



Zenoss Resource Manager Upgrade Guide

Release 5.1.3

Zenoss, Inc.

www.zenoss.com

Zenoss Resource Manager Upgrade Guide

Copyright © 2016 Zenoss, Inc. All rights reserved.

Zenoss and the Zenoss logo are trademarks or registered trademarks of Zenoss, Inc., in the United States and other countries. All other trademarks, logos, and service marks are the property of Zenoss or other third parties. Use of these marks is prohibited without the express written consent of Zenoss, Inc., or the third-party owner.

Flash is a registered trademark of Adobe Systems Incorporated.

Oracle, the Oracle logo, Java, and MySQL are registered trademarks of the Oracle Corporation and/or its affiliates.

Linux is a registered trademark of Linus Torvalds.

RabbitMQ is a trademark of VMware, Inc.

SNMP Informant is a trademark of Garth K. Williams (Informant Systems, Inc.).

Sybase is a registered trademark of Sybase, Inc.

Tomcat is a trademark of the Apache Software Foundation.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

All other companies and products mentioned are trademarks and property of their respective owners.

Part Number: 1092.16.146

Zenoss, Inc.
11305 Four Points Drive
Bldg 1 - Suite 300
Austin, Texas 78726

Contents

About this guide.....	7
Chapter 1: Supported software and upgrade paths.....	9
Release dates and versions.....	9
Supported upgrade paths.....	9
Part I: Upgrading only Control Center.....	12
Chapter 1: Upgrading Control Center with internet access.....	13
Stopping applications.....	13
Upgrading the master host.....	13
Upgrading resource pool hosts.....	14
Chapter 2: Upgrading Control Center without internet access.....	16
Downloading the RPM package.....	16
Staging the RPM package.....	16
Upgrading the master host.....	17
Upgrading resource pool hosts.....	17
Chapter 3: Upgrading high-availability deployments with internet access.....	19
Stopping applications.....	19
Stopping Control Center.....	19
Upgrading the master nodes.....	20
Upgrading resource pool hosts.....	21
Restarting Control Center.....	21
Chapter 4: Upgrading high-availability deployments without internet access.....	23
Downloading the RPM package.....	23
Staging the RPM package.....	23
Stopping applications.....	24
Stopping Control Center.....	24
Upgrading the master nodes.....	25
Upgrading resource pool hosts.....	25
Restarting Control Center.....	26

Part II: Upgrading only Resource Manager.....27

Chapter 1: Upgrading Resource Manager with internet access..... 28

Stopping applications.....28
Upgrading Resource Manager.....28

Chapter 2: Upgrading Resource Manager without internet access.....30

Downloading files for offline installation..... 30
Staging archive files.....30
Upgrading Resource Manager.....31

Chapter 3: Upgrading high-availabilty deployments with internet access..... 33

Stopping applications.....33
Upgrading Resource Manager.....33

Chapter 4: Upgrading high-availabilty deployments without internet access..... 35

Downloading files for offline installation.....35
Staging archive files.....35
Stopping applications.....36
Upgrading Resource Manager.....36

Chapter 5: After upgrading..... 38

Deleting the pre-upgrade snapshot.....38

Part III: Upgrading only the Pacemaker resource agent for Control Center..... 39

Chapter 1: Upgrading high-availabilty deployments with internet access..... 40

Identifying the Pacemaker resource agent version.....40
Stopping applications.....41
Upgrading the Pacemaker resource agent.....41

Chapter 2: Upgrading high-availabilty deployments without internet access..... 42

Identifying the Pacemaker resource agent version.....42
Identifying the currently-available version of the Pacemaker resource agent.....42
Staging the RPM package file.....43
Stopping applications.....43
Upgrading the Pacemaker resource agent.....43

Part IV: Upgrading Control Center and Resource Manager.....	45
Chapter 1: Preparing to upgrade.....	46
Storage changes.....	46
Retaining a customized Control Center web server port.....	47
Ensuring localhost resolution.....	48
Important information and recommendations.....	48
Chapter 2: Upgrading Resource Manager with internet access.....	50
Upgrading Control Center on the master host.....	50
Upgrading Control Center on resource pool hosts.....	56
Converting the data storage driver.....	61
ZooKeeper ensemble configuration.....	64
Master host isolation.....	71
Upgrading Resource Manager.....	73
Chapter 3: Upgrading Resource Manager without internet access.....	74
Upgrading Control Center on the master host.....	74
Upgrading Control Center on resource pool hosts.....	81
Converting the data storage driver.....	86
ZooKeeper ensemble configuration.....	89
Master host isolation.....	95
Upgrading Resource Manager.....	97
Chapter 4: Upgrading high-availability deployments with internet access.....	99
Key variables used in this chapter.....	99
Upgrading Control Center on the master nodes.....	100
Updating high-availability components.....	105
Upgrading Control Center on resource pool hosts.....	112
Converting the data storage driver.....	117
ZooKeeper ensemble configuration.....	121
Upgrading Resource Manager.....	126
Chapter 5: Upgrading high-availability deployments without internet access.....	128
Key variables used in this chapter.....	128
Downloading files for offline installation.....	129
Upgrading Control Center on the master nodes.....	129
Updating high-availability components.....	136
Upgrading Control Center on resource pool hosts.....	142
Converting the data storage driver.....	147
ZooKeeper ensemble configuration.....	151
Upgrading Resource Manager.....	156
Chapter 6: After upgrading.....	158
Preventing OpenTSDB health check failures.....	158
Checking ZooKeeper quorum keys.....	159

Correcting a CentralQuery configuration file.....	159
Configuring ZooKeeper log purging.....	160
Increasing Zope server threads.....	160
Updating the daily maintenance script.....	160
Creating a weekly maintenance script.....	161
Deleting the pre-upgrade snapshot.....	161
Deleting the pre-upgrade registry.....	161

Appendix A: Using Zenoss Toolbox..... 163

Zenoss Toolbox tools.....	163
Downloading Zenoss Toolbox with internet access.....	163
Downloading Zenoss Toolbox without internet access.....	164
Installing Zenoss Toolbox.....	164
Running Zenoss Toolbox tools.....	165

About this guide

Zenoss Resource Manager Upgrade Guide provides detailed instructions for upgrading Zenoss Resource Manager (Resource Manager) from one minor or micro version to a more recent version. All supported deployment configurations are included in this guide.

Note Zenoss strongly recommends reviewing the *Zenoss Resource Manager Planning Guide* carefully before using this guide.

Related publications

Title	Description
<i>Zenoss Resource Manager Administration Guide</i>	Provides an overview of Resource Manager architecture and features, as well as procedures and examples to help use the system.
<i>Zenoss Resource Manager Configuration Guide</i>	Provides required and optional configuration procedures for Resource Manager, to prepare your deployment for monitoring in your environment.
<i>Zenoss Resource Manager Installation Guide</i>	Provides detailed information and procedures for creating deployments of Control Center and Resource Manager.
<i>Zenoss Resource Manager Planning Guide</i>	Provides both general and specific information for preparing to deploy Resource Manager.
<i>Zenoss Resource Manager Reference Guide</i>	Provides information about a variety of Resource Manager technologies and options.
<i>Zenoss Resource Manager Release Notes</i>	Describes known issues, fixed issues, and late-breaking information not already provided in the published documentation set.
<i>Zenoss Resource Manager Upgrade Guide</i>	Provides detailed information and procedures for upgrading deployments of Resource Manager.

Additional information and comments

If you have technical questions about this product that are not answered in this guide, please visit the [Zenoss Support](#) site or contact Zenoss Support.

Zenoss welcomes your comments and suggestions regarding our documentation. To share your comments, please send an email to docs@zenoss.com. In the email, include the document title and part number. The part number appears at the end of the list of trademarks, at the front of this guide.

Change history

The following list associates document part numbers and the important changes to this guide since the previous release. Some of the changes involve features or content, but others do not. For information about new or changed features, refer to the *Zenoss Resource Manager Release Notes*.

1092.16.146

Add new version numbers.

1092.16.126

Add a new part for upgrading only the Pacemaker resource agents for Control Center.

Refine the procedure for creating the application data thin pool.

1092.16.118

Add a new part for upgrading only Resource Manager.

Modify supported upgrade paths for Resource Manager 5.1.2.

Add a substep to create the docker override directory.

1092.16.111

Add information about Control Center 1.1.3.

1092.16.102

Replace a placeholder instruction with the correct instruction.

Add the `install-registry` image to the list of offline images.

Include instructions to create a `systemd` override file instead of editing the Docker service definition file.

Correct the `Btrfs` filesystem cleanup command.

Add a symlink to `/tmp` in `/var/lib/docker`.

Update the commands for starting and testing a ZooKeeper ensemble.

Add a procedure for updating the `SERVICED_ZK` value on resource pool hosts that are not members of a ZooKeeper ensemble.

Add a reference topic for the ZooKeeper variables required on hosts in a Control Center cluster.

Improve high-availability procedures to ensure DRBD integrity.

Improve titles in the post-upgrade chapter.

1092.16.068

Add a list of links to the overview of Part I.

After converting the storage driver, perform a full backup.

The Docker configuration file needs a longer startup timeout value, to work around a known Docker issue with the `devicemapper` driver. All Docker configuration steps now include adding `TimeoutSec=300`.

1092.16.067

A new part is added, for upgrading only Control Center. Both parts are renamed to reflect the addition.

The scope of supported upgrade paths is changed to reflect the micro release of Control Center.

All Docker configuration steps now add the storage driver flag (`-s devicemapper`) to the `/etc/sysconfig/docker` file.

All resource pool host upgrade procedures include a step to unmount the distributed file system before restarting `serviced`.

A link the the post-upgrade chapter is added to the end of upgrade procedures, if one is available.

1092.16.060.1

Upgrades are grouped in parts by scope. Each part contains a preparation chapter, chapters for the supported upgrade paths, and a post-upgrade chapter. Only the latest scope is in this version of the guide; previous scopes are in earlier versions.

New procedures are included, for upgrading without internet access and for upgrading high-availability deployments.

A description of Zenoss Toolbox is included as an appendix.

1

Supported software and upgrade paths

Beginning with version 5.0.0, distributions of Resource Manager include an additional component, Control Center. Each component is developed and maintained separately, and each has its own version number. This chapter identifies the combinations of component versions that Zenoss supports, and the supported upgrade paths between the combinations.

Release dates and versions

Release Date	Control Center	Resource Manager
24 May 2016	1.1.4	5.1.3
27 Apr 2016	1.1.3	5.1.2
20 Apr 2016	1.1.3	5.1.1
4 Mar 2016	1.1.2	5.1.1
29 Feb 2016	1.1.1	5.1.1
20 Feb 2016	1.0.10	5.0.10
02 Dec 2015	1.0.9	5.0.9
16 Nov 2015	1.0.8	5.0.8
10 Oct 2015	1.0.7	5.0.7
14 Sep 2015	1.0.6	5.0.6
05 Aug 2015	1.0.5	5.0.5
10 Jul 2015	1.0.4	5.0.4
27 May 2015	1.0.3	5.0.3
20 Apr 2015	1.0.2	5.0.2
03 Apr 2015	1.0.1	5.0.1
24 Feb 2015	1.0.0	5.0.0

Supported upgrade paths

For questions about performing an upgrade or for assistance, please contact Zenoss Support.

Upgrade only Control Center

From	To
Control Center 1.1.1	Control Center 1.1.4
Control Center 1.1.2	Control Center 1.1.4
Control Center 1.1.3	Control Center 1.1.4

Upgrade only Resource Manager

From	To
Resource Manager 5.1.1	Resource Manager 5.1.3
Resource Manager 5.1.2	Resource Manager 5.1.3

Upgrade only the Pacemaker resource agent for Control Center

This upgrade path is valid only for high-availability deployments. The procedures in this part are valid for upgrading from any version of the resource agent to a newer (latest) version.

Upgrade Control Center and Resource Manager

From combination	To combination
Control Center 1.0.6 and Resource Manager 5.0.6	Control Center 1.1.4 and Resource Manager 5.1.3
Control Center 1.0.7 and Resource Manager 5.0.7	Control Center 1.1.4 and Resource Manager 5.1.3
Control Center 1.0.8 and Resource Manager 5.0.8	Control Center 1.1.4 and Resource Manager 5.1.3
Control Center 1.0.9 and Resource Manager 5.0.9	Control Center 1.1.4 and Resource Manager 5.1.3
Control Center 1.0.10 and Resource Manager 5.0.10	Control Center 1.1.4 and Resource Manager 5.1.3

Recent upgrade paths

The following tables identify upgrade paths that are supported but documented only in previous editions of this guide.

Table 1: Upgrade to 1.0.10 / 5.0.10

From combination	To combination
Control Center 1.0.3 and Resource Manager 5.0.3	Control Center 1.0.10 and Resource Manager 5.0.10

From combination	To combination
Control Center 1.0.4 and Resource Manager 5.0.4	Control Center 1.0.10 and Resource Manager 5.0.10
Control Center 1.0.5 and Resource Manager 5.0.5	Control Center 1.0.10 and Resource Manager 5.0.10
Control Center 1.0.6 and Resource Manager 5.0.6	Control Center 1.0.10 and Resource Manager 5.0.10
Control Center 1.0.7 and Resource Manager 5.0.7	Control Center 1.0.10 and Resource Manager 5.0.10
Control Center 1.0.8 and Resource Manager 5.0.8	Control Center 1.0.10 and Resource Manager 5.0.10
Control Center 1.0.9 and Resource Manager 5.0.9	Control Center 1.0.10 and Resource Manager 5.0.10

Table 2: Upgrade to 1.0.3 / 5.0.3

From combination	To combination
Control Center 1.0.0 and Resource Manager 5.0.0	Control Center 1.0.3 and Resource Manager 5.0.3
Control Center 1.0.1 and Resource Manager 5.0.1	Control Center 1.0.3 and Resource Manager 5.0.3
Control Center 1.0.2 and Resource Manager 5.0.2	Control Center 1.0.3 and Resource Manager 5.0.3

Part I: Upgrading only Control Center

The chapters in this part provide instructions for upgrading Control Center without upgrading Resource Manager.

Note Before upgrading only Control Center, make sure that you are upgrading to a supported combination of Control Center and Resource Manager.

The following table identifies the supported upgrades of Control Center alone.

From	To
Control Center 1.1.1	Control Center 1.1.4
Control Center 1.1.2	Control Center 1.1.4
Control Center 1.1.3	Control Center 1.1.4

To perform an upgrade, select one of the following chapters:

Upgrading Control Center with internet access on page 13

Upgrading Control Center without internet access on page 16

Upgrading high-availability deployments with internet access on page 19

Upgrading high-availability deployments without internet access on page 23

Upgrading Control Center with internet access

1

This chapter includes detailed procedures for upgrading Control Center cluster hosts that have internet access. For hosts that do not have internet access, or that are configured for high-availability, use one of the other chapters in this part.

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the `STATUS` column reads `Stopped`:

```
serviced service status Zenoss.resmgr
```

Upgrading the master host

This procedure upgrades the Control Center master host to version 1.1.4.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Verify that accidental upgrades of Docker are disabled.

- a Check the Docker repository.

```
grep enabled /etc/yum.repos.d/docker.repo
```

If the result is `enabled=1`, perform the following substeps.

- b Open `/etc/yum.repos.d/docker.repo` with a text editor.
- c Change the value of the `enabled` key from 1 to 0.
- d Save the file and close the text editor.

- 3 Stop Control Center.

```
systemctl stop serviced
```

- 4 Install the new version of Control Center.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 5 Delete the new Control Center configuration file.

There are no configuration file changes in this release.

```
rm /etc/default/serviced.rpmnew
```

- 6 Start Control Center.

```
systemctl start serviced
```

- **Single-host deployments:** You may log in to the Control Center browser interface, and then restart Resource Manager, if desired.
- **Multi-host deployments:** Update all resource pool hosts (the next procedure).

Upgrading resource pool hosts

This procedure upgrades Control Center resource pool hosts to version 1.1.4.

Perform this procedure on each resource pool host in your deployment.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Verify that accidental upgrades of Docker are disabled.
 - a Check the Docker repository.

```
grep enabled /etc/yum.repos.d/docker.repo
```

If the result is `enabled=1`, perform the following substeps.

- b Open `/etc/yum.repos.d/docker.repo` with a text editor.
- c Change the value of the `enabled` key from 1 to 0.
- d Save the file and close the text editor.

- 3 Stop Control Center and Docker.

```
systemctl stop serviced && systemctl stop docker
```

- 4 Unmount the distributed file system (DFS).

- a Identify the file system specification to unmount.

```
mount | awk '/serviced/ { print $1 }'
```

- b Unmount the DFS.

Replace *DFS-Mount* with the file system specification returned in the previous substep:

```
umount DFS-Mount
```

- 5 Install the new version of Control Center.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the *serviced* configuration file, and installs the new one as */etc/default/serviced.rpmnew*.

- 6 Delete the new Control Center configuration file.

There are no configuration file changes in this release.

```
rm /etc/default/serviced.rpmnew
```

- 7 Start Control Center.

```
systemctl start serviced
```

Upgrading Control Center without internet access

2

This chapter includes detailed procedures for upgrading Control Center cluster hosts that do not have internet access. For hosts that do have internet access, or that are configured for high-availability, use one of the other chapters in this part.

Downloading the RPM package

This procedure describes how to download the `serviced` RPM package to your workstation.

To perform this procedure, you need:

- A workstation with internet access.
 - A portable storage medium, such as a USB flash drive, with at least 20MB of free space.
 - Permission to download the required files from the [Zenoss Enterprise Software Downloads](#) site. You may request permission by filing a ticket at the [Zenoss Support](#) site.
- 1 In a web browser, navigate to the [Zenoss Enterprise Software Downloads](#) site.
 - 2 Click **File Portal - Zenoss Enterprise Software Downloads**.
 - 3 Log in with the account provided by Zenoss Support.
 - 4 Download the `serviced` RPM package.

Replace *Version* with the most recent version number available on the download page:

```
serviced-Version.x86_64.rpm
```

- 5 Copy the file to your portable storage medium.

Staging the RPM package

To perform this procedure, you need the portable storage medium that contains the `serviced` RPM package file.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Copy `serviced-*.x86_64.rpm` from your portable storage medium to `/tmp`.
- 3 Change the access permissions of the package file.

```
chmod 777 /tmp/serviced-*.x86_64.rpm
```


Upgrading the master host

This procedure upgrades the Control Center master host to version 1.1.4.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Stop Control Center.

```
systemctl stop serviced
```

- 3 Install the new version of Control Center.

```
yum install -y /tmp/serviced-*.x86_64.rpm
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 4 Delete the new Control Center configuration file.
There are no configuration file changes in this release.

```
rm /etc/default/serviced.rpmnew
```

- 5 Start Control Center.

```
systemctl start serviced
```

- **Single-host deployments:** You may log in to the Control Center browser interface, and then restart Resource Manager, if desired.
- **Multi-host deployments:** Update all resource pool hosts (the next procedure).

Upgrading resource pool hosts

This procedure upgrades Control Center resource pool hosts to version 1.1.4.

Perform this procedure on each resource pool host in your deployment.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Copy the `serviced` RPM file from the master host.
Replace *Master-Host-IP* with the IP address of the master host:

```
scp root@Master-Host-IP:/tmp/serviced-*.x86_64.rpm /tmp
```

- 3 Stop Control Center and Docker.

```
systemctl stop serviced && systemctl stop docker
```

- 4 Unmount the distributed file system (DFS).
 - a Identify the file system specification to unmount.

```
mount | awk '/serviced/ { print $1 }'
```

- b Unmount the DFS.
Replace *DFS-Mount* with the file system specification returned in the previous substep:

```
umount DFS-Mount
```

- 5 Install the new version of Control Center.

```
yum install -y /tmp/serviced-*.x86_64.rpm
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 6 Delete the new Control Center configuration file.

There are no configuration file changes in this release.

```
rm /etc/default/serviced.rpmnew
```

- 7 Start Control Center.

```
systemctl start serviced
```

Upgrading high-availability deployments with internet access

3

This chapter includes detailed procedures for upgrading Control Center in high-availability deployments that have internet access. For hosts that do not have internet access, or that are not configured for high-availability, use one of the other chapters in this part.

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\./ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the `STATUS` column reads `Stopped`:

```
serviced service status Zenoss.resmgr
```

Stopping Control Center

This procedure stops Control Center on the master host, and on resource pool hosts.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.

- 2 Display the public hostname of the current node.

```
uname -n
```

Make a note of which node (primary or secondary) is the current node, for use in a subsequent procedure.

- 3 Stop Control Center with the cluster management tool.

```
pcs cluster standby --all
```

- 4 Monitor the status of cluster resources.

```
watch pcs status
```

Monitor the status until all resources report Stopped. Resolve any issues before continuing.

- 5 Stop Control Center on all resource pool hosts.

Repeat this step on each resource pool host in your deployment.

- a Log in to the resource pool host as `root`, or as a user with superuser privileges.
- b Stop Control Center.

```
systemctl stop serviced
```

Upgrading the master nodes

This procedure upgrades the Control Center master nodes to version 1.1.4.

Perform this procedure on the primary node and on the secondary node.

- 1 Log in to the master node as `root`, or as a user with superuser privileges.
- 2 Verify that accidental upgrades of Docker are disabled.
 - a Check the Docker repository.

```
grep enabled /etc/yum.repos.d/docker.repo
```

If the result is `enabled=1`, perform the following substeps.

- b Open `/etc/yum.repos.d/docker.repo` with a text editor.
 - c Change the value of the `enabled` key from 1 to 0.
 - d Save the file and close the text editor.
- 3 Install the new version of Control Center.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 4 Delete the new Control Center configuration file.

There are no configuration file changes in this release.

```
rm /etc/default/serviced.rpmnew
```

- 5 Disable `serviced`.

The cluster management software controls `serviced`.

```
systemctl disable serviced
```

Upgrading resource pool hosts

This procedure upgrades Control Center resource pool hosts to version 1.1.4.

Perform this procedure on each resource pool host in your deployment.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Verify that accidental upgrades of Docker are disabled.
 - a Check the Docker repository.

```
grep enabled /etc/yum.repos.d/docker.repo
```

If the result is `enabled=1`, perform the following substeps.

- b Open `/etc/yum.repos.d/docker.repo` with a text editor.
 - c Change the value of the `enabled` key from 1 to 0.
 - d Save the file and close the text editor.
- 3 Stop Control Center and Docker.

```
systemctl stop serviced && systemctl stop docker
```

- 4 Unmount the distributed file system (DFS).
 - a Identify the file system specification to unmount.

```
mount | awk '/serviced/ { print $1 }'
```

- b Unmount the DFS.
Replace *DFS-Mount* with the file system specification returned in the previous substep:

```
umount DFS-Mount
```

- 5 Install the new version of Control Center.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 6 Delete the new Control Center configuration file.
There are no configuration file changes in this release.

```
rm /etc/default/serviced.rpmnew
```

Restarting Control Center

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
For this procedure, the primary node is the node that was the current node when you stopped Control Center.
- 2 Take the cluster out of standby mode.

```
pcs cluster unstandby --all
```

- 3 Monitor the status of cluster resources.

```
watch pcs status
```

Monitor the status until all resources report `Started`. Resolve any issues before continuing.

4 Start Control Center on all resource pool hosts.

Repeat this step on each resource pool host in your deployment.

- a** Log in to the resource pool host as `root`, or as a user with superuser privileges.
- b** Start Docker and Control Center.

```
systemctl start docker && systemctl start serviced
```

Log in to the Control Center browser interface, and then start applications.

Upgrading high-availability deployments without internet access

4

This chapter includes detailed procedures for upgrading Control Center in high-availability deployments that do not have internet access. For hosts that do have internet access, or that are not configured for high-availability, use one of the other chapters in this part.

Downloading the RPM package

This procedure describes how to download the `serviced` RPM package to your workstation.

To perform this procedure, you need:

- A workstation with internet access.
 - A portable storage medium, such as a USB flash drive, with at least 20MB of free space.
 - Permission to download the required files from the [Zenoss Enterprise Software Downloads](#) site. You may request permission by filing a ticket at the [Zenoss Support](#) site.
- 1 In a web browser, navigate to the [Zenoss Enterprise Software Downloads](#) site.
 - 2 Click **File Portal - Zenoss Enterprise Software Downloads**.
 - 3 Log in with the account provided by Zenoss Support.
 - 4 Download the `serviced` RPM package.

Replace *Version* with the most recent version number available on the download page:

```
serviced-Version.x86_64.rpm
```

- 5 Copy the file to your portable storage medium.

Staging the RPM package

To perform this procedure, you need the portable storage medium that contains the `serviced` RPM package file.

Perform this procedure on the primary node and on the secondary node.

- 1 Log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Copy `serviced-*.x86_64.rpm` from your portable storage medium to `/tmp`.
- 3 Change the access permissions of the package file.

```
chmod 777 /tmp/serviced-*.x86_64.rpm
```

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.

Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Stopping Control Center

This procedure stops Control Center on the master host, and on resource pool hosts.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Display the public hostname of the current node.

```
uname -n
```

Make a note of which node (primary or secondary) is the current node, for use in a subsequent procedure.

- 3 Stop Control Center with the cluster management tool.

```
pcs cluster standby --all
```

- 4 Monitor the status of cluster resources.

```
watch pcs status
```

Monitor the status until all resources report Stopped. Resolve any issues before continuing.

- 5 Stop Control Center on all resource pool hosts.

Repeat this step on each resource pool host in your deployment.

- a Log in to the resource pool host as `root`, or as a user with superuser privileges.
 - b Stop Control Center.

```
systemctl stop serviced
```


Upgrading the master nodes

This procedure upgrades the Control Center master nodes to version 1.1.4.

Perform this procedure on the primary node and on the secondary node.

- 1 Log in to the master node as `root`, or as a user with superuser privileges.
- 2 Install the new version of Control Center.

```
yum install -y /tmp/serviced-*.x86_64.rpm
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 3 Delete the new Control Center configuration file.
There are no configuration file changes in this release.

```
rm /etc/default/serviced.rpmnew
```

- 4 Disable `serviced`.
The cluster management software controls `serviced`.

```
systemctl disable serviced
```

- 5 Start Control Center with the cluster management tool.

```
pcs cluster unstandby --all
```

Upgrading resource pool hosts

This procedure upgrades Control Center resource pool hosts to version 1.1.4.

Perform this procedure on each resource pool host in your deployment.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Copy the `serviced` RPM file from a master node.
Replace *Master-Host-IP* with the IP address of a master node:

```
scp root@Master-Host-IP:/tmp/serviced-*.x86_64.rpm /tmp
```

- 3 Stop Control Center and Docker.

```
systemctl stop serviced && systemctl stop docker
```

- 4 Unmount the distributed file system (DFS).
 - a Identify the file system specification to unmount.

```
mount | awk '/serviced/ { print $1 }'
```

- b Unmount the DFS.
Replace *DFS-Mount* with the file system specification returned in the previous substep:

```
umount DFS-Mount
```

- 5 Install the new version of Control Center.

```
yum install -y /tmp/serviced-*.x86_64.rpm
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 6 Delete the new Control Center configuration file.

There are no configuration file changes in this release.

```
rm /etc/default/serviced.rpmnew
```

Restarting Control Center

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
For this procedure, the primary node is the node that was the current node when you stopped Control Center.
- 2 Take the cluster out of standby mode.

```
pcs cluster unstandby --all
```

- 3 Monitor the status of cluster resources.

```
watch pcs status
```

Monitor the status until all resources report `Started`. Resolve any issues before continuing.

- 4 Start Control Center on all resource pool hosts.
Repeat this step on each resource pool host in your deployment.
 - a Log in to the resource pool host as `root`, or as a user with superuser privileges.
 - b Start Docker and Control Center.

```
systemctl start docker && systemctl start serviced
```

Log in to the Control Center browser interface, and then start applications.

Part II: Upgrading only Resource Manager

The chapters in this part provide instructions for upgrading Resource Manager without upgrading Control Center.

Note Before upgrading only Resource Manager, make sure that you are upgrading to a supported combination of Control Center and Resource Manager.

The following table identifies the supported upgrades of Resource Manager alone.

From	To
Resource Manager 5.1.1	Resource Manager 5.1.3
Resource Manager 5.1.2	Resource Manager 5.1.3

To perform an upgrade, select one of the following chapters:

Upgrading Resource Manager with internet access on page 28

Upgrading Resource Manager without internet access on page 30

Upgrading high-availability deployments with internet access on page 33

Upgrading high-availability deployments without internet access on page 35

Upgrading Resource Manager with internet access

1

This chapter includes detailed procedures for upgrading Resource Manager on hosts that have internet access. For hosts that do not have internet access, or that are configured for high-availability, use one of the other chapters in this part.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before performing an upgrade or installing a ZenPack. For more information, see [Using Zenoss Toolbox](#) on page 163.

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the `STATUS` column reads `Stopped`:

```
serviced service status Zenoss.resmgr
```

Upgrading Resource Manager

This procedure upgrades Resource Manager.

Before performing this procedure, stop Resource Manager.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Download the primary Docker image of Resource Manager for this release
The download takes approximately 10-20 minutes.

```
docker run -it --rm -v /root:/mnt/root \
  zenoss/resmgr_5.1:5.1.3_1 rsync -a /root/5.1.x /mnt/root
```

When the download completes, the `rsync` command copies scripts that perform the upgrade to `/root/5.1.x`.

