Citrix CloudPlatform (powered by Apache CloudStack) Version 4.2.1-7 Release Notes

Revised August 20, 2015 07:00 am IST

Citrix CloudPlatform (powered by Apache CloudStack) Version 4.2.1-7 Release Notes Revised August 20, 2015 07:00 am IST

Release Notes for Citrix CloudPlatform version 4.2.1-7

Submitting Feedback and Getting Help

The support team is available to help customers plan and execute their installations. To contact the support team, log in to *the Support Portal*[1] by using the account credentials you received when you purchased your support contract.

[1] http://support.citrix.com/cms/kc/cloud-home/

Support Matrix

This section describes the operating systems, browsers, and hypervisors that have been newly tested and certified compatible with the CloudPlatform 4.2.1-7 Maintenance Release. Most of the earlier OS and hypervisor versions are also supported for use with the 4.2.1-7 Maintenance Release.

2.1. Supported OS Versions for Management Server

- RHEL versions 5.5, 6.2, 6.3, and 6.4
- CentOS versions 5.5, 6.2, 6.3, and 6.4

2.2. Supported Hypervisor Versions

The following hypervisora are supported:

- XenServer versions 5.6 SP2, 6.0, 6.0.2, 6.1, and 6.2
- KVM versions 5.5, 5.6, 5.7, 6.1, and 6.3
- VMware versions 4.1, 5.0.1 Update B, 5.0, and 5.1 Update 2
- Bare metal hosts are supported, which have no hypervisor. These hosts can run the following operating systems:
    - RHEL or CentOS, v6.2 or 6.3

      > **Note**
      >
      > Use libvirt version 0.9.10 for CentOS 6.3

    - Fedora 17
    - Ubuntu 12.04

2.3. Supported External Devices

- Netscaler VPX and MPX versions 9.3 and 10.e
- Netscaler SDX version 9.3
- SRX (Model srx100b) versions 10.3 or higher
- F5 10.1.0 (Build 3341.1084)

2.4. Supported Browsers

- Internet Explorer versions 8 and 9
- Firefox version 25 and 40
- Google Chrome versions 17, 20.0.1132.47m, and 44.0.2403.130 m

- Safari 5

Upgrade Instructions

3.1. Upgrade from 4.2.x.x to 4.2.1-7

Perform the following to upgrade from version 4.2.x.x to version 4.2.1-7.

1. Download the latest System VM templates:

   The System VM templates includes fixes for the OpenSSL vulnerability issues reported in *http://support.citrix.com/article/CTX140876*.

Hypervisor	Description
XenServer	Name: systemvm-xenserver-4.2.1-b Description: systemvm-xenserver-4.2.1-b URL (if using 32-bit system VM template): *http://download.cloud.com/templates/4.2/systemvmtemplate-2015-08-20-4.2.1-xen.vhd.bz2* URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/systemvmtemplate64-2015-08-20-4.2.1-xen.vhd.bz2* Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the XenServer zones. Hypervisor: XenServer Format: VHD OS Type: Debian GNU/Linux 7.0 (32-bit and 64-bit) (or the highest Debian release number available in the dropdown) Extractable: no Password Enabled: no Public: no Featured: no
KVM	Name: systemvm-kvm-4.2.1-b Description: systemvm-kvm-4.2.1-b URL (if using 32-bit system VM template): *http://download.cloud.com/templates/4.2/systemvmtemplate-2015-08-20-4.2.1-kvm.qcow2.bz2*

Hypervisor	Description
	URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/systemvmtemplate64-2015-08-20-4.2.1-kvm.qcow2.bz2* Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, choose All Zones to make the template available in all the KVM zones. Hypervisor: KVM Format: QCOW2 OS Type: Debian GNU/Linux 7.0 (32-bit and 64-bit) (or the highest Debian release number available in the dropdown) Extractable: no Password Enabled: no Public: no Featured: no
VMware	Name: systemvm-vmware-4.2.1-b Description: systemvm-vmware-4.2.1-b URL (if using 32-bit system VM template on earlier VMware version):*http://download.cloud.com/templates/4.2/systemvmtemplate-2015-08-20-4.2.1-vmware.ova* URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/systemvmtemplate64-2015-08-20-4.2.1-vmware.ova* Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, choose All Zones to make the template available in all the VMware zones. Hypervisor: VMware Format: OVA OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown)

Hypervisor	Description
	Extractable: no Password Enabled: no Public: no Featured: no

2. By using the **`prepareTemplate`** API, download the latest System VM to all the primary storages.

3. Check whether you already have the templates given at *http://support.citrix.com/article/CTX200024*.

   The article include the necessary Python scripts to update the version format entires to x.x.x.x in the database to successfully perform the SystemVM upgrade. If you are not using the tempates, follow the instructions:

   > **Note**
   >
   > The systemVM template upgrade considers only the first three digits in the release number, and therefore the template versions are considred the same for 4.2.0.0 and 4.2.0.1, which is incorrect. To make the necessary database changes to accept the four digit version format, x.x.x.x, run the Python script as given below.

   a. Download the Python script from *http://support.citrix.com/article/CTX200024*.

   b. Stop all the Management Servers.

   c. Backup the database.

   d. Execute the following python script on the Management Servers to update DB entries:

   ```
   # python sys-tmpl-upgrade-4.2.1.py -i <db host name/ip> -u <db user name> -p <db
    password>
   ```

   e. Start all the Management Servers.

   f. In the CloudPlatform UI, stop and start all the SSVM.

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

   Run the following commands on every KVM host.

   a. Download the CloudPlatform 4.2.1-7 RHEL 6.3 binaries from *https://www.citrix.com/English/ss/downloads/*

   b. Extract binaries:

   ```
   # cd /root
   ```

```
# tar xvf CloudPlatform-4.2.1-7-rhel6.3.tar.gz
```

c. Create a CloudPlatform 4.2.1-7 qemu repo:

```
# cd CloudPlatform-4.2.1-7-rhel6.3/6.3
# createrepo
```

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

```
[upgrade]
name=rhel63
baseurl=url-of-your-rhel6.3-repo
enabled=1
gpgcheck=0
[cloudstack]
name=cloudstack
baseurl=file:///root/CloudPlatform-4.2.1-7-rhel6.3/6.3
enabled=1
gpgcheck=0
```

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

```
yum upgrade
```

5. Stop all Usage Servers if running. Run the following command on all Usage Server hosts.

```
# service cloudstack-usage stop
```

6. Stop the Management Servers. Run the following command on all Management Server hosts.

```
# service cloudstack-management stop
```

7. On the MySQL master, backup the MySQL databases. Citrix recommends you to perform this step even in test upgrades. This will assist in debugging issues.

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

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

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

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

9. Download CloudPlatform 4.2.1-7 onto the management server host where it will run. Get the software from the following link:

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

To download this file from this link, you need a *My Citrix Account*[1].

10. Upgrade the CloudPlatform packages. You must have a file with the name as “CloudPlatform-4.2.1-N-OSVERSION.tar.gz”. Untar the file. Then, run the **`install.sh`** script from this file. Replace the file and directory names below with those you are using:

    ```
    # tar xzf CloudPlatform-4.2.1-N-OSVERSION.tar.gz
    # cd CloudPlatform-4.2.1-N-OSVERSION
    # ./install.sh
    ```

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

11. Choose "U" to upgrade the package

    ```
    >U
    ```

    You can see the output as the upgrade proceeds that ends with a message "Complete! Done."

12. If you have made changes to your existing copy of the configuration files db.properties or server.xml in your previous-version CloudPlatform installation, the changes will be preserved in the upgrade. However, you need to do the following steps to place these changes in a new version of the file which is compatible with version 4.2.1.

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

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

    a. Make a backup copy of your previous version file. For example: (substitute the file name in these commands as needed)

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

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

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

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

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

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

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

14. Start the first Management Server.

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

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

    Note

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

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

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

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

    Note

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

    a. Copy the CloudPlatform 4.2.1-7.tgz download to the host, untar it, and change to the resulting directory.

    b. Stop the running agent.

       ```
       # service cloudstack-agent stop
       ```

    c. Update the agent software.

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

    d. Choose "U" to update the packages.

   e. Upgrade all the existing bridge names to new bridge names by running this script:

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

   f. Install a libvirt hook with the following commands:

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

   g. Restart libvirtd.

   ```
   # service libvirtd restart
   ```

   h. Start the agent.

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

17. Log in to the CloudPlatform UI as administrator, and check the status of the hosts. All hosts must come to the Up state (except those that you know to be offline). You may need to wait for 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 the Up state. If the hosts do not come to the Up state, contact support.

18. (VMware only) Log in to the CloudPlatform UI.

