



Control Center Upgrade Guide for High-Availability Deployments

Release 1.5.1

Zenoss, Inc.

www.zenoss.com

Control Center Upgrade Guide for High-Availability Deployments

Copyright © 2018 Zenoss, Inc. All rights reserved.

Zenoss, Own IT, 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.

Linux is a registered trademark of Linus Torvalds.

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

Part Number: 1902.18.158.34

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

Contents

About this guide.....	8
Supported clients and browsers.....	8
Related publications.....	8
Change history.....	9
Chapter 1: Documented upgrade paths.....	10
Release dates.....	10
Upgrade paths included in this document.....	12
Chapter 2: Downloading and staging required files.....	13
Downloading Control Center files.....	13
Installing the repository mirror.....	14
Staging the Pacemaker resource agents package file.....	16
Staging Docker image files.....	16
Staging a Docker image file on ZooKeeper ensemble nodes.....	16
Part I: Upgrading 1.5.0 to 1.5.1.....	18
Scope.....	18
Chapter 3: Before upgrading 1.5.0 to 1.5.1.....	19
New features that affect upgrades.....	19
Upgrade best practices.....	19
Chapter 4: Stopping a Control Center deployment.....	21
Stopping a master host node.....	21
Stopping a delegate host.....	22
Chapter 5: Upgrading master hosts from 1.5.0 to 1.5.1.....	25
Updating Control Center on master hosts.....	25
Chapter 6: Upgrading delegates from 1.5.0 to 1.5.1.....	27
Updating Control Center on delegate hosts.....	27
Chapter 7: After upgrading.....	29
Setting the connection timeout of a resource pool.....	29
Removing unused images.....	29
Part II: Upgrading 1.4.x to 1.5.1.....	31
Scope.....	31

Chapter 8: Before upgrading 1.4.x to 1.5.1.....	32
New features that affect upgrades.....	32
Upgrade best practices.....	32
Chapter 9: Stopping a Control Center deployment.....	34
Stopping a master host node.....	34
Stopping a delegate host.....	35
Chapter 10: Upgrading master hosts from 1.4.x to 1.5.1.....	38
Updating Docker.....	38
Loading image files.....	40
Updating the Pacemaker resource agents.....	40
Updating Control Center on master hosts.....	40
Chapter 11: Upgrading delegates from 1.4.x to 1.5.1.....	42
Updating Docker.....	42
Configuring NFS 4.0.....	44
Updating Control Center on delegate hosts.....	44
Updating the ZooKeeper image on ensemble nodes.....	45
Chapter 12: After upgrading.....	46
Setting the connection timeout of a resource pool.....	46
Removing unused images.....	46
Part III: Upgrading 1.3.x to 1.5.1.....	48
Scope.....	48
Chapter 13: Before upgrading 1.3.x to 1.5.1.....	49
New features that affect upgrades from release 1.3.x.....	49
Upgrade best practices.....	50
Chapter 14: Stopping a Control Center deployment.....	51
Stopping a master host node.....	51
Stopping a delegate host.....	52
Chapter 15: Upgrading master hosts from 1.3.x to 1.5.1.....	55
Adding storage for audit logging.....	55
Updating Docker.....	55
Loading image files.....	57
Updating the Pacemaker resource agents.....	58
Updating Control Center on master hosts.....	58
Chapter 16: Upgrading delegates from 1.3.x to 1.5.1.....	59

Updating Docker.....	59
Configuring NFS 4.0.....	61
Updating Control Center on delegate hosts.....	61
Updating the ZooKeeper image on ensemble nodes.....	62
Chapter 17: After upgrading.....	63
Setting the connection timeout of a resource pool.....	63
Removing unused images.....	63
Part IV: Upgrading 1.2.x to 1.5.1.....	65
Scope.....	65
Chapter 18: Before upgrading 1.2.x to 1.5.1.....	66
New features that affect upgrades from release 1.2.x.....	66
Upgrade best practices.....	67
Chapter 19: Stopping a Control Center deployment.....	68
Stopping a master host node.....	68
Stopping a delegate host.....	69
Chapter 20: Upgrading master hosts from 1.2.x to 1.5.1.....	72
Adding storage for audit logging.....	72
Updating Docker.....	72
Loading image files.....	74
Updating the Pacemaker resource agents.....	74
Updating Control Center on master hosts.....	75
Chapter 21: Upgrading delegates from 1.2.x to 1.5.1.....	76
Updating Docker.....	76
Configuring NFS 4.0.....	78
Updating the ZooKeeper image on ensemble nodes.....	78
Updating Control Center on delegate hosts.....	79
Chapter 22: After upgrading from 1.2.x to 1.5.1.....	80
Setting the connection timeout of a resource pool.....	80
Removing unused images.....	80
Part V: Upgrading 1.1.x to 1.5.1.....	82
Scope.....	82
Chapter 23: Before upgrading 1.1.x to 1.5.1.....	83
New features that affect upgrades from 1.1.x.....	83
Upgrade best practices.....	84

Chapter 24: Stopping a Control Center deployment.....	86
Stopping a master host node.....	86
Stopping a delegate host.....	87
Chapter 25: Upgrading master hosts from 1.1.x to 1.5.1.....	90
Adding storage for audit logging.....	90
Updating to Docker CE.....	90
Configuring Docker.....	92
Loading image files.....	94
Updating the Pacemaker resource agents.....	95
Updating Control Center on the master host.....	95
Chapter 26: Upgrading delegates from 1.1.x to 1.5.1.....	97
Updating to Docker CE.....	97
Configuring Docker.....	99
Configuring NFS 4.0.....	101
Updating Control Center on delegate hosts.....	101
Updating the ZooKeeper image on ensemble nodes.....	102
Chapter 27: Starting an upgraded deployment.....	104
Registering the master host nodes.....	104
Updating delegate hosts with authentication.....	106
Chapter 28: After upgrading from 1.1.x to 1.5.1.....	109
Updating resource pool permissions.....	109
Setting the connection timeout of a resource pool.....	110
Removing unused images.....	110
Removing orphaned snapshot devices.....	111
Updating the OpenTSDB time-to-live value.....	112
Managing maintenance scripts.....	113
Appendix A: Control Center application data storage requirements.....	115
Examining application data storage status.....	115
About adding space to a tenant volume.....	125
Appendix B: Upgrading only the Pacemaker resource agents.....	126
Identifying the Pacemaker resource agents version.....	126
Identifying the available version of the resource agents.....	126
Updating the Pacemaker resource agents.....	127
Appendix C: Storage management utility.....	128
serviced-storage.....	128
Appendix D: Starting and stopping Control Center deployments.....	132
Stopping Control Center.....	132

Starting Control Center.....	135
Appendix E: Resolving package dependency conflicts.....	137
Resolving device mapper dependency conflicts.....	137
Resolving other dependency conflicts.....	138
Appendix F: User access control.....	140
Adding users to the default browser interface access group.....	140
Configuring a regular group as the Control Center browser interface access group.....	141
Enabling use of the command-line interface.....	142
Appendix G: Updating the cluster management software.....	143
Downloading and staging cluster software.....	143
Updating cluster software.....	144
Appendix H: Control Center configuration variables.....	146
Removed variables (at version 1.2.0).....	146
New variables (at version 1.2.0).....	147
Master host configuration variables.....	148
Delegate host configuration variables.....	151
Universal configuration variables.....	153
Best practices for configuration files.....	155
Control Center configuration file.....	155

About this guide

The *Control Center Upgrade Guide for High-Availability Deployments* provides detailed instructions for upgrading a high-availability deployment of Control Center. To perform the procedures in this guide, users must be able to download packages from the Zenoss portal site, which is available only to Zenoss customers.

Supported clients and browsers

The following table identifies the supported combinations of client operating systems and web browsers.

Client OS	Tested browsers
Windows 7, 10	Internet Explorer 11*
	Firefox 56 and later
	Chrome 61 and later
macOS 10.12.3, 10.13	Firefox 56 and later
	Chrome 61 and later
Ubuntu 14.04 LTS	Firefox 56 and later
	Chrome 61 and later

Related publications

Title	Description
<i>Control Center Release Notes</i>	Describes known issues, fixed issues, and late-breaking information not included in other publications.
<i>Control Center Installation Guide</i>	Provides detailed procedures for installing and configuring Control Center.
<i>Control Center Installation Guide for High-Availability Deployments</i>	Provides detailed procedures for installing and configuring Control Center in a high-availability deployment.
<i>Control Center Reference Guide</i>	Provides information and procedures for managing Control Center. This information is also available as online help in the Control Center browser interface.
<i>Control Center Upgrade Guide</i>	Provides detailed procedures for updating a Control Center deployment to the latest release.
<i>Control Center Upgrade Guide for High-Availability Deployments</i>	Provides detailed procedures for updating a high-availability deployment of Control Center to the latest release.

Documentation feedback

To provide feedback about this document, or to report an error or omission, please send an email to docs@controlcenter.io. In the email, please include the document title (*Control Center Upgrade Guide for High-Availability Deployments*) and part number (1902.18.158.34) and as much information as possible about the context of your feedback.

* Enterprise mode only; compatibility mode is not tested.

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 *Control Center Release Notes*.

1902.18.158.34 (1.5.1)

Add chapters for updating 1.5.0 deployments.

Add procedure for setting the connection timeout to resource pools to the post-upgrade chapter of each upgrade path.

1902.17.349 (1.5.0)

Specify Docker version in `yum install` command.

1902.17.331 (1.5.0)

Replace Leapfile.net with `delivery.zenoss.com`.

Add steps for importing the Zenoss GPG key.

Add options to the DFS resize section (*Control Center application data storage requirements* on page 115).

1902.17.311 (1.5.0)

Add chapters for updating 1.4.x deployments.

Add step for updating RHEL/CentOS to 7.4, if necessary.

Add steps for a delegate host workaround of an NFS issue on RHEL/CentOS 7.4.

1302.17.268

Update release number (1.4.1).

1302.17.235

Update release number (1.4.0).

Add a new storage requirement, for audit logs.

Remove the `-y` parameter from all `yum` command invocations.

1302.17.206

Update the `SERVICED_DOCKER_REGISTRY` configuration steps to ensure the correct value is set.

1302.17.172

Update release number (1.3.3).

Replace Docker 1.12.1 with Docker CE 17.03.1.

Remove step for disabling SELinux.

1302.17.122

Update release number (1.3.2).

1302.17.100

New document for release 1.3.1.

1

Documented upgrade paths

This chapter includes the dates of Control Center releases and the upgrade paths that are documented in this guide.

Release dates

Table 1: Release 1.5

Release	Date
Control Center 1.5.1	12 Jun 2018
Control Center 1.5.0	08 Nov 2017

Table 2: Release 1.4

Release	Date
Control Center 1.4.1	25 Sep 2017
Control Center 1.4.0	17 Aug 2017

Table 3: Release 1.3

Release	Date
Control Center 1.3.3	20 Jun 2017
Control Center 1.3.2	03 May 2017
Control Center 1.3.1	13 Apr 2017
Control Center 1.3.0	09 Mar 2017

Table 4: Release 1.2

Release	Date
Control Center 1.2.3	27 Feb 2017
Control Center 1.2.2	25 Jan 2017

Release	Date
Control Center 1.2.1	16 Dec 2016
Control Center 1.2.0	14 Nov 2016

Table 5: Release 1.1

Release	Date
Control Center 1.1.10	23 Nov 2016
Control Center 1.1.9	17 Oct 2016
Control Center 1.1.8	20 Sep 2016
Control Center 1.1.7	20 Jul 2016
Control Center 1.1.6	28 Jun 2016
Control Center 1.1.5	01 Jun 2016
Control Center 1.1.4	24 May 2016
Control Center 1.1.3	20 Apr 2016
Control Center 1.1.2	04 Mar 2016
Control Center 1.1.1	29 Feb 2016

Table 6: Release 1.0

Release	Date
Control Center 1.0.10	20 Feb 2016
Control Center 1.0.9	02 Dec 2015
Control Center 1.0.8	16 Nov 2015
Control Center 1.0.7	10 Oct 2015
Control Center 1.0.6	14 Sep 2015
Control Center 1.0.5	05 Aug 2015
Control Center 1.0.4	10 Jul 2015
Control Center 1.0.3	27 May 2015
Control Center 1.0.2	20 Apr 2015
Control Center 1.0.1	03 Apr 2015
Control Center 1.0.0	24 Feb 2015

Upgrade paths included in this document

Release 1.5.x upgrades

From	To
Control Center 1.5.0	Control Center 1.5.1

Release 1.4.x upgrades

From	To
Control Center 1.4.1	Control Center 1.5.1
Control Center 1.4.0	Control Center 1.5.1

Release 1.3.x upgrades

From	To
Control Center 1.3.3	Control Center 1.5.1
Control Center 1.3.2	Control Center 1.5.1
Control Center 1.3.1	Control Center 1.5.1
Control Center 1.3.0	Control Center 1.5.1

Release 1.2.x upgrades

From	To
Control Center 1.2.3	Control Center 1.5.1
Control Center 1.2.2	Control Center 1.5.1
Control Center 1.2.1	Control Center 1.5.1
Control Center 1.2.0	Control Center 1.5.1

Release 1.1.x upgrades

From	To
Control Center 1.1.10	Control Center 1.5.1
Control Center 1.1.9	Control Center 1.5.1
Control Center 1.1.8	Control Center 1.5.1
Control Center 1.1.7	Control Center 1.5.1
Control Center 1.1.5	Control Center 1.5.1
Control Center 1.1.4	Control Center 1.5.1
Control Center 1.1.3	Control Center 1.5.1
Control Center 1.1.2	Control Center 1.5.1
Control Center 1.1.1	Control Center 1.5.1

2

Downloading and staging required files

This chapter describes how to download and install or stage Control Center software and its operating system dependencies. The procedures in this chapter are required to perform an upgrade.

The following table identifies where to perform each procedure in this chapter.

Procedure	Where to perform
<i>Downloading Control Center files</i> on page 13	A workstation with internet access
<i>Installing the repository mirror</i> on page 14	All Control Center hosts
<i>Staging the Pacemaker resource agents package file</i> on page 16	Both master host nodes
<i>Staging Docker image files</i> on page 16	Both master host nodes
<i>Staging a Docker image file on ZooKeeper ensemble nodes</i> on page 16	Delegate hosts that are ZooKeeper ensemble nodes

Downloading Control Center files

To perform this procedure, you need:

- A workstation with internet access.
- Permission to download files from delivery.zenoss.com. Customers can request permission by filing a ticket at the [Zenoss Support](#) site.
- A secure network copy program.

Use this procedure to

- download the required files to a workstation
- copy the files to the hosts that need them

Perform these steps:

- 1 In a web browser, navigate to the download site, and then log in.
The download site is delivery.zenoss.com.
- 2 Download the self-installing Docker image files.

```
install-zenoss-serviced-isvcs-v61.run
```

```
install-zenoss-isvcs-zookeeper-v10.run
```

- 3 Download the Control Center RPM file.

```
serviced-1.5.1-1.x86_64.rpm
```

- 4 Identify the operating system release on Control Center hosts.

Enter the following command on each Control Center host in your deployment, if necessary. All Control Center hosts should be running the same operating system release and kernel.

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

- 5 Download the RHEL/CentOS repository mirror file for your deployment.

The download site provides a repository mirror file for each tested release of RHEL/CentOS. Each mirror file contains the release-specific packages that Control Center requires.

```
yum-mirror-centos7.2-1511-serviced-1.5.1.x86_64.rpm
```

```
yum-mirror-centos7.3-1611-serviced-1.5.1.x86_64.rpm
```

```
yum-mirror-centos7.4-1708-serviced-1.5.1.x86_64.rpm
```

- 6 Download the Pacemaker resource agents for Control Center.

```
serviced-resource-agents-1.1.0-1.x86_64.rpm
```

- 7 Optional: Download the Zenoss GNU Privacy Guard (GPG) key, if desired.

You can use the Zenoss GPG key to verify Zenoss RPM files and the yum metadata of the repository mirror.

- a Download the key.

```
curl --location -o /tmp/tmp.html \
'http://keys.gnupg.net/pks/lookup?op=get&search=0xED0A5FD2AA5A1AD7'
```

- b Determine whether the download succeeded.

```
grep -Ec '^\\-\\-\\-\\-BEGIN PGP' /tmp/tmp.html
```

- If the result is 0, return to the previous substep.
- If the result is 1, proceed to the next substep.

- c Extract the key.

```
awk '/^-----BEGIN PGP.*$/,/^-----END PGP.*$/' \
/tmp/tmp.html > ./RPM-GPG-KEY-Zenoss
```

- 8 Use a secure copy program to copy the files to Control Center hosts.

- Copy all files to both master nodes.
- Copy the RHEL/CentOS RPM file, the Control Center RPM file, and the Zenoss GPG key file to all delegate hosts.
- Copy the Docker image file for ZooKeeper to delegate hosts that are ZooKeeper ensemble nodes.

Installing the repository mirror

Use this procedure to install the Zenoss repository mirror on a Control Center host. Repeat this procedure on each host in your deployment.

- 1 Log in to the target host as `root`, or as a user with superuser privileges.
- 2 Move the RPM files and the Zenoss GPG key file to `/tmp`.

3 Optional: Remove the existing repository mirror, if necessary.

- a**
- Search for the mirror.

```
yum list --disablerepo=* | awk '/^yum-mirror/ { print $1}'
```

- b**
- Remove the mirror.

Replace *Old-Mirror* with the name of the Zenoss repository mirror returned in the previous substep:

```
yum remove Old-Mirror
```

4 Install the repository mirror.

```
yum install /tmp/yum-mirror-*.rpm
```

The yum command copies the contents of the RPM file to /opt/zenoss-repo-mirror.

5 Optional: Install the Zenoss GPG key, and then test the package files, if desired.

- a**
- Move the Zenoss GPG key to the mirror directory.

```
mv /tmp/RPM-GPG-KEY-Zenoss /opt/zenoss-repo-mirror
```

- b**
- Install the key.

```
rpm --import /opt/zenoss-repo-mirror/RPM-GPG-KEY-Zenoss
```

- c**
- Test the repository mirror package file.

```
rpm -K /tmp/yum-mirror-*.rpm
```

On success, the result includes the file name and the following information:

```
(sha1) dsa sha1 md5 gpg OK
```

- d**
- Test the Control Center package file.

```
rpm -K /tmp/serviced-1.5.1-1.x86_64.rpm
```

6 Optional: Update the configuration file of the Zenoss repository mirror to enable GPG key verification, if desired.

- a**
- Open the repository mirror configuration file (/etc/yum.repos.d/zenoss-mirror.repo) with a text editor, and then add the following lines to the end of the file.

```
repo_gpgcheck=1
gpgkey=file:///opt/zenoss-repo-mirror/RPM-GPG-KEY-Zenoss
```

- b**
- Save the file, and then close the editor.

- c**
- Update the yum metadata cache.

```
yum makecache fast
```

The cache update process includes the following prompt:

```
Retrieving key from file:///opt/zenoss-repo-mirror/RPM-GPG-KEY-
Zenoss
Importing GPG key 0xAA5A1AD7:
Userid      : "Zenoss, Inc. <dev@zenoss.com>"
Fingerprint: f31f fd84 6a23 b3d5 981d a728 ed0a 5fd2 aa5a 1ad7
```

```
From      : /opt/zenoss-repo-mirror/RPM-GPG-KEY-Zenoss
Is this ok [y/N]:
```

Enter y.

- 7 Move the Control Center package file to the mirror directory.

```
mv /tmp/serviced-1.5.1-1.x86_64.rpm /opt/zenoss-repo-mirror
```

- 8 Optional: Delete the mirror package file, if desired.

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

Staging the Pacemaker resource agents package file

Use this procedure to stage the RPM file for the Pacemaker resource agents for Control Center on the master host nodes.

- 1 Log in to a master host node as `root`, or as a user with superuser privileges.
- 2 In a separate terminal session, log in to the other master host node as `root`, or as a user with superuser privileges.
- 3 Create a directory for the RPM file, if necessary.

The directory is used for the Zenoss repository mirror.

```
test -d /opt/zenoss-repo-mirror || mkdir -p /opt/zenoss-repo-mirror
```

- 4 Move the Pacemaker resource agents RPM file to the mirror directory.

```
mv /tmp/serviced-resource-agents-1.1.0-1.x86_64.rpm \
/opt/zenoss-repo-mirror
```

Staging Docker image files

Before performing this procedure, verify that approximately 640MB of temporary space is available on the file system where `/root` is located.

Use this procedure to copy Docker image files to a Control Center host. The files are used when Docker is fully configured.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Copy or move the archive files to `/root`.
- 3 Add execute permission to the files.

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

Staging a Docker image file on ZooKeeper ensemble nodes

Before performing this procedure, verify that approximately 170MB of temporary space is available on the file system where `/root` is located.

Use this procedure to add a Docker image file to the Control Center delegate hosts that are ZooKeeper ensemble nodes. Delegate hosts that are not ZooKeeper ensemble nodes do not need the file.

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

- 2 Copy or move the `install-zenoss-isvcs-zookeeper-v10.run` file to `/root`.
- 3 Add execute permission to the file.

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

Upgrading 1.5.0 to 1.5.1 Scope

The chapters in this part provide instructions for performing an upgrade of a high-availability deployment of Control Center, from version 1.5.0 to 1.5.1.

The following table identifies the upgrades of Control Center that are documented in this part.

From	To
Control Center 1.5.0	Control Center 1.5.1

3

Before upgrading 1.5.0 to 1.5.1

This chapter provides information and procedures to prepare your high-availability Control Center deployment for an upgrade.

New features that affect upgrades

This release includes no new features or requirements that affect the upgrade process.

Note The Docker image files for Control Center are unchanged for this release. You only need to download and install the new `serviced` RPM package and the operating system mirror package.

For more information about this release, refer to the *Control Center Release Notes*.

Upgrade best practices

The following list outlines recommended best practices for upgrading high-availability Control Center deployments:

- 1 Download a copy of the *Control Center Release Notes* for this release and review its contents. The latest information is included in that document.
- 2 Upgrade only what needs upgrading. For example, you can update some components independently. For more information, see the following appendices:
 - [Upgrading only the Pacemaker resource agents](#) on page 126
 - [Updating the cluster management software](#) on page 143
- 3 Compare the Docker images that accompany this release with the images that accompany the installed release, and determine whether the image files need to be downloaded and installed. For more information, see [Releases and image tags](#).
- 4 The contents of `/etc/default/serviced` on the master nodes must be identical. Use a utility like `sum` to compare the files quickly.
- 5 On delegate hosts, most of the upgrade steps are identical. Use `screen`, `tmux`, or a similar program to establish sessions on each delegate host and perform the steps at the same time.
- 6 Review and verify the settings in delegate host configuration files (`/etc/default/serviced`) before starting the upgrade. Ideally, the settings on all delegate hosts are identical, except on ZooKeeper nodes and delegate hosts that do not mount the DFS.
- 7 Review the procedures in this guide before performing them. Every effort is made to avoid mistakes and anticipate needs; nevertheless, the instructions may be incorrect or inadequate for some requirements or environments.

- 8 Download and stage the required files for your upgrade. For more information, see [Downloading and staging required files](#) on page 13.

4

Stopping a Control Center deployment

This chapter includes procedures for stopping a high-availability Control Center deployment.

Note The procedures in this chapter assume that Control Center is the only source of Docker containers that are run on Control Center hosts.

Perform the procedures in this chapter in order.

Stopping a master host node

Use this procedure to stop the Control Center service (*serviced*) on the master host in a high-availability deployment.

- 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 step.

- 3 Stop the top-level service *serviced* is managing, if necessary.
 - a Show the status of running services.

```
serviced service status
```

The top-level service is the service listed immediately below the headings line.

- If the status of the top-level service and all child services is `stopped`, proceed to the next step.
 - If the status of the top-level service and all child services is **not** `stopped`, perform the remaining substeps.
- b Stop the top-level service.
Replace *Service* with the name or identifier of the top-level service:

```
serviced service stop Service
```

- c Monitor the stop.

```
serviced service status
```

When the status of the top-level service and all child services is stopped, proceed to the next step.

- 4 Stop Control Center with the cluster management tool.

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

- 5 Monitor the status of cluster resources.

```
watch pcs status
```

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

- 6 Ensure that no containers remain in the local repository.

- a Start the Docker service.

```
systemctl start docker
```

- b Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

If the command returns a result, enter the following command:

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- c Stop the Docker service.

```
systemctl stop docker
```

- 7 To ensure that no containers remain in both Docker repositories, log in to the other master node as `root`, or as a user with superuser privileges, and then perform the preceding step.

Stopping a delegate host

Use this procedure to stop the Control Center service (`serviced`) on a delegate host in a multi-host deployment. Repeat this procedure on each delegate host in your deployment.

