem: The maximum guest limit set for the hypervisor is checked during VM deployment. However, the check seems to fail if the limit is reduced to a lower value. Root Cause: The code checks whether the configured max_guest_limit has been equaled. If yes, it skips the host. However, if the host already has VMs more than this limit, the host passes through this check since the equals check is not relevant here. For example, you can easily verify this by setting the limit to 39 and having 39 VMs on host H1. In this case, this host will not be used anymore for deploying the 40th VM. Solution: Modify the code to check if the VMs are >= the limit in number. If so, the host will be skipped for deployment. Also, this check needs to be used for starting a stopped VM.
CS-38220	Problem: Creating shared network fails when the network CIDR matches the existing guest CIDR in a zone. Root Cause: The issue occurs because CloudPlatform on creating a shared network checks if there is any guest network already implemented with same CIDR and throws the error without checking if they have the same VLANs. Creating same CIDR shared network with different VLAN should be allowed. Solution: When creating a shared network, if CloudPlatform identifies an existing guest network with the same CIDR, CloudPlatform checks if they have the same VLAN. If they have same VLAN, CloudPlatform do not allow creating it. If they are same CIDR with different VLANs, CloudPlatform allows creating the network.
CS-38168	Problem: CloudPlatform API does not return the tags **`vmstopped`** or **`vmrunning`** when their value is zero. Root Cause: Not adding the **`vmstopped`** tag and the **`vmrunning`** tag to the response XML when the corresponding entries are zero. Solution: Fixed the **`listProjects`** response by adding these tags even if there are no running or stopped VMs.
CS-37800	Problem: CloudPlatform fails to expunge instances. Root Cause: CloudPlatform fails to expunge VMs that have rogue volumes left behind in the 'Expunging' state due to a failed volume migration. Solution: If a volume has not been created on storage, CloudPlatform should not send expunge request to the underlying hypervisor resource. This can be achieved by setting the destinations volume's path only after a successful migration.
CS-37740	Problem: Displays template registration errors when the URL is HTTPS. Root Cause: Default java keystore is override by **`realhostip.keystore`** and does not have necessary certificates. Solution: SSVM start up script has been modified to import java certificates to **`realhostip.keystore`**.

Issue ID	Description
CS-35216	Problem: The **`ListUsageRecords`** API returns null usage_id for migrated volumes. Root Cause: When a volume is migrated, a new entry is created in the **`volumes`** table and the old volume entry is marked as removed. Usage events for new volume creation are not logged. Solution: Moved volume usage events to be in sync with volume state machine. Correct volume usage events are logged during migration.
CS-34835	Problem: One ESXi host stuck disconnects frequently. Root Cause: Even though CloudPlatform is able to ping a host, the host is found to be behind on ping. This is because the Management Server's ping map is not being updated for the host. When CloudPlatform pings a host, it updates Management Server's ping map for the host. CloudPlatform also scans and updates all VMs of the host that are stuck in a transitional state and are missing from the power report. This task always fails in case of the host under observation because it has VMs that are in a transitional state. That is, 'Starting' but have been marked as removed. Because the ping task is never successfully completed, Management Server's ping map does not get updated for the host. Because of this, the host is always behind on ping. Solution: During ping task, in the host that are stuck in a transitional state, scan and update only for VMs that are not removed.
CS-34556	Problem: CloudPlatform does not accept the value set to the global configuration parameter: **`system.vm.default.hypervisor`**, which indicates the hypervisor type for System VMs. Root Cause: CloudPlatform does not accept the value set to the global configuration parameter: **`system.vm.default.hypervisor`**. Instead, CloudPlatform accepts the hypervisor type of User VM's destination host. Solution: Add the hypervisor type of System VMs that is provided as the value for the global configuration parameter: **`system.vm.default.hypervisor`** to the list of supported hypervisors at zeroth index. Then, add the hypervisor type of the user VM's destination host. This will enable CloudPlatform to first deploy systemVM's hypervisor type before it deploys the hypervisor type of the userVM's destination host.
CS-34532	Problem: In the Project view, site-to-site VPN gateways are not listed within a VPC. Root Cause: The projectId is not considered during ACL check while listing S2S VPN gateways. This returns empty response. Solution: Included projectId during ACL check while listing S2S VPN gateways.
CS-33740	Problem: Incorrect HTTPS configuration in the default **`server.xml`** file.