19. Destroy both the Secondary Storage VM (SSVM) and Console Proxy VM (CPVM).

20. (VMware) Run the following script to destroy and re-create all remaining System VMs.

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

    ```
    # nohup cloudstack-sysvmadm -d 192.168.1.5 -u cloud -p password -n -v > 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:

```
nohup: ignoring input
Restarting 4 networks...
Done restarting networks.
Restarting 2 vpcs...
INFO: Restarting vpc with id 2
INFO: Restarting vpc with id 1
INFO: Successfully restarted vpc with id 1
INFO: Successfully restarted vpc with id 2
Done restarting vpcs.
```

21. (XenServer or KVM) Run the following script to stop, then start, all System VMs including Secondary Storage VMs, Console Proxy VMs, and virtual routers.

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

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

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

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

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

    The content should be like the following:

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

22. (XenServer only) Upgrade all existing XenServer clusters to a version supported by CloudPlatform 4.2.1 and apply any required hotfixes.

    You can find the instructions for upgrading XenServer software and applying hotfixes in see *Section 3.7.2, “Applying Hotfixes to a XenServer Cluster”*.

23. (VMware only) After upgrade, if you want to change a Standard vSwitch zone to a VMware dvSwitch Zone, perform the following:

    a. Ensure that the Public and Guest traffics are not on the same network as the Management and Storage traffic.

    b. Set vmware.use.dvswitch to true.

    c. Access the physical network for the Public and guest traffic, then change the traffic labels as given below:

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

For example: dvSwitch18,,vmwaredvs

VLANID is optional.

d. Stop the Management server.

e. Start the Management server.

f. Add the new VMware dvSwitch-enabled cluster to this zone.

Post-Upgrade Considerations

Consider the following:

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

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

- (XenServer) A cluster-level global parameter, xen.vm.vcpu.max, has been added to configure the number of non-Windows VMs in a cluster. After upgrade, the parameter takes the default value of 16. However, you can change the default value by using the Global Settings or Cluster Settings tab in the UI if you manually apply the following database schema:

  ```
  INSERT IGNORE INTO `cloud`.`configuration` VALUES ('Advanced', 'DEFAULT',
   'managementserver',
  'xen.vm.vcpu.max', '16', 'Maximum number of VCPUs that VM can get in XenServer.');
  ```

  Warning

  The schema change might break the future CloudPlatform upgrades. Therefore, back port the changes in 4.2.1 before upgrading.

- Manually update **`systemvm.iso`** as given in *Section 3.6, “Updating SystemVM.ISO”*.

  In the previous 4.x releases, the Management Server version stored in the database version table is in x.x.x format. For example, 4.2.0 and 4.2.0.1 are stored as 4.2.0 as only the first three digits are considered as release version. Therefore, because the Management Server version number is the same for both the releases, the latest systemvm.iso files are not pushed after upgrade. Therefore, you must manually push systemvm.iso after upgrade.

3.2. Upgrade from 3.0.x to 4.2.1-7

Perform the following to upgrade from version 3.0.5, 3.0.6, 3.0.7 Patch E, or 3.0.7 Patch G to version 4.2.1-7.

1. While running the 3.0.x system, log in to the UI as root administrator.

2. Using the UI, add a new System VM template for each hypervisor type that is used in your cloud. In each zone, add a system VM template for each hypervisor used in that zone.

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

   a. In the left navigation bar, click Templates.

   b. In Select view, click Templates.

   c. Click Register template.

      The Register template dialog box is displayed.

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

      The System VM templates includes fixes for the OpenSSL HeartBleed vulnerability issues.

Hypervisor	Description
XenServer	Name: systemvm-xenserver-4.2.1-b Description: systemvm-xenserver-4.2.1-b URL (if using 32-bit system VM template): *http://download.cloud.com/templates/4.2/systemvmtemplate-2015-08-20-4.2.1-xen.vhd.bz2* URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/systemvmtemplate64-2015-08-20-4.2.1-xen.vhd.bz2* Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the XenServer zones. Hypervisor: XenServer Format: VHD

Hypervisor	Description
	OS Type: Debian GNU/Linux 7.0 (32-bit and 64-bit) (or the highest Debian release number available in the dropdown) Extractable: no Password Enabled: no Public: no Featured: no
KVM	Name: systemvm-kvm-4.2.1-b Description: systemvm-kvm-4.2.1-b URL (if using 32-bit system VM template): *http://download.cloud.com/templates/4.2/systemvmtemplate-2015-08-20-4.2.1-kvm.qcow2.bz2* URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/systemvmtemplate64-2015-08-20-4.2.1-kvm.qcow2.bz2* Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, choose All Zones to make the template available in all the KVM zones. Hypervisor: KVM Format: QCOW2 OS Type: Debian GNU/Linux 7.0 (32-bit and 64-bit) (or the highest Debian release number available in the dropdown) Extractable: no Password Enabled: no Public: no Featured: no
VMware	Name: systemvm-vmware-4.2.1-b Description: systemvm-vmware-4.2.1-b URL (if using 32-bit system VM template on earlier VMware version):*http://download.cloud.com/templates/4.2/*

Hypervisor	Description
	systemvmtemplate-2015-08-20-4.2.1-vmware.ova URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/systemvmtemplate64-2015-08-20-4.2.1-vmware.ova* Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, choose All Zones to make the template available in all the VMware zones. Hypervisor: VMware Format: OVA OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown) Extractable: no Password Enabled: no Public: no Featured: no

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

f. If you use more than one type of hypervisor in your cloud, repeat these steps to download the system VM template for each hypervisor type.

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

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

   Run the following commands on every KVM host.

   a. Download the CloudPlatform 4.2.1-7 RHEL 6.3 binaries from *https://www.citrix.com/English/ss/downloads/*.

      Extract the binaries:

```
# cd /root
# tar xvf CloudPlatform-4.2.1-7-rhel6.3.tar.gz
```

b. Create a CloudPlatform 4.2.1-7 qemu repo:

```
# cd CloudPlatform-4.2.1-7-rhel6.3/6.3
# createrepo
```

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

```
[upgrade]
name=rhel63
baseurl=url-of-your-rhel6.3-repo
enabled=1
gpgcheck=0
[cloudstack]
name=cloudstack
baseurl=file:///root/CloudPlatform-4.2.1-7-rhel6.3/6.3
enabled=1
gpgcheck=0
```

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

```
yum upgrade
```

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

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

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

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

6. On the MySQL master, take a backup of the MySQL databases. We recommend performing this step even in test upgrades. This will assist in debugging issues.

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

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

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

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

8. Download CloudPlatform 4.2.1-7 onto the management server host where it will run. Get the software from the following link:

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

You need a *My Citrix Account*[2].

9. Upgrade the CloudPlatform packages. You should have a file in the form of “CloudPlatform-4.2.1-N-OSVERSION.tar.gz”. Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

   ```
   # tar xzf CloudPlatform-4.2.1-N-OSVERSION.tar.gz
   # cd CloudPlatform-4.2.1-N-OSVERSION
   # ./install.sh
   ```

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

10. Choose "U" to upgrade the package

    ```
    >U
    ```

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

11. If you have made changes to your existing copy of the configuration files components.xml, db.properties, or server.xml in your previous version of CloudPlatform installation, the changes will be preserved in the upgrade. However, you need to do the following steps to place these changes in a new version of the file which is compatible with version 4.2.1.

    Note

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

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

    a. Backup your previous version file. For example: (substitute the file name components.xml, db.properties, or server.xml in these commands as needed)

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

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

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

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

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

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

12. Repeat steps *7* - *11* on each management server node.

13. Start the first Management Server.

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

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

    ### Note

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

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

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

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

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

    a. Stop the Management Server:

    ```
    service cloudstack-management stop
    ```

    b. Perform the following on each VMware cluster:

    i. Encrypt the vCenter password:

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

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

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

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

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

iii. Update the existing password with the encrypted one:

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

iv. Confirm that the table is updated:

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

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

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

vi. Change the existing password to the encrypted one:

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

vii. Confirm that the table is updated:

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

c. Start the CloudPlatform Management server

```
service cloudstack-management start
```

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

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

a. Copy the CloudPlatform 4.2.1-N-OSVERSION.tgz download to the host, untar it. Then, change to the resulting directory.

b. Stop the running agent.

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

c. Update the agent software.

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

d. Choose "U" to update the packages.

e. Edit **`/etc/cloudstack/agent/agent.properties`** to change the resource parameter from **`com.cloud.agent.resource.computing.LibvirtComputingResource`** to **`com.cloud.hypervisor.kvm.resource.LibvirtComputingResource`**.

f. Upgrade all the existing bridge names to new bridge names by running this script:

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

g. Install a libvirt hook with the following commands:

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

h. Restart libvirtd.

```
# service libvirtd restart
```

i. Start the agent.