- 3 Pull additional images for Resource Manager from Docker Hub.

```
/root/5.1.x/pull-docker-images.sh
```

- 4 Start the upgrade script.

The script to start depends on whether Service Impact is installed.

- If Service Impact is installed, enter the following command.

```
/root/5.1.x/upgrade-impact-5.1.x.sh
```

Note The script upgrades Resource Manager, but does not upgrade Service Impact. For more information about upgrading Service Impact, refer to the *Zenoss Service Impact Installation Guide for Resource Manager 5.x*.

- If Service Impact is not installed, enter the following command.

```
/root/5.1.x/upgrade-resmgr-5.1.x.sh
```

- 5 Restart Resource Manager.

Some Resource Manager services are started during the upgrade, and they need to be restarted.

```
serviced service restart Zenoss.resmgr
```

Proceed to [After upgrading](#) on page 38.

Upgrading Resource Manager without internet access

2

This chapter includes detailed procedures for upgrading Resource Manager on hosts that do not have internet access. For hosts that do have internet access, or that are configured for high-availability, use one of the other chapters in this part.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before performing an upgrade or installing a ZenPack. For more information, see [Using Zenoss Toolbox](#) on page 163.

Downloading files for offline installation

This procedure describes how to download Docker image files to your workstation.

To perform this procedure, you need:

- A workstation with internet access.
 - A portable storage medium, such as a USB flash drive, with at least 5 GB of free space.
 - Permission to download the required files from the [File Portal - Download Zenoss Enterprise Software](#) site. You may request permission by filing a ticket at the [Zenoss Support](#) site.
- 1 In a web browser, navigate to the [File Portal - Download Zenoss Enterprise Software](#) site.
 - 2 Log in with the account provided by Zenoss Support.
 - 3 Download archive files to your workstation.

Replace *Version* with the most recent version number available on the download page:

- `install-zenoss-hbase:vVersion.run`
 - `install-zenoss-opentsdb:vVersion.run`
 - `install-zenoss-resmgr_5.1:5.1Version.run`
- 4 Copy the files to your portable storage medium.

Staging archive files

To perform this procedure, you need the portable storage medium that contains the archive files.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Copy `*.run` from your portable storage medium to `/root`.

- 3 Add execute permission to the archive files.

```
chmod +x /root/*.run
```

Upgrading Resource Manager

This procedure upgrades Resource Manager.

Before performing this procedure, stop Resource Manager.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Install the new Resource Manager images.
 - a Change directory to `/root`.

```
cd /root
```

- b Install the HBase image.

```
./install-zenoss-hbase*.run
```

- c Install the OpenTSDB image.

```
./install-zenoss-opentsdb*.run
```

- d Install the Resource Manager image.

```
./install-zenoss-resmgr*.run
```

- e Optional: Delete the self-extracting image files, if desired.

```
rm ./install-zenoss-*.run
```

- 3 Extract files from the Resource Manager image.

```
docker run -it --rm -v /root:/mnt/root \
  zenoss/resmgr_5.1:5.1.3_1 rsync -a /root/5.1.x /mnt/root
```

The preceding commands copy upgrade scripts to `/root/5.1.x`.

- 4 Start the upgrade script.

The script to start depends on whether Service Impact is installed.

- If Service Impact is installed, enter the following command.

```
/root/5.1.x/upgrade-impact-5.1.x.sh
```

Note The script upgrades Resource Manager, but does not upgrade Service Impact. For more information about upgrading Service Impact, refer to the *Zenoss Service Impact Installation Guide for Resource Manager 5.x*.

- If Service Impact is not installed, enter the following command.

```
/root/5.1.x/upgrade-resmgr-5.1.x.sh
```

- 5 Restart Resource Manager.

Some Resource Manager services are started during the upgrade, and they need to be restarted.

```
serviced service restart Zenoss.resmgr
```

Proceed to [After upgrading](#) on page 38.

Upgrading high-availability deployments with internet access

3

This chapter includes detailed procedures for upgrading Resource Manager in high-availability deployments that have internet access. For deployments that do not have internet access, or that are not configured for high-availability, use one of the other chapters in this part.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before performing an upgrade or installing a ZenPack. For more information, see [Using Zenoss Toolbox](#) on page 163.

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the `STATUS` column reads `Stopped`:

```
serviced service status Zenoss.resmgr
```

Upgrading Resource Manager

This procedure upgrades Resource Manager.

Before performing this procedure, stop Resource Manager. For more information, see [Stopping applications](#) on page 19.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Download the primary Docker image of Resource Manager for this release

The download takes approximately 10-20 minutes.

```
docker run -it --rm -v /root:/mnt/root \
  zenoss/resmgr_5.1:5.1.3_1 rsync -a /root/5.1.x /mnt/root
```

When the download completes, the `rsync` command copies scripts that perform the upgrade to `/root/5.1.x`.

- 3 Pull additional images for Resource Manager from Docker Hub.

```
/root/5.1.x/pull-docker-images.sh
```

- 4 Start the upgrade script.

The script to start depends on whether Service Impact is installed.

- If Service Impact is installed, enter the following command.

```
/root/5.1.x/upgrade-impact-5.1.x.sh
```

Note The script upgrades Resource Manager, but does not upgrade Service Impact. For more information about upgrading Service Impact, refer to the *Zenoss Service Impact Installation Guide for Resource Manager 5.x*.

- If Service Impact is not installed, enter the following command.

```
/root/5.1.x/upgrade-resmgr-5.1.x.sh
```

- 5 Restart Resource Manager.

Some Resource Manager services are started during the upgrade, and they need to be restarted.

```
serviced service restart Zenoss.resmgr
```

Proceed to [After upgrading](#) on page 38.

Upgrading high-availability deployments without internet access

4

This chapter includes detailed procedures for upgrading Resource Manager in high-availability deployments that do not have internet access. For deployments that do have internet access, or that are not configured for high-availability, use one of the other chapters in this part.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before performing an upgrade or installing a ZenPack. For more information, see [Using Zenoss Toolbox](#) on page 163.

Downloading files for offline installation

This procedure describes how to download Docker image files to your workstation.

To perform this procedure, you need:

- A workstation with internet access.
 - A portable storage medium, such as a USB flash drive, with at least 5 GB of free space.
 - Permission to download the required files from the [File Portal - Download Zenoss Enterprise Software](#) site. You may request permission by filing a ticket at the [Zenoss Support](#) site.
- 1 In a web browser, navigate to the [File Portal - Download Zenoss Enterprise Software](#) site.
 - 2 Log in with the account provided by Zenoss Support.
 - 3 Download archive files to your workstation.

Replace *Version* with the most recent version number available on the download page:

- `install-zenoss-hbase:vVersion.run`
 - `install-zenoss-opentsdb:vVersion.run`
 - `install-zenoss-resmgr_5.1:5.1Version.run`
- 4 Copy the files to your portable storage medium.

Staging archive files

To perform this procedure, you need the portable storage medium that contains the archive files.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Copy `*.run` from your portable storage medium to `/root`.

- 3 Add execute permission to the archive files.

```
chmod +x /root/*.run
```

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Upgrading Resource Manager

This procedure upgrades Resource Manager.

Before performing this procedure, stop Resource Manager.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Install the new Resource Manager images.
 - a Change directory to `/root`.

```
cd /root
```

- b Install the HBase image.

```
./install-zenoss-hbase*.run
```

- c Install the OpenTSDB image.

```
./install-zenoss-opentsdb*.run
```

- d Install the Resource Manager image.

```
./install-zenoss-resmgr*.run
```

- e Optional: Delete the self-extracting image files, if desired.

```
rm ./install-zenoss-*.run
```

- 3 Extract files from the Resource Manager image.

```
docker run -it --rm -v /root:/mnt/root \
  zenoss/resmgr_5.1:5.1.3_1 rsync -a /root/5.1.x /mnt/root
```

The preceding commands copy upgrade scripts to `/root/5.1.x`.

- 4 Start the upgrade script.

The script to start depends on whether Service Impact is installed.

- If Service Impact is installed, enter the following command.

```
/root/5.1.x/upgrade-impact-5.1.x.sh
```

Note The script upgrades Resource Manager, but does not upgrade Service Impact. For more information about upgrading Service Impact, refer to the *Zenoss Service Impact Installation Guide for Resource Manager 5.x*.

- If Service Impact is not installed, enter the following command.

```
/root/5.1.x/upgrade-resmgr-5.1.x.sh
```

- 5 Restart Resource Manager.

Some Resource Manager services are started during the upgrade, and they need to be restarted.

```
serviced service restart Zenoss.resmgr
```

Proceed to [After upgrading](#) on page 38.

After upgrading

This chapter includes information about what to do after upgrading Resource Manager.

Deleting the pre-upgrade snapshot

Before the Resource Manager upgrade begins, the upgrade script creates and tags a snapshot of the system. Tagged snapshots persist until they are explicitly removed. When you are satisfied the new release is working properly, delete the pre-upgrade snapshot.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Display a list of all Control Center snapshots, with their tags.

```
serviced snapshot list -t
```

Example result:

Snapshot	Description	Tags
xm5mtezbyo2_20160211-220535.480		preupgrade-resmgr-5.1.2

- 3 Delete the pre-upgrade snapshot.

Replace *Snapshot-ID* with the identifier of the pre-upgrade snapshot returned in the previous step:

```
serviced snapshot remove Snapshot-ID
```

Part III: Upgrading only the Pacemaker resource agent for Control Center

The chapters in this part provide instructions for upgrading only the Pacemaker resource agent for Control Center. This upgrade path is valid only for high-availability deployments.

The procedures in this part are valid for upgrading from any version of the resource agent to a newer (latest) version. To perform an upgrade, select one of the following chapters:

Upgrading high-availability deployments with internet access on page 40

Upgrading high-availability deployments without internet access on page 42

Upgrading high-availability deployments with internet access

1

This chapter includes detailed procedures for upgrading the Pacemaker resource agent for Control Center in high-availability deployments that have internet access. For deployments that do not have internet access, use the other chapter in this part.

Identifying the Pacemaker resource agent version

This procedure identifies the installed version of the Pacemaker resource agent for Control Center. Perform this procedure to determine whether to upgrade the resource agent.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 3 On both nodes, identify the installed version of the Pacemaker resource agent for Control Center.

```
rpm -qa | grep serviced-resource-agents
```

The result includes the version number; in this example, 0.0.5:

```
serviced-resource-agents-0.0.5-1.x86_64
```

- 4 On either node, identify the currently-available version.

```
yum list all --enablerepo=zenoss-stable serviced-resource-agents
```

The result includes the version number; in this example, 0.0.6:

```
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
* base: centos.mirror.nac.net
* epel: mirror.es.its.nyu.edu
* extras: mirror.cisp.com
* updates: mirrors.advancedhosters.com
Available Packages
serviced-resource-agents.x86_64          0.0.6-1          zenoss-stable
```

If the number of the currently-available version is greater than the number of the installed version, proceed to the next topic.

Note Version 0.0.6 includes an important fix, and is recommended for all deployments.

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\./ { print $1 }'
```

- b Identify the applications that are running.

Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Upgrading the Pacemaker resource agent

This procedure upgrades the Pacemaker resource agent for Control Center.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Stop Control Center.

```
pcs resource disable serviced
```

- 3 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 4 On both nodes, install the latest version of the resource agent for Control Center.

```
yum install -y --enablerepo=zenoss-stable serviced-resource-agents
```

- 5 On either node, start Control Center.

```
pcs resource enable serviced
```

- 6 Start Resource Manager.

```
serviced service start Zenoss.resmgr
```

Upgrading high-availability deployments without internet access

2

This chapter includes detailed procedures for upgrading the Pacemaker resource agent for Control Center in high-availability deployments that do not have internet access. For deployments that do have internet access, use the other chapter in this part.

Identifying the Pacemaker resource agent version

This procedure identifies the installed version of the Pacemaker resource agent for Control Center.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 3 On both nodes, identify the installed version of the Pacemaker resource agent for Control Center.

```
rpm -qa | grep serviced-resource-agents
```

The result includes the version number; in this example, `0.0.5`:

```
serviced-resource-agents-0.0.5-1.x86_64
```

Perform the next procedure to determine what is the currently-available version of the resource agents package.

Identifying the currently-available version of the Pacemaker resource agent

This procedure describes how to identify the currently-available version of the Pacemaker resource agent for Control Center, and how to download its RPM package, if necessary.

To perform this procedure, you need:

- A workstation with internet access.
- A portable storage medium, such as a USB flash drive, with at least 1 MB of free space.
- Permission to download the required files from the [File Portal - Download Zenoss Enterprise Software](#) site. You may request permission by filing a ticket at the [Zenoss Support](#) site.

- 1 In a web browser, navigate to the [File Portal - Download Zenoss Enterprise Software](#) site.
- 2 Log in with the account provided by Zenoss Support.
- 3 Identify the version number of the resource agents RPM file.

The file name is `serviced-resource-agents-Version.x86_64.rpm`.

- If the version number of the RPM package file is greater than the version number of the packages installed on your master nodes, download the package file.
 - If the version number of the RPM package file is equal to the version number of the packages installed on your master nodes, do not download the package file.
- 4 Copy the RPM package file to your portable storage medium.

Staging the RPM package file

To perform this procedure, you need the portable storage medium that contains the RPM package file.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 3 On both nodes, copy RPM package file from your portable storage medium to `/root`.

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\./ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Upgrading the Pacemaker resource agent

This procedure upgrades the Pacemaker resource agent for Control Center.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Stop Control Center.

```
pcs resource disable serviced
```

- 3 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.

- 4 On both nodes, install the latest version of the resource agent for Control Center.

```
yum install -y -/root/serviced-resource-agents-*.rpm
```

- 5 On either node, start Control Center.

```
pcs resource enable serviced
```

- 6 Start Resource Manager.

```
serviced service start Zenoss.resmgr
```

Part IV: Upgrading Control Center and Resource Manager

The chapters in this part provide instructions for upgrading the combinations of Control Center and Resource Manager shown in the following table.

From combination	To combination
Control Center 1.0.6 and Resource Manager 5.0.6	Control Center 1.1.4 and Resource Manager 5.1.3
Control Center 1.0.7 and Resource Manager 5.0.7	Control Center 1.1.4 and Resource Manager 5.1.3
Control Center 1.0.8 and Resource Manager 5.0.8	Control Center 1.1.4 and Resource Manager 5.1.3
Control Center 1.0.9 and Resource Manager 5.0.9	Control Center 1.1.4 and Resource Manager 5.1.3
Control Center 1.0.10 and Resource Manager 5.0.10	Control Center 1.1.4 and Resource Manager 5.1.3

For information about upgrading other combinations, see [Supported upgrade paths](#) on page 9.

To perform an upgrade, first review the information in [Preparing to upgrade](#) on page 46, and then select one of the following chapters:

[Upgrading Resource Manager with internet access](#) on page 50

[Upgrading Resource Manager without internet access](#) on page 74

[Upgrading high-availability deployments with internet access](#) on page 99

[Upgrading high-availability deployments without internet access](#) on page 128

When you have completed the upgrade, proceed to [After upgrading](#) on page 158.

Preparing to upgrade

This chapter includes information about upgrading your deployment of Control Center and Resource Manager, and procedures that prepare your deployment for the upgrade. The information and procedures in this chapter are independent of the specific upgrade procedures, which are detailed in subsequent chapters.

Storage changes

This release of Control Center includes a new storage driver for application data, named `devicemapper`. The new driver is based on *the Docker devicemapper storage driver*, which in turn is based on the *device mapper framework of the Linux kernel*.

The key feature of the drivers is their use of thin provisioning, a virtualization method that allocates data blocks only when data is written. (The traditional method is to allocate data blocks when a file system is created, before any data is written.) Thin provisioning enables snapshots, a time-efficient and space-efficient method of copying data, and enables making a device appear to have more physical data blocks than are actually available (as long as some blocks are unfilled). Also, thin-provisioned storage can be extended without having to move data from one physical partition to another.

For Docker and for Control Center, Logical Volume Manager (LVM) tools are used to create thin-provisioned storage, in thin pools. A thin pool includes an area for metadata (a small percentage of the total) and an area for the data itself. To simplify creating thin pools, this release of Control Center includes `serviced-storage`, a utility that calls the LVM tools. The utility may be used to create thin pools for use by Docker's `devicemapper` storage driver as well as for the Control Center `devicemapper` storage driver.

For Docker data storage, the recommended storage layout is a device mapper thin pool on one or more primary partitions. Docker's `devicemapper` storage driver may be used in `loop-lvm` mode, on loopback-mounted sparse files. However, `loop-lvm` mode is not recommended for production use, and Zenoss strongly recommends using a device mapper thin pool for Docker storage, in all deployment scenarios.

With this release, the recommended (and default) storage driver for Control Center application data is `devicemapper`. The recommended layout for application data storage is a device mapper thin pool on one or more primary partitions. Support for the `btrfs` storage driver is deprecated. The change is reflected in the Control Center configuration file, `/etc/default/serviced`: The default value of the `SERVICED_FS_TYPE` variable is `devicemapper`.

To convert application data from `Btrfs` file system storage to device mapper thin pool storage, the Control Center master host requires additional block storage. For more information about storage requirements in general, refer to the *Zenoss Resource Manager Planning Guide*. To determine how much data is stored in a `Btrfs` volume, see the next topic.

Note Control Center does not support restoring application data backups made with one storage driver to another driver. For example, a backup of Resource Manager data that was stored on a Btrfs partition can not be restored to a `devicemapper` or `rsync` partition. In addition, snapshots and backups made with Control Center 1.0.x are not compatible with Control Center 1.1.x. Likewise, snapshots and backups made with Control Center 1.1.x are not compatible with Control Center 1.0.x.

Control Center stores metadata for its `devicemapper` storage driver in `/opt/serviced/var/volumes`. The amount of storage required for the metadata rarely exceeds 1GB, so the directory no longer requires a separate, non-root file system (except for high-availability deployments, to mirror the metadata).

Preparing to convert the data storage driver

Perform this procedure to determine how much new storage is required to migrate application data from a Btrfs file system to a device mapper thin pool.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Determine the amount of data stored in your Btrfs volume.

The following command may not work for Btrfs volumes backed by RAID systems. For more information about determining the amount of used space on Btrfs volumes, refer to the [Btrfs FAQ](#).

```
btrfs filesystem df /opt/serviced/var/volumes
```

Example result:

```
Data, single: total=29.01GiB, used=23.8GiB
System, single: total=4.00MiB, used=16.00KiB
Metadata, single: total=264.00MiB, used=67.8MiB
GlobalReserve, single: total=16.00MiB, used=0.00
```

In this example, the used space for data and metadata total approximately 23.8G.

Use the size information to add storage to the Control Center master host. For more information, refer to the *Zenoss Resource Manager Planning Guide*.


Retaining a customized Control Center web server port

If you are using a port other than 443 for the Control Center web server, perform this procedure to retain the customization.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the port number of the Control Center web server.

```
egrep '^[^#]*SERVICED_UI_PORT' /etc/default/serviced
```

- If the command returns no result, the web server port is the default port, 443. Do not perform the remaining steps in this procedure.
 - If the command returns a result, the web server port is the value to which the variable is set. Perform the remaining steps in this procedure.
- 3 Log in to the Control Center browser interface.
 - 4 In the **Applications** table, click **Zenoss.resmgr**.
 - 5 In the application title line, click **Edit Variables**.

Figure 1: The Edit Variables dialog

- 6 In the **Edit Variables** dialog, add the `controlplane-port` key.

Replace *Port-Number* with the port number of the `SERVICED_UI_PORT` variable in the Control Center configuration file, displayed previously.

```
controlplane-port Port-Number
```

- 7 In the **Edit Variables** dialog, click **Save Changes**.

Ensuring localhost resolution

Control Center requires an entry for localhost in `/etc/hosts`.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Determine whether `127.0.0.1` is mapped to localhost.

```
grep 127.0.0.1 /etc/hosts | grep localhost
```

If the preceding commands return no result, perform the following step.

- 3 Add an entry to `/etc/hosts` for localhost.

```
echo "127.0.0.1 localhost" >> /etc/hosts
```

Important information and recommendations

With this release, Control Center supports adding tags to snapshots. Tagged snapshots are not affected by the snapshot time-to-live setting—instead, tagged snapshots persist until explicitly removed. Before the Resource Manager upgrade begins, the upgrade script creates and tags a snapshot. You can roll back to the pre-upgrade snapshot at any time.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before performing an upgrade or installing a ZenPack. For more information, see [Using Zenoss Toolbox](#) on page 163.

This release requires RHEL/CentOS 7.1 or 7.2. Zenoss strongly recommends upgrading the operating system only at the step specified in each upgrade procedure. For hosts that do not have internet access, you must create and use a local mirror, or other media, to upgrade the operating system.

The procedures in this guide include instructions for the following configuration options, which may or may not be part of your deployment:

- Isolate the Control Center master host in its own resource pool.
- Create a ZooKeeper ensemble.

Both options improve reliability and require a multi-host deployment, and are strongly recommended. For more information, refer to the *Zenoss Resource Manager Planning Guide*.

For optimum results, Zenoss recommends reviewing the upgrade procedures for your deployment before performing the upgrade.

Upgrading Resource Manager with internet access

2

This chapter includes detailed procedures for upgrading Control Center cluster hosts that have internet access. For hosts that do not have internet access, or that are configured for high-availability, use one of the other chapters in this part.

The upgrade process includes the following general steps:

- 1 Stop the application or applications that Control Center is managing, and then stop Control Center on the master host.
- 2 Upgrade Control Center on the master host.
- 3 Upgrade resource pool hosts.

Upgrading Control Center on the master host

This section describes how to upgrade Control Center on the master host in a Control Center cluster.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before proceeding. For more information, see [Using Zenoss Toolbox](#) on page 163.

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\./ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the `STATUS` column reads `Stopped`:

```
serviced service status Zenoss.resmgr
```

Saving copies of Resource Manager images

This procedure ensures that Resource Manager images are present and up-to-date in the Control Center registry.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Determine whether Control Center is running.

```
systemctl status serviced | grep Active
```

If the result does not include `active (running)`, enter the following command:

```
systemctl start serviced
```

- 3 Determine whether the `SERVICED_REGISTRY` variable is set.

```
egrep '^[^#]*SERVICED_REGISTRY' /etc/default/serviced
```

- If the result is `SERVICED_REGISTRY=1`, stop performing this procedure, and continue to the next.
 - If the result is `SERVICED_REGISTRY=0`, or if the command returns no result, complete the remaining steps of this procedure before continuing to the next.
- 4 Edit the Control Center configuration file.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the `SERVICED_REGISTRY` declaration, and then set its value to 1.
 - c Remove the number sign character (`#`) from the beginning of the line.
 - d Save the file, and then close the editor.
 - e Restart Control Center.

```
systemctl stop serviced && systemctl start serviced
```

- 5 Synchronize the Docker and Control Center registries.

```
serviced docker sync
```

The synchronization may take approximately 20-30 minutes.

Upgrading Docker

This procedure upgrades Docker from version 1.5 or 1.8.2 to 1.9.0.

Note This release of Control Center requires RHEL/CentOS 7.1 or 7.2, and this procedure includes steps for upgrading the operating system, if necessary or desired. Zenoss recommends upgrading the operating system only at the step specified in this procedure.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Stop Control Center.

```
systemctl stop serviced
```

- 3 Determine which version of Docker is installed.

```
rpm -qa | grep docker
```

- 4 **Note** Perform this step **only** if `zenoss-docker-1.5` is installed.

Remove `zenoss-docker`, and enable installation of the new version of Docker.

- a Remove Docker, without removing Control Center.

```
rpm -e --nodeps zenoss-docker-1.5.0-2
```

- b Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- 5 **Note** Perform this step **only** if `docker-engine-1.8` is installed.

Remove Docker, and then verify that the Docker repository is enabled.

- a Remove Docker, without removing Control Center.

```
rpm -e --nodeps docker-engine-1.8.2
```

- b Check the Docker repository.

```
grep enabled /etc/yum.repos.d/docker.repo
```

If the result is `enabled=0`, perform the following substeps.

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.
 - d Change the value of the `enabled` key from 0 to 1.
 - e Save the file and close the text editor.
- 6 Remove the Docker data partition.

- a Identify the partition where `/var/lib/docker` is mounted.

```
mount | awk '/\var\/lib\/docker/ { print $1 }'
```

- b Unmount the partition.

Replace *Partition* with the device returned in the previous substep:

```
umount Partition
```

- c Erase the XFS file system on the partition.

Replace *Partition* with the device returned previously:

```
wipefs -a Partition
```

The partition is now ready for use as a device mapper thin pool.

- d Open `/etc/fstab` with a text editor.
- e Remove the entry for `/var/lib/docker`.

7 Upgrade the operating system, if necessary.

- a Determine which release is installed.

```
cat /etc/redhat-release
```

If the result includes 7.0, perform the following substeps.

- b Disable the serviced service.

```
systemctl disable serviced
```

- c Upgrade the operating system.

```
yum clean all && yum update -y
```

- d Restart the operating system.

```
reboot
```

- e Log in to the Control Center cluster host as `root`, or as a user with superuser privileges.

- f Enable the `serviced` service.

```
systemctl enable serviced
```

8 Install Docker 1.9.0, and then disable accidental upgrades.

- a Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- b Install Docker 1.9.0.

```
yum clean all && yum makecache fast
yum install -y docker-engine-1.9.0
```

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.

- d Change the value of the `enabled` key from 1 to 0.

- e Save the file and close the text editor.

9 Create a symbolic link for the Docker temporary directory.

Docker uses its temporary directory to spool images. The default directory is `/var/lib/docker/tmp`. The following command specifies the same directory that Control Center uses, `/tmp`. You can specify any directory that has a minimum of 10GB of unused space.

- a Create the `docker` directory in `/var/lib`.

```
mkdir /var/lib/docker
```

- b Create the link to `/tmp`.

```
ln -s /tmp /var/lib/docker/tmp
```

10 Create a `systemd` override file for the Docker service definition.

- a Create the override directory.

```
mkdir -p /etc/systemd/system/docker.service.d
```

- b Create the override file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/docker daemon \${OPTIONS} -H fd://
EOF
```

- c Reload the systemd manager configuration.

```
systemctl daemon-reload
```

- 11 Configure the docker service to start when the system starts.

```
systemctl enable docker
```

- 12 Install Control Center.

Control Center includes a utility that simplifies the process of creating a device mapper thin pool.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the serviced configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 13 Create a device mapper thin pool for Docker data.

- a Identify the primary partition for the thin pool to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b Create the thin pool.

Replace *Path-To-Device* with the path of an unused primary partition:

```
serviced-storage create-thin-pool docker Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 14 Configure and start the Docker service.

- a Create variables for adding arguments to the Docker configuration file.

The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

Replace *Thin-Pool-Device* with the name of the thin pool device created in the previous step:

```
myDriver="-s devicemapper"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myFlag="--storage-opt dm.thinpooldev"
myPool="Thin-Pool-Device"
```

- b Add the arguments to the Docker configuration file.

```
echo 'OPTIONS="'$myDriver $myFix $myFlag'='$myPool'' \
>> /etc/sysconfig/docker
```

- c Start or restart Docker.

```
systemctl restart docker
```

The initial startup takes up to a minute, and may fail. If the startup fails, repeat the previous command.

- 15 Configure name resolution in containers.

Each time it starts, `docker` selects an IPv4 subnet for its virtual Ethernet bridge. The selection can change; this step ensures consistency.

- a Identify the IPv4 subnet and netmask `docker` has selected for its virtual Ethernet bridge.

```
ip addr show docker0 | grep inet
```

- b Open `/etc/sysconfig/docker` in a text editor.
- c Add the following flags to the end of the `OPTIONS` declaration.

Replace *Bridge-Subnet* with the IPv4 subnet `docker` selected for its virtual bridge, and replace *Bridge-Netmask* with the netmask `docker` selected:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/Bridge-Netmask
```

For example, if the bridge subnet and netmask is 172.17.0.1/16, the flags to add are `--dns=172.17.0.1 --bip=172.17.0.1/16`.

Note Leave a blank space after the end of the thin pool device name, and make sure the double quote character (") is at the end of the line.

- d Restart the Docker service.

```
systemctl restart docker
```

Configuring Control Center on the master host

This procedure upgrades Control Center on the master host.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Manage the Control Center configuration files.
 - a Change directory to the configuration file directory.

```
cd /etc/default
```

- b Rename the existing configuration file, as a backup.

```
mv serviced serviced.pre-1.1.4
```

- c Rename the new configuration file, and then create a copy of it.

```
mv serviced.rpmnew serviced.orig-1.1.4
cp serviced.orig-1.1.4 serviced
```

- 3 Copy settings from the previous Control Center configuration file to the new one.
 - a Identify the customized variables in the pre-upgrade configuration file.

```
egrep '^[^#]*SERVICED' serviced.pre-1.1.4
```

- b Open `/etc/default/serviced` with a text editor, and then customize the same variables that were customized in the pre-upgrade configuration file.

The following variables are deprecated, and are not needed in the new configuration file:

- `SERVICED_REGISTRY`
 - `SERVICED_VARPATH`
- c Add `SERVICED_DOCKER_REGISTRY` to the file.
- The variable specifies the host and port at which the local Docker registry is available.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
SERVICED_DOCKER_REGISTRY=Hostname-Or-IP:5000
```

- d Save the file and close the text editor.
- 4 Start Control Center and watch its log file.

