

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.2 Release Notes

Revised October 2, 2013 3:00 pm Pacific



Citrix CloudPlatform (powered by Apache CloudStack) Version 4.2 Release Notes

Revised October 2, 2013 3:00 pm Pacific

© 2013 Citrix Systems, Inc. All rights reserved. Specifications are subject to change without notice. Citrix Systems, Inc., the Citrix logo, Citrix XenServer, Citrix XenCenter, and CloudPlatform are trademarks or registered trademarks of Citrix Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders.

Release notes for Citrix CloudPlatform version 4.2.

| | |
|-----------------------------------------------------------------------------|-----------|
| 1. Submitting Feedback and Getting Help | 1 |
| 2. Changes to Supported Operating Systems and Hypervisors | 3 |
| 2.1. Hypervisor Versions Supported | 3 |
| 2.2. New Supported OS for Guest Virtual Machines | 3 |
| 3. Name and Location Changes to Services and Scripts | 5 |
| 4. Upgrade Instructions | 7 |
| 4.1. Upgrade from 3.0.x to 4.2 | 7 |
| 4.2. Upgrade from 2.2.x to 4.2 | 16 |
| 4.3. Upgrade from 2.1.x to 4.2 | 25 |
| 4.4. Upgrading and Hotfixing XenServer Hypervisor Hosts | 25 |
| 4.4.1. Upgrading to a New XenServer Version | 25 |
| 4.4.2. Applying Hotfixes to a XenServer Cluster | 27 |
| 5. About This New Release | 31 |
| 6. What's New in 4.2 | 33 |
| 6.1. Features to Support Heterogeneous Workloads | 33 |
| 6.1.1. Regions | 33 |
| 6.1.2. Object Storage Plugin Architecture | 33 |
| 6.1.3. Zone-Wide Primary Storage | 34 |
| 6.1.4. Mapping VMware Datacenter to a CloudPlatform Zone | 34 |
| 6.2. Networking Enhancements | 34 |
| 6.2.1. IPv6 (Technical Preview) | 34 |
| 6.2.2. Portable IPs | 34 |
| 6.2.3. N-Tier Applications | 34 |
| 6.2.4. Assigning VLANs to Isolated Networks | 37 |
| 6.2.5. Persistent Networks | 37 |
| 6.2.6. Cisco VNMC Support | 37 |
| 6.2.7. VMware vNetwork Distributed vSwitch | 38 |
| 6.2.8. IP Reservation in Isolated Guest Networks and VPC | 38 |
| 6.2.9. Dedicated Resources: Public IP Addresses and VLANs Per Account | 38 |
| 6.2.10. Juniper SRX Support for Egress Firewall Rules | 38 |
| 6.2.11. Configuring the Default Egress Policy | 39 |
| 6.2.12. Non-Contiguous VLAN Ranges | 39 |
| 6.2.13. Isolation in Advanced Zone Using Private VLAN | 39 |
| 6.2.14. Configuring Multiple IP Addresses on a Single NIC | 40 |
| 6.2.15. Adding Multiple IP Ranges | 40 |
| 6.2.16. Support for Multiple Networks in VMs | 40 |
| 6.2.17. Global Server Load Balancing | 41 |
| 6.2.18. Load Balancing Using External Provider on Shared VLANs | 41 |
| 6.2.19. Health Checks for Load Balanced Instances | 41 |
| 6.3. Host and Virtual Machine Enhancements | 41 |
| 6.3.1. VMware DRS Support | 42 |
| 6.3.2. Windows 8 and Windows Server 2012 as VM Guest OS | 42 |
| 6.3.3. Change Account Ownership of Virtual Machines | 42 |
| 6.3.4. Dedicated Pod, Cluster, or Host | 42 |
| 6.3.5. Resizing Volumes | 42 |
| 6.3.6. VMware Volume Snapshot Improved Performance | 43 |
| 6.3.7. Live Storage Migration for XenServer and VMware | 43 |
| 6.3.8. Configuring Usage of Linked Clones on VMware | 43 |
| 6.3.9. VM Deployment Rules | 43 |
| 6.3.10. CPU and Memory Scaling for Running VMs | 44 |
| 6.3.11. CPU and Memory Over-Provisioning | 44 |

- 6.3.12. Kickstart Installation for Baremetal Provisioning 44
- 6.3.13. Enhanced Bare Metal Support on Cisco UCS 45
- 6.3.14. Changing a VM's Base Image 45
- 6.3.15. Reset VM on Reboot 45
- 6.3.16. Virtual Machine Snapshots for VMware 45
- 6.3.17. Increased Userdata Size When Deploying a VM 45
- 6.3.18. Set VMware Cluster Size Limit Depending on VMware Version 46
- 6.3.19. Limiting Resource Usage 46
- 6.4. Monitoring, Maintenance, and Operations Enhancements 46
 - 6.4.1. Publish and Subscribe for Event Notification 46
 - 6.4.2. Deleting and Archiving Events and Alerts 47
 - 6.4.3. Increased Granularity for Configuration Parameters 47
 - 6.4.4. API Request Throttling 47
 - 6.4.5. Sending Alerts to External SNMP and Syslog Managers 48
 - 6.4.6. Changing the Default Password Encryption 48
 - 6.4.7. Log Collection Utility cloud-bugtool 48
- 6.5. Third-Party UI Plugin Framework 49
- 7. Fixed Issues 51**
- 8. Known Issues 55**
- 9. API Changes from 3.0 to 4.2 61**
 - 9.1. Added API Commands in 4.2 61
 - 9.1.1. Secondary Storage 61
 - 9.1.2. VM Snapshot 61
 - 9.1.3. Load Balancer Health Check 61
 - 9.1.4. Egress Firewall Rules 61
 - 9.1.5. SSH Key 61
 - 9.1.6. Bare Metal 62
 - 9.1.7. NIC 62
 - 9.1.8. Regions 62
 - 9.1.9. User 62
 - 9.1.10. API Throttling 62
 - 9.1.11. Locking 62
 - 9.1.12. VM Scaling 63
 - 9.1.13. Migrate Volume 63
 - 9.1.14. Dedicated IP and VLAN 63
 - 9.1.15. Port Forwarding 63
 - 9.1.16. Scale System VM 63
 - 9.1.17. Deployment Planner 63
 - 9.1.18. Archive and Delete Events and Alerts 63
 - 9.1.19. Host Reservation 63
 - 9.1.20. Resize Volume 63
 - 9.1.21. Egress Firewall Rules 64
 - 9.1.22. Network ACL 64
 - 9.1.23. Resource Detail 64
 - 9.1.24. Nicira Integration 64
 - 9.1.25. BigSwitch VNS 64
 - 9.1.26. Simulator 64
 - 9.1.27. API Discovery 64
 - 9.1.28. Global Load Balancer 65
 - 9.1.29. Load Balancer 65
 - 9.1.30. Affinity Group 65
 - 9.1.31. Portable IP 65

| | |
|-----------------------------------------|----|
| 9.1.32. Internal Load Balancer VM | 65 |
| 9.1.33. Network Isolation | 66 |
| 9.1.34. Dedicated Resources | 66 |
| 9.2. Changed API Commands in 4.2 | 66 |
| 9.3. Deprecated APIs | 79 |

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/>

Changes to Supported Operating Systems and Hypervisors

This section describes the operating systems and hypervisors that have been newly tested and certified compatible with CloudPlatform 4.2, as well as which previously supported versions have been retired. Most earlier OS and hypervisor versions are also still supported for use with 4.2. For a complete list, see the System Requirements section of the CloudPlatform 4.2 Installation Guide.

2.1. Hypervisor Versions Supported

The following new hypervisor support has been added:

- XenServer 6.2 is supported with fresh CloudPlatform 4.2 installations. Deployments that are upgrading from a previous version of CloudPlatform can not support XenServer 6.2.
- KVM on RHEL 6.3
- Bare metal hosts are supported, which have no hypervisor. These hosts can run the following operating systems:
 - RHEL or CentOS, v6.2 or 6.3



Note

Use libvirt version 0.9.10 for CentOS 6.3

- Fedora 17
- Ubuntu 12.04

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

2.2. New Supported OS for Guest Virtual Machines

Windows 8 and Windows Server 2012 are now supported for guest virtual machines on KVM, XenServer, and VMware. See [Section 6.3.2, “Windows 8 and Windows Server 2012 as VM Guest OS”](#).

All previously supported guest VM operating systems are still supported.

Name and Location Changes to Services and Scripts

Service and script names and directory paths have changed in CloudPlatform 4.2. If you have written scripts based on the old commands, take a look and update them where appropriate.

| Old | New | Comments |
|--------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| cloud-management | cloudstack-management | Management Server service name |
| cloud-usage | cloudstack-usage | Usage Server service name |
| cloud-setup-databases | cloudstack-setup-databases | Script used during installation |
| /etc/cloud/management/ | /etc/cloudstack/management/ | Management Server configuration and properties files |
| /usr/lib64/cloud/agent/scripts/storage/secondary/cloud-install-sys-tmplt | /usr/share/cloudstack-common/scripts/storage/secondary/cloud-install-sys-tmplt | Script for seeding system VM template |
| /var/log/cloud/management/management-server.log | /var/log/cloudstack/management/management-server.log | Management Server log |
| /usr/lib64/cloud/common/scripts/vm/hypervisor/xenserver/ /usr/lib64/cloud/common/scripts/vm/hypervisor/xenserver/xenserver60/ | /usr/share/cloudstack-common/scripts/vm/hypervisor/xenserver /usr/share/cloudstack-common/scripts/vm/hypervisor/xenserver/xenserver60 | (XenServer) Management Server files to copy during XenServer host upgrade |
| /usr/lib64/cloud/common/scripts/vm/hypervisor/xenserver/cloud-setup-bonding.sh | /usr/share/cloudstack-common/scripts/vm/hypervisor/xenserver/cloud-setup-bonding.sh | (XenServer) Script used when upgrading XenServer hosts |
| cloud-agent | cloudstack-agent | (KVM) CloudPlatform agent service name |
| /etc/cloud/agent/agent.properties | /etc/cloudstack/agent/agent.properties | (KVM) Properties file used on KVM hosts during network configuration |
| /var/log/cloud/agent | /var/log/cloudstack/agent/ | (KVM) CloudPlatform KVM agent log file |
| /etc/cloud/agent | /etc/cloudstack/agent | (KVM) CloudPlatform KVM agent configuration and properties files |

Upgrade Instructions

4.1. Upgrade from 3.0.x to 4.2

Perform the following to upgrade from version 3.0.0, 3.0.1, 3.0.2, 3.0.3, 3.0.4, 3.0.5, 3.0.6, or 3.0.7 to version 4.2.

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

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

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



Note

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

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

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

| Hypervisor | Description |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| XenServer | <p>Name: systemvm-xenserver-4.2</p> <p>Description: systemvm-xenserver-4.2</p> <p>URL: http://download.cloud.com/templates/4.2/systemvmtemplate-2013-07-12-master-xen.vhd.bz2</p> <p>Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the XenServer zones.</p> <p>Hypervisor: XenServer</p> <p>Format: VHD</p> |

| Hypervisor | Description |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <p>OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p> |
| KVM | <p>Name: systemvm-kvm-4.2</p> <p>Description: systemvm-kvm-4.2</p> <p>URL: http://download.cloud.com/templates/4.2/systemvmtemplate-2013-06-12-master-kvm.qcow2.bz2</p> <p>Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, choose All Zones to make the template available in all the KVM zones.</p> <p>Hypervisor: KVM</p> <p>Format: QCOW2</p> <p>OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p> |
| VMware | <p>Name: systemvm-vmware-4.2</p> <p>Description: systemvm-vmware-4.2</p> <p>URL: http://download.cloud.com/templates/4.2/systemvmtemplate-4.2-vh7.ova</p> <p>Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, choose All Zones to make the template available in all the VMware zones.</p> <p>Hypervisor: VMware</p> <p>Format: OVA</p> <p>OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> |

| Hypervisor | Description |
|------------|----------------------|
| | Password Enabled: no |
| | Public: no |
| | Featured: no |

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



Warning

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

4. (KVM on RHEL 6.0/6.1 only) If your existing CloudPlatform deployment includes one or more clusters of KVM hosts running RHEL 6.0 or RHEL 6.1, you must first upgrade the operating system version on those hosts before upgrading CloudPlatform itself.

Run the following commands on every KVM host.

- a. Download the CloudPlatform 4.2.0 RHEL 6.3 binaries from <https://www.citrix.com/English/ss/downloads/>.
- b. Extract the binaries:

```
# cd /root
# tar xvf CloudPlatform-4.2.0-1-rhel6.3.tar.gz
```

- c. Create a CloudPlatform 4.2 qemu repo:

```
# cd CloudPlatform-4.2.0-1-rhel6.3/6.3
# createrepo .
```

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

```
[upgrade]
name=rhel63
baseurl=url-of-your-rhel6.3-repo
enabled=1
gpgcheck=0
[cloudstack]
name=cloudstack
baseurl=file:///root/CloudPlatform-4.2.0-1-rhel6.3/6.3
enabled=1
gpgcheck=0
```

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

```
yum upgrade
```

Chapter 4. Upgrade Instructions

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

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

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

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

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

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

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

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

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

9. Download CloudPlatform 4.2 onto the management server host where it will run. Get the software from the following link:

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

You need a [My Citrix Account](#)¹.

10. Upgrade the CloudPlatform packages. You should have a file in the form of "CloudPlatform-4.2-N-OSVERSION.tar.gz". Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-4.2-N-OSVERSION.tar.gz
# cd CloudPlatform-4.2-N-OSVERSION
# ./install.sh
```

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

11. Choose "U" to upgrade the package

```
>U
```

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

12. If you have made changes to your existing copy of the configuration files components.xml, db.properties, or server.xml in your previous-version CloudPlatform installation, the changes will be preserved in the upgrade. However, you need to do the following steps to place these changes in a new version of the file which is compatible with version 4.2.

¹ <http://www.citrix.com/lang/English/publicindex.asp?destURL=%2FEnglish%2FmyCitrix%2Findex.asp%3F#>

**Note**

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

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

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

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

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

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

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

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

13. Repeat steps 8 - 12 on each management server node.

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

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

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

**Note**

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

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

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



Note

After upgrade from 3.0.4 to 4.2, if the usage server fails to restart then copy db.properties from /etc/cloudstack/management to /etc/cloudstack/usage. Then start the Usage Server.

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

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

- a. Stop the Management Server:

```
service cloudstack-management stop
```

- b. Perform the following on each VMware cluster:

- i. Encrypt the vCenter password:

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

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

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

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

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

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

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

- iv. Confirm that the table is updated:

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

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

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

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

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

- vii. Confirm that the table is updated:

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

- c. Start the CloudPlatform Management server

```
service cloudstack-management start
```

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



Note

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

- Copy the CloudPlatform 4.2 .tgz download to the host, untar it, and cd into the resulting directory.
- Stop the running agent.

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

- Update the agent software.

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

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

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

- Install a libvirt hook with the following commands:

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

- Restart libvirtd.

