Citrix CloudPlatform (powered by Apache CloudStack) version 3.0.4 Release Notes

Revised July 18, 2012 5:48 pm Pacific

Citrix CloudPlatform (powered by Apache CloudStack) version 3.0.4 Release Notes
Revised July 18, 2012 5:48 pm Pacific
Edition 1

Release notes for the patch release to support the XenServer hotfixes XS602E003, XS602E004, and XS602E005.

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/

Upgrade Instructions

2.1. Upgrade from 3.0.x to 3.0.4

Perform the following to upgrade from version 3.0.0, 3.0.1, 3.0.2, or 3.0.3 to version 3.0.4.

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. See *bug CS-8222*[1]). 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. Stop all Usage Servers if running. Run this on all Usage Server hosts.

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

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

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

4. 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
   ```

5. Download CloudPlatform 3.0.4 onto 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*[2].

6. Upgrade the CloudPlatform packages. You should have a file in the form of “CloudStack-3.0.4-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 CloudStack-3.0.4-N-OSVERSION.tar.gz
   # cd CloudStack-3.0.4-N-OSVERSION
   # ./install.sh
   ```

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

7. Choose "U" to upgrade the package

[1] http://bugs.cloudstack.org/browse/CS-8222
[2] http://www.citrix.com/lang/English/publicindex.asp?destURL=%2FEnglish%2FmyCitrix%2Findex.asp%3F#

```
>U
```

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

8. If you have made changes to your existing copy of the file components.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 3.0.4.

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 components.xml, and you need to merge the two files:

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

   a. Make a backup copy of your /etc/cloud/management/components.xml file. For example:

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

   b. Copy /etc/cloud/management/components.xml.rpmnew to create a new /etc/cloud/management/components.xml:

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

   c. Merge your changes from the backup file into the new components.xml file.

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

9. Repeat steps 5 - 8 on each management server node.

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

```
# service cloud-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. Having the Management Server restarted without any issues indicates that the upgrade is successfully completed.

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

```
# service cloud-usage start
```

12. (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.

    a. Copy the CloudPlatform 3.0.4 tar file to the host, untar it, and change directory to the resulting directory.

    b. Stop the running agent.

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

    c. Update the agent software.

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

    d. Choose "U" to update the packages.

    e. Start the agent.

```
# service cloud-agent start
```

13. 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.

14. 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 cloud-sysvmadm script at the end of the upgrade procedure. For information about how to set this parameter, see “Edit the Global Configuration Settings” 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

15. Run the following script to stop, then start, all Secondary Storage VMs, Console Proxy VMs, and virtual routers. Run the script once on one management server. The script requires the 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 cloud-sysvmadm -d 192.168.1.5 -u cloud -p password -a > sysvm.log 2>&1 &
# tail -f sysvm.log
```

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

16. In order to deploy AWS API on its new port (7080), you need to deploy it under a separate webapps folder and make some changes to port settings.

a. Create the new webapps folder:

```
# mkdir -p /usr/share/cloud/management/webapps7080
```

b. Create a symbolic link:

```
# ln -s /usr/share/cloud/bridge/webapps/awsapi /usr/share/cloud/management/
webapps7080/awsapi
```

c. Remove the old folder:

```
# rm /usr/share/cloud/management/webapps/awsapi
```

d. Open port 7080:

```
# iptables -I INPUT -p tcp -m tcp --dport 7080 -j ACCEPT
```

e. If you have made any modifications in server.xml on your existing CloudPlatform installation, back it up:

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

Then replace with the new server.xml file:

```
# cp /etc/cloud/management/server.xml.rpmnew /etc/cloud/management/server.xml
```

Merge any changes from the backup file into the new server.xml file.

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

f. If you have made any modifications in server-nonssl.xml on your existing CloudPlatform installation, back it up:

```
# mv /etc/cloud/management/server-nonssl.xml /etc/cloud/management/server-nonssl.xml-backup
```

Then replace with the new server-nonssl.xml file:

```
# cp /etc/cloud/management/server-nonssl.xml.rpmnew /etc/cloud/management/server-nonssl.xml
```

Merge any changes from the backup file into the new server-nonssl.xml file.

```
# vi /etc/cloud/management/server-nonssl.xml
```

g. If you have used SSL authentication, and made any modifications in server-ssl.xml on your existing CloudPlatform installation, back it up:

```
# mv /etc/cloud/management/server-ssl.xml /etc/cloud/management/server-ssl.xml-backup
```

Then replace with the new server-ssl.xml file:

```
# cp /etc/cloud/management/server-ssl.xml.rpmnew /etc/cloud/management/server-ssl.xml
```

Merge any changes from the backup file into the new server-ssl.xml file.

```
# vi /etc/cloud/management/server-ssl.xml
```

h. Restart the Management Server to put the new settings into effect.

17. If needed, upgrade all Citrix XenServer hypervisor hosts in your cloud to a version supported by CloudPlatform 3.0.4. The supported versions are XenServer 5.6 SP2 and 6.0.2. Instructions for upgrade can be found in the CloudPlatform 3.0.3 - 3.0.4 Installation Guide.

18. Now apply the XenServer hotfix to the XenServer v6.0.2 hypervisor hosts. (Support for new hotfixes XS602E003, XS602E004, and XS602E005 is the reason for release 3.0.4.)

a. Disconnect the XenServer cluster from CloudPlatform.

In the left navigation bar of the CloudPlatform UI, select Infrastructure. Under Clusters, click View All. Select the XenServer cluster and click Actions - Unmanage.

This may fail if there are hosts not in one of the states Up, Down, Disconnected, or Alert. You may need to fix that before unmanaging this cluster.

Wait until the status of the cluster has reached Unmanaged. Use the CloudPlatform UI to check on the status. When the cluster is in the unmanaged state, there is no connection to the hosts in the cluster.

b. To clean up the VLAN, log in to one XenServer host and run:

```
/opt/xensource/bin/cloud-clean-vlan.sh
```

c. Now prepare the upgrade by running the following on one XenServer host:

```
/opt/xensource/bin/cloud-prepare-upgrade.sh
```

If you see a message like "can't eject CD", log in to the VM and umount the CD, then run this script again.

d. Upload the hotfix to the XenServer hosts. Always start with the Xen pool master, then the slaves. Using your favorite file copy utility (e.g. WinSCP), copy the hotfixes to the host. Place them in a temporary folder such as /root or /tmp.

On the Xen pool master, upload the hotfix with this command:

```
xe patch-upload file-name=XS602E003.xsupdate
```

Make a note of the output from this command, which is a UUID for the hotfix file. You'll need it in another step later.

Note

(Optional) If you are applying other hotfixes as well, you can repeat the commands in this section with the appropriate hotfix number: XS602E004.xsupdate and XS602E005.xsupdate.

e. Manually live migrate all VMs on this host to another host. First, get a list of the VMs on this host:

```
# xe vm-list
```

Then use this command to migrate each VM. Replace the example host name and VM name with your own:

```
# xe vm-migrate live=true host=<host-name> vm=<VM-name>
```

Troubleshooting: If you see a message like "You attempted an operation on a VM which requires PV drivers to be installed but the drivers were not detected. vm: b6cf79c8-02ee-050b-922f-49583d9f1a14 (i-2-8-VM)," run /opt/xensource/bin/make_migratable.sh b6cf79c8-02ee-050b-922f-49583d9f1a14.

f. Apply the hotfix. First, get the UUID of this host:

```
# xe host-list
```

Then use the following command to apply the hotfix. Replace the example host UUID with the current host ID, and replace the hotfix UUID with the output from the patch-upload command you ran on this machine earlier. You can also get the hotfix UUID by running xe patch-list.

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

g. Copy the following files from the CloudPlatform Management Server to the host.

Copy from here...	...to here
/usr/lib64/cloud/agent/scripts/vm/hypervisor/xenserver/xenserver60/NFSSR.py	/opt/xensource/sm/NFSSR.py
/usr/lib64/cloud/agent/scripts/vm/hypervisor/xenserver/setupxenserver.sh	/opt/xensource/bin/setupxenserver.sh
/usr/lib64/cloud/agent/scripts/vm/hypervisor/xenserver/make_migratable.sh	/opt/xensource/bin/make_migratable.sh

h. Reboot this XenServer host.

i. Run the following:

```
/opt/xensource/bin/setupxenserver.sh
```

Note

If the message "mv: cannot stat `/etc/cron.daily/logrotate': No such file or directory" appears, you can safely ignore it.

j. Run the following:

```
for pbd in `xe pbd-list currently-attached=false| grep ^uuid | awk '{print $NF}'`; do
 xe pbd-plug uuid=$pbd ; done
```

k. On each slave host in the Xen pool, repeat these steps, starting from "manually live migrate VMs."

l. Remove the host tags by running the following on a host:

```
for host in $(xe host-list | grep ^uuid | awk '{print $NF}') ; do xe host-param-clear
 uuid=$host param-name=tags; done;
```

m. Connect the cluster to CloudPlatform.

   a. In the CloudPlatform UI, click Manage in the action list of the cluster.

CloudPlatform starts to connect to the hosts in this cluster. It might take several minutes to reconnect to the hosts.