```
systemctl start serviced && journalctl -flu serviced
```

Control Center pulls images from Docker Hub to complete its update, and then copies Resource Manager images from its registry to the Docker library.

- **Single-host deployments:** You may log in to the Control Center browser interface, and then restart Resource Manager, if desired.
- **Multi-host deployments:** Update all resource pool hosts (the next procedure).

Upgrading Control Center on resource pool hosts

This section describes how to upgrade Control Center on resource pool hosts in a Control Center cluster.

Upgrading Docker

This procedure upgrades Docker from version 1.5 or 1.8.2 to 1.9.0.

Note This release of Control Center requires RHEL/CentOS 7.1 or 7.2, and this procedure includes steps for upgrading the operating system, if necessary or desired. Zenoss recommends upgrading the operating system only at the step specified in this procedure.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Stop Control Center.

```
systemctl stop serviced
```

- 3 Determine which version of Docker is installed.

```
rpm -qa | grep docker
```

- 4 **Note** Perform this step **only** if `zenoss-docker-1.5` is installed.

Remove `zenoss-docker`, and enable installation of the new version of Docker.

- a Remove Docker, without removing Control Center.

```
rpm -e --nodeps zenoss-docker-1.5.0-2
```


- b Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- 5 **Note** Perform this step **only** if `docker-engine-1.8` is installed.

Remove Docker, and then verify that the Docker repository is enabled.

- a Remove Docker, without removing Control Center.

```
rpm -e --nodeps docker-engine-1.8.2
```

- b Check the Docker repository.

```
grep enabled /etc/yum.repos.d/docker.repo
```

If the result is `enabled=0`, perform the following substeps.

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.
 - d Change the value of the `enabled` key from 0 to 1.
 - e Save the file and close the text editor.
- 6 Remove the Docker data partition.

- a Identify the partition where `/var/lib/docker` is mounted.

```
mount | awk '/\var/lib/docker/ { print $1 }'
```

- b Unmount the partition.

Replace *Partition* with the device returned in the previous substep:

```
umount Partition
```

- c Erase the XFS file system on the partition.

Replace *Partition* with the device returned previously:

```
wipefs -a Partition
```

The partition is now ready for use as a device mapper thin pool.

- d Open `/etc/fstab` with a text editor.
 - e Remove the entry for `/var/lib/docker`.
- 7 Upgrade the operating system, if necessary.
- a Determine which release is installed.

```
cat /etc/redhat-release
```

If the result includes 7.0, perform the following substeps.

- b Disable the `serviced` service.

```
systemctl disable serviced
```

- c Upgrade the operating system.

```
yum clean all && yum update -y
```

- d Restart the operating system.

```
reboot
```

- e Log in to the Control Center cluster host as `root`, or as a user with superuser privileges.
- f Enable the `serviced` service.

```
systemctl enable serviced
```

- 8 Install Docker 1.9.0, and then disable accidental upgrades.

- a Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- b Install Docker 1.9.0.

```
yum clean all && yum makecache fast
yum install -y docker-engine-1.9.0
```

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.
 - d Change the value of the `enabled` key from 1 to 0.
 - e Save the file and close the text editor.
- 9 Create a symbolic link for the Docker temporary directory.

Docker uses its temporary directory to spool images. The default directory is `/var/lib/docker/tmp`. The following command specifies the same directory that Control Center uses, `/tmp`. You can specify any directory that has a minimum of 10GB of unused space.

- a Create the `docker` directory in `/var/lib`.

```
mkdir /var/lib/docker
```

- b Create the link to `/tmp`.

```
ln -s /tmp /var/lib/docker/tmp
```

- 10 Create a `systemd` override file for the Docker service definition.

- a Create the override directory.

```
mkdir -p /etc/systemd/system/docker.service.d
```

- b Create the override file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
```

```
ExecStart=
ExecStart=/usr/bin/docker daemon \${OPTIONS} -H fd://
EOF
```

- c Reload the systemd manager configuration.

```
systemctl daemon-reload
```

- 11 Configure the docker service to start when the system starts.

```
systemctl enable docker
```

- 12 Install Control Center.

Control Center includes a utility that simplifies the process of creating a device mapper thin pool.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the serviced configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 13 Create a device mapper thin pool for Docker data.

- a Identify the primary partition for the thin pool to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b Create the thin pool.

Replace *Path-To-Device* with the path of an unused primary partition:

```
serviced-storage create-thin-pool docker Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 14 Configure and start the Docker service.

- a Create variables for adding arguments to the Docker configuration file.

The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

Replace *Thin-Pool-Device* with the name of the thin pool device created in the previous step:

```
myDriver="--s devicemapper"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myFlag="--storage-opt dm.thinpooldev"
myPool="Thin-Pool-Device"
```

- b Add the arguments to the Docker configuration file.

```
echo 'OPTIONS="'$myDriver $myFix $myFlag'='$myPool''' \
>> /etc/sysconfig/docker
```

- c Start or restart Docker.

```
systemctl restart docker
```

The initial startup takes up to a minute, and may fail. If the startup fails, repeat the previous command.

- 15 Configure name resolution in containers.

Each time it starts, `docker` selects an IPv4 subnet for its virtual Ethernet bridge. The selection can change; this step ensures consistency.

- a Identify the IPv4 subnet and netmask docker has selected for its virtual Ethernet bridge.

```
ip addr show docker0 | grep inet
```

- b Open `/etc/sysconfig/docker` in a text editor.
- c Add the following flags to the end of the `OPTIONS` declaration.

Replace `Bridge-Subnet` with the IPv4 subnet docker selected for its virtual bridge, and replace `Bridge-Netmask` with the netmask docker selected:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/Bridge-Netmask
```

For example, if the bridge subnet and netmask is `172.17.0.1/16`, the flags to add are `--dns=172.17.0.1`
`--bip=172.17.0.1/16`.

Note Leave a blank space after the end of the thin pool device name, and make sure the double quote character (") is at the end of the line.

- d Restart the Docker service.

```
systemctl restart docker
```

Configuring Control Center on resource pool hosts

Perform this procedure on each resource pool host in a Control Center cluster.

- 1 Manage the Control Center configuration files.
 - a Change directory to the configuration file directory.

```
cd /etc/default
```

- b Rename the existing configuration file, as a backup.

```
mv serviced serviced.pre-1.1.4
```

- c Rename the new configuration file, and then create a copy of it.

```
mv serviced.rpmnew serviced.orig-1.1.4
cp serviced.orig-1.1.4 serviced
```

- 2 Copy settings from the previous Control Center configuration file to the new one.
 - a Identify the customized variables in the pre-upgrade configuration file.

```
egrep '^[^#]*SERVICED' serviced.pre-1.1.4
```

- b Open `/etc/default/serviced` with a text editor, and then customize the same variables that were customized in the pre-upgrade configuration file.

The following variables are deprecated, and are not needed in the new configuration file:

- `SERVICED_REGISTRY`
 - `SERVICED_VARPATH`
- c Add `SERVICED_DOCKER_REGISTRY` to the file.

The variable specifies the host and port at which the local Docker registry is available.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
SERVICED_DOCKER_REGISTRY=Hostname-Or-IP:5000
```

- d Save the file and close the text editor.
- 3 Unmount the distributed file system (DFS).
 - a Identify the file system specification to unmount.

```
mount | awk '/serviced/ { print $1 }'
```

- b Unmount the DFS.
Replace *DFS-Mount* with the file system specification returned in the previous substep:

```
umount DFS-Mount
```

- 4 Start Control Center.

```
systemctl start serviced
```

When all resource pool hosts are upgraded, you may log in to the Control Center browser interface, and then restart Resource Manager, if desired.

Converting the data storage driver

At this point in the upgrade, Control Center is upgraded on the master host and on all of the resource pool hosts. You may continue using Resource Manager to monitor your infrastructure during the initial (and longest) part of the conversion process. However, Zenoss strongly recommends running Resource Manager in this configuration only as long as it takes to migrate Btrfs data to the new thin pool.

The procedures in this section migrate Resource Manager data from a Btrfs volume to a device mapper thin pool volume.

Starting the conversion

To perform this procedure, you need an unused primary partition that is large enough for your data. For more information, see [Preparing to convert the data storage driver](#) on page 47.

This procedure converts the majority of Resource Manager data stored in a Btrfs file system and copies it to a device mapper thin pool. This procedure may be performed while Resource Manager is monitoring your infrastructure.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Create a thin pool for Resource Manager data storage.
 - a Identify the primary partition for the thin pool to create, and the amount of space available on the primary partition.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

For more information about the output of the `lsblk` command and primary partitions, refer to the *Zenoss Resource Manager Planning Guide*.

- b Create a variable for 50% of the space available on the primary partition for the thin pool to create.
The thin pool stores application data and snapshots of the data. You can add storage to the pool at any time.

Replace *Half-Of-Available-Space* with 50% of the space available in the primary partition, in gigabytes. Include the symbol for gigabytes (G) after the numeric value.

```
myFifty=Half-Of-Available-SpaceG
```

- c Create the thin pool.

Replace *Path-To-Device* with the path of the target primary partition:

```
serviced-storage create-thin-pool -o dm.basesize=$myFifty \  
serviced Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 3 Initialize `serviced-storage` for the `btrfs` storage driver.

```
serviced-storage init /opt/serviced/var/volumes btrfs
```

- 4 Synchronize the contents of the current data storage with the new thin pool.

This step copies the current data to the new thin pool. A small amount of metadata is stored in the temporary directory you specify, which must have approximately 1GB of available space.

Replace *Thin-Pool-Name* with the name of the thin pool created previously, and replace *Temporary-Directory* with the path of a temporary directory (for example, `/tmp/tmpdata`):

```
serviced-storage sync -c -t devicemapper \  
-o dm.thinpooldev=Thin-Pool-Name \  
/opt/serviced/var/volumes Temporary-Directory
```

After a pause to compute space requirements, the `serviced-storage` command displays both detail and summary information about its work. Depending on the amount of data to convert, this step may take several hours.

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\./ { print $1 }'
```

- b Identify the applications that are running.

Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Finalizing the conversion

This procedure completes the conversion started previously.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Stop the Control Center service.

```
systemctl stop serviced
```

- 3 Synchronize the contents of the current data storage with the new thin pool.
This step copies data created since the previous sync operation.

Repeat the command you used before stopping the Control Center service:

```
serviced-storage sync -c -t devicemapper \
-o dm.thinpooldev=Thin-Pool-Name \
/opt/serviced/var/volumes Temporary-Directory
```

- 4 Unmount the Btrfs file system.

```
umount /opt/serviced/var/volumes
```

- 5 Unmount the export of the Btrfs file system.

```
umount -fl /exports/serviced_volumes_v2/*
systemctl restart nfs
```

- 6 Move the new metadata to `/opt/serviced/var/volumes`.

- a Disable the new thin pool.

Replace *Thin-Pool-Name* with the name of the thin pool created previously, and replace *Temporary-Directory* with the path of your temporary directory:

```
serviced-storage disable Temporary-Directory \
-o dm.thinpooldev=Thin-Pool-Name
```

- b Move the metadata contents to `/opt/serviced/var/volumes`.

Replace *Temporary-Directory* with the path of your temporary directory:

```
mv Temporary-Directory/.devicemapper \
Temporary-Directory/* /opt/serviced/var/volumes
```

- 7 Remove entries for Btrfs volumes from `/etc/fstab`, if necessary.

- a Determine whether `/etc/fstab` includes entries for Btrfs volumes.

```
awk '!/^.*#/ { print $2, $3 }' /etc/fstab
```

If the result includes `/opt/serviced/var/volumes` or `/exports/serviced_var_volumes_v2/`, perform the the following substeps.

- b Open `/etc/fstab` with a text editor.

- c Remove or comment the entries that include `/opt/serviced/var/volumes` and `/exports/serviced_var_volumes_v2/`.
- d Save the file, and then close the text editor.
- 8 Update the Control Center configuration file.
 - a Open `/etc/default/serviced` in a text editor.
 - b Find the `SERVICED_FS_TYPE` variable declaration, and then change the value from `btrfs` to `devicemapper`.
 - c Remove the number sign character (`#`) from the beginning of the line, if necessary.
 - d Add `SERVICED_DM_THINPOOLDEV` immediately after `SERVICED_FS_TYPE`. Replace `Thin-Pool-Name` with the name of the thin pool created previously,

```
SERVICED_DM_THINPOOLDEV=Thin-Pool-Name
```

- e Save the file, and then close the editor.
- 9 Start the Control Center service.

```
systemctl start serviced
```

Verifying the conversion

This procedure verifies that the new data storage works properly with Resource Manager version 1.0.x.

- 1 Log in to the Control Center browser interface.
- 2 Start Resource Manager and all related applications.
- 3 Perform database integrity checks.
For more information, see [Using Zenoss Toolbox](#) on page 163.
- 4 Create a backup.
Previous backups are incompatible with the new storage driver.

When you are satisfied, continue with the Resource Manager upgrade procedure.

ZooKeeper ensemble configuration

Control Center relies on [Apache ZooKeeper](#) to coordinate its services. The procedures in this section create a ZooKeeper ensemble of 3 nodes. To perform these procedures, you need a Control Center master host and a minimum of two resource pool hosts. Each resource pool host requires a separate primary partition for Control Center internal services, and each should have a static IP address. For more information about storage requirements, refer to the *Zenoss Resource Manager Planning Guide*.

Note Zenoss strongly recommends configuring a ZooKeeper ensemble for all production deployments.

A ZooKeeper ensemble requires a minimum of 3 nodes, and 3 nodes is sufficient for most deployments. A 5-node configuration improves failover protection during maintenance windows. Ensembles larger than 5 nodes are not necessary. An odd number of nodes is recommended, and an even number of nodes is strongly discouraged.

Note The Control Center ZooKeeper service requires consistently fast storage. Ideally, the primary partition for Control Center internal services is on a separate, high-performance device that has only one primary partition.

Control Center variables for ZooKeeper

This tables in this section associates the ZooKeeper-related Control Center variables to set in `/etc/default/serviced` with the roles that hosts play in a Control Center cluster.

Table 3: Control Center master host

<i>SERVICED_ISVCS_ZOOKEEPER_ID</i>
The unique identifier of a ZooKeeper ensemble node. Value: 1
<i>SERVICED_ISVCS_ZOOKEEPER_QUORUM</i>
The ZooKeeper node ID, IP address, peer communications port, and leader communications port of each host in an ensemble. Each quorum definition must be unique, so the IP address of the "current" host is 0.0.0.0. Value: <i>ZooKeeper-ID@IP-Address:2888:3888, ...</i>
<i>SERVICED_ZK</i>
The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it. Value: <i>IP-Address:2181, ...</i>

Table 4: Control Center resource pool host and ZooKeeper ensemble node

<i>SERVICED_ISVCS_ZOOKEEPER_ID</i>
The unique identifier of a ZooKeeper ensemble node. Value: 2 or 3
<i>SERVICED_ISVCS_ZOOKEEPER_QUORUM</i>
The ZooKeeper node ID, IP address, peer communications port, and leader communications port of each host in an ensemble. Each quorum definition must be unique, so the IP address of the "current" host is 0.0.0.0. Value: <i>ZooKeeper-ID@IP-Address:2888:3888, ...</i>
<i>SERVICED_ISVCS_START</i>
The list of Control Center internal services to start and run on hosts other than the master host. Value: <i>zookeeper</i>
<i>SERVICED_ZK</i>
The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it. Value: <i>IP-Address:2181, ...</i>

Table 5: Control Center resource pool host

<i>SERVICED_ZK</i>
The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it. Value: <i>IP-Address:2181, ...</i>

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Configuring the master host as ZooKeeper node

This procedure configures the Control Center master host as a member of the ZooKeeper ensemble.

Note For accuracy, this procedure constructs Control Center configuration variables in the shell and appends them to `/etc/default/serviced`. The last step is to move the variables from the end of the file to more appropriate locations.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Create a variable for each Control Center host to include in the ZooKeeper ensemble.
The variables are used in subsequent steps.

Note Define the variables identically on the master host and on each resource pool host.

Replace *Master-Host-IP* with the IP address of the Control Center master host, and replace *Pool-Host-A-IP* and *Pool-Host-B-IP* with the IP addresses of the Control Center resource pool hosts to include in the ensemble:

```
node1=Master-Host-IP
node2=Pool-Host-A-IP
node3=Pool-Host-B-IP
```

Note ZooKeeper requires IP addresses for ensemble configuration.

- 3 Set the ZooKeeper node ID to 1.

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=1" >> /etc/default/serviced
```

- 4 Specify the nodes in the ZooKeeper ensemble.
You may copy the following text and paste it in your console:

```
echo "SERVICED_ZK=${node1}:2181,${node2}:2181,${node3}:2181" \
```

```
>> /etc/default/serviced
```

- 5 Specify the nodes in the ZooKeeper quorum.

ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address of the current node with 0.0.0.0.

You may copy the following of text and paste it in your console:

```
q1="1@0.0.0.0:2888:3888"
q2="2@${node2}:2888:3888"
q3="3@${node3}:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

- 6 Clean up the Control Center configuration file.
 - a Open `/etc/default/serviced` with a text editor.
 - b Navigate to the end of the file, and cut the line that contains the `SERVICED_ZK` variable declaration at that location.

The value of this declaration specifies 3 hosts.
 - c Locate the `SERVICED_ZK` variable near the beginning of the file, and then delete the line it is on.

The value of this declaration is just the master host.
 - d Paste the `SERVICED_ZK` variable declaration from the end of the file in the location of the just-deleted declaration.
 - e Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration at that location.
 - f Locate the `SERVICED_ISVCS_ZOOKEEPER_ID` variable near the end of the file, and then delete the line it is on.

This declaration is commented out.
 - g Paste the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration from the end of the file in the location of the just-deleted declaration.
 - h Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration at that location.
 - i Locate the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable near the end of the file, and then delete the line it is on.

This declaration is commented out.
 - j Paste the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration from the end of the file in the location of the just-deleted declaration.
 - k Save the file, and then close the text editor.
- 7 Verify the ZooKeeper environment variables.

```
egrep '^[^#]*SERVICED' /etc/default/serviced | egrep '(_ZOO|_ZK)'
```

Configuring a resource pool host as a ZooKeeper node

To perform this procedure, you need a resource pool host with an XFS file system on a separate partition. For more information, see the next topic.

This procedure configures a ZooKeeper ensemble on a resource pool host. Repeat this procedure on each Control Center resource pool host to add to the ZooKeeper ensemble.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Create a variable for each Control Center host to include in the ZooKeeper ensemble.

The variables are used in subsequent steps.

Note Define the variables identically on the master host and on each resource pool host.

Replace *Master-Host-IP* with the IP address of the Control Center master host, and replace *Pool-Host-A-IP* and *Pool-Host-B-IP* with the IP addresses of the Control Center resource pool hosts to include in the ensemble:

```
node1=Master-Host-IP
node2=Pool-Host-A-IP
node3=Pool-Host-B-IP
```

Note ZooKeeper requires IP addresses for ensemble configuration.

- 3 Set the ID of this node in the ZooKeeper ensemble.

For *Pool-Host-A-IP* (**node2**), use the following command:

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=2" >> /etc/default/serviced
```

For *Pool-Host-B-IP* (**node3**), use the following command:

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=3" >> /etc/default/serviced
```

- 4 Specify the nodes in the ZooKeeper ensemble.
You may copy the following text and paste it in your console:

```
echo "SERVICED_ZK=${node1}:2181,${node2}:2181,${node3}:2181" \
>> /etc/default/serviced
```

- 5 Specify the nodes in the ZooKeeper quorum.

ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address of the current node with 0.0.0.0.

For *Pool-Host-A-IP* (**node2**), use the following commands:

```
q1="1@${node1}:2888:3888"
q2="2@0.0.0.0:2888:3888"
q3="3@${node3}:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

For *Pool-Host-B-IP* (**node3**), use the following commands:

```
q1="1@${node1}:2888:3888"
q2="2@${node2}:2888:3888"
q3="3@0.0.0.0:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

- 6 Set the *SERVICED_ISVCS_START* variable, and clean up the Control Center configuration file.
 - a Open */etc/default/serviced* with a text editor.
 - b Locate the *SERVICED_ISVCS_START* variable, and then delete all but *zookeeper* from its list of values.
 - c Remove the number sign character (#) from the beginning of the line.
 - d Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK* variable declaration at that location.

The value of this declaration specifies 3 hosts.

- e Locate the `SERVICED_ZK` variable near the beginning of the file, and then delete the line it is on.
The value of this declaration is just the master host.
 - f Paste the `SERVICED_ZK` variable declaration from the end of the file in the location of the just-deleted declaration.
 - g Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration at that location.
 - h Locate the `SERVICED_ISVCS_ZOOKEEPER_ID` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - i Paste the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration from the end of the file in the location of the just-deleted declaration.
 - j Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration at that location.
 - k Locate the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - l Paste the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration from the end of the file in the location of the just-deleted declaration.
 - m Save the file, and then close the text editor.
- 7 Verify the ZooKeeper environment variables.

```
egrep '^[^#]*SERVICED' /etc/default/serviced \
| egrep '(_ZOO|_ZK|_STA)'
```

- 8 Pull the required Control Center ZooKeeper image from the master host.
- a Identify the image to pull.

```
serviced version | grep IsvcsImages
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v40 zenoss/isvcs-zookeeper:v3]
```

- b Pull the Control Center ZooKeeper image.

Replace `Isvcs-ZK-Image` with the name and version number of the ZooKeeper image from the previous substep:

```
docker pull Isvcs-ZK-Image
```

Creating a file system for internal services

This procedure creates an XFS file system on a primary partition. For more information about primary partitions, refer to the *Zenoss Resource Manager Planning Guide*.

Note Control Center internal services include ZooKeeper, which requires consistently fast storage. Zenoss recommends using a separate, high-performance storage resource for Control Center internal services. For example, a drive that is configured with only one primary partition, which eliminates contention by other services.

- 1 Log in to the target host as `root`, or as a user with superuser privileges.

- 2 Identify the target primary partition for the file system to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

For more information about the output of the `lsblk` command, and about creating primary partitions, refer to the *Zenoss Resource Manager Planning Guide*.

- 3 Create an XFS file system.
Replace *Partition* with the path of the target primary partition:

```
mkfs -t xfs Partition
```

- 4 Add an entry to the `/etc/fstab` file.
Replace *Partition* with the path of the primary partition used in the previous step:

```
echo "Partition \  
/opt/serviced/var/isvcs xfs defaults 0 0" >> /etc/fstab
```

- 5 Create the mount point for internal services data.

```
mkdir -p /opt/serviced/var/isvcs
```

- 6 Mount the file system, and then verify it mounted correctly.

```
mount -a && mount | grep isvcs
```

Example result:

```
/dev/xvdb1 on /opt/serviced/var/isvcs type xfs  
(rw,relatime,seclabel,attr2,inode64,noquota)
```

Starting a ZooKeeper ensemble

This procedure starts a ZooKeeper ensemble.

The window of time for starting a ZooKeeper ensemble is relatively short. The goal of this procedure is to restart Control Center on each ensemble node at about the same time, so that each node can participate in electing the leader.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the second node of the ZooKeeper ensemble (*Pool-Host-A-IP*).
- 3 In another separate window, log in to the third node of the ZooKeeper ensemble (*Pool-Host-B-IP*).
- 4 On the master host, stop and start `serviced`.

```
systemctl stop serviced && systemctl start serviced
```

- 5 On both resource pool hosts, stop and start `serviced`.

```
systemctl stop serviced && systemctl start serviced
```

- 6 On the master host, check the status of the ZooKeeper ensemble.

```
{ echo stats; sleep 1; } | nc localhost 2181 | grep Mode  
{ echo stats; sleep 1; } | nc Pool-Host-A-IP 2181 | grep Mode  
{ echo stats; sleep 1; } | nc Pool-Host-B-IP 2181 | grep Mode
```

If `nc` is not available, you can use `telnet` with [interactive ZooKeeper commands](#).

- 7 Optional: Log in to the Control Center browser interface, and then start Resource Manager and related applications, if desired.

The next procedure requires stopping Resource Manager.

Updating resource pool hosts

The default configuration of resource pool hosts sets the value of the `SERVICED_ZK` variable to the master host only. This procedure updates the setting to include the full ZooKeeper ensemble.

Perform this procedure on each resource pool host in your Control Center cluster.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Update the variable.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the `SERVICED_ZK` declaration, and then replace its value with the same value used in the ZooKeeper ensemble nodes.
 - c Save the file, and then close the editor.
- 3 Restart Control Center.

```
systemctl restart serviced
```

Master host isolation

Control Center enables or just performs rapid recovery from application service failures. When Control Center internal services and application services share a host, application failures can limit recovery options. Zenoss strongly recommends isolating the Control Center master host in a separate resource pool.

Note A master host in a separate resource pool requires fewer RAM and CPU resources than a master host that runs application services. For more information, refer to the *Zenoss Resource Manager Planning Guide*.

To perform the steps in this section, you need a Control Center master host and a minimum of one resource pool host. If you are upgrading a single-host deployment, skip this section.

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\./ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Moving the master host to new resource pool

This procedure creates a new resource pool for the Control Center master host, and then moves the master host to the new pool.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Create a new resource pool named `master`.

```
serviced pool add master
```

- 3 Display the `serviced` identifiers of all Control Center cluster hosts.
`serviced host list`
 The first column of the results table contains the `serviced` identifiers.
- 4 Remove the master host from the cluster.

Replace *Serviced-ID* with the ID of the master host:

```
serviced host remove Serviced-ID
```

- 5 Add the master host to the cluster, in the `master` resource pool.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
serviced host add --memory=0 Hostname-Or-IP:4979 master
```

If you enter a hostname, all hosts in your Control Center cluster must be able to resolve the name, either through an entry in `/etc/hosts`, or through a nameserver on your network.

When the value of the memory flag is 0, Control Center uses up to 100% of non-system RAM for its processes, even if the host's memory increases or decreases. When the value is 100 or any other value, Control Center uses a fixed amount of memory; the amount is based on the percentage of available memory when the host is added to Control Center.

- 6 Update the Control Center configuration file.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the `SERVICED_AGENT` declaration.
 - c Change the value from 1 to 0.
 - d Remove the number sign character (`#`) from the beginning of the line, if necessary.
 - e Locate the `SERVICED_MASTER_POOLID` declaration.
 - f Change the value from `default` to `master`.
 - g Remove the number sign character (`#`) from the beginning of the line, if necessary.
 - h Save the file, and then close the editor.
- 7 Restart the Control Center service.

```
systemctl restart serviced
```

- 8 Optional: After a few minutes, log in to the Control Center browser interface, and then start Resource Manager, if desired.

The next procedure requires stopping Resource Manager.

Upgrading Resource Manager

Perform this procedure to upgrade Resource Manager.

Before performing this procedure, stop Resource Manager. For more information, see [Stopping applications](#) on page 13.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Download the primary Docker image of Resource Manager for this release

The download takes approximately 10-20 minutes.

```
docker run -it --rm -v /root:/mnt/root \
  zenoss/resmgr_5.1:5.1.3_1 rsync -a /root/5.1.x /mnt/root
```

When the download completes, the `rsync` command copies scripts that perform the upgrade to `/root/5.1.x`.

- 3 Pull additional images for Resource Manager from Docker Hub.

```
/root/5.1.x/pull-docker-images.sh
```

- 4 Start the upgrade script.

The script to start depends on whether Service Impact is installed.

- If Service Impact is installed, enter the following command.

```
/root/5.1.x/upgrade-impact-5.1.x.sh
```

Note The script upgrades Resource Manager, but does not upgrade Service Impact. For more information about upgrading Service Impact, refer to the *Zenoss Service Impact Installation Guide for Resource Manager 5.x*.

- If Service Impact is not installed, enter the following command.

```
/root/5.1.x/upgrade-resmgr-5.1.x.sh
```

- 5 Restart Resource Manager.

Some Resource Manager services are started during the upgrade, and they need to be restarted.

```
serviced service restart Zenoss.resmgr
```

Proceed to [After upgrading](#) on page 158.

Upgrading Resource Manager without internet access

3

This chapter includes detailed procedures for upgrading Control Center cluster hosts that do not have internet access. For hosts that do have internet access, or that are configured for high-availability, use one of the other chapters in this part.

Upgrading Control Center on the master host

This section describes how to upgrade Control Center on the master host of a Control Center cluster. The upgrade process includes the following procedures:

- 1 Copy and stage upgrade source files.
- 2 Stop the application or applications that Control Center is managing, and then stop Control Center.
- 3 Save copies of Resource Manager images, if necessary.
- 4 Upgrade Docker and the operating system, if necessary.
- 5 Configure Control Center.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before proceeding. For more information, see [Using Zenoss Toolbox](#) on page 163.

Downloading files for offline installation

This procedure describes how to download RPM packages and Docker image files to your workstation.

To perform this procedure, you need:

- A workstation with internet access.
 - A portable storage medium, such as a USB flash drive, with at least 5 GB of free space.
 - Permission to download the required files from the [File Portal - Download Zenoss Enterprise Software](#) site. You may request permission by filing a ticket at the [Zenoss Support](#) site.
- 1 In a web browser, navigate to the [File Portal - Download Zenoss Enterprise Software](#) site.
 - 2 Log in with the account provided by Zenoss Support.
 - 3 Download archive files to your workstation.