```
# service libvirtd restart
```

- i. Start the agent.

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

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



Note

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

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

19. If you are upgrading from 3.0.1 or 3.0.2, perform the following:
 - a. Ensure that the admin port is set to 8096 by using the "integration.api.port" global parameter.

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

- b. Restart the Management Server.



Note

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

20. Run the following script to stop, then start, all System VMs including Secondary Storage VMs, Console Proxy VMs, and virtual routers.
 - a. Run the script once on one management server. Substitute your own IP address of the MySQL instance, the MySQL user to connect as, and the password to use for that user. In addition to those parameters, provide the "-a" argument. For example:

```
# nohup cloudstack-sysvmadm -d 192.168.1.5 -u cloud -p password -a > sysvm.log 2>&1 &
```

This might take up to an hour or more to run, depending on the number of accounts in the system.

- b. After the script terminates, check the log to verify correct execution:

```
# tail -f sysvm.log
```

The content should be like the following:

```
Stopping and starting 1 secondary storage vm(s)...
Done stopping and starting secondary storage vm(s)
Stopping and starting 1 console proxy vm(s)...
Done stopping and starting console proxy vm(s).
Stopping and starting 4 running routing vm(s)...
Done restarting router(s).
```

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

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

XenServer or KVM:

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

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

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

The output should be like the following:

```
Cloudstack Release 4.2 Mon Aug 12 15:10:04 PST 2013
```

ESXi

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

Run the following commands on the Management Server:

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

The output should be like the following:

```
Cloudstack Release 4.2 Mon Sep 24 15:10:04 PST 2012
```

21. If you want to close the admin port again (recommended in production systems), set `integration.api.port` to null. Then restart the Management Server. For information about how to set `integration.api.port`, see “Setting Configuration Parameters” in the Installation Guide.
22. (XenServer only) If needed, upgrade all Citrix XenServer hypervisor hosts in your cloud to a version supported by CloudPlatform 4.2 and apply any required hotfixes. Instructions for upgrading XenServer software and applying hotfixes can be found in [Section 4.4, “Upgrading and Hotfixing XenServer Hypervisor Hosts”](#).

23. (VMware only) After upgrade, if you want to change a Standard vSwitch zone to a VMware dvSwitch Zone, perform the following:
- Ensure that the Public and Guest traffics are not on the same network as the Management and Storage traffic.
 - Set `vmware.use.dvswitch` to `true`.
 - Access the physical network for the Public and guest traffic, then change the traffic labels as given below:

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

For example: `dvSwitch18,,vmwaredvs`

VLANID is optional.

- Stop the Management server.
 - Start the Management server.
 - Add the new VMware dvSwitch-enabled cluster to this zone.
24. (VMware only) If your existing cloud includes any deployed data centers, you should set the global configuration setting `vmware.create.full.clone` to `false`. Then restart the Management Server. For information about how to set `vmware.create.full.clone`, see “Setting Configuration Parameters” in the Installation Guide. For information about how CloudPlatform supports full and linked clones, see “Configuring Usage of Linked Clones on VMware” in the CloudPlatform Administration Guide.



Note

Troubleshooting tip: If passwords which you know to be valid appear not to work after upgrade, or other UI issues are seen, try clearing your browser cache and reloading the UI page.



Note

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

4.2. Upgrade from 2.2.x to 4.2

- Ensure that you query your IP address usage records and process them; for example, issue invoices for any usage that you have not yet billed users for.

Starting in 3.0.2, the usage record format for IP addresses is the same as the rest of the usage types. Instead of a single record with the assignment and release dates, separate records are generated per aggregation period with start and end dates. After upgrading to 4.2, any existing IP address usage records in the old format will no longer be available.

2. If you are using version 2.2.0 - 2.2.13, first upgrade to 2.2.14 by using the instructions in the [2.2.14 Release Notes](#)².



Note

(KVM only) If KVM hypervisor is used in your cloud, be sure you completed the step to insert a valid username and password into the host_details table on each KVM node as described in the 2.2.14 Release Notes. This step is critical, as the database will be encrypted after the upgrade to 4.2.

3. While running the 2.2.x system (which by this step should be at version 2.2.14 or greater), log in to the UI as root administrator.
4. Using the UI, add a new System VM template for each hypervisor type that is used in your cloud. In each zone, add a system VM template for each hypervisor used in that zone.



Note

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

- a. In the left navigation bar, click Templates.
- b. In Select view, click Templates.
- c. Click Register template.
The Register template dialog box is displayed.
- d. In the Register template dialog box, specify the following values depending on the hypervisor type (do not change these):

| Hypervisor | Description |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| XenServer | Name: systemvm-xenserver-4.2 Description: systemvm-xenserver-4.2 URL: http://download.cloud.com/templates/4.2/systemvmtemplate-2013-07-12-master-xen.vhd.bz2 |

² <http://download.cloud.com/releases/2.2.0/CloudStack2.2.14ReleaseNotes.pdf>

| Hypervisor | Description |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <p>Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the XenServer zones.</p> <p>Hypervisor: XenServer</p> <p>Format: VHD</p> <p>OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p> |
| KVM | <p>Name: systemvm-kvm-4.2</p> <p>Description: systemvm-kvm-4.2</p> <p>URL: http://download.cloud.com/templates/4.2/systemvmtemplate-2013-06-12-master-kvm.qcow2.bz2</p> <p>Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, choose All Zones to make the template available in all the KVM zones.</p> <p>Hypervisor: KVM</p> <p>Format: QCOW2</p> <p>OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown)</p> <p>Extractable: no</p> <p>Password Enabled: no</p> <p>Public: no</p> <p>Featured: no</p> |
| VMware | <p>Name: systemvm-vmware-4.2</p> <p>Description: systemvm-vmware-4.2</p> <p>URL: http://download.cloud.com/templates/4.2/systemvmtemplate-4.2-vh7.ova</p> <p>Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running</p> |

| Hypervisor | Description |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | VMware, choose All Zones to make the template available in all the VMware zones. Hypervisor: VMware Format: OVA OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown) Extractable: no Password Enabled: no Public: no Featured: no |

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



Warning

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

5. (KVM on RHEL 6.0, 6.1) If your existing CloudPlatform deployment includes one or more clusters of KVM hosts running RHEL 6.0 or RHEL 6.1, you must first upgrade the operating system version on those hosts before upgrading CloudPlatform itself.

Run the following commands on every KVM host.

- a. Download the CloudPlatform 4.2.0 RHEL 6.3 binaries from <https://www.citrix.com/English/ss/downloads/>.
- b. Extract the binaries:

```
# cd /root
# tar xvf CloudPlatform-4.2.0-1-rhel6.3.tar.gz
```

- c. Create a CloudPlatform 4.2 qemu repo:

```
# cd CloudPlatform-4.2.0-1-rhel6.3/6.3
# createrepo .
```

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

```
[upgrade]
name=rhel63
baseurl=url-of-your-rhel6.3-repo
```

```
enabled=1
gpgcheck=0
[cloudstack]
name=cloudstack
baseurl=file:///root/CloudPlatform-4.2.0-1-rhel6.3/6.3
enabled=1
gpgcheck=0
```

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

```
yum upgrade
```

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

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

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

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

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

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

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

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

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

10. Download CloudPlatform 4.2 onto the management server host where it will run. Get the software from the following link:

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

You need a [My Citrix Account](#)³.

11. Upgrade the CloudPlatform packages. You should have a file in the form of “CloudPlatform-4.2-N-OSVERSION.tar.gz”. Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
# tar xzf CloudPlatform-4.2-N-OSVERSION.tar.gz
# cd CloudPlatform-4.2-N-OSVERSION
# ./install.sh
```

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

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

12. Choose "U" to upgrade the package.

```
> U
```

13. If you have made changes to your existing copy of the configuration files `components.xml`, `db.properties`, or `server.xml` in your previous-version CloudPlatform installation, the changes will be preserved in the upgrade. However, you need to do the following steps to place these changes in a new version of the file which is compatible with version 4.2.



Note

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

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

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

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

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

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

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

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

14. On the management server node, run the following command. It is recommended that you use the command-line flags to provide your own encryption keys. See [Password and Key Encryption in the Installation Guide](#).

```
# cloudstack-setup-encryption -e <encryption_type> -m <management_server_key> -k
<database_key>
```

When used without arguments, as in the following example, the default encryption type and keys will be used:

- (Optional) For `encryption_type`, use `file` or `web` to indicate the technique used to pass in the database encryption password. Default: `file`.
- (Optional) For `management_server_key`, substitute the default key that is used to encrypt confidential parameters in the properties file. Default: `password`. It is highly recommended that you replace this with a more secure value

Chapter 4. Upgrade Instructions

- (Optional) For `database_key`, substitute the default key that is used to encrypt confidential parameters in the CloudPlatform database. Default: `password`. It is highly recommended that you replace this with a more secure value.

15. Repeat steps 9 - 14 on every management server node. If you provided your own encryption key in step 14, use the same key on all other management servers.
16. Start the first Management Server. Do not start any other Management Server nodes yet.

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

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

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

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

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



Note

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

- a. Copy the CloudPlatform 4.2 .tgz download to the host, untar it, and `cd` into the resulting directory.
- b. Stop the running agent.

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

- c. Update the agent software.

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

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

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

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

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

- h. Restart libvirtd.

```
# service libvirtd restart
```

- i. Start the agent.

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

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

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

20. Run the following script to stop, then start, all System VMs including Secondary Storage VMs, Console Proxy VMs, and virtual routers.

- a. Run the command once on one management server. Substitute your own IP address of the MySQL instance, the MySQL user to connect as, and the password to use for that user. In addition to those parameters, provide the "-c" and "-r" arguments. For example:

```
# nohup cloudstack-sysvmadm -d 192.168.1.5 -u cloud -p password -a > sysvm.log 2>&1 &
```

This might take up to an hour or more to run, depending on the number of accounts in the system.

- b. After the script terminates, check the log to verify correct execution:

```
# tail -f sysvm.log
```

The content should be like the following:

```
Stopping and starting 1 secondary storage vm(s)...
Done stopping and starting secondary storage vm(s)
Stopping and starting 1 console proxy vm(s)...
Done stopping and starting console proxy vm(s).
Stopping and starting 4 running routing vm(s)...
Done restarting router(s).
```

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

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

XenServer or KVM:

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

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

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

The output should be like the following:

```
Cloudstack Release 4.2 Mon Aug 12 15:10:04 PST 2013
```

ESXi

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

Run the following commands on the Management Server:

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

The output should be like the following:

```
Cloudstack Release 4.2 Mon Aug 12 15:10:04 PST 2012
```

21. (XenServer only) If needed, upgrade all Citrix XenServer hypervisor hosts in your cloud to a version supported by CloudPlatform 4.2 and apply any required hotfixes. Instructions for upgrading and applying hotfixes can be found in [Section 4.4, “Upgrading and Hotfixing XenServer Hypervisor Hosts”](#).
22. (VMware only) If your existing cloud includes any deployed data centers, you should set the global configuration setting `vmware.create.full.clone` to `false`. Then restart the Management Server. For information about how to set `vmware.create.full.clone`, see “Setting Configuration Parameters” in the Installation Guide. For information about how CloudPlatform supports full and linked clones, see “Configuring Usage of Linked Clones on VMware” in the CloudPlatform Administration Guide.

**Note**

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

4.3. Upgrade from 2.1.x to 4.2

Direct upgrades from version 2.1.0 - 2.1.10 to 4.2 are not supported. CloudPlatform must first be upgraded to version 2.2.14. For information on how to upgrade from 2.1.x to 2.2.14, see the CloudPlatform 2.2.14 Release Notes.

4.4. Upgrading and Hotfixing XenServer Hypervisor Hosts

In CloudPlatform 4.2, you can upgrade XenServer hypervisor host software without having to disconnect the XenServer cluster. You can upgrade XenServer 5.6 GA, 5.6 FP1, or 5.6 SP2 to any newer version that is supported by CloudPlatform. The actual upgrade is described in XenServer documentation, but there are some additional steps you must perform before and after the upgrade.

4.4.1. Upgrading to a New XenServer Version

To upgrade XenServer hosts when running CloudPlatform 4.2:

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

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

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

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

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

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

```
# xe pool-list
```

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

```
# xe host-list
```

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



Note

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

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

5. Disconnect the XenServer cluster from CloudPlatform. It will remain disconnected only long enough to upgrade one host.
 - a. Log in to the CloudPlatform UI as root.
 - b. Navigate to the XenServer cluster, and click Actions – Unmanage.
 - c. Watch the cluster status until it shows Unmanaged.
6. Upgrade the XenServer software on the master host:
 - a. Insert the XenServer CD.
 - b. Reboot the host.
 - c. Upgrade to the newer version of XenServer. Use the steps in XenServer documentation.
7. Cancel the maintenance mode on the master host.
8. Reconnect the XenServer cluster to CloudPlatform.
 - a. Log in to the CloudPlatform UI as root.
 - b. Navigate to the XenServer cluster, and click Actions – Manage.
 - c. Watch the status to see that all the hosts come up.
9. Upgrade the slave hosts in the cluster:
 - a. Put a slave host into maintenance mode.

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

- If you upgraded from XenServer 5.6 SP2 to XenServer 6.0.2 or higher, change any VMs that have the OS type CentOS 5.6 (32-bit), CentOS 5.7 (32-bit), Oracle Enterprise Linux 5.6 (32-bit), Oracle Enterprise Linux 5.7 (32-bit), Red Hat Enterprise Linux 5.6 (32-bit), or Red Hat Enterprise Linux 5.7 (32-bit) to Other Linux (32-bit). Change any VMs that have the 64-bit versions of these same OS types to Other Linux (64-bit).
- If you upgraded from XenServer 5.6 to XenServer 6.0.2 or higher, do all of the above.

4.4.2. Applying Hotfixes to a XenServer Cluster

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

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

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

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

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

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

```
# xe pool-list
```

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

```
# xe host-list
```

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

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

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

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

6. Hotfix the master host:

- a. Add the XenServer hot fixes to the master host.
 - i. Assign a UUID to the update file:

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

The command displays the UUID of the update file:

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

- ii. Repeat the `xe patch-upload` command for all other XenServer updates: `XS602E004.xsupdate`, `XS602E005.xsupdate`.

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

- b. Apply XenServer hot fixes to master host:

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

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

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

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

9. Hotfix the slave hosts in the cluster:

- a. Put a slave host into maintenance mode.

Wait until all the VMs are migrated to other hosts.

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

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

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

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

- e. Restart the slave hosts.