Note

Upgrade all hosts to the same XenServer version before you connect this cluster to CloudPlatform. XenServer downgrade is not supported.

n. After all the hosts in this cluster are up, remove the VLAN in this cluster. Run the following on a host:

```
# /opt/xensource/bin/cloud-clean-vlan.sh
```

2.2. Upgrade from 2.2.x to 3.0.4

1. Ensure that you query your IPaddress 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. See *CS-8222*[3]. 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 3.0.4, 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.

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 3.0.4.

3. While running the 2.2.14 system, 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

   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.

[3] http://bugs.cloudstack.org/browse/CS-8222

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-3.0.0 Description: systemvm-xenserver-3.0.0 URL: http://download.cloud.com/templates/acton/acton-systemvm-02062012.vhd.bz2 Zone: Choose the zone where this hypervisor is used Hypervisor: XenServer Format: VHD OS Type: Debian GNU/Linux 5.0 (32-bit) Extractable: no Password Enabled: no Public: no Featured: no
KVM	Name: systemvm-kvm-3.0.0 Description: systemvm-kvm-3.0.0 URL: http://download.cloud.com/templates/acton/acton-systemvm-02062012.qcow2.bz2 Zone: Choose the zone where this hypervisor is used Hypervisor: KVM Format: QCOW2 OS Type: Debian GNU/Linux 5.0 (32-bit) Extractable: no Password Enabled: no Public: no Featured: no
VMware	Name: systemvm-vmware-3.0.0 Description: systemvm-vmware-3.0.0 URL: http://download.cloud.com/templates/acton/acton-systemvm-02062012.ova Zone: Choose the zone where this hypervisor is used

Hypervisor	Description
	Hypervisor: VMware Format: OVA OS Type: Debian GNU/Linux 5.0 (32-bit) Extractable: no Password Enabled: no Public: no Featured: no

5. Watch the screen to be sure that the template downloads successfully and enters the READY state. Do not proceed until this is successful

6. **WARNING**: If you use more than one type of hypervisor in your cloud, be sure you have repeated these steps to download the system VM template for each hypervisor type. Otherwise, the upgrade will fail.

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

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

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

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

9. 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
```

10. Download CloudPlatform 3.0.4 onto 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*[4].

11. Upgrade the CloudPlatform packages. You should have a file in the form of “CloudStack-3.0.4-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 CloudStack-3.0.4-N-OSVERSION.tar.gz
```

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

```
# cd CloudStack-3.0.4-N-OSVERSION
# ./install.sh
```

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

12. Choose "U" to upgrade the package

```
>U
```

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

13. If you have made changes to your existing copy of the file components.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 3.0.4.

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 components.xml, and you need to merge the two files:

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

a. Make a backup copy of your /etc/cloud/management/components.xml file. For example:

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

b. Copy /etc/cloud/management/components.xml.rpmnew to create a new /etc/cloud/management/components.xml:

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

c. Merge your changes from the backup file into the new components.xml file.

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

14. If you have made changes to your existing copy of the /etc/cloud/management/db.properties file 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 3.0.4.

a. Make a backup copy of your file /etc/cloud/management/db.properties. For example:

```
# mv /etc/cloud/management/db.properties /etc/cloud/management/db.properties-backup
```

b. Copy /etc/cloud/management/db.properties.rpmnew to create a new /etc/cloud/management/db.properties:

```
# cp -ap /etc/cloud/management/db.properties.rpmnew etc/cloud/management/
db.properties
```

c. Merge your changes from the backup file into the new db.properties file.

```
# vi /etc/cloud/management/db.properties
```

15. 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.

```
# cloud-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
- (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.

16. Repeat steps 10 - 15 on every management server node. If you provided your own encryption key in step 15, use the same key on all other management servers.

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

```
# service cloud-management start
```

Wait until the databases are upgraded. Ensure that the database upgrade is complete. You should see a message like "Complete! Done." After confirmation, start the other Management Servers one at a time by running the same command on each node.

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

```
# service cloud-usage start
```

19. (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.

a. Copy the CloudStack 3.0.4 .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. Start the agent.

```
# service cloud-agent start
```

20. 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.

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

    a. Run the command once on one management server. Provide the IP address of the MySQL instance, the MySQL user name, and the database password for that user. In addition to those parameters, provide the "-a" argument. For example:

    ```
    # nohup cloud-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).
    ```

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

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

XenServer or KVM:

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

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

```
# ssh -i <private-key-path> <link-local-ip> -p 3922
# cat /etc/cloudstack-release
```

The output should be like the following:

```
Cloudstack Release 3.0 Mon Feb 6 15:10:04 PST 2012
```

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 <private-key-path> <private-ip> -p 3922
# cat /etc/cloudstack-release
```

The output should be like the following:

```
Cloudstack Release 3.0 Mon Feb 6 15:10:04 PST 2012
```

23. In order to deploy AWS API on its new port (7080), you need to deploy it under a separate webapps folder.

    a. Create the new webapps folder:

    ```
    # mkdir -p /usr/share/cloud/management/webapps7080
    ```

    b. Create a symbolic link:

    ```
    # ln -s /usr/share/cloud/bridge/webapps/awsapi /usr/share/cloud/management/
    webapps7080/awsapi
    ```

    c. Remove the old folder:

    ```
    # rm /usr/share/cloud/management/webapps/awsapi
    ```

    d. Open port 7080:

    ```
    # iptables -I INPUT -p tcp -m tcp --dport 7080 -j ACCEPT
    ```

e. If you have made any modifications in server.xml on your existing CloudPlatform installation, back it up:

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

Then replace with the new server.xml file:

```
# cp /etc/cloud/management/server.xml.rpmnew /etc/cloud/management/server.xml
```

Merge any changes from the backup file into the new server.xml file.

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

f. If you have made any modifications in server-nonssl.xml on your existing CloudPlatform installation, back it up:

```
# mv /etc/cloud/management/server-nonssl.xml /etc/cloud/management/server-nonssl.xml-
backup
```

Then replace with the new server-nonssl.xml file:

```
# cp /etc/cloud/management/server-nonssl.xml.rpmnew /etc/cloud/management/server-
nonssl.xml
```

Merge any changes from the backup file into the new server-nonssl.xml file.

```
# vi /etc/cloud/management/server-nonssl.xml
```

g. If you have used SSL authentication, and made any modifications in server-ssl.xml on your existing CloudPlatform installation, back it up:

```
# mv /etc/cloud/management/server-ssl.xml /etc/cloud/management/server-ssl.xml-backup
```

Then replace with the new server-ssl.xml file:

```
# cp /etc/cloud/management/server-ssl.xml.rpmnew /etc/cloud/management/server-ssl.xml
```

Merge any changes from the backup file into the new server-ssl.xml file.

```
# vi /etc/cloud/management/server-ssl.xml
```

h. Restart the Management Server to put the new settings into effect.

24. If needed, upgrade all Citrix XenServer hypervisor hosts in your cloud to a version supported by CloudPlatform 3.0.4. The supported versions are XenServer 5.6 SP2 and 6.0.2. Instructions for upgrade can be found in the CloudPlatform 3.0.3 - 3.0.4 Installation Guide.

25. Now apply the XenServer hotfix to XenServer v6.0.2 hypervisor hosts. (Support for new hotfixes XS602E003, XS602E004, and XS602E005 is the reason for release 3.0.4.)

    a. Disconnect the XenServer cluster from CloudPlatform.

In the left navigation bar of the CloudPlatform UI, select Infrastructure. Under Clusters, click View All. Select the XenServer cluster and click Actions - Unmanage.

This may fail if there are hosts not in one of the states Up, Down, Disconnected, or Alert. You may need to fix that before unmanaging this cluster.

Wait until the status of the cluster has reached Unmanaged. Use the CloudPlatform UI to check on the status. When the cluster is in the unmanaged state, there is no connection to the hosts in the cluster.

b. To clean up the VLAN, log in to one XenServer host and run:

```
/opt/xensource/bin/cloud-clean-vlan.sh
```

c. Now prepare the upgrade by running the following on one XenServer host:

```
/opt/xensource/bin/cloud-prepare-upgrade.sh
```

If you see a message like "can't eject CD", log in to the VM and umount the CD, then run this script again.

d. Upload the hotfix to the XenServer hosts. Always start with the Xen pool master, then the slaves. Using your favorite file copy utility (e.g. WinSCP), copy the hotfixes to the host. Place them in a temporary folder such as /root or /tmp.

On the Xen pool master, upload the hotfix with this command:

```
xe patch-upload file-name=XS602E003.xsupdate
```

Make a note of the output from this command, which is a UUID for the hotfix file. You'll need it in another step later.

Note

(Optional) If you are applying other hotfixes as well, you can repeat the commands in this section with the appropriate hotfix number: XS602E004.xsupdate and XS602E005.xsupdate.

e. Manually live migrate all VMs on this host to another host. First, get a list of the VMs on this host:

```
# xe vm-list
```

Then use this command to migrate each VM. Replace the example host name and VM name with your own:

```
# xe vm-migrate live=true host=<host-name> vm=<VM-name>
```

Troubleshooting: If you see a message like "You attempted an operation on a VM which requires PV drivers to be installed but the drivers were not detected.vm:

b6cf79c8-02ee-050b-922f-49583d9f1a14 (i-2-8-VM)," run /opt/xensource/bin/make_migratable.sh b6cf79c8-02ee-050b-922f-49583d9f1a14.

f. Apply the hotfix. First, get the UUID of this host:

```
# xe host-list
```

Then use the following command to apply the hotfix. Replace the example host UUID with the current host ID, and replace the hotfix UUID with the output from the patch-upload command you ran on this machine earlier. You can also get the hotfix UUID by running xe patch-list.