Replace *Version* with the most recent version number available on the download page:

- `install-registry:Version.run`
- `install-zenoss-hbase:vVersion.run`
- `install-zenoss-isvcs-zookeeper:vVersion.run`

- `install-zenoss-opentsdb:vVersion.run`
 - `install-zenoss-resmgr_5.1:5.1Version.run`
 - `install-zenoss-serviced-isvcs:vVersion.run`
 - `serviced-resource-agents-Version.x86_64.rpm`
- 4 Download the RHEL/CentOS mirror package for your upgrade.

Note If you are planning to upgrade the operating system during your Control Center and Resource Manager upgrade, choose the mirror package that matches the RHEL/CentOS release to which you are upgrading, not the release that is installed now.

Replace *Version* with the most recent version number available on the download page, and replace *Release* with the version of RHEL/CentOS appropriate for your environment:

```
yum-mirror-centos7.Release-Version.x86_64.rpm
```

- 5 Copy the files to your portable storage medium.

Staging files for offline installation

Before performing this procedure, verify that approximately 4GB of temporary space is available on the file system where `/root` is located.

This procedure adds files for offline installation to the Control Center master host. The staged files are required in subsequent procedures.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Copy the archive files from your portable storage medium to `/root`.
- 3 Set the file permissions of the self-extracting archive files to execute.

```
chmod +x /root/*.run
```

- 4 Change directory to `/root`.

```
cd /root
```

- 5 Install the Resource Manager repository mirror.

```
yum install -y ./yum-mirror-*.x86_64.rpm
```

- 6 Optional: Delete the package file, if desired.

```
rm ./yum-mirror-*.x86_64.rpm
```

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\/./ { print $1 }'
```

- b Identify the applications that are running.

Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.

- a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Saving copies of Resource Manager images

This procedure ensures that Resource Manager images are present and up-to-date in the Control Center registry.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Determine whether Control Center is running.

```
systemctl status serviced | grep Active
```

If the result does not include `active (running)`, enter the following command:

```
systemctl start serviced
```

- 3 Determine whether the `SERVICED_REGISTRY` variable is set.

```
egrep '^[^#]*SERVICED_REGISTRY' /etc/default/serviced
```

- If the result is `SERVICED_REGISTRY=1`, stop performing this procedure, and continue to the next.
 - If the result is `SERVICED_REGISTRY=0`, or if the command returns no result, complete the remaining steps of this procedure before continuing to the next.
- 4 Edit the Control Center configuration file.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the `SERVICED_REGISTRY` declaration, and then set its value to 1.
 - c Remove the number sign character (`#`) from the beginning of the line.
 - d Save the file, and then close the editor.
 - e Restart Control Center.

```
systemctl stop serviced && systemctl start serviced
```

- 5 Synchronize the Docker and Control Center registries.

```
serviced docker sync
```

The synchronization may take approximately 20-30 minutes.

Upgrading Docker

This procedure upgrades Docker from version 1.5 or 1.8.2 to 1.9.0.

Note This release of Control Center requires RHEL/CentOS 7.1 or 7.2, and this procedure includes steps for upgrading the operating system, if necessary or desired. Zenoss recommends upgrading the operating system only at the step specified in this procedure. However, instructions for performing an upgrade without internet access are not included.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Stop Control Center.

```
systemctl stop serviced
```

- 3 Determine which version of Docker is installed, and then remove it, without removing Control Center.

```
rpm -qa | grep docker
```

- If the result includes `zenoss-docker`, enter the following command:

```
rpm -e --nodeps zenoss-docker-1.5.0-2
```

- If the result includes `docker-engine`, enter the following command:

```
rpm -e --nodeps docker-engine-1.8.2
```

- 4 Remove the Docker data partition.
 - a Identify the partition where `/var/lib/docker` is mounted.

```
mount | awk '/\var\/lib\/docker/ { print $1 }'
```

- b Unmount the partition.

Replace *Partition* with the device returned in the previous substep:

```
umount Partition
```

- c Erase the XFS file system on the partition.

Replace *Partition* with the device returned previously:

```
wipefs -a Partition
```

The partition is now ready for use as a device mapper thin pool.

- d Open `/etc/fstab` with a text editor.
 - e Remove the entry for `/var/lib/docker`.
- 5 Upgrade the operating system, if necessary.

- a Determine which release is installed.

```
cat /etc/redhat-release
```

If the result includes `7.0`, perform the following substeps.

- b Disable the `serviced` service.

```
systemctl disable serviced
```

- c Upgrade the operating system.
You may use a local mirror or other upgrade media.

- d Restart the operating system.

```
reboot
```

- e Log in to the Control Center cluster host as `root`, or as a user with superuser privileges.
- f Enable the `serviced` service.

```
systemctl enable serviced
```

- 6 Install Docker 1.9.0.

```
yum clean all && yum makecache fast
yum --enablerepo=zenoss-mirror install -y docker-engine
```

The result may include a warning about `serviced` requirements, which may be ignored.

- 7 Create a symbolic link for the Docker temporary directory.

Docker uses its temporary directory to spool images. The default directory is `/var/lib/docker/tmp`. The following command specifies the same directory that Control Center uses, `/tmp`. You can specify any directory that has a minimum of 10GB of unused space.

- a Create the `docker` directory in `/var/lib`.

```
mkdir /var/lib/docker
```

- b Create the link to `/tmp`.

```
ln -s /tmp /var/lib/docker/tmp
```

- 8 Create a `systemd` override file for the Docker service definition.

- a Create the override directory.

```
mkdir -p /etc/systemd/system/docker.service.d
```

- b Create the override file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/docker daemon $OPTIONS -H fd://
EOF
```

- c Reload the `systemd` manager configuration.

```
systemctl daemon-reload
```

- 9 Configure the `docker` service to start when the system starts.

```
systemctl enable docker
```

- 10 Install Control Center.

Control Center includes a utility that simplifies the process of creating a device mapper thin pool.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 11 Create a device mapper thin pool for Docker data.
 - a Identify the primary partition for the thin pool to create.

```
lsblk --output=NAME, SIZE, TYPE, FSTYPE, MOUNTPOINT
```

- b Create the thin pool.

Replace *Path-To-Device* with the path of an unused primary partition:

```
serviced-storage create-thin-pool docker Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 12 Configure and start the Docker service.

- a Create variables for adding arguments to the Docker configuration file.

The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

Replace *Thin-Pool-Device* with the name of the thin pool device created in the previous step:

```
myDriver="--s devicemapper"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myFlag="--storage-opt dm.thinpooldev"
myPool="Thin-Pool-Device"
```

- b Add the arguments to the Docker configuration file.

```
echo 'OPTIONS="'$myDriver $myFix $myFlag'='$myPool'"' \
>> /etc/sysconfig/docker
```

- c Start or restart Docker.

```
systemctl restart docker
```

The initial startup takes up to a minute, and may fail. If the startup fails, repeat the previous command.

- 13 Configure name resolution in containers.

Each time it starts, `docker` selects an IPv4 subnet for its virtual Ethernet bridge. The selection can change; this step ensures consistency.

- a Identify the IPv4 subnet and netmask `docker` has selected for its virtual Ethernet bridge.

```
ip addr show docker0 | grep inet
```

- b Open `/etc/sysconfig/docker` in a text editor.
 - c Add the following flags to the end of the `OPTIONS` declaration.

Replace *Bridge-Subnet* with the IPv4 subnet `docker` selected for its virtual bridge, and replace *Bridge-Netmask* with the netmask `docker` selected:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/Bridge-Netmask
```

For example, if the bridge subnet and netmask is `172.17.0.1/16`, the flags to add are `--dns=172.17.0.1`
`--bip=172.17.0.1/16`.

Note Leave a blank space after the end of the thin pool device name, and make sure the double quote character (") is at the end of the line.

- d Restart the Docker service.

```
systemctl restart docker
```

Configuring Control Center on the master host

Perform this procedure to upgrade Control Center on the master host.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Manage the Control Center configuration files.
 - a Change directory to the configuration file directory.

```
cd /etc/default
```

- b Rename the existing configuration file, as a backup.

```
mv serviced serviced.pre-1.1.4
```

- c Rename the new configuration file, and then create a copy of it.

```
mv serviced.rpmnew serviced.orig-1.1.4
cp serviced.orig-1.1.4 serviced
```

- 3 Copy settings from the previous Control Center configuration file to the new one.
 - a Identify the customized variables in the pre-upgrade configuration file.

```
egrep '^[^#]*SERVICED' serviced.pre-1.1.4
```

- b Open `/etc/default/serviced` with a text editor, and then customize the same variables that were customized in the pre-upgrade configuration file.

The following variables are deprecated, and are not needed in the new configuration file:

- `SERVICED_REGISTRY`
- `SERVICED_VARPATH`

- c Add `SERVICED_DOCKER_REGISTRY` to the file.

The variable specifies the host and port at which the local Docker registry is available.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
SERVICED_DOCKER_REGISTRY=Hostname-Or-IP:5000
```

- d Save the file and close the text editor.
- 4 Install the new Control Center images.

- a Change directory to `/root`.

```
cd /root
```

- b Install the Control Center ZooKeeper image.

```
./install-zenoss-isvcs-zookeeper-*.run
```

- c Install the Control Center image.

```
./install-zenoss-serviced-isvcs-*.run
```


- d Install the registry image.

```
./install-registry*.run
```

- 5 Start Control Center and watch its log file.

```
systemctl start serviced && journalctl -flu serviced
```

Control Center copies Resource Manager images from its registry to the Docker library.

- **Single-host deployments:** Log in to the Control Center browser interface, and then restart Resource Manager.
- **Multi-host deployments:** Update all resource pool hosts (the next section).

Upgrading Control Center on resource pool hosts

This section provides procedures for upgrading Control Center on resource pool hosts. Repeat these procedures on each resource pool host in your Control Center cluster.

Note Before performing these procedures, upgrade Control Center on the master host.

Staging files for offline installation

To perform this procedure, you need the portable storage medium that contains the archive files used in installing the master host.

This procedure adds files for offline installation to a resource pool host. The files are required in subsequent procedures.

Perform this procedure on each resource pool host in your deployment.

- 1 Log in to the target host as `root`, or as a user with superuser privileges.
- 2 Copy `yum-mirror-*.x86_64.rpm` from your portable storage medium to `/tmp`.
- 3 Install the Resource Manager repository mirror.

```
yum install -y /tmp/yum-mirror-*.x86_64.rpm
```

- 4 Optional: Delete the package file, if desired.

```
rm /tmp/yum-mirror-*.x86_64.rpm
```

Upgrading Docker

This procedure upgrades Docker from version 1.5 or 1.8.2 to 1.9.0.

Note This release of Control Center requires RHEL/CentOS 7.1 or 7.2, and this procedure includes steps for upgrading the operating system, if necessary or desired. Zenoss recommends upgrading the operating system only at the step specified in this procedure. However, instructions for performing an upgrade without internet access are not included.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Stop Control Center.

```
systemctl stop serviced
```

- 3 Determine which version of Docker is installed, and then remove it, without removing Control Center.

```
rpm -qa | grep docker
```

- If the result includes `zenoss-docker`, enter the following command:

```
rpm -e --nodeps zenoss-docker-1.5.0-2
```

- If the result includes `docker-engine`, enter the following command:

```
rpm -e --nodeps docker-engine-1.8.2
```

- 4 Remove the Docker data partition.

- a Identify the partition where `/var/lib/docker` is mounted.

```
mount | awk '/\var\/lib\/docker/ { print $1 }'
```

- b Unmount the partition.

Replace *Partition* with the device returned in the previous substep:

```
umount Partition
```

- c Erase the XFS file system on the partition.

Replace *Partition* with the device returned previously:

```
wipefs -a Partition
```

The partition is now ready for use as a device mapper thin pool.

- d Open `/etc/fstab` with a text editor.
- e Remove the entry for `/var/lib/docker`.

- 5 Upgrade the operating system, if necessary.

- a Determine which release is installed.

```
cat /etc/redhat-release
```

If the result includes `7.0`, perform the following substeps.

- b Disable the `serviced` service.

```
systemctl disable serviced
```

- c Upgrade the operating system.

You may use a local mirror or other upgrade media.

- d Restart the operating system.

```
reboot
```

- e Log in to the Control Center cluster host as `root`, or as a user with superuser privileges.
- f Enable the `serviced` service.

```
systemctl enable serviced
```

6 Install Docker 1.9.0.

```
yum clean all && yum makecache fast
yum --enablerepo=zenoss-mirror install -y docker-engine
```

The result may include a warning about `serviced` requirements, which may be ignored.

7 Create a symbolic link for the Docker temporary directory.

Docker uses its temporary directory to spool images. The default directory is `/var/lib/docker/tmp`. The following command specifies the same directory that Control Center uses, `/tmp`. You can specify any directory that has a minimum of 10GB of unused space.

- a** Create the `docker` directory in `/var/lib`.

```
mkdir /var/lib/docker
```

- b** Create the link to `/tmp`.

```
ln -s /tmp /var/lib/docker/tmp
```

8 Create a `systemd` override file for the Docker service definition.

- a** Create the override directory.

```
mkdir -p /etc/systemd/system/docker.service.d
```

- b** Create the override file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/docker daemon \$OPTIONS -H fd://
EOF
```

- c** Reload the `systemd` manager configuration.

```
systemctl daemon-reload
```

9 Configure the `docker` service to start when the system starts.

```
systemctl enable docker
```

10 Install Control Center.

Control Center includes a utility that simplifies the process of creating a device mapper thin pool.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

11 Create a device mapper thin pool for Docker data.

- a** Identify the primary partition for the thin pool to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b** Create the thin pool.

Replace *Path-To-Device* with the path of an unused primary partition:

```
serviced-storage create-thin-pool docker Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

12 Configure and start the Docker service.

- a Create variables for adding arguments to the Docker configuration file.

The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

Replace *Thin-Pool-Device* with the name of the thin pool device created in the previous step:

```
myDriver="-s devicemapper"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myFlag="--storage-opt dm.thinpooldev"
myPool="Thin-Pool-Device"
```

- b Add the arguments to the Docker configuration file.

```
echo 'OPTIONS="'$myDriver $myFix $myFlag'='$myPool'"' \
>> /etc/sysconfig/docker
```

- c Start or restart Docker.

```
systemctl restart docker
```

The initial startup takes up to a minute, and may fail. If the startup fails, repeat the previous command.

13 Configure name resolution in containers.

Each time it starts, `docker` selects an IPv4 subnet for its virtual Ethernet bridge. The selection can change; this step ensures consistency.

- a Identify the IPv4 subnet and netmask `docker` has selected for its virtual Ethernet bridge.

```
ip addr show docker0 | grep inet
```

- b Open `/etc/sysconfig/docker` in a text editor.
- c Add the following flags to the end of the `OPTIONS` declaration.

Replace *Bridge-Subnet* with the IPv4 subnet `docker` selected for its virtual bridge, and replace *Bridge-Netmask* with the netmask `docker` selected:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/Bridge-Netmask
```

For example, if the bridge subnet and netmask is `172.17.0.1/16`, the flags to add are `--dns=172.17.0.1`
`--bip=172.17.0.1/16`.

Note Leave a blank space after the end of the thin pool device name, and make sure the double quote character (") is at the end of the line.

- d Restart the Docker service.

```
systemctl restart docker
```

Configuring Control Center on resource pool hosts

Perform this procedure after upgrading Docker.

- 1 Log in to the Control Center resource pool host as `root`, or as a user with superuser privileges.
- 2 Manage the Control Center configuration files.
 - a Change directory to the configuration file directory.

```
cd /etc/default
```

- b Rename the existing configuration file, as a backup.

```
mv serviced serviced.pre-1.1.4
```

- c Rename the new configuration file, and then create a copy of it.

```
mv serviced.rpmnew serviced.orig-1.1.4
cp serviced.orig-1.1.4 serviced
```

- 3 Copy settings from the previous Control Center configuration file to the new one.
 - a Identify the customized variables in the pre-upgrade configuration file.

```
egrep '^[^#]*SERVICED' serviced.pre-1.1.4
```

- b Open `/etc/default/serviced` with a text editor, and then customize the same variables that were customized in the pre-upgrade configuration file.

The following variables are deprecated, and are not needed in the new configuration file:

- `SERVICED_REGISTRY`
- `SERVICED_VARPATH`

- c Add `SERVICED_DOCKER_REGISTRY` to the file.

The variable specifies the host and port at which the local Docker registry is available.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
SERVICED_DOCKER_REGISTRY=Hostname-Or-IP:5000
```

- d Save the file and close the text editor.
- 4 Unmount the distributed file system (DFS).
 - a Identify the file system specification to unmount.

```
mount | awk '/serviced/ { print $1 }'
```

- b Unmount the DFS.

Replace *DFS-Mount* with the file system specification returned in the previous substep:

```
umount DFS-Mount
```

- 5 Start Control Center.

```
systemctl start serviced
```

When all resource pool hosts are upgraded, log in to the Control Center browser interface, and then restart Resource Manager.

Converting the data storage driver

At this point in the upgrade, Control Center is upgraded on the master host and on all of the resource pool hosts. You may continue using Resource Manager to monitor your infrastructure during the initial (and longest) part of the conversion process. However, Zenoss strongly recommends running Resource Manager in this configuration only as long as it takes to migrate Btrfs data to the new thin pool.

The procedures in this section migrate Resource Manager data from a Btrfs volume to a device mapper thin pool volume.

Starting the conversion

To perform this procedure, you need an unused primary partition that is large enough for your data. For more information, see [Preparing to convert the data storage driver](#) on page 47.

This procedure converts the majority of Resource Manager data stored in a Btrfs file system and copies it to a device mapper thin pool. This procedure may be performed while Resource Manager is monitoring your infrastructure.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Create a thin pool for Resource Manager data storage.
 - a Identify the primary partition for the thin pool to create, and the amount of space available on the primary partition.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

For more information about the output of the `lsblk` command and primary partitions, refer to the *Zenoss Resource Manager Planning Guide*.

- b Create a variable for 50% of the space available on the primary partition for the thin pool to create. The thin pool stores application data and snapshots of the data. You can add storage to the pool at any time. Replace *Half-Of-Available-Space* with 50% of the space available in the primary partition, in gigabytes. Include the symbol for gigabytes (G) after the numeric value.

```
myFifty=Half-Of-Available-SpaceG
```

- c Create the thin pool.

Replace *Path-To-Device* with the path of the target primary partition:

```
serviced-storage create-thin-pool -o dm.basesize=$myFifty \  
serviced Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 3 Initialize `serviced-storage` for the `btrfs` storage driver.

```
serviced-storage init /opt/serviced/var/volumes btrfs
```

- 4 Synchronize the contents of the current data storage with the new thin pool.

This step copies the current data to the new thin pool. A small amount of metadata is stored in the temporary directory you specify, which must have approximately 1GB of available space.

Replace *Thin-Pool-Name* with the name of the thin pool created previously, and replace *Temporary-Directory* with the path of a temporary directory (for example, `/tmp/tmpdata`):

```
serviced-storage sync -c -t devicemapper \  
-o dm.thinpooldev=Thin-Pool-Name \  
-o dm.thinpooldev=Thin-Pool-Name \  
-o dm.thinpooldev=Thin-Pool-Name \
```

```
/opt/serviced/var/volumes Temporary-Directory
```

After a pause to compute space requirements, the `serviced-storage` command displays both detail and summary information about its work. Depending on the amount of data to convert, this step may take several hours.

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\./ { print $1 }'
```

- b Identify the applications that are running.

Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the `STATUS` column reads `Stopped`:

```
serviced service status Zenoss.resmgr
```

Finalizing the conversion

This procedure completes the conversion started previously.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Stop the Control Center service.

```
systemctl stop serviced
```

- 3 Synchronize the contents of the current data storage with the new thin pool.
This step copies data created since the previous sync operation.

Repeat the command you used before stopping the Control Center service:

```
serviced-storage sync -c -t devicemapper \  
-o dm.thinpooldev=Thin-Pool-Name \  
/opt/serviced/var/volumes Temporary-Directory
```

- 4 Unmount the Btrfs file system.

```
umount /opt/serviced/var/volumes
```

- 5 Unmount the export of the Btrfs file system.

```
umount -fl /exports/serviced_volumes_v2/*
```

```
systemctl restart nfs
```

- 6 Move the new metadata to `/opt/serviced/var/volumes`.
 - a Disable the new thin pool.

Replace *Thin-Pool-Name* with the name of the thin pool created previously, and replace *Temporary-Directory* with the path of your temporary directory:

```
serviced-storage disable Temporary-Directory \
-o dm.thinpooldev=Thin-Pool-Name
```

- b Move the metadata contents to `/opt/serviced/var/volumes`.

Replace *Temporary-Directory* with the path of your temporary directory:

```
mv Temporary-Directory/.devicemapper \
Temporary-Directory/* /opt/serviced/var/volumes
```

- 7 Remove entries for Btrfs volumes from `/etc/fstab`, if necessary.
 - a Determine whether `/etc/fstab` includes entries for Btrfs volumes.

```
awk '!/^.*#/ { print $2, $3 }' /etc/fstab
```

If the result includes `/opt/serviced/var/volumes` or `/exports/serviced_var_volumes_v2/`, perform the the following substeps.

- b Open `/etc/fstab` with a text editor.
 - c Remove or comment the entries that include `/opt/serviced/var/volumes` and `/exports/serviced_var_volumes_v2/`.
 - d Save the file, and then close the text editor.
- 8 Update the Control Center configuration file.
 - a Open `/etc/default/serviced` in a text editor.
 - b Find the `SERVICED_FS_TYPE` variable declaration, and then change the value from `btrfs` to `devicemapper`.
 - c Remove the number sign character (`#`) from the beginning of the line, if necessary.
 - d Add `SERVICED_DM_THINPOOLDEV` immediately after `SERVICED_FS_TYPE`. Replace *Thin-Pool-Name* with the name of the thin pool created previously,

```
SERVICED_DM_THINPOOLDEV=Thin-Pool-Name
```

- e Save the file, and then close the editor.
- 9 Start the Control Center service.

```
systemctl start serviced
```

Verifying the conversion

This procedure verifies that the new data storage works properly with Resource Manager version 1.0.x.

- 1 Log in to the Control Center browser interface.
- 2 Start Resource Manager and all related applications.
- 3 Perform database integrity checks.

For more information, see [Using Zenoss Toolbox](#) on page 163.
- 4 Create a backup.

Previous backups are incompatible with the new storage driver.

When you are satisfied, continue with the Resource Manager upgrade procedure.

ZooKeeper ensemble configuration

Control Center relies on *Apache ZooKeeper* to coordinate its services. The procedures in this section create a ZooKeeper ensemble of 3 nodes. To perform these procedures, you need a Control Center master host and a minimum of two resource pool hosts. Each resource pool host requires a separate primary partition for Control Center internal services, and each should have a static IP address. For more information about storage requirements, refer to the *Zenoss Resource Manager Planning Guide*.

Note Zenoss strongly recommends configuring a ZooKeeper ensemble for all production deployments.

A ZooKeeper ensemble requires a minimum of 3 nodes, and 3 nodes is sufficient for most deployments. A 5-node configuration improves failover protection during maintenance windows. Ensembles larger than 5 nodes are not necessary. An odd number of nodes is recommended, and an even number of nodes is strongly discouraged.

Note The Control Center ZooKeeper service requires consistently fast storage. Ideally, the primary partition for Control Center internal services is on a separate, high-performance device that has only one primary partition.

Control Center variables for ZooKeeper

This tables in this section associates the ZooKeeper-related Control Center variables to set in `/etc/default/serviced` with the roles that hosts play in a Control Center cluster.

Table 6: Control Center master host

SERVICED_ISVCS_ZOOKEEPER_ID

The unique identifier of a ZooKeeper ensemble node.

Value: 1

SERVICED_ISVCS_ZOOKEEPER_QUORUM

The ZooKeeper node ID, IP address, peer communications port, and leader communications port of each host in an ensemble. Each quorum definition must be unique, so the IP address of the "current" host is 0.0.0.0.

Value: *ZooKeeper-ID@IP-Address:2888:3888, . . .*

SERVICED_ZK

The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it.

Value: *IP-Address:2181, . . .*

Table 7: Control Center resource pool host and ZooKeeper ensemble node

SERVICED_ISVCS_ZOOKEEPER_ID

The unique identifier of a ZooKeeper ensemble node.

Value: 2 or 3

SERVICED_ISVCS_ZOOKEEPER_QUORUM

The ZooKeeper node ID, IP address, peer communications port, and leader communications port of each host in an ensemble. Each quorum definition must be unique, so the IP address of the "current" host is 0.0.0.0.

Value: *ZooKeeper-ID@IP-Address:2888:3888,...*

SERVICED_ISVCS_START

The list of Control Center internal services to start and run on hosts other than the master host.

Value: *zookeeper*

SERVICED_ZK

The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it.

Value: *IP-Address:2181,...*

Table 8: Control Center resource pool host

SERVICED_ZK

The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it.

Value: *IP-Address:2181,...*

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\./ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Configuring the master host as ZooKeeper node

This procedure configures the Control Center master host as a member of the ZooKeeper ensemble.

Note For accuracy, this procedure constructs Control Center configuration variables in the shell and appends them to `/etc/default/serviced`. The last step is to move the variables from the end of the file to more appropriate locations.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Create a variable for each Control Center host to include in the ZooKeeper ensemble.
The variables are used in subsequent steps.

Note Define the variables identically on the master host and on each resource pool host.

Replace *Master-Host-IP* with the IP address of the Control Center master host, and replace *Pool-Host-A-IP* and *Pool-Host-B-IP* with the IP addresses of the Control Center resource pool hosts to include in the ensemble:

```
node1=Master-Host-IP
node2=Pool-Host-A-IP
node3=Pool-Host-B-IP
```

Note ZooKeeper requires IP addresses for ensemble configuration.

- 3 Set the ZooKeeper node ID to 1.

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=1" >> /etc/default/serviced
```

- 4 Specify the nodes in the ZooKeeper ensemble.
You may copy the following text and paste it in your console:

```
echo "SERVICED_ZK=${node1}:2181,${node2}:2181,${node3}:2181" \
>> /etc/default/serviced
```

- 5 Specify the nodes in the ZooKeeper quorum.

ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address of the current node with `0.0.0.0`.

You may copy the following of text and paste it in your console:

```
q1="1@0.0.0.0:2888:3888"
q2="2@${node2}:2888:3888"
q3="3@${node3}:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

- 6 Clean up the Control Center configuration file.
 - a Open `/etc/default/serviced` with a text editor.
 - b Navigate to the end of the file, and cut the line that contains the `SERVICED_ZK` variable declaration at that location.
The value of this declaration specifies 3 hosts.
 - c Locate the `SERVICED_ZK` variable near the beginning of the file, and then delete the line it is on.
The value of this declaration is just the master host.
 - d Paste the `SERVICED_ZK` variable declaration from the end of the file in the location of the just-deleted declaration.
 - e Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration at that location.
 - f Locate the `SERVICED_ISVCS_ZOOKEEPER_ID` variable near the end of the file, and then delete the line it is on.

This declaration is commented out.

- g** Paste the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration from the end of the file in the location of the just-deleted declaration.
 - h** Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration at that location.
 - i** Locate the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - j** Paste the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration from the end of the file in the location of the just-deleted declaration.
 - k** Save the file, and then close the text editor.
- 7 Verify the ZooKeeper environment variables.

```
egrep '^[^#]*SERVICED' /etc/default/serviced | egrep '(_ZOO|_ZK)'
```

Configuring a resource pool host as a ZooKeeper node

To perform this procedure, you need a resource pool host with an XFS file system on a separate partition. For more information, see the next topic.

This procedure configures a ZooKeeper ensemble on a resource pool host. Repeat this procedure on each Control Center resource pool host to add to the ZooKeeper ensemble.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Create a variable for each Control Center host to include in the ZooKeeper ensemble.

The variables are used in subsequent steps.

Note Define the variables identically on the master host and on each resource pool host.

Replace *Master-Host-IP* with the IP address of the Control Center master host, and replace *Pool-Host-A-IP* and *Pool-Host-B-IP* with the IP addresses of the Control Center resource pool hosts to include in the ensemble:

```
node1=Master-Host-IP
node2=Pool-Host-A-IP
node3=Pool-Host-B-IP
```

Note ZooKeeper requires IP addresses for ensemble configuration.

- 3 Set the ID of this node in the ZooKeeper ensemble.

For *Pool-Host-A-IP* (node2), use the following command:

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=2" >> /etc/default/serviced
```

For *Pool-Host-B-IP* (node3), use the following command:

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=3" >> /etc/default/serviced
```

- 4 Specify the nodes in the ZooKeeper ensemble.
You may copy the following text and paste it in your console:

```
echo "SERVICED_ZK=${node1}:2181,${node2}:2181,${node3}:2181" \
>> /etc/default/serviced
```

5 Specify the nodes in the ZooKeeper quorum.

ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address of the current node with 0.0.0.0.

For *Pool-Host-A-IP (node2)*, use the following commands:

```
q1="1@${node1}:2888:3888"
q2="2@0.0.0.0:2888:3888"
q3="3@${node3}:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

For *Pool-Host-B-IP (node3)*, use the following commands:

```
q1="1@${node1}:2888:3888"
q2="2@${node2}:2888:3888"
q3="3@0.0.0.0:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

6 Set the *SERVICED_ISVCS_START* variable, and clean up the Control Center configuration file.

- a Open `/etc/default/serviced` with a text editor.
- b Locate the *SERVICED_ISVCS_START* variable, and then delete all but `zookeeper` from its list of values.
- c Remove the number sign character (`#`) from the beginning of the line.
- d Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK* variable declaration at that location.
The value of this declaration specifies 3 hosts.
- e Locate the *SERVICED_ZK* variable near the beginning of the file, and then delete the line it is on.
The value of this declaration is just the master host.
- f Paste the *SERVICED_ZK* variable declaration from the end of the file in the location of the just-deleted declaration.
- g Navigate to the end of the file, and cut the line that contains the *SERVICED_ISVCS_ZOOKEEPER_ID* variable declaration at that location.
- h Locate the *SERVICED_ISVCS_ZOOKEEPER_ID* variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
- i Paste the *SERVICED_ISVCS_ZOOKEEPER_ID* variable declaration from the end of the file in the location of the just-deleted declaration.
- j Navigate to the end of the file, and cut the line that contains the *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable declaration at that location.
- k Locate the *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
- l Paste the *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable declaration from the end of the file in the location of the just-deleted declaration.
- m Save the file, and then close the text editor.

7 Verify the ZooKeeper environment variables.