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

10. Cancel the maintenance mode on the slave hosts.

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

About This New Release

CloudPlatform 4.2 release focuses on the following three major areas:

- Improved support for both traditional and cloud-era workloads
- New third-party UI plug-in architecture
- Networking enhancements

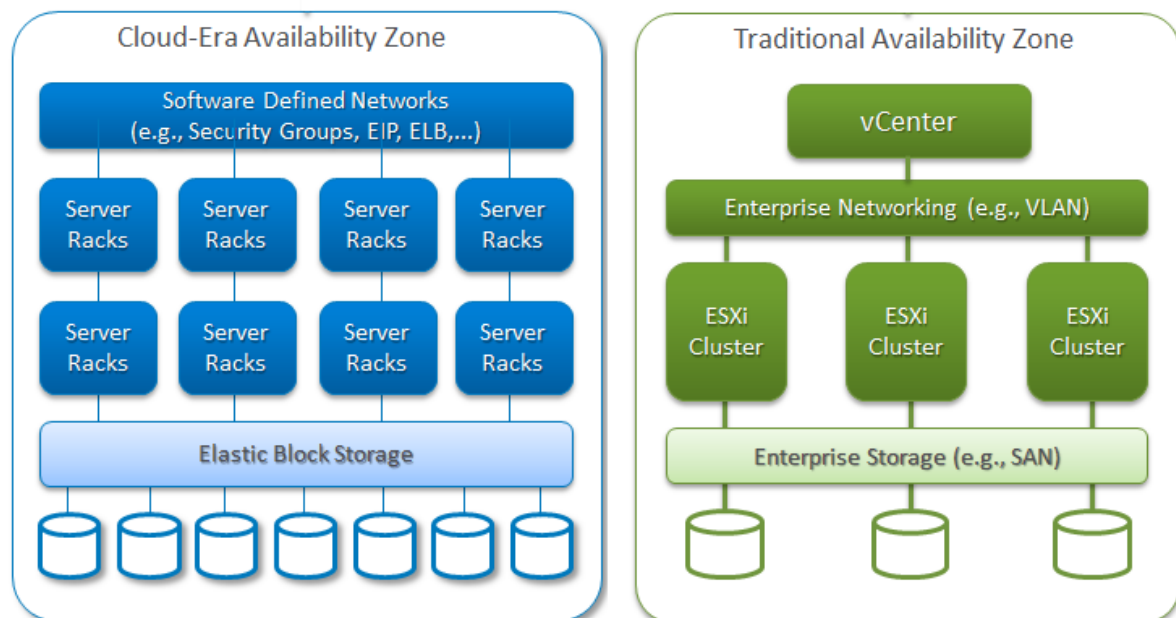
In addition to these major new areas of functionality, CloudPlatform 4.2 provides many additional enhancements to enable more powerful ongoing management and use of CloudPlatform:

- Host and virtual machine enhancements
- Monitoring, maintenance, and operations enhancements

To support these areas, CloudPlatform 4.2 includes more than 70 new features and enhancements. All of the new features are summarized later in this Release Note. First, here is a summary of the release as a whole.

Heterogeneous Cloud Workloads

CloudPlatform now more fully supports two different styles of workload, and you can combine both in a single cloud. In pragmatic terms, this means you can have two different styles of zone. One is a zone that follows a server virtualization model; often, this means simply putting an existing setup under CloudPlatform management. The second is a zone which is built “from the cloud up” to take advantage of cloud architecture. The hardware and software are set up differently inside each type of zone.



Third-Party UI Plugins

Customers and partners can now add capabilities to CloudPlatform. The new 3rd-party plugin architecture means that anyone can add features which will appear right alongside in-house features. The code for the plugin is simply placed in a special directory within CloudPlatform’s installed code at any time after CloudPlatform installation. The new plugin appears only when it is enabled by the cloud administrator. See [Section 6.5, “Third-Party UI Plugin Framework”](#).



Networking Enhancements

Version 4.2 adds several new networking features as part of the ongoing forward development of CloudPlatform.

- nTier VPC applications
- Persistent networks without running VMs
- Global Server Load Balancing support
- IP addresses and networks dedicated to accounts
- IP address and IP range enhancements
- VMware vSphere Distributed Switch (VDS) support
- More types of network elements you can include in your cloud: Cisco ASA 1000v Firewall and more.

See [Section 6.2, “Networking Enhancements”](#).

Additional Enhancements

In addition to these three major new areas of functionality, CloudPlatform 4.2 includes various enhancements and new features. Some of the categories in which you will find new capabilities:

- New ways to work with VMs and hosts. See [Section 6.3, “Host and Virtual Machine Enhancements”](#).
- New capabilities for monitoring, maintenance, and general ops. See [Section 6.4, “Monitoring, Maintenance, and Operations Enhancements”](#).

What's New in 4.2

CloudPlatform 4.2 includes the following new features.

6.1. Features to Support Heterogeneous Workloads

The following new features help CloudPlatform 4.2 better support both legacy and cloud-era style zones.

6.1.1. Regions

To increase reliability of the cloud, you can optionally group resources into geographic regions. A region is the largest available infrastructure unit within a cloud deployment. A region is made up of several availability zones, where each zone is equivalent to a datacenter. Each region is controlled by its own cluster of Management Servers, running in one of the zones. The zones in a region are typically located in close geographical proximity. Regions are a useful technique for providing fault tolerance and disaster recovery.

By grouping zones into regions, the cloud can achieve higher availability and scalability. User accounts can span regions, so that users can deploy VMs in multiple, widely-dispersed regions. Even if one of the regions becomes unavailable, the services are still available to the end-user through VMs deployed in another region. And by grouping communities of zones under their own nearby Management Servers, the latency of communications within the cloud is reduced compared to managing widely-dispersed zones from a single central Management Server.

Usage records can also be consolidated and tracked at the region level, creating reports or invoices for each geographic region.

6.1.2. Object Storage Plugin Architecture

Artifacts such as templates, ISOs and snapshots are kept in storage which CloudPlatform refers to as secondary storage. To improve scalability and performance, as when a number of hosts access secondary storage concurrently, object storage can be used for secondary storage. Object storage can also provide built-in high availability capability. When using object storage, access to secondary storage data can be made available across multiple zones in a region. This is a huge benefit, as it is no longer necessary to copy templates, snapshots etc. across zones as would be needed in an NFS-only environment.

Object storage is provided through third-party software such as Amazon Simple Storage Service (S3) or any other object storage that supports the S3 interface. These third party object storages can be integrated with CloudPlatform by writing plugin software that uses the object storage plugin capability introduced in CloudPlatform 4.2. Several new pluggable service interfaces are available so that different storage providers can develop vendor-specific plugins based on the well-defined contracts that can be seamlessly managed by CloudPlatform.



Note

CloudPlatform will only support either NFS-based secondary storage or Object Storage and not both concurrently.

6.1.3. Zone-Wide Primary Storage

(Supported on KVM and VMware)

In CloudPlatform 4.2, you can provision primary storage on a per-zone basis. Data volumes in the primary storage can be attached to any VM on any host in the zone.

In previous CloudPlatform versions, each cluster had its own primary storage. Data in the primary storage was directly available only to VMs within that cluster. If a VM in a different cluster needed some of the data, it must be copied from one cluster to another, using the zone's secondary storage as an intermediate step. This operation was unnecessarily time-consuming.

6.1.4. Mapping VMware Datacenter to a CloudPlatform Zone

In order to support zone-wide functions for VMware, changes have been made so that CloudPlatform is now aware of VMware Datacenters and can map each Datacenter to a CloudPlatform zone. Previously, CloudPlatform was only aware of VMware Clusters, a smaller organizational unit than Datacenters. This implies that a single CloudPlatform zone could possibly contain clusters from different VMware Datacenters. In order for zone-wide functions, such as zone-wide primary storage, to work for VMware hosts, CloudPlatform has to make sure that a zone contains only a single VMware Datacenter. Therefore, when you are creating a new CloudPlatform zone, you will now be able to select a VMware Datacenter for the zone. If you are provisioning multiple VMware Datacenters, each one will be set up as a single zone in CloudPlatform.



Note

If you are upgrading from a previous CloudPlatform version, and your existing deployment contains a zone with clusters from multiple VMware Datacenters, that zone will not be forcibly migrated to the new model. It will continue to function as before. However, any new zone-wide operations, such as zone-wide primary storage, will not be available in that zone.

6.2. Networking Enhancements

The following new features provide additional networking functionality in CloudPlatform 4.2.

6.2.1. IPv6 (Technical Preview)

CloudPlatform 4.2 introduces initial support for IPv6. This feature is provided as a technical preview only. Full support is planned for a future release.

6.2.2. Portable IPs

Portable IPs in CloudPlatform are elastic IPs that can be transferred across geographically separated zones. As an administrator, you can provision a pool of portable IPs at region level and are available for user consumption. The users can acquire portable IPs if admin has provisioned portable public IPs at the region level they are part of. These IPs can be used for any service within an advanced zone. You can also use portable IPs for EIP service in Basic zones. Additionally, a portable IP can be transferred from one zone to another.

6.2.3. N-Tier Applications

In the earlier versions of CloudPlatform, a functionality was added to allow users to create a multi-tier application connected to a single instance of a Virtual Router that supports inter-VLAN routing. Such a

multi-tier application is called a virtual private cloud (VPC). Users were also able to connect their multi-tier applications to a private Gateway or a Site-to-Site VPN tunnel and route certain traffic to those gateways. For CloudPlatform 4.2, additional features are implemented to enhance VPC applications.

- [Section 6.2.3.1, “Support for KVM”](#)
- [Section 6.2.3.2, “Load Balancing Support for VPC”](#)
- [Section 6.2.3.3, “NetScaler Support for VPC”](#)
- [Section 6.2.3.4, “Enhanced Access Control List”](#)
- [Section 6.2.3.4.1, “ACL on Private Gateway”](#)
- [Section 6.2.3.4.2, “Allow ACL on All Layer 4 Protocols”](#)
- [Section 6.2.3.4.3, “Support for ACL Deny Rules”](#)
- [Section 6.2.3.5, “Deploying VMs to a VPC Tier and Shared Networks”](#)
- [Section 6.2.3.6, “Adding a Private Gateway to a VPC”](#)
- [Section 6.2.3.6.1, “Source NAT on Private Gateway”](#)
- [Section 6.2.3.6.2, “VPN Gateways”](#)
- [Section 6.2.3.6.3, “Creating a Static Route”](#)
- [Section 6.2.3.6.4, “Blacklisting Routes”](#)

6.2.3.1. Support for KVM

VPC is now supported on KVM hypervisors.

6.2.3.2. Load Balancing Support for VPC

In a VPC, you can configure two types of load balancing—external LB and internal LB. External LB is nothing but a LB rule created to redirect the traffic received at a public IP of the VPC virtual router. The traffic is load balanced within a tier based on your configuration. Citrix NetScaler and VPC virtual router are supported for external LB. When you use internal LB service, traffic received at a tier is load balanced across different VMs within that tier. For example, traffic reached at Web tier is redirected to another VM in that tier. External load balancing devices are not supported for internal LB. The service is provided by a internal LB VM configured on the target tier.

Load Balancing Public Traffic (Public LB)

A CloudPlatform user or administrator may create load balancing rules that balance traffic received at a public IP to one or more VMs that belong to a network tier that provides load balancing service in a VPC. A user creates a rule, specifies an algorithm, and assigns the rule to a set of VMs within a tier.

Load Balancing Tier-to-Tier Traffic (Internal LB)

CloudPlatform supports sharing workload across different tiers within your VPC. Assume that multiple tiers are set up in your environment, such as Web tier and Application tier. Traffic to each tier is balanced on the VPC virtual router on the public side. If you want the traffic coming from the Web tier to the Application tier to be balanced, use the internal load balancing feature offered by CloudPlatform.

6.2.3.3. NetScaler Support for VPC

Citrix NetScaler is supported for external LB. Certified version for this feature is NetScaler 10.0 Build 74.4006.e.

6.2.3.4. Enhanced Access Control List

Network Access Control List (ACL) on the VPC virtual router is enhanced. The network ACLs can be created for the tiers only if the NetworkACL service is supported. In CloudPlatform terminology, Network ACL is a group of Network ACL items. Network ACL items are nothing but numbered rules that are evaluated in order, starting with the lowest numbered rule. These rules determine whether traffic is allowed in or out of any tier associated with the network ACL. You need to add the Network ACL items to the Network ACL, then associate the Network ACL with a tier. Network ACL is associated with a VPC and can be assigned to multiple VPC tiers within a VPC. A Tier is associated with a Network ACL at all the times. Each tier can be associated with only one ACL.

The default Network ACL is used when no ACL is associated. Default behavior is all incoming traffic to guest networks is blocked and all outgoing traffic from guest networks is allowed. Default network ACL cannot be removed or modified.

6.2.3.4.1. ACL on Private Gateway

The traffic on the VPC private gateway is controlled by creating both ingress and egress network ACL rules. The ACLs contains both allow and deny rules. As per the rule, all the ingress traffic to the private gateway interface and all the egress traffic out from the private gateway interface are blocked. You can change this default behaviour while creating a private gateway.

6.2.3.4.2. Allow ACL on All Layer 4 Protocols

In addition to the existing protocol support for ICMP, TCP, UDP, support for All Layer 4 protocols is added. The protocol numbers from 0 to 255 are supported.

6.2.3.4.3. Support for ACL Deny Rules

In addition to the existing support for ACL Allow rules, support for ACL Deny rules has been added in CloudPlatform 4.2. As part of this, two operations are supported: Number and Action. You can configure a rule, allow or deny, by using action. Use Number to add a rule number.

6.2.3.5. Deploying VMs to a VPC Tier and Shared Networks

CloudPlatform allows you to deploy VMs on a VPC tier and one or more shared networks. With this feature, the VMs deployed in a multi-tier application can receive services offered by a service provider over the shared network. One example of such a service is monitoring service.

6.2.3.6. Adding a Private Gateway to a VPC

A private gateway can be added by the root admin only. The VPC private network has 1:1 relationship with the NIC of the physical network. You can configure multiple private gateways to a single VPC. No gateways with duplicated VLAN and IP are allowed in the same data center.

6.2.3.6.1. Source NAT on Private Gateway

You might want to deploy multiple VPCs with the same super CIDR and guest tier CIDR. Therefore, multiple guest VMs from different VPCs can have the same IPs to reach a enterprise data center through the private gateway. In such cases, a NAT service need to be configured on the private

gateway. If Source NAT is enabled, the guest VMs in VPC reaches the enterprise network via private gateway IP address by using the NAT service.

The Source NAT service on a private gateway can be enabled while adding the private gateway. On deletion of a private gateway, source NAT rules specific to the private gateway are deleted.

6.2.3.6.2. VPN Gateways

Support up to 8 VPN Gateways is added.

6.2.3.6.3. Creating a Static Route

CloudPlatform enables you to specify routing for the VPN connection you create. You can enter one or CIDR addresses to indicate which traffic is to be routed back to the gateway.