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

```
systemctl stop serviced
```

- 3 Ensure that no containers remain in the local repository.

- a Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, proceed to the next step.
- If the command returns a result, perform the following substeps.

- b Remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, the most likely cause is an NFS conflict. Perform the following substeps.

- c** Stop the NFS and Docker services.

```
systemctl stop nfs && systemctl stop docker
```

- d** Start the NFS and Docker services.

```
systemctl start nfs && systemctl start docker
```

- e** Repeat the attempt to remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, perform the remaining substeps.

- f** Disable the automatic startup of `serviced`.

```
systemctl disable serviced
```

- g** Reboot the host.

```
reboot
```

- h** Log in to the delegate host as `root`, or as a user with superuser privileges.

- i** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

4 Dismount all filesystems mounted from the Control Center master host.

This step ensures no stale mounts remain when the storage on the master host is replaced.

- a** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- b** Force the filesystems to dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
  umount -f $FS
done
```

- c** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- d** Perform a lazy dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
    umount -f -l $FS
done
```

- e** Restart the NFS service.

```
systemctl restart nfs
```

- f** Determine whether any filesystems remain mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the remaining substeps.

- g** Disable the automatic startup of serviced.

```
systemctl disable serviced
```

- h** Reboot the host.

```
reboot
```

- i** Log in to the delegate host as `root`, or as a user with superuser privileges.

- j** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```


5

Upgrading master hosts from 1.5.0 to 1.5.1

Perform the procedures in this chapter to upgrade both Control Center 1.5.0 master hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Note The Docker image files for Control Center are unchanged for this release. You only need to download and install the new `serviced` RPM package and the operating system mirror package.

Updating Control Center on master hosts

Use this procedure to update Control Center on both master hosts to version 1.5.1.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.
 - a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \  
/opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.
 - a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

For more information about configuring the master host, see [Control Center configuration variables](#) on page 146.

6

Upgrading delegates from 1.5.0 to 1.5.1

Perform the procedures in this chapter to upgrade Control Center 1.5.0 delegate hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Note The Docker image files for Control Center are unchanged for this release. You only need to download and install the new `serviced` RPM package and the operating system mirror package.

Updating Control Center on delegate hosts

This procedure updates Control Center on delegate hosts to version 1.5.1.

Perform this procedure on each delegate host in your deployment.

- 1 Log in to a delegate host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.
 - a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \
/opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.
 - a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.
For more information about configuring a delegate host, see [Control Center configuration variables](#) on page 146.

When all delegate hosts are updated, start the cluster. For more information, see [Starting and stopping Control Center deployments](#) on page 132.

7

After upgrading

Perform the procedures in this chapter after Control Center is upgraded.

Setting the connection timeout of a resource pool

Use this procedure to set the length of time the services scheduler waits for a disconnected delegate host to reconnect before moving its services to a different host in the resource pool. This affects all multi-host resource pools. Zenoss recommends using a minimum value of 15 seconds for all multi-host pools. For resource pools that are connected through high-latency, wide-area networks, Zenoss recommends values greater than 15 seconds.

- 1 Log in to the master host as a user with `serviced` CLI privileges.
- 2 Display the list of resource pools and their connection timeout values.

```
serviced pool list -v | grep -E 'ID|ConnectionTimeout'
```

- 3 Optional: Set the connection timeout value of a resource pool.

This command accepts the following units identifiers:

- ms (milliseconds)
- s (seconds)
- m (minutes)
- h (hours)

Replace *Pool-ID* with a resource pool identifier, and replace *Timeout+Units* with an integer followed by a units identifier:

```
serviced pool set-conn-timeout Pool-ID Timeout+Units
```

Removing unused images

Use this procedure to identify and remove unused Control Center images.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the images associated with the installed version of `serviced`.

```
serviced version | grep Images
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v61 zenoss/isvcs-zookeeper:v10]
```

3 Start Docker.

```
systemctl start docker
```

4 Display the serviced images in the local repository.

```
docker images | awk '/REPO|isvcs/'
```

Example result (edited to fit):

REPOSITORY	TAG	IMAGE ID
zenoss/serviced-isvcs	v40	88cd6c24cc82
zenoss/serviced-isvcs	v61	0aab5a2123f2
zenoss/isvcs-zookeeper	v3	46fa0a2fc4bf
zenoss/isvcs-zookeeper	v10	0ff3b3117fb8

The example result shows the current versions and one set of previous versions. Your result may include additional previous versions and will show different images IDs.

5 Remove unused images.

Replace *Image-ID* with the image ID of an image for a previous version.

```
docker rmi Image-ID
```

Repeat this command for each unused image.

6 Stop Docker.

```
systemctl stop docker
```

Upgrading 1.4.x to 1.5.1

Scope

The chapters in this part provide instructions for performing an upgrade of a high-availability deployment of Control Center, from version 1.4.x to 1.5.1.

The following table identifies the upgrades of Control Center that are documented in this part.

From	To
Control Center 1.4.1	Control Center 1.5.1
Control Center 1.4.0	Control Center 1.5.1

Before upgrading 1.4.x to 1.5.1

This chapter provides information and procedures to prepare your high-availability Control Center deployment for an upgrade.

New features that affect upgrades

This release includes new features and new requirements that affect the upgrade process. The following list provides an overview of the changes that are addressed during this upgrade.

- RHEL/CentOS 7.4 is added; RHEL/CentOS 7.1 is withdrawn. The upgrade process includes a step for updating the operating system.
- On RHEL/CentOS 7.4, there may be a file locking defect in NFS 4.1. To avoid the issue, delegate hosts are configured to use NFS 4.0.
- The `serviced-zenossdbpack` maintenance script is moved from `/etc/cron.weekly` to `/opt/serviced/bin` and a new cron job, `/etc/cron.d/cron_zenossdbpack`, is installed to invoke the maintenance script. By default, the script runs every Sunday at midnight.
- This release replaces Docker Community Edition (CE) 17.03.1 with Docker CE 17.09.0.

For more information about this release, refer to the *Control Center Release Notes*.

Upgrade best practices

The following list outlines recommended best practices for upgrading high-availability Control Center deployments:

- 1 Download a copy of the *Control Center Release Notes* for this release and review its contents. The latest information is included in that document.
- 2 Upgrade only what needs upgrading. For example, you can update some components independently. For more information, see the following appendices:
 - [Upgrading only the Pacemaker resource agents](#) on page 126
 - [Updating the cluster management software](#) on page 143
- 3 Compare the Docker images that accompany this release with the images that accompany the installed release, and determine whether the image files need to be downloaded and installed. For more information, see [Releases and image tags](#).
- 4 The contents of `/etc/default/serviced` on the master nodes must be identical. Use a utility like `sum` to compare the files quickly.

- 5 On delegate hosts, most of the upgrade steps are identical. Use `screen`, `tmux`, or a similar program to establish sessions on each delegate host and perform the steps at the same time.
- 6 Review and verify the settings in delegate host configuration files (`/etc/default/serviced`) before starting the upgrade. Ideally, the settings on all delegate hosts are identical, except on ZooKeeper nodes and delegate hosts that do not mount the DFS.
- 7 Review the procedures in this guide before performing them. Every effort is made to avoid mistakes and anticipate needs; nevertheless, the instructions may be incorrect or inadequate for some requirements or environments.
- 8 Download and stage the required files for your upgrade. For more information, see [Downloading and staging required files](#) on page 13.

Stopping a Control Center deployment

This chapter includes procedures for stopping a high-availability Control Center deployment.

Note The procedures in this chapter assume that Control Center is the only source of Docker containers that are run on Control Center hosts.

Perform the procedures in this chapter in order.

Stopping a master host node

Use this procedure to stop the Control Center service (*serviced*) on the master host in a high-availability deployment.

- 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 step.

- 3 Stop the top-level service *serviced* is managing, if necessary.
 - a Show the status of running services.

```
serviced service status
```

The top-level service is the service listed immediately below the headings line.

- If the status of the top-level service and all child services is `stopped`, proceed to the next step.
- If the status of the top-level service and all child services is **not** `stopped`, perform the remaining substeps.

- b Stop the top-level service.

Replace *Service* with the name or identifier of the top-level service:

```
serviced service stop Service
```

- c Monitor the stop.

```
serviced service status
```

When the status of the top-level service and all child services is stopped, proceed to the next step.

- 4 Stop Control Center with the cluster management tool.

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

- 5 Monitor the status of cluster resources.

```
watch pcs status
```

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

- 6 Ensure that no containers remain in the local repository.

- a Start the Docker service.

```
systemctl start docker
```

- b Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

If the command returns a result, enter the following command:

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- c Stop the Docker service.

```
systemctl stop docker
```

- 7 To ensure that no containers remain in both Docker repositories, log in to the other master node as `root`, or as a user with superuser privileges, and then perform the preceding step.

Stopping a delegate host

Use this procedure to stop the Control Center service (`serviced`) on a delegate host in a multi-host deployment. Repeat this procedure on each delegate host in your deployment.

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

```
systemctl stop serviced
```

- 3 Ensure that no containers remain in the local repository.

- a Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, proceed to the next step.
- If the command returns a result, perform the following substeps.

- b Remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, the most likely cause is an NFS conflict. Perform the following substeps.

- c** Stop the NFS and Docker services.

```
systemctl stop nfs && systemctl stop docker
```

- d** Start the NFS and Docker services.

```
systemctl start nfs && systemctl start docker
```

- e** Repeat the attempt to remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, perform the remaining substeps.

- f** Disable the automatic startup of `serviced`.

```
systemctl disable serviced
```

- g** Reboot the host.

```
reboot
```

- h** Log in to the delegate host as `root`, or as a user with superuser privileges.

- i** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

4 Dismount all filesystems mounted from the Control Center master host.

This step ensures no stale mounts remain when the storage on the master host is replaced.

- a** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- b** Force the filesystems to dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
  umount -f $FS
done
```

- c** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- d** Perform a lazy dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
    umount -f -l $FS
done
```

- e** Restart the NFS service.

```
systemctl restart nfs
```

- f** Determine whether any filesystems remain mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the remaining substeps.

- g** Disable the automatic startup of serviced.

```
systemctl disable serviced
```

- h** Reboot the host.

```
reboot
```

- i** Log in to the delegate host as `root`, or as a user with superuser privileges.
- j** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

Upgrading master hosts from 1.4.x to 1.5.1

10

Perform the procedures in this chapter to upgrade both Control Center 1.4.x master hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Updating Docker

Use this procedure to update Docker to version 17.09.0.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Remove the Docker repository description file, if present.

Upgrades no longer require external repositories.

```
rm -f /etc/yum.repos.d/docker.repo
```

- 3 Update the operating system, if necessary.

- a Determine which release is installed.

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

- If the result is 7.2, 7.3, or 7.4, proceed to the next step.
- If the result is 7.1, perform the remaining substeps.

- b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

- c Update the operating system, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update && reboot
```

- d Log in as `root`, or as a user with superuser privileges.
- e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

- 4 Update the Linux kernel, if necessary.

- a** Determine which kernel version is installed.

```
uname -r
```

If the result is lower than 3.10.0-327.22.2.el7.x86_64, perform the following substeps.

- b** Disable automatic start of serviced.

```
systemctl disable serviced
```

- c** Update the kernel, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update kernel && reboot
```

- d** Log in as root, or as a user with superuser privileges.

- e** Enable automatic start of serviced.

```
systemctl enable serviced
```

- 5** Stop the Docker service.

```
systemctl stop docker
```

- 6** Remove Docker 17.03.1.

- a** Remove without checking dependencies.

```
rpm -e --nodeps docker-ce
```

- b** Clean the yum databases.

```
yum clean all
```

- 7** Install Docker CE 17.09.0.

```
yum install --enablerepo=zenoss-mirror docker-ce-17.09.0.ce
```

If yum returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 8** Start the Docker service.

```
systemctl start docker
```

- 9** Remove a previous workaround for an NFS rpcbind issue.

- a** Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

- b** Reload the systemd manager configuration.

```
systemctl daemon-reload
```

Loading image files

Use this procedure to load images into the local registry.

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

```
systemctl start docker
```

- 3 Change directory to `/root`.

```
cd /root
```

- 4 Load the images.

```
for image in install-zenoss-*.run
do
  /bin/echo -en "\nLoading $image..."
  yes | ./$image
done
```

- 5 List the images in the registry.

```
docker images
```

The result should show one image for each archive file.

- 6 Optional: Delete the archive files, if desired.

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

- 7 Stop the Docker service.

```
systemctl stop docker
```

Updating the Pacemaker resource agents

Use this procedure to update the Pacemaker resource agents for Control Center.

- 1 Log in to a master host node as `root`, or as a user with superuser privileges.
- 2 In a separate terminal session, log in to the other master host node as `root`, or as a user with superuser privileges.
- 3 On both nodes, install the latest version of the resource agents package.

```
myPath=/opt/zenoss-repo-mirror
yum install \
  $myPath/serviced-resource-agents-1.1.0-1.x86_64.rpm
```

Updating Control Center on master hosts

Use this procedure to update Control Center on both master hosts to version 1.5.1.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.

- a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \  
/opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.

- a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^\\b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

For more information about configuring the master host, see [Control Center configuration variables](#) on page 146.

Upgrading delegates from 1.4.x to 1.5.1

Perform the procedures in this chapter to upgrade Control Center 1.4.x delegate hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Updating Docker

Use this procedure to update Docker to version 17.09.0.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Remove the Docker repository description file, if present.

Upgrades no longer require external repositories.

```
rm -f /etc/yum.repos.d/docker.repo
```

- 3 Update the operating system, if necessary.

- a Determine which release is installed.

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

- If the result is 7.2, 7.3, or 7.4, proceed to the next step.
- If the result is 7.1, perform the remaining substeps.

- b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

- c Update the operating system, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update && reboot
```

- d Log in as `root`, or as a user with superuser privileges.

- e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

- 4 Update the Linux kernel, if necessary.

- a** Determine which kernel version is installed.

```
uname -r
```

If the result is lower than 3.10.0-327.22.2.el7.x86_64, perform the following substeps.

- b** Disable automatic start of serviced.

```
systemctl disable serviced
```

- c** Update the kernel, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update kernel && reboot
```

- d** Log in as root, or as a user with superuser privileges.

- e** Enable automatic start of serviced.

```
systemctl enable serviced
```

- 5** Stop the Docker service.

```
systemctl stop docker
```

- 6** Remove Docker 17.03.1.

- a** Remove without checking dependencies.

```
rpm -e --nodeps docker-ce
```

- b** Clean the yum databases.

```
yum clean all
```

- 7** Install Docker CE 17.09.0.

```
yum install --enablerepo=zenoss-mirror docker-ce-17.09.0.ce
```

If yum returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 8** Start the Docker service.

```
systemctl start docker
```

- 9** Remove a previous workaround for an NFS rpcbind issue.

- a** Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

- b** Reload the systemd manager configuration.

```
systemctl daemon-reload
```

Configuring NFS 4.0

Use this procedure to configure NFS 4.0 on delegate hosts if the operating system release is 7.4. There may be a file locking defect in NFS 4.1 with RHEL/CentOS 7.4.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Determine which release is installed.

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

- If the result includes 7.4, perform the remaining steps of this procedure.
 - If the result includes 7.2 or 7.3, continue to the next procedure.
- 3 Change the NFS configuration file.
 - a Open `/etc/nfsmount.conf` with a text editor.
 - b Locate the `Defaultvers` directive.
 - c Remove the number sign character (`#`) from the beginning of the line.
 - d Change the value from 4 to 4.0.
The line should appear as follows:

```
Defaultvers=4.0
```

- e Save the file, and then close the editor.
- 4 Restart the NFS server.

```
systemctl restart nfs-server
```

Updating Control Center on delegate hosts

This procedure updates Control Center on delegate hosts to version 1.5.1.

Perform this procedure on each delegate host in your deployment.

- 1 Log in to a delegate host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.
 - a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \
/opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.

- a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^\\b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

For more information about configuring a delegate host, see [Control Center configuration variables](#) on page 146.

When all delegate hosts are updated, start the cluster. For more information, see [Starting and stopping Control Center deployments](#) on page 132.

Updating the ZooKeeper image on ensemble nodes

Use this procedure to install a new Docker image for ZooKeeper into the local repository of delegate hosts that are ZooKeeper ensemble nodes.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the hosts in the ZooKeeper ensemble.

```
grep -E '^\\b*SERVICED_ZK=' /etc/default/serviced
```

The result is a list of 3 or 5 hosts, separated by the comma character (,). The master host is always a node in the ZooKeeper ensemble.

- 3 Log in to a ZooKeeper ensemble node as `root`, or as a user with superuser privileges.
- 4 If you are upgrading Control Center 1.3.3, start the Docker service.

```
systemctl start docker
```

- 5 Change directory to `/root`.

```
cd /root
```

- 6 Extract the ZooKeeper image.

```
./install-zenoss-isvcs-zookeeper-v10.run
```

At the prompt, press `y`.

- 7 Optional: Delete the archive file, if desired.

```
rm -i ./install-zenoss-isvcs-zookeeper-v10.run
```

- 8 Repeat the preceding five steps on each ZooKeeper node in the ensemble.

After upgrading

Perform the procedures in this chapter after Control Center is upgraded.

Setting the connection timeout of a resource pool

Use this procedure to set the length of time the services scheduler waits for a disconnected delegate host to reconnect before moving its services to a different host in the resource pool. This affects all multi-host resource pools. Zenoss recommends using a minimum value of 15 seconds for all multi-host pools. For resource pools that are connected through high-latency, wide-area networks, Zenoss recommends values greater than 15 seconds.

- 1 Log in to the master host as a user with `serviced` CLI privileges.
- 2 Display the list of resource pools and their connection timeout values.

```
serviced pool list -v | grep -E 'ID|ConnectionTimeout'
```

- 3 Optional: Set the connection timeout value of a resource pool.

This command accepts the following units identifiers:

- ms (milliseconds)
- s (seconds)
- m (minutes)
- h (hours)

Replace *Pool-ID* with a resource pool identifier, and replace *Timeout+Units* with an integer followed by a units identifier:

```
serviced pool set-conn-timeout Pool-ID Timeout+Units
```

Removing unused images

Use this procedure to identify and remove unused Control Center images.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the images associated with the installed version of `serviced`.

```
serviced version | grep Images
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v61 zenoss/isvcs-zookeeper:v10]
```

3 Start Docker.

```
systemctl start docker
```

4 Display the serviced images in the local repository.

```
docker images | awk '/REPO|isvcs/'
```

Example result (edited to fit):

REPOSITORY	TAG	IMAGE ID
zenoss/serviced-isvcs	v40	88cd6c24cc82
zenoss/serviced-isvcs	v61	0aab5a2123f2
zenoss/isvcs-zookeeper	v3	46fa0a2fc4bf
zenoss/isvcs-zookeeper	v10	0ff3b3117fb8

The example result shows the current versions and one set of previous versions. Your result may include additional previous versions and will show different images IDs.

5 Remove unused images.

Replace *Image-ID* with the image ID of an image for a previous version.

```
docker rmi Image-ID
```

Repeat this command for each unused image.

6 Stop Docker.

```
systemctl stop docker
```

Upgrading 1.3.x to 1.5.1 Scope

The chapters in this part provide instructions for performing an upgrade of a high-availability deployment of Control Center, from version 1.3.x to 1.5.1.

The following table identifies the upgrades of Control Center that are documented in this part.

From	To
Control Center 1.3.3	Control Center 1.5.1
Control Center 1.3.2	Control Center 1.5.1
Control Center 1.3.1	Control Center 1.5.1
Control Center 1.3.0	Control Center 1.5.1

13

Before upgrading 1.3.x to 1.5.1

This chapter provides information and procedures to prepare your high-availability Control Center deployment for an upgrade.

New features that affect upgrades from release 1.3.x

This release includes new features and new requirements that affect the upgrade process. The following list provides an overview of the changes that are addressed during this upgrade.

- RHEL/CentOS 7.4 is added; RHEL/CentOS 7.1 is withdrawn. The upgrade process includes a step for updating the operating system.
- On RHEL/CentOS 7.4, there may be a file locking defect in NFS 4.1. To avoid the issue, delegate hosts are configured to use NFS 4.0.
- The `serviced-zenossdbpack` maintenance script is moved from `/etc/cron.weekly` to `/opt/serviced/bin` and a new cron job, `/etc/cron.d/cron_zenossdbpack`, is installed to invoke the maintenance script. By default, the script runs every Sunday at midnight.
- A new remote storage location is now required for audit logs. The default mount point is `/var/log/serviced`, and the default amount of space required is 10GB. The master node upgrade process includes a procedure for ensuring that a remote storage location is configured.
- The Pacemaker resource agents are changed so that a master host does not fail over when Control Center is unresponsive or its health checks are failing.
- This release replaces Docker 1.12.1 with Docker Community Edition (CE) 17.09.0.
- The following new configuration variables are available in `/etc/default/serviced`, for tuning TCP/IP communications between ZooKeeper ensemble hosts and Control Center:

```
SERVICED_ZK_CONNECT_TIMEOUT
SERVICED_ZK_PER_HOST_CONNECT_DELAY
SERVICED_ZK_RECONNECT_START_DELAY
SERVICED_ZK_RECONNECT_MAX_DELAY
```

For more information, see [Control Center configuration file](#) on page 155.

- The Control Center RPM package includes a script that adds the `serviced` user group to a host, if necessary. Users must be members of the `serviced` group to use the command-line interface. For more information, see [User access control](#) on page 140.

For more information about this release, refer to the *Control Center Release Notes*.

Upgrade best practices

The following list outlines recommended best practices for upgrading high-availability Control Center deployments:

- 1 Download a copy of the *Control Center Release Notes* for this release and review its contents. The latest information is included in that document.
- 2 Upgrade only what needs upgrading. For example, you can update some components independently. For more information, see the following appendices:
 - [Upgrading only the Pacemaker resource agents](#) on page 126
 - [Updating the cluster management software](#) on page 143
- 3 Compare the Docker images that accompany this release with the images that accompany the installed release, and determine whether the image files need to be downloaded and installed. For more information, see [Releases and image tags](#).
- 4 The contents of `/etc/default/serviced` on the master nodes must be identical. Use a utility like `sum` to compare the files quickly.
- 5 On delegate hosts, most of the upgrade steps are identical. Use `screen`, `tmux`, or a similar program to establish sessions on each delegate host and perform the steps at the same time.
- 6 Review and verify the settings in delegate host configuration files (`/etc/default/serviced`) before starting the upgrade. Ideally, the settings on all delegate hosts are identical, except on ZooKeeper nodes and delegate hosts that do not mount the DFS.
- 7 Review the procedures in this guide before performing them. Every effort is made to avoid mistakes and anticipate needs; nevertheless, the instructions may be incorrect or inadequate for some requirements or environments.
- 8 Download and stage the required files for your upgrade. For more information, see [Downloading and staging required files](#) on page 13.

14

Stopping a Control Center deployment

This chapter includes procedures for stopping a high-availability Control Center deployment.

Note The procedures in this chapter assume that Control Center is the only source of Docker containers that are run on Control Center hosts.

Perform the procedures in this chapter in order.

Stopping a master host node

Use this procedure to stop the Control Center service (*serviced*) on the master host in a high-availability deployment.

- 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 step.

- 3 Stop the top-level service *serviced* is managing, if necessary.
 - a Show the status of running services.

```
serviced service status
```

The top-level service is the service listed immediately below the headings line.

- If the status of the top-level service and all child services is `stopped`, proceed to the next step.
 - If the status of the top-level service and all child services is **not** `stopped`, perform the remaining substeps.
- b Stop the top-level service.
Replace *Service* with the name or identifier of the top-level service:

```
serviced service stop Service
```

- c Monitor the stop.

```
serviced service status
```

When the status of the top-level service and all child services is stopped, proceed to the next step.

- 4 Stop Control Center with the cluster management tool.

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

- 5 Monitor the status of cluster resources.

```
watch pcs status
```

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

- 6 Ensure that no containers remain in the local repository.

- a Start the Docker service.

```
systemctl start docker
```

- b Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

If the command returns a result, enter the following command:

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- c Stop the Docker service.

```
systemctl stop docker
```

- 7 To ensure that no containers remain in both Docker repositories, log in to the other master node as `root`, or as a user with superuser privileges, and then perform the preceding step.

Stopping a delegate host

Use this procedure to stop the Control Center service (`serviced`) on a delegate host in a multi-host deployment. Repeat this procedure on each delegate host in your deployment.

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

```
systemctl stop serviced
```

- 3 Ensure that no containers remain in the local repository.

- a Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, proceed to the next step.
- If the command returns a result, perform the following substeps.

- b Remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, the most likely cause is an NFS conflict. Perform the following substeps.