```
egrep '^[^#]*SERVICED' /etc/default/serviced \
| egrep '(_ZOO|_ZK|_STA)'
```

8 Pull the required Control Center ZooKeeper image from the master host.

- a Identify the image to pull.

```
serviced version | grep IsvcsImages
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v40 zenoss/isvcs-zookeeper:v3]
```

- b Pull the Control Center ZooKeeper image.

Replace *Isvcs-ZK-Image* with the name and version number of the ZooKeeper image from the previous substep:

```
docker pull Isvcs-ZK-Image
```

Creating a file system for internal services

This procedure creates an XFS file system on a primary partition. For more information about primary partitions, refer to the *Zenoss Resource Manager Planning Guide*.

Note Control Center internal services include ZooKeeper, which requires consistently fast storage. Zenoss recommends using a separate, high-performance storage resource for Control Center internal services. For example, a drive that is configured with only one primary partition, which eliminates contention by other services.

- 1 Log in to the target host as `root`, or as a user with superuser privileges.
- 2 Identify the target primary partition for the file system to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

For more information about the output of the `lsblk` command, and about creating primary partitions, refer to the *Zenoss Resource Manager Planning Guide*.

- 3 Create an XFS file system.
Replace *Partition* with the path of the target primary partition:

```
mkfs -t xfs Partition
```

- 4 Add an entry to the `/etc/fstab` file.
Replace *Partition* with the path of the primary partition used in the previous step:

```
echo "Partition \  
/opt/serviced/var/isvcs xfs defaults 0 0" >> /etc/fstab
```

- 5 Create the mount point for internal services data.

```
mkdir -p /opt/serviced/var/isvcs
```

- 6 Mount the file system, and then verify it mounted correctly.

```
mount -a && mount | grep isvcs
```

Example result:

```
/dev/xvdb1 on /opt/serviced/var/isvcs type xfs  
(rw,relatime,seclabel,attr2,inode64,noquota)
```

Starting a ZooKeeper ensemble

This procedure starts a ZooKeeper ensemble.

The window of time for starting a ZooKeeper ensemble is relatively short. The goal of this procedure is to restart Control Center on each ensemble node at about the same time, so that each node can participate in electing the leader.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the second node of the ZooKeeper ensemble (*Pool-Host-A-IP*).
- 3 In another separate window, log in to the third node of the ZooKeeper ensemble (*Pool-Host-B-IP*).
- 4 On the master host, stop and start `serviced`.

```
systemctl stop serviced && systemctl start serviced
```

- 5 On both resource pool hosts, stop and start `serviced`.

```
systemctl stop serviced && systemctl start serviced
```

- 6 On the master host, check the status of the ZooKeeper ensemble.

```
{ echo stats; sleep 1; } | nc localhost 2181 | grep Mode
{ echo stats; sleep 1; } | nc Pool-Host-A-IP 2181 | grep Mode
{ echo stats; sleep 1; } | nc Pool-Host-B-IP 2181 | grep Mode
```

If `nc` is not available, you can use `telnet` with [interactive ZooKeeper commands](#).

- 7 Optional: Log in to the Control Center browser interface, and then start Resource Manager and related applications, if desired.

The next procedure requires stopping Resource Manager.

Updating resource pool hosts

The default configuration of resource pool hosts sets the value of the `SERVICED_ZK` variable to the master host only. This procedure updates the setting to include the full ZooKeeper ensemble.

Perform this procedure on each resource pool host in your Control Center cluster.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Update the variable.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the `SERVICED_ZK` declaration, and then replace its value with the same value used in the ZooKeeper ensemble nodes.
 - c Save the file, and then close the editor.
- 3 Restart Control Center.

```
systemctl restart serviced
```

Master host isolation

Control Center enables or just performs rapid recovery from application service failures. When Control Center internal services and application services share a host, application failures can limit recovery options. Zenoss strongly recommends isolating the Control Center master host in a separate resource pool.

Note A master host in a separate resource pool requires fewer RAM and CPU resources than a master host that runs application services. For more information, refer to the *Zenoss Resource Manager Planning Guide*.

To perform the steps in this section, you need a Control Center master host and a minimum of one resource pool host. If you are upgrading a single-host deployment, skip this section.

Stopping applications

This procedure stops all Control Center applications.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.

Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Moving the master host to new resource pool

This procedure creates a new resource pool for the Control Center master host, and then moves the master host to the new pool.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Create a new resource pool named `master`.

```
serviced pool add master
```

- 3 Display the `serviced` identifiers of all Control Center cluster hosts.


```
serviced host list
```

The first column of the results table contains the `serviced` identifiers.
- 4 Remove the master host from the cluster.

Replace *Serviced-ID* with the ID of the master host:

```
serviced host remove Serviced-ID
```

- 5 Add the master host to the cluster, in the `master` resource pool.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
serviced host add --memory=0 Hostname-Or-IP:4979 master
```

If you enter a hostname, all hosts in your Control Center cluster must be able to resolve the name, either through an entry in `/etc/hosts`, or through a nameserver on your network.

When the value of the memory flag is 0, Control Center uses up to 100% of non-system RAM for its processes, even if the host's memory increases or decreases. When the value is 100 or any other value, Control Center uses a fixed amount of memory; the amount is based on the percentage of available memory when the host is added to Control Center.

- 6 Update the Control Center configuration file.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the `SERVICED_AGENT` declaration.
 - c Change the value from 1 to 0.
 - d Remove the number sign character (#) from the beginning of the line, if necessary.
 - e Locate the `SERVICED_MASTER_POOLID` declaration.
 - f Change the value from `default` to `master`.
 - g Remove the number sign character (#) from the beginning of the line, if necessary.
 - h Save the file, and then close the editor.
- 7 Restart the Control Center service.

```
systemctl restart serviced
```

- 8 Optional: After a few minutes, log in to the Control Center browser interface, and then start Resource Manager, if desired.
The next procedure requires stopping Resource Manager.

Upgrading Resource Manager

This procedure upgrades Resource Manager.

Before performing this procedure, stop Resource Manager. For more information, see [Stopping applications](#) on page 13.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Install the new Resource Manager images.
 - a Change directory to `/root`.

```
cd /root
```

- b Install the HBase image.

```
./install-zenoss-hbase*.run
```

- c Install the OpenTSDB image.

```
./install-zenoss-opentsdb*.run
```

- d Install the Resource Manager image.

```
./install-zenoss-resmgr*.run
```

- e Optional: Delete the self-extracting image files, if desired.

```
rm ./install-zenoss-*.run
```

- 3 Extract files from the Resource Manager image.

```
docker run -it --rm -v /root:/mnt/root \
  zenoss/resmgr_5.1:5.1.3_1 rsync -a /root/5.1.x /mnt/root
```

The preceding commands copy upgrade scripts to `/root/5.1.x`.

- 4 Start the upgrade script.

The script to start depends on whether Service Impact is installed.

- If Service Impact is installed, enter the following command.

```
/root/5.1.x/upgrade-impact-5.1.x.sh
```

Note The script upgrades Resource Manager, but does not upgrade Service Impact. For more information about upgrading Service Impact, refer to the *Zenoss Service Impact Installation Guide for Resource Manager 5.x*.

- If Service Impact is not installed, enter the following command.

```
/root/5.1.x/upgrade-resmgr-5.1.x.sh
```

- 5 Restart Resource Manager.

Some Resource Manager services are started during the upgrade, and they need to be restarted.

```
serviced service restart Zenoss.resmgr
```

Proceed to [After upgrading](#) on page 158.

Upgrading high-availability deployments with internet access

4

This chapter includes detailed procedures for upgrading high-availability hosts that have internet access. For hosts that do not have internet access, or that are not configured for high-availability, use one of the other chapters in this part.

Key variables used in this chapter

The following tables associate important features of a high-availability deployment with the variables used in this chapter.

Feature	Variable Name	
	Primary Node	Secondary Node
Public IP address of master node (static; known to all machines in the Control Center cluster)	<i>Primary-Public-IP</i>	<i>Secondary-Public-IP</i>
Public hostname of master node (returned by <code>uname -n</code> ; resolves to the public IP address)	<i>Primary-Public-Name</i>	<i>Secondary-Public-Name</i>
Private IP address of master node (static; dual-NIC systems only)	<i>Primary-Private-IP</i>	<i>Secondary-Private-IP</i>
Private hostname of master node (resolves to the private IP address; dual-NIC systems only)	<i>Primary-Private-Name</i>	<i>Secondary-Private-Name</i>
Feature	Variable Name	
Virtual IP address of the high-availability cluster (static; known enterprise-wide)	<i>HA-Virtual-IP</i>	
Virtual hostname of the high-availability cluster (known enterprise-wide)	<i>HA-Virtual-Name</i>	
Public IP address of resource pool host A (static; for ZooKeeper ensemble)	<i>Pool-Host-A-IP</i>	
Public IP address of resource pool host B (static; for ZooKeeper ensemble)	<i>Pool-Host-B-IP</i>	
Primary partition for Btrfs file system (mirrored)	<i>Btrfs-Partition</i>	
Primary partition for Docker data (not mirrored)	<i>Docker-Partition</i>	

Feature	Variable Name
Primary partition for Control Center internal services data (mirrored)	<i>Isvcs-Partition</i>
Primary partition for Control Center metadata (mirrored)	<i>Metadata-Partition</i>
Primary partition for Control Center application data (mirrored)	<i>App-Data-Partition</i>
Primary partition for Control Center backups (not mirrored)	<i>Backups-Partition</i>

Upgrading Control Center on the master nodes

This section describes how to upgrade Control Center on the master nodes in a Control Center deployment. The upgrade process includes the following general steps:

- 1 Stop the application or applications that Control Center is managing, and then stop Control Center on the master node.
- 2 Save copies of Resource Manager images, if necessary.
- 3 Stop Control Center on all hosts.
- 4 Upgrade Docker on the master nodes, and the operating system, if necessary.
- 5 Upgrade Control Center on the master nodes.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before proceeding. For more information, see [Using Zenoss Toolbox](#) on page 163.

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Stopping Control Center

This procedure stops Control Center on the master host, and on resource pool hosts.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Display the public hostname of the current node.

```
uname -n
```

Make a note of which node (primary or secondary) is the current node, for use in a subsequent procedure.

- 3 Stop Control Center with the cluster management tool.

```
pcs cluster standby --all
```

- 4 Monitor the status of cluster resources.

```
watch pcs status
```

Monitor the status until all resources report `Stopped`. Resolve any issues before continuing.

- 5 Stop Control Center on all resource pool hosts.

Repeat this step on each resource pool host in your deployment.

- a Log in to the resource pool host as `root`, or as a user with superuser privileges.
- b Stop Control Center.

```
systemctl stop serviced
```

Upgrading Docker on the master nodes

This procedure upgrades Docker from version 1.5 or 1.8.2 to 1.9.0.

Note This release of Control Center requires RHEL/CentOS 7.1 or 7.2, and this procedure includes steps for upgrading the operating system, if necessary or desired. Zenoss recommends upgrading the operating system only at the step specified in this procedure.

Perform this procedure on the primary node and on the secondary node.

- 1 Log in to the master node as `root`, or as a user with superuser privileges.
- 2 Determine which version of Docker is installed.

```
rpm -qa | grep docker
```

- 3 **Note** Perform this step **only** if `zenoss-docker-1.5` is installed.

Remove `zenoss-docker`, and enable installation of the new version of Docker.

- a Remove Docker, without removing Control Center.

```
rpm -e --nodeps zenoss-docker-1.5.0-2
```

- b Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- 4 **Note** Perform this step **only** if `docker-engine-1.8` is installed.

Remove Docker, and then verify that the Docker repository is enabled.

- a Remove Docker, without removing Control Center.

```
rpm -e --nodeps docker-engine-1.8.2
```

- b Check the Docker repository.

```
grep enabled /etc/yum.repos.d/docker.repo
```

If the result is `enabled=0`, perform the following substeps.

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.
 d Change the value of the `enabled` key from 0 to 1.
 e Save the file and close the text editor.
- 5 Install Docker 1.9.0, and then disable accidental upgrades.

- a Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- b Install Docker 1.9.0.

```
yum clean all && yum makecache fast
yum install -y docker-engine-1.9.0
```

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.
 d Change the value of the `enabled` key from 1 to 0.
 e Save the file and close the text editor.
- 6 Create a symbolic link for the Docker temporary directory.

Docker uses its temporary directory to spool images. The default directory is `/var/lib/docker/tmp`. The following command specifies the same directory that Control Center uses, `/tmp`. You can specify any directory that has a minimum of 10GB of unused space.

- a Create the `docker` directory in `/var/lib`.

```
mkdir /var/lib/docker
```

- b Create the link to `/tmp`.

```
ln -s /tmp /var/lib/docker/tmp
```

- 7 Create a `systemd` override file for the Docker service definition.

- a Create the override directory.

```
mkdir -p /etc/systemd/system/docker.service.d
```

- b Create the override file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
```

```
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/docker daemon \${OPTIONS} -H fd://
EOF
```

- c Reload the systemd manager configuration.

```
systemctl daemon-reload
```

- 8 Install Control Center.

Control Center includes a utility that simplifies the process of creating a device mapper thin pool.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 9 Create a device mapper thin pool for Docker data.

- a Identify the primary partition for the thin pool to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b Create the thin pool.

Replace *Path-To-Device* with the path of an unused primary partition:

```
serviced-storage create-thin-pool docker Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 10 Configure and start the Docker service.

- a Create variables for adding arguments to the Docker configuration file.

The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

Replace *Thin-Pool-Device* with the name of the thin pool device created in the previous step:

```
myDriver="-s devicemapper"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myFlag="--storage-opt dm.thinpooldev"
myPool="Thin-Pool-Device"
```

- b Add the arguments to the Docker configuration file.

```
echo 'OPTIONS="'$myDriver $myFix $myFlag'='$myPool''' \
>> /etc/sysconfig/docker
```

- c Start or restart Docker.

```
systemctl restart docker
```

The initial startup takes up to a minute, and may fail. If the startup fails, repeat the previous command.

- 11 Configure name resolution in containers.

Each time it starts, `docker` selects an IPv4 subnet for its virtual Ethernet bridge. The selection can change; this step ensures consistency.

- a Identify the IPv4 subnet and netmask docker has selected for its virtual Ethernet bridge.

```
ip addr show docker0 | grep inet
```

- b Open `/etc/sysconfig/docker` in a text editor.
- c Add the following flags to the end of the `OPTIONS` declaration.

Replace `Bridge-Subnet` with the IPv4 subnet docker selected for its virtual bridge, and replace `Bridge-Netmask` with the netmask docker selected:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/Bridge-Netmask
```

For example, if the bridge subnet and netmask is 172.17.0.1/16, the flags to add are `--dns=172.17.0.1`
`--bip=172.17.0.1/16`.

Note Leave a blank space after the end of the thin pool device name, and make sure the double quote character (") is at the end of the line.

- d Restart the Docker service.

```
systemctl restart docker
```

Upgrading Control Center

This procedure pulls new images for Control Center internal services and updates the configuration files on the master nodes.

Perform this procedure on the primary node and on the secondary node.

- 1 Log in to the master node as `root`, or as a user with superuser privileges.
- 2 Manage the Control Center configuration files.
 - a Change directory to the configuration file directory.

```
cd /etc/default
```

- b Rename the existing configuration file, as a backup.

```
mv serviced serviced.pre-1.1.4
```

- c Rename the new configuration file, and then create a copy of it.

```
mv serviced.rpmnew serviced.orig-1.1.4
cp serviced.orig-1.1.4 serviced
```

- 3 Copy settings from the previous Control Center configuration file to the new one.
 - a Identify the customized variables in the pre-upgrade configuration file.

```
egrep '^[^#]*SERVICED' serviced.pre-1.1.4
```

- b Open `/etc/default/serviced` with a text editor, and then customize the same variables that were customized in the pre-upgrade configuration file.

The following variables are deprecated, and are not needed in the new configuration file:

- `SERVICED_REGISTRY`
- `SERVICED_VARPATH`

- c Add `SERVICED_DOCKER_REGISTRY` to the file.

The variable specifies the host and port at which the local Docker registry is available.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
SERVICED_DOCKER_REGISTRY=Hostname-Or-IP:5000
```

- d Save the file and close the text editor.
- 4 Pull the required Control Center images from Docker Hub.
 - a Identify the images to pull.

```
serviced version | grep IsvcsImages
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v40 zenoss/isvcs-zookeeper:v3]
```

- b Pull Control Center images.

Replace *Isvcs-Image-Name* with one of the images named in the output of the previous substep:

```
docker pull Isvcs-Image-Name
```

Repeat the command for each required image.

- 5 Install the Pacemaker resource agent for Control Center.

Pacemaker uses resource agents (scripts) to implement a standardized interface for managing arbitrary resources in a cluster. Zenoss provides a Pacemaker resource agent to manage the Control Center master host in a high-availability cluster.

```
yum --enablerepo=zenoss-stable install -y serviced-resource-agents
```

- 6 Disable the Docker and Control Center services.

The cluster management software controls the services.

- a Stop and disable the Docker service.

```
systemctl stop docker && systemctl disable docker
```

- b Disable the Control Center service.

```
systemctl disable serviced
```

Updating high-availability components

The procedures in this section update the DRBD and Pacemaker/Corosync components with new resources.

Updating the DRBD configuration

This procedure updates DRBD with volumes for Control Center metadata and application data, which are new with this release of Control Center.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.

- 3 On both nodes, identify the primary partitions to use for the new volumes.

```
lsblk --output=NAME,SIZE
```

Record the paths of the primary partitions to use in the following table. The information is needed in subsequent steps and procedures.

Node	<i>Metadata-Partition</i>	<i>App-Data-Partition</i>

- 4 On both nodes, remove any file system signature from the partitions.

Replace *Metadata-Partition* with the path of the primary partition designated for Control Center metadata, and replace *App-Data-Partition* with the path of the primary partition designated for Control Center application data:

```
wipefs -a Metadata-Partition
wipefs -a App-Data-Partition
```

- 5 On both nodes, add entries for the Control Center metadata and application data volumes.

- a Open `/etc/drbd.d/serviced-dfs.res` with a text editor.
- b Add entries for the volumes.

Replace *Metadata-Partition* with the path of the primary partition designated for Control Center metadata, and replace *App-Data-Partition* with the path of the primary partition designated for Control Center application data:

```
volume 2 {
    device /dev/drbd2;
    disk Metadata-Partition;
    meta-disk internal;
}
volume 3 {
    device /dev/drbd3;
    disk App-Data-Partition;
    meta-disk internal;
}
```

- c Save the file, and then close the editor.
- 6 On both nodes, update the DRBD configuration, and start DRBD.
 - a Display the commands for updating volumes individually.

```
drbdadm create-md --dry-run serviced-dfs
```

Example result:

```
drbdmeta 0 v08 /dev/sdb1 internal create-md
drbdmeta 1 v08 /dev/sdb2 internal create-md
drbdmeta 2 v08 /dev/sdb3 internal create-md
drbdmeta 3 v08 /dev/sdb4 internal create-md
```

In this example, the paths of the primary partitions for the new volumes are `/dev/sdb3` and `/dev/sdb4`.

- b Initialize the new volumes.

Note Do not initialize existing volumes.

Replace *Metadata-Partition* with the path of the primary partition designated for Control Center metadata, and replace *App-Data-Partition* with the path of the primary partition designated for Control Center application data:

```
drbdmeta 2 v08 Metadata-Partition internal create-md
drbdmeta 3 v08 App-Data-Partition internal create-md
```

- c Start DRBD.

```
drbdadm up serviced-dfs && drbdadm adjust all
```

- 7 On the node that was the current node when you stopped Control Center, start synchronization, and then monitor it.

- a Start synchronization.

```
drbdadm primary --force serviced-dfs
```

- b Monitor the process.

```
watch drbd-overview
```

Do not proceed until the status is UpToDate/UpToDate, as in the following example output:

```
0:serviced-dfs/0 Connected Primary/Secondary UpToDate/UpToDate
1:serviced-dfs/1 Connected Primary/Secondary UpToDate/UpToDate
2:serviced-dfs/2 Connected Primary/Secondary UpToDate/UpToDate
3:serviced-dfs/3 Connected Primary/Secondary UpToDate/UpToDate
```

The Primary/Secondary values show that the command was run on the primary node; otherwise, the values are Secondary/Primary. Likewise, the first value in the UpToDate/UpToDate field is the status of the node on which the command is run, and the second value is the status of the remote node.

- c Format the partition for Control Center metadata.

The following command uses the path of the DRBD devices defined previously, not the path of the primary partition.

```
mkfs.xfs /dev/drbd2
```

The command creates an XFS file system on the current node, and DRBD mirrors the file system to the remote node.

Configuring Logical Volume Manager

Control Center application data is managed by a device mapper thin pool created with Logical Volume Manager (LVM). This procedure adjusts the LVM configuration for mirroring by DRBD.

Perform this procedure on both master nodes.

- 1 Log in to the node as `root`, or as a user with superuser privileges.
- 2 Edit the LVM configuration file.
 - a Open `/etc/lvm/lvm.conf` with a text editor.
 - b Locate the filter directive.

The line is in the `devices` section, and is commented out.

```
# filter = [ "a|.*|/" ]
```

- c** Add a filter directive to exclude the partition for Control Center application data.

Place the directive immediately below the commented filter directive.

Replace *App-Data-Partition* with the absolute path of the primary partition designated for Control Center application data.

```
filter = ["r|App-Data-Partition|"]
```

- d** Disable caching and the metadata daemon.

Set the values of the `write_cache_state` and `use_lvmetad` directives to 0.

```
write_cache_state = 0
use_lvmetad = 0
```

- e** Save the file and close the editor.

- 3** Verify the new directives.

```
egrep '^[^#]*(filter|write_cache|use_lvm)' /etc/lvm/lvm.conf
```

The output should match the changes made in the previous step.

- 4** Delete any stale cache entries.

```
rm -f /etc/lvm/cache/.cache
```

- 5** Restart the host.

```
reboot
```

Creating a thin pool for application data

This procedure creates and mirrors a device mapper thin pool for Control Center application data.

- 1** Log in to the primary node as `root`, or as a user with superuser privileges.

For this procedure, the primary node is the node that was the current node when you stopped Control Center.

- 2** In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
3 On both nodes, start DRBD.

```
drbdadm up all
```

- 4** On the primary node, start synchronization.

```
drbdadm primary --force serviced-dfs
```

- 5** On the primary node, create a thin pool for application data.

- a** Identify the amount of space available on the primary partition for application data.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b** Create a variable for 50% of the space available on the thin pool partition.

The thin pool stores application data and snapshots of the data. You can add storage to the pool at any time.

Replace *Half-Of-Available-Space* with 50% of the space available in the primary partition, in gigabytes. Include the symbol for gigabytes (G) after the numeric value.

```
myFifty=Half-Of-Available-SpaceG
```

- c Create the thin pool.

The following command uses the path of the DRBD device defined previously, not the path of the primary partition for application data.

```
serviced-storage create-thin-pool -o dm.basesize=$myFifty \  
serviced /dev/drbd3 -v
```

On success, the `serviced-storage` command displays the name of the thin pool it creates. Make a note of the name, which is needed in a subsequent procedure.

- 6 Stop DRBD.

- a On the primary node, enter the following commands:

```
vgchange -a n serviced && drbdadm down all
```

- b On the secondary node, enter the following command:

```
drbdadm down all
```

Updating the cluster configuration

This procedure adds new resources to the Control Center cluster configuration.

- 1 Log in to a master node as `root`, or as a user with superuser privileges.
For `pcs` commands, the nodes are equivalent.
- 2 Start the cluster.

```
pcs cluster start
```

- 3 Define the LVM resource.

```
pcs resource create serviced-lvm ocf:heartbeat:LVM \  
volgrpname=serviced
```

- 4 Define the storage resource for Control Center application data.

This ensures the the thin pool is deactivated and unmounted correctly.

```
pcs resource create serviced-storage \  
ocf:zenoss:serviced-storage
```

- 5 Add the new resources to the Control Center resource group.

```
pcs resource group add serviced-group \  
serviced-lvm --after serviced-volumes  
pcs resource group add serviced-group \  
serviced-storage --after serviced-lvm
```

- 6 Verify Control Center resource agent timeouts.

With this release, the recommended resource agent timeouts for starting and stopping Control Center have changed to 360 and 130 seconds, respectively.

- a Display the current values.

```
pcs resource show serviced
```

- b Update the values, if necessary.

```
pcs resource update serviced op start timeout=360s
pcs resource update serviced op stop timeout=130s
```

- 7 Disable Pacemaker monitoring of NFS health.

During normal operations, Control Center occasionally stops and restarts NFS, which could be misinterpreted by Pacemaker and trigger an unwanted failover.

```
pcs resource op remove nfs monitor interval=60s
pcs resource op add nfs monitor interval=0s
```

- 8 Verify the cluster configuration.

```
pcs resource show --groups
```

The correct result shows serviced-group with 8 members, in the following order:

```
serviced-group: serviced-isvcs serviced-volumes serviced-lvm serviced-
storage VirtualIP docker nfs serviced
```

Verifying cluster startup

This procedure verifies the new configuration by attempting to start the resources on one node only. With the other node in standby mode, Pacemaker does not automatically fail over to the other node.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
For this procedure, the primary node is the node that was the current node when you stopped Control Center.
- 2 Determine which node is the primary DRBD node.

```
pcs status
```

Example result:

```
Cluster name: serviced-ha
Last updated: Mon Feb 22 11:37:58 2016 Last change: Mon Feb 22
11:35:19 2016 by root via crm_attribute on Secondary-Public-Name
Stack: corosync
Current DC: Primary-Public-Name (version 1.1.13-a14efad) - partition
with quorum
2 nodes and 10 resources configured

Node Primary-Public-Name: standby
Node Secondary-Public-Name: standby

Full list of resources:

Master/Slave Set: DFSMaster [DFS]
Stopped: [ Primary-Public-Name Secondary-Public-Name ]
Resource Group: serviced-group
  serviced-isvcs (ocf::heartbeat:Filesystem): Stopped
  serviced-volumes (ocf::heartbeat:Filesystem): Stopped
  serviced-lvm (ocf::heartbeat:LVM): Stopped
```

```

serviced-storage (ocf::zenoss:serviced-storage): Stopped
VirtualIP (ocf::heartbeat:IPaddr2): Stopped
docker (systemd:docker): Stopped
nfs (systemd:nfs): Stopped
serviced (ocf::zenoss:serviced): Stopped

```

PCSD Status:

```

Primary-Public-Name: Offline
Secondary-Public-Name: Offline

```

Daemon Status:

```

corosync: active/disabled
pacemaker: active/enabled
pcsd: active/enabled

```

The line that begins with `Current DC` identifies the primary node. Review all of the command output for errors.

3 Start DRBD.

- a** On the secondary node, enter the following command:

```
drbdadm up all
```

- b** On the primary node, enter the following commands:

```
drbdadm up all && drbdadm primary serviced-dfs
```

4 Take the current node out of standby mode.

- a** Display the public hostname of the current node.

```
uname -n
```

- b** Clean up before bringing the node out of standby mode.

```
pcs resource cleanup --all
```

- c** Put the current node back online.

Replace *Primary-Public-Name* with the hostname of the current node:

```
pcs cluster unstandby Primary-Public-Name
```

5 Monitor the status of cluster resources.

```
watch pcs status
```

Monitor the status until all resources report `Started`. Resolve any issues before continuing.

Verifying cluster failover

This procedure simulates a failover.

- 1** Log in to the primary node as `root`, or as a user with superuser privileges.
- 2** Enable the secondary node.
 - a** Take the secondary node out of standby mode.

Replace *Secondary-Public-Name* with the public hostname of the secondary node:

```
pcs cluster unstandby Secondary-Public-Name
```

- b** Monitor the status of the secondary node.

```
watch pcs status
```

Do not continue until the status of the secondary node is Online.

- 3** Verify that DRBD has completely synchronized all four volumes on the secondary node.

```
drbd-overview
```

Example result:

```
0:serviced-dfs/0 Connected Primary/Secondary UpToDate/UpToDate
1:serviced-dfs/1 Connected Primary/Secondary UpToDate/UpToDate
2:serviced-dfs/2 Connected Primary/Secondary UpToDate/UpToDate
3:serviced-dfs/3 Connected Primary/Secondary UpToDate/UpToDate
```

- 4** Force a failover.

Pacemaker initiates a failover when the primary node is put in standby mode.

Replace *Primary-Public-Name* with the public hostname of the primary node:

```
pcs cluster standby Primary-Public-Name
```

- 5** Monitor the cluster status.

```
watch pcs status
```

Wait until all resources report a status of Started. Resolve any issues before continuing.