6.2.3.6.4. Blacklisting Routes

CloudPlatform enables you to block a list of routes so that they are not assigned to any of the VPC private gateways. Specify the list of routes that you want to blacklist in the `blacklisted.routes` global parameter. Note that the parameter update affects only new static route creations. If you block an existing static route, it remains intact and continue functioning. You cannot add a static route if the route is blacklisted for the zone.

6.2.4. Assigning VLANs to Isolated Networks

CloudPlatform provides you the ability to control VLAN assignment to Isolated networks. You can assign a VLAN ID when a network is created, just the way it's done for Shared networks. The former behaviour also is supported — VLAN is randomly allocated to a network from the VNET range of the physical network when the network turns to Implemented state. The VLAN is released back to the VNET pool when the network shuts down as a part of the Network Garbage Collection. The VLAN can be re-used either by the same network when it is implemented again, or by any other network. On each subsequent implementation of a network, a new VLAN can be assigned.

You cannot change a VLAN once it's assigned to the network. The VLAN remains with the network for its entire life cycle.

6.2.5. Persistent Networks

CloudPlatform 4.2 supports Persistent Networks. The network that you can provision without having to deploy any VMs on it is called a Persistent Network. A Persistent Network can be part of a VPC or a non-VPC environment. With the addition of this feature, you will have the ability to create a network in CloudPlatform in which physical devices can be deployed without having to run any VMs. Additionally, you can deploy physical devices on that network. Another advantages is that you can create a VPC with a tier that consists only physical devices. For example, you might create a VPC for a three-tier application, deploy VMs for Web and Application tier, and use physical machines for the Database tier. Another use case is that if you are providing services by using physical hardware, you can define the network as persistent and therefore even if all its VMs are destroyed the services will not be discontinued.

6.2.6. Cisco VNMC Support

Cisco Virtual Network Management Center (VNMC) provides centralized multi-device and policy management for Cisco Network Virtual Services. When Cisco VNMC is integrated with ASA 1000v Cloud Firewall and Cisco Nexus 1000v dvSwitch in CloudPlatform you will be able to:

- Configure Cisco ASA 1000v Firewalls

- Create and apply security profiles that contain ACL policy sets for both ingress and egress traffic, and NAT policy sets

CloudPlatform supports Cisco VNMC on Cisco Nexus 1000v dvSwitch-enabled VMware hypervisors.

6.2.7. VMware vNetwork Distributed vSwitch

CloudPlatform supports VMware vSphere Distributed Switch (VDS) for virtual network configuration in a VMware vSphere environment. Each vCenter server instance can support up to 128 VDSs and each VDS can manage up to 500 VMware hosts. CloudPlatform supports configuring virtual networks in a deployment with a mix of Virtual Distributed Switch, Standard Virtual Switch and Nexus 1000v Virtual Switch.

6.2.8. IP Reservation in Isolated Guest Networks and VPC

In Isolated guest networks and VPC in CloudPlatform 4.2, a part of the guest IP address space can be reserved for non-CloudPlatform VMs or physical servers. To do so, you configure a range of Reserved IP addresses by specifying the CIDR when a guest network is in Implemented state. The advantage of having this feature is that if your customers wish to have non-CloudPlatform controlled VMs or physical servers on the same network, they can use a part of the IP address space that is primarily provided to the guest network. When IP reservation is configured, the administrator can add additional VMs or physical servers that are not part of CloudPlatform to the same network and assign them the Reserved IP addresses. CloudPlatform guest VMs cannot acquire IPs from the Reserved IP Range.

6.2.9. Dedicated Resources: Public IP Addresses and VLANs Per Account

CloudPlatform provides you the ability to reserve a set of public IP addresses and VLANs exclusively for an account. During zone creation, you can continue to define a set of VLANs and multiple public IP ranges. This feature extends the functionality to enable you to dedicate a fixed set of VLANs and guest IP addresses for a tenant.

This feature provides you the following capabilities:

- Reserve a VLAN range and public IP address range from an Advanced zone and assign it to an account
- Disassociate a VLAN and public IP address range from an account

6.2.10. Juniper SRX Support for Egress Firewall Rules



Note

Egress firewall rules are not supported on Shared networks. They are supported only on Isolated guest networks.

Egress firewall rules were previously supported on virtual routers, and now they are also supported on Juniper SRX external networking devices.

Egress traffic originates from a private network to a public network, such as the Internet. By default, the egress traffic is blocked, so no outgoing traffic is allowed from a guest network to the Internet.

However, you can control the egress traffic in an Advanced zone by creating egress firewall rules. When an egress firewall rule is applied, the traffic specific to the rule is allowed and the remaining traffic is blocked. When all the firewall rules are removed the default policy, Block, is applied.

6.2.11. Configuring the Default Egress Policy

The default egress policy for Isolated guest network can be configured by using Network offering. Use the create network offering option to determine whether the default policy should be block or allow all the traffic to the public network from a guest network. Use this network offering to create the network. If no policy is specified, by default all the traffic is allowed from the guest network that you create by using this network offering.

You have two options: Allow and Deny.

If you select Allow for a network offering, by default egress traffic is allowed. However, when an egress rule is configured for a guest network, rules are applied to block the specified traffic and rest are allowed. If no egress rules are configured for the network, egress traffic is accepted. If you select Deny for a network offering, by default egress traffic for the guest network is blocked. However, when an egress rule is configured for a guest network, rules are applied to allow the specified traffic. While implementing a guest network, CloudPlatform adds the firewall egress rule specific to the default egress policy for the guest network.

This feature is supported only on virtual router and Juniper SRX.

6.2.12. Non-Contiguous VLAN Ranges

CloudPlatform provides you with the flexibility to add non contiguous VLAN ranges to your network. The administrator can either update an existing VLAN range or add multiple non contiguous VLAN ranges while creating a zone. You can also use the UpdatephysicalNetwork API to extend the VLAN range.

6.2.13. Isolation in Advanced Zone Using Private VLAN

Isolation of guest traffic in shared networks can be achieved by using Private VLANs (PVLAN). PVLANS provide Layer 2 isolation between ports within the same VLAN. In a PVLAN-enabled shared network, a user VM cannot reach other user VM though they can reach the DHCP server and gateway, this would in turn allow users to control traffic within a network and help them deploy multiple applications without communication between application as well as prevent communication with other users' VMs.

- Isolate VMs in a shared networks by using Private VLANs.
- Supported on KVM, XenServer, and VMware hypervisors.
- PVLAN-enabled shared network can be a part of multiple networks of a guest VM.

For further reading:

- [Understanding Private VLANs](#)¹
- [Cisco Systems' Private VLANs: Scalable Security in a Multi-Client Environment](#)²

¹ http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2_25_see/configuration/guide/swpvlan.html#wp1038379

² <http://tools.ietf.org/html/rfc5517>

- [Private VLAN \(PVLAN\) on vNetwork Distributed Switch - Concept Overview \(1010691\)](#)³

6.2.14. Configuring Multiple IP Addresses on a Single NIC

(Supported on XenServer, KVM, and VMware hypervisors)

CloudPlatform now provides you the ability to associate multiple private IP addresses per guest VM NIC. This feature is supported on all the network configurations—Shared, Isolated, and VPC. Security Groups, Static NAT and Port forwarding services are supported on these additional IPs. In addition to the primary IP, you can assign additional IPs to the guest VM NIC. Up to 256 IP addresses are allowed per NIC.

As always, you can specify an IP from the guest subnet; if not specified, an IP is automatically picked up from the guest VM subnet. You can view the IPs associated with for each guest VM NICs on the UI. You can apply NAT on these additional guest IPs by using firewall configuration in the CloudPlatform UI. You must specify the NIC to which the IP should be associated.

6.2.15. Adding Multiple IP Ranges

(Supported on KVM, XenServer, and VMware hypervisors)



Note

The feature can only be implemented on IPv4 addresses.

CloudPlatform 4.2 provides you with the flexibility to add guest IP ranges from different subnets in Basic zones and security groups-enabled Advanced zones. For security groups-enabled Advanced zones, it implies multiple subnets can be added to the same VLAN. With the addition of this feature, you will be able to add IP address ranges from the same subnet or from a different one when IP address are exhausted. This would in turn allows you to employ higher number of subnets and thus reduce the address management overhead.

Ensure that you manually configure the gateway of the new subnet before adding the IP range. Note that CloudPlatform supports only one gateway for a subnet; overlapping subnets are not currently supported.

You can also delete IP ranges. This operation fails if an IP from the remove range is in use. If the remove range contains the IP address on which the DHCP server is running, CloudPlatform acquires a new IP from the same subnet. If no IP is available in the subnet, the remove operation fails.

6.2.16. Support for Multiple Networks in VMs

(Supported on XenServer, VMware and KVM hypervisors)

CloudPlatform 4.2 provides you the ability to add and remove multiple networks to a VM. You can remove a network from a VM and add a new network. You can also change the default network of a VM. For this feature to work on VMware, ensure that vm-tools are running on guest VMs.

³ <http://kb.vmware.com>

6.2.17. Global Server Load Balancing

CloudPlatform 4.2 supports Global Server Load Balancing (GSLB) functionalities to provide business continuity by load balancing traffic to an instance on active zones only in case of zone failures . CloudPlatform achieve this by extending its functionality of integrating with NetScaler Application Delivery Controller (ADC), which also provides various GSLB capabilities, such as disaster recovery and load balancing. The DNS redirection technique is used to achieve GSLB in CloudPlatform. In order to support this functionality, region level services and service provider are introduced. A new service 'GSLB' is introduced as a region level service. The GSLB service provider is introduced that will provider the GSLB service. Currently, NetScaler is the supported GSLB provider in CloudPlatform. GSLB functionality works in an Active-Active data center environment.

6.2.18. Load Balancing Using External Provider on Shared VLANs

Network services like Firewall, Load Balancing, and NAT are now supported in shared networks created in an advanced zone. In effect, the following network services shall be made available to a VM in a shared network: Source NAT, Static NAT, Port Forwarding, Firewall and Load balancing. Subset of these service can be chosen while creating a network offering for shared networks. Services available in a shared network is defined by the network offering and the service chosen in the network offering. For example, if network offering for a shared network has source NAT service enabled, a public IP shall be provisioned and source NAT is configured on the firewall device to provide public access to the VMs on the shared network. Static NAT, Port Forwarding, Load Balancing, and Firewall services shall be available only on the acquired public IPs associated with a shared network. Additionally, Netscaler and Juniper SRX firewall device can be configured inline or side-by-side mode

Note that Source NAT per zone is not supported when the service provider is virtual router. However, Source NAT per account is supported with virtual router in a Shared Network..

6.2.19. Health Checks for Load Balanced Instances



Note

This feature is supported only on NetScaler version 10.0 and beyond.

(NetScaler load balancer only) A load balancer rule distributes requests among a pool of services (a service in this context means an application running on a virtual machine). When creating a load balancer rule, you can specify a health check which will ensure that the rule forwards requests only to services that are healthy (running and available). When a health check is in effect, the load balancer will stop forwarding requests to any resources that it has found to be unhealthy. If the resource later becomes available again, the periodic health check (periodicity is configurable) will discover it and the resource will once again be made available to the load balancer.

To configure how often the health check is performed by default, use the global configuration setting `healthcheck.update.interval`. This default applies to all the health check policies in the cloud. You can override this value for an individual health check policy.

6.3. Host and Virtual Machine Enhancements

The following new features expand the ways you can use hosts and virtual machines.

6.3.1. VMware DRS Support

The VMware vSphere Distributed Resources Scheduler (DRS) is supported.

6.3.2. Windows 8 and Windows Server 2012 as VM Guest OS

(Supported on XenServer, VMware, and KVM)

Windows 8 and Windows Server 2012 can now be used as OS types on guest virtual machines. The OS would be made available the same as any other, by uploading an ISO or a template. When you register a Windows Server 2012 template in CloudStack and choose Windows Server 2012 (64-Bit) as OS type, CloudStack can successfully deploy a Windows Server 2012 vm from the registered template. But on vCenter General panel, Guest OS description for this VM will be shown as Other (64-bit) because VMware 5.1 SDK does not provide an OS enum type for Windows 2012. The detailed instructions for uploading ISOs and templates are given in the Administrator's Guide.



Note

Limitation: When used with VMware hosts, this feature works only for the following versions: vSphere ESXi 5.1 and ESXi 5.0 Patch 4.

6.3.3. Change Account Ownership of Virtual Machines

A root administrator can now change the ownership of any virtual machine from one account to any other account. A domain or sub-domain administrator can do the same for VMs within the domain from one account to any other account in the domain.

6.3.4. Dedicated Pod, Cluster, or Host

Dedicating pod, cluster or host to a specific domain/account means that the domain/account will have sole access to the dedicated pod, cluster or hosts such that scalability, security and manageability within a domain/account can be improved. The resources which belong to that tenant will be placed into that dedicated pod, cluster or host.

6.3.5. Resizing Volumes

CloudPlatform provides the ability to resize data disks; CloudPlatform controls volume size by using disk offerings. This provides CloudPlatform administrators with the flexibility to choose how much space they want to make available to the end users. Volumes within the disk offerings with the same storage tag can be resized. For example, if you only want to offer 10, 50, and 100 GB offerings, the allowed resize should stay within those limits. That implies if you define a 10 GB, a 50 GB and a 100 GB disk offerings, a user can upgrade from 10 GB to 50 GB, or 50 GB to 100 GB. If you create a custom-sized disk offering, then you have the option to resize the volume by specifying a new, larger size. Additionally, using the `resizeVolume` API, a data volume can be moved from a static disk offering to a custom disk offering with the size specified. This functionality allows those who might be billing by certain volume sizes or disk offerings to stick to that model, while providing the flexibility to migrate to whatever custom size necessary. This feature is supported on KVM, XenServer, and VMware hosts. However, shrinking volumes is not supported on VMware hosts.

6.3.6. VMware Volume Snapshot Improved Performance

When you take a snapshot of a data volume on VMware, CloudPlatform will now use a more efficient storage technique to improve performance.

Previously, every snapshot was immediately exported from vCenter to a mounted NFS share and packaged into an OVA file format. This operation consumed time and resources. Starting from 4.2, the original file formats, such as VMDK, provided by vCenter will be retained. An OVA file will only be created as needed, on demand.

The new process applies only to newly created snapshots after upgrade to CloudPlatform 4.2. Snapshots that have already been taken and stored in OVA format will continue to exist in that format, and will continue to work as expected.

6.3.7. Live Storage Migration for XenServer and VMware

(Supported on XenServer and VMware)

Storage migration allows VMs to be moved from one host to another, where the VMs are not located on storage shared between the two hosts. It provides the option to live migrate a VM's disks along with the VM itself. It is now possible to migrate a VM from one XenServer resource pool / VMware cluster to another, or to migrate a VM whose disks are on local storage, or even to migrate a VM's disks from one storage repository to another, all while the VM is running.

6.3.8. Configuring Usage of Linked Clones on VMware

(For ESX hypervisor in conjunction with vCenter)

In CloudPlatform 4.2, the creation of VMs as full clones is allowed. In previous versions, only linked clones were possible.