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

g. Copy the following files from the CloudPlatform Management Server to the host.

Copy from here...	...to here
/usr/lib64/cloud/agent/scripts/vm/hypervisor/xenserver/xenserver60/NFSSR.py	/opt/xensource/sm/NFSSR.py
/usr/lib64/cloud/agent/scripts/vm/hypervisor/xenserver/setupxenserver.sh	/opt/xensource/bin/setupxenserver.sh
/usr/lib64/cloud/agent/scripts/vm/hypervisor/xenserver/make_migratable.sh	/opt/xensource/bin/make_migratable.sh

h. Reboot this XenServer host.

i. Run the following:

```
/opt/xensource/bin/setupxenserver.sh
```

Note

If the message "mv: cannot stat `/etc/cron.daily/logrotate': No such file or directory" appears, you can safely ignore it.

j. Run the following:

```
for pbd in `xe pbd-list currently-attached=false| grep ^uuid | awk '{print $NF}'`; do
 xe pbd-plug uuid=$pbd ; done
```

k. On each slave host in the Xen pool, repeat these steps, starting from "manually live migrate VMs."

l. Remove the host tags by running the following on a host:

```
for host in $(xe host-list | grep ^uuid | awk '{print $NF}') ; do xe host-param-clear
 uuid=$host param-name=tags; done;
```

m. Connect the cluster to CloudPlatform.

   a. In the CloudPlatform UI, click Manage in the action list of the cluster.

      CloudPlatform starts to connect to the hosts in this cluster. It might take several minutes to reconnect to the hosts.

Note

Upgrade all hosts to the same XenServer version before you connect this cluster to CloudPlatform. XenServer downgrade is not supported.

n. After all the hosts are up in this cluster, remove the VLAN in this cluster. Run the following on a host:

```
# /opt/xensource/bin/cloud-clean-vlan.sh
```

2.3. Upgrade from 2.1.x to 3.0.4

Direct upgrades from version 2.1.0 - 2.1.10 to 3.0.4 are not supported. It 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 version 2.2.14 Release Notes.

Version 3.0.4

3.1. What's New in 3.0.4

CloudPlatform 3.0.4 is the first maintenance patch for CloudPlatform 3.0.3. CloudPlatform 3.0.4 supports the XenServer patch XS602E003, XS602E004, and XS602E005. This release includes no new features. For a list of the major fixed items, see Issues Fixed in 3.0.4.

3.2. Issues Fixed in 3.0.4

Issue ID	Description
CS-15300, CS-15340	The admin accounts of a domain now honour the limits imposed on that domain just like the regular accounts do. A domain admin now is not allowed to create an unlimited number of instances, volumes, snapshots, and so on.
CS-15323, CS-15524	CloudPlatform supports the following Citrix XenServer hotfixes: XS602E003, XS602E004, and XS602E005.
CS-15376, CS-15373	The AWS APIs (EC2 and S3) now listen on the 7080 port and send request to CloudPlatform on the 8080 port just as any other clients of CloudPlatform.
CS-15382	During 2.2.14 to 3.0.4 upgrade, the hosts no longer go to the Alert state if destroyed networks existed with non-existent tags prior to upgrade.
CS-15396	The CloudPlatform database now contain the UUD information after the 2.2.14 to 3.0.4 upgrade.
CS-15404	For both fresh installations and the 2.2.14 to 3.0.4 upgrade, physical network id now starts from 200.
CS-15406	After the 2.214 to 3.0.14 upgrade, the hosts no longer go to the Alert state if the xen.guest.network.device parameter is set to null.
CS-15414	After the 2.2.14 to 3.0.4 upgrade, the value of the global parameter xen.guest.network.device is now decrypted before setting the traffic label.
CS-15429	While creating an instance with data volume, disk offering also is considered while checking the account limit on volume resources.
CS-15430	Create snapshot now fails gracefully if creating a snapshot exceeds the snapshot resource limit for a domain admin or a user account.
CS-15441	No MySQL exception is observed in the usage.log when starting the Usage server.
CS-15444, CS-15443	Adding a new firewall rule no longer throws any error.
CS-15448	Permission issue on the AWS API is fixed. User registration works as expected.
CS-15449	Running cloudstack-aws-api-register no longer fails with the "User registration failed with error: [Errno 113] No route to host" error.
CS-15450	Upgrade from 2.2.14 to 3.0.4 no longer fails on a VMware host.

Issue ID	Description
CS-15455	The iptable rules are configured to open the awsapi port (7080) as part of the installation.
CS-15475	The Add GuestNetwork tab in Network is now correctly displayed in a multi-zone basic/advance setup.
CS-15494	The VPN users can be successfully added to the Advanced zone.
CS-15495	Upgrade from 2.2.14 to 3.0.4 no longer fails with CloudRuntimeException.
CS-15516	Upgrade from 2.2.14 to 3.0.4 no longer fails while adding the physical network information.
CS-15535, CS-13944	The CloudPlatform 2.2.x to 3.0.x database upgrade for multiple physical networks is now supported.
CS-15578	Upgrade from 2.2.14 to 3.04 no longer fails if the setup has a zone deleted and the corresponding router state removed.

3.3. Known Issues in 3.0.4

Issue ID	Description

CS-15407	After the 2.2.14 to 3.0.4 upgrade, VLAN allocation on multiple physical networks does not happen as expected.
	
	To workaround this issue, follow the instructions given below:
	
	1. Revert to your 2.2.14 setup.
	
	2. Stop all the VMs with the isolated virtual networks in your cloud setup.
	
	3. Run following query to find if any networks still have the NICs allocated:
	
	a. Check if any virtual guest networks have the NICs allocated:
	
	`#SELECT DISTINCT op.id from `cloud`.`op_networks` op JOIN `cloud`.`networks` n on op.id=n.id WHERE nics_count != 0 AND guest_type = 'Virtual';`
	
	b. If this returns any network IDs, then ensure the following:
	
	i. All the VMs are stopped.
	
	ii. No new VM is started.
	
	iii. Shutdown the Management Server.
	
	c. Remove the NICs count for the virtual network IDs returned in step (a), and set the NIC count to 0:
	
	`UPDATE `cloud`.`op_networks` SET nics_count = 0 WHERE id = <enter id of virtual network>`

Issue ID	Description
	d. Restart the Management Server, and wait for all the networks to shut down.  **Note** Networks shutdown is determined by the network.gc.interval and network.gc.wait parameters. 4. Ensure that all the networks are shut down and all the guest VNETs are free. 5. Run the upgrade script. This allocates all your guest VNET ranges to the first physical network. 6. By using the updatePhysicalNetwork API, reconfigure the VNET ranges for each physical network as desired. 7. Start all the VMs.
CS-14680	CloudPlatform and LDAP user validation cannot happen simultaneously because the user password is hashed and stored in the database, and LDAP requires the passwords in plain text. To work with the LDAP user, the MD5 hash should be disabled in the login process by commenting the md5HashedLogin variable in sharedFunctions.js file available at /usr/share/cloud/management/webapps/client/scripts, and restart the cloud-management service. `var md5HashedLogin = false;` However, if md5HashedLogin is set to false, the end user can login with the LDAP credentials but not with the CloudPlatform user credentials.
CS-15476	The 2.2.14 to 3.0.4 upgrade fails if multiple untagged physical networks exist before the upgrade.

Version 3.0.3

4.1. What's New in 3.0.3

4.1.1. Secure Console Access on XenServer

With the addition of Secure Console feature, users can now securely access the VM consoles on the XenServer hypervisor. You can either SSH or use the View Console option in the Management Server to securely connect to the VMs on the XenServer host. The Management Server uses the xapi API to stream the VM consoles. However, there is no change in the way you can access the console of a VM. This feature is supported on XenServer 5.6 and 6.0 versions.

4.1.2. Stopped VM

This release supports creating VMs without starting them on the backend. You can determine whether the VM needs to be started as part of the VM deployment. A VM can be deployed in two ways: create and start a VM (the default method); create a VM and leave it in the stopped state.

A new request parameter, startVM, is introduced in the deployVm API to support the stopped VM feature. The possible values are:

- true - The VM starts as a part of the VM deployment
- false - The VM is left in stopped state at the end of the VM deployment

4.1.3. Uploading an Existing Volume to a Virtual Machine

Existing data can now be made accessible to a virtual machine. This is called uploading a volume to the VM. For example, this is useful to upload data from a local file system and attach it to a VM. Root administrators, domain administrators, and end users can all upload existing volumes to VMs. The upload is performed by using HTTP. The uploaded volume is placed in the zone's secondary storage.

This functionality is supported for the following hypervisors:

- Hypervisor : Disk Image Format
- XenServer : VHD
- VMware : OVA
- KVM : QCOW2
- OVM : RAW

4.1.4. Dedicated High-Availability Hosts

One or more hosts can now be designated for use only by high-availability (HA) enabled VMs that are restarted due to a host failure. Setting up a pool of such dedicated HA hosts as the recovery destination for all HA-enabled VMs make it easier to determine which VMs are restarted as part of the high-availability function. You can designate a host as a dedicated-HA restart node only if the Dedicated HA Hosts feature is enabled by setting the appropriate global configuration parameter.