- 6** Restore cluster services.

This starts the secondary node and Control Center internal services, but does not start Resource Manager services, which is intentional. Resource pool hosts have to be upgraded first.

```
pcs cluster unstandby --all
```

Upgrading Control Center on resource pool hosts

Perform the procedures in this section on each resource pool host in your Control Center deployment.

Upgrading Docker

This procedure upgrades Docker from version 1.5 or 1.8.2 to 1.9.0.

Note This release of Control Center requires RHEL/CentOS 7.1 or 7.2, and this procedure includes steps for upgrading the operating system, if necessary or desired. Zenoss recommends upgrading the operating system only at the step specified in this procedure.

- 1** Log in to the resource pool host as `root`, or as a user with superuser privileges.

- 2 Stop Control Center.

```
systemctl stop serviced
```

- 3 Determine which version of Docker is installed.

```
rpm -qa | grep docker
```

- 4 **Note** Perform this step **only** if `zenoss-docker-1.5` is installed.

Remove `zenoss-docker`, and enable installation of the new version of Docker.

- a Remove Docker, without removing Control Center.

```
rpm -e --nodeps zenoss-docker-1.5.0-2
```

- b Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- 5 **Note** Perform this step **only** if `docker-engine-1.8` is installed.

Remove Docker, and then verify that the Docker repository is enabled.

- a Remove Docker, without removing Control Center.

```
rpm -e --nodeps docker-engine-1.8.2
```

- b Check the Docker repository.

```
grep enabled /etc/yum.repos.d/docker.repo
```

If the result is `enabled=0`, perform the following substeps.

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.
- d Change the value of the `enabled` key from 0 to 1.
- e Save the file and close the text editor.

- 6 Remove the Docker data partition.

- a Identify the partition where `/var/lib/docker` is mounted.

```
mount | awk '/\var\/lib\/docker/ { print $1 }'
```

- b Unmount the partition.

Replace *Partition* with the device returned in the previous substep:

```
umount Partition
```

- c Erase the XFS file system on the partition.

Replace *Partition* with the device returned previously:

```
wipefs -a Partition
```

The partition is now ready for use as a device mapper thin pool.

- d Open `/etc/fstab` with a text editor.
 - e Remove the entry for `/var/lib/docker`.
- 7 Upgrade the operating system, if necessary.
- a Determine which release is installed.

```
cat /etc/redhat-release
```

If the result includes 7.0, perform the following substeps.

- b Disable the `serviced` service.

```
systemctl disable serviced
```

- c Upgrade the operating system.

```
yum clean all && yum update -y
```

- d Restart the operating system.

```
reboot
```

- e Log in to the Control Center cluster host as `root`, or as a user with superuser privileges.
- f Enable the `serviced` service.

```
systemctl enable serviced
```

- 8 Install Docker 1.9.0, and then disable accidental upgrades.
- a Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- b Install Docker 1.9.0.

```
yum clean all && yum makecache fast
yum install -y docker-engine-1.9.0
```

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.
 - d Change the value of the `enabled` key from 1 to 0.
 - e Save the file and close the text editor.
- 9 Create a `systemd` override file for the Docker service definition.
- a Create the override directory.

```
mkdir -p /etc/systemd/system/docker.service.d
```

- b** Create the override file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/docker daemon \${OPTIONS} -H fd://
EOF
```

- c** Reload the systemd manager configuration.

```
systemctl daemon-reload
```

- 10** Configure the docker service to start when the system starts.

```
systemctl enable docker
```

- 11** Install Control Center.

Control Center includes a utility that simplifies the process of creating a device mapper thin pool.

```
yum --enablerepo=zenoss-stable install -y serviced-1.1.4
```

The installation preserves the existing version of the serviced configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 12** Create a device mapper thin pool for Docker data.

- a** Identify the primary partition for the thin pool to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b** Create the thin pool.

Replace *Path-To-Device* with the path of an unused primary partition:

```
serviced-storage create-thin-pool docker Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 13** Configure and start the Docker service.

- a** Create variables for adding arguments to the Docker configuration file.

The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

Replace *Thin-Pool-Device* with the name of the thin pool device created in the previous step:

```
myDriver="-s devicemapper"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myFlag="--storage-opt dm.thinpooldev"
myPool="Thin-Pool-Device"
```

- b** Add the arguments to the Docker configuration file.

```
echo 'OPTIONS="'$myDriver $myFix $myFlag'='$myPool'' \
>> /etc/sysconfig/docker
```

- c** Start or restart Docker.

```
systemctl restart docker
```

The initial startup takes up to a minute, and may fail. If the startup fails, repeat the previous command.

14 Configure name resolution in containers.

Each time it starts, `docker` selects an IPv4 subnet for its virtual Ethernet bridge. The selection can change; this step ensures consistency.

- a Identify the IPv4 subnet and netmask `docker` has selected for its virtual Ethernet bridge.

```
ip addr show docker0 | grep inet
```

- b Open `/etc/sysconfig/docker` in a text editor.
- c Add the following flags to the end of the `OPTIONS` declaration.

Replace *Bridge-Subnet* with the IPv4 subnet `docker` selected for its virtual bridge, and replace *Bridge-Netmask* with the netmask `docker` selected:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/Bridge-Netmask
```

For example, if the bridge subnet and netmask is 172.17.0.1/16, the flags to add are `--dns=172.17.0.1 --bip=172.17.0.1/16`.

Note Leave a blank space after the end of the thin pool device name, and make sure the double quote character (") is at the end of the line.

- d Restart the Docker service.

```
systemctl restart docker
```

Upgrading Control Center

This procedure updates the configuration files on resource pool hosts.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Manage the Control Center configuration files.
 - a Change directory to the configuration file directory.

```
cd /etc/default
```

- b Rename the existing configuration file, as a backup.

```
mv serviced serviced.pre-1.1.4
```

- c Rename the new configuration file, and then create a copy of it.

```
mv serviced.rpmnew serviced.orig-1.1.4
cp serviced.orig-1.1.4 serviced
```

- 3 Copy settings from the previous Control Center configuration file to the new one.
 - a Identify the customized variables in the pre-upgrade configuration file.

```
egrep '^[^#]*SERVICED' serviced.pre-1.1.4
```

- b Open `/etc/default/serviced` with a text editor, and then customize the same variables that were customized in the pre-upgrade configuration file.

The following variables are deprecated, and are not needed in the new configuration file:

- `SERVICED_REGISTRY`

- `SERVICED_VARPATH`
- c Add `SERVICED_DOCKER_REGISTRY` to the file.
The variable specifies the host and port at which the local Docker registry is available.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
SERVICED_DOCKER_REGISTRY=Hostname-Or-IP:5000
```

- d Save the file and close the text editor.
- 4 Unmount the distributed file system (DFS).
 - a Identify the file system specification to unmount.

```
mount | awk '/serviced/ { print $1 }'
```

- b Unmount the DFS.
Replace *DFS-Mount* with the file system specification returned in the previous substep:

```
umount DFS-Mount
```

- 5 Start Control Center.

```
systemctl start serviced
```

Converting the data storage driver

At this point in the upgrade, Control Center is upgraded on the master nodes and on all of the resource pool hosts. You may continue using Resource Manager to monitor your infrastructure during the initial (and longest) part of the conversion process. However, Zenoss strongly recommends running Resource Manager in this configuration only as long as it takes to migrate Btrfs data to the new thin pool.

The procedures in this section migrate Resource Manager data from a Btrfs volume to a device mapper thin pool volume.

Starting the conversion

To perform this procedure, you need an unused primary partition that is large enough for your data. For more information, see [Preparing to convert the data storage driver](#) on page 47.

This procedure converts the majority of Resource Manager data stored in a Btrfs file system and copies it to a device mapper thin pool. This procedure may be performed while Resource Manager is monitoring your infrastructure. Control Center application data is moved to its new thin pool volume, and Control Center metadata is moved to its new XFS volume.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Verify that cluster services are running.

```
watch pcs status
```

Do not proceed until both master nodes are online and all resources are started.

- 3 Initialize `serviced-storage` for the `btrfs` storage driver.

```
serviced-storage init /opt/serviced/var/volumes btrfs
```

- 4 Mount the DRBD volume for Control Center metadata at a temporary location.
The device path is `/dev/drbd2`.

- a Create a mount point for the temporary location.

Replace *Temporary-Directory* with the path of a temporary directory:

```
mkdir -p Temporary-Directory
```

- b Mount the DRBD volume.

```
mount /dev/drbd2 Temporary-Directory
```

- 5 Synchronize the contents of the Btrfs volume with the new storage.

This step copies current data to the new thin pool, and current metadata to the new XFS volume.

Replace *Thin-Pool-Name* with the name of the thin pool created in a previous procedure, and replace *Temporary-Directory* with the path of the temporary location created in the previous step:

```
serviced-storage sync -c -t devicemapper \
-o dm.thinpooldev=Thin-Pool-Name \
/opt/serviced/var/volumes Temporary-Directory
```

After a pause to compute space requirements, the `serviced-storage` command displays both detail and summary information about its work. Depending on the amount of data to convert, this step may take several hours.

Note Since both the thin pool and the XFS file system are DRBD volumes, the data copied in this step is automatically mirrored to the secondary node.

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.

Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.

- a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Finalizing the conversion

This procedure completes the conversion started previously.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Stop Control Center on each resource pool host.
 - a Log in to each resource pool host as `root`, or as a user with superuser privileges.
 - b Stop Control Center

```
systemctl stop serviced
```

- 3 On the primary node, stop the Control Center service.
 - a Stop the resource.

```
pcs resource disable serviced
```

- b Monitor the shutdown.

```
watch pcs status
```

Do not proceed until the `serviced` resource status is `Stopped`.

- 4 On the primary node, synchronize the contents of the Btrfs volume with the new storage. This step copies data created since the previous sync operation.

Repeat the command you used before stopping the Control Center service:

```
serviced-storage sync -c -t devicemapper \  
-o dm.thinpooldev=Thin-Pool-Name \  
/opt/serviced/var/volumes Temporary-Directory
```

- 5 On the primary node, disable the thin pool.

```
serviced-storage disable Temporary-Directory \  
-o dm.thinpooldev=Thin-Pool-Name
```

- 6 On the primary node, unmount the DRBD volume for Control Center metadata.

```
umount Temporary-Directory
```

- 7 On the primary node, stop all cluster services.
 - a Stop all resources.

```
pcs cluster standby --all
```

- b Monitor the shutdown.

```
watch pcs status
```

Do not proceed until the status of all resources is `Stopped`.

- 8 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 9 On both nodes, update the Control Center configuration file with new storage settings.
 - a Open `/etc/default/serviced` with a text editor.
 - b Locate the `SERVICED_FS_TYPE` declaration, and change the value from `btrfs` to `devicemapper`.
 - c Add `SERVICED_DM_THINPOOLDEV` to the file, immediately after `SERVICED_FS_TYPE`.

The variable specifies the name of the thin pool device for Control Center application data.

Replace *Thin-Pool-Name* with the name of the thin pool created used to synchronize data:

```
SERVICED_DM_THINPOOLDEV=Thin-Pool-Name
```

- d Save the file and close the text editor.
- 10 On both nodes, update the DRBD resource definition, and then update DRBD.
 - a Open `/etc/drbd.d/serviced-dfs.res` with a text editor.
 - b Remove the entry for the Btrfs file system.

The entry to remove should be for volume 1, which references `/dev/drbd1`.

In the following example, the value of *Btrfs-Partition* is unique in your environment:

```
volume 1 {
    device /dev/drbd1;
    disk Btrfs-Partition;
    meta-disk internal;
}
```

- c Save the file and close the text editor.
- d Update DRBD.

```
drbdadm up all && drbdadm adjust all && drbdadm down all
```

- 11 On the primary node, update the Pacemaker configuration.

This change adds the Control Center metadata volume, DRBD device `/dev/drbd2`.

```
pcs resource update serviced-volumes \
    device=/dev/drbd/by-res/serviced-dfs/2 \
    fstype=xfs
```

- 12 On the primary node, start all cluster services.

```
pcs resource enable serviced && pcs cluster unstandby --all
```

- 13 Monitor the startup.

Do not proceed until all resources are started, DRBD is finished synchronizing, and both master nodes are online.

- a On the primary node, open two separate shell sessions or windows.
- b In one session, monitor the cluster startup.

```
watch pcs status
```

- c In another session, monitor the DRBD synchronization startup.

```
watch drbd-overview
```

- 14 Start Control Center on each resource pool host.

- a Log in to each resource pool host as `root`, or as a user with superuser privileges.
- b Start Control Center

```
systemctl start serviced
```


Verifying the conversion

This procedure verifies that the new data storage works properly with Resource Manager version 1.0.x.

- 1 Log in to the Control Center browser interface.
- 2 Start Resource Manager and all related applications.
- 3 Perform database integrity checks.
For more information, see [Using Zenoss Toolbox](#) on page 163.
- 4 Create a backup.
Previous backups are incompatible with the new storage driver.

When you are satisfied, continue with the Resource Manager upgrade procedure.

ZooKeeper ensemble configuration

Control Center relies on [Apache ZooKeeper](#) to coordinate its services. The configuration steps in this section create a ZooKeeper ensemble of 3 nodes.

A ZooKeeper ensemble requires a minimum of 3 nodes, and 3 nodes is sufficient for most deployments. A 5-node configuration improves failover protection during maintenance windows. Ensembles larger than 5 nodes are not necessary. An odd number of nodes is recommended, and an even number of nodes is strongly discouraged.

Control Center variables for ZooKeeper

This tables in this section associates the ZooKeeper-related Control Center variables to set in `/etc/default/serviced` with the roles that hosts play in a Control Center cluster.

Table 9: Control Center master host

<i>SERVICED_ISVCS_ZOOKEEPER_ID</i>
The unique identifier of a ZooKeeper ensemble node. Value: 1
<i>SERVICED_ISVCS_ZOOKEEPER_QUORUM</i>
The ZooKeeper node ID, IP address, peer communications port, and leader communications port of each host in an ensemble. Each quorum definition must be unique, so the IP address of the "current" host is 0.0.0.0. Value: <i>ZooKeeper-ID@IP-Address:2888:3888, . . .</i>
<i>SERVICED_ZK</i>
The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it. Value: <i>IP-Address:2181, . . .</i>

Table 10: Control Center resource pool host and ZooKeeper ensemble node

<i>SERVICED_ISVCS_ZOOKEEPER_ID</i>
The unique identifier of a ZooKeeper ensemble node. Value: 2 or 3
<i>SERVICED_ISVCS_ZOOKEEPER_QUORUM</i>

The ZooKeeper node ID, IP address, peer communications port, and leader communications port of each host in an ensemble. Each quorum definition must be unique, so the IP address of the "current" host is 0.0.0.0.

Value: *ZooKeeper-ID@IP-Address:2888:3888,...*

SERVICED_ISVCS_START

The list of Control Center internal services to start and run on hosts other than the master host.

Value: *zookeeper*

SERVICED_ZK

The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it.

Value: *IP-Address:2181,...*

Table 11: Control Center resource pool host

SERVICED_ZK

The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it.

Value: *IP-Address:2181,...*

Configuring master nodes

This procedure configures both Control Center master nodes as members of the ZooKeeper ensemble.

Note For accuracy, this procedure constructs Control Center configuration variables in the shell and appends them to `/etc/default/serviced`. The last step is to move the variables from the end of the file to more appropriate locations.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 3 On both nodes, create a variable for each Control Center host to include in the ZooKeeper ensemble.

The variables are used in subsequent steps.

Note Define the variables identically on both the primary and the secondary nodes, and on each resource pool host.

Replace *HA-Virtual-IP* with the virtual IP address of the high-availability cluster, and replace *Pool-Host-A-IP* and *Pool-Host-B-IP* with the IP addresses of the Control Center resource pool hosts to include in the ensemble:

```
node1=HA-Virtual-IP
node2=Pool-Host-A-IP
node3=Pool-Host-B-IP
```

Note ZooKeeper requires IP addresses for ensemble configuration.

- 4 On both nodes, set the ZooKeeper node ID to 1.

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=1" >> /etc/default/serviced
```

- 5 On both nodes, specify the nodes in the ZooKeeper ensemble.
You may copy the following text and paste it in your console:

```
echo "SERVICED_ZK=${node1}:2181,${node2}:2181,${node3}:2181" \
>> /etc/default/serviced
```

- 6 On both nodes, specify the nodes in the ZooKeeper quorum.
ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address of the current node with 0.0.0.0.
You may copy the following of text and paste it in your console:

```
q1="1@0.0.0.0:2888:3888"
q2="2@${node2}:2888:3888"
q3="3@${node3}:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

- 7 On both nodes, clean up the Control Center configuration file.
- a Open `/etc/default/serviced` with a text editor.
 - b Navigate to the end of the file, and cut the line that contains the `SERVICED_ZK` variable declaration at that location.
The value of this declaration specifies 3 hosts.
 - c Locate the `SERVICED_ZK` variable near the beginning of the file, and then delete the line it is on.
The value of this declaration is just the master node.
 - d Paste the `SERVICED_ZK` variable declaration from the end of the file in the location of the just-deleted declaration.
 - e Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration at that location.
 - f Locate the `SERVICED_ISVCS_ZOOKEEPER_ID` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - g Paste the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration from the end of the file in the location of the just-deleted declaration.
 - h Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration at that location.
 - i Locate the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - j Paste the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration from the end of the file in the location of the just-deleted declaration.
 - k Save the file, and then close the text editor.
- 8 On both hosts, verify the ZooKeeper environment variables.

```
egrep '^[^#]*SERVICED' /etc/default/serviced | egrep '(_ZOO|_ZK)'
```

Configuring a resource pool host as a ZooKeeper node

To perform this procedure, you need a resource pool host with an XFS file system on a separate partition.

This procedure configures a ZooKeeper ensemble on a resource pool host. Repeat this procedure on each Control Center resource pool host to add to the ZooKeeper ensemble.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.

- 2 Create a variable for each Control Center host to include in the ZooKeeper ensemble.

Replace *HA-Virtual-IP* with the virtual IP address of the high-availability cluster, and replace *Pool-Host-A-IP* and *Pool-Host-B-IP* with the IP addresses of the Control Center resource pool hosts to include in the ensemble:

```
node1=HA-Virtual-IP
node2=Pool-Host-A-IP
node3=Pool-Host-B-IP
```

- 3 Set the ID of this node in the ZooKeeper ensemble.

For *Pool-Host-A-IP* (node2), use the following command:

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=2" >> /etc/default/serviced
```

For *Pool-Host-B-IP* (node3), use the following command:

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=3" >> /etc/default/serviced
```

- 4 Specify the nodes in the ZooKeeper ensemble.
You may copy the following text and paste it in your console:

```
echo "SERVICED_ZK=${node1}:2181,${node2}:2181,${node3}:2181" \
>> /etc/default/serviced
```

- 5 Specify the nodes in the ZooKeeper quorum.

ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address of the current node with 0.0.0.0.

For *Pool-Host-A-IP* (node2), use the following commands:

```
q1="1@${node1}:2888:3888"
q2="2@0.0.0.0:2888:3888"
q3="3@${node3}:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

For *Pool-Host-B-IP* (node3), use the following commands:

```
q1="1@${node1}:2888:3888"
q2="2@${node2}:2888:3888"
q3="3@0.0.0.0:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

- 6 Set the *SERVICED_ISVCS_START* variable, and clean up the Control Center configuration file.
 - a Open */etc/default/serviced* with a text editor.
 - b Locate the *SERVICED_ISVCS_START* variable, and then delete all but *zookeeper* from its list of values.
 - c Remove the number sign character (#) from the beginning of the line.
 - d Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK* variable declaration at that location.
The value of this declaration specifies 3 hosts.
 - e Locate the *SERVICED_ZK* variable near the beginning of the file, and then delete the line it is on.
The value of this declaration is just the master node.

- f** Paste the `SERVICED_ZK` variable declaration from the end of the file in the location of the just-deleted declaration.
 - g** Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration at that location.
 - h** Locate the `SERVICED_ISVCS_ZOOKEEPER_ID` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - i** Paste the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration from the end of the file in the location of the just-deleted declaration.
 - j** Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration at that location.
 - k** Locate the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - l** Paste the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration from the end of the file in the location of the just-deleted declaration.
 - m** Save the file, and then close the text editor.
- 7** Verify the ZooKeeper environment variables.

```
egrep '^[^#]*SERVICED' /etc/default/serviced \
| egrep '(_ZOO|_ZK|_STA)'
```

- 8** Pull the required Control Center ZooKeeper image from the master host.
- a** Identify the image to pull.

```
serviced version | grep IsvcsImages
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v40 zenoss/isvcs-zookeeper:v3]
```

- b** Pull the Control Center ZooKeeper image.

Replace `Isvcs-ZK-Image` with the name and version number of the ZooKeeper image from the previous substep:

```
docker pull Isvcs-ZK-Image
```

Starting a ZooKeeper ensemble

This procedure starts a ZooKeeper ensemble.

The window of time for starting a ZooKeeper ensemble is relatively short. The goal of this procedure is to restart Control Center on each ensemble node at about the same time, so that each node can participate in electing the leader.

- 1** Use the virtual hostname (*HA-Virtual-Name*) or virtual IP address (*HA-Virtual-IP*) of the high-availability cluster to start a Bash shell on the Control Center master host as `root`, or as a user with superuser privileges.
- 2** Display the public hostname of the current node.

```
uname -n
```

The result is either *Primary-Public-Name* or *Secondary-Public-Name*.

- 3** Place the other node in standby mode.

This avoids potential conflicts and errors in the event of an unexpected `serviced` shutdown during the ZooKeeper startup.

Replace `Other-Node-Hostname` with the public hostname of the other master node:

```
pcs cluster standby Other-Node-Hostname
```

- 4 In a separate window, log in to the second node of the ZooKeeper ensemble (`Pool-Host-A-IP`).
- 5 In another separate window, log in to the third node of the ZooKeeper ensemble (`Pool-Host-B-IP`).
- 6 On the primary node, stop and start `serviced`.

```
pcs resource disable serviced && pcs resource enable serviced
```

- 7 On both resource pool hosts, stop and start `serviced`.

```
systemctl stop serviced && systemctl start serviced
```

- 8 On the primary node, check the status of the ZooKeeper ensemble.

```
{ echo stats; sleep 1; } | nc localhost 2181 | grep Mode
{ echo stats; sleep 1; } | nc Pool-Host-A-IP 2181 | grep Mode
{ echo stats; sleep 1; } | nc Pool-Host-B-IP 2181 | grep Mode
```

If `nc` is not available, you can use `telnet` with [interactive ZooKeeper commands](#).

- 9 Restore the cluster.

Replace `Other-Node-Hostname` with the public hostname of the other master node:

```
pcs cluster unstandby Other-Node-Hostname
```

Updating resource pool hosts

The default configuration of resource pool hosts sets the value of the `SERVICED_ZK` variable to the master host only. This procedure updates the setting to include the full ZooKeeper ensemble.

Perform this procedure on each resource pool host in your Control Center cluster.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Update the variable.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the `SERVICED_ZK` declaration, and then replace its value with the same value used in the ZooKeeper ensemble nodes.
 - c Save the file, and then close the editor.
- 3 Restart Control Center.

```
systemctl restart serviced
```

Upgrading Resource Manager

This procedure upgrades Resource Manager.

Before performing this procedure, stop Resource Manager. For more information, see [Stopping applications](#) on page 19.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.

- 2 Download the primary Docker image of Resource Manager for this release

The download takes approximately 10-20 minutes.

```
docker run -it --rm -v /root:/mnt/root \
  zenoss/resmgr_5.1:5.1.3_1 rsync -a /root/5.1.x /mnt/root
```

When the download completes, the `rsync` command copies scripts that perform the upgrade to `/root/5.1.x`.

- 3 Pull additional images for Resource Manager from Docker Hub.

```
/root/5.1.x/pull-docker-images.sh
```

- 4 Start the upgrade script.

The script to start depends on whether Service Impact is installed.

- If Service Impact is installed, enter the following command.

```
/root/5.1.x/upgrade-impact-5.1.x.sh
```

Note The script upgrades Resource Manager, but does not upgrade Service Impact. For more information about upgrading Service Impact, refer to the *Zenoss Service Impact Installation Guide for Resource Manager 5.x*.

- If Service Impact is not installed, enter the following command.

```
/root/5.1.x/upgrade-resmgr-5.1.x.sh
```

- 5 Restart Resource Manager.

Some Resource Manager services are started during the upgrade, and they need to be restarted.

```
serviced service restart Zenoss.resmgr
```

Proceed to [After upgrading](#) on page 158.

Upgrading high-availability deployments without internet access

5

This chapter includes detailed procedures for upgrading high-availability hosts that do not have internet access. For hosts that do have internet access, or that are not configured for high-availability, use one of the other chapters in this part.

Key variables used in this chapter

The following tables associate important features of a high-availability deployment with the variables used in this chapter.

Feature	Variable Name	
	Primary Node	Secondary Node
Public IP address of master node (static; known to all machines in the Control Center cluster)	<i>Primary-Public-IP</i>	<i>Secondary-Public-IP</i>
Public hostname of master node (returned by <code>uname -n</code> ; resolves to the public IP address)	<i>Primary-Public-Name</i>	<i>Secondary-Public-Name</i>
Private IP address of master node (static; dual-NIC systems only)	<i>Primary-Private-IP</i>	<i>Secondary-Private-IP</i>
Private hostname of master node (resolves to the private IP address; dual-NIC systems only)	<i>Primary-Private-Name</i>	<i>Secondary-Private-Name</i>

Feature	Variable Name
Virtual IP address of the high-availability cluster (static; known enterprise-wide)	<i>HA-Virtual-IP</i>
Virtual hostname of the high-availability cluster (known enterprise-wide)	<i>HA-Virtual-Name</i>
Public IP address of resource pool host A (static; for ZooKeeper ensemble)	<i>Pool-Host-A-IP</i>
Public IP address of resource pool host B (static; for ZooKeeper ensemble)	<i>Pool-Host-B-IP</i>
Primary partition for Btrfs file system (mirrored)	<i>Btrfs-Partition</i>
Primary partition for Docker data (not mirrored)	<i>Docker-Partition</i>

Feature	Variable Name
Primary partition for Control Center internal services data (mirrored)	<i>Isvcs-Partition</i>
Primary partition for Control Center metadata (mirrored)	<i>Metadata-Partition</i>
Primary partition for Control Center application data (mirrored)	<i>App-Data-Partition</i>
Primary partition for Control Center backups (not mirrored)	<i>Backups-Partition</i>

Downloading files for offline installation

This procedure describes how to download RPM packages and Docker image files to your workstation.

To perform this procedure, you need:

- A workstation with internet access.
 - A portable storage medium, such as a USB flash drive, with at least 5 GB of free space.
 - Permission to download the required files from the [File Portal - Download Zenoss Enterprise Software](#) site. You may request permission by filing a ticket at the [Zenoss Support](#) site.
- 1 In a web browser, navigate to the [File Portal - Download Zenoss Enterprise Software](#) site.
 - 2 Log in with the account provided by Zenoss Support.
 - 3 Download archive files to your workstation.

Replace *Version* with the most recent version number available on the download page:

- `install-registry:Version.run`
 - `install-zenoss-hbase:vVersion.run`
 - `install-zenoss-isvcs-zookeeper:vVersion.run`
 - `install-zenoss-opentsdb:vVersion.run`
 - `install-zenoss-resmgr_5.1:5.1Version.run`
 - `install-zenoss-serviced-isvcs:vVersion.run`
 - `serviced-resource-agents-Version.x86_64.rpm`
- 4 Download the RHEL/CentOS mirror package for your upgrade.

Note If you are planning to upgrade the operating system during your Control Center and Resource Manager upgrade, choose the mirror package that matches the RHEL/CentOS release to which you are upgrading, not the release that is installed now.

Replace *Version* with the most recent version number available on the download page, and replace *Release* with the version of RHEL/CentOS appropriate for your environment:

```
yum-mirror-centos7.Release-Version.x86_64.rpm
```

- 5 Copy the files to your portable storage medium.

Upgrading Control Center on the master nodes

This section describes how to upgrade Control Center on the master nodes in a Control Center deployment. The upgrade process includes the following general steps:

- 1 Stop the application or applications that Control Center is managing, and then stop Control Center on the master node.
- 2 Save copies of Resource Manager images, if necessary.
- 3 Stop Control Center on all hosts.

- 4 Upgrade Docker on the master nodes, and the operating system, if necessary.
- 5 Upgrade Control Center on the master nodes.

Note Zenoss strongly recommends checking the integrity of Resource Manager databases before proceeding. For more information, see [Using Zenoss Toolbox](#) on page 163.

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.
Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.
 - a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.
Repeat the following command until the `STATUS` column reads `Stopped`:

```
serviced service status Zenoss.resmgr
```

Stopping Control Center

This procedure stops Control Center on the master host, and on resource pool hosts.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Display the public hostname of the current node.

```
uname -n
```

Make a note of which node (primary or secondary) is the current node, for use in a subsequent procedure.

- 3 Stop Control Center with the cluster management tool.

```
pcs cluster standby --all
```

- 4 Monitor the status of cluster resources.

```
watch pcs status
```

Monitor the status until all resources report `Stopped`. Resolve any issues before continuing.

- 5 Stop Control Center on all resource pool hosts.
Repeat this step on each resource pool host in your deployment.