- c** Stop the NFS and Docker services.

```
systemctl stop nfs && systemctl stop docker
```

- d** Start the NFS and Docker services.

```
systemctl start nfs && systemctl start docker
```

- e** Repeat the attempt to remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, perform the remaining substeps.

- f** Disable the automatic startup of `serviced`.

```
systemctl disable serviced
```

- g** Reboot the host.

```
reboot
```

- h** Log in to the delegate host as `root`, or as a user with superuser privileges.

- i** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

4 Dismount all filesystems mounted from the Control Center master host.

This step ensures no stale mounts remain when the storage on the master host is replaced.

- a** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- b** Force the filesystems to dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
  umount -f $FS
done
```

- c** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- d** Perform a lazy dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
    umount -f -l $FS
done
```

- e** Restart the NFS service.

```
systemctl restart nfs
```

- f** Determine whether any filesystems remain mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the remaining substeps.

- g** Disable the automatic startup of serviced.

```
systemctl disable serviced
```

- h** Reboot the host.

```
reboot
```

- i** Log in to the delegate host as `root`, or as a user with superuser privileges.

- j** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

Upgrading master hosts from 1.3.x to 1.5.1

15

Perform the procedures in this chapter to upgrade both Control Center 1.3.x master hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Adding storage for audit logging

The default configuration of Control Center audit logging requires 10GB of storage in `/var/log/serviced` on master host nodes. For high-availability deployments, Zenoss recommends mounting remote storage instead of adding local storage to master nodes. Use this procedure to create the mount point.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Create a mount point for audit logs.

The default location for `serviced` audit logs is `/var/log/serviced`. The location is determined by the `SERVICED_LOG_PATH` variable in `/etc/default/serviced`.

```
mkdir -p /var/log/serviced
```

- 3 Mount a remote filesystem at `/var/log/serviced`.
For more information about mounting a remote filesystem, refer to your operating system documentation.

Updating Docker

Use this procedure to update Docker to version 17.09.0.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Remove the Docker repository description file, if present.
Upgrades no longer require external repositories.

```
rm -f /etc/yum.repos.d/docker.repo
```

- 3 Update the operating system, if necessary.
 - a Determine which release is installed.

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

- If the result is 7.2, 7.3, or 7.4, proceed to the next step.

- If the result is 7.1, perform the remaining substeps.

b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

c Update the operating system, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update && reboot
```

d Log in as `root`, or as a user with superuser privileges.

e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

4 Update the Linux kernel, if necessary.

a Determine which kernel version is installed.

```
uname -r
```

If the result is lower than `3.10.0-327.22.2.el7.x86_64`, perform the following substeps.

b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

c Update the kernel, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update kernel && reboot
```

d Log in as `root`, or as a user with superuser privileges.

e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

5 Stop the Docker service.

```
systemctl stop docker
```

6 Remove Docker.

a Remove without checking dependencies.

Control Center 1.3.0, 1.3.1, or 1.3.2:

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

Control Center 1.3.3:

```
rpm -e --nodeps docker-ce
```

b Clean the yum databases.

```
yum clean all
```


7 Install Docker CE 17.09.0.

```
yum install --enablerepo=zenoss-mirror docker-ce-17.09.0.ce
```

If yum returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

8 Start the Docker service.

```
systemctl start docker
```

9 Remove a previous workaround for an NFS rpcbind issue.**a** Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

b Reload the systemd manager configuration.

```
systemctl daemon-reload
```

Loading image files

Use this procedure to load images into the local registry.

- 1** Log in to the master host as `root`, or as a user with superuser privileges.
- 2** If you are upgrading Control Center 1.3.3, start the Docker service.

```
systemctl start docker
```

3 Change directory to `/root`.

```
cd /root
```

4 Load the images.

```
for image in install-zenoss-*.run
do
  /bin/echo -en "\nLoading $image..."
  yes | ./$image
done
```

5 List the images in the registry.

```
docker images
```

The result should show one image for each archive file.

6 Optional: Delete the archive files, if desired.

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

7 Stop the Docker service.

```
systemctl stop docker
```

Updating the Pacemaker resource agents

Use this procedure to update the Pacemaker resource agents for Control Center.

- 1 Log in to a master host node as `root`, or as a user with superuser privileges.
- 2 In a separate terminal session, log in to the other master host node as `root`, or as a user with superuser privileges.
- 3 On both nodes, install the latest version of the resource agents package.

```
myPath=/opt/zenoss-repo-mirror
yum install \
  $myPath/serviced-resource-agents-1.1.0-1.x86_64.rpm
```

Updating Control Center on master hosts

Use this procedure to update Control Center on both master hosts to version 1.5.1.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.
 - a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \
  /opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.
 - a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^\b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

For more information about configuring the master host, see [Control Center configuration variables](#) on page 146.

Upgrading delegates from 1.3.x to 1.5.1

16

Perform the procedures in this chapter to upgrade Control Center 1.3.x delegate hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Updating Docker

Use this procedure to update Docker to version 17.09.0.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Remove the Docker repository description file, if present.

Upgrades no longer require external repositories.

```
rm -f /etc/yum.repos.d/docker.repo
```

- 3 Update the operating system, if necessary.

- a Determine which release is installed.

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

- If the result is 7.2, 7.3, or 7.4, proceed to the next step.
- If the result is 7.1, perform the remaining substeps.

- b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

- c Update the operating system, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update && reboot
```

- d Log in as `root`, or as a user with superuser privileges.
- e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

- 4 Update the Linux kernel, if necessary.

- a Determine which kernel version is installed.

```
uname -r
```

If the result is lower than 3.10.0-327.22.2.el7.x86_64, perform the following substeps.

- b Disable automatic start of serviced.

```
systemctl disable serviced
```

- c Update the kernel, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update kernel && reboot
```

- d Log in as root, or as a user with superuser privileges.

- e Enable automatic start of serviced.

```
systemctl enable serviced
```

- 5 Stop the Docker service.

```
systemctl stop docker
```

- 6 Remove Docker.

- a Remove without checking dependencies.

Control Center 1.3.0, 1.3.1, or 1.3.2:

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

Control Center 1.3.3:

```
rpm -e --nodeps docker-ce
```

- b Clean the yum databases.

```
yum clean all
```

- 7 Install Docker CE 17.09.0.

```
yum install --enablerepo=zenoss-mirror docker-ce-17.09.0.ce
```

If yum returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 8 Start the Docker service.

```
systemctl start docker
```

- 9 Remove a previous workaround for an NFS rpcbind issue.

- a Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

- b** Reload the `systemd` manager configuration.

```
systemctl daemon-reload
```

Configuring NFS 4.0

Use this procedure to configure NFS 4.0 on delegate hosts if the operating system release is 7.4. There may be a file locking defect in NFS 4.1 with RHEL/CentOS 7.4.

- 1** Log in to the host as `root`, or as a user with superuser privileges.
- 2** Determine which release is installed.

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

- If the result includes 7.4, perform the remaining steps of this procedure.
 - If the result includes 7.2 or 7.3, continue to the next procedure.
- 3** Change the NFS configuration file.
 - a** Open `/etc/nfsmount.conf` with a text editor.
 - b** Locate the `Defaultvers` directive.
 - c** Remove the number sign character (`#`) from the beginning of the line.
 - d** Change the value from 4 to 4.0.
The line should appear as follows:

```
Defaultvers=4.0
```

- e** Save the file, and then close the editor.
- 4** Restart the NFS server.

```
systemctl restart nfs-server
```

Updating Control Center on delegate hosts

This procedure updates Control Center on delegate hosts to version 1.5.1.

Perform this procedure on each delegate host in your deployment.

- 1** Log in to a delegate host as `root`, or as a user with superuser privileges.
- 2** Save the current `serviced` configuration file as a reference.
 - a** Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b** Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3** Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \
  /opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.
 - a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^\\b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

For more information about configuring a delegate host, see [Control Center configuration variables](#) on page 146.

When all delegate hosts are updated, start the cluster. For more information, see [Starting and stopping Control Center deployments](#) on page 132.

Updating the ZooKeeper image on ensemble nodes

Use this procedure to install a new Docker image for ZooKeeper into the local repository of delegate hosts that are ZooKeeper ensemble nodes.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the hosts in the ZooKeeper ensemble.

```
grep -E '^\\b*SERCIVED_ZK=' /etc/default/serviced
```

The result is a list of 3 or 5 hosts, separated by the comma character (,). The master host is always a node in the ZooKeeper ensemble.

- 3 Log in to a ZooKeeper ensemble node as `root`, or as a user with superuser privileges.
- 4 If you are upgrading Control Center 1.3.3, start the Docker service.

```
systemctl start docker
```

- 5 Change directory to `/root`.

```
cd /root
```

- 6 Extract the ZooKeeper image.

```
./install-zenoss-isvcs-zookeeper-v10.run
```

At the prompt, press `y`.

- 7 Optional: Delete the archive file, if desired.

```
rm -i ./install-zenoss-isvcs-zookeeper-v10.run
```

- 8 Repeat the preceding five steps on each ZooKeeper node in the ensemble.

After upgrading

Perform the procedures in this chapter after Control Center is upgraded.

Setting the connection timeout of a resource pool

Use this procedure to set the length of time the services scheduler waits for a disconnected delegate host to reconnect before moving its services to a different host in the resource pool. This affects all multi-host resource pools. Zenoss recommends using a minimum value of 15 seconds for all multi-host pools. For resource pools that are connected through high-latency, wide-area networks, Zenoss recommends values greater than 15 seconds.

- 1 Log in to the master host as a user with `serviced` CLI privileges.
- 2 Display the list of resource pools and their connection timeout values.

```
serviced pool list -v | grep -E 'ID|ConnectionTimeout'
```

- 3 Optional: Set the connection timeout value of a resource pool.

This command accepts the following units identifiers:

- ms (milliseconds)
- s (seconds)
- m (minutes)
- h (hours)

Replace *Pool-ID* with a resource pool identifier, and replace *Timeout+Units* with an integer followed by a units identifier:

```
serviced pool set-conn-timeout Pool-ID Timeout+Units
```

Removing unused images

Use this procedure to identify and remove unused Control Center images.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the images associated with the installed version of `serviced`.

```
serviced version | grep Images
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v61 zenoss/isvcs-zookeeper:v10]
```

3 Start Docker.

```
systemctl start docker
```

4 Display the serviced images in the local repository.

```
docker images | awk '/REPO|isvcs/'
```

Example result (edited to fit):

REPOSITORY	TAG	IMAGE ID
zenoss/serviced-isvcs	v40	88cd6c24cc82
zenoss/serviced-isvcs	v61	0aab5a2123f2
zenoss/isvcs-zookeeper	v3	46fa0a2fc4bf
zenoss/isvcs-zookeeper	v10	0ff3b3117fb8

The example result shows the current versions and one set of previous versions. Your result may include additional previous versions and will show different images IDs.

5 Remove unused images.

Replace *Image-ID* with the image ID of an image for a previous version.

```
docker rmi Image-ID
```

Repeat this command for each unused image.

6 Stop Docker.

```
systemctl stop docker
```


Upgrading 1.2.x to 1.5.1

Scope

The chapters in this part provide instructions for performing an upgrade of a high-availability deployment of Control Center, from version 1.2.x to 1.5.1.

The following table identifies the upgrades of Control Center that are documented in this part.

From	To
Control Center 1.2.3	Control Center 1.5.1
Control Center 1.2.2	Control Center 1.5.1
Control Center 1.2.1	Control Center 1.5.1
Control Center 1.2.0	Control Center 1.5.1

Before upgrading 1.2.x to 1.5.1

This chapter provides information and procedures to prepare your high-availability Control Center deployment for an upgrade.

New features that affect upgrades from release 1.2.x

This release includes new features and new requirements that affect the upgrade process. The following list provides an overview of the changes that are addressed during this upgrade.

Note Zenoss Resource Manager may require additional steps to prepare to use this release. For more information, refer to Zenoss Resource Manager documentation.

- RHEL/CentOS 7.4 is added; RHEL/CentOS 7.1 is withdrawn. The upgrade process includes a step for updating the operating system.
- On RHEL/CentOS 7.4, there may be a file locking defect in NFS 4.1. To avoid the issue, delegate hosts are configured to use NFS 4.0.
- The `serviced-zenossdbpack` maintenance script is moved from `/etc/cron.weekly` to `/opt/serviced/bin` and a new cron job, `/etc/cron.d/cron_zenossdbpack`, is installed to invoke the maintenance script. By default, the script runs every Sunday at midnight.
- A new remote storage location is now required for audit logs. The default mount point is `/var/log/serviced`, and the default amount of space required is 10GB. The master node upgrade process includes a procedure for ensuring that a remote storage location is configured.
- The Pacemaker resource agents are changed so that a master host does not fail over when Control Center is unresponsive or its health checks are failing.
- When storage becomes critically low, Control Center initiates an *emergency shutdown* of applications and services while sufficient resources remain to take action to avoid data loss. This feature requires specific minimum amounts of space in the Control Center thin pool for application data.

Note Before upgrading, determine whether the thin pool is adequate and if necessary, resize it. For more information, see [Control Center application data storage requirements](#) on page 115.

- Previously, Zenoss recommended using the master host nodes for Control Center services only. Starting with this release, Zenoss recommends using the master host nodes for Control Center services and for two database services. The services requires additional RAM and CPU resources. For more information about the amount of RAM and CPU resources to add to master host nodes, please contact your Zenoss representative.
- This release replaces Docker 1.12.1 with Docker Community Edition (CE) 17.09.0.

- The following new configuration variables are available in `/etc/default/serviced`, for tuning TCP/IP communications between ZooKeeper ensemble hosts and Control Center:

```
SERVICED_ZK_CONNECT_TIMEOUT
SERVICED_ZK_PER_HOST_CONNECT_DELAY
SERVICED_ZK_RECONNECT_START_DELAY
SERVICED_ZK_RECONNECT_MAX_DELAY
```

For more information, see [Control Center configuration file](#) on page 155.

- The Control Center RPM package includes a script that adds the `serviced` user group to a host, if necessary. Users must be members of the `serviced` group to use the command-line interface. For more information, see [User access control](#) on page 140.

For more information about this release, refer to the [Control Center Release Notes](#).

Upgrade best practices

The following list outlines recommended best practices for upgrading high-availability Control Center deployments:

- 1 Download a copy of the [Control Center Release Notes](#) for this release and review its contents. The latest information is included in that document.
- 2 Upgrade only what needs upgrading. For example, you can update some components independently. For more information, see the following appendices:
 - [Upgrading only the Pacemaker resource agents](#) on page 126
 - [Updating the cluster management software](#) on page 143
- 3 Compare the Docker images that accompany this release with the images that accompany the installed release, and determine whether the image files need to be downloaded and installed. For more information, see [Releases and image tags](#).
- 4 The contents of `/etc/default/serviced` on the master nodes must be identical. Use a utility like `sum` to compare the files quickly.
- 5 On delegate hosts, most of the upgrade steps are identical. Use `screen`, `tmux`, or a similar program to establish sessions on each delegate host and perform the steps at the same time.
- 6 Review and verify the settings in delegate host configuration files (`/etc/default/serviced`) before starting the upgrade. Ideally, the settings on all delegate hosts are identical, except on ZooKeeper nodes and delegate hosts that do not mount the DFS.
- 7 Review the procedures in this guide before performing them. Every effort is made to avoid mistakes and anticipate needs; nevertheless, the instructions may be incorrect or inadequate for some requirements or environments.
- 8 Download and stage the required files for your upgrade. For more information, see [Downloading and staging required files](#) on page 13.

Stopping a Control Center deployment

This chapter includes procedures for stopping a high-availability Control Center deployment.

Note The procedures in this chapter assume that Control Center is the only source of Docker containers that are run on Control Center hosts.

Perform the procedures in this chapter in order.

Stopping a master host node

Use this procedure to stop the Control Center service (*serviced*) on the master host in a high-availability deployment.

- 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 step.

- 3 Stop the top-level service *serviced* is managing, if necessary.
 - a Show the status of running services.

```
serviced service status
```

The top-level service is the service listed immediately below the headings line.

- If the status of the top-level service and all child services is `stopped`, proceed to the next step.
- If the status of the top-level service and all child services is **not** `stopped`, perform the remaining substeps.

- b Stop the top-level service.

Replace *Service* with the name or identifier of the top-level service:

```
serviced service stop Service
```

- c Monitor the stop.

```
serviced service status
```

When the status of the top-level service and all child services is stopped, proceed to the next step.

- 4 Stop Control Center with the cluster management tool.

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

- 5 Monitor the status of cluster resources.

```
watch pcs status
```

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

- 6 Ensure that no containers remain in the local repository.

- a Start the Docker service.

```
systemctl start docker
```

- b Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

If the command returns a result, enter the following command:

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- c Stop the Docker service.

```
systemctl stop docker
```

- 7 To ensure that no containers remain in both Docker repositories, log in to the other master node as `root`, or as a user with superuser privileges, and then perform the preceding step.

Stopping a delegate host

Use this procedure to stop the Control Center service (`serviced`) on a delegate host in a multi-host deployment. Repeat this procedure on each delegate host in your deployment.

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

```
systemctl stop serviced
```

- 3 Ensure that no containers remain in the local repository.

- a Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, proceed to the next step.
- If the command returns a result, perform the following substeps.

- b Remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, the most likely cause is an NFS conflict. Perform the following substeps.

- c** Stop the NFS and Docker services.

```
systemctl stop nfs && systemctl stop docker
```

- d** Start the NFS and Docker services.

```
systemctl start nfs && systemctl start docker
```

- e** Repeat the attempt to remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, perform the remaining substeps.

- f** Disable the automatic startup of `serviced`.

```
systemctl disable serviced
```

- g** Reboot the host.

```
reboot
```

- h** Log in to the delegate host as `root`, or as a user with superuser privileges.

- i** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

4 Dismount all filesystems mounted from the Control Center master host.

This step ensures no stale mounts remain when the storage on the master host is replaced.

- a** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- b** Force the filesystems to dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
    umount -f $FS
done
```

- c** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- d** Perform a lazy dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
    umount -f -l $FS
done
```

- e** Restart the NFS service.

```
systemctl restart nfs
```

- f** Determine whether any filesystems remain mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the remaining substeps.

- g** Disable the automatic startup of serviced.

```
systemctl disable serviced
```

- h** Reboot the host.

```
reboot
```

- i** Log in to the delegate host as `root`, or as a user with superuser privileges.

- j** Enable the automatic startup of serviced.

```
systemctl enable serviced
```

Upgrading master hosts from 1.2.x to 1.5.1

20

Perform the procedures in this chapter to upgrade both Control Center 1.2.x master hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Adding storage for audit logging

The default configuration of Control Center audit logging requires 10GB of storage in `/var/log/serviced` on master host nodes. For high-availability deployments, Zenoss recommends mounting remote storage instead of adding local storage to master nodes. Use this procedure to create the mount point.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Create a mount point for audit logs.

The default location for `serviced` audit logs is `/var/log/serviced`. The location is determined by the `SERVICED_LOG_PATH` variable in `/etc/default/serviced`.

```
mkdir -p /var/log/serviced
```

- 3 Mount a remote filesystem at `/var/log/serviced`.
For more information about mounting a remote filesystem, refer to your operating system documentation.

Updating Docker

Use this procedure to update Docker to version 17.09.0.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Remove the Docker repository description file, if present.
Upgrades no longer require external repositories.

```
rm -f /etc/yum.repos.d/docker.repo
```

- 3 Update the operating system, if necessary.
 - a Determine which release is installed.

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

- If the result is 7.2, 7.3, or 7.4, proceed to the next step.

- If the result is 7.1, perform the remaining substeps.

b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

c Update the operating system, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update && reboot
```

d Log in as `root`, or as a user with superuser privileges.

e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

4 Update the Linux kernel, if necessary.

a Determine which kernel version is installed.

```
uname -r
```

If the result is lower than `3.10.0-327.22.2.el7.x86_64`, perform the following substeps.

b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

c Update the kernel, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update kernel && reboot
```

d Log in as `root`, or as a user with superuser privileges.

e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

5 Stop the Docker service.

```
systemctl stop docker
```

6 Remove Docker.

a Remove without checking dependencies.

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

b Clean the yum databases.

```
yum clean all
```

7 Install Docker CE 17.09.0.

```
yum install --enablerepo=zenoss-mirror docker-ce-17.09.0.ce
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 8 Start the Docker service.

```
systemctl start docker
```

- 9 Remove a previous workaround for an NFS `rpcbind` issue.

- a Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

- b Reload the `systemd` manager configuration.

```
systemctl daemon-reload
```

Loading image files

Use this procedure to load images into the local registry.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Change directory to `/root`.

```
cd /root
```

- 3 Load the images.

```
for image in install-zenoss-*.run
do
  /bin/echo -en "\nLoading $image..."
  yes | ./$image
done
```

- 4 List the images in the registry.

```
docker images
```

The result should show one image for each archive file.

- 5 Optional: Delete the archive files, if desired.

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

- 6 Stop the Docker service.

```
systemctl stop docker
```

Updating the Pacemaker resource agents

Use this procedure to update the Pacemaker resource agents for Control Center.

- 1 Log in to a master host node as `root`, or as a user with superuser privileges.
- 2 In a separate terminal session, log in to the other master host node as `root`, or as a user with superuser privileges.
- 3 On both nodes, install the latest version of the resource agents package.

```
myPath=/opt/zenoss-repo-mirror
```

```
yum install \
  $myPath/serviced-resource-agents-1.1.0-1.x86_64.rpm
```

Updating Control Center on master hosts

Use this procedure to update Control Center on both master hosts to version 1.5.1.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.
 - a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \
  /opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.
 - a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^\b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

For more information about configuring the master host, see [Control Center configuration variables](#) on page 146.

Upgrading delegates from 1.2.x to 1.5.1

21

Perform the procedures in this chapter to upgrade Control Center 1.2.x delegate hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Updating Docker

Use this procedure to update Docker to version 17.09.0.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Remove the Docker repository description file, if present.

Upgrades no longer require external repositories.

```
rm -f /etc/yum.repos.d/docker.repo
```

- 3 Update the operating system, if necessary.

- a Determine which release is installed.

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

- If the result is 7.2, 7.3, or 7.4, proceed to the next step.
- If the result is 7.1, perform the remaining substeps.

- b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

- c Update the operating system, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update && reboot
```

- d Log in as `root`, or as a user with superuser privileges.

- e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

- 4 Update the Linux kernel, if necessary.

- a** Determine which kernel version is installed.

```
uname -r
```

If the result is lower than 3.10.0-327.22.2.el7.x86_64, perform the following substeps.

- b** Disable automatic start of serviced.

```
systemctl disable serviced
```

- c** Update the kernel, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update kernel && reboot
```

- d** Log in as root, or as a user with superuser privileges.

- e** Enable automatic start of serviced.

```
systemctl enable serviced
```

- 5** Stop the Docker service.

```
systemctl stop docker
```

- 6** Remove Docker.

- a** Remove without checking dependencies.

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

- b** Clean the yum databases.

```
yum clean all
```

- 7** Install Docker CE 17.09.0.

```
yum install --enablerepo=zenoss-mirror docker-ce-17.09.0.ce
```

If yum returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 8** Start the Docker service.

```
systemctl start docker
```

- 9** Remove a previous workaround for an NFS rpcbind issue.

- a** Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

- b** Reload the systemd manager configuration.

```
systemctl daemon-reload
```

Configuring NFS 4.0

Use this procedure to configure NFS 4.0 on delegate hosts if the operating system release is 7.4. There may be a file locking defect in NFS 4.1 with RHEL/CentOS 7.4.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Determine which release is installed.

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

- If the result includes 7.4, perform the remaining steps of this procedure.
 - If the result includes 7.2 or 7.3, continue to the next procedure.
- 3 Change the NFS configuration file.
 - a Open `/etc/nfsmount.conf` with a text editor.
 - b Locate the `Defaultvers` directive.
 - c Remove the number sign character (`#`) from the beginning of the line.
 - d Change the value from 4 to 4.0.
The line should appear as follows:

```
Defaultvers=4.0
```

- e Save the file, and then close the editor.
- 4 Restart the NFS server.

```
systemctl restart nfs-server
```

Updating the ZooKeeper image on ensemble nodes

Use this procedure to install a new Docker image for ZooKeeper into the local repository of delegate hosts that are ZooKeeper ensemble nodes.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the hosts in the ZooKeeper ensemble.

```
grep -E '^\\b*SERVICED_ZK=' /etc/default/serviced
```

The result is a list of 3 or 5 hosts, separated by the comma character (`,`). The master host is always a node in the ZooKeeper ensemble.

- 3 Log in to a ZooKeeper ensemble node as `root`, or as a user with superuser privileges.
- 4 Change directory to `/root`.

```
cd /root
```

- 5 Extract the ZooKeeper image.