For a full description of clone types, refer to VMware documentation. In summary: A full clone is a copy of an existing virtual machine which, once created, does not depend in any way on the original virtual machine. A linked clone is also a copy of an existing virtual machine, but it has ongoing dependency on the original. A linked clone shares the virtual disk of the original VM, and retains access to all files that were present at the time the clone was created.

A new global configuration setting has been added, `vmware.create.full.clone`. When the administrator sets this to true, end users can create guest VMs only as full clones. The default value is true.



Warning

For customers upgrading from a previous version of CloudPlatform: If your cloud includes any deployed data centers, you should set `vmware.create.full.clone` to false after your upgrade is complete.

6.3.9. VM Deployment Rules

Rules can be set up to ensure that particular VMs are not placed on the same physical host. These "anti-affinity rules" can increase the reliability of applications by ensuring that the failure of a single host can not take down the entire group of VMs supporting a given application. See Affinity Groups in the CloudPlatform 4.2 Administration Guide.

6.3.10. CPU and Memory Scaling for Running VMs

(Supported on VMware and XenServer)

You can now change the CPU and RAM values for a running virtual machine. In previous versions of CloudPlatform, this could only be done on a stopped VM.

It is not always possible to accurately predict the CPU and RAM requirements when you first deploy a VM. You might need to increase or decrease these resources at any time during the life of a VM. With the new ability to dynamically modify CPU and RAM levels, you can change these resources for a running VM without incurring any downtime.

Dynamic CPU and RAM scaling can be used in the following cases:

- New VMs that are created after the installation of CloudPlatform 4.2. If you are upgrading from a previous version of CloudPlatform, your existing VMs created with previous versions will not have the dynamic scaling capability.
- User VMs on hosts running VMware and XenServer.
- System VMs on VMware.
- VM Tools or XenServer Tools must be installed on the virtual machine.
- The new requested CPU and RAM values must be within the constraints allowed by the hypervisor and the VM operating system.

6.3.11. CPU and Memory Over-Provisioning

(Supported for XenServer, KVM, and VMware)

In CloudPlatform 4.2, CPU and memory (RAM) over-provisioning factors can be set for each cluster to change the number of VMs that can run on each host in the cluster. This helps optimize the use of resources. By increasing the over-provisioning ratio, more resource capacity will be used. If the ratio is set to 1, no over-provisioning is done.

In previous releases, CloudPlatform did not perform memory over-provisioning. It performed CPU over-provisioning based on a ratio configured by the administrator in the global configuration setting `cpu.overprovisioning.factor`. Starting in 4.2, the administrator can specify a memory over-provisioning ratio, and can specify both CPU and memory over-provisioning ratios on a per-cluster basis, rather than only on a global basis.

In any given cloud, the optimum number of VMs for each host is affected by such things as the hypervisor, storage, and hardware configuration. These may be different for each cluster in the same cloud. A single global over-provisioning setting could not provide the best utilization for all the different clusters in the cloud. It had to be set for the lowest common denominator. The new per-cluster setting provides a finer granularity for better utilization of resources, no matter where the CloudPlatform placement algorithm decides to place a VM.

6.3.12. Kickstart Installation for Baremetal Provisioning

CloudPlatform 4.2 supports the kickstart installation method for RPM-based Linux operating systems on Baremetal hosts in basic zones. Users can provision a Baremetal host managed by CloudPlatform as long as they have the kick start file and corresponding OS installation ISO ready.

Tested on CentOS 5.5, CentOS 6.2, CentOS 6.3, Ubuntu 12.04.

For more information, see the Baremetal Installation Guide.

6.3.13. Enhanced Bare Metal Support on Cisco UCS

You can now more easily provision new Cisco UCS server blades into CloudPlatform for use as bare metal hosts. The goal is to enable easy expansion of the cloud by leveraging the programmability of the UCS converged infrastructure and CloudPlatform's knowledge of the cloud architecture and ability to orchestrate. With this new feature, CloudPlatform can automatically understand the UCS environment, server profiles, etc. to make it easy to deploy a bare metal OS on a Cisco UCS.

6.3.14. Changing a VM's Base Image

Every VM is created from a base image, which is a template or ISO which has been created and stored in CloudPlatform. Both cloud administrators and end users can create and modify templates, ISOs, and VMs.

In CloudPlatform 4.2, there is a new way to modify an existing VM. You can change an existing VM from one base image to another. For example, suppose there is a template based on a particular operating system, and the OS vendor releases a software patch. The administrator or user naturally wants to apply the patch and then make sure existing VMs start using it. Whether a software update is involved or not, it's also possible to simply switch a VM from its current template to any other desired template.

6.3.15. Reset VM on Reboot

In CloudPlatform 4.2, you can specify that you want to discard the root disk and create a new one whenever a given VM is rebooted. This is useful for secure environments that need a fresh start on every boot and for desktops that should not retain state. The IP address of the VM will not change due to this operation.

6.3.16. Virtual Machine Snapshots for VMware

(VMware hosts only) In addition to the existing CloudPlatform ability to snapshot individual VM volumes, you can now take a VM snapshot to preserve all the VM's data volumes as well as (optionally) its CPU/memory state. This is useful for quick restore of a VM. For example, you can snapshot a VM, then make changes such as software upgrades. If anything goes wrong, simply restore the VM to its previous state using the previously saved VM snapshot.

The snapshot is created using the VMware native snapshot facility. The VM snapshot includes not only the data volumes, but optionally also whether the VM is running or turned off (CPU state) and the memory contents. The snapshot is stored in CloudPlatform's primary storage.

VM snapshots can have a parent/child relationship. Each successive snapshot of the same VM is the child of the snapshot that came before it. Each time you take an additional snapshot of the same VM, it saves only the differences between the current state of the VM and the state stored in the most recent previous snapshot. The previous snapshot becomes a parent, and the new snapshot is its child. It is possible to create a long chain of these parent/child snapshots, which amount to a "redo" record leading from the current state of the VM back to the original.

6.3.17. Increased Userdata Size When Deploying a VM

You can now specify up to 32KB of userdata when deploying a virtual machine through the CloudPlatform UI or the `deployVirtualMachine` API call.

6.3.18. Set VMware Cluster Size Limit Depending on VMware Version

The maximum number of hosts in a vSphere cluster is determined by the VMware hypervisor software. For VMware versions 4.2, 4.1, 5.0, and 5.1, the limit is 32 hosts.

For CloudPlatform 4.2, the global configuration setting `vmware.percluster.host.max` has been removed. The maximum number of hosts in a VMware cluster is now determined by the underlying hypervisor software.



Note

Best Practice: It is advisable for VMware clusters in CloudPlatform to be smaller than the VMware hypervisor's maximum size. A cluster size of up to 8 hosts has been found optimal for most real-world situations.

6.3.19. Limiting Resource Usage

Previously in CloudPlatform, resource usage limit was imposed based on the resource count, that is, restrict a user or domain on the basis of the number of VMs, volumes, or snapshots used. In CloudPlatform 4.2, a new set of resource types has been added to the existing pool of resources (VMs, Volumes, and Snapshots) to support the customization model—need-basis usage, such as large VM or small VM. The new resource types are now broadly classified as CPU, RAM, Primary storage, and Secondary storage. CloudPlatform 4.2 allows the root administrator to impose resource usage limit by the following resource types for Domain, Project and Accounts.

- CPUs
- Memory (RAM)
- Primary Storage (Volumes)
- Secondary Storage (Snapshots, Templates, ISOs)

6.4. Monitoring, Maintenance, and Operations Enhancements

6.4.1. Publish and Subscribe for Event Notification

An event is essentially a significant or meaningful change in the state of both virtual and physical resources associated with a cloud environment. In CloudPlatform an event could be a state change of virtual or physical resources, an action performed by an user (action events), or policy based events (alerts). In CloudPlatform 4.2, a new event notification framework has been added. This framework provides a means for the Management Server components to publish and subscribe to CloudPlatform events. Event notification is achieved by implementing the concept of event bus abstraction in the Management Server.

A new event for state change, resource state change, is introduced as part of Event notification framework. Every resource, such as user VM, volume, NIC, network, public IP, snapshot, and template, is associated with a state machine and generates events as part of the state change. That implies that a change in the state of a resource results in a state change event, and the event

is published in the corresponding state machine on the event bus. All the CloudPlatform events (alerts, action events, usage events) and the additional category of resource state change events, are published on to the events bus.

6.4.2. Deleting and Archiving Events and Alerts

In addition to viewing a list of events and alerts in the UI, the administrator can now delete and archive them. In order to support deleting and archiving alerts, the following global parameters have been added:

- **alert.purge.delay**: The alerts older than specified number of days are purged. Set the value to 0 to never purge alerts automatically.
- **alert.purge.interval**: The interval in seconds to wait before running the alert purge thread. The default is 86400 seconds (one day).



Note

Archived alerts or events cannot be viewed in the UI, or by using the API. They are maintained in the database for auditing or compliance purposes.

6.4.3. Increased Granularity for Configuration Parameters

Some configuration parameters which were previously available only at the global level of the cloud can now be set for smaller components of the cloud, such as at the zone level. To set these parameters, look for the new Settings tab in the UI. You will find it on the detail page for an account, cluster, zone, or primary storage.

The account level parameters are: `remote.access.vpn.client.iprange`, `allow.public.user.templates`, `use.system.public.ips`, and `use.system.guest.vlans`

The cluster level parameters are

`cluster.storage.allocated.capacity.notificationthreshold`,
`cluster.storage.capacity.notificationthreshold`,
`cluster.cpu.allocated.capacity.notificationthreshold`,
`cluster.memory.allocated.capacity.notificationthreshold`,
`cluster.cpu.allocated.capacity.disablethreshold`,
`cluster.memory.allocated.capacity.disablethreshold`,
`cpu.overprovisioning.factor`, `mem.overprovisioning.factor`, `vmware.reserve.cpu`,
and `vmware.reserve.mem`.

The zone level parameters are `pool.storage.allocated.capacity.disablethreshold`,
`pool.storage.capacity.disablethreshold`, `storage.overprovisioning.factor`,
`network.throttling.rate`, `guest.domain.suffix`, `router.template.xen`,
`router.template.kvm`, `router.template.vmware`, `router.template.hyperv`,
`router.template.lxc`, `enable.dynamic.scale.vm`, `use.external.dns`, and
`blacklisted.routes`.

6.4.4. API Request Throttling

In CloudPlatform 4.2, you can limit the rate at which API requests can be placed for each account. This is useful to avoid malicious attacks on the Management Server, prevent performance degradation, and provide fairness to all accounts.

If the number of API calls exceeds the threshold, an error message is returned for any additional API calls. The caller will have to retry these API calls at another time.

To control the API request throttling, use the following new global configuration settings:

- `api.throttling.enabled` - Enable/Disable API throttling. By default, this setting is false, so API throttling is not enabled.
- `api.throttling.interval` (in seconds) - Time interval during which the number of API requests is to be counted. When the interval has passed, the API count is reset to 0.
- `api.throttling.max` - Maximum number of APIs that can be placed within the `api.throttling.interval` period.
- `api.throttling.cachesize` - Cache size for storing API counters. Use a value higher than the total number of accounts managed by the cloud. One cache entry is needed for each account, to store the running API total for that account within the current time window.

6.4.5. Sending Alerts to External SNMP and Syslog Managers

In addition to showing administrator alerts on the Dashboard in the CloudPlatform UI and sending them in email, CloudPlatform now can also send the same alerts to external SNMP or Syslog management software. This is useful if you prefer to use an SNMP or Syslog manager to monitor your cloud.

The supported protocol is SNMP version 2.

6.4.6. Changing the Default Password Encryption

Passwords are encoded when creating or updating users. The new default preferred encoder, replacing MD5, is SHA256. It is more secure than MD5 hashing. If you take no action to customize password encryption and authentication, SHA256 Salt will be used.

If you prefer a different authentication mechanism, CloudPlatform 4.2 provides a way for you to determine the default encoding and authentication mechanism for admin and user logins. Two new configurable lists have been introduced: `userPasswordEncoders` and `userAuthenticators`. `userPasswordEncoders` allow you to configure the order of preference for encoding passwords, and `userAuthenticator` allows you to configure the order in which authentication schemes are invoked to validate user passwords.

The plain text user authenticator has been modified not to convert supplied passwords to their md5 sums before checking them with the database entries. It performs a simple string comparison between retrieved and supplied login passwords instead of comparing the retrieved md5 hash of the stored password against the supplied md5 hash of the password, because clients no longer hash the password.

6.4.7. Log Collection Utility `cloud-bugtool`

CloudPlatform provides a command-line utility called `cloud-bugtool` to make it easier to collect the logs and other diagnostic data required for troubleshooting. This is especially useful when interacting with Citrix Technical Support.

You can use `cloud-bugtool` to collect the following:

- Basic system and environment information and network configuration including IP addresses, routing, and name resolver settings

- Information about running processes
- Management Server logs
- System logs in /var/log/
- Dump of the cloud database



Warning

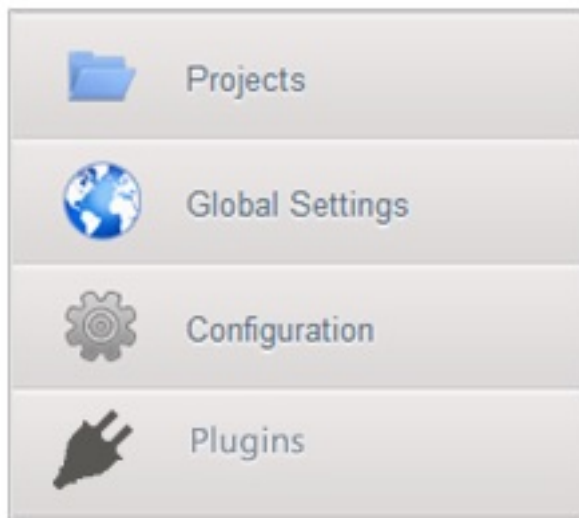
cloud-bugtool collects information which might be considered sensitive and confidential. Using the `--noDb` option to avoid the cloud database can reduce this concern, though it is not guaranteed to exclude all sensitive data.

6.5. Third-Party UI Plugin Framework

Using the new third-party plugin framework, you can write and install extensions to CloudPlatform. The installed and enabled plugins will appear in the UI alongside the Citrix-provided features.

The basic procedure for adding a UI plugin is explained in the Developer Guide. In summary, the plugin developer creates the plugin code itself (in Javascript), a thumbnail image, the plugin listing, and a CSS file. The CloudPlatform administrator adds the folder containing the plugin code under the CloudPlatform PLUGINS folder and adds the plugin name to a configuration file (plugins.js).

The next time the user refreshes the UI in the browser, the plugin will appear under the Plugins button in the left navigation bar.