- a Log in to the resource pool host as `root`, or as a user with superuser privileges.
- b Stop Control Center.

```
systemctl stop serviced
```

Staging files for offline installation

Before performing this procedure, complete all of the steps in [Downloading files for offline installation](#) on page 74. In addition, verify that approximately 4GB of temporary space is available on the file system where `/root` is located.

This procedure adds files for offline installation to the master node. The staged files are required in subsequent procedures.

Perform this procedure on the primary node and on the secondary node.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Copy the archive files from your portable storage medium to `/root`.
- 3 Set the file permissions of the self-extracting archive files to execute.

```
chmod +x /root/*.run
```

- 4 Change directory to `/root`.

```
cd /root
```

- 5 Install the Resource Manager repository mirror.

```
yum install -y ./yum-mirror-*.x86_64.rpm
```

- 6 Optional: Delete the package file, if desired.

```
rm ./yum-mirror-*.x86_64.rpm
```

Upgrading Docker on the master nodes

This procedure upgrades Docker from version 1.5 or 1.8.2 to 1.9.0.

Note This release of Control Center requires RHEL/CentOS 7.1 or 7.2, and this procedure includes steps for upgrading the operating system, if necessary or desired. Zenoss recommends upgrading the operating system only at the step specified in this procedure.

Perform this procedure on the primary node and on the secondary node.

- 1 Log in to the master node as `root`, or as a user with superuser privileges.
- 2 Determine which version of Docker is installed, and then remove it, without removing Control Center.

```
rpm -qa | grep docker
```

- If the result includes `zenoss-docker`, enter the following command:

```
rpm -e --nodeps zenoss-docker-1.5.0-2
```

- If the result includes `docker-engine`, enter the following command:

```
rpm -e --nodeps docker-engine-1.8.2
```

3 Remove the Docker data partition.

- a Identify the partition where `/var/lib/docker` is mounted.

```
mount | awk '/\/var\/lib\/docker/ { print $1 }'
```

- b Unmount the partition.

Replace *Partition* with the device returned in the previous substep:

```
umount Partition
```

- c Erase the XFS file system on the partition.

Replace *Partition* with the device returned previously:

```
wipefs -a Partition
```

The partition is now ready for use as a device mapper thin pool.

- d Open `/etc/fstab` with a text editor.
- e Remove the entry for `/var/lib/docker`.

4 Upgrade the operating system, if necessary.

- a Determine which release is installed.

```
cat /etc/redhat-release
```

If the result includes `7.0`, perform the following substeps.

- b Disable the `serviced` service.

```
systemctl disable serviced
```

- c Upgrade the operating system.
You may use a local mirror or other upgrade media.
- d Restart the operating system.

```
reboot
```

- e Log in to the Control Center cluster host as `root`, or as a user with superuser privileges.
- f Enable the `serviced` service.

```
systemctl enable serviced
```

5 Install Docker 1.9.0.

```
yum clean all && yum makecache fast
yum --enablerepo=zenoss-mirror install -y docker-engine
```

The result may include a warning about `serviced` requirements, which may be ignored.

6 Create a symbolic link for the Docker temporary directory.

Docker uses its temporary directory to spool images. The default directory is `/var/lib/docker/tmp`. The following command specifies the same directory that Control Center uses, `/tmp`. You can specify any directory that has a minimum of 10GB of unused space.

- a Create the docker directory in `/var/lib`.

```
mkdir /var/lib/docker
```

- b Create the link to `/tmp`.

```
ln -s /tmp /var/lib/docker/tmp
```

- 7 Create a systemd override file for the Docker service definition.

- a Create the override directory.

```
mkdir -p /etc/systemd/system/docker.service.d
```

- b Create the override file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/docker daemon $OPTIONS -H fd://
EOF
```

- c Reload the systemd manager configuration.

```
systemctl daemon-reload
```

- 8 Install the new version of Control Center.

Control Center includes a command that simplifies the process of creating a device mapper thin pool.

```
yum --enablerepo=zenoss-mirror install -y serviced
```

The installation preserves the existing version of the `serviced` configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 9 Create a device mapper thin pool for Docker data.

- a Identify the primary partition for the thin pool to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b Create the thin pool.

Replace *Path-To-Device* with the path of an unused primary partition:

```
serviced-storage create-thin-pool docker Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 10 Configure and start the Docker service.

- a Create variables for adding arguments to the Docker configuration file.

The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

Replace *Thin-Pool-Device* with the name of the thin pool device created in the previous step:

```
myDriver="-s devicemapper"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myFlag="--storage-opt dm.thinpooldev"
myPool="Thin-Pool-Device"
```

- b Add the arguments to the Docker configuration file.

```
echo 'OPTIONS="'$myDriver $myFix $myFlag'='$myPool'"' \
>> /etc/sysconfig/docker
```

- c Start or restart Docker.

```
systemctl restart docker
```

The initial startup takes up to a minute, and may fail. If the startup fails, repeat the previous command.

11 Configure name resolution in containers.

Each time it starts, `docker` selects an IPv4 subnet for its virtual Ethernet bridge. The selection can change; this step ensures consistency.

- a Identify the IPv4 subnet and netmask `docker` has selected for its virtual Ethernet bridge.

```
ip addr show docker0 | grep inet
```

- b Open `/etc/sysconfig/docker` in a text editor.
- c Add the following flags to the end of the `OPTIONS` declaration.

Replace *Bridge-Subnet* with the IPv4 subnet `docker` selected for its virtual bridge, and replace *Bridge-Netmask* with the netmask `docker` selected:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/Bridge-Netmask
```

For example, if the bridge subnet and netmask is `172.17.0.1/16`, the flags to add are `--dns=172.17.0.1 --bip=172.17.0.1/16`.

Note Leave a blank space after the end of the thin pool device name, and make sure the double quote character (") is at the end of the line.

- d Restart the Docker service.

```
systemctl restart docker
```

Upgrading Control Center

This procedure pulls new images for Control Center internal services and updates the configuration files on the master nodes.

Perform this procedure on the primary node and on the secondary node.

- 1 Log in to the master node as `root`, or as a user with superuser privileges.
- 2 Manage the Control Center configuration files.
 - a Change directory to the configuration file directory.

```
cd /etc/default
```

- b Rename the existing configuration file, as a backup.

```
mv serviced serviced.pre-1.1.4
```

- c Rename the new configuration file, and then create a copy of it.

```
mv serviced.rpmnew serviced.orig-1.1.4
```

```
cp serviced.orig-1.1.4 serviced
```

- 3 Copy settings from the previous Control Center configuration file to the new one.
 - a Identify the customized variables in the pre-upgrade configuration file.

```
egrep '^[^#]*SERVICED' serviced.pre-1.1.4
```

- b Open `/etc/default/serviced` with a text editor, and then customize the same variables that were customized in the pre-upgrade configuration file.

The following variables are deprecated, and are not needed in the new configuration file:

- `SERVICED_REGISTRY`
- `SERVICED_VARPATH`

- c Add `SERVICED_DOCKER_REGISTRY` to the file.

The variable specifies the host and port at which the local Docker registry is available.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
SERVICED_DOCKER_REGISTRY=Hostname-Or-IP:5000
```

- d Save the file and close the text editor.
- 4 Install the new Control Center images.
 - a Change directory to `/root`.

```
cd /root
```

- b Install the Control Center ZooKeeper image.

```
./install-zenoss-isvcs-zookeeper-*.run
```

- c Install the Control Center image.

```
./install-zenoss-serviced-isvcs-*.run
```

- d Install the registry image.

```
./install-registry*.run
```

- 5 Install the Pacemaker resource agent for Control Center.

Pacemaker uses resource agents (scripts) to implement a standardized interface for managing arbitrary resources in a cluster. Zenoss provides a Pacemaker resource agent to manage the Control Center master host in a high-availability cluster.

```
yum install -y /root/serviced-resource-agents-*.rpm
```

- 6 Disable the Docker and Control Center services.

The cluster management software controls the services.

- a Stop and disable the Docker service.

```
systemctl stop docker && systemctl disable docker
```

- b Disable the Control Center service.

```
systemctl disable serviced
```

Updating high-availability components

The procedures in this section update the DRBD and Pacemaker/Corosync components with new resources.

Updating the DRBD configuration

This procedure updates DRBD with volumes for Control Center metadata and application data, which are new with this release of Control Center.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 3 On both nodes, identify the primary partitions to use for the new volumes.

```
lsblk --output=NAME,SIZE
```

Record the paths of the primary partitions to use in the following table. The information is needed in subsequent steps and procedures.

Node	<i>Metadata-Partition</i>	<i>App-Data-Partition</i>
------	---------------------------	---------------------------

- 4 On both nodes, remove any file system signature from the partitions.

Replace *Metadata-Partition* with the path of the primary partition designated for Control Center metadata, and replace *App-Data-Partition* with the path of the primary partition designated for Control Center application data:

```
wipefs -a Metadata-Partition
wipefs -a App-Data-Partition
```

- 5 On both nodes, add entries for the Control Center metadata and application data volumes.
 - a Open `/etc/drbd.d/serviced-dfs.res` with a text editor.
 - b Add entries for the volumes.

Replace *Metadata-Partition* with the path of the primary partition designated for Control Center metadata, and replace *App-Data-Partition* with the path of the primary partition designated for Control Center application data:

```
volume 2 {
    device /dev/drbd2;
    disk Metadata-Partition;
    meta-disk internal;
}
volume 3 {
    device /dev/drbd3;
    disk App-Data-Partition;
    meta-disk internal;
}
```

- c Save the file, and then close the editor.
- 6 On both nodes, update the DRBD configuration, and start DRBD.
 - a Display the commands for updating volumes individually.

```
drbdadm create-md --dry-run serviced-dfs
```


Example result:

```
drbdmeta 0 v08 /dev/sdb1 internal create-md
drbdmeta 1 v08 /dev/sdb2 internal create-md
drbdmeta 2 v08 /dev/sdb3 internal create-md
drbdmeta 3 v08 /dev/sdb4 internal create-md
```

In this example, the paths of the primary partitions for the new volumes are `/dev/sdb3` and `/dev/sdb4`.

- b** Initialize the new volumes.

Note Do not initialize existing volumes.

Replace *Metadata-Partition* with the path of the primary partition designated for Control Center metadata, and replace *App-Data-Partition* with the path of the primary partition designated for Control Center application data:

```
drbdmeta 2 v08 Metadata-Partition internal create-md
drbdmeta 3 v08 App-Data-Partition internal create-md
```

- c** Start DRBD.

```
drbdadm up serviced-dfs && drbdadm adjust all
```

- 7** On the node that was the current node when you stopped Control Center, start synchronization, and then monitor it.

- a** Start synchronization.

```
drbdadm primary --force serviced-dfs
```

- b** Monitor the process.

```
watch drbd-overview
```

Do not proceed until the status is `UpToDate/UpToDate`, as in the following example output:

```
0:serviced-dfs/0 Connected Primary/Secondary UpToDate/UpToDate
1:serviced-dfs/1 Connected Primary/Secondary UpToDate/UpToDate
2:serviced-dfs/2 Connected Primary/Secondary UpToDate/UpToDate
3:serviced-dfs/3 Connected Primary/Secondary UpToDate/UpToDate
```

The `Primary/Secondary` values show that the command was run on the primary node; otherwise, the values are `Secondary/Primary`. Likewise, the first value in the `UpToDate/UpToDate` field is the status of the node on which the command is run, and the second value is the status of the remote node.

- c** Format the partition for Control Center metadata.

The following command uses the path of the DRBD devices defined previously, not the path of the primary partition.

```
mkfs.xfs /dev/drbd2
```

The command creates an XFS file system on the current node, and DRBD mirrors the file system to the remote node.

Configuring Logical Volume Manager

Control Center application data is managed by a device mapper thin pool created with Logical Volume Manager (LVM). This procedure adjusts the LVM configuration for mirroring by DRBD.

Perform this procedure on both master nodes.

- 1 Log in to the node as `root`, or as a user with superuser privileges.
- 2 Edit the LVM configuration file.
 - a Open `/etc/lvm/lvm.conf` with a text editor.
 - b Locate the filter directive.

The line is in the `devices` section, and is commented out.

```
# filter = [ "a|.*|/" ]
```

- c Add a filter directive to exclude the partition for Control Center application data.

Place the directive immediately below the commented filter directive.

Replace *App-Data-Partition* with the absolute path of the primary partition designated for Control Center application data.

```
filter = ["r|App-Data-Partition|"]
```

- d Disable caching and the metadata daemon.

Set the values of the `write_cache_state` and `use_lvmetad` directives to 0.

```
write_cache_state = 0
use_lvmetad = 0
```

- e Save the file and close the editor.
- 3 Verify the new directives.

```
egrep '^[^#]*(filter|write_cache|use_lvmetad)' /etc/lvm/lvm.conf
```

The output should match the changes made in the previous step.

- 4 Delete any stale cache entries.

```
rm -f /etc/lvm/cache/.cache
```

- 5 Restart the host.

```
reboot
```

Creating a thin pool for application data

This procedure creates and mirrors a device mapper thin pool for Control Center application data.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
For this procedure, the primary node is the node that was the current node when you stopped Control Center.
- 2 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 3 On both nodes, start DRBD.

```
drbdadm up all
```

- 4 On the primary node, start synchronization.

```
drbdadm primary --force serviced-dfs
```

- 5 On the primary node, create a thin pool for application data.

- a Identify the amount of space available on the primary partition for application data.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b Create a variable for 50% of the space available on the thin pool partition.

The thin pool stores application data and snapshots of the data. You can add storage to the pool at any time. Replace *Half-Of-Available-Space* with 50% of the space available in the primary partition, in gigabytes. Include the symbol for gigabytes (G) after the numeric value.

```
myFifty=Half-Of-Available-SpaceG
```

- c Create the thin pool.

The following command uses the path of the DRBD device defined previously, not the path of the primary partition for application data.

```
serviced-storage create-thin-pool -o dm.basesize=$myFifty \  
serviced /dev/drbd3 -v
```

On success, the `serviced-storage` command displays the name of the thin pool it creates. Make a note of the name, which is needed in a subsequent procedure.

- 6 Stop DRBD.

- a On the primary node, enter the following commands:

```
vgchange -a n serviced && drbdadm down all
```

- b On the secondary node, enter the following command:

```
drbdadm down all
```

Updating the cluster configuration

This procedure adds new resources to the Control Center cluster configuration.

- 1 Log in to a master node as `root`, or as a user with superuser privileges.
For `pcs` commands, the nodes are equivalent.
- 2 Start the cluster.

```
pcs cluster start
```

- 3 Define the LVM resource.

```
pcs resource create serviced-lvm ocf:heartbeat:LVM \  
volgrpname=serviced
```

- 4 Define the storage resource for Control Center application data.

This ensures the the thin pool is deactivated and unmounted correctly.

```
pcs resource create serviced-storage \  
ocf:zenoss:serviced-storage
```

- 5 Add the new resources to the Control Center resource group.

```
pcs resource group add serviced-group \
  serviced-lvm --after serviced-volumes
pcs resource group add serviced-group \
  serviced-storage --after serviced-lvm
```

- 6 Verify Control Center resource agent timeouts.

With this release, the recommended resource agent timeouts for starting and stopping Control Center have changed to 360 and 130 seconds, respectively.

- a Display the current values.

```
pcs resource show serviced
```

- b Update the values, if necessary.

```
pcs resource update serviced op start timeout=360s
pcs resource update serviced op stop timeout=130s
```

- 7 Disable Pacemaker monitoring of NFS health.

During normal operations, Control Center occasionally stops and restarts NFS, which could be misinterpreted by Pacemaker and trigger an unwanted failover.

```
pcs resource op remove nfs monitor interval=60s
pcs resource op add nfs monitor interval=0s
```

- 8 Verify the cluster configuration.

```
pcs resource show --groups
```

The correct result shows `serviced-group` with 8 members, in the following order:

```
serviced-group: serviced-isvcs serviced-volumes serviced-lvm serviced-
storage VirtualIP docker nfs serviced
```

Verifying cluster startup

This procedure verifies the new configuration by attempting to start the resources on one node only. With the other node in standby mode, Pacemaker does not automatically fail over to the other node.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
For this procedure, the primary node is the node that was the current node when you stopped Control Center.
- 2 Determine which node is the primary DRBD node.

```
pcs status
```

Example result:

```
Cluster name: serviced-ha
Last updated: Mon Feb 22 11:37:58 2016 Last change: Mon Feb 22
 11:35:19 2016 by root via crm_attribute on Secondary-Public-Name
Stack: corosync
Current DC: Primary-Public-Name (version 1.1.13-a14efad) - partition
with quorum
2 nodes and 10 resources configured
```

```

Node Primary-Public-Name: standby
Node Secondary-Public-Name: standby

Full list of resources:

Master/Slave Set: DFSMaster [DFS]
Stopped: [ Primary-Public-Name Secondary-Public-Name ]
Resource Group: serviced-group
  serviced-isvcs (ocf::heartbeat:Filesystem): Stopped
  serviced-volumes (ocf::heartbeat:Filesystem): Stopped
  serviced-lvm (ocf::heartbeat:LVM): Stopped
  serviced-storage (ocf::zenoss:serviced-storage): Stopped
  VirtualIP (ocf::heartbeat:IPaddr2): Stopped
  docker (systemd:docker): Stopped
  nfs (systemd:nfs): Stopped
  serviced (ocf::zenoss:serviced): Stopped

PCSD Status:
  Primary-Public-Name: Offline
  Secondary-Public-Name: Offline

Daemon Status:
  corosync: active/disabled
  pacemaker: active/enabled
  pcsd: active/enabled

```

The line that begins with `Current DC` identifies the primary node. Review all of the command output for errors.

3 Start DRBD.

- a** On the secondary node, enter the following command:

```
drbdadm up all
```

- b** On the primary node, enter the following commands:

```
drbdadm up all && drbdadm primary serviced-dfs
```

4 Take the current node out of standby mode.

- a** Display the public hostname of the current node.

```
uname -n
```

- b** Clean up before bringing the node out of standby mode.

```
pcs resource cleanup --all
```

- c** Put the current node back online.

Replace *Primary-Public-Name* with the hostname of the current node:

```
pcs cluster unstandby Primary-Public-Name
```

5 Monitor the status of cluster resources.

```
watch pcs status
```

Monitor the status until all resources report `Started`. Resolve any issues before continuing.

Verifying cluster failover

This procedure simulates a failover.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Enable the secondary node.
 - a Take the secondary node out of standby mode.

Replace *Secondary-Public-Name* with the public hostname of the secondary node:

```
pcs cluster unstandby Secondary-Public-Name
```

- b Monitor the status of the secondary node.

```
watch pcs status
```

Do not continue until the status of the secondary node is `Online`.

- 3 Verify that DRBD has completely synchronized all four volumes on the secondary node.

```
drbd-overview
```

Example result:

```
0:serviced-dfs/0 Connected Primary/Secondary UpToDate/UpToDate
1:serviced-dfs/1 Connected Primary/Secondary UpToDate/UpToDate
2:serviced-dfs/2 Connected Primary/Secondary UpToDate/UpToDate
3:serviced-dfs/3 Connected Primary/Secondary UpToDate/UpToDate
```

- 4 Force a failover.
Pacemaker initiates a failover when the primary node is put in standby mode.

Replace *Primary-Public-Name* with the public hostname of the primary node:

```
pcs cluster standby Primary-Public-Name
```

- 5 Monitor the cluster status.

```
watch pcs status
```

Wait until all resources report a status of `Started`. Resolve any issues before continuing.

- 6 Restore cluster services.

This starts the secondary node and Control Center internal services, but does not start Resource Manager services, which is intentional. Resource pool hosts have to be upgraded first.

```
pcs cluster unstandby --all
```

Upgrading Control Center on resource pool hosts

Perform the procedures in this section on each resource pool host in your Control Center deployment.

Staging files for offline installation

To perform this procedure, you need the portable storage medium that contains the archive files used in installing the master host.

This procedure adds files for offline installation to a resource pool host. The files are required in subsequent procedures.

Perform this procedure on each resource pool host in your deployment.

- 1 Log in to the target host as `root`, or as a user with superuser privileges.
- 2 Copy `yum-mirror-*.x86_64.rpm` from your portable storage medium to `/tmp`.
- 3 Install the Resource Manager repository mirror.

```
yum install -y /tmp/yum-mirror-*.x86_64.rpm
```

- 4 Optional: Delete the package file, if desired.

```
rm /tmp/yum-mirror-*.x86_64.rpm
```

Upgrading Docker

This procedure upgrades Docker from version 1.5 or 1.8.2 to 1.9.0.

Note This release of Control Center requires RHEL/CentOS 7.1 or 7.2, and this procedure includes steps for upgrading the operating system, if necessary or desired. Zenoss recommends upgrading the operating system only at the step specified in this procedure. However, instructions for performing an upgrade without internet access are not included.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Stop Control Center.

```
systemctl stop serviced
```

- 3 Determine which version of Docker is installed, and then remove it, without removing Control Center.

```
rpm -qa | grep docker
```

- If the result includes `zenoss-docker`, enter the following command:

```
rpm -e --nodeps zenoss-docker-1.5.0-2
```

- If the result includes `docker-engine`, enter the following command:

```
rpm -e --nodeps docker-engine-1.8.2
```

- 4 Remove the Docker data partition.
 - a Identify the partition where `/var/lib/docker` is mounted.

```
mount | awk '/\/var\/lib\/docker/ { print $1 }'
```

- b Unmount the partition.

Replace *Partition* with the device returned in the previous substep:

```
umount Partition
```

- c Erase the XFS file system on the partition.

Replace *Partition* with the device returned previously:

```
wipefs -a Partition
```

The partition is now ready for use as a device mapper thin pool.

- d Open `/etc/fstab` with a text editor.
 - e Remove the entry for `/var/lib/docker`.
- 5 Upgrade the operating system, if necessary.
- a Determine which release is installed.

```
cat /etc/redhat-release
```

If the result includes 7.0, perform the following substeps.

- b Disable the `serviced` service.

```
systemctl disable serviced
```

- c Upgrade the operating system.
You may use a local mirror or other upgrade media.
- d Restart the operating system.

```
reboot
```

- e Log in to the Control Center cluster host as `root`, or as a user with superuser privileges.
- f Enable the `serviced` service.

```
systemctl enable serviced
```

- 6 Install Docker 1.9.0, and then disable accidental upgrades.
- a Add the Docker repository to the host's repository list.

```
cat > /etc/yum.repos.d/docker.repo <<-EOF
[dockerrepo]
name=Docker Repository
baseurl=https://yum.dockerproject.org/repo/main/centos/7
enabled=1
gpgcheck=1
gpgkey=https://yum.dockerproject.org/gpg
EOF
```

- b Install Docker 1.9.0.

```
yum clean all && yum makecache fast
yum install -y docker-engine-1.9.0
```

- c Open `/etc/yum.repos.d/docker.repo` with a text editor.
 - d Change the value of the `enabled` key from 1 to 0.
 - e Save the file and close the text editor.
- 7 Create a `systemd` override file for the Docker service definition.
- a Create the override directory.

```
mkdir -p /etc/systemd/system/docker.service.d
```


- b** Create the override file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/docker daemon \${OPTIONS} -H fd://
EOF
```

- c** Reload the systemd manager configuration.

```
systemctl daemon-reload
```

- 8** Configure the docker service to start when the system starts.

```
systemctl enable docker
```

- 9** Install the new version of Control Center.

Control Center includes a command that simplifies the process of creating a device mapper thin pool.

```
yum --enablerepo=zenoss-mirror install -y serviced
```

The installation preserves the existing version of the serviced configuration file, and installs the new one as `/etc/default/serviced.rpmnew`.

- 10** Create a device mapper thin pool for Docker data.

- a** Identify the primary partition for the thin pool to create.

```
lsblk --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
```

- b** Create the thin pool.

Replace *Path-To-Device* with the path of an unused primary partition:

```
serviced-storage create-thin-pool docker Path-To-Device
```

On success, the result includes the name of the thin pool, which always starts with `/dev/mapper`.

- 11** Configure and start the Docker service.

- a** Create variables for adding arguments to the Docker configuration file.

The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

Replace *Thin-Pool-Device* with the name of the thin pool device created in the previous step:

```
myDriver="-s devicemapper"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myFlag="--storage-opt dm.thinpooldev"
myPool="Thin-Pool-Device"
```

- b** Add the arguments to the Docker configuration file.

```
echo 'OPTIONS="'$myDriver $myFix $myFlag'='$myPool'' \
>> /etc/sysconfig/docker
```

- c** Start or restart Docker.

```
systemctl restart docker
```

The initial startup takes up to a minute, and may fail. If the startup fails, repeat the previous command.

12 Configure name resolution in containers.

Each time it starts, `docker` selects an IPv4 subnet for its virtual Ethernet bridge. The selection can change; this step ensures consistency.

- a Identify the IPv4 subnet and netmask `docker` has selected for its virtual Ethernet bridge.

```
ip addr show docker0 | grep inet
```

- b Open `/etc/sysconfig/docker` in a text editor.
- c Add the following flags to the end of the `OPTIONS` declaration.

Replace *Bridge-Subnet* with the IPv4 subnet `docker` selected for its virtual bridge, and replace *Bridge-Netmask* with the netmask `docker` selected:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/Bridge-Netmask
```

For example, if the bridge subnet and netmask is 172.17.0.1/16, the flags to add are `--dns=172.17.0.1 --bip=172.17.0.1/16`.

Note Leave a blank space after the end of the thin pool device name, and make sure the double quote character (") is at the end of the line.

- d Restart the Docker service.

```
systemctl restart docker
```

Upgrading Control Center

This procedure updates the configuration files on resource pool hosts.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Manage the Control Center configuration files.
 - a Change directory to the configuration file directory.

```
cd /etc/default
```

- b Rename the existing configuration file, as a backup.

```
mv serviced serviced.pre-1.1.4
```

- c Rename the new configuration file, and then create a copy of it.

```
mv serviced.rpmnew serviced.orig-1.1.4
cp serviced.orig-1.1.4 serviced
```

- 3 Copy settings from the previous Control Center configuration file to the new one.
 - a Identify the customized variables in the pre-upgrade configuration file.

```
egrep '^[^#]*SERVICED' serviced.pre-1.1.4
```

- b Open `/etc/default/serviced` with a text editor, and then customize the same variables that were customized in the pre-upgrade configuration file.

The following variables are deprecated, and are not needed in the new configuration file:

- `SERVICED_REGISTRY`

- `SERVICED_VARPATH`
- c Add `SERVICED_DOCKER_REGISTRY` to the file.
The variable specifies the host and port at which the local Docker registry is available.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host:

```
SERVICED_DOCKER_REGISTRY=Hostname-Or-IP:5000
```

- d Save the file and close the text editor.
- 4 Unmount the distributed file system (DFS).
 - a Identify the file system specification to unmount.

```
mount | awk '/serviced/ { print $1 }'
```

- b Unmount the DFS.
Replace *DFS-Mount* with the file system specification returned in the previous substep:

```
umount DFS-Mount
```

- 5 Start Control Center.

```
systemctl start serviced
```

Converting the data storage driver

At this point in the upgrade, Control Center is upgraded on the master nodes and on all of the resource pool hosts. You may continue using Resource Manager to monitor your infrastructure during the initial (and longest) part of the conversion process. However, Zenoss strongly recommends running Resource Manager in this configuration only as long as it takes to migrate Btrfs data to the new thin pool.

The procedures in this section migrate Resource Manager data from a Btrfs volume to a device mapper thin pool volume.

Starting the conversion

To perform this procedure, you need an unused primary partition that is large enough for your data. For more information, see [Preparing to convert the data storage driver](#) on page 47.

This procedure converts the majority of Resource Manager data stored in a Btrfs file system and copies it to a device mapper thin pool. This procedure may be performed while Resource Manager is monitoring your infrastructure. Control Center application data is moved to its new thin pool volume, and Control Center metadata is moved to its new XFS volume.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Verify that cluster services are running.

```
watch pcs status
```

Do not proceed until both master nodes are online and all resources are started.

- 3 Initialize `serviced-storage` for the `btrfs` storage driver.

```
serviced-storage init /opt/serviced/var/volumes btrfs
```

- 4 Mount the DRBD volume for Control Center metadata at a temporary location.
The device path is `/dev/drbd2`.

- a Create a mount point for the temporary location.

Replace *Temporary-Directory* with the path of a temporary directory:

```
mkdir -p Temporary-Directory
```

- b Mount the DRBD volume.

```
mount /dev/drbd2 Temporary-Directory
```

- 5 Synchronize the contents of the Btrfs volume with the new storage.

This step copies current data to the new thin pool, and current metadata to the new XFS volume.

Replace *Thin-Pool-Name* with the name of the thin pool created in a previous procedure, and replace *Temporary-Directory* with the path of the temporary location created in the previous step:

```
serviced-storage sync -c -t devicemapper \
-o dm.thinpooldev=Thin-Pool-Name \
/opt/serviced/var/volumes Temporary-Directory
```

After a pause to compute space requirements, the `serviced-storage` command displays both detail and summary information about its work. Depending on the amount of data to convert, this step may take several hours.

Note Since both the thin pool and the XFS file system are DRBD volumes, the data copied in this step is automatically mirrored to the secondary node.

Stopping applications

This procedure stops all Control Center applications.