4.1.5. Support for Amazon Web Services API

This release supports Amazon Web Services APIs, including Elastic Compute Cloud (EC2) API. Fidelity with the EC2 API and the installation experience for this functionality are both enhanced. In prior releases, users were required to install a separate component called CloudBridge, in addition to installing the Management Server. For new installations of CloudPlatform 3.0.3, this software is installed automatically along with CloudPlatform and runs in a more closely integrated fashion. The feature is disabled by default, but can be easily enabled by setting the appropriate global configuration parameter and performing a few setup steps.

4.1.6. Support for Cisco Nexus 1000v Virtual Switch

In addition to standard vSwitch, Cisco Nexus 1000v is now supported for virtual network configuration in VMware vSphere deployments. With this, guest traffic isolation through VLANs is possible in VMware environment. Using Nexus virtual switch simplifies configuring and monitoring virtual networks that span across a large number of hosts, and facilitates live migration in VMWare-based cloud deployment.

4.2. Summary of New Features by Bug Number

Issue ID	Feature
CS-14385	Dedicated High Availability Host
CS-14378	Stopped VM
CS-10789	Uploading an Existing Volume to a Virtual Machine
CS-6061	Secure Console Access on XenServer
CS-9909	Cisco Nexus 1000v Support
CS-14435	Support for Amazon Web Services API

4.3. Issues Fixed in 3.0.3

Defects	Description
CS-14256	Virtual Router no longer remains in starting state for subdomain or user on a KVM 3.0.1 prerlease host on RHEL 6.2.
CS-7495	Implemented a variety of Xen management host improvements.
CS-8105	NFS v4 for primary storage now works as expected on KVM hosts.
CS-9989	The error messages returned during VM deployment failure will have much more details than before.
CS-12584	You can no longer add security groups not supported by the hypervisor in use.
CS-12705	When creating a Network offering by using SRX as the service provider for SourceNAT servcies, an option is given in the CloudPlatform UI now to set the source_nat type to "per Zone"/"per account".
CS-12782	Assigning a VM from Basic to Advanced zone no longer ignores the network ID. A warning message is displayed for VM movements across zones.
CS-12591	Broadcast Address on the Second Public IP NIC is now corrected.

Defects	Description
CS-13272	When a user is deleted, all the associated properties, such as IPs and virtual routers, are now deleted.
CS-13377	Creating template from a root disk of a stopped instance now provides an option to make it a "Featured template".
CS-13500	Reaching the first guest VM by using its public IP from the second guest VM no longer fails.
CS-13853	The default gateway can no longer be 0.0.0.0 in the Secondary Storage VM (SSVM).
CS-13863	The queryAsyncJobResult command in XML format now returns the correct UUIDs.
CS-13867	Corrected CSP xenserver-cloud-supp.tgz for XenServer 5.6 and 6.0.
CS-13904	Labels and values for the service offerings CPU and memory are now consistent.
CS-13998	The SSVM kernel panic issue is fixed on XenServer.
CS-14090	The issue is fixed where running the VMware snapshots randomly fails with the ArrayIndexOutOfBoundsException error.
CS-14021	The java.lang.OutOfMemoryError is fixed on the Management Server.
CS-14025	The Python Eggs are provided to easily package the test client for each branch of CloudPlatform.
CS-14068	Resetting the VM password through the CloudPlatform UI no longer causes any error.
CS-14156	The pod which has the administrator's virtual router is no longer selected while creating the virtual routers for guests.
CS-14182	The users can now delete their ISOs as normal users.
CS-14185	The listOSTypes API now filters out the types of operating system by using the keywords.
CS-14204	The cloud-setup-bonding.sh command no longer generates the "command not found" error.
CS-14214	The Specify VLAN option cannot be enabled now for an isolated Network offering with SourceNAT enabled.
CS-14234	Sending project invite email to an account now requires SMTP configured in CloudPlatform.
CS-14237	The garbage collector of the primary storage no longer fails when the first host in the cluster is not up.
CS-14241	Custom Volume Disk Offering is now matching the Global configuration value.
CS-14270	The listNetworks API no longer assumes that the broadcast type is always VLAN.
CS-14319	The internal name of the VM is no longer present in the error message that is displayed to a domain administrator.

Defects	Description
CS-14321	The listVolumes API call now returns a valid value for the isExtractable parameter for the ISO-derived disk and data disk volumes.
CS-14323	Invalid API calls will now give valid response in json/xml format.
CS-14339	Custom Disk Size will now allow values larger than 100GB.
CS-14357	The ConsoleProxyLoadReportCommand is no longer fired continuously.
CS-14421	Fixed the issue of virtual router deployments. The DHCP entries can now be assigned to the router.
CS-14555	Unzipped downloaded template MD5SUM will no longer override the zipped template MD5SUM in the database.
CS-14598	The complete screen of the running VM is now displayed in the console proxy.
CS-14600	Windows or Linux based consoles are no longer lost upon rebooting VMs.
CS-14784	Multiple subnets with the same VLAN now work as expected.
CS-13303, 14874, 13897, 13944, 14088, 14190	A variety of upgrade issues have been fixed in release 3.0.3.
CS-15080	Setting a private network on a VLAN for VMWare environment is now supported.
CS-15168	The console proxy now works as expected and no exception is shown in the log after upgrading from version 2.2.14 to 3.0.2.
CS-15172	Version 3.0.2 now accepts the valid public key.

4.4. Known Issues in 3.0.3

Defects	Description
CS-14346	The UpdateVirtualMachine API call does not check whether the VM is stopped. Therefore, stop the VM manually before issuing this call.
CS-14361	On KVM hosts, the volume size allocated in primary storage is stored incorrectly in the database or displayed incorrectly in the Dashboard. The value displayed in the Dashboard is very small compared to the original size.
CS-14452	Data disk volumes are not automatically copied from one cluster to another.
CS-14770	The API does not return the keypair information when a VM is deployed with sshkey. This affects the API commands related to virtual machines (deployVirtualMachine, listVirtualMachines, ... *VirtualMachine), as well as the corresponding AWS APIs.
CS-14780	You are allowed to ping the elastic IP address of the VM even though no ingress rule is set that allows the ICMP protocol.
CS-14796	Deploying a VM is allowed even if the user data is not Base 64 encoded. Using the ec2-run-instances API with -d or -f option with user data in plain does not return any error.

Defects	Description
CS-14879	When a user VM is stopped or terminated, the static NAT associated with this VM is not disabled. This public IP address is still owned by this account and cannot be associated to any other user VM.
CS-14939	Adding a VMware cluster is not supported when the Management Network is migrated to the Distributed Virtual Switch environment.
CS-14952	The vCenter IP Address and the datacenter information is not present in the "virtual_supervisor_module" table. The Nexus virtual switch credentials are not encrypted.
CS-15009	The port_profile table will not be populated with port profile information. In this release, CloudPlatform directly connects to the VSM for all the port profile operations; therefore, no port profile information is cached.
CS-15037	Hairpin NAT is not supported when NetScaler is used for EIP.
CS-15092	Connecting to the guest VMs through SSH is extremely slow, and it results in connection timeout.
CS-15105	The cloud-sysvmadm script does not work if the integration.api.port parameter is set to any port other than 8096.
CS-15117	In a deployment with Nexus 1000v virtual switch, disable/enable operation of the Nexus virtual switch is not working as expected. The Nexus 1000v virtual switch continues to be used to create network or edit network operations even after disabling the switch.
CS-15118	In a deployment with Nexus 1000v virtual switch, zone VLAN range is not validated against the reserved list of VLANs for Nexus 1000v.
CS-15120	No actions are listed in the Action column of the Volumes page in the CloudPlatform UI.
CS-15124	Mixed switch environment is not supported. The zone can either be deployed as Standard vSwitch based or Nexus virtual switch based.
CS-15163	The minimum limit is not honored when there is not enough capacity to deploy all the VMs and the ec2-run-instances command with the -n >n1 -n2> option is used to deploy multiple VMs.
CS-15167	The Amazon Web Services API does not allow the admin user to view or act on the resources owned by the regular users.
CS-15198	Peak bandwidth (PIR) and burst size shaping policies are not applied on Nexus 1000v virtual switch interface.
CS-15218	You might find the term "CloudStack" when you expect "CloudPlatform" in scripts, file names, etc. The use of the new product name CloudPlatform is not yet fully implemented.
CS-15256	If cluster addition fails in a zone using the Cisco Nexus 1000v virtual switch, a subsequent retry will not succeed in adding the cluster. To work around:

Defects	Description

	1. Find the VSM id that was attempted to be added along with the cluster when the cluster creation failed. To do this, log into the MySQL database and execute a select query:
	
	`# mysql -uroot -p <password_for_mysql_db>;`
	`mysql> use cloud;`
	`mysql> select id from`
	`` `cloud`.`virtual_supervisor_module` where``
	`ipaddr="<vsm_ipaddress>";`
	
	2. Delete the cluster to VSM mapping in the cluster_vsm_map table for this vsm id:
	
	`mysql> delete from` `` `cloud`.`virtual_supervisor_module` ``
	`where vsm_id=<the id returned in step1>;`
	
	3. Try again to add the cluster.

4.5. API Changes from 3.0.2 to 3.0.3

4.5.1. New API Commands

API Commands	Description
listCiscoNexusVSMs	Retrieves information of a Cisco Nexus 1000v virtual switch associated with a cluster. It lists the control VLAN ID, packet VLAN ID, and data VLAN ID, as well as the IP address of the Nexus virtual switch.
enableCiscoNexusVSM	Enables a Cisco Nexus VSM device.
disableCiscoNexusVSM	Disables a Cisco Nexus VSM device.
deleteCiscoNexusVSM	Deletes a Cisco Nexus VSM device.
markDefaultZoneForAccount	Marks a default zone for the current account.
uploadVolume	Uploads a data disk.