Issue ID	Description
	Root Cause: Path to the **`cloud-localhost.pk12`** file was incorrect because of the '\' (backslash). Solution: Replace '\' (backslash) with '/' (forward slash) in the path to the **`cloud-localhost.pk12`** file.
CS-32841	Problem: The **`listVolumes`** API fails for a particular domain with NPE. Root Cause: This occurs when the volume associated VM instance has NULL or invalid state. Solution: Fix the code to guard this situation because this should not block volume listing.
CS-32556	Problem: When the agent LB starts up, many of the hosts are rebalanced successfully. However, certain hosts are not being rebalanced and get stuck in the op_host_transfer queue forever. Root Cause: Hosts that are not being rebalanced, are not Routing hosts. However, the entries for Console Proxy and SSVM VMs are available in the hosts table. Such hosts should not be considered for rebalancing. Solution: SSVM and CPVMs are excluded from agent rebalancing.
CS-32525	Problem: The **`listCapacity`** API does not list a few capacity types as capacity_states were disabled in database. Root Cause: If the capacitychecker thread runs during zone creation process, it sets the capacity_state as disabled (allocation state of the zone under creation) and never gets enabled. Solution: For zone dependent capacities, whenever capacitychecker thread runs, it updates the capacity_state depending on the allocation state of the zone.
CS-31872	Problem: In RVR deployment, the master router will get loss of the internet connectivity even though backup router restarts. Root Cause: In RVR setup, when temporarily connection issue occurs, upstream switch would send packets to the router which sends arping later, mostly BACKUP because ESXi will send RARP (Reverse-ARP) If an interface of VM is down or up, it indicates that the port of L2 switch that has backup router owns MAC address, which is same as that of vNIC of the master router. Solution: No code has been modified for resolving this issue. Disable RARP on all vSwitches in the hosts of the zone where the RVR could move to. the following KB article from VMware shows the steps to disable RARP in vSwitch. RVR, in this issue, suffered the same problem that MS NLB experienced in the case described in the KB. *http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1556*
CS-31829	Problem: Irrespective of the number of volumes snapshots scheduled at the same time, only five of them are being processed actively at the same time.

Issue ID	Description
	Root Cause: Thread pool, Indirect agent communication thread pool is used for agent communication in Management Sever and in SSVM, CPVM, and KVM hosts. In Management Server, size of the above thread pool is **`workers`**+10 where **`workers`** is a global configuration parameter with five as the default value. But in SSVM, the size of the thread pool is not configurable using this global configuration parameter. With five as default, only five snapshots can be backed up to secondary storage simultaneously. Solution: Allow the number of threads on SSVM agent for connection with Management Sserver (Agent->NioClient) to be configurable using the global configuration parameter: **`workers`**.
CS-30962	Problem: Planner allocates new storage pools for Ready data disks within the same cluster. This triggers data migration that can be avoided Root Cause: StoragePool reuse logic starts only if ROOT volume is Ready. However, for restoreVM, ROOT is re-created. So this logic does not get invoked for restoreVM. Solution: Reuse the storage pools for any Ready disk if the pool fits the deployment plan (dc/pod/cluster).
CS-30607	Problem: Root administrator cannot view or manage the LB rules created by the domain administrator in the CloudPlatform UI. Root Cause: The **`listloadbalancerrules`** API does not contain the **`listall`** parameter. Because of this, the root administrators are not able to view the load balancer rules created by domain administrators. Solution: Added the **`listall`** parameter to the **`listloadbalancerrules`** API.
CS-30372	Problem: If multiple snapshots of a volume are queued up for backing up in the secondary storage, the latest snapshot fails to back up. Root Cause: If multiple snapshots of a volume are queued up for backing up to secondary storage, CloudPlatform backs up the snapshot that is at the top of the queue and deletes the other snapshot's VDI from hypervisors. This occurs irrespective of whether the snapshots are backed up or not. You can view the following error message: **`The uuid you supplied was invalid.`** Solution: Now, CloudPlatform verifies snapshot's timestamp. CloudPlatform deletes the snapshot only if it is created before the snapshot that is backed up in the secondary storage is created.
CS-30097	Problem: VM are stopped minutes after they are started in CloudPlatform Root Cause: HA work, which is "Done", is scheduled for some HA jobs. This retries infinitely. Solution: You do not schedule HA work, which is "Done. By default, it will retry five times. You can change the global configuration parameter: max.retries to set a different retry time.