- 1 Use the virtual hostname or virtual IP address of the high-availability cluster to log in to the Control Center master node as `root`, or as a user with superuser privileges.
- 2 Identify the applications to stop.
 - a Identify the applications that Control Center is managing.

```
serviced service list | awk '/Zenoss\.\/ { print $1 }'
```

- b Identify the applications that are running.

Replace *Application* with the name of each application returned from the preceding commands:

```
serviced service status Application
```

- 3 Stop Resource Manager, and then verify it is stopped.

- a Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

- b Verify the application is stopped.

Repeat the following command until the STATUS column reads Stopped:

```
serviced service status Zenoss.resmgr
```

Finalizing the conversion

This procedure completes the conversion started previously.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 Stop Control Center on each resource pool host.
 - a Log in to each resource pool host as `root`, or as a user with superuser privileges.
 - b Stop Control Center

```
systemctl stop serviced
```

- 3 On the primary node, stop the Control Center service.
 - a Stop the resource.

```
pcs resource disable serviced
```

- b Monitor the shutdown.

```
watch pcs status
```

Do not proceed until the `serviced` resource status is `Stopped`.

- 4 On the primary node, synchronize the contents of the Btrfs volume with the new storage. This step copies data created since the previous sync operation.

Repeat the command you used before stopping the Control Center service:

```
serviced-storage sync -c -t devicemapper \  
-o dm.thinpooldev=Thin-Pool-Name \  
/opt/serviced/var/volumes Temporary-Directory
```

- 5 On the primary node, disable the thin pool.

```
serviced-storage disable Temporary-Directory \  
-o dm.thinpooldev=Thin-Pool-Name
```

- 6 On the primary node, unmount the DRBD volume for Control Center metadata.

```
umount Temporary-Directory
```

- 7 On the primary node, stop all cluster services.
 - a Stop all resources.

```
pcs cluster standby --all
```

- b Monitor the shutdown.

```
watch pcs status
```

Do not proceed until the status of all resources is `Stopped`.

- 8 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 9 On both nodes, update the Control Center configuration file with new storage settings.
 - a Open `/etc/default/serviced` with a text editor.
 - b Locate the `SERVICED_FS_TYPE` declaration, and change the value from `btrfs` to `devicemapper`.
 - c Add `SERVICED_DM_THINPOOLDEV` to the file, immediately after `SERVICED_FS_TYPE`.

The variable specifies the name of the thin pool device for Control Center application data.

Replace *Thin-Pool-Name* with the name of the thin pool created used to synchronize data:

```
SERVICED_DM_THINPOOLDEV=Thin-Pool-Name
```

- d Save the file and close the text editor.
- 10 On both nodes, update the DRBD resource definition, and then update DRBD.

- a Open `/etc/drbd.d/serviced-dfs.res` with a text editor.
- b Remove the entry for the Btrfs file system.

The entry to remove should be for volume 1, which references `/dev/drbd1`.

In the following example, the value of *Btrfs-Partition* is unique in your environment:

```
volume 1 {
    device /dev/drbd1;
    disk Btrfs-Partition;
    meta-disk internal;
}
```

- c Save the file and close the text editor.
- d Update DRBD.

```
drbdadm up all && drbdadm adjust all && drbdadm down all
```

- 11 On the primary node, update the Pacemaker configuration.

This change adds the Control Center metadata volume, DRBD device `/dev/drbd2`.

```
pcs resource update serviced-volumes \
    device=/dev/drbd/by-res/serviced-dfs/2 \
    fstype=xfs
```

- 12 On the primary node, start all cluster services.

```
pcs resource enable serviced && pcs cluster unstandby --all
```

- 13 Monitor the startup.

Do not proceed until all resources are started, DRBD is finished synchronizing, and both master nodes are online.

- a On the primary node, open two separate shell sessions or windows.
- b In one session, monitor the cluster startup.

```
watch pcs status
```

- c In another session, monitor the DRBD synchronization startup.

```
watch drbd-overview
```

- 14 Start Control Center on each resource pool host.

- a Log in to each resource pool host as `root`, or as a user with superuser privileges.
- b Start Control Center

```
systemctl start serviced
```

Verifying the conversion

This procedure verifies that the new data storage works properly with Resource Manager version 1.0.x.

- 1 Log in to the Control Center browser interface.
- 2 Start Resource Manager and all related applications.
- 3 Perform database integrity checks.
For more information, see [Using Zenoss Toolbox](#) on page 163.
- 4 Create a backup.
Previous backups are incompatible with the new storage driver.

When you are satisfied, continue with the Resource Manager upgrade procedure.

ZooKeeper ensemble configuration

Control Center relies on [Apache ZooKeeper](#) to coordinate its services. The configuration steps in this section create a ZooKeeper ensemble of 3 nodes.

A ZooKeeper ensemble requires a minimum of 3 nodes, and 3 nodes is sufficient for most deployments. A 5-node configuration improves failover protection during maintenance windows. Ensembles larger than 5 nodes are not necessary. An odd number of nodes is recommended, and an even number of nodes is strongly discouraged.

Control Center variables for ZooKeeper

This tables in this section associates the ZooKeeper-related Control Center variables to set in `/etc/default/serviced` with the roles that hosts play in a Control Center cluster.

Table 12: Control Center master host

<code>SERVICED_ISVCS_ZOOKEEPER_ID</code>
The unique identifier of a ZooKeeper ensemble node. Value: 1
<code>SERVICED_ISVCS_ZOOKEEPER_QUORUM</code>
The ZooKeeper node ID, IP address, peer communications port, and leader communications port of each host in an ensemble. Each quorum definition must be unique, so the IP address of the "current" host is 0.0.0.0. Value: <code>ZooKeeper-ID@IP-Address:2888:3888, ...</code>
<code>SERVICED_ZK</code>
The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it. Value: <code>IP-Address:2181, ...</code>

Table 13: Control Center resource pool host and ZooKeeper ensemble node

<code>SERVICED_ISVCS_ZOOKEEPER_ID</code>
The unique identifier of a ZooKeeper ensemble node. Value: 2 or 3
<code>SERVICED_ISVCS_ZOOKEEPER_QUORUM</code>

The ZooKeeper node ID, IP address, peer communications port, and leader communications port of each host in an ensemble. Each quorum definition must be unique, so the IP address of the "current" host is 0.0.0.0.

Value: *ZooKeeper-ID@IP-Address:2888:3888, . . .*

SERVICED_ISVCS_START

The list of Control Center internal services to start and run on hosts other than the master host.

Value: *zookeeper*

SERVICED_ZK

The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it.

Value: *IP-Address:2181, . . .*

Table 14: Control Center resource pool host

SERVICED_ZK

The list of endpoints in the Control Center ZooKeeper ensemble, separated by the comma character (,). Each endpoint includes the IP address of the ensemble node, and the port that Control Center uses to communicate with it.

Value: *IP-Address:2181, . . .*

Configuring master nodes

This procedure configures both Control Center master nodes as members of the ZooKeeper ensemble.

Note For accuracy, this procedure constructs Control Center configuration variables in the shell and appends them to `/etc/default/serviced`. The last step is to move the variables from the end of the file to more appropriate locations.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 3 On both nodes, create a variable for each Control Center host to include in the ZooKeeper ensemble.

The variables are used in subsequent steps.

Note Define the variables identically on both the primary and the secondary nodes, and on each resource pool host.

Replace *HA-Virtual-IP* with the virtual IP address of the high-availability cluster, and replace *Pool-Host-A-IP* and *Pool-Host-B-IP* with the IP addresses of the Control Center resource pool hosts to include in the ensemble:

```
node1=HA-Virtual-IP
node2=Pool-Host-A-IP
node3=Pool-Host-B-IP
```

Note ZooKeeper requires IP addresses for ensemble configuration.

- 4 On both nodes, set the ZooKeeper node ID to 1.

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=1" >> /etc/default/serviced
```


- 5 On both nodes, specify the nodes in the ZooKeeper ensemble.
You may copy the following text and paste it in your console:

```
echo "SERVICED_ZK=${node1}:2181,${node2}:2181,${node3}:2181" \
>> /etc/default/serviced
```

- 6 On both nodes, specify the nodes in the ZooKeeper quorum.
ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address of the current node with 0.0.0.0.
You may copy the following of text and paste it in your console:

```
q1="1@0.0.0.0:2888:3888"
q2="2@${node2}:2888:3888"
q3="3@${node3}:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

- 7 On both nodes, clean up the Control Center configuration file.
- a Open `/etc/default/serviced` with a text editor.
 - b Navigate to the end of the file, and cut the line that contains the `SERVICED_ZK` variable declaration at that location.
The value of this declaration specifies 3 hosts.
 - c Locate the `SERVICED_ZK` variable near the beginning of the file, and then delete the line it is on.
The value of this declaration is just the master node.
 - d Paste the `SERVICED_ZK` variable declaration from the end of the file in the location of the just-deleted declaration.
 - e Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration at that location.
 - f Locate the `SERVICED_ISVCS_ZOOKEEPER_ID` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - g Paste the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration from the end of the file in the location of the just-deleted declaration.
 - h Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration at that location.
 - i Locate the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - j Paste the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration from the end of the file in the location of the just-deleted declaration.
 - k Save the file, and then close the text editor.
- 8 On both hosts, verify the ZooKeeper environment variables.

```
egrep '^[^#]*SERVICED' /etc/default/serviced | egrep '(_ZOO|_ZK)'
```

Configuring a resource pool host as a ZooKeeper node

To perform this procedure, you need a resource pool host with an XFS file system on a separate partition.

This procedure configures a ZooKeeper ensemble on a resource pool host. Repeat this procedure on each Control Center resource pool host to add to the ZooKeeper ensemble.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.

- 2 Create a variable for each Control Center host to include in the ZooKeeper ensemble.

Replace *HA-Virtual-IP* with the virtual IP address of the high-availability cluster, and replace *Pool-Host-A-IP* and *Pool-Host-B-IP* with the IP addresses of the Control Center resource pool hosts to include in the ensemble:

```
node1=HA-Virtual-IP
node2=Pool-Host-A-IP
node3=Pool-Host-B-IP
```

- 3 Set the ID of this node in the ZooKeeper ensemble.

For *Pool-Host-A-IP* (node2), use the following command:

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=2" >> /etc/default/serviced
```

For *Pool-Host-B-IP* (node3), use the following command:

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=3" >> /etc/default/serviced
```

- 4 Specify the nodes in the ZooKeeper ensemble.
You may copy the following text and paste it in your console:

```
echo "SERVICED_ZK=${node1}:2181,${node2}:2181,${node3}:2181" \
>> /etc/default/serviced
```

- 5 Specify the nodes in the ZooKeeper quorum.

ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address of the current node with 0.0.0.0.

For *Pool-Host-A-IP* (node2), use the following commands:

```
q1="1@${node1}:2888:3888"
q2="2@0.0.0.0:2888:3888"
q3="3@${node3}:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

For *Pool-Host-B-IP* (node3), use the following commands:

```
q1="1@${node1}:2888:3888"
q2="2@${node2}:2888:3888"
q3="3@0.0.0.0:2888:3888"
echo "SERVICED_ISVCS_ZOOKEEPER_QUORUM=${q1},${q2},${q3}" \
>> /etc/default/serviced
```

- 6 Set the *SERVICED_ISVCS_START* variable, and clean up the Control Center configuration file.
 - a Open */etc/default/serviced* with a text editor.
 - b Locate the *SERVICED_ISVCS_START* variable, and then delete all but *zookeeper* from its list of values.
 - c Remove the number sign character (#) from the beginning of the line.
 - d Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK* variable declaration at that location.
The value of this declaration specifies 3 hosts.
 - e Locate the *SERVICED_ZK* variable near the beginning of the file, and then delete the line it is on.
The value of this declaration is just the master node.

- f** Paste the `SERVICED_ZK` variable declaration from the end of the file in the location of the just-deleted declaration.
 - g** Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration at that location.
 - h** Locate the `SERVICED_ISVCS_ZOOKEEPER_ID` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - i** Paste the `SERVICED_ISVCS_ZOOKEEPER_ID` variable declaration from the end of the file in the location of the just-deleted declaration.
 - j** Navigate to the end of the file, and cut the line that contains the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration at that location.
 - k** Locate the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable near the end of the file, and then delete the line it is on.
This declaration is commented out.
 - l** Paste the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable declaration from the end of the file in the location of the just-deleted declaration.
 - m** Save the file, and then close the text editor.
- 7** Verify the ZooKeeper environment variables.

```
egrep '^[^#]*SERVICED' /etc/default/serviced \
| egrep '(_ZOO|_ZK|_STA)'
```

- 8** Pull the required Control Center ZooKeeper image from the master host.
- a** Identify the image to pull.

```
serviced version | grep IsvcsImages
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v40 zenoss/isvcs-zookeeper:v3]
```

- b** Pull the Control Center ZooKeeper image.

Replace `Isvcs-ZK-Image` with the name and version number of the ZooKeeper image from the previous substep:

```
docker pull Isvcs-ZK-Image
```

Starting a ZooKeeper ensemble

This procedure starts a ZooKeeper ensemble.

The window of time for starting a ZooKeeper ensemble is relatively short. The goal of this procedure is to restart Control Center on each ensemble node at about the same time, so that each node can participate in electing the leader.

- 1** Use the virtual hostname (*HA-Virtual-Name*) or virtual IP address (*HA-Virtual-IP*) of the high-availability cluster to start a Bash shell on the Control Center master host as `root`, or as a user with superuser privileges.
- 2** Display the public hostname of the current node.

```
uname -n
```

The result is either *Primary-Public-Name* or *Secondary-Public-Name*.

- 3** Place the other node in standby mode.

This avoids potential conflicts and errors in the event of an unexpected `serviced` shutdown during the ZooKeeper startup.

Replace `Other-Node-Hostname` with the public hostname of the other master node:

```
pcs cluster standby Other-Node-Hostname
```

- 4 In a separate window, log in to the second node of the ZooKeeper ensemble (`Pool-Host-A-IP`).
- 5 In another separate window, log in to the third node of the ZooKeeper ensemble (`Pool-Host-B-IP`).
- 6 On the primary node, stop and start `serviced`.

```
pcs resource disable serviced && pcs resource enable serviced
```

- 7 On both resource pool hosts, stop and start `serviced`.

```
systemctl stop serviced && systemctl start serviced
```

- 8 On the primary node, check the status of the ZooKeeper ensemble.

```
{ echo stats; sleep 1; } | nc localhost 2181 | grep Mode
{ echo stats; sleep 1; } | nc Pool-Host-A-IP 2181 | grep Mode
{ echo stats; sleep 1; } | nc Pool-Host-B-IP 2181 | grep Mode
```

If `nc` is not available, you can use `telnet` with [interactive ZooKeeper commands](#).

- 9 Restore the cluster.

Replace `Other-Node-Hostname` with the public hostname of the other master node:

```
pcs cluster unstandby Other-Node-Hostname
```

Updating resource pool hosts

The default configuration of resource pool hosts sets the value of the `SERVICED_ZK` variable to the master host only. This procedure updates the setting to include the full ZooKeeper ensemble.

Perform this procedure on each resource pool host in your Control Center cluster.

- 1 Log in to the resource pool host as `root`, or as a user with superuser privileges.
- 2 Update the variable.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the `SERVICED_ZK` declaration, and then replace its value with the same value used in the ZooKeeper ensemble nodes.
 - c Save the file, and then close the editor.
- 3 Restart Control Center.

```
systemctl restart serviced
```

Upgrading Resource Manager

This procedure upgrades Resource Manager.

Before performing this procedure, stop Resource Manager. For more information, see [Stopping applications](#) on page 19.

- 1 Log in to the primary node as `root`, or as a user with superuser privileges.

2 Install the new Resource Manager images.

- a Change directory to /root.

```
cd /root
```

- b Install the HBase image.

```
./install-zenoss-hbase*.run
```

- c Install the OpenTSDB image.

```
./install-zenoss-opentsdb*.run
```

- d Install the Resource Manager image.

```
./install-zenoss-resmgr*.run
```

- e Optional: Delete the self-extracting image files, if desired.

```
rm ./install-zenoss-*.run
```

3 Extract files from the Resource Manager image.

```
docker run -it --rm -v /root:/mnt/root \
  zenoss/resmgr_5.1:5.1.3_1 rsync -a /root/5.1.x /mnt/root
```

The preceding commands copy upgrade scripts to /root/5.1.x.

4 Start the upgrade script.

The script to start depends on whether Service Impact is installed.

- If Service Impact is installed, enter the following command.

```
/root/5.1.x/upgrade-impact-5.1.x.sh
```

Note The script upgrades Resource Manager, but does not upgrade Service Impact. For more information about upgrading Service Impact, refer to the *Zenoss Service Impact Installation Guide for Resource Manager 5.x*.

- If Service Impact is not installed, enter the following command.

```
/root/5.1.x/upgrade-resmgr-5.1.x.sh
```

5 Restart Resource Manager.

Some Resource Manager services are started during the upgrade, and they need to be restarted.

```
serviced service restart Zenoss.resmgr
```

Proceed to *After upgrading* on page 158.

After upgrading

This chapter includes information about what to do after upgrading your deployment of Control Center and Resource Manager.

Preventing OpenTSDB health check failures

Deployments that were upgraded to the combination of Control Center 1.0.3 and Resource Manager 5.0.3 in the past could experience false negative OpenTSDB health checks. This procedure determines whether your deployment is affected, and prevents the health check failures.

- 1 Log in to the Control Center browser interface.
- 2 In the **Applications** table, click **Zenoss.resmgr**.
- 3 Scroll down to the **Services** table, and then click **Infrastructure > opentsdb > reader**.
- 4 In the **Actions** column of the **Configuration Files** table, click **Edit**.
- 5 In the **Edit Configuration** dialog, review the `tsd.storage.hbase.*` keys.

The following list shows the incorrect and correct values for each affected key:

tsd.storage.bbase.data_table

Incorrect: `{{ (parent (parent .)).ID}}-tsdb`

Correct: `{{ (parent (parent (parent .))).ID}}-tsdb`

tsd.storage.bbase.uid_table

Incorrect: `{{ (parent (parent .)).ID}}-tsdb-uid`

Correct: `{{ (parent (parent (parent .))).ID}}-tsdb-uid`

tsd.storage.bbase.meta_table

Incorrect: `{{ (parent (parent .)).ID}}-tsdb-meta`

Correct: `{{ (parent (parent (parent .))).ID}}-tsdb-meta`

tsd.storage.bbase.tree_table

Incorrect: `{{ (parent (parent .)).ID}}-tsdb-tree`

Correct: `{{ (parent (parent (parent .))).ID}}-tsdb-tree`

- If the values cause health check failures, perform the remaining steps of this procedure.
 - If the values do not cause health check failures, no additional steps are required.
- 6 Update the `tsd.storage.hbase.*` keys.
Use the correct values listed in the previous step.
 - 7 In the **Edit Configuration** dialog, click the **Save** button.
 - 8 Repeat steps 3 through 7 for **Infrastructure > opentsdb > writer**

Checking ZooKeeper quorum keys

Some upgrades may have missed an upgrade to an OpenTSDB key. This procedure determines whether your deployment is affected, and corrects the oversight.

- 1 Log in to the Control Center browser interface.
- 2 In the **Applications** table, click **Zenoss.resmgr**.
- 3 Scroll down to the **Services** table, and then click **Infrastructure > opentsdb > reader**.
- 4 In the **Actions** column of the **Configuration Files** table, click **Edit**.
- 5 In the **Edit Configuration** dialog, review the value of the `tsd.storage.hbase.zk_quorum` key.

In the following example, the value includes `localhost`:

```
tsd.storage.hbase.zk_quorum = {{with $zks := (child (child (parent
  (parent .)) "HBase") "ZooKeeper").Instances }}{{ range (each
  $zks) }}localhost:{{plus 2181 .}}{{if ne (plus 1 .) $zks}},{{end}}
  {{end}}}}}}
```

The correct value of the `tsd.storage.hbase.zk_quorum` key replaces `localhost` with `127.0.0.1`.

- If the value includes `localhost`, perform the remaining steps of this procedure.
 - If the value does not include `localhost`, no additional steps are required.
- 6 Update the value of the `tsd.storage.hbase.zk_quorum` key.
Replace `localhost` with `127.0.0.1`.
 - 7 In the **Edit Configuration** dialog, click the **Save** button.
 - 8 Repeat steps 3 through 7 for **Infrastructure > > writer**

Correcting a CentralQuery configuration file

Some upgrades may have an incorrect configuration for logging in a CentralQuery configuration file. This procedure determines whether your deployment is affected, and corrects the oversight.

- 1 Log in to the Control Center browser interface.
- 2 In the **Applications** table, click **Zenoss.resmgr**.
- 3 Scroll down to the **Services** table, and then click **Zenoss > Metrics > CentralQuery**.
- 4 In the **Actions** column of the **Configuration Files** table, click the **Edit** control of the `/opt/zenoss/etc/central-query/configuration.yaml` file.
- 5 In the **Edit Configuration** dialog, review the declaration for logging.

The following example shows the correct declaration:

```
logging:
  level: INFO
  loggers:
    "org.zenoss": INFO
```

- If the declaration for logging is not correct, perform the remaining steps of this procedure.
 - If the declaration for logging is correct, no additional steps are required.
- 6 Update the declaration for logging.

Use the declaration in the following example:

```
logging:
  level: INFO
  loggers:
```

```
"org.zenoss": INFO
```

- 7 In the **Edit Configuration** dialog, click the **Save** button.

Configuring ZooKeeper log purging

Some upgrades may be missing two key-value pairs in a ZooKeeper configuration file. This procedure determines whether your deployment is affected, and corrects the oversight.

- 1 Log in to the Control Center browser interface.
- 2 In the **Applications** table, click **Zenoss.resmgr**.
- 3 Scroll down to the **Services** table, and then click **Infrastructure > HBase > ZooKeeper**.
- 4 In the **Actions** column of the **Configuration Files** table, click the **Edit** control of the `/etc/zookeeper.cfg` file.
- 5 In the **Edit Configuration** dialog, review the contents.

The following example shows the key-value pairs that may be missing:

```
autopurge.snapRetainCount=3
autopurge.purgeInterval=1
```

- If the key-value pairs are missing, perform the remaining steps of this procedure.
 - If the key-value pairs are present, no additional steps are required.
- 6 Update the contents of the configuration file.
Add the content of the following example:

```
autopurge.snapRetainCount=3
autopurge.purgeInterval=1
```

- 7 In the **Edit Configuration** dialog, click the **Save** button.

Increasing Zope server threads

Some upgrades may have the value of the Zope server threads key set to 1 instead of the default, 4. This procedure determines whether your deployment is affected, and corrects the oversight.

- 1 Log in to the Control Center browser interface.
- 2 In the **Applications** table, click **Zenoss.resmgr**.
- 3 Scroll down to the **Services** table, and then click **Zenoss > User Interface > Zope**.
- 4 In the **Actions** column of the **Configuration Files** table, click the **Edit** control of the `/opt/zenoss/etc/zope.conf` file.
- 5 In the **Edit Configuration** dialog, determine whether the `zserver-threads` key is set.
Typically, the `zserver-threads` key is about 20% past the beginning of the file.
 - If the `zserver-threads` key is set and the value is 1, perform the remaining steps of this procedure.
 - If the `zserver-threads` key is set and the value is 4, no additional steps are required.
 - If the `zserver-threads` key is not set, the value defaults to 4, and no additional steps are required.
- 6 Change the value of the `zserver-threads` key from 1 to 4.
- 7 In the **Edit Configuration** dialog, click the **Save** button.

Updating the daily maintenance script

The `/etc/cron.daily/serviced` script installed with previous releases of Control Center invokes the `/opt/serviced/bin/serviced-container-cleanup` script to maintain the Btrfs file system of /

`var/lib/docker`. If `/etc/cron.daily/serviced` includes the cleanup script invocation, remove the invocation.

Creating a weekly maintenance script

The Resource Manager databases require regular maintenance to perform optimally. This procedure creates a script for cron to run once a week, to perform the required maintenance.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Create a shell script for cron to invoke.
 - a Open `/etc/cron.weekly/serviced` with a text editor.
The file is empty.
 - b Add the following content to the file.

```
#!/bin/sh
/bin/serviced service run zope zenosdbpack
```

- c Save the file, and then close the text editor.
- 3 Set file permissions.

```
chmod 0755 /etc/cron.weekly/serviced
```

Note The command fails when Resource Manager is stopped. When a failure occurs, start Resource Manager, and then run the command manually.

Deleting the pre-upgrade snapshot

Before the Resource Manager upgrade begins, the upgrade script creates and tags a snapshot of the system. Tagged snapshots persist until they are explicitly removed. When you are satisfied the new release is working properly, delete the pre-upgrade snapshot.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Display a list of all Control Center snapshots, with their tags.

```
serviced snapshot list -t
```

Example result:

Snapshot	Description	Tags
xm5mtezbyo2_20160211-220535.480		preupgrade-resmgr-5.1.2

- 3 Delete the pre-upgrade snapshot.

Replace *Snapshot-ID* with the identifier of the pre-upgrade snapshot returned in the previous step:

```
serviced snapshot remove Snapshot-ID
```

Deleting the pre-upgrade registry

This release of Control Center includes a new registry for application images. When you are satisfied the new release is working properly, delete the old registry.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.

- 2 Change directory to the Control Center registry location.

```
cd /opt/serviced/var/isvcs/docker-registry
```

- 3 Check the sizes of the two registries.

```
du -sh ./*
```

Each registry contains about 2GB or 3GB of images.

- 4 Delete the pre-upgrade registry.

```
rm -rf ./registry
```

A

Using Zenoss Toolbox

This appendix describes how to install and use Zenoss Toolbox.

Zenoss Toolbox tools

The Zenoss Toolbox tools examine key Resource Manager components for common issues affecting data integrity. Zenoss recommends running the following tools, in order, before upgrading Resource Manager:

- 1 The `zodbscan` tool quickly scans the Zope Object Database (ZODB) to provide a preliminary indication of the health of the database, and to determine whether the database needs to be compressed with `zenosdbpack` before upgrading.
- 2 The `findposkeyerror` tool checks objects and their relationships, and provides options for fixing errors.
- 3 The `zenrelationscan` tool checks only ZenRelations between objects.
- 4 The `zencatalogscan` tool checks ZODB object catalogs, which speed up web interface access.

The tools are run inside a Zope container, and the log files for each command are found in `$ZENHOME/log/toolbox`.

Downloading Zenoss Toolbox with internet access

This procedure describes how to download Zenoss Toolbox to a Control Center master host that has internet access.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Create a temporary directory, and change the current working directory to the temporary directory. The directory must be local (not mounted).

```
mkdir /tmp/toolbox && cd /tmp/toolbox
```

- 3 Download Zenoss Toolbox.

```
myUrl=https://github.com/zenoss/zenoss.toolbox/archive/master.zip
curl -sL --insecure -o master.zip $myUrl
```

- 4 Change the directory and file permissions.

The directory and file must be readable, writable, and executable by all users.

```
chmod -R 777 /tmp/toolbox
```

Downloading Zenoss Toolbox without internet access

This procedure downloads Zenoss Toolbox to a Control Center master host that does not have internet access.

- 1 Log onto a system that has internet access.
- 2 Start a web browser, and then navigate to [the Zenoss Toolbox releases page](#).
- 3 Download the latest version of the Zenoss Toolbox source code ZIP file.
The name of the file is `zenoss.toolbox-Version.zip`.
- 4 Use your operating system to rename the file to `master.zip`.
- 5 Use a file transfer utility such as [WinSCP](#) to copy the file to the Control Center master host.
- 6 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 7 Create a temporary directory, and change the current working directory to the temporary directory.
The directory must be local (not mounted).

```
mkdir /tmp/toolbox && cd /tmp/toolbox
```

- 8 Copy the Zenoss Toolbox ZIP file to the temporary directory.

Replace *Path-to-File* with the location of the `master.zip` file.

```
cp Path-to-File /tmp/toolbox
```

- 9 Change the directory and file permissions.
The directory and file must be readable, writable, and executable by all users.

```
chmod -R 777 /tmp/toolbox
```

Installing Zenoss Toolbox

This procedure describes how to install Zenoss Toolbox for use in Control Center Zope containers.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Start a shell as the `zenoss` user in a Zope container.
 - a Change directory to the temporary location of the Zenoss Toolbox `master.zip` file.

```
cd /tmp/toolbox
```

- b Start an interactive shell in a Zope container and save a snapshot named `InstallZenossToolbox`.

```
mySnap=InstallZenossToolbox
serviced service shell -i -s $mySnap zope bash
```

- c Switch user to `zenoss`.

```
su - zenoss
```

- 3 Install Zenoss Toolbox, and then exit the container.

- a Install Zenoss Toolbox.

```
easy_install /mnt/pwd/master.zip
```

- b Exit the zenoss user account.

```
exit
```

- c Exit the Zope container.

```
exit
```

- 4 Commit the named snapshot.

```
serviced snapshot commit $mySnap
```

- 5 Restart the Zope service.

```
serviced service restart zope
```

Running Zenoss Toolbox tools

- 1 Log in to the Control Center master host as a user with serviced CLI privileges.
- 2 Start an interactive session in a Zope container.

```
serviced service attach zope/0
```

- 3 Switch user to zenoss.

```
su - zenoss
```

- 4 Run the Zenoss Toolbox tools, in order.
For more information about the tools, see [Zenoss Toolbox tools](#) on page 163.
- 5 Exit the zenoss user account.

```
exit
```

- 6 Exit the Zope container.

```
exit
```