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

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

Note

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

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

18. (VMware only) Log in to the CloudPlatform UI. Destroy both the Secondary Storage VM (SSVM) and Console Proxy VM (CPVM).

19. (VMware) Run the following script to destroy and re-create all remaining System VMs.

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

```
# nohup cloudstack-sysvmadm -d 192.168.1.5 -u cloud -p password -n -v > 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:

```
nohup: ignoring input
Restarting 4 networks...
Done restarting networks.
Restarting 2 vpcs...
INFO: Restarting vpc with id 2
INFO: Restarting vpc with id 1
INFO: Successfully restarted vpc with id 1
INFO: Successfully restarted vpc with id 2
Done restarting vpcs.
```

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

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

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

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

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

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

    The content should be like the following:

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

21. 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.

    Depending on the hypervisor, use one of the techniques mentioned in the following notes:

XenServer or KVM:

SSH in by using the link local IP address of the system VM. For example, in the command below, substitute your own path to the private key used to log in to the system VM and your own link local IP. Run the following commands on the XenServer or KVM host on which the system VM is present:

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

The output should be like the following:

```
Cloudstack Release 4.2.0 Tue August 18 15:10:04 PST 2015
```

ESXi

SSH in using the private IP address of the system VM. For example, in the command below, substitute your own path to the private key used to log in to the system VM and your own private IP. Run the following commands on the Management Server:

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

The output should be like the following:

```
Cloudstack Release 4.2.0 Tue August 18 15:10:04 PST 2015
```

22. If you want to close the admin port again (recommended in production systems), set integration.api.port to null. Then restart the Management Server.

23. (XenServer only) If needed, upgrade all Citrix XenServer hypervisor hosts in your cloud to a version supported by CloudPlatform 4.2.1-7 and apply any required hotfixes. see “Setting Configuration Parameters” in the Installation Guide.

24. (VMware only) After upgrade, if you want to change a Standard vSwitch zone to a VMware dvSwitch Zone, perform the following:

    a. Ensure that the Public and Guest traffics are not on the same network as the Management and Storage traffic.

    b. Set vmware.use.dvswitch to true.

    c. Access the physical network for the Public and guest traffic, then change the traffic labels as given below:

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

    For example, dvSwitch18,vmwaredvs, and VLANID is optional.

    d. Stop the Management server.

    e. Start the Management server.

    f. Add the new VMware dvSwitch-enabled cluster to this zone.

Note

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

Note

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

3.3. Upgrade from 2.2.x to 4.2.1-7

Direct upgrade from 2.2.x to 4.2.1-7 is not supported. You must do the following:

1. Upgrade to 4.2.1

2. Upgrade from 4.2.1 to 4.2.1-7

Upgrade to 4.2.1:

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

   Starting in 3.0.2, the usage record format for IP addresses is the same as the rest of the usage types. Instead of a single record with the assignment and release dates, separate records are generated per aggregation period with start and end dates. After upgrading to 4.2.1, any existing IP address usage records in the old format will no longer be available.

2. If you are using version 2.2.0 - 2.2.14, first upgrade to 2.2.16 by using the instructions in the *2.2.14 Release Notes3*[3].

[3] http://download.cloud.com/releases/2.2.0/CloudStack2.2.14ReleaseNotes.pdf

Note

(KVM only) If KVM hypervisor is used in your cloud, be sure you completed the step to insert a valid username and password into the host_details table on each KVM node as described in the 2.2.14 Release Notes. This step is critical, as the database will be encrypted after the upgrade to 4.2.1.

3. While running the 2.2.x system (which by this step should be at version 2.2.14 or greater), log in to the UI as root administrator.

4. Using the UI, add a new System VM template for each hypervisor type that is used in your cloud. In each zone, add a system VM template for each hypervisor used in that zone.

   Note

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

   a. In the left navigation bar, click Templates.

   b. In Select view, click Templates.

   c. Click Register template.

      The Register template dialog box is displayed.

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

      The System VM templates includes fixes for the OpenSSL HeartBleed vulnerability issues.

Hypervisor	Description
XenServer	Name: systemvm-xenserver-4.2.1-b Description: systemvm-xenserver-4.2.1-b URL (if using 32-bit system VM template): *http://download.cloud.com/templates/4.2/systemvmtemplate-2015-08-20-4.2.1-xen.vhd.bz2* URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/*

Hypervisor	Description
	systemvmtemplate64-2015-08-20-4.2.1-xen.vhd.bz2 Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running XenServer, choose All Zones to make the template available in all the XenServer zones. Hypervisor: XenServer Format: VHD OS Type: Debian GNU/Linux 7.0 (32-bit and 64-bit) (or the highest Debian release number available in the dropdown) Extractable: no Password Enabled: no Public: no Featured: no
KVM	Name: systemvm-kvm-4.2.1-b Description: systemvm-kvm-4.2.1-b URL (if using 32-bit system VM template): *http://download.cloud.com/templates/4.2/systemvmtemplate-2015-08-20-4.2.1-kvm.qcow2.bz2* URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/systemvmtemplate64-2015-08-20-4.2.1-kvm.qcow2.bz2* Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running KVM, choose All Zones to make the template available in all the KVM zones. Hypervisor: KVM Format: QCOW2 OS Type: Debian GNU/Linux 7.0 (32-bit and 64-bit) (or the highest Debian release number available in the dropdown) Extractable: no

Hypervisor	Description
	Password Enabled: no Public: no Featured: no
VMware	Name: systemvm-vmware-4.2.1-b Description: systemvm-vmware-4.2.1-b URL (if using 32-bit system VM template on earlier VMware version):*http://download.cloud.com/templates/4.2/systemvmtemplate-2015-08-20-4.2.1-vmware.ova* URL (if using 64-bit system VM template): *http://download.cloud.com/templates/4.2/64bit/systemvmtemplate64-2015-08-20-4.2.1-vmware.ova* Zone: Choose the zone where this hypervisor is used. If your CloudPlatform deployment includes multiple zones running VMware, choose All Zones to make the template available in all the VMware zones. Hypervisor: VMware Format: OVA OS Type: Debian GNU/Linux 7.0 (32-bit) (or the highest Debian release number available in the dropdown) Extractable: no Password Enabled: no Public: no Featured: no

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

f. If you use more than one type of hypervisor in your cloud, repeat these steps to download the system VM template for each hypervisor type.

Warning

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

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

   Run the following commands on every KVM host.

   a. Download the CloudPlatform 4.2.1-7 RHEL 6.3 binaries from *https://www.citrix.com/English/ss/downloads/*.

      Extract the binaries:

      ```
      # cd /root
      # tar xvf CloudPlatform-4.2.1-7-rhel6.3.tar.gz
      ```

   b. Create a CloudPlatform 4.2.1-7 qemu repo:

      ```
      # cd CloudPlatform-4.2.1-7-rhel6.3/6.3
      # createrepo
      ```

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

      ```
      [upgrade]
      name=rhel63
      baseurl=url-of-your-rhel6.3-repo
      enabled=1
      gpgcheck=0
      [cloudstack]
      name=cloudstack
      baseurl=file:///root/CloudPlatform-4.2.1-7-rhel6.3/6.3
      enabled=1
      gpgcheck=0
      ```

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

      ```
      yum upgrade
      ```

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

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

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

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

8. On the MySQL master, take a backup of the MySQL databases. We recommend performing this step even in test upgrades. This will assist in debugging issues.

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

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

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

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

10. Download CloudPlatform 4.2.1-7 onto the management server host where it will run. Get the software from the following link:

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

    You need a *My Citrix Account*[4].

11. Upgrade the CloudPlatform packages. You should have a file in the form of “CloudPlatform-4.2.1-N-OSVERSION.tar.gz”. Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

    ```
    # tar xzf CloudPlatform-4.2.1-N-OSVERSION.tar.gz
    # cd CloudPlatform-4.2.1-N-OSVERSION
    # ./install.sh
    ```

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

12. Choose "U" to upgrade the package

    ```
    >U
    ```

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

13. If you have made changes to your existing copy of the configuration files components.xml, db.properties, or server.xml in your previous version of CloudPlatform installation, the changes will be preserved in the upgrade. However, you need to do the following steps to place these changes in a new version of the file which is compatible with version 4.2.1.

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

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

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

a. Backup your previous version file. For example: (substitute the file name components.xml, db.properties, or server.xml in these commands as needed)

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

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

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

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

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

14. On the management server node, run the following command. It is recommended that you use the command-line flags to provide your own encryption keys.

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

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

- (Optional) For encryption_type, use file or web to indicate the technique used to pass in the database encryption password. Default: file.

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

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

15. Repeat steps *9* - *13* on each management server node. If you provided your own encryption key in step *14*, use the same key on all other Management Servers.

16. Start the first Management Server.

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

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

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

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

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

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

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

a. Copy the CloudPlatform 4.2.1-N-OSVERSION.tgz download to the host, untar it. Then, change to the resulting directory.

b. Stop the running agent.

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

c. Update the agent software.

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

d. Choose "U" to update the packages.

e. Edit **`/etc/cloudstack/agent/agent.properties`** to change the resource parameter from **`com.cloud.agent.resource.computing.LibvirtComputingResource`** to **`com.cloud.hypervisor.kvm.resource.LibvirtComputingResource`**.

f. Upgrade all the existing bridge names to new bridge names by running this script:

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

g. Install a libvirt hook with the following commands:

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

h. Restart libvirtd.