```
./install-zenoss-isvcs-zookeeper-v10.run
```

At the prompt, press `y`.

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

```
rm -i ./install-zenoss-isvcs-zookeeper-v10.run
```

- 7 Repeat the preceding four steps on each ZooKeeper node in the ensemble.

Updating Control Center on delegate hosts

Use this procedure to update Control Center on delegate hosts to version 1.5.1.

Perform this procedure on each delegate host in your deployment.

- 1 Log in to a delegate host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.
 - a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \  
/opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.
 - a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^\\b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

For more information about configuring a delegate host, see [Control Center configuration variables](#) on page 146.

When all delegate hosts are updated, start the cluster. For more information, see [Starting and stopping Control Center deployments](#) on page 132.

After upgrading from 1.2.x to 1.5.1

Perform the procedures in this chapter after Control Center is upgraded.

Setting the connection timeout of a resource pool

Use this procedure to set the length of time the services scheduler waits for a disconnected delegate host to reconnect before moving its services to a different host in the resource pool. This affects all multi-host resource pools. Zenoss recommends using a minimum value of 15 seconds for all multi-host pools. For resource pools that are connected through high-latency, wide-area networks, Zenoss recommends values greater than 15 seconds.

- 1 Log in to the master host as a user with `serviced` CLI privileges.
- 2 Display the list of resource pools and their connection timeout values.

```
serviced pool list -v | grep -E 'ID|ConnectionTimeout'
```

- 3 Optional: Set the connection timeout value of a resource pool.

This command accepts the following units identifiers:

- ms (milliseconds)
- s (seconds)
- m (minutes)
- h (hours)

Replace *Pool-ID* with a resource pool identifier, and replace *Timeout+Units* with an integer followed by a units identifier:

```
serviced pool set-conn-timeout Pool-ID Timeout+Units
```

Removing unused images

Use this procedure to identify and remove unused Control Center images.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the images associated with the installed version of `serviced`.

```
serviced version | grep Images
```


Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v61 zenoss/isvcs-zookeeper:v10]
```

3 Start Docker.

```
systemctl start docker
```

4 Display the serviced images in the local repository.

```
docker images | awk '/REPO|isvcs/'
```

Example result (edited to fit):

REPOSITORY	TAG	IMAGE ID
zenoss/serviced-isvcs	v40	88cd6c24cc82
zenoss/serviced-isvcs	v61	0aab5a2123f2
zenoss/isvcs-zookeeper	v3	46fa0a2fc4bf
zenoss/isvcs-zookeeper	v10	0ff3b3117fb8

The example result shows the current versions and one set of previous versions. Your result may include additional previous versions and will show different images IDs.

5 Remove unused images.

Replace *Image-ID* with the image ID of an image for a previous version.

```
docker rmi Image-ID
```

Repeat this command for each unused image.

6 Stop Docker.

```
systemctl stop docker
```

Upgrading 1.1.x to 1.5.1

Scope

The chapters in this part provide instructions for performing an upgrade of a high-availability deployment of Control Center, from version 1.1.x to 1.5.1.

The following table identifies the upgrades of Control Center that are documented in this part.

From	To
Control Center 1.1.10	Control Center 1.5.1
Control Center 1.1.9	Control Center 1.5.1
Control Center 1.1.8	Control Center 1.5.1
Control Center 1.1.7	Control Center 1.5.1
Control Center 1.1.5	Control Center 1.5.1
Control Center 1.1.4	Control Center 1.5.1
Control Center 1.1.3	Control Center 1.5.1
Control Center 1.1.2	Control Center 1.5.1
Control Center 1.1.1	Control Center 1.5.1

Before upgrading 1.1.x to 1.5.1

This chapter provides information and procedures to prepare your high-availability Control Center deployment for an upgrade.

New features that affect upgrades from 1.1.x

This release includes new features and new requirements that affect the upgrade process. The following list provides an overview of the changes that are addressed during this upgrade.

Note Zenoss Resource Manager may require additional steps to prepare to use this release. For more information, refer to Zenoss Resource Manager documentation.

- RHEL/CentOS 7.4 is added; RHEL/CentOS 7.1 is withdrawn. The upgrade process includes a step for updating the operating system.
- On RHEL/CentOS 7.4, there may be a file locking defect in NFS 4.1. To avoid the issue, delegate hosts are configured to use NFS 4.0.
- The `serviced-zenossdbpack` maintenance script is moved from `/etc/cron.weekly` to `/opt/serviced/bin` and a new cron job, `/etc/cron.d/cron_zenossdbpack`, is installed to invoke the maintenance script. By default, the script runs every Sunday at midnight.
- A new remote storage location is now required for audit logs. The default mount point is `/var/log/serviced`, and the default amount of space required is 10GB. The master node upgrade process includes a procedure for ensuring that a remote storage location is configured.
- The Pacemaker resource agents are changed so that a master host does not fail over when Control Center is unresponsive or its health checks are failing.
- When storage becomes critically low, Control Center initiates an *emergency shutdown* of applications and services while sufficient resources remain to take action to avoid data loss. This feature requires specific minimum amounts of space in the Control Center thin pool for application data.

Note Before upgrading, determine whether the thin pool is adequate and if necessary, resize it. For more information, see [Control Center application data storage requirements](#) on page 115.

- Previously, Zenoss recommended using the master host nodes for Control Center services only. Starting with this release, Zenoss recommends using the master host nodes for Control Center services and for two database services. The services requires additional RAM and CPU resources. For more information about the amount of RAM and CPU resources to add to master host nodes, please contact your Zenoss representative.
- This release replaces Docker 1.9.0 with Docker Community Edition (CE) 17.09.0.

- The following new configuration variables are available in `/etc/default/serviced`, for tuning TCP/IP communications between ZooKeeper ensemble hosts and Control Center:

```
SERVICED_ZK_CONNECT_TIMEOUT
SERVICED_ZK_PER_HOST_CONNECT_DELAY
SERVICED_ZK_RECONNECT_START_DELAY
SERVICED_ZK_RECONNECT_MAX_DELAY
```

For more information, see [Control Center configuration file](#) on page 155.

- The Control Center RPM package includes a script that adds the `serviced` user group to a host, if necessary. Users must be members of the `serviced` group to use the command-line interface. For more information, see [User access control](#) on page 140.
- The minimum kernel version for Control Center hosts is 3.10.0-327.36.2. For optimal results, the most recent kernel is recommended. To prevent dependency issues, updating the kernel or operating system is a step in the Docker update procedure.
- The `serviced` configuration file includes many new variables since version 1.1.1, and some deprecated variables. For more information, see [Control Center configuration variables](#) on page 146.
- All delegate communications are authenticated. To enable this feature, all existing hosts must install unique credentials, which are generated on the master host. The installation steps are included in the startup procedures.
- With delegate authentication, Control Center can control administrative and DFS access permissions at the resource pool level. During the upgrade, all existing resource pools are given both administrative and DFS access permissions. The post-upgrade chapter includes an optional procedure for removing permissions from a resource pool.
- In previous releases, the `SERVICED_NFS_CLIENT` variable was set on delegate hosts to prevent access to the DFS. In this release, `SERVICED_NFS_CLIENT` is deprecated in favor of setting DFS access permission at the resource pool level. To ease the transition to the new functionality, delegate host configurations that include the `SERVICED_NFS_CLIENT` variable are still supported.
- This release includes a new resource pool feature, the ability to set the length of time the scheduler waits for a disconnected delegate host to rejoin its pool before moving the services scheduled for the delegate to a different host in the pool. This feature is useful for remote resource pools that are connected through a high-latency, wide-area network.
- Among other changes since version 1.9.0, Docker 17.09.0 includes a new storage subsystem. The initial startup takes a little longer, as the old layout is replaced with the new layout.

For more information about this release, refer to the *Control Center Release Notes*.

Upgrade best practices

The following list outlines recommended best practices for upgrading high-availability Control Center deployments:

- 1 Download a copy of the *Control Center Release Notes* for this release and review its contents. The latest information is included in that document.
- 2 Upgrade only what needs upgrading. For example, you can update some components independently. For more information, see the following appendices:
 - [Upgrading only the Pacemaker resource agents](#) on page 126
 - [Updating the cluster management software](#) on page 143
- 3 Compare the Docker images that accompany this release with the images that accompany the installed release, and determine whether the image files need to be downloaded and installed. For more information, see [Releases and image tags](#).
- 4 The contents of `/etc/default/serviced` on the master nodes must be identical. Use a utility like `sum` to compare the files quickly.

- 5 On delegate hosts, most of the upgrade steps are identical. Use `screen`, `tmux`, or a similar program to establish sessions on each delegate host and perform the steps at the same time.
- 6 Review and verify the settings in delegate host configuration files (`/etc/default/serviced`) before starting the upgrade. Ideally, the settings on all delegate hosts are identical, except on ZooKeeper nodes and delegate hosts that do not mount the DFS.
- 7 Review the procedures in this guide before performing them. Every effort is made to avoid mistakes and anticipate needs; nevertheless, the instructions may be incorrect or inadequate for some requirements or environments.
- 8 Download and stage the required files for your upgrade. For more information, see [Downloading and staging required files](#) on page 13.

Stopping a Control Center deployment

This chapter includes procedures for stopping a high-availability Control Center deployment.

Note The procedures in this chapter assume that Control Center is the only source of Docker containers that are run on Control Center hosts.

Perform the procedures in this chapter in order.

Stopping a master host node

Use this procedure to stop the Control Center service (*serviced*) on the master host in a high-availability deployment.

- 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 step.

- 3 Stop the top-level service *serviced* is managing, if necessary.
 - a Show the status of running services.

```
serviced service status
```

The top-level service is the service listed immediately below the headings line.

- If the status of the top-level service and all child services is `stopped`, proceed to the next step.
 - If the status of the top-level service and all child services is **not** `stopped`, perform the remaining substeps.
- b Stop the top-level service.
Replace *Service* with the name or identifier of the top-level service:

```
serviced service stop Service
```

- c Monitor the stop.

```
serviced service status
```

When the status of the top-level service and all child services is stopped, proceed to the next step.

- 4 Stop Control Center with the cluster management tool.

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

- 5 Monitor the status of cluster resources.

```
watch pcs status
```

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

- 6 Ensure that no containers remain in the local repository.

- a Start the Docker service.

```
systemctl start docker
```

- b Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

If the command returns a result, enter the following command:

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- c Stop the Docker service.

```
systemctl stop docker
```

- 7 To ensure that no containers remain in both Docker repositories, log in to the other master node as `root`, or as a user with superuser privileges, and then perform the preceding step.

Stopping a delegate host

Use this procedure to stop the Control Center service (`serviced`) on a delegate host in a multi-host deployment. Repeat this procedure on each delegate host in your deployment.

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

```
systemctl stop serviced
```

- 3 Ensure that no containers remain in the local repository.

- a Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, proceed to the next step.
- If the command returns a result, perform the following substeps.

- b Remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, the most likely cause is an NFS conflict. Perform the following substeps.

- c** Stop the NFS and Docker services.

```
systemctl stop nfs && systemctl stop docker
```

- d** Start the NFS and Docker services.

```
systemctl start nfs && systemctl start docker
```

- e** Repeat the attempt to remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, perform the remaining substeps.

- f** Disable the automatic startup of `serviced`.

```
systemctl disable serviced
```

- g** Reboot the host.

```
reboot
```

- h** Log in to the delegate host as `root`, or as a user with superuser privileges.

- i** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

4 Dismount all filesystems mounted from the Control Center master host.

This step ensures no stale mounts remain when the storage on the master host is replaced.

- a** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- b** Force the filesystems to dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
  umount -f $FS
done
```

- c** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- d** Perform a lazy dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
    umount -f -l $FS
done
```

- e** Restart the NFS service.

```
systemctl restart nfs
```

- f** Determine whether any filesystems remain mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the remaining substeps.

- g** Disable the automatic startup of serviced.

```
systemctl disable serviced
```

- h** Reboot the host.

```
reboot
```

- i** Log in to the delegate host as `root`, or as a user with superuser privileges.
- j** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

Upgrading master hosts from 1.1.x to 1.5.1

25

Perform the procedures in this chapter to upgrade both Control Center 1.1.x master hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Adding storage for audit logging

The default configuration of Control Center audit logging requires 10GB of storage in `/var/log/serviced` on master host nodes. For high-availability deployments, Zenoss recommends mounting remote storage instead of adding local storage to master nodes. Use this procedure to create the mount point.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Create a mount point for audit logs.

The default location for `serviced` audit logs is `/var/log/serviced`. The location is determined by the `SERVICED_LOG_PATH` variable in `/etc/default/serviced`.

```
mkdir -p /var/log/serviced
```

- 3 Mount a remote filesystem at `/var/log/serviced`.
For more information about mounting a remote filesystem, refer to your operating system documentation.

Updating to Docker CE

Use this procedure to update Docker to Docker Community Edition (Docker CE) version 17.09.0.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Remove the Docker repository description file.
Upgrades no longer require external repositories.

```
rm -f /etc/yum.repos.d/docker.repo
```

- 3 Update the operating system, if necessary.
 - a Determine which release is installed.

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

- If the result is 7.2, 7.3, or 7.4, proceed to the next step.

- If the result is 7.1, perform the remaining substeps.

b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

c Update the operating system, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update && reboot
```

d Log in as `root`, or as a user with superuser privileges.

e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

4 Update the Linux kernel, if necessary.

a Determine which kernel version is installed.

```
uname -r
```

If the result is lower than `3.10.0-327.22.2.el7.x86_64`, perform the following substeps.

b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

c Update the kernel, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update kernel && reboot
```

d Log in as `root`, or as a user with superuser privileges.

e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

5 Identify the name of the LVM thin pool for Docker.

a Start the Docker service.

```
systemctl start docker
```

b Display the name of the LVM thin pool for Docker.

```
docker info 2>/dev/null | grep 'Pool Name'
```

Example result:

```
Pool Name: docker-docker--pool
```

Record the name for use in a subsequent step.

c Stop the Docker service.

```
systemctl stop docker
```

- 6 Back up the Docker environment file.

```
test -f /etc/sysconfig/docker \
  && mv /etc/sysconfig/docker /etc/sysconfig/docker.bak
```

- 7 Stop the Docker service.

```
systemctl stop docker
```

- 8 Remove Docker 1.9.0.

- a Remove without checking dependencies.

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

- b Clean the yum databases.

```
yum clean all
```

- 9 Install Docker CE 17.09.0.

```
yum install --enablerepo=zenoss-mirror docker-ce-17.09.0.ce
```

If yum returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 10 Start the Docker service.

```
systemctl start docker
```

- 11 Remove a previous workaround for an NFS rpcbind issue.

- a Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
  && rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

- b Reload the systemd manager configuration.

```
systemctl daemon-reload
```

Configuring Docker

Use this procedure to configure Docker.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Create a new Docker drop-in file.
 - a Create a directory for the drop-in file, if necessary.

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

- b Create a backup of the drop-in file, if it exists.

```
test -f /etc/systemd/system/docker.service.d/docker.conf \
  && cp -p /etc/systemd/system/docker.service.d/docker.conf \
  /etc/systemd/system/docker.service.d/docker.conf.bak
```

- c Create the new file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=~/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/dockerd \${OPTIONS}
TasksMax=infinity
EOF
```

- 3 Reload the systemd manager configuration.

```
systemctl daemon-reload
```

- 4 Configure and start the Docker service.

- a Create a variable for the name of the Docker thin pool.

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

```
myPool="Thin-Pool-Device"
```

- b Create variables for adding arguments to the Docker environment file. The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

```
myDriver="--storage-driver devicemapper"
myLog="--log-level=error"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myMount="--storage-opt dm.mountopt=discard"
myFlag="--storage-opt dm.thinpooldev=$myPool"
```

- c Add the arguments to the Docker configuration file.

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

- d Start or restart Docker.

```
systemctl restart docker
```

The startup may take up to a minute as Docker updates the storage layout.

- 5 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:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/16
```

For example, if the bridge subnet is 172.17.0.1, add the following flags:

```
--dns=172.17.0.1 --bip=172.17.0.1/16
```

Note Use a space character () to separate flags, and make sure the double quote character (") delimits the declaration of *OPTIONS*.

- d Save the file, and then close the editor.
- e Restart the Docker service.

```
systemctl restart docker
```

- 6 Remove a previous workaround for an NFS rpcbind issue.

- a Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

- b Reload the systemd manager configuration.

```
systemctl daemon-reload
```

- 7 Compare the previous version of the Docker environment file with the new version, and add any customizations for your deployment to the new version.

```
diff /etc/sysconfig/docker /etc/sysconfig/docker.bak
```

If you change this file, restart Docker with `systemctl restart docker`.

Loading image files

Use this procedure to load images into the local registry.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Change directory to `/root`.

```
cd /root
```

- 3 Load the images.

```
for image in install-zenoss-*.run
do
  /bin/echo -en "\nLoading $image..."
  yes | ./$image
done
```

- 4 List the images in the registry.

```
docker images
```

The result should show one image for each archive file.

- 5 Optional: Delete the archive files, if desired.

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

- 6 Stop the Docker service.

```
systemctl stop docker
```

Updating the Pacemaker resource agents

Use this procedure to update the Pacemaker resource agents for Control Center.

- 1 Log in to a master host node as `root`, or as a user with superuser privileges.
- 2 In a separate terminal session, log in to the other master host node as `root`, or as a user with superuser privileges.
- 3 On both nodes, install the latest version of the resource agents package.

```
myPath=/opt/zenoss-repo-mirror
yum install \
  $myPath/serviced-resource-agents-1.1.0-1.x86_64.rpm
```

Updating Control Center on the master host

Use this procedure to update Control Center on the master host to version 1.5.1.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.
 - a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \
  /opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.
 - a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

For more information about configuring the master host, see [Control Center configuration variables](#) on page 146.

Upgrading delegates from 1.1.x to 1.5.1

26

Perform the procedures in this chapter to upgrade Control Center 1.1.x delegate hosts to 1.5.1. First, download and stage the required files. For more information, see [Downloading and staging required files](#) on page 13.

Updating to Docker CE

Use this procedure to update Docker to Docker Community Edition (Docker CE) version 17.09.0.

- 1 Log in as `root`, or as a user with superuser privileges.
- 2 Remove the Docker repository description file.

Upgrades no longer require external repositories.

```
rm -f /etc/yum.repos.d/docker.repo
```

- 3 Update the operating system, if necessary.

- a Determine which release is installed.

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

- If the result is 7.2, 7.3, or 7.4, proceed to the next step.
- If the result is 7.1, perform the remaining substeps.

- b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

- c Update the operating system, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update && reboot
```

- d Log in as `root`, or as a user with superuser privileges.
- e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

- 4 Update the Linux kernel, if necessary.

- a Determine which kernel version is installed.

```
uname -r
```

If the result is lower than 3.10.0-327.22.2.el7.x86_64, perform the following substeps.

- b Disable automatic start of `serviced`.

```
systemctl disable serviced
```

- c Update the kernel, and then restart the host.

The following commands require internet access or a local mirror of operating system packages.

```
yum makecache fast && yum update kernel && reboot
```

- d Log in as `root`, or as a user with superuser privileges.
- e Enable automatic start of `serviced`.

```
systemctl enable serviced
```

- 5 Identify the name of the LVM thin pool for Docker.

- a Start the Docker service.

```
systemctl start docker
```

- b Display the name of the LVM thin pool for Docker.

```
docker info 2>/dev/null | grep 'Pool Name'
```

Example result:

```
Pool Name: docker-docker--pool
```

Record the name for use in a subsequent step.

- c Stop the Docker service.

```
systemctl stop docker
```

- 6 Back up the Docker environment file.

```
test -f /etc/sysconfig/docker \
&& mv /etc/sysconfig/docker /etc/sysconfig/docker.bak
```

- 7 Stop the Docker service.

```
systemctl stop docker
```

- 8 Remove Docker 1.9.0.

- a Remove without checking dependencies.

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

- b Clean the yum databases.

```
yum clean all
```

9 Install Docker CE 17.09.0.

```
yum install --enablerepo=zenoss-mirror docker-ce-17.09.0.ce
```

If yum returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

10 Start the Docker service.

```
systemctl start docker
```

11 Remove a previous workaround for an NFS rpcbind issue.**a** Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

b Reload the systemd manager configuration.

```
systemctl daemon-reload
```

Configuring Docker

Use this procedure to configure Docker.

1 Log in as root, or as a user with superuser privileges.**2** Create a new Docker drop-in file.**a** Create a directory for the drop-in file, if necessary.

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

b Create a backup of the drop-in file, if it exists.

```
test -f /etc/systemd/system/docker.service.d/docker.conf \
&& cp -p /etc/systemd/system/docker.service.d/docker.conf \
/etc/systemd/system/docker.service.d/docker.conf.bak
```

c Create the new file.

```
cat <<EOF > /etc/systemd/system/docker.service.d/docker.conf
[Service]
TimeoutSec=300
EnvironmentFile=-/etc/sysconfig/docker
ExecStart=
ExecStart=/usr/bin/dockerd $OPTIONS
TasksMax=infinity
EOF
```

3 Reload the systemd manager configuration.

```
systemctl daemon-reload
```

4 Configure and start the Docker service.**a** Create a variable for the name of the Docker thin pool.

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

```
myPool="Thin-Pool-Device"
```

- b** Create a variable for the master host.

Replace *Master* with the fully-qualified domain name (*HA-Virtual-Name*) or IPv4 address (*HA-Virtual-IP*) of the high-availability cluster:

```
myCluster="Master"
```

The value of this variable must match the value of the *SERVICED_DOCKER_REGISTRY* variable in `/etc/default/serviced`.

- c** Create variables for adding arguments to the Docker environment file. The `--exec-opt` argument is a workaround for [a Docker issue](#) on RHEL/CentOS 7.x systems.

```
myDriver="--storage-driver devicemapper"
myLog="--log-level=error"
myFix="--exec-opt native.cgroupdriver=cgroupfs"
myMount="--storage-opt dm.mountopt=discard"
myFlag="--storage-opt dm.thinpooldev=$myPool"
myReg="--insecure-registry=$myCluster:5000"
```

- d** Add the arguments to the Docker configuration file.

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

- e** Start or restart Docker.

```
systemctl restart docker
```

The startup may take up to a minute as Docker updates the storage layout.

5 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:

```
--dns=Bridge-Subnet --bip=Bridge-Subnet/16
```

For example, if the bridge subnet is 172.17.0.1, add the following flags:

```
--dns=172.17.0.1 --bip=172.17.0.1/16
```

Note Use a space character () to separate flags, and make sure the double quote character (") delimits the declaration of *OPTIONS*.

- d Save the file, and then close the editor.
- e Restart the Docker service.

```
systemctl restart docker
```

- 6 Remove a previous workaround for an NFS `rpcbind` issue.
 - a Remove the NFS service drop-in file, if it exists.

```
test -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf \
&& rm -f /etc/systemd/system/nfs-server.service.d/nfs-server.conf
```

- b Reload the `systemd` manager configuration.

```
systemctl daemon-reload
```

- 7 Compare the previous version of the Docker environment file with the new version, and add any customizations for your deployment to the new version.

```
diff /etc/sysconfig/docker /etc/sysconfig/docker.bak
```

If you change this file, restart Docker with `systemctl restart docker`.

Configuring NFS 4.0

Use this procedure to configure NFS 4.0 on delegate hosts if the operating system release is 7.4. There may be a file locking defect in NFS 4.1 with RHEL/CentOS 7.4.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Determine which release is installed.

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

- If the result includes 7.4, perform the remaining steps of this procedure.
- If the result includes 7.2 or 7.3, continue to the next procedure.

- 3 Change the NFS configuration file.

- a Open `/etc/nfsmount.conf` with a text editor.
- b Locate the `Defaultvers` directive.
- c Remove the number sign character (`#`) from the beginning of the line.
- d Change the value from 4 to 4.0.
The line should appear as follows:

```
Defaultvers=4.0
```

- e Save the file, and then close the editor.
- 4 Restart the NFS server.

```
systemctl restart nfs-server
```

Updating Control Center on delegate hosts

This procedure updates Control Center on delegate hosts to version 1.5.1.

Perform this procedure on each delegate host in your deployment.

- 1 Log in to a delegate host as `root`, or as a user with superuser privileges.
- 2 Save the current `serviced` configuration file as a reference.
 - a Rename the file.

```
mv /etc/default/serviced /etc/default/serviced-pre-1.5.1
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-pre-1.5.1
```

- 3 Install the new version of Control Center.

```
yum install --enablerepo=zenoss-mirror \
/opt/zenoss-repo-mirror/serviced-1.5.1-1.x86_64.rpm
```

If `yum` returns an error due to dependency issues, see [Resolving package dependency conflicts](#) on page 137 for potential resolutions.

- 4 Make a backup copy of the new configuration file.
 - a Copy the file.