Fixed Issues

| Issue ID | Description |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CS-9018 | Domain router performance enhancement: Adding NAT rules now only specifically adds NAT rules to the interface that the port forward rule deals with. Lowering CPU and memory usage for processing packets. In previous versions, the NAT rules were programmed on all interfaces, incurring a performance hit. |
| CS-12707 | Template and ISO entries are no longer repeated in the UI or API result. |
| CS-13192 | (F5/Netscaler) IpAssoc command is no longer sent to router before and after Lb rules are created. This applies in cases where LB service is provided by F5/Netscaler. |
| CS-14342 | Added the ability to pack/unpack OVAs on secondary storage. This removes the overhead of unpacking when deploying to the primary storage. |
| CS-15184 | When attempting to delete a host from a cluster that still has hosts present, the administrator is no longer presented with an invalid error message. |
| CS-15609 | Volumes no longer can be created as a part of VM creation, where unallocated space is insufficient on primary storage. |
| CS-16087 | A network cannot be created from a network offering where some network services are associated with no providers. This is by design. |
| CS-16324 | Enable Baremetal providers when creating a basic zone that has Baremetal network offering. |
| CS-16495 | Make DNS/DHCP service more robust in virtual router. |
| CS-16503 | Improved statistic and count reporting in list API commands. |
| CS-16702 | Can start console proxy VM with Hypervisor VMware 5.0. |
| CS-16738 | The listZones API now accepts pagesize parameter correctly. |
| CS-16798 | (XenServer) Hosts no longer enter alert state when upgrading the hypervisor from XenServer 6.0.2 to 6.1. |
| CS-16835 | User-entered time offsets are now respected, and persisted when VMs stop and start. |

| Issue ID | Description |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CS-16880 | After all public IPs are released, the corresponding NIC on virtual routers is now unplugged. |
| CS-16884 | List Zone count is now displayed in UI. |
| CS-16921 | The ManageSnapshot command handles long Snapshot chain. |
| CS-16940 | UI: stop VM now properly returns an error when the hosts are not managed by any MS in the cluster. |
| CS-16974 | In a Security Group-enabled Advanced Zone, virtual router can be started after upgrade from 2.2.14 to 3.0.6. |
| CS-17003 | Network statistics are now correctly collected and presented for VMs residing on VMware hypervisors. |
| CS-17022 | Modified qemu to make CloudPlatform compatible with RHEL 6.3. |
| CS-17052 | SRX-F5 - Now able to add an SRX device after it was removed. |
| CS-17053 | When editing a pod in the CloudPlatform UI, there is now a Save button. |
| CS-17058 | The field current_bytes_sent is now updated in the user_statistics table in case of load balancing, when an SRX is present. |
| CS-17059 | Improved billing related to statistics collected using the SRX firewall. |
| CS-17085 | On Juniper SRX device, the ability to collect usage statistics is integrated in shared networks. |
| CS-17088 | The OS Type information is now consistent on upgraded setups. |
| CS-17121 | Instead of a hard-coded value in VM_DETAILS_BATCH_SIZE, new global configuration variable, detail.batch.query.size, is now used. This enhances the ability for performance tuning. |
| CS-17139 | (KVM) Upgraded setups now properly delete and clean up root and data volumes from primary storage that were created prior to upgrade and marked for cleanup. |
| CS-17150 | (XenServer) CPU cap is correctly set when using XenServer |
| CS-17169 | The administrator is no longer allowed to add a shared network with a subnet that is already associated with another VLAN. |

| Issue ID | Description |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CS-17183 | Response time for list* API calls for page size beyond 1000 improved. The response time earlier was more than 10 minutes. |
| CS-17225 | In the templates list in the UI, the template will no longer disappear after you reorder one of the templates multiple times. |
| CS-17272 | If you cancel the Recurring Snapshots dialog by pressing Esc key, the UI continues to function normally. Previously, it would freeze. |
| CS-17337 | Can now tell whether a compute offering is public or private in the List view or Detail view. |
| CS-17342 | Restart Network now properly creates Router VM in case of Multiple Hypervisors present in the environment. |
| CS-17607 | Added checks for volume names which had spaces. The API previously allowed creation of volumes with spaces in names, which led to problems when creating snapshots. |
| CS-17620 | Worker VMs are no longer orphaned after snapshot. |
| CS-17625, CS-17635 | Zero bytes size ISO template import no longer causes infinite loop on SSVM. |
| CS-17632 | Duplicate IP addresses can no longer be assigned in a single isolated network. |
| CS-17655 | (VMware) Added the ability to parallelize commands sent to the SSVM to improve performance. |
| CS-17765 | KVM host no longer reboots itself after deleting primary storage. |
| CS-17769 | Can create more than one isolated network with same CIDR, which is different from zone CIDR. |
| CS-17781 | Local Storage Disk status is no longer out of sync between database and XenServer. You can now attach new disk. |
| CS-17941 | After a shared guest network is deleted, you can now create a new shared guest network with IP range overlapping deleted network. |
| CS-17989 | You can create a private template from a snapshot. |
| CS-18068 | (KVM) Security group rules are no longer cleaned for a VM during snapshotting. |
| CS-18070 | With multiple subnets per VLAN, VM metadata is now working on the additional subnets. |
| CS-18117 | In a redundant virtual router setup, if the master router goes down, public IP addresses are correctly configured on the new master. |

Chapter 7. Fixed Issues

| Issue ID | Description |
|----------|----------------------------------------------------------------------------------|
| CS-18121 | Resource counts are now reflecting accurately, where user accounts were deleted. |

Known Issues

| Issue ID | Description |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CS-16008 | <p>In a clustered management server deployment, hosts are not load balanced across management servers in cluster. This is by design.</p> <p>Workaround: All Management server in cluster must be synced by running:</p> <pre data-bbox="855 600 1441 667"># ntpdate 0.xenserver.pool.ntp.org</pre> <pre data-bbox="855 685 1441 752"># service ntpd start</pre> |
| CLOUDSTACK-1717 | <p>Local region entry that gets added by default should not include "/api" for its end_point. Also the endpoint should have the actual hostname instead of localhost.</p> |
| CLOUDSTACK-1960 | <p>The euro symbol € does not work while accessing the guest virtual machine consoles on a UK keyboard by using console proxy.</p> |
| CLOUDSTACK-1964 | <p>In Simplified Chinese, some combination keys used to switch IME cannot work well.</p> <p>Workaround: For "Ctrl+Shift" and "Ctrl+Space", click the input style of IME to select the input style and switch keyboard layout. For "Ctrl+Dot", click the "Chinese/Western Punctuation (Ctrl +.)" in the IME Toolbar to switch the punctuation between full-width and half-width.</p> |
| CLOUDSTACK-1986 | <p>The Japanese keyboard keys ¥, \, , Muhenkan, Henkan, and Hiragana/Katakana are not working even after possible key translations tried.</p> <p>Workaround:</p> <p>For keys: _</p> <p>Set the console proxy keyboard layout to "Standard (US) Keyboard". Add English Keyboard layout to the Japanese guest VM from "Regional Setting" option from Control Panel (in case of Windows). Set the Japanese guest OS keyboard layout to "EN". Try the keyboard keys _ using Japanese keyboard in localized environment.</p> <p>For Muhenkan key:</p> <p>You can use F6, F7 and F8 instead of the Muhenkan key. F6 key converts the string into Hiragana. F7 key converts the string into</p> |

| Issue ID | Description |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <p>Katakana. F8 key converts the string into Hankaku-Katakana. Muhenkan keys toggles the string Hiragana, Katakana and Hankaku-Katakana.</p> <p>Henkan key:</p> <p>You can use space bar (key) instead of the Henkan key.</p> <p>Hiragana/Katakana key:</p> <p>We have to use IME menu below to change IME input mode in case the Hiragana/Katakana key is unavailable.</p> |
| CLOUDSTACK-2112 | VM will go into stopped state after live migration failed during a scale up VMs operation. Need to be manually restarted. |
| CLOUDSTACK-2293 | DeletePhysicalNetworkCmd is not deleting the external devices. |
| CLOUDSTACK-2562 | <p>When a virtual router on a VMware ESX is restarted out of band by VMware HA, CloudPlatform does not re-program the port forwarding, NAT, and load balancer rules.</p> <p>Workaround: Restart the virtual router using the CloudPlatform UI or API.</p> |
| CLOUDSTACK-2646 | <p>When firewall and LB service providers are different, CloudPlatform incorrectly allows both the rules on the same public IP. Workaround: Admin should not create network offering with different service providers for firewall and LB, while keeping conserve mode on.</p> |
| CLOUDSTACK-2709 | Egress rules are not supported for shared networks. |
| CLOUDSTACK-2910 | <p>Ctrl combined with >. is not working on SC IME.</p> <p>Workaround: Click the “Chinese/Western Punctuation(Ctrl+.)” in the IME tool bar to switch the punctuation between full-width and half-width.</p> |
| CLOUDSTACK-3111 | Volume listing screen shows Hypervisor column as empty if the volumes are attached to instances running in KVM Hypervisor. |
| CLOUDSTACK-3154 | When deleting the VMware DataCenter from the Zone the DataCenter gets deleted but UI screen stays in processing state. |
| CLOUDSTACK-3212 | Default guest network can now have multiple subnets per VLAN, but the IP range list page |

| Issue ID | Description |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <p>does not display the netmask and gateway for each subnet.</p> <p>Workaround: Use the API listVlanIPRanges to get the complete details.</p> |
| CLOUDSTACK-3223 | <p>(VMware) During system virtual machine deployment, you might see an error message like "Cannot complete operation due to concurrent modification by another operation" in VMware vCenter console while trying to configure virtual network for the management network. This is a concurrency issue observed when two or more system virtual machines try to configure the same virtual network on the vCenter. This operation will not have any operational impact, and so can be ignored.</p> |
| CLOUDSTACK-3252 | <p>An instance deployed using explicit or implicit dedication doesn't generate a usage event.</p> |
| CLOUDSTACK-3317 | <p>Management and storage network traffic cannot be configured to use VMware Distributed vSwitch (DVS). Continue to use standard vSwitch.</p> |
| CLOUDSTACK-3466 | <p>VM Migration across VMware clusters which are added with different switches (Standard Switch, VMware DVS, Cisco Nexus 1000v) is not supported.</p> |
| CLOUDSTACK-3680 | <p>(KVM on CentOS 5.5, 5.6) While accessing console view of a guest virtual machine, the keystrokes tab, ctrl, \, tilde, single quote, double quote, and caret ^ do not work on CentOS 5.5\5.6 running on KVM. This is due to a known bug in CentOS (see http://www.centos.org/modules/newbb/viewtopic.php?topic_id=33233&forum=55¹).</p> |
| CLOUDSTACK-3968 | <p>Distributed port groups on DV Switch are not removed when the associated account from CloudPlatform is removed.</p> |
| CLOUDSTACK-4016 | <p>listPublicIpAddresses lists the portable IP that was already transferred to a different Isolated network.</p> |
| CLOUDSTACK-4128 | <p>When using S3 storage, it is required to have staging secondary storage for the Zone but the UI does not enforce it. This will result in System VMs not coming up when zone is enabled.</p> <p>Workaround: Add staging secondary storage prior to enabling the zone.</p> |

¹ http://www.centos.org/modules/newbb/viewtopic.php?topic_id=33233&forum=55

| Issue ID | Description |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CLOUDSTACK-4139 | <p>(VMware) The volumes created from snapshots on VMware deployments cannot be resized when attached to a running VM. The volume is created with IDE disk instead of SCSI disk which cannot be resized.</p> <p>Workaround: Detach the volume created from a snapshot and resize it, and then reattach it to the VM.</p> |
| CLOUDSTACK-4207 | <p>The following exception is observed when the Management Server is started after upgrade from any older versions to CloudPlatform 4.2.</p> <p>jsonParseException: The JsonSerializer com.cloud.agent.transport.ArrayTypeAdaptor@2426e26f failed to deserialize json object</p> <p>Ignore this exception, this would stop after you upgrade the System VM. However, if you want to prevent this, stop system VM from the hypervisor before upgrade.</p> |
| CLOUDSTACK-4300 | <p>System VM fails to start on KVM Host running 6.0 after CloudPlatform upgrade to 4.2. Workaround: first upgrade the KVM host to 6.2+ before upgrading CloudPlatform.</p> |
| CLOUDSTACK-4327 | <p>When primary storage is put into maintenance, the System VMs are put again on the same pool and its maintenance state is not honored.</p> <p>Workaround: Shutdown the storage.</p> |
| CLOUDSTACK-4364 | <p>Restore VM needs to log usage event for volume so that it is correctly charged for usage.</p> |
| CLOUDSTACK-4402 | <p>Cannot delete primary storage if the associated host is already removed. Workaround: Unmount the primary storage first before deleting the host.</p> |
| CLOUDSTACK-4475 | <p>If cluster-wide and zone-wide primary storage are mixed together, the data disk by default will be created on cluster wide primary storage.</p> <p>Workaround: If admin wants data disk to be created on zone-wide primary storage, then create a disk offering with the tag on zone-wide primary storage.</p> |
| CLOUDSTACK-4492 | <p>Uploaded volume state was not set to "Uploaded" in CloudPlatform 3.0.6. After upgrade to 4.x, volume attach fails because of volume being in incorrect state. Workaround: Upload and attach volume after the upgrade.</p> |

| Issue ID | Description |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CLOUDSTACK-4504 | VM creation is failing using the Ubuntu ISO with XenServer 6.1 and 6.2. |
| CLOUDSTACK-4517 | Deployment of VM using CentOS 6.2 template registered before upgrade is failing. |
| CLOUDSTACK-4578 | <p>(VMware) If the host where the SSVM is running goes down, the SSVM is not being recreated on another host in the cluster.</p> <p>Workaround: Forcefully stop the SSVM through the CloudPlatform API call stopSystemVm. Then the new SSVM will be created on a second host.</p> |
| CLOUDSTACK-4593 | <p>Live Storage Migration and VM Snapshot features are not fully functional after upgrade.</p> <p>Workaround: Stop and then start the VM post upgrade.</p> |
| CLOUDSTACK-4594 | After upgrading CloudPlatform if you perform Live Storage Migration of a VM, reverting that VM to an older snapshot fails if the snapshot is created before the Live Storage Migration operation. |
| CLOUDSTACK-4596 | CloudPlatform allows the same IP range to be defined in different VLANs across public and portable IP ranges. |
| CLOUDSTACK-4622 | If a VM from a guest network is added to a network tier of a VPC, then IP reservation allows the CIDR to be the superset of Network CIDR for that VPC tier. |
| CLOUDSTACK-4707 | The <code>sourceTemplateId</code> field is not getting set for derived templates. |

API Changes from 3.0 to 4.2

9.1. Added API Commands in 4.2

9.1.1. Secondary Storage

- `addImageStore` (Adds all types of secondary storage providers, S3/Swift/NFS)
- `createSecondaryStagingStore` (Adds a staging secondary storage in each zone)
- `listImageStores` (Lists all secondary storages, S3/Swift/NFS)
- `listSecondaryStagingStores` (Lists all staging secondary storages)
- `addS3` (Adds a Amazon Simple Storage Service instance.) It is recommended to use `addImageStore` instead.
- `listS3s` (Lists all the Amazon Simple Storage Service instances.) It is recommended to use `listImageStores` instead.