```
# service libvirtd restart
```

i. Start the agent.

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

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

Note

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

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

20. (VMware only) Log in to the CloudPlatform UI. Destroy both the Secondary Storage VM (SSVM) and Console Proxy VM (CPVM).

21. (VMware) Run the following script to destroy and re-create all remaining System VMs.

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

```
# nohup cloudstack-sysvmadm -d 192.168.1.5 -u cloud -p password -n -v > 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:

```
nohup: ignoring input
Restarting 4 networks...
Done restarting networks.
Restarting 2 vpcs...
INFO: Restarting vpc with id 2
INFO: Restarting vpc with id 1
INFO: Successfully restarted vpc with id 1
INFO: Successfully restarted vpc with id 2
Done restarting vpcs.
```

22. (XenServer or KVM) Run the following script to stop, then start, all System VMs including Secondary Storage VMs, Console Proxy VMs, and virtual routers.

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

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

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

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

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

    The content should be like the following:

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

23. 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.

    Depending on the hypervisor, use one of the techniques mentioned in the following notes:

XenServer or KVM:

SSH in by using the link local IP address of the system VM. For example, in the command below, substitute your own path to the private key used to log in to the system VM and your own link local IP. Run the following commands on the XenServer or KVM host on which the system VM is present:

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

The output should be like the following:

```
Cloudstack Release 4.2.0 Tue August 18 15:10:04 PST 2015
```

ESXi

SSH in using the private IP address of the system VM. For example, in the command below, substitute your own path to the private key used to log in to the system VM and your own private IP. Run the following commands on the Management Server:

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

The output should be like the following:

```
Cloudstack Release 4.2.0 Tue August 18 15:10:04 PST 2015
```

24. (XenServer only) If needed, upgrade all Citrix XenServer hypervisor hosts in your cloud to a version supported by CloudPlatform 4.2.1-7 and apply any required hotfixes. see “Setting Configuration Parameters” in the Installation Guide.

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

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

After you complete this procedure, upgrade from 4.2.1 to 4.2.1-7.

For more information, see *Section 3.1, “Upgrade from 4.2.x.x to 4.2.1-7”*

3.4. Upgrading from 2.1.x to 4.2.1-7

Direct upgrades from version 2.1.0 - 2.1.10 to 4.2.1-7 are not supported. CloudPlatform must first be upgraded to version 2.2.16, then to 4.2.1. From 4.2.1, you can upgrade to 4.2.1-7. For information on how to upgrade from 2.1.x to 2.2.16, see the CloudPlatform 2.2.14 Release Notes.

3.5. Upgrade CloudPlatform Baremetal Agent on PXE and DHCP Servers

If you installed bare metal clusters using a previous version of CloudPlatform, use the following steps to upgrade the baremetal agent in order to get the latest bug fixes for 4.2.1.

1. Logon as root to the host or virtual machine running the Baremetal PXE server and DHCP server.

2. Download CloudPlatform 4.2.1-7 onto the PXE or DHCP server. Get the software from the following link:

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

   You need a *My Citrix Account*[5].

3. Upgrade the CloudPlatform packages. You should have a file in the form of “CloudPlatform-4.2.1-N-OSVERSION.tar.gz”. Untar the file, then run the install.sh script inside it. Replace the file and directory names below with those you are using:

   ```
   # tar xzf CloudPlatform-4.2.1-N-OSVERSION.tar.gz
   # cd CloudPlatform-4.2.1-N-OSVERSION
   # ./install.sh
   ```

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

4. Choose "U" to upgrade the package

   ```
   >U
   ```

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

5. Run the bare metal setup script:

   ```
   cloudstack-setup-baremetal
   ```

3.6. Updating SystemVM.ISO

- On CloudPlatform versions 3.0.5.x and 3.0.7.x **`systemvm.iso`** will get propagated automatically; therefore, no separate procedure is required.

- On CloudPlatform versions 4.2.1.x, perform the following based on the hypervisor that you use:

  - XenServer: No action is required.

  - KVM

    a. On the KVM host, stop the CloudPlatform agent.

    b. Upgrade the CloudPlatform agent.

    c. Restart the CloudPlatform agent.

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

   d. Stop and Start SystemVMs.

- VMware

   a. Stop all the Management Servers.

   b. Remove the old **`systemvm<version>.iso`** file from the **`systemvm`** directory, **`\\<secondary_storage_path>\systemvm\`**.

      Where <version> denotes the Management Server version number.

   c. Start the Management Server.

      Verify if the new **`systemvm.iso`** is pushed to the **`systemvm`** folder in the Secondary Storage directory.

   d. Stop and Start SystemVMs.

3.7. Upgrading and Hotfixing XenServer Hypervisor Hosts

In CloudPlatform 4.2.1, you can upgrade XenServer hypervisor host software without having to disconnect the XenServer cluster. You can upgrade XenServer 5.6 GA, 5.6 FP1, or 5.6 SP2 to any newer version that is supported by CloudPlatform. The actual upgrade is described in XenServer documentation, but there are some additional steps you must perform before and after the upgrade.

3.7.1. Upgrading to a New XenServer Version

To upgrade XenServer hosts when running CloudPlatform 4.2.1:

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

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

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

   ```
   # service cloudstack-management restart
   ```

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

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

   ```
   # xe pool-list
   ```

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

   ```
   # xe host-list
   ```

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

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

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

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

   a. Logon to the CloudPlatform UI as root.

   b. Navigate to the XenServer cluster, and click Actions – Unmanage.

   c. Watch the cluster status until it shows Unmanaged.

6. Upgrade the XenServer software on the master host:

   a. Insert the XenXerver CD.

   b. Reboot the host.

   c. Upgrade to the newer version of XenServer. Use the steps in XenServer documentation.

7. Cancel the maintenance mode on the master host.

8. Reconnect the XenServer cluster to CloudPlatform.

   a. Log in to the CloudPlatform UI as root.

   b. Navigate to the XenServer cluster, and click Actions – Manage.

   c. Watch the status to see that all the hosts come up.

9. Upgrade the slave hosts in the cluster:

   a. Put a slave host into maintenance mode.

      Wait until all the VMs are migrated to other hosts.

   b. Upgrade the XenServer software on the slave.

   c. Cancel maintenance mode for the slave.

   d. Repeat steps *a* through *c* for each slave host in the XenServer pool.

10. You might need to change the OS type settings for VMs running on the upgraded hosts, if any of the following apply:

    - If you upgraded from XenServer 5.6 GA to XenServer 5.6 SP2, change any VMs that have the OS type CentOS 5.5 (32-bit), Oracle Enterprise Linux 5.5 (32-bit), or Red Hat Enterprise Linux

5.5 (32-bit) to Other Linux (32-bit). Change any VMs that have the 64-bit versions of these same OS types to Other Linux (64-bit).

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

- If you upgraded from XenServer 5.6 to XenServer 6.0.2 or higher, do all of the above.

3.7.2. Applying Hotfixes to a XenServer Cluster

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

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

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

   ```
   # service cloudstack-management restart
   ```

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

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

      ```
      # xe pool-list
      ```

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

      ```
      # xe host-list
      ```

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

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

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

   a. Logon to the CloudPlatform UI as root.

   b. Navigate to the XenServer cluster, and click Actions – Unmanage.

   c. Watch the cluster status until it shows Unmanaged.

6. Hotfix the master host:

   a. Add the XenServer hot fixes to the master host.

      i. Assign a UUID to the update file:

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

The command displays the UUID of the update file:

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

ii. Repeat the xe patch-upload command for all other XenServer updates: XS62ESP1005.xsupdate, XS62ESP1003.xsupdate.

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

b. Apply XenServer hot fixes to master host:

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

c. Repeat xe patch-apply command for all the hot fixes.

d. Install the required CSP files.

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

e. Restart the master host.

7. Cancel the maintenance mode on the master host.

8. Reconnect the XenServer cluster to CloudPlatform.

   a. Logon to the CloudPlatform UI as root.

   b. Navigate to the XenServer cluster, and click Actions – Manage.

   c. Watch the status to see that all the hosts come up.