```
cp /etc/default/serviced /etc/default/serviced-1.5.1-orig
```

- b Set permissions to read-only.

```
chmod 0440 /etc/default/serviced-1.5.1-orig
```

- 5 Display the settings of the reference configuration file.

```
grep -E '^b*[A-Z_]+' /etc/default/serviced-pre-1.5.1
```

- 6 Open the new configuration file with a text editor, and then update the file for your environment.

Among other changes, the `SERVICED_NFS_CLIENT` variable is deprecated. Make note of the delegates that use the setting, and then rescind DFS access permission from the resource pool to which they belong. For more information, see [Updating resource pool permissions](#) on page 109.

For more information about configuring a delegate host, see [Control Center configuration variables](#) on page 146.

Updating the ZooKeeper image on ensemble nodes

Use this procedure to install a new Docker image for ZooKeeper into the local repository of delegate hosts that are ZooKeeper ensemble nodes.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the hosts in the ZooKeeper ensemble.

```
grep -E '^b*SERVICED_ZK=' /etc/default/serviced
```

The result is a list of 3 or 5 hosts, separated by the comma character (,). The master host is always a node in the ZooKeeper ensemble.

- 3 Log in to a ZooKeeper ensemble node as `root`, or as a user with superuser privileges.

- 4 Change directory to `/root`.

```
cd /root
```

- 5 Extract the ZooKeeper image.

```
./install-zenoss-isvcs-zookeeper-v10.run
```

At the prompt, press **y**.

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

```
rm -i ./install-zenoss-isvcs-zookeeper-v10.run
```

- 7 Repeat the preceding four steps on each ZooKeeper node in the ensemble.

Starting an upgraded deployment

This chapter includes the procedure for starting a high-availability Control Center deployment after upgrading to version 1.5.1.

Registering the master host nodes

Use this procedure to register authentication credentials for the master host nodes.

- 1 Log in to a master host node as `root`, or as a user with superuser privileges.
- 2 In a separate window, log in to the other master host node as `root`, or as a user with superuser privileges.
- 3 On both nodes, verify the settings in the `serviced` configuration file.

```
grep -E '^\b*[A-Z_]+' /etc/default/serviced | sort
```

The settings must be identical on both hosts.

- 4 Start each delegate host that is a member of the ZooKeeper ensemble.
The ensemble is needed to start the master host instance of Control Center.

- a Identify the hosts.

```
grep -E '^\b*SERVICED_ZK=' /etc/default/serviced
```

The result is a list of 3 or 5 hosts, separated by the comma character (,). The master host is always a node in the ZooKeeper ensemble.

- b Log in to each delegate host in the ensemble as `root`, or as a user with superuser privileges.
- c On each delegate host, start Control Center.

```
systemctl start serviced
```

- 5 On either node, start the cluster, and then monitor the startup process.
 - a Start the cluster.

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

- b Monitor the startup process.

```
watch pcs status
```

Do not proceed until all resources report `Started`.

- c Identify the active node in the cluster.
The information is included in the output of `pcs status`.
- 6 On the active node, display the host ID, and then register the host.
 - a Display the host ID of the active host.

```
hostid
```

- b Register the active master host.
Replace *Host-ID* with the host ID of the active node:

```
serviced key reset --register Host-ID
```

- 7 On the active node, fail over to the passive node, and then monitor the process.
 - a Display the status of mirrored storage areas.

```
drbd-overview
```

Do not proceed until the status of all devices is `UpToDate/UpToDate`.

- b Fail over to the passive node.
Replace *Active-Node* with the IP address or hostname of the active master host:

```
pcs cluster standby Active-Node
```

- c Monitor the failover process.

```
watch pcs status
```

The failover process takes several minutes. Do not proceed until all resources report `Started`.

- d Display the status of mirrored storage areas.

```
drbd-overview
```

Do not proceed until the status of all devices is `UpToDate/UpToDate`.

- 8 On the formerly-passive node, register the host.
 - a Display the host ID of the now-active host.

```
hostid
```

- b Register the now-active master host.
Replace *Host-ID* with the host ID of the now-active node:

```
serviced key reset --register Host-ID
```

- 9 On either node, restore the cluster.
 - a Restore the cluster.
Replace *Former-Active-Node* with the IP address or hostname of the now-passive master host:

```
pcs cluster unstandby Former-Active-Node
```

- b Monitor the startup process.

```
watch pcs status
```

Do not proceed until all resources report `Started`.

- c Display the status of mirrored storage areas.

```
drbd-overview
```

Do not proceed until the status of all devices is UpToDate/UpToDate.

Updating delegate hosts with authentication

Starting with version 1.3.0, Control Center requires authentication tokens for all delegate communications. The tokens are based on RSA key pairs created by the master `serviced` instance. When you create a key pair for a delegate, `serviced` bundles its public key with the delegate's private key. The `serviced` instance on the delegate installs the credentials and uses them to sign messages with the required unique tokens.

Credentials are installed by using an SSH connection or a file.

- The command to create a key pair can initiate an SSH connection with a delegate and install credentials. This option is the most secure, because no file is created. However, it requires either public key authentication or password authentication between the master and delegate hosts.
- When no SSH connection is requested, the command to create a key pair creates a file containing the credentials. You can move the credentials file to the delegate host with any file transfer method, and then install it on the delegate.

The following procedures demonstrate how to create credentials and install them on a delegate.

Starting and registering a delegate using SSH

To succeed, the following statements about the login account used to perform this procedure must be true:

- The account exists on both the master host and on the delegate host.
- The account has `serviced` CLI privileges.
- The account has either public key authentication or password authentication enabled on the master host and on the delegate host.

Use this procedure to start a delegate host after upgrading Control Center to version 1.5.1. This procedure also includes steps to create and register the authentication credentials the delegate needs, through an SSH connection.

- 1 Log in to the delegate host as `root`, or as a user with superuser privileges.
- 2 Determine whether `serviced` is configured to start when the system starts.

```
systemctl is-enabled serviced
```

- If the result is `enabled`, proceed to the next step.
- If the result is `disabled`, enter the following command:

```
systemctl enable serviced
```

- 3 Verify the settings in the `serviced` configuration file.

```
grep -E '^b*[A-Z_]+' /etc/default/serviced
```

- 4 Start `serviced`, and then monitor the startup.

During this startup, `serviced` invokes `docker pull` to retrieve its updated images.

```
systemctl daemon-reload && systemctl start serviced \
  && journalctl -u serviced -f -o cat
```

Do not proceed to the next step until the following message is displayed:

```
Host Agent successfully started
```

- 5 Log out of the delegate host.
- 6 Log in to the master host as `root`, or as a user with superuser privileges.
- 7 Obtain the host ID of the delegate host started previously.
 - a Display the host IDs of all Control Center hosts.

```
serviced host list | cut -c-85
```

- b Record the host ID of the delegate host.
- 8 Create authentication credentials for the delegate host, and register the credentials.

If the master and delegate host are configured for key-based access, the following command does not prompt you to add the delegate to the list of known hosts or to provide the password of the remote user account.

Replace *Host-ID* with the host ID of the delegate host started previously:

```
serviced key reset --register Host-ID
```

Starting and registering a delegate using a file

Use this procedure to start a delegate host after upgrading Control Center to version 1.5.1. This procedure also includes steps to create a credentials file and to use the file to register the delegate.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Obtain the host ID of the delegate host to start and register.
 - a Display the host IDs of all Control Center hosts.

```
serviced host list | cut -c-85
```

- b Record the host ID of the delegate host.
- 3 Create authentication credentials for the delegate host.

Replace *Host-ID* with the host ID of the delegate host identified in the preceding step:

```
serviced key reset Host-ID
```

The command creates a unique credentials file in the local directory.

- 4 Use a file transfer utility such as `scp` to copy the credentials file to the delegate host.

Once copied to the delegate host, the credentials file is not needed on the master host and can be deleted.
- 5 Log in to the Control Center delegate host as `root`, or as a user with superuser privileges.
- 6 Determine whether `serviced` is configured to start when the system starts.

```
systemctl is-enabled serviced
```

- If the result is enabled, proceed to the next step.

- If the result is disabled, enter the following command:

```
systemctl enable serviced
```

- 7 Verify the settings in the `serviced` configuration file.

```
grep -E '^\\b*[A-Z_]+' /etc/default/serviced
```

- 8 Start `serviced`, and then monitor the startup.

During this startup, `serviced` invokes `docker pull` to retrieve its updated images.

```
systemctl daemon-reload && systemctl start serviced \  
&& journalctl -u serviced -f -o cat
```

Do not proceed to the next step until the following message is displayed:

```
Host Agent successfully started
```

- 9 Install the credentials.

Replace *Credentials-File* with the pathname of the credentials file:

```
serviced host register Credentials-File
```

- 10 Delete the credentials file.

The file is no longer needed on the delegate host.

Replace *Credentials-File* with the pathname of the credentials file:

```
rm Credentials-File
```

After upgrading from 1.1.x to 1.5.1

Perform the procedures in this chapter after Control Center is upgraded.

Updating resource pool permissions

Use this procedure to identify and change the permissions associated with one or more resource pools. During the upgrade to version 1.5.1, existing resource pools are assigned both administrative and DFS permissions.

In this release, the command that displays information about pools uses integer values for the Permissions field. The following list associates the values with the permissions they represent:

- 1, administrative permission
- 2, DFS access permission
- 3, both administrative and DFS access permissions

- 1 Log in to the master host as a user with `serviced` CLI privileges.
- 2 Display the list of resource pools and their permissions.

```
serviced pool list -v | grep -E 'ID|Permissions'
```

Example result:

```
"ID": "default",
"Permissions": 3,
"ID": "master",
"Permissions": 3,
```

In the preceding example, the `default` and `master` resource pools have administrative permission and DFS access permission.

- 3 Optional: Remove DFS access permission from a pool.

If you intend to remove both DFS access and administrative access permissions from a resource pool, you must remove DFS access permissions first.

Replace *Pool-Name* with the name of a resource pool from the previous step:

```
serviced pool set-permission --dfs=false Pool-Name
```

- 4 Optional: Remove administrative permission from a pool.

If you intend to remove both DFS access and administrative access permissions from a resource pool, you must remove DFS access permissions first.

Replace *Pool-Name* with the name of a resource pool from a previous step:

```
serviced pool set-permission --admin=false Pool-Name
```

Setting the connection timeout of a resource pool

Use this procedure to set the length of time the services scheduler waits for a disconnected delegate host to reconnect before moving its services to a different host in the resource pool. This affects all multi-host resource pools. Zenoss recommends using a minimum value of 15 seconds for all multi-host pools. For resource pools that are connected through high-latency, wide-area networks, Zenoss recommends values greater than 15 seconds.

- 1 Log in to the master host as a user with `serviced` CLI privileges.
- 2 Display the list of resource pools and their connection timeout values.

```
serviced pool list -v | grep -E 'ID|ConnectionTimeout'
```

- 3 Optional: Set the connection timeout value of a resource pool.

This command accepts the following units identifiers:

- ms (milliseconds)
- s (seconds)
- m (minutes)
- h (hours)

Replace *Pool-ID* with a resource pool identifier, and replace *Timeout+Units* with an integer followed by a units identifier:

```
serviced pool set-conn-timeout Pool-ID Timeout+Units
```

Removing unused images

Use this procedure to identify and remove unused Control Center images.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the images associated with the installed version of `serviced`.

```
serviced version | grep Images
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v61 zenoss/isvcs-zookeeper:v10]
```

- 3 Start Docker.

```
systemctl start docker
```

- 4 Display the `serviced` images in the local repository.

```
docker images | awk '/REPO|isvcs/'
```

Example result (edited to fit):

```
REPOSITORY          TAG          IMAGE ID
zenoss/serviced-isvcs v40         88cd6c24cc82
zenoss/serviced-isvcs v61         0aab5a2123f2
zenoss/isvcs-zookeeper v3          46fa0a2fc4bf
zenoss/isvcs-zookeeper v10        0ff3b3117fb8
```

The example result shows the current versions and one set of previous versions. Your result may include additional previous versions and will show different images IDs.

5 Remove unused images.

Replace *Image-ID* with the image ID of an image for a previous version.

```
docker rmi Image-ID
```

Repeat this command for each unused image.

6 Stop Docker.

```
systemctl stop docker
```

Removing orphaned snapshot devices

Use this procedure to identify orphaned snapshot devices in the LVM thin pool for application data, and to remove them.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Identify the location of tenant volumes.

```
grep -E '^\\b*SERVICED_VOLUMES_PATH' /etc/default/serviced
```

- If the command returns a result, the location of tenant volumes is the value of the `SERVICED_VOLUMES_PATH` variable.
- If the command does not return a result, the location of tenant volumes is the default value of `SERVICED_VOLUMES_PATH`, `/opt/serviced/var/volumes`.

- 3 Identify the device of the `serviced` thin pool belongs.

```
ls /dev/mapper | grep serviced
```

Example result:

```
serviced-serviced--pool
serviced-serviced--pool_tdata
serviced-serviced--pool_tmeta
```

The first result is the thin pool device. The other results represent the data and metadata portions of the device.

- 4 Check for orphaned snapshot devices.

Replace *Thin-Pool-Device* with the name of the thin pool device from the previous step, and replace *Volumes-Path* with the location of tenant volumes:

```
serviced-storage -o dm.thinpooldev=/dev/mapper/Thin-Pool-Device \
check Volumes-Path
```

- If the result is `No orphaned devices found`, stop. There are no orphaned snapshot devices to remove.
 - If the result is `Orphaned devices were found`, perform the next step.
- 5 Remove orphaned snapshot devices.
Replace *Thin-Pool-Device* with the name of the thin pool device from the previous step, and replace *Volumes-Path* with the location of tenant volumes:

```
serviced-storage -o dm.thinpooldev=/dev/mapper/Thin-Pool-Device \  
check -c Volumes-Path
```

Updating the OpenTSDB time-to-live value

Older versions of Control Center used a longer time-to-live (TTL) value for data maintained by the OpenTSDB database. The correct value for this version is 2592000 seconds (30 days). Use this procedure to update the TTL value, if necessary.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Start an interactive shell in the OpenTSDB container.

```
docker exec -it serviced-isvcs_opentsdb bash
```

- 3 Start an interactive HBase shell.

```
/opt/hbase/bin/hbase shell
```

Example result:

```
HBase Shell; enter 'help<RETURN>' for list of supported commands.  
Type "exit<RETURN>" to leave the HBase Shell  
Version 0.94.16, r1557241, Fri Jan 10 20:43:03 UTC 2014
```

```
hbase(main):001:0>
```

- 4 Display the current settings of the `tsdb` table.

```
describe 'tsdb'
```

- If the result includes `TTL => '2592000'`, stop. Use the `exit` command twice, to end the HBase shell and then the shell in the OpenTSDB container.
 - If the result includes a larger value for the TTL setting, perform the remaining steps.
- 5 Disable the `tsdb` table.

```
disable 'tsdb'
```

- 6 Set the TTL value to 2592000 seconds (30 days).

```
alter 'tsdb', {NAME=>'t', TTL=>'2592000'}
```

- 7 Enable the `tsdb` table.

```
enable 'tsdb'
```


- 8 Display the current settings.

```
describe 'tsdb'
```

- If the result includes `TTL => '2592000'`, proceed to the next step.
 - If the result includes a different value for the TTL setting, repeat the preceding steps.
- 9 End the interactive HBase shell.

```
exit
```

- 10 End the interactive shell in the OpenTSDB container.

```
exit
```

Managing maintenance scripts

The following topics provide information about the maintenance scripts that are installed when Control Center is installed on a host.

Control Center maintenance scripts on the master host

The scripts in the following list are installed when Control Center is installed, and are started either daily or weekly by `anacron`.

`/etc/cron.hourly/serviced`

This script invokes `logrotate` hourly, to manage the files in `/var/log/serviced`.

This script is required on the master host only.

`/etc/cron.weekly/serviced-fstrim`

This script invokes `fstrim` weekly, to reclaim unused blocks in the application data thin pool.

The life span of a solid-state drive (SSD) degrades when `fstrim` is run too frequently. If the block storage of the application data thin pool is an SSD, you can reduce the frequency at which this script is invoked, as long as the thin pool never runs out of free space. An identical copy of this script is located in `/opt/serviced/bin`.

This script is required on the master host only.

`/etc/cron.d/cron_zenossdbpack`

This script invokes `/opt/serviced/bin/serviced-zenossdbpack`, the database maintenance script for a Zenoss application, every Sunday at midnight. If the Zenoss application is not installed or is offline, the command fails. You can change the day of the week and time of day when the maintenance script is invoked by editing `/etc/cron.d/cron_zenossdbpack`.

This script is required on one host in the resource pool in which the Zenoss database services run. For more information, refer to the configuration guide for your Zenoss application.

Control Center maintenance scripts on delegate hosts

The scripts in the following list are installed when Control Center is installed, and are started either daily or weekly by `anacron`.

Of these scripts, only the first is required on delegate hosts. The others should be removed.

/etc/cron.hourly/serviced

This script invokes `logrotate` hourly, to manage the `/var/log/serviced-audit.log` and `/var/log/application-audit.log` files.

This script should be removed from delegate hosts.

/etc/cron.weekly/serviced-fstrim

This script should be removed from delegate hosts.

/etc/cron.d/cron_zenosdbpack

This script is required on one host in the resource pool in which the Zenoss database services run. For more information, refer to the configuration guide for your Zenoss application.

Control Center application data storage requirements



Control Center uses an LVM thin pool to store tenant (application) data. LVM thin pools include separate storage areas for data and for metadata. For each tenant it manages, Control Center maintains a separate virtual device (a volume) in the data storage area of its LVM thin pool. Also, Control Center creates virtual devices in the data storage area for snapshots of tenant data.

To ensure consistency, Control Center requires the following, minimum amount of free space in its thin pool:

- 3GiB available for each tenant volume
- 3GiB available in the data storage area
- 62MiB available in the metadata storage area
- The total amount of metadata storage must be 1% of total data storage

Examining application data storage status

Beginning with release 1.3.0, Control Center initiates an emergency shutdown when the minimum required amounts of free space are not available in the `serviced` thin pool or tenant volumes. For more information, see [Control Center application data storage requirements](#) on page 115.

Use this procedure to display the amount of free space in a Control Center thin pool and tenant volumes, to determine how much space is available in the LVM volume group that contains the thin pool, and to determine whether additional steps are required.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Display the amount of space available in the `serviced` thin pool.

```
serviced volume status
```

Figure 1: Example output with highlighted values

```

Status for volume /opt/serviced/var/volumes:

Driver:      devicemapper
Driver Type: direct-lvm
Volume Path: /opt/serviced/var/volumes

Thin Pool
-----
Logical Volume:      serviced-serviced--pool
Metadata (total/used/avail):  2.2 GiB / 64.18 MiB (2.85%) / 2.137 GiB (97.15%)
Data (total/used/avail):      200 GiB / 30.576 GiB (15%) / 169.424 GiB (85%)

3ir8lrglh8b09ipowynz3qx2f Application Data
-----
Volume Mount Point:      /opt/serviced/var/volumes/3ir8lrglh8b09ipowynz3qx2f
Filesystem (total/used/avail): 98.31 GiB / 1.511 GiB (1.5%) / 91.78 GiB (93%)
Virtual device size:      100 GiB

```

The result includes detailed information about the `serviced` thin pool and each tenant volume.

- If the amount of free space in the `serviced` thin pool is sufficient, stop. No further action is required.
 - If the amount of free space in the data or metadata portions of the `serviced` thin pool is not sufficient, perform the following steps.
- 3 Identify the volume group to which the `serviced` thin pool belongs.

```
lvs --options=lv_name,vg_name,lv_size
```

The volume group associated with `serviced-pool` contains the `serviced` thin pool.

- 4 Display the amount of free space in the volume group that contains the `serviced` thin pool. Replace *Volume-Group* with the name of the volume group identified in the previous step:

```
vgs --no-headings --options=vg_free Volume-Group
```

- If the amount of free space in the volume group is not sufficient to increase one or both of the storage areas of the `serviced` thin pool to their required minimums, add physical or logical storage to the volume group. For more information, see [Adding space to the data area of a Control Center thin pool](#) on page 117.
 - If the amount of free space in the volume group is sufficient to increase one or both of the storage areas of the `serviced` thin pool to their required minimums, proceed to the next step.
- 5 Increase the free space in one or both areas of the `serviced` thin pool.
 - To increase the amount of space in the metadata area, proceed to [Adding space to the metadata area of a Control Center thin pool](#) on page 116.
 - To increase the amount of space in the data area, proceed to [Adding space to the data area of a Control Center thin pool](#) on page 117.

Adding space to the metadata area of a Control Center thin pool

Use this procedure to increase the amount of space in the metadata area of a Control Center thin pool.

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

- 2 Display information about LVM logical volumes on the host.

```
lvs --options=lv_name,vg_name,lv_size
```

Typically, the logical volume is `serviced-pool` and the containing volume group is `serviced`.

- 3 Add space to the metadata storage area of the `serviced` thin pool.
In the following command:

- Replace *Size* with the amount of space to add (in megabytes) and the units identifier (M).
- Replace *Volume-Group* with the name of the LVM volume group identified in the previous step.
- Replace *Logical-Volume* with the name of the logical volume identified in the previous step.

```
lvextend -L+SizeM Volume-Group/Logical-Volume_tmeta
```

Adding space to the data area of a Control Center thin pool

Use the procedures in this section to increase the amount of space in the data area of a Control Center thin pool. For more information about required minimums, see [Control Center application data storage requirements](#) on page 115. Perform the procedures in this section in order.

Note Before proceeding, use the Control Center backup feature to perform a complete backup. Creating a backup is required. One of the procedures in this section manipulates the volume group partition, which could cause a loss of data if performed incorrectly. You must have a valid backup before proceeding.

Identifying node roles

Use this procedure to identify and record the role of each node in the high-availability cluster.

- 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 status of cluster resources.

```
pcs status resources
```

The first section shows the status of the nodes.

- 3 Record the node roles or restart the cluster.
 - If the status of the nodes includes a list named `Masters` and a list named `Slaves`, record the information. In the procedures that follow, the node in the list named `Masters` is represented by the variable name *Active-Node*, and the node in the list named `Slaves` is represented by the variable name *Passive-Node*.
 - If the status of the nodes includes a list named `Stopped`, do not attempt to perform the procedures that follow. Instead, restart the cluster, and then repeat this procedure.

Stopping the cluster

Use this procedure to stop the applications that `serviced` is managing, `serviced` itself, and cluster services.

- 1 Log in to *Active-Node* as `root`, or as a user with superuser privileges.
- 2 Stop the top-level service `serviced` is managing, if necessary.
 - a Show the status of running services.

```
serviced service status
```

The top-level service is the service listed immediately below the headings line.

- If the status of the top-level service and all child services is `stopped`, proceed to the next step.
- If the status of the top-level service and all child services is **not** `stopped`, perform the remaining substeps.

b Stop the top-level service.

Replace *Service* with the name or identifier of the top-level service:

```
serviced service stop Service
```

c Monitor the stop.

```
serviced service status
```

When the status of the top-level service and all child services is `stopped`, proceed to the next step.

3 Disable `serviced`.

```
pcs resource disable serviced
```

4 Stop `serviced` on all delegate hosts.

For more information, see [Stopping a delegate host](#) on page 22.

5 Disable additional resources.

a Disable storage resources.

```
pcs resource disable serviced-storage
```

b Disable the `serviced` group.

```
pcs resource disable serviced-group
```

6 Confirm that all resources are stopped, and identify the master node.

a Confirm that all resources are stopped.

```
watch pcs status
```

7 Disable replication.

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

8 In a separate terminal session, log in to *Passive-Node* as `root`, or as a user with superuser privileges.

9 On both nodes, confirm that replication is disabled.

```
drbd-overview
```

The output should match the following example:

```
0:serviced-dfs/0 Unconfigured . .
1:serviced-dfs/1 Unconfigured . .
2:serviced-dfs/2 Unconfigured . .
```

10 On both nodes, display the size of the `serviced` thin pool device.

Typically, the `serviced` thin pool device is associated with `/dev/drbd2`. To verify the configuration in your environment, review `/etc/drbd.d/serviced-dfs.res`.

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

None of devices listed in the output should include subordinate DRBD devices.

Note When this procedure is complete, add storage to the device that contains the `serviced` thin pool. For more information, refer to your operating system or hypervisor documentation. Perform the procedure on both master host nodes.

Resizing the partition

Use this procedure to resize the partition that contains the `serviced` thin pool.

- 1 Log in to *Active-Node* as `root`, or as a user with superuser privileges.
- 2 In a separate terminal session, log in to *Passive-Node* as `root`, or as a user with superuser privileges.
- 3 On both nodes, inform the kernel of the partition table change that occurred when the size of the `serviced` thin pool device was increased.