9.1.2. VM Snapshot

- `createVMSnapshot` (Creates a virtual machine snapshot; see [Section 6.3.16, “Virtual Machine Snapshots for VMware”](#))
- `deleteVMSnapshot` (Deletes a virtual machine snapshot)
- `listVMSnapshot` (Shows a virtual machine snapshot)
- `revertToVMSnapshot` (Returns a virtual machine to the state and data saved in a given snapshot)

9.1.3. Load Balancer Health Check

- `createLBHealthCheckPolicy` (Creates a new health check policy for a load balancer rule; see [Section 6.2.19, “Health Checks for Load Balanced Instances”](#))
- `deleteLBHealthCheckPolicy` (Deletes an existing health check policy from a load balancer rule)
- `listLBHealthCheckPolicies` (Displays the health check policy for a load balancer rule)

9.1.4. Egress Firewall Rules

- `createEgressFirewallRules` (Creates an egress firewall rule on the guest network; see [Section 6.2.10, “Juniper SRX Support for Egress Firewall Rules”](#))
- `deleteEgressFirewallRules` (Deletes a egress firewall rule on the guest network.)
- `listEgressFirewallRules` (Lists the egress firewall rules configured for a guest network.)

9.1.5. SSH Key

- `resetSSHKeyForVirtualMachine` (Resets the SSHkey for virtual machine.)

9.1.6. Bare Metal

- `addBaremetalHost` (Adds a new host. Technically, this API command was present in v3.0.6, but its functionality was disabled. See [Section 6.3.12, “Kickstart Installation for Baremetal Provisioning”](#))
- `addBaremetalDhcp` (Adds a DHCP server for bare metal hosts)
- `addBaremetalPxePingServer` (Adds a PXE PING server for bare metal hosts)
- `addBaremetalPxeKickStartServer` (Adds a PXE server for bare metal hosts)
- `listBaremetalDhcp` (Shows the DHCP servers currently defined for bare metal hosts)
- `listBaremetalPxePingServer` (Shows the PXE PING servers currently defined for bare metal hosts)

9.1.7. NIC

- `addNicToVirtualMachine` (Adds a new NIC to the specified VM on a selected network; see [Section 6.2.14, “Configuring Multiple IP Addresses on a Single NIC”](#))
- `removeNicFromVirtualMachine` (Removes the specified NIC from a selected VM.)
- `updateDefaultNicForVirtualMachine` (Updates the specified NIC to be the default one for a selected VM.)
- `addIpToNic` (Assigns secondary IP to a NIC.)
- `removeIpFromNic` (Assigns secondary IP to a NIC.)
- `listNics` (Lists the NICs associated with a VM.)

9.1.8. Regions

- `addRegion` (Registers a Region into another Region; see [Section 6.1.1, “Regions”](#))
- `updateRegion` (Updates Region details: ID, Name, Endpoint, User API Key, and User Secret Key.)
- `removeRegion` (Removes a Region from current Region.)
- `listRegions` (Get all the Regions. They can be filtered by using the ID or Name.)

9.1.9. User

- `getUser` (This API can only be used by the Admin. Get user account details by using the API Key.)

9.1.10. API Throttling

- `getApiLimit` (Show number of remaining APIs for the invoking user in current window)
- `resetApiLimit` (For root admin, if `accountId` parameter is passed, it will reset count for that particular account, otherwise it will reset all counters)
- `resetApiLimit` (Reset the API count.)

9.1.11. Locking

- `lockAccount` (Locks an account)

- lockUser (Locks a user account)

9.1.12. VM Scaling

- scaleVirtualMachine (Scales the virtual machine to a new service offering.)

9.1.13. Migrate Volume

- migrateVirtualMachineWithVolume (Attempts migrating VM with its volumes to a different host.)
- listStorageProviders (Lists storage providers.)
- findStoragePoolsForMigration (Lists storage pools available for migrating a volume.)

9.1.14. Dedicated IP and VLAN

- dedicatePublicIpRange (Dedicates a Public IP range to an account.)
- releasePublicIpRange (Releases a Public IP range back to the system pool.)
- dedicateGuestVlanRange (Dedicates a guest VLAN range to an account.)
- releaseDedicatedGuestVlanRange (Releases a dedicated guest VLAN range to the system.)
- listDedicatedGuestVlanRanges (Lists dedicated guest VLAN ranges.)

9.1.15. Port Forwarding

- updatePortForwardingRule (Updates a port forwarding rule. Only the private port and the VM can be updated.)

9.1.16. Scale System VM

- scaleSystemVm (Scale the service offering for a systemVM, console proxy, or secondary storage.)

9.1.17. Deployment Planner

- listDeploymentPlanners (Lists all the deployment planners available.)

9.1.18. Archive and Delete Events and Alerts

- archiveEvents (Archive one or more events.)
- deleteEvents (Delete one or more events.)
- archiveAlerts (Archive one or more alerts.)
- deleteAlerts (Delete one or more alerts.)

9.1.19. Host Reservation

- releaseHostReservation (Releases host reservation.)

9.1.20. Resize Volume

- resizeVolume (Resizes a volume.)

- updateVolume (Updates the volume.)

9.1.21. Egress Firewall Rules

- createEgressFirewallRule (Creates a egress firewall rule for a given network.)
- deleteEgressFirewallRule (Deletes an egress firewall rule.)
- listEgressFirewallRules (Lists all egress firewall rules for network.)

9.1.22. Network ACL

- updateNetworkACLItem (Updates ACL item with specified ID.)
- createNetworkACLList (Creates a Network ACL for the given VPC.)
- deleteNetworkACLList (Deletes a Network ACL.)
- replaceNetworkACLList (Replaces ACL associated with a Network or private gateway.)
- listNetworkACLLists (Lists all network ACLs.)

9.1.23. Resource Detail

- addResourceDetail (Adds detail for the Resource.)
- removeResourceDetail (Removes detail for the Resource.)
- listResourceDetails (List resource details.)

9.1.24. Nicira Integration

- addNiciraNvpDevice (Adds a Nicira NVP device.)
- deleteNiciraNvpDevice (Deletes a Nicira NVP device.)
- listNiciraNvpDevices (Lists Nicira NVP devices.)
- listNiciraNvpDeviceNetworks (Lists network that are using a Nicira NVP device.)

9.1.25. BigSwitch VNS

- addBigSwitchVnsDevice (Adds a BigSwitch VNS device.)
- deleteBigSwitchVnsDevice (Deletes a BigSwitch VNS device.)
- listBigSwitchVnsDevices (Lists BigSwitch VNS devices.)

9.1.26. Simulator

- configureSimulator (Configures a simulator.)

9.1.27. API Discovery

- listApis (Lists all the available APIs on the server, provided by the API Discovery plugin.)

9.1.28. Global Load Balancer

- createGlobalLoadBalancerRule (Creates a global load balancer rule.)
- deleteGlobalLoadBalancerRule (Deletes a global load balancer rule.)
- updateGlobalLoadBalancerRule (update global load balancer rules.)
- listGlobalLoadBalancerRules (Lists load balancer rules.)
- assignToGlobalLoadBalancerRule (Assign load balancer rule or list of load balancer rules to a global load balancer rules.)
- removeFromGlobalLoadBalancerRule (Removes a load balancer rule association with global load balancer rule)

9.1.29. Load Balancer

- createLoadBalancer (Creates a Load Balancer)
- listLoadBalancers (Lists Load Balancers)
- deleteLoadBalancer (Deletes a load balancer)
- configureInternalLoadBalancerElement (Configures an Internal Load Balancer element.)
- createInternalLoadBalancerElement (Create an Internal Load Balancer element.)
- listInternalLoadBalancerElements (Lists all available Internal Load Balancer elements.)

9.1.30. Affinity Group

- createAffinityGroup (Creates an affinity or anti-affinity group.)
- deleteAffinityGroup (Deletes an affinity group.)
- listAffinityGroups (Lists all the affinity groups.)
- updateVMAffinityGroup (Updates the affinity or anti-affinity group associations of a VM. The VM has to be stopped and restarted for the new properties to take effect.)
- listAffinityGroupTypes (Lists affinity group types available.)

9.1.31. Portable IP

- createPortableIpRange (Adds a range of portable IP ranges to a Region.)
- deletePortableIpRange (Deletes a range of portable IP ranges associated with a Region.)
- listPortableIpRanges (Lists portable IP ranges.)

9.1.32. Internal Load Balancer VM

- stopInternalLoadBalancerVM (Stops an Internal LB VM.)
- startInternalLoadBalancerVM (Starts an existing Internal LB VM.)
- listInternalLoadBalancerVMs (List internal LB VMs.)

9.1.33. Network Isolation

- listNetworkIsolationMethods (Lists supported methods of network isolation.)

9.1.34. Dedicated Resources

- dedicateZone (Dedicates a zone.)
- dedicatePod (Dedicates a pod.)
- dedicateCluster (Dedicate an existing cluster.)
- dedicateHost (Dedicates a host.)
- releaseDedicatedZone (Release dedication of zone.)
- releaseDedicatedPod (Release dedication for the pod.)
- releaseDedicatedCluster (Release dedication for cluster.)
- releaseDedicatedHost (Release dedication for host.)
- listDedicatedZones (List dedicated zones.)
- listDedicatedPods (Lists dedicated pods.)
- listDedicatedClusters (Lists dedicated clusters.)
- listDedicatedHosts (Lists dedicated hosts.)

9.2. Changed API Commands in 4.2

| API Commands | Description |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| listNetworkACLs | The following new request parameters are added: aclid (optional), action (optional), protocol (optional) The following new response parameters are added: aclid, action, number |
| copyTemplate | The following new response parameters are added: isdynamicallyscalable, sshkeyenabled |
| listRouters | The following new response parameters are added: ip6dns1, ip6dns2, role |
| updateConfiguration | The following new request parameters are added: accountid (optional), clusterid (optional), storageid (optional), zoneid (optional) The following new response parameters are added: id, scope |
| listVolumes | The following request parameter is removed: details The following new response parameter is added: displayvolume |

| API Commands | Description |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| suspendProject | The following new response parameters are added: cpuavailable, cpulimit, cputotal, ipavailable, iplimit, iptotal, memoryavailable, memorylimit, memorytotal, networkavailable, networklimit, networktotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal, snapshotavailable, snapshotlimit, snapshottotal, templateavailable, templatelimit, templatetotal, vmavailable, vmlimit, vmrunning, vmstopped, vmtotal, volumeavailable, volumelimit, volumetotal, vpcavailable, vpclimit, vpcotal |
| listRemoteAccessVpns | The following new response parameters are added: id |
| registerTemplate | The following new request parameters are added: imagestoreuuid (optional), isdynamicallyscalable (optional), isrouting (optional) The following new response parameters are added: isdynamicallyscalable, sshkeyenabled |
| addTrafficMonitor | The following response parameters are removed: privateinterface, privatezone, publicinterface, publiczone, usageinterface, username |
| createTemplate | The following response parameters are removed: clusterid, clustername, disksizeallocated, disksizetotal, disksizeused, ipaddress, path, podid, podname, state, tags, type The following new response parameters are added: account, accountid, bootable, checksum, crossZones, details, displaytext, domain, domainid, format, hostid, hostname, hypervisor, isdynamicallyscalable, isextractable, isfeatured, ispublic, isready, ostypeid, ostypeid, passwordenabled, project, projectid, removed, size, sourcetemplateid, sshkeyenabled, status, templatetag, templatetype, tags |
| listLoadBalancerRuleInstances | The following new response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| migrateVolume | The following new request parameters is added: livemigrate (optional) The following new response parameters is added: displayvolume |
| createAccount | The following new request parameters are added: accountid (optional), userid (optional) |

| API Commands | Description |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <p>The following new response parameters are added: accountdetails, cpuavailable, cpulimit, cputotal, defaultzoneid, ipavailable, iplimit, iptotal, iscleanuprequired, isdefault, memoryavailable, memorylimit, memorytotal, name, networkavailable, networkdomain, networklimit, networktotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, projectavailable, projectlimit, projecttotal, receivedbytes, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal, sentbytes, snapshotavailable, snapshotlimit, snapshottotal, templateavailable, templatelimit, templatetotal, vmavailable, vmlimit, vmrunning, vmstopped, vmtotal, volumeavailable, volumelimit, volumetotal, vpcavailable, vpclimit, vpctotal, user</p> <p>The following parameters are removed: account, accountid, apikey, created, email, firstname, lastname, secretkey, timezone, username</p> |
| updatePhysicalNetwork | The following new request parameters is added: removevlan (optional) |
| listTrafficMonitors | The following response parameters are removed: privateinterface, privatezone, publicinterface, publiczone, usageinterface, username |
| attachIso | The following new response parameters are added: diskiread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| listProjects | The following new request parameters are added: cpuavailable, cpulimit, cputotal, ipavailable, iplimit, iptotal, memoryavailable, memorylimit, memorytotal, networkavailable, networklimit, networktotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal, snapshotavailable, snapshotlimit, snapshottotal, templateavailable, templatelimit, templatetotal, vmavailable, vmlimit, vmrunning, vmstopped, vmtotal, volumeavailable, volumelimit, volumetotal, vpcavailable, vpclimit, vpctotal |
| enableAccount | The following new response parameters are added: cpuavailable, cpulimit, cputotal, isdefault, memoryavailable, memorylimit, memorytotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal |

| API Commands | Description |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| listPublicIpsAddresses | The following new response parameters are added: isportable, vmipaddress |
| enableStorageMaintenance | The following new response parameters are added: hypervisor, scope, suitableformigration |
| listLoadBalancerRules | The following new request parameters is added: networkkid (optional) The following new response parameters is added: networkkid |
| stopRouter | The following new response parameters are added: ip6dns1, ip6dns2, role |
| listClusters | The following new response parameters are added: cpuovercommitratio, memoryovercommitratio |
| attachVolume | The following new response parameter is added: displayvolume |
| updateVPCOffering | The following request parameters is made mandatory: id |
| resetSSHKeyForVirtualMachine | The following new request parameter is added: keypair (required) The following parameter is removed: name The following new response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| updateCluster | The following new request parameters are added: cpuovercommitratio (optional), memoryovercommitratio (optional) The following new response parameters are added: cpuovercommitratio, memoryovercommitratio |
| listPrivateGateways | The following new response parameters are added: aclid, sourcenatsupported |
| ldapConfig | The following new request parameters are added: listall (optional) The following parameters has been made optional: searchbase, hostname, queryfilter The following new response parameter is added: ssl |
| listTemplates | The following new response parameters are added: isdynamicallyscalable, sshkeyenabled |
| listNetworks | The following new response parameters are added: aclid, displaynetwork, ip6cidr, |