9. Hotfix the slave hosts in the cluster:

   a. Put a slave host into maintenance mode.

      Wait until all the VMs are migrated to other hosts.

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

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

   c. Repeat Step a through b for each slave host in the XenServer pool.

   d. Install the required CSP files.

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

   e. Restart the slave hosts.

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

10. Cancel the maintenance mode on the slave hosts.

11. You might need to change the OS type settings for VMs running on the upgraded hosts, if any of the following apply:

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

    - If you upgraded from XenServer 5.6 GA or 5.6 FP1 to XenServer 6.0.2, change any VMs that have the OS type CentOS 5.5 (32-bit), CentOS 5.6 (32-bit), CentOS 5.7 (32-bit), Oracle Enterprise Linux 5.5 (32-bit), Oracle Enterprise Linux 5.6 (32-bit), Oracle Enterprise Linux 5.7 (32-bit), Red Hat Enterprise Linux 5.5 (32-bit), Red Hat Enterprise Linux 5.6 (32-bit) , or Red Hat Enterprise Linux 5.7 (32-bit) to Other Linux (32-bit). Change any VMs that have the 64-bit versions of these same OS types to Other Linux (64-bit).

About This Release

This section describes the following:

- Issues that are resolved as part of this release.
- Known issues that are reported in this release.
- Addendum that describes the latest updates and corections to the information available with the CloudPlatform (powered by Apache CloudStack) 4.2.1-6 Administrator's Guide.

4.1. Fixed Issues

Issue ID	Description
CS-42674	Problem: SSL v3 does not get disabled on Console Proxy VM (CPVM). Solution: Initializing **`SSLConext.getInstance`**("TLS") is not enough. You must disable the SSLv3.
CS-40194	Problem: When SSVM is destroyed while a snapshot operation is in progress, worker VM snapshot created is left behind. Solution: Storage garbage collector has been enhanced to handle this. While deleting 'ERROR' snapshots, CloudPlatform will also check whether the volume is attached to a VM. If yes, the worker VM snapshot taken on the VM will be deleted.
CS-40139	Problem: The order in which templates are returned in the API response is inconsistent for a given page and pagesize. Solution: Added additional order-by clause in the sql query on the temp_zone_pair column made fetch templates.
CS-39918	Problem: Deploying a VM with RHEL 5.5 to 6.4 template. After the VM is deployed, the OS type is sent to "Other" in the vCenter. Root Cause: **`VMWareGuestOsMapper`** missed the correct entries for RHEL 5.5 to 6.4 OS and CloudPlatform uses that mapping table to pass correct information to invoke vCenter API. Solution: Added the correct mappings for RHEL 5.5 up to RHEL 6.4 OS in VmwareGuestOsMapper. Now, users are able to view the correct OS for VM provisioned in vCenter.
CS-39348	Problem: The size of the volume differs from the value specified by the user.

Issue ID	Description
	Root Cause: Volume gets created with the size mentioned in the custom disk offering instead of the size specified by the user. For customized offering, do not expect the size of the volume to be available with the offering. If it is available with the offering, use it instead of the size of the volume specified by the user. Solution: Do not allow the custom disk size disk offering creation with disk size.
CS-38549	Problem: VM's network and disk statistics shows N/A for an XenServer pool that was upgraded from 6.1 to 6.2. Root Cause: This occurs for all VMs if one of the VBD metrics ref is not valid. Solution: Add the VBD validity check. If it fails, log the message.
CS-38340	Problem: Some operations on VM with VM snapshots result in the corruption of disks. Root Cause: Allowing operations documented in the Restriction on VM with VM snapshots section at *https://cwiki.apache.org/confluence/display/CLOUDSTACK/VM+Snapshots* could result in unpredictable code path execution. This could result in disk corruption. Solution: Do not allow the operations on VM with VM snapshots that would result in unpredictable results. These changes have been ported from Goleta to 4.3.0.2.
CS-34857	Problem: Instance creation on VMware fails when **`vm.instancename.flag`** is set tot true. Root Cause: In 4.2.0.0 MR1, VM’s name in vCenter is **`InternalName-VM`** by default (for example, **`i-2-3-VM`**). However, if the CloudPlatform global configuration **`vm.instancename.flag`** is set to true, a VM's name in vCenter will be the DisplayName of that VM. Solution: The DisplayName of a VM appears as its name in vCenter.
CS-34548	Problem: The global configuration parameter system.vm.default.hypervisor does not take effect. Root Cause: With the first userVM deployment for an account, CloudPlatform creates a router VM. For the router VM deployment,

Issue ID	Description
	CloudPlatform currently considers the same hypervisorType as that of the destination host of userVM and does not honour the value set by the global configuration parameter. Solution: Make sure that if the global configuration paramter system.vm.default.hypervisor is set, it is honored, and system/router VMs are deployed on it.
CS-33312	Problem: Incorrect Japanese keyboard mapping in CloudPlatform. Root Cause: Japanese keys in the javascript key mapping file was not correctly mapped to the output characters in the VM console. Solution: The Japanese keyboard mapping issue has been fixed for VMs with Windows or centOS on VMware hypervisor. Ensure that VMs are deployed with "keyboard=jp" parameter and the keyboard language selected in the console proxy keyboard dropdown menu is Japanese.
CS-33086	Problem: After running the **`removeIpFromNic`** API from the Cludplatform GUI, the IP Address is not freed up properly in the DB. Root Cause: IP address status is not getting updated in the DB when IP Address is released. Solution: This issue is fixed such that after running the **`removeIPfromNic`** API, the IP Address is freed up and the status of IP Adress is properly updated in the DB.
CS-30370	Problem: A scheduled snapshot export fails due to another snapshot export task deleting a required VDI. Root Cause: If there are two or more snapshots for the same volume that are queued up for backing up to secondary storage, CloudPlatform backs up the recent snapshot that is at the top of the queue and deletes the other snapshots VDI from hypervisors. This occurs irrespective of the backing up of the snapshot and this leads to the error message: "The uuid you supplied was invalid". Solution: CloudPlatform checks for snapshot creation timestamp and deletes only those snapshots that got created before the latest, backed up snapshot.

Issue ID	Description
CS-27487	Problem: ListVirtualMachine is slow if VM has many entries in the user_vm_details table and contains many resource tags. Root Cause: user_vm_view is created by joining the user_vm_details tables. This duplicates resource tags many times for each VM and reduces performance. Solution: Do not join the user_vm_details tables to create user_vm_view. This is because the code always find details from the user_vm_details table directly. Also, avoid retrieving duplicate tags to control duplicate resource tag entry.
CS-25882	Problem: The VM state in the database and the agent gets out of sync. Root Cause: VM deployment may involve starting VR. Also, the HA proceess may also try to start the same VR. In case of concurrent operations to start a VR, the VM deployment job does not fail immediately. Instead, the retry logic is applied for checking the op_it_work table to see if an outstanding job has been working on the same VR. If there is any issue with a pending op_it_work item, the retry takes more than an hour and then fails. This occurs even though the VR has already been started by HA process as the current VM state is not verified using the while loop of check op_it_work items, but purely rely on the op_it_work table updates. Solution: In the while loop of check op_it_work items, add logic to check the current VM state. If the VM state appears as Running in the database, exit from the work item check loop and do another retry.
CS-25804	Problem: source_template_id of the templates is null when it is created from snapshot with its corresponding volume removed. Solution: When looking for the source template, search for the volume that is removed. Also clean up the code because it had boilerplate code and was updating the database after creating the snapshot record, instead of doing it together.
CS-25793	Problem: Zoneid is missing in the response when the listSnapshots API is used to list all the snapshots.

Issue ID	Description
	Root Cause: Zoneid corresponding to the snapshot is not set in the response object. Solution: Added the zoneid in the snapshot response object.
CS-25575	Problem: An isolated network with egress policy allow is created. When a router is configured in this network, the router allows the egress traffic by default. Do not create any egress rule on the network. After rebooting the router, it is not allowing the egress traffic. Root Cause: On router reboot without any egress rule created on the network, the iptables rules to allow egress traffic is not created. Solution: On router reboot, configuring rules on router to add egress default rules for allowing traffic.
CS-25422	Problem: After a volume is migrated, the usage table shows the old volume id. Root Cause: After migrating a volume (does not apply for Live Migration), CloudPlatform destroys old volume and creates the new volume with the same details. As part of this process, it should publish Usage Event for the old volume removal and should publish Volume Create event for creating the new volume. However, this was not happening and the usage table was displaying the old volume id. Solution: Fixed this by using Volume State Machine to publish the events instead of publishing at operation level.
CS-24970	Problem: Failed to create snapshot from volume when the task is performed repeatedly in zone wide Primary Storage. Root Cause: While taking snapshot of a volume, CloudPlatform determines the storage pool that contains the volume. Then, CloudPlatform chooses the endpoint to perform snapshot backup operation by selecting any host that has the storage pool mounted on it. In case of zone-wide primary storage, since every host present in the zone will have the storage mounted, a random host will be chosen. If the host chosen is not in the cluster that contains the VM that the volume is attached to, the snapshot operation fails.