4.5.2. Changed API Commands

API Commands	Description
reconnectHost	A new response parameter is added: hahost.
addCluster	The following request parameters are added: • vsmipaddress (optional) • vsmpassword (optional) • vsmusername (optional) The following parameter is made mandatory: podid
listVolumes	A new response parameter is added: status

API Commands	Description
migrateVolume	A new response parameter is added: status
prepareHostForMaintenance	A new response parameter is added: hahost.
addSecondaryStorage	A new response parameter is added: hahost.
enableAccount	A new response parameter is added: defaultzoneid
attachVolume	A new response parameter is added: status
cancelHostMaintenance	A new response parameter is added: hahost
addSwift	A new response parameter is added: hahost
listSwifts	A new response parameter is added: hahost
listExternalLoadBalancers	A new response parameter is added: hahost
createVolume	A new response parameter is added: status
listCapabilities	A new response parameter is added: customdiskofferingmaxsize
disableAccount	A new response parameter is added: defaultzoneid
deployVirtualMachine	A new request parameter is added: startvm (optional)
deleteStoragePool	A new request parameter is added: forced (optional)
updateAccount	A new response parameter is added: defaultzoneid
addHost	A new response parameter is added: hahost
updateHost	A new response parameter is added: hahost
detachVolume	A new response parameter is added: status
listAccounts	A new response parameter is added: defaultzoneid
listHosts	A new response parameter is added: hahost A new request parameter is added: hahost (optional)

Version 3.0.2

5.1. New Upgrade Path

Starting with version 3.0.2, existing 2.2.x installations can be upgraded. See Upgrade from 2.2.x to 3.0.2.

5.2. What's New in 3.0.2

Version 3.0.2 includes no new product features. However, the following changes were introduced in version 3.0.2.

CS-14588	When creating a virtual machine in the UI, the name of the VM specified will now directly affect its DNS name, unlike previous releases. DNS names can only be specified during creation. You can continue to modify the Display Name, but this will have no effect on the DNS name.

5.3. Issues Fixed in 3.0.2

Defect	Description
CS-13398	Pressing ESC while in any UI dialog will no longer freeze the UI.
CS-14305	You can now upgrade the service offerings for system VMs.
CS-14333	Fixed an issue that prevented any users from upgrading their service offerings.
CS-14341	Fixed an issue where the orders of the service/disk offerings were not being honored in the UI.
CS-14364	Ubuntu 10.04 now works with the Quick Basic Install.
CS-14369	Quick Basic Install for KVM hypervisors can now be done on a single host.
CS-14372	Async Job timeout counter is now properly reset across calls.
CS-14406	Putting a host into maintenance will now properly live migrate the VM rather than stopping it.
CS-14433	Fixed an issue where System VMs were not correctly releasing their IP when destroyed.
CS-14441	VMWare: Networks labels are not correctly honored.
CS-14447	Registering an ISO/Template with “All Zones” will now correctly propagate the templates to all zones.
CS-14447	Deleting an account will now properly remove its associated Virtual Router.
CS-14512	The DeployVirtualMachine command will accept UUID for the “iptonetworklist” parameter.
CS-14542	Fixes issues with labeling networks for KVM.
CS-14593	Fixed an issue with where you hit a “EncryptionOperationNotPossibleException” when attempting to validate user credentials against a LDAP server.

Defect	Description
CS-14614	Fixed issues with attempting to boot off of a CentOS 6.2 Live CD.
CS-14622	Fixed issues with dedicating public IP ranges to an account.
CS-8222	IP address usage records are now created for each day the IP address is assigned, in order to be consistent with virtual machine and storage usage records.

5.4. Known Issues in 3.0.2

Issue ID	Description
CS-14103	Global settings such as *.network.device (for example, xen.guest.network.device) are not cleared during upgrade from 2.2.x to 3.0.x. If you had manually set the global configuration parameter *.network.device (for example, xen.guest.network.device) to a custom network label, you will need to set the network traffic labels manually each time you set up a new zone after upgrade.
CS-14201 (was 14430)	VMWare: Template sizes are being reported different depending on whether the primary storage is using ISCSI or NFS.
CS-14237 (was 14468)	Primary Storage GC cannot occur when the first host in a given cluster is not in an “Up” state.
CS-14275 (was 14506)	F5: Unable to properly remove a F5 device.
CS-14291 (was 14523)	The EIP/ELB network offering for basic zones does not support multiple NetScalers.
CS-14296 (was 14530)	OVM: Network traffic labels are not supported.
CS-14303 (was 14537)	The IP addresses for a shared network are still being consumed even if no services are defined for that network.
CS-14655	XenServer known issue: If using network bonding on XenServer 6.0.x along with the CloudStack XenServer Support Package (CSP), interfaces, IPs, and networking can fail after adding hosts to the cloud and rebooting them. Workaround: patch each XenServer node, using the following steps. This patch is required only if using the CSP, and is a temporary requirement until XenServer issues the official hotfix for this issue. Before installing the CSP package, install the newer biosdevname rpm, selecting the correct patch for your version of XenServer: XenServer patch XS759 for XenServer 6.0.2 or XenServer patch XS753 for XenServer 6.0.0. 1. Download three files: • *http://download.cloud.com/releases/3.0.2/60-net.rules.template* • *http://download.cloud.com/releases/3.0.2/net-rename-sideways.sh* • *http://download.cloud.com/releases/3.0.2/biosdevname-0.3.7-1.xs753.i386.rpm* (for XenServer

Issue ID	Description
	6.0.0) or *http://download.cloud.com/releases/3.0.2/biosdevname-0.3.7-1.xs759.i386.rpm* (for XenServer 6.0.2) 2. Run the following. Substitute the filename of the patch you are using: `rpm -ivh <packageName>.rpm -force` 3. Replace /etc/udev/scripts/net-rename-sideways.sh with the new version. 4. Replace /etc/sysconfig/network-scripts/interface-rename-data/60-net.rules.template with the newer version. 5. Delete /etc/udev/rules.d/60-net.rules.
CS-14756	Installing KVM on RHEL 6.2 will result in unreliable network performance. Workaround: blacklist vhost-net. Edit /etc/modprobe.d/blacklist-kvm.conf and include vhost-net.

5.5. API Changes from 3.0.1 to 3.0.2

Added the following API command:

changeServiceForSystemVm

Changes the service offering for a system VM (console proxy or secondary storage). The system VM must be in a "Stopped" state for this command to take effect.

Version 3.0.1

6.1. New Software License

Starting with version 3.0.1, the software license has changed from the GNU Public License Version 3 to the Apache License Version 2.0. For the text of the license, see *Apache License*[1].

6.2. What's New in 3.0.1

Version 3.0.1 includes no new product features. For a list of the major fixed items, see Issues Fixed in 3.0.1. The following changes were introduced in version 3.0.1.

Componenets	Description
XenServer Version	XenServer 6.0.2 is now fully supported.
vSphere 5.0	VMWare vSphere 5.0 is now fully supported.
KVM Basic Install	In the basic installation, which is one that follows the prompts in the guided tour, the use of the KVM hypervisor is now tested and officially supported in addition to XenServer.
Adding a Zone	Some fields and controls in the Add Zone wizard are different than in 3.0.0. These changes will allow the first cluster and host to work properly immediately after you click "Enable" at the end of the Add Zone wizard: You must now choose the hypervisor before adding the first cluster. This applies only to the first cluster; you can still add clusters with different hypervisors later. You can configure network traffic labels for the different traffic types. This fixes a known issue from release 3.0.0.
14379	Port 8096, which allows API calls without authentication, is closed and disabled by default on any fresh 3.0.1 installations. You can enable 8096 (or another port) for this purpose as follows: After the first Management Server is installed and running, set the global configuration parameter integration.api.port to the desired port, and restart the Management Server. Then, on the Management Server host machine, create an iptables rule allowing access to that port.