```
partprobe
```

- 4 On both nodes, confirm the size increase of the raw device.

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

Subordinate device sizes are unchanged.

- 5 On both nodes, recreate the partition of the raw device.

The following substeps assume that the raw device has a single partition that contains all of the space in the device. If you are using a different layout, adjust the procedure as required.

- a Start the disk partition utility.
Replace *Raw-Device* with the path of the raw device associated with `/dev/drbd2` or its equivalent in your environment:

```
fdisk Raw-Device
```

- b Display the partition table.

```
p
```

- c Delete the partition table.

```
d
```

- d Create a new partition table.

```
n
```

Accept the defaults to create a new primary partition, and for the starting and ending cylinders of the partition. The starting cylinder should match the starting cylinder displayed previously. The ending cylinder should be greater than the ending cylinder displayed previously.

- e Write the new partition table, and then exit the disk partition utility.

```
w
```

- 6 On both nodes, confirm the size increase of the partition.

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

The subordinate device size should be larger.

Starting the cluster

Use this procedure to start cluster services.

- 1 Log in to *Active-Node* as `root`, or as a user with superuser privileges.
- 2 In a separate terminal session, log in to *Passive-Node* as `root`, or as a user with superuser privileges.
- 3 Start DRBD.
 - a On *Active-Node*, start DRBD.

```
drbdadm up all
```

The following example output is expected:

```
Moving the internal meta data to its proper location
Internal drbd meta data successfully moved.
```

- b On *Passive-Node*, start DRBD.

```
drbdadm up all
```

The output of this command should match the output on *Active-Node*.

- c On both nodes, confirm that the DRBD device for the `serviced` thin pool is larger.

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

- 4 On *Active-Node*, wait for disks to synchronize.

```
watch drbd-overview
```

Do not proceed until the status of all devices is `UpToDate/UpToDate`.

- 5 On *Active-Node*, start cluster services for the distributed file system.

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

- 6 On either node, confirm the identity of the primary node.

```
pcs status resources
```

The first section shows the status of the nodes.

- 7 On *Active-Node*, resize the volume group that contains the `serviced` thin pool.

- a Identify the volume group that contains the `serviced` thin pool.

```
lvs --options=lv_name, vg_name, lv_size
```

Typically, the logical volume is `serviced-pool` and the containing volume group is `serviced`.

- b Display information about the volume group.

Replace *Volume-Group* with the name of the volume group identified in the previous substep:

```
vgdisplay Volume-Group
```


- c Identify the DRBD disk associated with the serviced thin pool.

Typically, the `serviced` thin pool device is associated with `/dev/drbd2`. To verify the configuration in your environment, review `/etc/drbd.d/serviced-dfs.res`.

- d Resize the volume group.

Replace *DRBD-Device* with the DRBD device associated with the `serviced` thin pool:

```
pvresize DRBD-Device
```

- e Display information about the volume group.

Replace *Volume-Group* with the name of the volume group identified in the previous substep:

```
vgdisplay Volume-Group
```

The size of the volume group should be larger.

- 8 On *Active-Node*, resize the `serviced` thin pool.

- a Start the `serviced` group cluster services.

```
pcs resource enable serviced-group
```

- b Add space to the data storage area of the `serviced` thin pool.

In the following command:

- Replace *Total-Size* with the sum of the existing device size plus the space to add to the device, in gigabytes. Include the units identifier, G.
- Replace *Volume-Group* with the name of the LVM volume group identified in the previous step.
- Replace *Logical-Volume* with the name of the logical volume identified in the previous step.

```
lvextend -L+Total-SizeG Volume-Group/Logical-Volume
```

- c Display information about LVM logical volumes on the host.

```
lvs --options=lv_name,vg_name,lv_size
```

The result should show the larger size of the logical volume.

- 9 On *Active-Node*, start the storage resource.

- a Start the storage service.

```
pcs resource enable serviced-storage
```

- b Confirm that the resource started correctly.

```
pcs status
```

- 10 Optional: On *Active-Node*, increase the size of the application tenant volume, if desired.

- a Display the device mapper name of the `serviced` thin pool.

```
grep -E '^\\b*SERVICED_DM_THINPOOLDEV' /etc/default/serviced \
| sed -e 's/.*/=/'
```

Typically, the name is `/dev/mapper/serviced-serviced--pool`.

- b Increase the size of the tenant device.

In the following command:

- Replace *Device-Mapper-Name* with the device mapper name of the thin pool.

- Replace *Tenant-ID* with the identifier of the tenant device.
- Replace *Total-Size* with the sum of the existing device size plus the space to add to the device, in gigabytes. Include the units identifier, G.

```
serviced-storage resize -d /opt/serviced/var/volumes \
-o dm.thinpooldev=Device-Mapper-Name Tenant-ID Total-SizeG
```

Starting Control Center and verifying the cluster

Use this procedure to start Control Center and verify the cluster.

- 1 Log in to *Active-Node* as root, or as a user with superuser privileges.
- 2 Identify the hosts in the ZooKeeper ensemble.

```
grep -E '^\b*SERVICED_ZK=' /etc/default/serviced
```

The result is a list of 1, 3, or 5 hosts, separated by the comma character (,). The master host is always a node in the ZooKeeper ensemble.

- 3 In separate windows, log in to each of the delegate hosts that are nodes in the ZooKeeper ensemble as root, or as a user with superuser privileges.
- 4 On *Active-Node*, start the Control Center resource.

```
pcs resource enable serviced
```

- 5 On the other ZooKeeper ensemble hosts, start *serviced*.

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

```
systemctl start serviced
```

- 6 On *Active-Node*, confirm that *serviced* starts successfully.

```
journalctl -flu serviced
```

- 7 On *Active-Node*, stop the cluster.

- a Stop the *serviced* resource.

```
pcs resource disable serviced
```

- b Confirm that *serviced* stops successfully.

```
journalctl -flu serviced
```

- c Stop the storage resource.

```
pcs resource disable serviced-storage
```

- d Stop the *serviced* group.

```
pcs resource disable serviced-group
```

- e Wait for all cluster resources to stop.

```
watch pcs status
```

- 8 On *Active-Node*, fail over to *Passive-Node*.

- a Place *Active-Node* in standby mode.

```
pcs cluster standby Active-Node
```

- b Confirm that *Active-Node* switches roles.

```
pcs status resources
```

- 9 Confirm that the storage shows the correct size.

- a Start the serviced group.

```
pcs resource enable serviced-group
```

- b Confirm that the `serviced-lvm` resource starts.

```
pcs status
```

- c Confirm that the logical volume of the `serviced` thin pool shows the correct size.

```
lvs --options=lv_name,vg_name,lv_size
```

Complete the startup by keeping or changing the current roles:

- To keep *Passive-Node* in its current role as the master/primary node, proceed to [Starting the cluster and keeping current roles](#) on page 123.
- To make *Active-Node* the master/primary node, proceed to [Starting the cluster and restoring previous roles](#) on page 124.

Starting the cluster and keeping current roles

Use this procedure to complete the startup and keep *Passive-Node* in its current role as the master/primary node.

- 1 Log in to *Passive-Node* as `root`, or as a user with superuser privileges.
- 2 Start the storage resource.

```
pcs resource enable serviced-storage
```

- 3 Confirm that the resource starts up.

```
pcs status
```

- 4 Take *Active-Node* out of standby mode.
Replace *Active-Node* with the hostname assigned in [Identifying node roles](#) on page 117:

```
pcs cluster unstandby Active-Node
```

- 5 Confirm that disk replication is started.

```
drbd-overview
```

Do not proceed until the status of all devices is `UpToDate/UpToDate`.

- 6 Start the Control Center resource.

```
pcs resource enable serviced
```

- 7 Start your application.
For example, the Zenoss application:

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

Starting the cluster and restoring previous roles

Use this procedure to complete the startup and make *Active-Node* the master/primary node.

- 1 Log in to *Passive-Node* as `root`, or as a user with superuser privileges.
- 2 Stop the `serviced` resource group.

```
pcs resource disable serviced-group
```

- 3 Take *Active-Node* out of standby mode.
Replace *Active-Node* with the hostname assigned in [Identifying node roles](#) on page 117:

```
pcs cluster unstandby Active-Node
```

- 4 Confirm that disk replication is started.

```
drbd-overview
```

Do not proceed until the status of all devices is `UpToDate/UpToDate`.

- 5 Put *Passive-Node* into standby mode.
Replace *Passive-Node* with the hostname assigned in [Identifying node roles](#) on page 117:

```
pcs cluster standby Passive-Node
```

- 6 Confirm that node roles are reversed.

```
pcs status resources
```

Do not proceed until *Active-Node* is in the list named `Masters` and *Passive-Node* is in the list named `Slaves`.

- 7 Take *Passive-Node* out of standby mode.
Replace *Passive-Node* with the hostname assigned in [Identifying node roles](#) on page 117:

```
pcs cluster unstandby Passive-Node
```

- 8 Confirm that disk replication is started.

```
drbd-overview
```

Do not proceed until the status of all devices is `UpToDate/UpToDate`.

- 9 Start cluster resources.
 - a Start the `serviced` group.

```
pcs resource enable serviced-group
```

- b** Start the storage group.

```
pcs resource enable serviced-storage
```

- c** Start `serviced`.

```
pcs resource enable serviced
```

- 10** Start your application.

For example, the Zenoss application:

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

About adding space to a tenant volume

The LVM thin pool for application data can be used to store multiple tenant volumes and multiple snapshots of tenant data.

- If you added space to the thin pool to create a new tenant volume, see [serviced-storage](#) on page 128.
- If you added space to the thin pool to provide additional storage for snapshots, then the tenant volume does not need to be resized.
- If you added space to the thin pool because the tenant devices were oversubscribed, then the tenant volume does not need to be resized. A tenant volume is oversubscribed when its size at creation exceeded the amount of space available in the thin pool.

Adding space to a tenant volume

Use this procedure to increase the size of a tenant volume in a Control Center thin pool.

- 1** Log in to the master host as `root`, or as a user with superuser privileges.
- 2** Identify the tenant device to resize.

```
serviced volume status
```

- 3** Display the device mapper name of the `serviced` thin pool.

```
grep -E '^\\b*SERVICED_DM_THINPOOLDEV' /etc/default/serviced \
| sed -e 's/.*=//'
```

Typically, the name is `/dev/mapper/serviced-serviced--pool`.

- 4** Increase the size of the tenant device.
In the following command:

- Replace *Device-Mapper-Name* with the device mapper name of the thin pool.
- Replace *Tenant-ID* with the identifier of the tenant device.
- Replace *Total-Size* with the sum of the existing device size plus the space to add to the device, in gigabytes. Include the units identifier, `G`.

```
serviced-storage resize -d /opt/serviced/var/volumes \
-o dm.thinpooldev=Device-Mapper-Name Tenant-ID Total-SizeG
```

Upgrading only the Pacemaker resource agents

B

This appendix includes procedures for upgrading the Pacemaker resource agents for Control Center in a high-availability deployment.

Identifying the Pacemaker resource agents version

Use this procedure to identify the installed version of the Pacemaker resource agents 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 agents for Control Center.

```
yum info serviced-resource-agents | grep Version
```

Perform the next procedure to determine the available version of the resource agents.

Identifying the available version of the resource agents

To perform this procedure, you need:

- A workstation with internet access.
- Permission to download files from the delivery.zenoss.com site. You may request permission by filing a ticket at the [Zenoss Support](https://zenoss.com/support) site.

Use this procedure to identify the available version of the Pacemaker resource agents for Control Center, and to download its package file, if necessary.

- 1 In a web browser, navigate to delivery.zenoss.com.
- 2 Log in with the account provided by Zenoss Support.
- 3 Identify the version number of the resource agents package file.

The file name is `serviced-resource-agents-1.1.0-1.x86_64.rpm`.

- If the version number of the package file is greater than the installed version number, download the package file.
 - If the version number of the package file is equal to the installed version number, stop. The installed version is up-to-date.
- 4 Use a secure copy program to copy the package file to each master host node.

The recommended storage location for the package file is `/opt/zenoss-repo-mirror`.

Updating the Pacemaker resource agents

Use this procedure to update the Pacemaker resource agents for Control Center.

- 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 Stop the top-level service `serviced` is managing, if necessary.

- a Show the status of running services.

```
serviced service status
```

The top-level service is the service listed immediately below the headings line.

- If the status of the top-level service and all child services is `stopped`, proceed to the next step.
 - If the status of the top-level service and all child services is **not** `stopped`, perform the remaining substeps.
- b Stop the top-level service.
Replace *Service* with the name or identifier of the top-level service:

```
serviced service stop Service
```

- c Monitor the stop.

```
serviced service status
```

When the status of the top-level service and all child services is `stopped`, proceed to the next step.

- 3 Stop Control Center.

```
pcs resource disable serviced
```

- 4 In a separate window, log in to the other master node as `root`, or as a user with superuser privileges.
- 5 On both nodes, install the latest version of the resource agents package.
For example, if the package file is located in the `/opt/zenoss-repo-mirror` directory:

```
yum install /opt/zenoss-repo-mirror/serviced-resource-  
agents-1.1.0-1.x86_64.rpm
```

- 6 On either node, start Control Center.

```
pcs resource enable serviced
```

- 7 On either node, start the application.

```
serviced service start Service
```

C

Storage management utility

This appendix includes a description of the `serviced-storage` command, the required utility for creating the Docker thin pool and creating and managing the Control Center application data thin pool.

serviced-storage

The `serviced-storage` command manages Control Center storage.

Use this command to create LVM thin pools for Docker and Control Center.

USAGE

```
serviced-storage [-h|--help] [-o DeviceMapperOption=Value] \
  [-v] Command [CommandOptions]
```

GLOBAL OPTIONS

--help, -h

Shows the help information.

-o DeviceMapperOption=Value

A device mapper option. Applies only to device mapper drivers.

-v

Displays verbose logging.

COMMANDS

check

Check for orphaned devices.

create

Create a volume on a driver.

create-thin-pool

Create an LVM thin pool.

disable

Disable a driver.

init

Initialize a driver.

list

Print volumes on a driver.

mount

Mount an existing volume from a driver.

remove

Remove an existing volume from a driver.

resize

Resize an existing volume.

set

Set the default driver.

status

Print the driver status

sync

Sync data from a volume to another volume.

unset

Unset the default driver.

version

Print the version and exit.

serviced-storage check

The `serviced-storage check` command searches for orphaned snapshot devices in the `serviced` application data thin pool and removes them, if requested. This command requires the path of `serviced` tenant volumes, which is determined by the `SERVICED_VOLUMES_PATH` variable in `/etc/default/serviced`. The default path is `/opt/serviced/var/volumes`.

Syntax:

```
serviced-storage [GlobalOptions] check [-c|--clean] Path
```

Command options:

[-c|--clean]

Remove orphaned snapshot devices.

serviced-storage create-thin-pool

The `serviced-storage create-thin-pool` command creates an LVM thin pool either for Docker data or for Control Center application data. When devices are specified, the command creates an LVM volume group.

Syntax:

```
serviced-storage [GlobalOptions] create-thin-pool \
  [-s|--size]=[Value] [G|%] [docker|serviced] \
  [DevicePath [DevicePath...] | VolumeGroupName]
```

Command options:

[-s|--size]=[Value][G|%]

The size of the thin pool to create. The size can be a fixed value (in gigabytes) or a relative value (a percentage) of the available storage. When this option is not used, the thin pool size defaults to 90% of the specified storage resource.

serviced-storage resize

The `serviced-storage resize` command increases the size of a serviced tenant device in its LVM thin pool. Like LVM thin pools, the size of a serviced tenant device can never decrease.

Syntax:

```
serviced-storage [GlobalOptions] resize \
  [-d|--driver]=Value TenantID NewSize
```

Command options:

[-d|--driver]=Value

The path of the tenant volume.

EXAMPLES

Create an LVM volume group named `zenoss` and use it for both thin pools:

```
vgcreate zenoss /dev/sdb /dev/sdc
serviced-storage create-thin-pool --size=50G docker zenoss
serviced-storage create-thin-pool --size=50% serviced zenoss
```

If you specify devices or partitions, `serviced-storage` creates an LVM volume group with the same name as the thin pool. The following example yields the same result as the previous, except the name of the volume group is `docker` instead of `zenoss`:

```
serviced-storage create-thin-pool docker /dev/sdb /dev/sdc
serviced-storage create-thin-pool serviced docker
```

Create thin pools on separate block devices:

```
serviced-storage create-thin-pool docker /dev/sdb
serviced-storage create-thin-pool serviced /dev/sdc
```

Create thin pools on separate partitions:

```
serviced-storage create-thin-pool docker /dev/sdb1
serviced-storage create-thin-pool serviced /dev/sdc3
```

Increase the size of the `serviced` LVM thin pool, and then increase the size of a serviced tenant device.

```
lvextend -L+300G zenoss/serviced-pool
serviced-storage -o dm.thinpooldev=/dev/mapper/zenoss-serviced--pool \
  resize -d /opt/serviced/var/volumes 58uuetj38draeu9alp6002b1y 200G
```

Identify the `serviced` application data thin pool, and then remove orphaned snapshot devices.

```
ls /dev/mapper | grep serviced  
serviced-storage -o dm.thinpooldev=/dev/mapper/zenoss-serviced--pool \  
check -c /opt/serviced/var/volumes
```

Starting and stopping Control Center deployments

D

This appendix includes procedures for stopping and starting a high-availability Control Center deployment.

Note The procedures in this appendix assume that Control Center is the only source of Docker containers that are run on a host.

Stopping Control Center

To stop Control Center in a high-availability deployment, perform the procedures in this section, in order.

Stopping a master host node

Use this procedure to stop the Control Center service (`serviced`) on the master host in a high-availability deployment.

- 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 step.

- 3 Stop the top-level service `serviced` is managing, if necessary.
 - a Show the status of running services.

```
serviced service status
```

The top-level service is the service listed immediately below the headings line.

- If the status of the top-level service and all child services is `stopped`, proceed to the next step.
 - If the status of the top-level service and all child services is **not** `stopped`, perform the remaining substeps.
- b Stop the top-level service.

Replace *Service* with the name or identifier of the top-level service:

```
serviced service stop Service
```

- c Monitor the stop.

```
serviced service status
```

When the status of the top-level service and all child services is stopped, proceed to the next step.

- 4 Stop Control Center with the cluster management tool.

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

- 5 Monitor the status of cluster resources.

```
watch pcs status
```

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

- 6 Ensure that no containers remain in the local repository.

- a Start the Docker service.

```
systemctl start docker
```

- b Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

If the command returns a result, enter the following command:

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- c Stop the Docker service.

```
systemctl stop docker
```

- 7 To ensure that no containers remain in both Docker repositories, log in to the other master node as `root`, or as a user with superuser privileges, and then perform the preceding step.

Stopping a delegate host

Use this procedure to stop the Control Center service (`serviced`) on a delegate host in a multi-host deployment. Repeat this procedure on each delegate host in your deployment.

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

```
systemctl stop serviced
```

- 3 Ensure that no containers remain in the local repository.

- a Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, proceed to the next step.
- If the command returns a result, perform the following substeps.

- b Remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
 - If the remove command does not complete, the most likely cause is an NFS conflict. Perform the following substeps.
- c** Stop the NFS and Docker services.

```
systemctl stop nfs && systemctl stop docker
```

- d** Start the NFS and Docker services.

```
systemctl start nfs && systemctl start docker
```

- e** Repeat the attempt to remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
 - If the remove command does not complete, perform the remaining substeps.
- f** Disable the automatic startup of `serviced`.

```
systemctl disable serviced
```

- g** Reboot the host.

```
reboot
```

- h** Log in to the delegate host as `root`, or as a user with superuser privileges.

- i** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

4 Dismount all filesystems mounted from the Control Center master host.

This step ensures no stale mounts remain when the storage on the master host is replaced.

- a** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
 - If the preceding command returns a result, perform the following substeps.
- b** Force the filesystems to dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
    umount -f $FS
done
```

- c** Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.

- If the preceding command returns a result, perform the following substeps.
- d** Perform a lazy dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs')
do
    umount -f -l $FS
done
```

- e** Restart the NFS service.

```
systemctl restart nfs
```

- f** Determine whether any filesystems remain mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts \
| grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
 - If the preceding command returns a result, perform the remaining substeps.
- g** Disable the automatic startup of `serviced`.

```
systemctl disable serviced
```

- h** Reboot the host.

```
reboot
```

- i** Log in to the delegate host as `root`, or as a user with superuser privileges.
- j** Enable the automatic startup of `serviced`.

```
systemctl enable serviced
```

Starting Control Center

Use this procedure to start Control Center in a high-availability deployment. The default configuration of the Control Center service (`serviced`) is to start when the host starts. This procedure is only needed after stopping `serviced` to perform maintenance tasks.

- 1** Log in to the primary node as `root`, or as a user with superuser privileges.
In this context, the primary node is the node that was the current node when you stopped Control Center.
- 2** Identify the hosts in the ZooKeeper ensemble.

```
grep -E '^\\b*SERVICED_ZK=' /etc/default/serviced
```

The result is a list of 1, 3, or 5 hosts, separated by the comma character (`,`). The master host is always a node in the ZooKeeper ensemble.

- 3** In separate windows, log in to each of the delegate hosts that are nodes in the ZooKeeper ensemble as `root`, or as a user with superuser privileges.
- 4** Take the cluster out of standby mode.

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

- 5 On the other ZooKeeper ensemble hosts, start `serviced`.

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

```
systemctl start serviced
```

- 6 Monitor the status of cluster resources.

```
watch pcs status
```

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

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

- a Attach to the container of the ZooKeeper service.

```
docker exec -it serviced-isvcs_zookeeper bash
```

- b Query the master host and identify its role in the ensemble.

Replace *Master* with the hostname or IP address of the master host:

```
{ echo stats; sleep 1; } | nc Master 2181 | grep Mode
```

The result includes `leader` or `follower`. When multiple hosts rely on the ZooKeeper instance on the master host, the result includes `standalone`.

- c Query the other delegate hosts to identify their role in the ensemble.

Replace *Delegate* with the hostname or IP address of a delegate host:

```
{ echo stats; sleep 1; } | nc Delegate 2181 | grep Mode
```

- d Detach from the container of the ZooKeeper service.

```
exit
```

If none of the nodes reports that it is the ensemble leader within a few minutes of starting `serviced`, reboot the ensemble hosts.

- 8 Log in to each of the delegate hosts that are not nodes in the ZooKeeper ensemble as `root`, or as a user with superuser privileges, and then start `serviced`.

```
systemctl start serviced
```

- 9 Optional: Monitor the startup, if desired.

```
journalctl -u serviced -f -o cat
```

Once Control Center is started, it is ready to start managing applications. For more information, refer to the documentation of your application.



Resolving package dependency conflicts

This appendix includes procedures for resolving common Docker CE and Control Center dependency conflicts.

Resolving device mapper dependency conflicts

To perform this procedure, you need:

- An RHEL/CentOS system with internet access and the same operating system release and kernel as the Control Center hosts in your deployment.
- A secure network copy program.

Use this procedure to resolve dependency issues in which the installed versions of device mapper libraries are newer than the versions included in the Zenoss mirror. The following example shows a typical yum error of this type:

```
Error: Package: 7:device-mapper-event-1.02.107-5.el7.x86_64 (zenoss-mirror)
Requires: device-mapper = 7:1.02.107-5.el7
Installed: 7:device-mapper-1.02.107-5.el7_2.5.x86_64 (@updates)
device-mapper = 7:1.02.107-5.el7_2.5
```

Follow these steps:

- 1 Display the version number of the installed device mapper package.

```
rpm -q device-mapper | cut -d - -f 3-
```

Example result:

```
1.02.135-1.el7_3.1.x86_64
```

Record the version number for subsequent use.

- 2 Log in to a compatible host that is connected to the internet as `root`, or as a user with superuser privileges. The host must have the same operating system (RHEL or CentOS) and release installed as the Control Center hosts in your deployment.
- 3 Install yum utilities, if necessary.

- a Determine whether the `yum` utilities package is installed.

```
rpm -qa | grep yum-utils
```

- If the command returns a result, the package is installed. Proceed to the next step.
- If the command does not return a result, the package is not installed. Perform the following substep.

- b Install the `yum-utils` package.

```
yum install yum-utils
```

- 4 Download the required dependencies, and then create a `tar` archive of the files.

- a Create a variable for the dependency version to download.

Replace *Device-Mapper-Version* with the version number displayed in a previous step:

```
myVersion=Device-Mapper-Version
```

- b Create a temporary directory for the dependencies.

```
mkdir /tmp/downloads
```

- c Download the dependencies to the temporary directory.

```
yum install --downloadonly --downloadaddir=/tmp/downloads \  
device-mapper-event-$myVersion
```

The `yum` command downloads two package files.

- d Create a `tar` archive of the temporary directory.