Issue ID	Description
CS-29420	Problem: When trying to launch a VM on a network that uses the network offering with the Loadbalancer that F5 provides, errors are observed in the logs and the VM creation ultimately fails. Root Cause: The **`listVLAN`** response of F5 BigIP has been modified in 11.x.All VLANs are prefixed with the partition/path information. Solution: The **`listVlan`** response handler has been modified to remove path information.
CS-27530	Problem: Hosts get over-provisioned instead of spreading out evenly choosing the host that has least capacity usage. Solution: Introduced a new allocator - **`firstfitleastconsumed`** - to find hosts or pools with least capacity usage within the cluster and use those resources first. You can use the **`vm.allocation.algorithm`** parameter to use this allocator. For hosts, the allocator decides the least consumed host through the **`host.capacityType.to.order.clusters`** setting to base the usage on CPU or RAM. This applies only for the shared storage and not for the local storage.
CS-27488	Problem: **`ListVirtualMachine`** is slow if VM has many entries in the **`user_vm_details`** table and many resource tags. Root Cause: The **`user_vm_view`** is created by joining the **`user_vm_details`** table, which will duplicate resource tags so many time for each VM, which caused performance slow. Solution: Must not join the **`user_vm_details`** table in **`user_vm_view`** at all since the code never used them. Always find details from the **`user_vm_details`** table directly. Also, handle duplicate resource tag entry by avoiding retrieving duplicate tags.
CS-25948	Problem: Security hole in bash package. Solution: Update the bash package because of the security vulnerability in our system VM templates.
CS-23214	Problem: After a VM is migrated, the volumes attached to the VM have the old disk chain information in the CloudPlatform database. Root Cause: CloudPlatform maintains the chain information of each volume in the database, where chain information is the complete volume path. When the volumes associated with a VM are migrated to a different storage pool, their chain information changes. However, this information is not reflected in the CloudPlatform database. Solution: Upon volume migration, VM migration, and VM migration with storage, update volume's chain_info in the database to reflect the complete path of a volume.
CS-22950	Problem: Can logon with the hash as a password from the database. Root Cause: It is possible to login with the hash from the database as a password. This is due to the fact that the ‘PLAINTEXT’ authenticator is enabled by default.

Issue ID	Description
	Solution: Removed PLAINTEXT from the default authenticator list.
CS-20759	Problem: Random selection of secondary storage by CloudPlatform when there are more than one secondary storage available. Root Cause: If there are more than one secondary storage available, CloudPlatform randomly selects one of them. CloudPlatform does not check whether the secondary storage is full before it selects the secondary storage. Solution: CloudPlatform now checks the size of the secondary storage. If there are more than one secondary storage, CloudPlatform does not select the secondary storage if its used percentage exceeds 90%. If the used percentages of all the available secondary storages exceed 90%, CloudPlatform selects one of them.
CS-20343	Problem: The **`cloudmanagement.keystore`** has lax file permissions. Solution: The fix sets the owner cloud to the file cloud.keystore with read write permission and no other user can access this file for security reasons.
CS-20088	Problem: Performance tests or scale tests show a number of VM deploy failures. Root Cause: The failures are mainly due to NetworkModel code introduced in 4.2.1 that iterates over all the vlans in the pod. This causedeach deployVM thread to hold the global lock on Network longer and cause delays. This delay in turn causes more threads to choose same host and fail since capacity is not available on that host. Solution: In Planner, remove the clusters that do not contain a host with matching service offering tag. This will save some iterations over clusters that do not have matching tagged host. In NetworkModel, do not query the vlans for the pod within the loop. Also optimized the logic to query the ip/ipv6 In DeploymentPlanningManagerImpl, do not process the affinity group if the plan has hostId provided.