The following new error codes have been added. If one of the errors occurs, check the error string for more information:

4250 : "com.cloud.utils.exception.CloudRuntimeException"

4255 : "com.cloud.utils.exception.ExceptionUtil"

4260 : "com.cloud.utils.exception.ExecutionException"

4265 : "com.cloud.utils.exception.HypervisorVersionChangedException"

4270 : "com.cloud.utils.exception.RuntimeCloudException"

[1] http://www.apache.org/licenses/LICENSE-2.0.txt

4275 : "com.cloud.exception.CloudException"

4280 : "com.cloud.exception.AccountLimitException"

4285 : "com.cloud.exception.AgentUnavailableException"

4290 : "com.cloud.exception.CloudAuthenticationException"

4295 : "com.cloud.exception.CloudExecutionException"

4300 : "com.cloud.exception.ConcurrentOperationException"

4305 : "com.cloud.exception.ConflictingNetworkSettingsException"

4310 : "com.cloud.exception.DiscoveredWithErrorException"

4315 : "com.cloud.exception.HAStateException"

4320 : "com.cloud.exception.InsufficientAddressCapacityException"

4325 : "com.cloud.exception.InsufficientCapacityException"

4330 : "com.cloud.exception.InsufficientNetworkCapacityException"

4335 : "com.cloud.exception.InsufficientServerCapacityException"

4340 : "com.cloud.exception.InsufficientStorageCapacityException"

4345 : "com.cloud.exception.InternalErrorException"

4350 : "com.cloud.exception.InvalidParameterValueException"

4355 : "com.cloud.exception.ManagementServerException"

4360 : "com.cloud.exception.NetworkRuleConflictException"

4365 : "com.cloud.exception.PermissionDeniedException"

4370 : "com.cloud.exception.ResourceAllocationException"

4375 : "com.cloud.exception.ResourceInUseException"

4380 : "com.cloud.exception.ResourceUnavailableException"

4385 : "com.cloud.exception.StorageUnavailableException"

4390 : "com.cloud.exception.UnsupportedServiceException"

4395 : "com.cloud.exception.VirtualMachineMigrationException"

4400 : "com.cloud.exception.AccountLimitException"

4405 : "com.cloud.exception.AgentUnavailableException"

4410 : "com.cloud.exception.CloudAuthenticationException"

4415 : "com.cloud.exception.CloudException"

4420 : "com.cloud.exception.CloudExecutionException"

4425 : "com.cloud.exception.ConcurrentOperationException"

4430 : "com.cloud.exception.ConflictingNetworkSettingsException"

4435 : "com.cloud.exception.ConnectionException"

4440 : "com.cloud.exception.DiscoveredWithErrorException"

4445 : "com.cloud.exception.DiscoveryException"

4450 : "com.cloud.exception.HAStateException"

4455 : "com.cloud.exception.InsufficientAddressCapacityException"

4460 : "com.cloud.exception.InsufficientCapacityException"

4465 : "com.cloud.exception.InsufficientNetworkCapacityException"

4470 : "com.cloud.exception.InsufficientServerCapacityException"

4475 : "com.cloud.exception.InsufficientStorageCapacityException"

4480 : "com.cloud.exception.InsufficientVirtualNetworkCapcityException"

4485 : "com.cloud.exception.InternalErrorException"

4490 : "com.cloud.exception.InvalidParameterValueException"

4495 : "com.cloud.exception.ManagementServerException"

4500 : "com.cloud.exception.NetworkRuleConflictException"

4505 : "com.cloud.exception.PermissionDeniedException"

4510 : "com.cloud.exception.ResourceAllocationException"

4515 : "com.cloud.exception.ResourceInUseException"

4520 : "com.cloud.exception.ResourceUnavailableException"

4525 : "com.cloud.exception.StorageUnavailableException"

4530 : "com.cloud.exception.UnsupportedServiceException"

4535 : "com.cloud.exception.VirtualMachineMigrationException"

In addition, there is special error code for ServerApiException when it is thrown in a standalone manner when failing to detect any of the above standard exceptions:

9999 : "com.cloud.api.ServerApiException"

6.3. Issues Fixed in 3.0.1

Defects	Description
Many	In the Add Zone wizard, added a step for configuring network traffic labels on the physical network(s).
13313	The Add Zone wizard will now skip adding an ESXi host if the cluster is VMWare.
13899	NetScaler is no longer a selectable provider for static NAT network service.
13966	Fixed issue with not cleaning up instance when it fails to acquire an EIP address.
14016	NetScaler – Deleting a load balancer rule will no longer delete other load balancer rules pointing at the same private port.
14023	UI – You can now update SSL certificates on system VMs.

Defects	Description
14042	Fixed issues where VMs are not able to access the public network when attached to an isolated guest network with source NAT enabled and a shared network.
14047	Login API no longer fails when using UUID for the domainId parameter.
14073	Zones will now be automatically be created as public zones and not dedicated to the ROOT domain.
14077	DestroyVirtualMachine API call will now work against VM when the VM's state is Starting.
14101	You can now specify the storage network when adding a basic zone.
14135	You can now specify the storage network when adding a basic zone.
14135	Windows 2003 is not reported as supported for XenServer.
14188	Fixed an issue where the pre-generated SSH keys are not properly updated to the system VMs.
14189	Fixed an issue where the secondary storage VM is not using the storage network to download templates.
14202	Non-bootable ISOs no longer show up in the Add Instance wizard.
14216	Fixed issues with KVM when adding multiple physical networks.
14239	Added ability for administrators to limit the number of guest networks.
14282	Fixed an issue where KVM is not able to reconnect the management server after a management server reboot.
14285	NetScaler SDX: can now create VPX instances when XVA image is not NSVPX-XEN-9.3-52.4_nc.xva
14313	Fixed an issue with the JSON builder that prevented a template copy across zones from working properly.
14332	You can now delete a host.
14336	Login API now returns the account's UUID.
14392	You can now add public IP ranges even though you did not add zone VLANs.
14484	Fixed issues where new port forward could not be added if you created it, then deleted, and attempted to recreate the port forward again.
14492	System VMs now work if you have configured multiple physical networks across zones.
14515	Snapshots are now properly cleaned up.

6.4. Known Issues in 3.0.1

Defects	Description
14430	VMWare: Template sizes are being reported different depending on whether the primary storage is using ISCSI or NFS.

Defects	Description
14468	Primary Storage GC cannot occur when the first host in a given cluster is not in an “Up” state.
14506	F5: Unable to properly remove a F5 device.
14523	The EIP/ELB network offering for basic zones does not support multiple NetScalers.
14530	OVM: Network traffic labels are not supported.
14537	IP addresses for a shared network are still being consumed even if no services are defined for that network.

6.5. API Changes from 3.0.0 to 3.0.1

6.5.1. Added API Commands

ldapRemove

Remove the LDAP context for this site.

6.5.2. Changed API Commands

API Commands	Description
addHost	Changed request parameters: podid (old version - optional, new version - required)
assignVirtualMachine	New response field: instancename
attachIso	New response field: instancename
changeServiceForVirtualMachine	New response field: instancename
deployVirtualMachine	New response field: instancename
destroyVirtualMachine	New response field: instancename
detachIso	New response field: instancename
disableAccount	New response fields: networkavailable, networklimit, networktotal, projectavailable, projectlimit, projecttotal
enableAccount	New response fields: networkavailable, networklimit, networktotal, projectavailable, projectlimit, projecttotal
listAccounts	New response fields: networkavailable, networklimit, networktotal, projectavailable, projectlimit, projecttotal
listLoadBalancerRuleInstances	New response field: instancename
listOsCategories	New request parameter: name (optional)
listOsTypes	New request parameter: description (optional)
listSystemVms	New request parameter: storageid (optional)
listVirtualMachines	New response field: instancename
migrateVirtualMachine	New response field: instancename
rebootVirtualMachine	New response field: instancename
recoverVirtualMachine	New response field: instancename

API Commands	Description
resetPasswordForVirtualMachine	New response field: instancename
restoreVirtualMachine	New response field: instancename
startVirtualMachine	New request parameter: hostid (optional) New response field: instancename
stopVirtualMachine	New response field: instancename
updateAccount	New response fields: networkavailable, networklimit, networktotal, projectavailable, projectlimit, projecttotal
updateVirtualMachine	New response field: instancename

Version 3.0.0

7.1. Overview of Major New Features in 3.0

Version 3.0 is a major new release. It provides several new features compared to version 2.2.x. This section provides overviews of the new features.

7.1.1. Redesigned User Interface

The user interface has been redesigned to provide easier navigation as well as a more intuitive workflow. Graphical displays of the infrastructure topology have replaced drill-down lists as the main way to access the various components such as zones, hosts, and networks. The main Dashboard now provides a more clear display of key information for managing the cloud. The end-user UI also benefits from this redesign, making is easier for users to manage their VMs and other resources. The new Project View lets users switch context from one set of resources to another, enabling a more efficient focus on the task at hand.

7.1.2. NetScaler Load Balancer

Citrix NetScaler is now supported as an external network element for load balancing. Set up an external load balancer when you want to provide load balancing through means other than the provided virtual router. The NetScaler can be set up in direct (outside the firewall) mode. It must be added before any load balancing rules are deployed on guest VMs in the zone.

Note

Limitations: NetScaler cannot yet be used as a firewall. It cannot currently be set up in in-line mode (behind the firewall).

7.1.3. Sticky Session Policies for Load Balancer Rules

Sticky sessions are used in Web-based applications to ensure continued availability of information across the multiple requests in a user's session. For example, if a shopper is filling a cart, you need to remember what has been purchased so far. The concept of “stickiness” is also referred to as persistence or maintaining state.

Any load balancer rule can have a stickiness policy. The policy consists of a name, stickiness method, and parameters. The stickiness method could be load balancer-generated cookie, application-generated cookie, or source-based. In the source-based method, the source IP address is used to identify the user and locate the user’s stored data. In the other methods, cookies are used. The cookie generated by the load balancer or application is included in request and response URLs to create persistence. A variety of options are provided to control the exact behavior of cookies, such as how they are generated and whether they are cached.