| API Commands | Description |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | ip6gateway, ispersistent, networkcidr, reservediprange |
| restartNetwork | The following new response parameters are added: isportable, vmipaddress |
| prepareTemplate | The following new response parameters are added: isdynamicallyscalable, sshkeyenabled |
| rebootVirtualMachine | The following new response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| changeServiceForRouter | The following new request parameters are added: aclid (optional), action (optional), protocol (optional) The following new response parameters are added: id, scope |
| updateZone | The following new request parameters are added: ip6dns1 (optional), ip6dns2 (optional) The following new response parameters are added: ip6dns1, ip6dns2 |
| ldapRemove | The following new response parameters are added: ssl |
| updateServiceOffering | The following new response parameters are added: deploymentplanner, isvolatile |
| updateStoragePool | The following new response parameters are added: hypervisor, scope, suitableformigration |
| listFirewallRules | The following request parameter is removed: traffictype The following new response parameters are added: networkid |
| updateUser | The following new response parameters are added: iscallerchilddomain, isdefault |
| updateProject | The following new response parameters are added: cpuavailable, cpulimit, cputotal, ipavailable, iplimit, iptotal, memoryavailable, memorylimit, memorytotal, networkavailable, networklimit, networktotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal, snapshotavailable, snapshotlimit, snapshottotal, templateavailable, templatelimit, templatetotal, vmavailable, vmlimit, vmrunning, vmstopped, vmtotal, volumeavailable, volumelimit, volumetotal, vpcavailable, vpcplimit, vpcptotal |

| API Commands | Description |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| updateTemplate | <p>The following new request parameters are added: isdynamicallyscalable (optional), isrouting (optional)</p> <p>The following new response parameters are added: isdynamicallyscalable, sshkeyenabled</p> |
| disableUser | <p>The following new response parameters are added: iscallerchilddomain, isdefault</p> |
| activateProject | <p>The following new response parameters are added: cpuavailable, cpulimit, cputotal, ipavailable, iplimit, iptotal, memoryavailable, memorylimit, memorytotal, networkavailable, networklimit, networktotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal, snapshotavailable, snapshotlimit, snapshottotal, templateavailable, templatelimit, templatetotal, vmavailable, vmlimit, vmrunning, vmstopped, vmtotal, volumeavailable, volumelimit, volumetotal, vpcavailable, vpclimit, vpcotal</p> |
| createNetworkACL | <p>The following new request parameters are added: aclid (optional), action (optional), number (optional)</p> <p>The following request parameter is now optional: networkkid</p> <p>The following new response parameters are added: aclid, action, number</p> |
| enableStaticNat | <p>The following new request parameters are added: vmguestip (optional)</p> |
| registerIso | <p>The following new request parameters are added: imagestoreuuid (optional), isdynamicallyscalable (optional)</p> <p>The following new response parameters are added: isdynamicallyscalable, sshkeyenabled</p> |
| createIpForwardingRule | <p>The following new response parameter is added: vmguestip</p> |
| resetPasswordForVirtualMachine | <p>The following new response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup</p> |
| createVolume | <p>The following new request parameter is added: displayvolume (optional)</p> <p>The following new response parameter is added: displayvolume</p> |

| API Commands | Description |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| startRouter | The following new response parameters are added: ip6dns1, ip6dns2, role |
| listCapabilities | The following new response parameters are added: apilimitinterval and apilimitmax. |
| createServiceOffering | <p>The following new request parameters are added: deploymentplanner (optional), isvolatile (optional), serviceofferingdetails (optional).</p> <p>isvolatile indicates whether the service offering includes Volatile VM capability, which will discard the VM's root disk and create a new one on reboot.</p> <p>The following new response parameters are added: deploymentplanner, isvolatile</p> |
| restoreVirtualMachine | <p>The following request parameter is added: templateID (optional). This is used to point to the new template ID when the base image is updated.</p> <p>The following response parameters are added: diskiread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup</p> |
| createNetwork | <p>The following new request parameters are added: aclid (optional), displaynetwork (optional), endipv6 (optional), ip6cidr (optional), ip6gateway (optional), isolatedpvlan (optional), startipv6 (optional)</p> <p>The following new response parameters are added: aclid, displaynetwork, ip6cidr, ip6gateway, ispersistent, networkcidr, reservediprange</p> |
| createVlanIpRange | <p>The following new request parameters are added: startipv6, endipv6, ip6gateway, ip6cidr</p> <p>Changed parameters: startip (is now optional)</p> <p>The following new response parameters are added: startipv6, endipv6, ip6gateway, ip6cidr</p> |
| CreateZone | <p>The following new request parameters are added: ip6dns1, ip6dns2</p> <p>The following new response parameters are added: ip6dns1, ip6dns2</p> |
| deployVirtualMachine | The following request parameters are added: affinitygroupids (optional), affinitygroupnames (optional), displayvm (optional), ip6address (optional) |

| API Commands | Description |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <p>The following request parameter is modified: iptonetworklist has a new possible value, ipv6</p> <p>The following new response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup</p> |
| createNetworkOffering | <p>The following request parameters are added: details (optional), egressdefaultpolicy (optional), ispersistent (optional)</p> <p>ispersistent determines if the network or network offering created or listed by using this offering are persistent or not.</p> <p>The following response parameters are added: details, egressdefaultpolicy, ispersistent</p> |
| listNetworks | <p>The following request parameters is added: isPersistent.</p> <p>This parameter determines if the network or network offering created or listed by using this offering are persistent or not.</p> |
| listNetworkOfferings | <p>The following request parameters is added: isPersistent.</p> <p>This parameter determines if the network or network offering created or listed by using this offering are persistent or not.</p> <p>For listNetworkOfferings, the following response parameter has been added: details, egressdefaultpolicy, ispersistent</p> |
| addF5LoadBalancer configureNetscalerLoadBalancer addNetscalerLoadBalancer listF5LoadBalancers configureF5LoadBalancer listNetscalerLoadBalancers | <p>The following response parameter is removed: inline.</p> |
| listRouters | <p>For nic responses, the following fields have been added.</p> <ul style="list-style-type: none"> • ip6address • ip6gateway • ip6cidr |

| API Commands | Description |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| listVirtualMachines | <p>The following request parameters are added: affinitygroupid (optional), vpcid (optional)</p> <p>The following response parameters are added: diskiread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup</p> |
| listRouters listZones | <p>For DomainRouter and DataCenter response, the following fields have been added.</p> <ul style="list-style-type: none"> • ip6dns1 • ip6dns2 <p>For listZones, the following optional request parameters are added: name, networktype</p> |
| listFirewallRules createFirewallRule | <p>The following request parameter is added: traffictype (optional).</p> <p>The following response parameter is added: networkid</p> |
| listUsageRecords | <p>The following response parameter is added: virtualsize.</p> |
| deletelso | <p>The following request parameter is removed: forced</p> |
| addCluster | <p>The following request parameters are added: cpuovercommitratio (optional), guestvswitchtype (optional), guestvswitchtype (optional), memoryovercommitratio (optional), publicvswitchtype (optional), publicvswitchtype (optional)</p> <p>The following request parameters are added: cpuovercommitratio, memoryovercommitratio</p> |
| updateCluster | <p>The following request parameters are added: cpuovercommitratio, ramovercommitratio</p> |
| createStoragePool | <p>The following request parameters are added: hypervisor (optional), provider (optional), scope (optional)</p> <p>The following request parameters have been made mandatory: podid, clusterid</p> <p>The following response parameter has been added: hypervisor, scope, suitableformigration</p> |
| listStoragePools | <p>The following request parameter is added: scope (optional)</p> <p>The following response parameters are added: hypervisor, scope, suitableformigration</p> |

| API Commands | Description |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| updateDiskOffering | The following response parameter is added: displayoffering |
| changeServiceForVirtualMachine | The following response parameter are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| recoverVirtualMachine | The following response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| listCapabilities | The following response parameters are added: apilimitinterval, apilimitmax |
| createRemoteAccessVpn | The following response parameters are added: id |
| startVirtualMachine | The following response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| detachIso | The following response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| updateVPC | The following request parameters has been made mandatory: id, name |
| associateIpAddress | The following request parameters are added: isportable (optional), regionid (optional) The following response parameters are added: isportable, vmipaddress |
| listProjectAccounts | The following response parameters are added: cpuavailable, cpulimit, cputotal, ipavailable, iplimit, iptotal, memoryavailable, memorylimit, memorytotal, networkavailable, networklimit, networktotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal, snapshotavailable, snapshotlimit, snapshottotal, templateavailable, templatelimit, templatetotal, vmavailable, vmlimit, vmrunning, vmstopped, vmtotal, volumeavailable, volumelimit, volumetotal, vpcavailable, vpclimit, vpcctotal |
| disableAccount | The following response parameters are added: cpuavailable, cpulimit, cputotal, isdefault, memoryavailable, memorylimit, memorytotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal |

| API Commands | Description |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| listPortForwardingRules | The following response parameters are added: vmguestip |
| migrateVirtualMachine | The following response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| cancelStorageMaintenance | The following response parameters are added: hypervisor, scope, suitableformigration |
| createPortForwardingRule | The following request parameter is added: vmguestip (optional) The following response parameter is added: vmguestip |
| addVpnUser | The following response parameter is added: state |
| createVPCOffering | The following request parameter is added: serviceproviderlist (optional) |
| assignVirtualMachine | The following response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| listConditions | The following response parameters are added: account, counter, domain, domainid, project, projectid, relationaloperator, threshold Removed response parameters: name, source, value |
| createPrivateGateway | The following request parameters are added: aclid (optional), sourcenatsupported (optional) The following response parameters are added: aclid, sourcenatsupported |
| updateVirtualMachine | The following request parameters are added: displayvm (optional), isdynamicallyscalable (optional) The following response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| destroyRouter | The following response parameters are added: ip6dns1, ip6dns2, role |
| listServiceOfferings | The following response parameters are added: deploymentplanner, isvolatile |
| listUsageRecords | The following response parameters are removed: virtualsize |
| createProject | The following response parameters are added: cpuavailable, cpulimit, |

| API Commands | Description |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <p>cputotal, ipavailable, iplimit, iptotal, memoryavailable, memorylimit, memorytotal, networkavailable, networklimit, networktotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal, snapshotavailable, snapshotlimit, snapshottotal, templateavailable, templatelimit, templatetotal, vmavailable, vmlimit, vmrunning, vmstopped, vmtotal, volumeavailable, volumelimit, volumetotal, vpcavailable, vpclimit, vpctotal</p> |
| enableUser | <p>The following response parameters are added: iscallerchilddomain, isdefault</p> |
| createLoadBalancerRule | <p>The following response parameter is added: networkid</p> |
| updateAccount | <p>The following response parameters are added: cpuavailable, cpulimit, cputotal, isdefault, memoryavailable, memorylimit, memorytotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal</p> |
| copyIso | <p>The following response parameters are added: isdynamicallyscalable, sshkeyenabled</p> |
| uploadVolume | <p>The following request parameters are added: imagestoreuuid (optional), projectid (optional)</p> <p>The following response parameters are added: displayvolume</p> |
| createDomain | <p>The following request parameter is added: domainid (optional)</p> |
| stopVirtualMachine | <p>The following response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup</p> |
| listAccounts | <p>The following response parameters are added: cpuavailable, cpulimit, cputotal, isdefault, memoryavailable, memorylimit, memorytotal, primarystorageavailable, primarystoragelimit, primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal</p> |
| createSnapshot | <p>The following response parameter is added: zoneid</p> |
| updateIso | <p>The following request parameters are added: isdynamicallyscalable (optional), isrouting (optional)</p> <p>The following response parameters are added: isdynamicallyscalable, sshkeyenabled</p> |

| API Commands | Description |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| listIpForwardingRules | The following response parameter is added: vmguestip |
| updateNetwork | The following request parameters are added: displaynetwork (optional), guestvmcidr (optional) The following response parameters are added: aclid, displaynetwork, ip6cidr, ip6gateway, ispersistent, networkcidr, reservediprange |
| destroyVirtualMachine | The following response parameters are added: diskioread, diskiowrite, diskkbsread, diskkbswrite, displayvm, isdynamicallyscalable, affinitygroup |
| createDiskOffering | The following request parameter is added: displayoffering (optional) The following response parameter is added: displayoffering |
| rebootRouter | The following response parameters are added: ip6dns1, ip6dns2, role |
| listConfigurations | The following request parameters are added: accountid (optional), clusterid (optional), storageid (optional), zoneid (optional) The following response parameters are added: id, scope |
| createUser | The following request parameter is added: userid (optional) The following response parameters are added: iscallerchilddomain, isdefault |
| listDiskOfferings | The following response parameter is added: displayoffering |
| detachVolume | The following response parameter is added: displayvolume |
| deleteUser | The following response parameters are added: displaytext, success Removed parameters: id, account, accountid, accounttype, apikey, created, domain, domainid, email, firstname, lastname, secretkey, state, timezone, username |
| listSnapshots | The following request parameter is added: zoneid (optional) The following response parameter is added: zoneid |
| markDefaultZoneForAccount | The following response parameters are added: cpuavailable, cpulimit, cputotal, isdefault, memoryavailable, memorylimit, memorytotal, primarystorageavailable, primarystoragelimit, |

| API Commands | Description |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | primarystoragetotal, secondarystorageavailable, secondarystoragelimit, secondarystoragetotal |
| restartVPC | The following request parameters are made mandatory: id |
| updateHypervisorCapabilities | The following response parameters are added: hypervisor, hypervisorversion, maxdatavolumeslimit, maxguestslimit, maxhostspercluster, securitygroupenabled, storagemotionenabled Removed parameters: cpunumber, cpuspeed, created, defaultuse, displaytext, domain, domainid, hosttags, issystem, limitcpuuse, memory, name, networkrate, offerha, storagetype, systemvmtype, tags |
| updateLoadBalancerRule | The following response parameter is added: networkid |
| listVlanIpRanges | The following response parameters are added: endipv6, ip6cidr, ip6gateway, startipv6 |
| listHypervisorCapabilities | The following response parameters are added: maxdatavolumeslimit, maxhostspercluster, storagemotionenabled |
| updateNetworkOffering | The following response parameters are added: details, egressdefaultpolicy, ispersistent |
| createVirtualRouterElement | The following request parameters are added: providertype (optional) |
| listVpnUsers | The following response parameter is added: state |
| listUsers | The following response parameters are added: iscallerchilddomain, isdefault |
| listSupportedNetworkServices | The following response parameter is added: provider |
| listIos | The following response parameters are added: isdynamicallyscalable, sshkeyenabled |

9.3. Deprecated APIs

- addExternalLoadBalancer (Adds F5 external load balancer appliance.)
- deleteExternalLoadBalancer (Deletes a F5 external load balancer appliance added in a zone.)
- listExternalLoadBalancers (Lists F5 external load balancer appliances added in a zone.)