Support Matrix

This section describes the operating systems, browsers, and hypervisors that have been newly tested and certified compatible with CloudPlatform 4.5.1 Most earlier OS and hypervisor versions are also supported for use with 4.5.1. For a complete list, see **Chapter 3 System Requirements** of the CloudPlatform 4.5.1 Installation Guide.

2.1. Supported Operating System Versions for Management Server

- RHEL versions 6.3, 6.5, 6.6 and 7
- CentOS versions 6.3, 6.5, and 7

Note

RHEL 7 and CentOS 7 are supported only for experimental use with Linux Containers (LXC) hypervisor. Also

Citrix recommends you to purchase a RHEL support license. Citrix support will not be helpful in resolving issues with the underlying RHEL operating system.

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.

- KVM 6.6
- VMware vCenter versions 5.0 Update 3a
- VMware vCenter versions 5.1 Update 2a
- VMware vCenter ver6sions 5.5 Update up to 2
- VMware vCenter versions 6.0

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
- VMware vCenter versions 6.0
- KVM 6.3

  The libvirt versions supported for KVM on RHEL 6.x are libvirt-0.10.2-41 and beyond.
- XenServer 6.5 SP1
- XenServer version 6.5
- 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 version 11.0
- 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-42189	Problem: After upgrading from CloudPlatform version 3.0.7 to versions 4.2 or later, the VMs that are created before upgrade (that is, in version 3.0.7) fail to migrate storage across clusters.
	
	Root Cause: When a VM is created on XenServer, the VM information is stored in the default VM record. This record stores the information of all VMs.
	
	Also, when a VM is created on a host from a template, it gets a mac_seed value. This value is stored in the default VM record for that host. A unique mac_seed value is generated for each template. So, each VM that is created on that host (or pool) using the same template has the same mac_seed value.
	
	The migration of the first VM to a host in another cluster completes successfully. However, when another migration is performed, it detects that there is already a VM, which has the same mac_seed value and the migration fails.
	
	Workaround: There are two workarounds:
	
	• Stop/Start the VM, then migrate it.
	
	• Do the following:
	
	1. Logon to the host of the VM.
	
	2. List the VMs:
	
	`# xe vm-list`
	
	3. Remove the mac_seed of the VM:
	
	`# xe vm-param-remove uuid=< uuid of the vm > param-name=other-config param-key=mac_seed`
	
	After performing this, the migration works.
	
	Note
	
	Note: The issue is not seen when live storage migration is done within the same pool/cluster.

CS-42172	Problem: Reusing the One post URL succeeds for the deleted volume within the upload timeout. However, the volume does not get deleted from secondary storage when the clean up thread runs.