7.1.4. Using an LDAP Server for User Authentication

In version 3.0, you can use an external LDAP server such as Microsoft Active Directory or ApacheDS for end-user authentication. Just map the accounts to the corresponding LDAP accounts using a query filter. The query filter is written using the query syntax of the particular LDAP server, and can include special wildcard characters for matching common values such as the user’s email address and name. The external LDAP directory tree is searched starting at a specified base directory and the

distinguished name (DN) and password of the matching user are returned. This information along with the given password is used to authenticate the user.

7.1.5. VM Storage Migration

The CloudPlatform administrator can move a virtual machine's root disk volume or any additional data disk from one storage pool to another in the same zone. You can use the storage migration feature to achieve some commonly desired administration goals, such as balancing the load on storage pools and increasing the reliability of virtual machines by moving them away from any storage pool that is experiencing issues. This functionality is supported in XenServer, KVM, and VMware.

7.1.6. Swift for Secondary Storage

In version 3.0, OpenStack Object Storage *Swift*[1] is supported for secondary storage. When using Swift, you configure Swift storage for the entire platform, then set up NFS secondary storage for each zone. The NFS storage in each zone acts as a staging area through which all templates and other secondary storage data pass before being forwarded to Swift. The Swift storage acts as a cloud-wide resource, making templates and other data available to any zone in the cloud. There is no hierarchy in the Swift storage, just one Swift container per storage object. Any secondary storage in the whole cloud can pull a container from Swift at need – no more copying templates and snapshots from one zone to another. Everything is available everywhere

7.1.7. Password and Key Encryption

CloudPlatform stores several sensitive passwords and secret keys that are used to provide security. Starting in version 3.0, these values are always automatically encrypted. These include the database secret key, database password, SSH keys, compute node root password, VPN password, user API secret key, and VNC password.

In version 3.0, the Java Simplified Encryption (JASYPT) library is used. The data values are encrypted and decrypted using a database secret key. Of course, the database secret key itself cannot be stored in the open – it must be encrypted. To read it, a second secret key must be provided from an external source during Management Server startup. This key can be provided in one of two ways: loaded from a file or provided by the CloudPlatform administrator. The encryption type, database secret key, and Management Server secret key are set by the administrator during the CloudPlatform installation.

7.1.8. Security Group Egress Rules

Security groups can be used to control network traffic to and from VMs. A security group is a group of VMs that filter their incoming and outgoing traffic according to a set of rules, called ingress and egress rules. These rules filter network traffic according to the IP address that is attempting to communicate with the VM.

In addition to ingress rules that control incoming network traffic to VMs in a given security group, starting in version 3.0 you can also define egress rules to control outgoing network traffic. If no egress rules are specified, then all traffic will be allowed out. Once egress rules are specified, the following types of traffic are allowed out: traffic specified in egress rules; queries to DNS servers; and responses to any traffic that has been allowed in through an ingress rule. An egress rule can be specified either by CIDR to specify IP addresses, or by account to allow traffic from another security group.

[1] http://swift.openstack.org

7.1.9. Using Projects to Organize Users and Resources

In version 3.0, users can group themselves into projects so they can collaborate and share virtual resources. The usage per project as well as per user is tracked, so the usage can be billed to either a user account or a project. For example, a private cloud within a software company might have all members of the QA department assigned to one project, so the company can track the resources used in testing while the project members can more easily isolate their efforts from other users of the same cloud. Per-project resource limits can be set.

You can configure CloudPlatform to allow any user to create a new project, or you can restrict that ability to just administrators. You can either add people directly to a project, or you have to send an invitation which the recipient must accept.

A user can be a member of any number of projects and can switch to a new Project View in the CloudPlatform UI to show only project-related information, such as project VMs, fellow project members, project-related alerts, and so on.

7.1.10. Providing Network Services for Users

People using cloud infrastructure have a variety of needs and preferences when it comes to the networking services provided by the cloud. Provisioning physical and virtual networks has always been supported in CloudPlatform. As an administrator, you can do the following additional things to set up networking for your users:

- Set up several different providers (also known as network elements) for the same service on a single physical network. For example, you can provide both Cisco and Juniper firewalls. You can have multiple instances of the same service provider in a network; for example, more than one Juniper SRX device.
- Bundle different types of network services into network offerings. When creating a new VM, the user chooses one of the available network offerings, and that determines which network services the VM can use. A network offering is a named set of network services, such as DHCP, source NAT, load balancing, firewall, VPN, port forwarding, and specific network service providers, such as Juniper SRX for the firewall. You can add new network offerings as time goes on so end users can upgrade to a better class of service on their network.
- Provide more ways for a network to be accessed by a user, such as through a project of which the user is a member.
- Set up two types of virtual networks: shared and isolated. An isolated network can be accessed only by virtual machines of a single account. A shared network can be accessed by virtual machines that belong to many different accounts. Network isolation on shared networks is accomplished using techniques such as security groups.
- ¬ More directly control the physical network, such as add/remove/update physical networks in a zone, configure VLANs on the physical network, specify properties like network speed, configure a name so the network can be recognized by hypervisors, configure the IP addresses trunked to a physical network, and specify what type of traffic is carried on the physical network (such as guest VM traffic vs. internal management traffic).

7.2. New Features in 3.0.0

Issue ID	Description
4282	Added nonce support in API.

Issue ID	Description
5510	Openstack Swift can now be used as an alternative to NFS storage for templates, ISO, and snapshots.
5822	All sensitive passwords are now properly encrypted in the database and any configuration files.
6745	UUIDs are now used in place of regular DB IDs. 3.0 API will support both.
6876	Netscaler MPX, VPX, and SDX is now supported.
7883	Templates, ISOs, Disk, and Service offerings can now be sorted to allow admins to more easily view them in the UI.
Many	Network as a Service feature.
8313	Basic LDAP authentication is now built in as an optional AUTH adapter.
8620	Projects feature.
8791	User dispersing allocator has now been added as an alternative algorithm for VM placement.
8962	Admins can now re-assign VM from one account to another.
9128	Network throttling is now controlled via network offerings.
9154	Redundant Router support has been added.
Many	Brand new 3.0 User Interface.
9949	Users can now revert a VM to the original template it was created from.
Many	State management now included to pod and cluster level from the original host and zone level support.
10405	API Version annotation supported.
10588	XenServer 6.0 is now supported.
10617	Egress rules for security groups now supported.
10657	Capacity now has two levels of threshold support. One threshold is used to alert. The other is to disable resource allocation.
10792	Added ability for admins to set ingress rules that cannot be removed by user.
10796	Sticky session now supported for load balancers.
11303	Added support in login API call to take in a map of parameters that can be passed into the authenticators.
11173	VPN usage is now added as a new usage record.
11598	MTU for secondary storage is now configurable via Global configuration.
11689	Templates now have a SSH enabled flag similar to password enabled flag.
Many	vSphere 5.0 now has Beta support.
Many	RHEL/Centos 6.2 (KVM) is now supported.

7.3. Issues Fixed in 3.0.0

Defects	Description
Many	VM Sync has been improved so that VM state should be better reflected between Management Server and Hypervisor.
8150	Template delete events are now recorded after being expunged.
8870	IPs from “Direct-Tagged” networks (shared guest networks) are no longer counted as part of the total number of public IPs.
9036	Migrated VMs will now have their consumed resources reflected properly in the capacity reports.
9842	Network Usage time range aggregation has been improved.
10043	Restarting the Management Server will no longer change a host status from Maintenance to Up.
10067	Extractable attribute can now be edited by administrators.
10195	Added a new VM state of “Unknown” if the host state is in “Alert”.
10217	Management Server installation will now check for FQDN hostname instead of stopping with an error condition.
10292	UI will no longer allow attaching volumes from one hypervisor to another.
10307	Network Usage will now account for more than one virtual interface.
10354	VMware ISO attach and detach events are now correctly registered.
10362	Disabling VPN will now work correctly even if the virtual router is in a STOPPED state.
10674	Management Server will now alert if it hits a snapshot limit quota.
10779	Port Ranges now work in the UI.
10831	Adding Secondary Storage URL with double slashes will now work.
11056	DHCP issues with Debian/Ubuntu guest OS have been resolved.
11131	VM scheduler will no longer retry in the same zone if that zone has been declared as non-allocable.
11193	Management Server will now alert a link-local IP capacity issue.
Many	Added new Global configurations to set limits for the number of guest VMs per hypervisor.
11273	Management Server will no longer attempt to program security group rules for non-reachable hypervisors.
11284	Administrators can now add a Basic Zone without security group support.
11311	Improved listVirtualMachine API call performance.
11387	Public IP of the secondary storage VM will now be correctly returned to the pool after being expunged.
11492	Volume limits are now checked when deploying a new VM.