Issue ID	Description

	During snapshot creation of a volume, the snapshot the VM that the volume is attached to is created first. In case a host that does not contain the volume is chosen as the endpoint, the look-up for the VM in the cluster is negative. Since the VM could not be located, a worker VM will be created and the disk for which we need to create a snapshot (data path for detached volumes) will be attached to it. This operation fails with the following error in vCenter.
	`File []/vmfs/volumes/10caf80b-6d0a3284/ 250e5eb75b9d43e49acdf4ec723e8313.vmdk was not found.`
	Additional information: In case of attached volume, during VM look-up before taking the VM snapshot, see 'Using owner VM for snapshot operation' in the SSVM logs. Sample:
	`2014-10-09 02:47:11,085 INFO [storage.resource.VmwareStorageProcessor] (agentRequest-Handler-2:null) Using owner VM i-6-782-VM for snapshot operation`
	If the above is missing in the SSVM logs for an attached volume, it indicates the wrong host.
	Solution: While creating the snapshot of a volume, CloudPlatform chooses the endpoint to perform backup snapshot operation by selecting any host that has the storage containing the volume mounted on it. Instead, if the volume is attached to a VM, the endpoint chosen by CloudPlatform should be the host that contains the VM.

CS-22759	Problem: VMs do not get expunged and results in the usage of storage. Root Cause: The thread local context, which is used to handle connections to a vCenter is being re-used intermittently. When it is re-used, if the existing context corresponds to the other vCenter, the operation that is being executed fails. The re-use is occurring because the ThreadLocal context is not cleared after the PingTask completion. Solution: Do the following: 1. Clear the ThreadLocal context after completing the PingTask.

Issue ID	Description
	2. If CloudPlatform ThreadLocal context is being re-used, validate that the context corresponds to the correct vCenter API session.
CS-22755	Problem: Data disk attach failed for the clusters with only zone-wide primary. Root Cause: When copying/creating volume, call the storage allocators to get the appropriate pool. First, call the **`ClusterScopeStoragePoolAllocator`**. It checks for the pod Id. If the pod is null, **`ZoneWideStoragePoolAllocator`** is used. Pod passed to **`ClusterScopeStoragePoolAllocator`** is not null, **`ClusterScopeStoragePoolAllocator`** is allocating the pool for Data disk, which is resulting in failure later. The pod, which is getting passed to **`ClusterScopeStoragePoolAllocator`**, is not null, because VM's pod ID is passing the vm's pod id, which is obviously not null and resulting in pool's get allocated from Cluster scope. Solution: Use pool's pod id instead of VM's pod id.
CS-22659	Problem: FTP inbound service is not working in isolated networks. Creating ingress rules for FTP on public IP and accessing FTP service in VM from the public network is failed. Root Cause: FTP connection tracking modules are missing in VR, due to which the FTP data connections are failing. Solution: Load nf_nat_ftp module in VR by using the modprobe command. The nf_nat_ftp and nf_conntrack_ftp modules are loaded.
CS-22560	Problem: Unable to create snapshot of a volume. Root Cause: Snapshot operation on Datadisk volumes DATA-20 and DATA-21 fails at the Management Server layer. This failure is not related to NPEs being thrown by SSVM. Each of the above mentioned DATA volumes has a snapshot that has inconsistent entries in the snapshots and **`snapshot_store_ref`** tables.

Issue ID	Description
	There is a snapshot for each of these volumes that is marked as removed in snapshots table, but in snapshot_store_ref table it appears as Ready for Primary storage and Creating for the Secondary storage. Because of this mismatch, CloudPlatform throws an NPE while trying to retrieve parent snapshot details during the snapshot operation. Solution: For both disks, remove the snapshots entries in the snapshot_store_ref table that are inconsistent. For example, `delete from snapshot_store_ref where snapshot_id in (34, 35);`
CS-22503	Problem: Not able to remove NIC from VM. Root Cause: While removing NIC, the query to check for PF rules only considers destination IP. It is possible to have same destination IPs in different networks. Solution: Include network id also as a part of the filter in the query. This will remove the IPs outside NIC's network to be considered.
CS-21346	Problem: After rebooting CloudPlatform, a delay is noticed in connecting hosts. Root Cause: The following points explain this issue: 1. The overall ping timeout is computed using the formula "total ping timeout = ping.interval * ping.timeout". 2. During restart, the Management Server picks up the hosts in the "Disconnected" state, whose 'last ping' timestamp is less than this cutoff timeout, for reconnecting (that is, "last ping < (current time - total ping timeout)") 3. At the time of restart, the Management Server resets the last ping timestamp of hosts by setting it to "current time - 10mins" based on an already existing logic. 4. From 2 and 3, it can be inferred that time taken for hosts to get reconnected is effectively "total ping timeout - 10mins". This cases the delay in connecting the hosts.

Issue ID	Description
	5. Also, it can be inferred that for "total ping timeout < 10 mins", the reconnect will happen immediately. Solution: Remove the hardcoding of 10 mins from the code. Instead, use the configurations ping.timeout and ping.interval.

CS-21335	Problem: When createTemplate is used to create templates from volumes, all of them get the same name. However, only one template gets synchronized during the template synchronization because template synchronization is performed based on the unique name. All otherv templates that carry the same name gets deleted. This occurs only for VMWare and for createTemplate from volume.
	Root Cause: The templates get the same name because UUID was generated using the name of template. Same UUID gets generated for same name templates. In all other hypervisors and other template creation operations the UUID is first generated by CloudPlatform and stored in the database. But, once template creation from snapshot/volume is performed by hypervisor, the unique name is changed to the name given by the hypervisor. This does not occur in VMWare for createTemplate from volume, where it takes the name generated from CloudPlatform itself.
	Solution: The unique name generation logic has been standardized used by registertemplate. This generates unique name as follows:
	`private static String`
	` generateUniqueName(long id, long userId,`
	` String displayName) {`
	`StringBuilder name = new StringBuilder();`
	`name.append(id);`
	`name.append("-");`
	`name.append(userId);`
	`name.append("-");`
	`name.append(UUID.nameUUIDFromBytes`
	` ((displayName +`
	` System.currentTimeMillis()) .getBytes()).`
	` toString());`
	`return name.toString();`

CS-21288	Problem: Unable to delete/archive events of a deleted user. Solution: To include deleted accounts, call customSearchIncludingremoved instead of customsearch in the getAccountIdsForDomains method.

Issue ID	Description
CS-21207	Problem: The Projects list box does not display all the projects. Solution: Keep on calling the listprojects API until all the projects are displayed.
CS-21106	Problem: Unable to add VMware host to an existing cluster when Nexus 1000v is used as backend. Root Cause: Unlike the addCluster API, the work flow does not supply the VSM metadatain the case of addHost API. This missing information results in failure while performing validation during the discovery process. Solution: Nexus VSM instance would have been validated and added to cluster while adding the cluster. So, Cloudplatform does not need to do this while adding a host to that cluster. Just check if the cluster already has a Nexus VSM instance associated with it. If yes, skip validation and attempt to add Nexus VSM instance to cluster.
CS-21072	Problem: CloudPlatform enters host name in the addHost string three times. Root Cause: AddHost API command fails if either the VMware DC name or the VMware Cluster name contains spaces or any other special characters. That is, AddHost API command fails if the URL encoded value for VMware DC or VMware Cluster is different from the non-decoded value. Solution: The URL decodes vCenter path during cluster discovery before trying to find a match.
CS-20846	Problem: SourceNAT,StaticNAT, and Portfowrding is not working when CloudPlatform is deployed with Vmware DVS. Root Cause: There is a change in vCenter API for vCenter 5.5 regarding virtual NIC association with distributed virtual switch, which breaks the existing functionality that works till vCenter 5.1. Fix this problem to ensure compatibility with 5.5 or later versions of vCenter. Solution: Modified the code that reads the value of the property VirtualE1000.deviceInfo.summary to ensure compatibility with 5.5 or later versions of vCenter.