Issue ID	Description
	Root Cause: In the current design, as soon as the upload operation is completed, the SSVM agent stops monitoring it. If the upload URL has not been expired and if it is used to upload data again, and SSVM accepts this as a new request and the data get uploaded again. Because the volume is already marked as removed in the Management Server, garbage collector thread does not pick it up for cleanup and the entry remains in SSVM. Workaround: There is no workaround for this issue. Users must manually delete the volume from the secondary storage because it cannot be used further as it is already marked as Removed.
CS-42050	Problem: Volume snapshot does not go to the Error state when the snapshot threshold is set to a low value. Root Cause: If volume.snapshot.job.cancel.threshold has a low value, there is a high chance of CloudPlatform returns error in creating volume snapshot while the job is still executing on hypervisor and become successful. This is due to limitation in current job framework. This will not create problems as the volume snapshot is marked in the proper BackedUp state if the volume snapshot is successful on hypervisor. This might be reported as a failure in CloudPlatform. Workaround: No workaround has been specified.
CS-41713	Problem: Adding a new host to DRS cluster does not participate in load balancing. Root Cause: When a new host is added to a cluster, CloudPlatform, by default, doesn't create all the port groups already available with the cluster on the new host. If the cluster is DRS/HA enabled, because the host does not have all the necessary networking port groups (existing VM's port groups), it is not eligible to participate in the cluster DRS load balancing or HA.. Workaround: When the CloudPlatform administrator adds a new host to a DRS/HA enabled cluster, for complete utilization of the functionalities, they must ensure that all the existing port groups of the cluster are available on the new host. For more information, refer to the articles available online that describe methods to copy port groups from one host to another..
CS-41631	Problem: CloudPlatform left behind a pair of stale system VMs in vCenter. Root Cause: The issue was triggered because of CloudPlatform's prolonged loss of connection to vCenter. Because of this problem, both the hosts present in the datacenter were 'Disconnected' and placed in 'Alert' state. Also during this connection loss, because VM sync reported the power state of the VMs as PowerReportMissing, all VMs including both the system VMs were marked as ‘Stopped’ in CloudPlatform. Once the connection to vCenter was restored, both the hosts come up. However, if the hosts that contain the system VMs take too long to connect, the root volume (VM in vCenter) expunge fails. Then, the new system VMs will be launched and stale system VMs will be left behind. Workaround: Identify the system VMs that CloudPlatform does not use currently and delete these system VMs from vCenter.

Issue ID	Description
CS-40905	Problem: NetScaler version 11.0 VPX fails to add to CloudPlatform 4.5.1 if the SSL key generated with 512 bit encryption. Workaround: Generate the SSL key with 1024 key strength as per the KB article available at *http://support.citrix.com/article/CTX135480*.
CS-40536	Problem: Snapshots fail for the root volume, which has hourly recurring snapshot scheduled before the upgrade. Workaround: Delete the stale VDIs entries from XenServer by using the **xe** command or by using XenCenter.
CS-40444	Problem: VM with GPU offering on RHEL 7 platform fails to start. Root Cause: XenCenter fails to provide the console view and relies on SSH and VNC to connect to the VM. There is no problem with the VM and it runs properly. Workaround: Citrix recommends you to first create a template and then, deploy a VM with GPU offering. Make sure that the template automatically gets the private IP address when it runs.
CS-39915	Problem: The polling thread functionality performed by a Management server goes down and the other Management Servers fail to take up this task. Because of this problem, the VMs fail to fetch IP addresses from hypervisors.. Root Cause: In a cluster Management Server setup, when a VM is deployed using a Management Server (for example, MS1) on a specific host, mgmt_server_id in the host table is set to MS1. Because MS1 is fetching VM IP, the other management servers do not perform this action. If MS1 is stopped, the other Management Servers do not perform the VM IP fetch action for that VM.. Workaround: Do one of the following to address this issue: • Stop and start the VM or restart the VM. • Bring the stopped Management Server to the Up state.
CS-39911	Problem: FireFox browser crashes when uploading a file of size one gigabyte (GB) or higher. Workaround: Use another browser, which CloudPlatform supports, to upload files of size one GB and higher.
CS-38773	Problem: When using an external DHCP provider for KVM, IP address allocated by the external DHCP may not be accurately fetched on RHEL 6.3. Root Cause: The **guestfish** library that is used for fetching the assigned IP address caches the data and reports stale information. Workaround: Citrix recommends you to use RHEL 6.5 and above.

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	• 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	• 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	“provisioningtype” option will be added to the following: • Create disk offering dialog • Disk offering details view
Tags for Security Group Rules	• 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/