Defects	Description
11542	Management Server will no longer allow same public IP ranged across zones.
11585	Multiple public VLANs are now correctly supported in VMware.
11616	Manual live migration of VMs for KVM is now supported.
11814	General guest VM options for VMware are now supported.
11838	Deleted VM template names can now be re-used.
11902	Added global configuration to allow different NIC drivers for VMware system VM.
11926	Installation of system template will now perform mount point validation before proceeding.
Many	Many passwords are no longer logged in the Management Server logs.
Many	Secondary storage VM has been hardened.
12139	Added a way for Administrators to specify the default system template to use on a global or per zone basis.
12113	Improved re-try algorithm when attempting to copy a template from secondary to primary during failure scenarios.
12162	CreateLoadBalancer public ID was incorrectly published as optional. It is now required.
12192	State NFS handle are now correctly handled for KVM.
12290	Security Groups improvements.
12476	DHCP anti-spoofing fixes.
12481	Account ID is now returned as part of listAccount API.
12705	Source NAT is no longer configured on additional IP of a different network interface.
12782	Capacity now reflects hosts in maintenance mode.
12820	KVM: Attached disks are no longer removed after a VM reboot.
12877	Pagesize = -1 now works correctly.
12848	Removed notion of setting a default network when adding a shared network. Default networks are now specified during VM deployment.
12929	Added domain ID to all events.
13201	Added global configuration to allow administrator to specify default network device drivers for system VMs.
13315	Added BASIC auth http proxy for secondary storage VM.
13396	Escaped double-quotes in JSON responses.
13537	Templates created from snapshots now work with NFSv4.
13777	VMware snapshot errors are now handled better.

7.4. Known Issues in 3.0.0

Defects	Description
11535	In-line mode for load balancer is not supported for all external devices.
12741	vSphere: maintenance mode will not live migrate system VM to another host.
12840	Capacity view is not available for pods or clusters.
13518	Security Groups are not supported in Advanced Networking
Many	F5 Known Issues: • Unable to create load balancer rule for port 22. • No support for changing algorithm once rule has been created. • Source algorithm is not supported. Setting a rule to source will prevent other rules from being created properly. • Virtual router upgrades do not migrate all sticky session parameters correctly.
Many	NetScaler Known Issues: • When a VM from a load balancer rules is removed, it will also get removed for other load balancer rules of the same port.. • Sticky session method “lbCookie” and “appCookie” do not work for any port other than 80. • Virtual router upgrades do not migrate all sticky session parameters correctly. • Once the public port 80 has been mapped to any private port, “A”, no other public port can be mapped to that private port, “A”..
13336	vSphere: cross cluster volume migration does not work properly.
13359	Programming F5/NetScaler rules can be better optimized.
Many	Network restart can fail under certain circumstances.
13883	Multiple NetScalers are not supported in Basic Networking.
13935	vSphere: detaching an ISO from a restored VM instance fails.
13963	vSphere: template download from templates created off of the root volume does not work properly.
Many	Disabling a pod or cluster does not prevent resource creation. Only zone level is supported right now.
14024	KVM: clustered LVM is not working properly.
Many	Bare metal host provisioning is not working properly.
Many	In the Add Zone wizard, there is no step for configuring network traffic labels on the physical network(s). Workaround: Don’t enable the zone in the last step of the wizard. Enable the zone only after configuring traffic labels for each traffic type, on each physical network, on each hypervisor in the zone. Set up the labels on

Defects	Description
	the hypervisor host, then configure matching labels through the CloudPlatform UI.

7.5. API Changes from 2.2.14 to 3.0

7.5.1. Change to Behavior of List Commands

There was a major change in how our List* API commands work in version 3.0 compared to 2.2.x. The rules below apply only for managed resources – those that belong to an account, domain, or project. They are irrelevant for the List* commands displaying unmanaged (system) resources, such as hosts, clusters, and external network resources. When no parameters are passed in to the call, the caller sees only resources owned by the caller (even when the caller is the administrator). Previously, the administrator saw everyone else's resources by default.

When accountName and domainId are passed in:

- The caller sees the resources dedicated to the account specified.
- If the call is executed by a regular user, the user is authorized to specify only the user's own account and domainId.
- If the caller is a domain administrator, an authorization check is performed to see whether the caller is permitted to view resources for the given account and domainId.

When projectId is passed in, only resources belonging to that project are listed.

When domainId is passed in, the call returns only resources belonging to the domain specified. To see the resources of subdomains, use the parameter isRecursive=true. Again, the regular user can see only resources owned by that user, the root administrator can list anything, and a domain administrator is authorized to see only resources of the administrator's own domain and subdomains.

To see all resources the caller is authorized to see, except for Project resources, use the parameter listAll=true.

To see all Project resources the caller is authorized to see, use the parameter projectId=-1.

There is one API command that doesn't fall under the rules above completely: the listTemplates command. This command has its own flags defining the list rules, as shown in the following table.

listTemplates Flag	Description
Featured	Returns templates that have been marked as featured and public.
Self	Returns templates that have been registered or created by the calling user.
selfexecutable	Same as self, but only returns templates that are ready to be deployed with.
sharedexecutable	Ready templates that have been granted to the calling user by another user.
executable	Templates that are owned by the calling user, or public templates, that can be used to deploy a new VM.
community	Returns templates that have been marked as public but not featured.

listTemplates Flag	Description
All	Returns all templates (only usable by admins).

7.5.2. Removed API commands

- createConfiguration (Adds configuration value)
- configureSimulator (Configures simulator)

7.5.3. Added API commands

- assignVirtualMachine (Move a user VM to another user under same domain.)
- restoreVirtualMachine (Restore a VM to original template or specific snapshot)
- createLBStickinessPolicy (Creates a Load Balancer stickiness policy)
- deleteLBStickinessPolicy (Deletes a LB stickiness policy.)
- listLBStickinessPolicies (Lists LBStickiness policies.)
- ldapConfig (Configure the LDAP context for this site.)
- addSwift (Adds Swift).
- listSwifts (List Swift.)
- migrateVolume (Migrate volume)
- updateStoragePool (Updates a storage pool.)
- authorizeSecurityGroupEgress (Authorizes a particular egress rule for this security group)
- revokeSecurityGroupEgress (Deletes a particular egress rule from this security group)
- createNetworkOffering (Creates a network offering.)
- deleteNetworkOffering (Deletes a network offering.)
- createProject (Creates a project)
- deleteProject (Deletes a project)
- updateProject (Updates a project)
- activateProject (Activates a project)
- suspendProject (Suspends a project)
- listProjects (Lists projects and provides detailed information for listed projects)
- addAccountToProject (Adds acoount to a project)
- deleteAccountFromProject (Deletes account from the project)
- listProjectAccounts (Lists project's accounts)
- listProjectInvitations (Lists an account's invitations to join projects)

- updateProjectInvitation (Accepts or declines project invitation)
- deleteProjectInvitation (Deletes a project invitation)
- updateHypervisorCapabilities (Updates a hypervisor capabilities.)
- listHypervisorCapabilities (Lists all hypervisor capabilities.)
- createPhysicalNetwork (Creates a physical network)
- deletePhysicalNetwork (Deletes a Physical Network.)
- listPhysicalNetworks (Lists physical networks)
- updatePhysicalNetwork (Updates a physical network)
- listSupportedNetworkServices (Lists all network services provided by CloudPlatform or for the given Provider.)
- addNetworkServiceProvider (Adds a network serviceProvider to a physical network)
- deleteNetworkServiceProvider (Deletes a Network Service Provider.)
- listNetworkServiceProviders (Lists network serviceproviders for a given physical network.)
- updateNetworkServiceProvider (Updates a network serviceProvider of a physical network)
- addTrafficType (Adds traffic type to a physical network)
- deleteTrafficType (Deletes traffic type of a physical network)
- listTrafficTypes (Lists traffic types of a given physical network.)
- updateTrafficType (Updates traffic type of a physical network)
- listTrafficTypeImplementors (Lists implementors of implementor of a network traffic type or implementors of all network traffic types)
- createStorageNetworkIpRange (Creates a Storage network IP range.)
- deleteStorageNetworkIpRange (Deletes a storage network IP Range.)
- listStorageNetworkIpRange (List a storage network IP range.)
- updateStorageNetworkIpRange (Update a Storage network IP range, only allowed when no IPs in this range have been allocated.)
- listUsageTypes (List Usage Types)
- addF5LoadBalancer (Adds a F5 BigIP load balancer device)
- configureF5LoadBalancer (configures a F5 load balancer device)
- deleteF5LoadBalancer (delete a F5 load balancer device)
- listF5LoadBalancers (lists F5 load balancer devices)
- listF5LoadBalancerNetworks (lists network that are using a F5 load balancer device)
- addSrxFirewall (Adds a SRX firewall device)

- deleteSrxFirewall (delete a SRX firewall device)
- configureSrxFirewall (Configures a SRX firewall device)
- listSrxFirewalls (lists SRX firewall devices in a physical network)
- listSrxFirewallNetworks (lists network that are using SRX firewall device)
- addNetscalerLoadBalancer (Adds a netscaler load balancer device)
- deleteNetscalerLoadBalancer (delete a netscaler load balancer device)
- configureNetscalerLoadBalancer (configures a netscaler load balancer device)
- listNetscalerLoadBalancers (lists netscaler load balancer devices)
- listNetscalerLoadBalancerNetworks (lists network that are using a netscaler load balancer device)
- createVirtualRouterElement (Create a virtual router element.)
- configureVirtualRouterElement (Configures a virtual router element.)
- listVirtualRouterElements (Lists all available virtual router elements.)