Issue ID	Description
CS-20783	Problem: On rebooting the virtual router from the CloudPlatform GUI, the passwd_server service attempts to start, but terminates with the exit code 137. Root Cause: This is caused by the delay in the execution of the awk command during certain load conditions such as the boot time. When running the **`ips=$(ip addr show dev eth0 | grep inet | grep eth0 | awk '{print $2} ' )`** command, a write to the pipe occurs even before the reading the **`awk`** command. As a result, the process receives a SIGPIPE signal. The default behaviour of the process on receiving a SIGPIPE is to terminate immediately. Solution: One option is to ignore the SIGPIPE signal and let the process continue. Also, you can try using a command other than **`awk`**.
CS-20695	Problem: VOLUME.DELETE usage event is missing for the VMs in the ERROR state. Root Cause: When deploying a virtual machine fails, a DB entry is updated for this VM as ERROR. Also, marks the volume entry for this VM as Destroyed. But in our code, we are not publish VOLUME.EVENT event in invoking internal destroyVolume command. Instead, it scatters VOLUME.EVENT publish code around in service layer code. Since we missed calling publish VOLUME.EVENT in this particular case, no VOLUME.EVENT is published when we failed to deploy a VM. Solution: Refactor code to always publish VOLUME.EVENT event in invoking internal destroyVolume routine to mark a volume as Destoyed.
CS-20670	Problem: Name-tracking in various CloudPlatfrom orchesteration flows is affected when vCenter performs snapshot operation and stacks a delta disk on top of the chain. Solution: While doing a disk look-up in vCenter by matching against its name, trim the postfix appended to the disk name by vCenter after snapshot operation.
CS-20653	Problem: Scroll to the end of a list view to load more list rows. The newly-loaded list rows do not have the secondary IP dropdown.

Issue ID	Description
	Root Cause: Outdated UI data is passed to subselect's data provider when infinite scroll event loads more data items. Thus, the widget cannot find out the network/VM ID to query for the secondary IPs. Solution: Pass the most up-to-date data to the list view widget.
CS-20626	Problem: The queryAsyncJobResult API function does not return the jobinstanceid response. Root Cause: This occurs because the asyncjob framework resets the instance_id and the instance_type fields in the async_job table after completing an async job. Solution: No need to reset the instance_id and the instance_type fields after completing an asyncjob. So, removed that part of code.
CS-20607	Problem: The value set in the Default field of Maximum guest limit for XenServer has no effect on XenServers that are added to CloudPlatform before setting the value. Only Xenservers that are added to CloudPlatform after setting the value respond to the new value. Root Cause: The logic to compare the number of running VMs to the maximum limit was not correct. Solution: Corrected the logic to compare to the maximum limit.
CS-20599	Problem: Events generated while creating snapshot from volume had the volume 'id'. This 'id' is internal to CloudPlatform. The volume 'UUID' is expected to be in the events. Root Cause: In Events, the internal entity 'id' has not been replaced with with 'UUID'. Solution: A specific fix made to use volume UUID in volume snapshot related events.
CS-20534	Problem: A quick view stays open after hovering the mouse off a row. This confises the users and they perform tasks from the wrong quickview. Root Cause: This was caused by extra margin padding around the quickview. Solution: Removed the extra padding around the quickview. After the mouse is out of the box, it is hidden immediately.

Issue ID	Description
CS-20531	Problem: VMWare worker VMs are left behind in vCenter and are not cleaned up properly. Root Cause: All worker VMs will not be recycled after restarting Management Servers that belong to a clustered Management Server pool. The worker VMs whose Management Server is the same as the one that owned by their hosts will be recycled. Solution: Code is modified to recycle all worker VMs that belong to the hosts managed by the Management Server irrespective of the original Management Server.
CS-20526	Problem: Static NAT does not work after failover in RVR. This issue apears only for the additional public subnets. Root Cause: There is no mechanism to add routes in **`enable_pubip.sh`** for additional public subnets in RVR when failover occurs. When the back up switches to master, **`enable_pubip.sh`** is called for bringing up public interfaces and adding routes for the interface. However, **`enable_pubip.sh`** fails to add route to eth3. Due to this, the ingress traffic that comes in via eth3 goes out via eth2. Solution: Add routes for additional subnets using the gateway and the device information in VR so that the gateway and the device information are maintained in **`/var/cache/cloud/ifaceGwIp`** in the VR. Use this information to add routes for additional public subnet interfaces in **`enable_pubip.sh`** when VR switches to master.
CS-20525	Problem: On an additional public subnet, removing public IP does not delete SNAT rules. The issue is observed when the first IP is added, but while removing the first IP also is removed. Root Cause: For additional public subnets (nonsourceNAT network) the first one is selected as the first IP from the list, which is retrieved from the database. When you have few IPs the delete/add operations on them changes the order of the IPs. A static NAT rule is configured on the last IP first, and later few static NAT rules are configured on other IPs. While removing, the last IP is removed first so that it is not selected as the first IP. Therefore, the SNAT rules configured are untouched on the VR. While

Issue ID	Description
	adding, source NAT rules are added for the first IP. Solution: To delete SNAT rules, while disabling static NAT on IP from the non source NAT network, set the source NAT flag to true. This is to make sure that the SNAT rules are got removed.
CS-20524	Problem: Enabling static NAT on additional public IP removes the SNAT rule. Root Cause: The SNAT rule, which is configured during ipassoc of this IP address, is getting removed. While configuring static NAT, the **`firewall.sh`** script incorrectly searches and finds (**`grep`**) the SNAT rule and removes it. Solution: Corrected the **`firewall.sh`** script to correctly search and find (**`grep`**) the static NAT rule related SNAT rule for removal.
CS-20478	Problem: User instance created before 4.2.1 GA fails to connect to the guest network if **`nicAdapter`** used by the instance is other than default **`nicAdapter`**, which is E1000. Root Cause: From 4.2.1 onwards, template level setting of **`nicAdapter`** has been used while creating a user instance. Also, **`nicAdapter`** property persisted in the **`user_vm_details`** table during the instance creation. All subsequent start operation would use it rather than depend upon template level setting of **`nicAdapter`**. Instances created before 4.2.1 would not have any information about **`nicAdapter`** in the **`user_vm_details`** table, which results in switching to default **`nicAdapter`** (which is E1000) upon a stop and start operation after upgrading to 4.2.1. Solution: This needs data migration during upgrade. This is a workaround for this issue. Update database (user_vm_details table) so that nicAdapter property (which is from template_details table whose value is not 'E1000') be stored. `> insert into cloud.user_vm_details (vm_id,name,value,display_detail) VALUES (<VM_ID>,'nicAdapter', <nicAdapter_Property_From_template_details_table>,1);` An Example SQL appears as follows:

Issue ID	Description

	`> insert into cloud.user_vm_details`
	`(vm_id,name,value,display_detail) VALUES`
	`(1111,'nicAdapter','Vmxnet3',1);`
	`Prior to 4.2.1 user`
	`instance used tonicAdapter`

CS-20122	Problem: When a VM Snapshot is stuck in the Creating state, observed connection issues (host does not connect to Management Server) with the host where the VM in question is available. Root Cause: A failed VM snapshot is not moving VM snapshot into the Error state, which is resulting in host connection issues. Solution: Handling timeout exceptions and other snapshot error scenarios such that unless the operation succeeds, VM snapshot will be marked as in the Error state.
CS-20100	Problem: Usage is generated for volumes even after the volume is destroyed and moved to the Expunged status. Root Cause: When a VM is reset, old ROOT volume is deleted and new volume is created for the VM. However, usage events for old volume deletion and new volume creation are missing. Due to this old volume usage does not stop. Solution: Add volume usage events during VM reset.
CS-20093	Problem: When trying to attach multiple data disks to a VM in quick succession, AttachVolumeCmds consistently fails. Root Cause: This occurs when attaching volumes to a VM in fast succession. In VMware, the following steps are performed while attaching disks to a VM: 1. Preparing the disk that needs to be attached. 2. Re-configuring the VM to attach the prepared disk device. Step 1 involves figuring out the device number on the controller key that the disk would be connected to. Step 1 and Step 2 are not synchronized in the CloudPlatform code. So, when attaching two disks in quick succession, while preparing the second disk, the device number of the first disk gets selected for the second disk too. This occurs if the VM is still being re-configured with the first disk (that is,

Issue ID	Description
	Step 2 is in progress for the first disk). This results in the Invalid configuration for device 'x' error, where x is the device number on the controller key that the first disk is connected to. Solution: Synchronize the tasks of disk preparation and reconfiguration of a VM with the disk. This ensures that CloudPlatform does not attach a disk to a VM with a wrong device number on the controller key.
CS-20084	Problem: The **`catalina.out`** log file (available at **`/var/log/cloudstack/management/catalina.out`**) keeps on growing without a proper daily rotation and compression. Root Cause: There is no log rotate configuration file for the **`catalina.out`** file created by CloudPlatform. So, this file keeps on growing to a large size and causes problems to the Management Server. Solution: Fixed this issue by creating tomcat-cloudstack logrotate configuration file at **`/etc/logrotate.d/tomcat-cloudstack`**. CloudPlatform management **`catalina.out`** log rotates as expected.
CS-20046	Problem: CloudPlatform allows creation of VMs with the same Display name when the vm.instancename.flag is set to 'true'. Root Cause: This occurs because CloudPlatform restricts two VMs from having the same name only if they are in the same network. But, if **`vm.instancename.flag`** is set to 'true', VM name in vCenter will be its display name. And vCenter does not allow two VMs to have the same name under the same DC, which results in the failure of VM start in CloudPlatform and is not handled correctly. Solution: During VM creation, if vm.instancename.flag is set to 'true' and hypervisor type is VMware, check if VM with the same hostname already exists in the zone.
CS-19950	Problem: For a clustered environment using HAProxy, in case of network glitch, SSVM/CPVM agents will reconnect back to Management Server, but may connect back to a Management Server node different from the one before network outage. Previous Management Server will still keep sending messages to those agents, causing the "Channel is closed" error.