```
cd /tmp && tar czf ./downloads.tgz ./downloads
```

- 5 Use a secure copy program to copy the archive file to the `/tmp` directory of the Control Center host or hosts that need the dependencies.

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

- 7 Install the device mapper dependencies.

- a Extract the packages from the `tar` archive.

```
cd /tmp && tar xzf ./downloads.tgz
```

- b Install the dependencies.

```
yum install $(ls /tmp/downloads/*.rpm)
```

Return to the procedure you were performing before turning to this appendix and retry the `yum install` command that failed previously.

Resolving other dependency conflicts

Use this procedure to resolve dependency issues in which the installed versions of one or more dependencies are newer than the versions included in the Zenoss mirror. The following example shows a typical `yum` error of this type:

```
Error: Package: policycoreutils-python-2.5-9.el7.x86_64 (zenoss-mirror)  
Requires: policycoreutils = 2.5-9.el7
```

```
Installed: policycoreutils-2.5-11.el7_3.x86_64 (@updates)
```

Follow these steps:

1 Install the older package.

Replace *Package-Name* with the name of the package displayed in the error message:

```
rpm -Uvh --oldpackage Package-Name
```

2 Clean all yum caches.

```
yum clean all
```

Return to the procedure you were performing before turning to this appendix and retry the `yum install` command that failed previously.

F

User access control

Control Center provides a browser interface and a command-line interface.

To gain access to the Control Center browser interface, users must have login accounts on the Control Center master host. In addition, users must be members of the Control Center browser interface access group, which by default is the system group, `wheel`. To enhance security, you may change the browser interface access group from `wheel` to any other group.

To use the Control Center command-line interface (CLI) on a Control Center host, a user must have login account on the host, and the account must be a member of the `serviced` group. The `serviced` group is created when the Control Center RPM package is installed.

Note You can use two different groups to control access to the browser interface and the CLI. You can enable access to both interfaces for the same users by choosing the `serviced` group as the browser interface access group.

Pluggable Authentication Modules (PAM) has been tested and is recommended for enabling access to both the browser interface and the command-line interface. However, the PAM configuration must include the `sudo` service. Control Center relies on the host's `sudo` configuration, and if no configuration is present, PAM defaults to the configuration for `other`, which is typically too restrictive for Control Center users. For more information about configuring PAM, refer to your operating system documentation.

Adding users to the default browser interface access group

Use this procedure to add users to the default browser interface access group of Control Center, `wheel`.

Note Perform this procedure or the next procedure, but not both.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Add a user to the `wheel` group.

Replace *User* with the name of a login account on the master host.

```
usermod -aG wheel User
```

Repeat the preceding command for each user to add.

Configuring a regular group as the Control Center browser interface access group

Use this procedure to change the default browser interface access group of Control Center from `wheel` to a non-system group.

The following Control Center variables are used in this procedure:

SERVICED_ADMIN_GROUP

Default: `wheel`

The name of the Linux group on the `serviced` master host whose members are authorized to use the `serviced` browser interface. You may replace the default group with a group that does not have superuser privileges.

SERVICED_ALLOW_ROOT_LOGIN

Default: `1 (true)`

Determines whether the `root` user account on the `serviced` master host may be used to gain access to the `serviced` browser interface.

Note Perform this procedure or the previous procedure, but not both.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Create a variable for the group to designate as the administrative group.
In this example, the group is `ccuser`. You may choose a different group, or choose the `serviced` group. (Choosing the `serviced` group allows all browser interface users to use the CLI.)

```
myGROUP=ccuser
```

- 3 Create a new group, if necessary.

```
groupadd $myGROUP
```

- 4 Add one or more existing users to the group.

Replace *User* with the name of a login account on the host:

```
usermod -aG $myGROUP User
```

Repeat the preceding command for each user to add.

- 5 Specify the new administrative group in the `serviced` configuration file.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the line for the `SERVICED_ADMIN_GROUP` variable, and then make a copy of the line, immediately below the original.
 - c Remove the number sign character (`#`) from the beginning of the line.
 - d Change the value from `wheel` to the name of the group you chose earlier.
 - e Save the file, and then close the editor.
- 6 Optional: Prevent the `root` user from gaining access to the Control Center browser interface, if desired.
 - a Open `/etc/default/serviced` in a text editor.
 - b Locate the line for the `SERVICED_ALLOW_ROOT_LOGIN` variable, and then make a copy of the line, immediately below the original.
 - c Remove the number sign character (`#`) from the beginning of the line.
 - d Change the value from `1` to `0`.

- e Save the file, and then close the editor.

Enabling use of the command-line interface

Use this procedure to enable a user to perform administrative tasks with the Control Center command-line interface.

- 1 Log in to the host as `root`, or as a user with superuser privileges.
- 2 Add a user to the `serviced` group.

Replace *User* with the name of a login account on the host.

```
usermod -aG serviced User
```

Repeat the preceding command for each user to add.

Updating the cluster management software



The procedures in this appendix describe how to update the Distributed Replicated Block Device (DRBD) and Pacemaker/Corosync software, if desired. Typically, the software is updated only to resolve issues caused by the currently-installed versions of the packages.

Note Zenoss supports DRBD 8.4, not DRBD 9.0. The procedures in this section download the latest version of release 8.4.

Downloading and staging cluster software

To perform this procedure, you need:

- An RHEL/CentOS system with internet access and the same operating system release and kernel as the master host nodes.
- A secure network copy program.

Use this procedure to download packages for Distributed Replicated Block Device (DRBD) and Pacemaker/Corosync, and to bundle them for installation on master host nodes.

- 1 Log in to a compatible host that is connected to the internet as `root`, or as a user with superuser privileges. The host must have the same operating system (RHEL or CentOS) and release installed, and the same version of the Linux kernel, as the master host nodes.
- 2 Install `yum` utilities, if necessary.
 - a Determine whether the `yum` utilities package is installed.

```
rpm -qa | grep yum-utils
```

- If the command returns a result, the package is installed. Proceed to the next step.
 - If the command does not return a result, the package is not installed. Perform the following substep.
- b Install the `yum-utils` package.

```
yum install yum-utils
```

- 3 Add the Enterprise Linux packages repository (ELRepo), if necessary.
 - a Determine whether the ELRepo repository is available.

```
yum repolist | grep elrepo
```

- If the command returns a result, the repository is available. Proceed to the next step.
 - If the command does not return a result, the repository is not available. Perform the following substeps.
- b** Import the public key for the repository.

```
rpm --import https://www.elrepo.org/RPM-GPG-KEY-elrepo.org
```

- c** Add the repository to the download host.

```
rpm -Uvh \
  http://www.elrepo.org/elrepo-release-7.0-2.el7.elrepo.noarch.rpm
```

- d** Clean and update the yum caches.

```
yum clean all && yum makecache fast
```

- 4** Download the required packages and their dependencies, and then create a tar archive of the package files.
 - a** Create a temporary directory for the packages.

```
mkdir /tmp/downloads
```

- b** Download the DRBD packages to the temporary directory.

```
repotrack -a x86_64 -r elrepo -p /tmp/downloads kmod-drbd84
```

- c** Download the Corosync/Pacemaker packages to the temporary directory.

```
repotrack -a x86_64 -p /tmp/downloads pcs
```

- d** Create a tar archive of the temporary directory.

```
cd /tmp && tar czf ./downloads.tgz ./downloads
```

- 5** Use a secure copy program to copy the packages archive to the /tmp directory of each master host node.

Updating cluster software

Use this procedure to update Distributed Replicated Block Device (DRBD) and Pacemaker/Corosync, if desired.

- 1** Stop all applications, and then stop the high-availability cluster.
For more information, see [Stopping a Control Center deployment](#) on page 21.
- 2** Log in to the primary node as `root`, or as a user with superuser privileges.
- 3** In a separate window, log in to the secondary node as `root`, or as a user with superuser privileges.
- 4** On both nodes, stop the cluster software.
 - a** Stop the PCS daemon.

```
systemctl stop pcsd.service
```

- b** Stop DRBD.

```
drbdadm down all
```

- 5** On both nodes, extract and install the cluster management packages.

- a** Extract the packages.

```
cd /tmp && tar xzf ./downloads.tgz
```

- b** Install the packages.

```
yum install ./downloads/*.rpm
```

- 6** On both nodes, install the Pacemaker resource agents for Control Center.

```
yum install \
/opt/zenoss-repo-mirror/serviced-resource-agents-1.1.0-1.x86_64.rpm
```

- 7** Determine which node is the primary node.

```
pcs status
```

- 8** On the primary node, start the cluster software.

- a** Start DRBD.

```
drbdadm up all
```

- b** Start DRBD and assign the primary role.

```
drbdadm up all && drbdadm primary --force serviced-dfs
```

- c** Start the PCS daemon.

```
systemctl start pcsd.service
```

- 9** On the secondary node, start the cluster software.

- a** Start DRBD.

```
drbdadm up all
```

- b** Start the PCS daemon.