Issue ID	Description
	Root Cause: The issue is caused by our AgentManager code has an internal cache for agents managed by a Management Server. In case of disconnect caused by network outage and reconnecting back to a different Management Server node, CloudPlatform code cannot robustly invalidate the old cache and the old Management Server node still sends messages to agents that are already not managed by it. Solution: Before sending message to agents, Management Server will check the database to see the node that is currently connected by agents. If the Management Server node ID is changed, it will forward the message to the new Management Server.
CS-19928	Problem: When running restartnetwork with cleanup set to 'true' in the Basic or the Advanced shared network, DNS of VMs in the network may failed to work. Root Cause: The IP address of the VR would change during re-creation because it is not limited to the first IP address in the network as in the Isolated or the VPC case. When this occurs, the DNS entries of old VMs would be pointed to old VR's IP address, which results in the DNS failure. Solution: Add a placeholder NIC for the IP address assigned to VR, when it get assigned for the first time, in the Basic or shared network. Then, use the same IP address later whenever the VR has been recreated. This help keeping the VR's IP address same for the lifetime of network.
CS-19913	Problem: Volume migration between storage pools times out on CloudPlatform, but the migration successfully completes on Xenserver. Root Cause: Live migration of a volume that is attached to a running VM from one storage pool to another is a long running task. By default, the agent waits for (30*2) minutes for any command to finish on the resource. If volume migration takes longer than that the agent assumes, the command does not finish and moves ahead. Later, the command completes on the resource, but the agent does not process the answer. Solution: Fixed this issue by making sure that the agent will wait for a period equal to two times the

Issue ID	Description
	migratewait interval for finishing the volume migration request. **migratewait** is a global configuration paramater, which defaults to 60 minutes.
CS-19876	Problem: For Basic zone setup, multiple router VMs were getting started for user VMs and the user VM deployment was ultimately failing. Root Cause: For Basic zone, contrary to the normal behaviour, router VM was not starting in the same Pod as the user VM. Now, if the router starts in a different Pod than the user VM, the VM fails to start because the DHCP service verification fails. When VM fails to start, retry three times and as a part of it three routers get started. This caused accumulation of routers on every VM start leading to many operational routers. PodId in which the router should get started was not being saved to the database. So, when planner loaded the router from the database, it always got podId as null and that would allow planner to deploy the router in any pod. Solution: Fixed the VO's setPodId method that was causing the failure of the DB save operation. The method's name was not according to expected format.
CS-19815	Problem: Log messages and alerts are created for SSVM disconnection for the running SSVMs. Root Cause: Sometimes, Management Server will not repond to agent's **ping** command on time. In KVM's case, libvirt will take time to respond, which will cause the agent reconnecting to Management Server and eventually failing all the on-going tasks on the host. Solution: Do not try reconnecting in such cases.
CS-19655	Problem: Failed to start an instance after restoring the running instance. Root Cause: If the new ROOT disk is not created in the same storage pool where the old disk was, value of the disk path contained in the disk information becomes invalid. Solution: If there are two or more primary storage pools in the same cluster, during restore it fails with an error message. Because of the mismatch in disk location between the old and the new

Issue ID	Description
	ROOT disk, the VMDK not found exception displayed.
CS-19533	Problem: Multiple threads run to collect statistics from VR in one Management Server. This leads to the consumption of direct agent threads unnecessarily. Root Cause: Same threads are being instantiated by the multiple managers (VirtualNetworkApplianceManagerImpl and VpcVirtualNetworkApplianceManagerImpl) in the Management Server. Because of this, two threads are observed for tasks such as network usage task, Network stats update task, and Check router task. Solution: Remove the duplicate task scheduled by VpcVirtualNetworkApplianceManagerImpl. After applying the fix, only one thread instantiated by VirtualNetworkApplianceManagerImpl is observed.
CS-19349	Problem: The VM.instancename global setting parameter for VSphere uses instancename without the prefix 'iX-'. Root Cause: By default, the format of VM's name in vCenter is 'InternalName-VM' (for example, i-2-3-VM). But, if CloudPlatform global configuration vm.instancename.flag is set to 'true', a VM's name in vCenter will be its DisplayName. Solution: VM's name in vCenter should be its DisplayName.
CS-19250	Problem: Iptables rules occasionally fail to program on XenServer host. Root Cause: XenServer uses a combination of VM instance name, port, and chain name to compose iptable chain name. This name often exceeds the maximum character limit of 28 characters and cause iptables programming failure. Solution: Checking iptables chain name before applying the chain to iptables. If the name exceeds 28 characters, truncate it.

4.2. Known Issues

The following table displays the known issues in this maintenance release:

Issue ID	Description
CS-43949	Problem: In CloudPlatform 4.2.1, it is possible to deploy VMs without specifying HV type. Note that the users can specify HV type explicitly as part of deployVirtualMachine API or obtain HV type from the template/ISO that is used to deploy the VM. In such cases, the HV type is set to 'None' or null value in the database. For such VMs, expunge operation fails with NPE (as observed in the CloudPlatform logs). Typically expunge operations run periodically as part of some background thread (check **`config expunge.interval`** for frequency). Workaround: Update the affected VM entries in the database with the correct HV type.
CS-43863	Problem: After registering a template in ALL ZONES in a multi-zone environment, the CloudPlatform UI does not display the template in all available zones. However, the database displays that the templates are available for all the zones. Workaround: The template is created only for one zone despite selecting the ALL ZONES option in the Register Template wizard. To address this problem, select the required zones while registering the template.

4.3. CloudPlatform 4.2.1-7 Release Notes Addendum

4.3.1. Behaviour of VDS configured with Cloudplatform

VMware vNetwork Distributed Switch (VDS) configured with CloudPlatform supports only the public and the guest traffic. VDS does not support the management and the storage traffic.

4.3.2. Providing a Display Name to the Guest Virtual Machines

> **Note**
>
> In the *CloudPlatform (powered by Apache CloudStack) 4.2.1-6 Administrator's Guide*, read the **11.6. Appending a Display Name to the Guest VM’s Internal Name** section as follows:

Every guest VM has an internal name. The host uses the internal name to identify the guest VMs. CloudPlatform gives you an option to provide a display name to guest VMs. You can add this display name to the internal name so that it is displayed in vCenter. This feature is intended to make the correlation between instance names and internal names easier in large datacenter deployments.

To provide display names to a VM, you need to set the global configuration parameter **`vm.instancename.flag`** to **`true`**. The default value of this parameter is **`false`**.

The default format of the internal name is **`i-<user_id>-<vm_id>-<instance.name>`**, where **`instance.name`** is a global parameter. After you set the *`vm.instancename.flag`* parameter to **`true`** and provide a display name during the creation of a guest VM, the display name will be displayed in vCenter for the guest.

> **Note**
>
> The VMs that are deployed for CloudPlatform version 3.0.7 will continue displaying the VM names in vCenter in the **`InternalName-DisplayName`** format. The VMs that are deployed for the higher versions of CloudPlatform will display their display name in vCenter.

The following table explains how a VM name is displayed in different scenarios.

In the following table, Display Name represents the user-supplied display name.

User-Provided Display Name	vm.instancename.flag	Host name on the VM	Name on vCenter
Yes	True	Display name	Display Name
No	True	<instance.name>-<UUID>	<instance.name>-<UUID>
Yes	False	Display name	i-<user_id>-<vm_id>-<instance.name>
No	False	<instance.name>-<UUID>	i-<user_id>-<vm_id>-<instance.name>

4.3.3. Limitation on Dynamically Assigning vCPUs on Windows VMs

Read the following limitation as part of the **11.11.5. Limitations** section in the *CloudPlatform (powered by Apache CloudStack) 4.2.1-6 Administrator's Guide*:

- Dynamically assigning vCPUs on Windows VMs does not work in case of XenServer. This is because the Windows VMs need to be restarted to make the changes take effect.