```
systemctl start pcsd.service
```

- 10** Start the high-availability cluster, and then start applications.

For more information, see [Starting Control Center](#) on page 135.

H

Control Center configuration variables

This appendix includes the content in the following list:

- 1 [Removed variables \(at version 1.2.0\)](#) on page 146
- 2 [New variables \(at version 1.2.0\)](#) on page 147
- 3 [Master host configuration variables](#) on page 148
- 4 [Delegate host configuration variables](#) on page 151
- 5 [Universal configuration variables](#) on page 153
- 6 [Best practices for configuration files](#) on page 155
- 7 [Control Center configuration file](#) on page 155

Many configuration choices depend on application requirements. Please review your application documentation before configuring hosts.

Removed variables (at version 1.2.0)

The variables in the following table were present in version 1.1.1 and are not present in version 1.2.0. The variables are shown in the order in which they appeared in `/etc/default/serviced`.

Variable	Explanation
<code>TMP</code>	Not POSIX compliant. Replaced by <code>TMPDIR</code> .
<code>SERVICED_AGENT</code>	<code>SERVICED_MASTER</code> is sufficient to specify the host role.
<code>SERVICED_VARPATH</code>	<code>SERVICED_ISVCS_PATH</code> , <code>SERVICED_VOLUMES_PATH</code> , and <code>SERVICED_BACKUPS_PATH</code> are more specific and preferred.
<code>SERVICED_MONITOR_DFS_MASTER_INTERVAL</code> <code>SERVICED_MONITOR_DFS_MASTER_RESTART</code> <code>SERVICED_MONITOR_DFS_REMOTE_UPDATE_INTERVAL</code>	Control Center now uses different mechanisms to monitor DFS status.

New variables (at version 1.2.0)

The variables in the following table were not present in version 1.1.1 and are present in version 1.2.0. The variables are shown in the order in which they appear in `/etc/default/serviced`.

Variable	Purpose
<code>TMPDIR</code>	Replacement for <code>TMP</code> .
<code>SERVICED_MASTER_IP</code>	The IP address of the master host.
<code>SERVICED_MASTER_POOLID</code>	The name of the resource pool to which the master host belongs.
<code>SERVICED_RPC_TLS_MIN_VERSION</code> <code>SERVICED_RPC_TLS_CIPHERS</code>	Additional options for encrypting RPC communications.
<code>SERVICED_UI_POLL_FREQUENCY</code>	The frequency at which browser interface clients poll the server.
<code>SERVICED_MUX_DISABLE_TLS</code> <code>SERVICED_MUX_TLS_MIN_VER</code> <code>SERVICED_MUX_TLS_CIPHERS</code>	Additional options for encrypting RPC communications.
<code>SERVICED_DM_ARGS</code> <code>SERVICED_DM_BASESIZE</code> <code>SERVICED_DM_LOOPDATASIZE</code> <code>SERVICED_DM_LOOPMETADATASIZE</code> <code>SERVICED_DM_THINPOOLDEV</code>	Options for the <code>devicemapper</code> storage driver.
<code>SERVICED_STORAGE_STATS_UPDATE_INTERVAL</code>	The number of seconds between polls of kernel statistics about the application data thin pool.
<code>SERVICED_LOGSTASH_CYCLE_TIME</code>	The amount of time between logstash purges.
<code>SERVICED_SVCSTATS_CACHE_TIMEOUT</code>	The lifetime of data in the browser interface cache for statistics about services.
<code>SERVICED_NFS_CLIENT</code>	DEPRECATED: Prevent a delegate host from mounting the DFS. In version 1.2.0, delegate host access to the DFS is controlled through resource pool permissions.
<code>SERVICED_DOCKER_DNS</code>	The Docker DNS configuration for containers.
<code>SERVICED_SNAPSHOT_USE_PERCENT</code>	The amount of free space in the application data thin pool, used to determine whether the pool can store a new snapshot.
<code>SERVICED_ZK_SESSION_TIMEOUT</code>	The number of seconds the ZooKeeper leader waits before flushing an inactive connection.
<code>SERVICED_ES_STARTUP_TIMEOUT</code>	The number of seconds to wait for the Elasticsearch service to start.
<code>SERVICED_RPC_DIAL_TIMEOUT</code>	The number of seconds until an RPC connection attempt times out.
<code>SERVICED_AUTH_TOKEN_EXPIRATION</code>	The lifetime of a delegate host revocation certificate.

Variable	Purpose
<i>SERVICED_CONTROLLER_BINARY</i>	The path of the <code>serviced-controller</code> binary.
<i>SERVICED_HOME</i>	The path of the home directory for <code>serviced</code> .
<i>SERVICED_ETC_PATH</i>	The path of the directory for <code>serviced</code> configuration files.

Master host configuration variables

The tables in this section provide an overview of the `serviced` configuration variables that apply to the Control Center master host. Set these variables as required for your environment or applications.

Storage variables

The variables in the following table are set only on the master host.

- Use one of the first two groups of variables but not both.
- Before starting the master host for the first time, you might need to change the defaults of the third group.
- Typically, the defaults of the last two groups of variables are not changed until Control Center has managed an application for a while and a need arises.

The *SERVICED_STORAGE_STATS_UPDATE_INTERVAL* variable sets the interval for collecting kernel statistics about the application data thin pool. Its default value is unlikely to require a change until a need arises.

Variable	Description	Purpose
<i>SERVICED_FS_TYPE</i> <i>SERVICED_DM_ARGS</i> <i>SERVICED_DM_BASESIZE</i> <i>SERVICED_DM_THINPOOLDEV</i>	The specifications of a <code>devicemapper</code> -based application data storage resource for production use.	Provide basic information about the data storage resource.
<i>SERVICED_FS_TYPE</i> <i>SERVICED_DM_LOOPDATASIZE</i> <i>SERVICED_DM_LOOPMETADATASIZE</i> <i>SERVICED_ALLOW_LOOP_BACK</i>	The specifications of a <code>devicemapper</code> -based application data storage resource for development use.	Provide basic information about the data storage resource.
<i>SERVICED_ISVCS_PATH</i> <i>SERVICED_VOLUMES_PATH</i> <i>SERVICED_BACKUPS_PATH</i>	The data storage paths of separate functional components of Control Center internal services.	Enable separate storage areas for one or more components. The default installation process puts all three components on the same device.
<i>SERVICED_SNAPSHOT_TTL</i> <i>SERVICED_SNAPSHOT_USE_PERCENT</i> <i>SERVICED_MAX_DFS_TIMEOUT</i>	The snapshot retention interval, the percentage of the data storage thin pool that is unused, and the snapshot attempt timeout interval.	Prevent the creation of snapshots that are too large to fit the thin pool.
<i>SERVICED_LOGSTASH_MAX_DAYS</i> <i>SERVICED_LOGSTASH_MAX_SIZE</i> <i>SERVICED_LOGSTASH_CYCLE_TIME</i>	The variables that manage the amount of space used by the application log storage service.	Prevent application logs from filling the storage device that <code>logstash</code> uses.

Internal services endpoint variables

The variables in the following table must be set identically on all Control Center delegate hosts.

The `SERVICED_AUTH_TOKEN_EXPIRATION` variable affects RPC, mux, and internal services endpoint traffic.

Variable	Endpoint	Description
<code>SERVICED_DOCKER_REGISTRY</code>	(varies)	The local Docker registry for Control Center internal services images and application images.
<code>SERVICED_ENDPOINT</code>	<i>Master-Host</i> : 4979	The <code>serviced</code> RPC server. The endpoint port number must match the value of <code>SERVICED_RPC_PORT</code> .
<code>SERVICED_LOG_ADDRESS</code>	<i>Master-Host</i> : 5042	The <code>logstash</code> service.
<code>SERVICED_LOGSTASH_ES</code>	<i>Master-Host</i> : 9100	The Elasticsearch service for logstash.
<code>SERVICED_STATS_PORT</code>	<i>Master-Host</i> : 8443	The <code>serviced</code> metrics consumer service.
<code>SERVICED_AUTH_TOKEN_EXPIRATION</code>	(none)	The length of time a delegate authentication token is valid.

RPC service variables

The variables in the following table must be set identically on all Control Center hosts, except:

- `SERVICED_RPC_PORT`, set only on the master
- `SERVICED_MAX_RPC_CLIENTS`, set only on delegates

By default, `serviced` uses TLS to encrypt all RPC traffic. The `SERVICED_KEY_FILE` and `SERVICED_CERT_FILE` variables identify the digital certificate used for RPC, mux, and HTTP traffic.

The `SERVICED_AUTH_TOKEN_EXPIRATION` variable affects RPC, mux, and internal services endpoint traffic.

Variable	Where to set
<code>SERVICED_ENDPOINT</code>	Master, delegates
<code>SERVICED_MAX_RPC_CLIENTS</code>	Delegates
<code>SERVICED_RPC_PORT</code>	Master
<code>SERVICED_RPC_CERT_VERIFY</code>	Master, delegates
<code>SERVICED_RPC_DISABLE_TLS</code>	Master, delegates
<code>SERVICED_RPC_TLS_MIN_VERSION</code>	Master, delegates
<code>SERVICED_RPC_TLS_CIPHERS</code>	Master, delegates
<code>SERVICED_KEY_FILE</code>	Master
<code>SERVICED_CERT_FILE</code>	Master
<code>SERVICED_RPC_DIAL_TIMEOUT</code>	Master, delegates

Variable	Where to set
<i>SERVICED_AUTH_TOKEN_EXPIRATION</i>	Master

Multiplexer variables

The variables in the following table must be set identically on all Control Center hosts.

By default, `serviced` uses TLS to encrypt all mux traffic. The *SERVICED_KEY_FILE* and *SERVICED_CERT_FILE* variables identify the digital certificate used for RPC, mux, and HTTP traffic.

The *SERVICED_AUTH_TOKEN_EXPIRATION* variable affects RPC, mux, and internal services endpoint traffic.

Variable	Where to set
<i>SERVICED_MUX_PORT</i>	Master, delegates
<i>SERVICED_MUX_DISABLE_TLS</i>	Master, delegates
<i>SERVICED_MUX_TLS_MIN_VERSION</i>	Master, delegates
<i>SERVICED_MUX_TLS_CIPHERS</i>	Master, delegates
<i>SERVICED_KEY_FILE</i>	Master
<i>SERVICED_CERT_FILE</i>	Master
<i>SERVICED_AUTH_TOKEN_EXPIRATION</i>	Master

HTTP server variables

The variables in the following table are set only on the master host, except the *SERVICED_UI_PORT* variable, which must be set identically on all Control Center hosts.

By default, `serviced` uses TLS to encrypt all HTTP traffic. The *SERVICED_KEY_FILE* and *SERVICED_CERT_FILE* variables identify the digital certificate used for RPC, mux, and HTTP traffic.

Variable	Description
<i>SERVICED_UI_PORT</i>	The port on which the HTTP server listens for requests.
<i>SERVICED_TLS_MIN_VERSION</i>	The minimum version of TLS that <code>serviced</code> accepts for HTTP traffic.
<i>SERVICED_TLS_CIPHERS</i>	The list TLS ciphers that <code>serviced</code> accepts for HTTP traffic.
<i>SERVICED_KEY_FILE</i>	The path of a digital certificate key file.
<i>SERVICED_CERT_FILE</i>	The path of a digital certificate file.

Browser interface variables (master host only)

The variables in the following table are set only on the master host.

Variable	Description
<i>SERVICED_UI_POLL_FREQUENCY</i>	The number of seconds between polls from browser interface clients.
<i>SERVICED_SVCSTATS_CACHE_TIMEOUT</i>	The number of seconds to cache statistics about services.
<i>SERVICED_ADMIN_GROUP</i>	The group on the master host whose members can use the browser interface.
<i>SERVICED_ALLOW_ROOT_LOGIN</i>	Determines whether <code>root</code> on the master host can use the browser interface.

Tuning variables (master host only)

Variable	Description
<i>SERVICED_ES_STARTUP_TIMEOUT</i>	The number of seconds to wait for the Elasticsearch service to start.
<i>SERVICED_MASTER_POOLID</i>	The name of the default resource pool. This variable is only used the first time <code>serviced</code> is started.

Delegate host configuration variables

The tables in this section provide an overview of the `serviced` configuration variables that apply to Control Center delegate hosts. Set these variables as required for your environment or applications.

Delegate variables

The following miscellaneous variables apply only to delegate hosts.

Variable	Description
<i>SERVICED_ZK</i>	The list of hosts in the ZooKeeper ensemble.
<i>SERVICED_STATS_PERIOD</i>	The frequency at which delegates gather metrics to send to the master host.
<i>SERVICED_IPTABLES_MAX_CONNECTIONS</i>	The maximum number of open connections to allow on a delegate. The number increases when the master is unavailable and decreases when the master returns to service.

Internal services endpoint variables

The variables in the following table must be set identically on all Control Center delegate hosts.

The *SERVICED_AUTH_TOKEN_EXPIRATION* variable affects RPC, mux, and internal services endpoint traffic.

Variable	Endpoint	Description
<i>SERVICED_DOCKER_REGISTRY</i>	(varies)	The local Docker registry for Control Center internal services images and application images.
<i>SERVICED_ENDPOINT</i>	<i>Master-Host</i> : 4979	The serviced RPC server. The endpoint port number must match the value of <i>SERVICED_RPC_PORT</i> .
<i>SERVICED_LOG_ADDRESS</i>	<i>Master-Host</i> : 5042	The <i>logstash</i> service.
<i>SERVICED_LOGSTASH_ES</i>	<i>Master-Host</i> : 9100	The Elasticsearch service for logstash.
<i>SERVICED_STATS_PORT</i>	<i>Master-Host</i> : 8443	The serviced metrics consumer service.
<i>SERVICED_AUTH_TOKEN_EXPIRATION</i>	(none)	The length of time a delegate authentication token is valid.

RPC service variables

The variables in the following table must be set identically on all Control Center hosts, except:

- *SERVICED_RPC_PORT*, set only on the master
- *SERVICED_MAX_RPC_CLIENTS*, set only on delegates

By default, *serviced* uses TLS to encrypt all RPC traffic. The *SERVICED_KEY_FILE* and *SERVICED_CERT_FILE* variables identify the digital certificate used for RPC, mux, and HTTP traffic.

The *SERVICED_AUTH_TOKEN_EXPIRATION* variable affects RPC, mux, and internal services endpoint traffic.

Variable	Where to set
<i>SERVICED_ENDPOINT</i>	Master, delegates
<i>SERVICED_MAX_RPC_CLIENTS</i>	Delegates
<i>SERVICED_RPC_PORT</i>	Master
<i>SERVICED_RPC_CERT_VERIFY</i>	Master, delegates
<i>SERVICED_RPC_DISABLE_TLS</i>	Master, delegates
<i>SERVICED_RPC_TLS_MIN_VERSION</i>	Master, delegates
<i>SERVICED_RPC_TLS_CIPHERS</i>	Master, delegates
<i>SERVICED_KEY_FILE</i>	Master
<i>SERVICED_CERT_FILE</i>	Master
<i>SERVICED_RPC_DIAL_TIMEOUT</i>	Master, delegates
<i>SERVICED_AUTH_TOKEN_EXPIRATION</i>	Master

Multiplexer variables

The variables in the following table must be set identically on all Control Center hosts.

By default, `serviced` uses TLS to encrypt all mux traffic. The `SERVICED_KEY_FILE` and `SERVICED_CERT_FILE` variables identify the digital certificate used for RPC, mux, and HTTP traffic.

The `SERVICED_AUTH_TOKEN_EXPIRATION` variable affects RPC, mux, and internal services endpoint traffic.

Variable	Where to set
<code>SERVICED_MUX_PORT</code>	Master, delegates
<code>SERVICED_MUX_DISABLE_TLS</code>	Master, delegates
<code>SERVICED_MUX_TLS_MIN_VERSION</code>	Master, delegates
<code>SERVICED_MUX_TLS_CIPHERS</code>	Master, delegates
<code>SERVICED_KEY_FILE</code>	Master
<code>SERVICED_CERT_FILE</code>	Master
<code>SERVICED_AUTH_TOKEN_EXPIRATION</code>	Master

Universal configuration variables

The tables in this section provide an overview of the `serviced` configuration variables that apply to all Control Center hosts. Set these variables as required for your environment or applications.

Role variable

Variable	Description
<code>SERVICED_MASTER</code>	Assigns the role of a <code>serviced</code> instance, either master or delegate. The master runs the application services scheduler and other internal services. Delegates run the application services assigned to the resource pool to which they belong.

Browser interface variable (all hosts)

Variable	Description
<code>SERVICED_UI_PORT</code>	The port on which the HTTP server listens for requests.

Networking variables

Variable	Description
<code>SERVICED_STATIC_IPS</code>	A list of one or more static IP addresses for IP assignment.
<code>SERVICED_OUTBOUND_IP</code>	The IP address of the network interface for <code>serviced</code> to use. When this variable is not set, <code>serviced</code> uses the IP address of the default network interface and assumes it has internet access. To prevent <code>serviced</code> from assuming it has internet access, set this variable.

Variable	Description
<i>SERVICED_VIRTUAL_ADDRESS_SUBNET</i>	The private network for containers that use virtual IP addresses. The default is 10.3.0.0/16, and the network can be unique on each host. A /29 network is sufficient.
<i>SERVICED_DOCKER_DNS</i>	A list of one or more DNS servers. The list is injected into all Docker containers.

Debugging variables

Variable	Description
<i>SERVICED_LOG_LEVEL</i>	The log level <code>serviced</code> uses when writing to the system log.
<i>SERVICED_DEBUG_PORT</i>	The port on which <code>serviced</code> listens for HTTP requests for the Go profiler .
<i>SERVICED_DOCKER_LOG_DRIVER</i>	The log driver for all Docker container logs.
<i>SERVICED_DOCKER_LOG_CONFIG</i>	Docker <code>--log-opt</code> options.

Tuning variables (all Control Center hosts)

Variable	Description
<i>GOMAXPROCS</i>	The maximum number of CPU cores that <code>serviced</code> uses.
<i>SERVICED_MAX_CONTAINER_AGE</i>	The number of seconds <code>serviced</code> waits before removing a stopped container.
<i>SERVICED_ISVCS_ENV_[0-9]+</i>	Startup arguments to pass to specific internal services.
<i>SERVICED_SERVICE_MIGRATION_TAG</i>	Overrides the default value for the service migration image.
<i>SERVICED_OPTS</i>	Startup arguments for <code>serviced</code> .
<i>SERVICED_CONTROLLER_BINARY</i>	The path of the <code>serviced-controller</code> binary.
<i>SERVICED_HOME</i>	The path of the home directory for <code>serviced</code> .
<i>SERVICED_ETC_PATH</i>	The path of the directory for <code>serviced</code> configuration files.
<i>SERVICED_VHOST_ALIASES</i>	A list of hostname aliases for a host; for example, <code>localhost</code> .
<i>SERVICED_ZK_CONNECT_TIMEOUT</i>	The number of seconds Control Center waits for a connection to the lead ZooKeeper host.
<i>SERVICED_ZK_PER_HOST_CONNECT_DELAY</i>	The number of seconds Control Center waits before attempting to connect to the next host in its round-robin list of ZooKeeper hosts.

Variable	Description
<i>SERVICED_ZK_RECONNECT_START_DELAY</i> , <i>SERVICED_ZK_RECONNECT_MAX_DELAY</i>	These are used together when Control Center is unable to re-establish a connection with the lead ZooKeeper host.

Best practices for configuration files

The Control Center configuration file, `/etc/default/serviced`, contains Bash environment variables that are read by the `serviced` daemon startup script. The following list describes recommended best practices for its use and maintenance:

- 1 When in doubt, make a backup. Before editing, making a backup copy of the configuration file is always the safest choice.
- 2 Copy a variable, then edit the copy. If you need to revert a variable to its default value, you don't have to leave the file to look it up.
- 3 Copy and edit a variable only if the default value needs to be changed. It's easier to troubleshoot problems when only non-default variables are copied and edited.
- 4 Put the first character of the variable declaration in the first column of its line. It's easier to `grep` for settings when each one starts a line.
- 5 Add customizations to the top of the file. Customizations at the end of the file or scattered throughout the file may be overlooked.
- 6 In high-availability deployments, the contents of `/etc/default/serviced` on the master nodes must be identical. Use a utility like `sum` to compare the files quickly.

Control Center configuration file

The Control Center configuration file, `/etc/default/serviced`, contains Bash environment variables that are read by the `serviced` daemon startup script. The order of the following list matches the order of the variables in the file.

HOME

Default: (the value of shell variable `HOME`)

The path Docker clients use to locate the `.docker/config.json` authentication file, which contains Docker Hub credentials.

TMPDIR

Default: (the value of shell variable `TMPDIR`)

The path `serviced` uses for temporary files.

GOMAXPROCS

Default: 2

The maximum number of CPU cores `serviced` uses.

SERVICED_MASTER

Default: 1 (true)

Assigns the role of a `serviced` instance, either master or delegate. The master runs the application services scheduler and other internal services. Delegates run the application services assigned to the resource pool to which they belong.

Only one `serviced` instance can be the master; all other instances must be delegates. The default value assigns the master role. To assign the delegate role, set the value to 0 (false). This variable must be explicitly set on all Control Center hosts.

SERVICED_MASTER_IP**Default:** 127.0.0.1

A convenience variable, for use in places where the IP address or hostname of the master host is required. This variable is unused unless it is both set here and referenced elsewhere. (For example, by replacing `{{SERVICED_MASTER_IP}}` with `$(SERVICED_MASTER_IP)`.)

SERVICED_MASTER_POOLID**Default:** default

The name of the default resource pool. This variable is only used the first time `serviced` is started.

SERVICED_ZK**Default:** (none)

The list of endpoints in the `serviced` ZooKeeper ensemble, separated by the comma character (`,`). Each endpoint identifies an ensemble node. Each Control Center server and in-container proxy uses `SERVICED_ZK` to create a randomized, round-robin list, and cycles through the list when it attempts to establish a connection with the lead ZooKeeper host.

SERVICED_DOCKER_REGISTRY**Default:** localhost:5000

The endpoint of the local Docker registry, which `serviced` uses to store internal services and application images.

If the default value is changed, the host's Docker configuration file must include the `--insecure-registry` flag with the same value as this variable.

The safest replacement for `localhost` is the IPv4 address of the registry host. Otherwise, the fully-qualified domain name of the host must be specified.

SERVICED_OUTBOUND_IP**Default:** (none)

The IPv4 address that delegates use to connect to the master host. When no address is specified, `serviced` attempts to discover its public IP address by pinging `google.com`.

This variable must be set on all Control Center hosts in either of the following scenarios:

- Control Center is deployed behind a firewall and `google.com` is not reachable. Set the value to the IPv4 address of the master host.
- Control Center is deployed in a high-availability cluster. Set the value to the virtual IPv4 address of the high-availability cluster (*HA-Virtual-IP*).

Note Setting the Docker `HTTP_PROXY` or `HTTPS_PROXY` environment variables prevents access to the IP address defined with this variable. To enable access, unset the Docker variables, and then reboot the host.

SERVICED_STATIC_IPS**Default:** (none)

A list of one or more static IP addresses that are available for IP assignment. Use the comma character (`,`) to separate addresses.

SERVICED_ENDPOINT**Default:** `{{SERVICED_MASTER_IP}}:4979`

The endpoint of the `serviced` RPC server. Replace `{{SERVICED_MASTER_IP}}` with the IP address or hostname of the `serviced` master host. The port number of this endpoint must match the value of the `SERVICED_RPC_PORT` variable defined on the `serviced` master host.

SERVICED_MAX_RPC_CLIENTS**Default:** 3

The preferred maximum number of simultaneous connections a `serviced` delegate uses for RPC requests. The value is used to create a pool of sockets, which are reused as needed. Increasing the value increases the number of open sockets and the use of socket-related operating system resources.

When the demand for connections exceeds the supply of open sockets, `serviced` opens more sockets.

When demand eases, `serviced` reduces the number of open sockets to the preferred maximum.

SERVICED_RPC_PORT**Default:** 4979

The port on which the `serviced` RPC server listens for connections. The value of this variable must match the port number defined for the `SERVICED_ENDPOINT` variable on all `serviced` delegate hosts.

SERVICED_RPC_CERT_VERIFY**Default:** false

Determines whether `serviced` performs TLS certificate verification for RPC connections. The certificate is defined by the `SERVICED_CERT_FILE` variable.

SERVICED_RPC_DISABLE_TLS**Default:** false

Determines whether `serviced` encrypts RPC traffic with TLS.

SERVICED_RPC_TLS_MIN_VERSION**Default:** VersionTLS10

The minimum version of TLS `serviced` accepts for RPC connections. Valid values include the default, VersionTLS11, and VersionTLS12.

SERVICED_RPC_TLS_CIPHERS**Default:** (list of ciphers)

The list of TLS ciphers `serviced` prefers for RPC connections, separated by the comma character (,):

- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256

Other ciphers are supported; the preceding ciphers provide strong security for relatively low processing overhead.

An instance of `serviced` is on both ends of an RPC connection, so both daemons use the first cipher in the list. To use a different cipher, put it first in the list, on all Control Center hosts.

SERVICED_UI_PORT**Default:** :443

The port on which the `serviced` HTTP server listens for requests for its internal services and for tenant services. The value may be expressed as follows:

IP-Address:Port-Number
:Port-Number
Port-Number

Tenant applications can specify alternative ports with the port public endpoint feature.

The value of this variable must be identical on all Control Center hosts in a deployment.

SERVICED_UI_POLL_FREQUENCY**Default:** 3

The number of seconds between polls from Control Center browser interface clients. The value is included in a JavaScript library that is sent to the clients.

SERVICED_MUX_PORT**Default:** 22250

The port `serviced` uses for traffic among Docker containers.

SERVICED_MUX_DISABLE_TLS**Default:** 0

Determines whether inter-host traffic among Docker containers is encrypted with TLS. Intra-host traffic among Docker containers is not encrypted. To disable encryption, set the value to 1.

SERVICED_MUX_TLS_MIN_VERSION**Default:** `VersionTLS10`

The minimum version of TLS `serviced` accepts for mux traffic. Valid values include the default, `VersionTLS11`, and `VersionTLS12`.

SERVICED_MUX_TLS_CIPHERS**Default:** (list of ciphers)

The list of TLS ciphers `serviced` prefers for mux traffic, separated by the comma character (,):

- `TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA`
- `TLS_RSA_WITH_AES_128_CBC_SHA`
- `TLS_RSA_WITH_AES_256_CBC_SHA`
- `TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA`
- `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`

Other ciphers are supported; the preceding ciphers provide strong security for relatively low processing overhead.

An instance of `serviced` is on both ends of a mux connection, so both daemons use the first cipher in the list. To use a different cipher, put it first in the list, on all Control Center hosts.

SERVICED_ISVCS_PATH**Default:** `/opt/serviced/var/isvcs`

The location of `serviced` internal services data.

SERVICED_VOLUMES_PATH**Default:** `/opt/serviced/var/volumes`

The location of `serviced` application data.

SERVICED_BACKUPS_PATH**Default:** `/opt/serviced/var/backups`

The location of `serviced` backup files.

SERVICED_LOG_PATH**Default:** `/var/log/serviced`

The location of `serviced` audit log files. Non-audit (operations) messages are written to `journald`.

SERVICED_KEY_FILE**Default:** `$TMPDIR/zenoss_key.[0-9]+`

The path of a digital certificate key file. Choose a location that is not modified during operating system updates, such as `/etc`.

This key file is used for all TLS-encrypted communications (RPC, mux, and HTTP). The default, insecure key file is created when the `serviced` web server first starts, and is based on a public key that is compiled into `serviced`.

SERVICED_CERT_FILE

Default: `$TMPDIR/zenoss_cert.[0-9]+`

The path of a digital certificate file. Choose a location that is not modified during operating system updates, such as `/etc`. Certificates with passphrases are not supported.

This certificate file is used for all TLS-encrypted communications (RPC, mux, and HTTP). The default, insecure certificate file is created when the `serviced` web server first starts, and is based on a public certificate that is compiled into `serviced`.

SERVICED_TLS_MIN_VERSION

Default: `VersionTLS10`

The minimum version of TLS that `serviced` accepts for HTTP traffic. Valid values include the default, `VersionTLS11`, and `VersionTLS12`.

SERVICED_TLS_CIPHERS

Default: (list of ciphers)

The list of TLS ciphers that `serviced` accepts for HTTP traffic, separated by the comma character (`,`):

- 1 `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`
- 2 `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`
- 3 `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`
- 4 `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`
- 5 `TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA`
- 6 `TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA`
- 7 `TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA`
- 8 `TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA`
- 9 `TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA`
- 10 `TLS_RSA_WITH_AES_256_CBC_SHA`
- 11 `TLS_RSA_WITH_AES_128_CBC_SHA`
- 12 `TLS_RSA_WITH_3DES_EDE_CBC_SHA`
- 13 `TLS_RSA_WITH_AES_128_GCM_SHA256`
- 14 `TLS_RSA_WITH_AES_256_GCM_SHA384`

To disable support for most ciphers, you can remove them from the list. The following rules apply to the list:

- The first cipher, `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, must always be present in the list of ciphers.
- The first four ciphers in the list must always precede any of the ciphers that appear after the first four. The first four ciphers are valid for HTTP/2, while the remaining ciphers are not.

SERVICED_FS_TYPE

Default: `devicemapper`

The driver to manage application data storage on the `serviced` master host. Only `devicemapper` is supported in production deployments.

The only supported storage layout for the `devicemapper` driver is an LVM thin pool. To create a thin pool, use the `serviced-storage` utility. To specify the name of the thin pool device, use the `SERVICED_DM_THINPOOLDEV` variable.

SERVICED_DM_ARGS**Default:** (none)Customized startup arguments for the `devicemapper` storage driver.***SERVICED_DM_BASESIZE*****Default:** 100G

The base size of virtual storage devices for tenants in the application data thin pool, in gigabytes. The units symbol (G) is required. This variable is used when `serviced` starts for the first time, to set the initial size of tenant devices, and when a backup is restored, to set the size of the restored tenant device.

The base size device is sparse device that occupies at most 1MB of space in the application data thin pool; its size has no immediate practical impact. However, the application data thin pool should have enough space for twice the size of each tenant device it supports, to store both the data itself and snapshots of the data. Since the application data thin pool is an LVM logical volume, its size can be increased at any time. Likewise, the size of a tenant device can be increased, as long as the available space in the thin pool can support the larger tenant device plus snapshots.

SERVICED_DM_LOOPDATASIZE**Default:** 100G

Specifies the size of the data portion of the loop-back file. This setting is ignored when `SERVICED_ALLOW_LOOP_BACK` is `false`.

SERVICED_DM_LOOPMETADATASIZE**Default:** 2G

Specifies the size of the metadata portion of the loop-back file. This setting is ignored when `SERVICED_ALLOW_LOOP_BACK` is `false`.

SERVICED_DM_THINPOOLDEV**Default:** (none)The name of the thin pool device to use with the `devicemapper` storage driver.***SERVICED_STORAGE_STATS_UPDATE_INTERVAL*****Default:** 300 (5 minutes)

The number of seconds between polls of kernel statistics about the application data thin pool.

This setting is ignored when the operating system kernel version is less than 3.10.0-366.

SERVICED_ALLOW_LOOP_BACK**Default:** `false`

Determines whether loop-back files can be used with the `devicemapper` storage driver. This option is not supported for production use.

SERVICED_MAX_CONTAINER_AGE**Default:** 86400 (24 hours)The number of seconds `serviced` waits before removing a stopped container.***SERVICED_VIRTUAL_ADDRESS_SUBNET*****Default:** 10.3.0.0/16

The private subnet for containers that use virtual IP addresses on a host. This value may be unique on each Control Center host, if necessary.

RFC 1918 restricts private networks to the 10.0/24, 172.16/20, and 192.168/16 address spaces. However, `serviced` accepts any valid IPv4 address space.

Specify the value in CIDR notation. A /29 network provides sufficient address space.

SERVICED_LOG_LEVEL**Default:** 0

The log level `serviced` uses when writing to the system log. Valid values are 0 (normal) and 2 (debug).

SERVICED_LOG_ADDRESS

Default: `{{SERVICED_MASTER_IP}}:5042`

The endpoint of the logstash service. Replace `{{SERVICED_MASTER_IP}}` with the IP address or hostname of the `serviced` master host.

SERVICED_LOGSTASH_ES

Default: `{{SERVICED_MASTER_IP}}:9100`

The endpoint of the Elasticsearch service for logstash. On delegate hosts, replace `{{SERVICED_MASTER_IP}}` with the IP address or hostname of the Elasticsearch host, which by default is the `serviced` master host.

SERVICED_LOGSTASH_MAX_DAYS

Default: 14

The maximum number of days to keep application logs in the logstash database before purging them.

SERVICED_LOGSTASH_MAX_SIZE

Default: 10

The maximum size of the logstash database, in gigabytes.

SERVICED_LOGSTASH_CYCLE_TIME

Default: 6

The amount of time between logstash purges, in hours.

SERVICED_STATS_PORT

Default: `{{SERVICED_MASTER_IP}}:8443`

The endpoint of the `serviced` metrics consumer service. Replace `{{SERVICED_MASTER_IP}}` with the IP address or hostname of the `serviced` master host.

SERVICED_STATS_PERIOD

Default: 10

The frequency, in seconds, at which delegates gather metrics to send to the `serviced` metrics consumer service on the master host.

SERVICED_SVCSTATS_CACHE_TIMEOUT

Default: 5

The number of seconds to cache statistics about services. The cache is used by Control Center browser interface clients.

SERVICED_DEBUG_PORT

Default: 6006

The port on which `serviced` listens for HTTP requests for the *Go profiler*. To stop listening for requests, set the value to `-1`.

SERVICED_ISVCS_ENV_[0-9]+

Default: (none)

Startup arguments to pass to internal services. You may define multiple arguments, each for a different internal service. The variables themselves, and their arguments, use the following syntax:

SERVICED_ISVCS_ENV_%d

Each variable name ends with a unique integer in place of `%d`.

Service-Name:Key=Value

The value of each variable includes the following elements, in order:

- 1 *Service-Name*, the internal service name. The following command returns the internal service names that may be used for *Service-Name*:

```
docker ps | awk '/serviced-isvcs:/{print $NF}'
```

- 2 The colon character (:).
- 3 *Key*, a variable to pass to the internal service.
- 4 The equals sign character (=).
- 5 *Value*, the definition of the variable to pass to the internal service.

The following example variable passes `ES_JAVA_OPTS=-Xmx4g` to the Elasticsearch internal service.

```
SERVICED_ISVCS_ENV_0=serviced-isvcs_elasticsearch-  
logstash:ES_JAVA_OPTS=-Xmx4g
```

SERVICED_ADMIN_GROUP

Default: wheel

The name of the Linux group on the `serviced` master host whose members are authorized to use the `serviced` browser interface. You may replace the default group with a group that does not have superuser privileges.

SERVICED_ALLOW_ROOT_LOGIN

Default: 1 (true)

Determines whether the `root` user account on the `serviced` master host may be used to gain access to the `serviced` browser interface.

SERVICED_IPTABLES_MAX_CONNECTIONS

Default: 655360

The default value of this variable ensures that a `serviced` delegate does not run out of connections if the `serviced` master goes down. The connections are automatically cleaned up by the kernel soon after the `serviced` master comes back online.

SERVICED_SNAPSHOT_TTL

Default: 12

The number of hours an application data snapshot is retained before removal. To disable snapshot removal, set the value to zero. The application data storage can fill up rapidly when this value is zero or too high.

SERVICED_NFS_CLIENT

Default: 1

DEPRECATED: Prevent a delegate host from mounting the DFS.

SERVICED_SERVICE_MIGRATION_TAG

Default: 1.0.2

Overrides the default value for the service migration image.

SERVICED_ISVCS_START

Default: (none)

Enables one or more internal services to run on a delegate host. Currently, only `zookeeper` has been tested.

SERVICED_ISVCS_ZOOKEEPER_ID

Default: (none)

The unique identifier of a ZooKeeper ensemble node. The identifier must be a positive integer.

SERVICED_ISVCS_ZOOKEEPER_QUORUM**Default:** (none)

The comma-separated list of nodes in a ZooKeeper ensemble. Each entry in the list specifies the ZooKeeper ID, IP address or hostname, peer communications port, and leader communications port of a node in the ensemble. Each quorum definition must be unique, so the IP address or hostname of the "current" host must be 0.0.0.0.

The following example shows the syntax of a node entry:

```
ZooKeeper-ID@Host-IP-Or-Name:2888:3888
```

SERVICED_DOCKER_LOG_DRIVER**Default:** json-file

The log driver for all Docker container logs, including containers for Control Center internal services. Valid values:

- json-file
- syslog
- journald
- gelf
- fluentd
- none

This is a direct port of the Docker `--log-driver` option.

SERVICED_DOCKER_LOG_CONFIG**Default:** max-file=5,max-size=10m

A comma-separated list of Docker `--log-opt` options as `key=value` pairs. To specify the default values for a log driver, or for drivers that need no additional options, such as `journald`, use a single comma character (,) as the value of this variable.

SERVICED_DOCKER_DNS**Default:** (empty)

The IP address of one or more DNS servers. The value of this variable is injected into each Docker container that `serviced` starts. Separate multiple values with the comma character (,).

SERVICED_OPTS**Default:** (empty)

Special options for the `serviced` startup command.

SERVICED_SNAPSHOT_USE_PERCENT**Default:** 20

The amount of free space in the thin pool specified with `SERVICED_DM_THINPOOLDEV`, expressed as a percentage the total size. This value is used to determine whether the thin pool can hold a new snapshot.

SERVICED_ZK_SESSION_TIMEOUT**Default:** 15

The number of seconds the lead ZooKeeper host waits before flushing an inactive connection.

SERVICED_ZK_CONNECT_TIMEOUT**Default:** 1

The number of seconds Control Center waits for a connection to the lead ZooKeeper host.

SERVICED_ZK_PER_HOST_CONNECT_DELAY**Default:** 0

The number of seconds Control Center waits before attempting to connect to the next host in its round-robin list of ZooKeeper hosts. For more information about the round-robin list, see *SERVICED_ZK*.

SERVICED_ZK_RECONNECT_START_DELAY

Default: 1

SERVICED_ZK_RECONNECT_START_DELAY and *SERVICED_ZK_RECONNECT_MAX_DELAY* are used together when Control Center is unable to re-establish a connection with the lead ZooKeeper host.

To prevent unnecessary spikes in TCP traffic, Control Center waits a randomized amount of time that is equal to plus or minus 20% of the value of *SERVICED_ZK_RECONNECT_START_DELAY*. If Control Center is unable to reconnect after contacting all of the hosts in its round-robin list of ZooKeeper hosts, the wait time is increased by a randomized value and the process of attempting to reconnect begins again. If the attempts fail again, the process repeats until the wait time reaches the value of *SERVICED_ZK_RECONNECT_MAX_DELAY*, and the wait time of subsequent reconnection attempts is capped at *SERVICED_ZK_RECONNECT_MAX_DELAY*. Once connection is re-established, the wait time is reset to *SERVICED_ZK_RECONNECT_START_DELAY*.

For more information about the round-robin list, see *SERVICED_ZK*.

SERVICED_ZK_RECONNECT_MAX_DELAY

Default: 1

See *SERVICED_ZK_RECONNECT_START_DELAY*.

SERVICED_ES_STARTUP_TIMEOUT

Default: 240

The number of seconds to wait for the Elasticsearch service to start.

SERVICED_MAX_DFS_TIMEOUT

Default: 300

The number of seconds until a DFS snapshot attempt times out.

SERVICED_RPC_DIAL_TIMEOUT

Default: 30

The number of seconds until an RPC connection attempt times out.

SERVICED_AUTH_TOKEN_EXPIRATION

Default: 3600 (1 hour)

The expiration time, in seconds, of delegate authentication tokens. This timeout affects RPC, mux, and *serviced* internal services endpoint communications.

SERVICED_CONTROLLER_BINARY

Default: `/opt/serviced/bin/serviced-controller`

The path of the *serviced-controller* binary, which runs in every container that *serviced* manages.

SERVICED_HOME

Default: `/opt/serviced`

The path of the home directory for *serviced*.

SERVICED_ETC_PATH

Default: `/opt/serviced/etc`

The path of the directory for *serviced* configuration files. The default is *SERVICED_HOME*/etc.

SERVICED_VHOST_ALIASES

Default: (none)

A list of hostname aliases for a host; for example, `localhost`. Separate multiple values with the comma character (,).