



Zenoss Core Installation Guide

Release 6.0.1

Zenoss, Inc.

www.zenoss.com

Zenoss Core Installation Guide

Copyright © 2017 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.

Amazon Web Services, AWS, and EC2 are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries.

Flash is a registered trademark of Adobe Systems Incorporated.

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

Linux is a registered trademark of Linus Torvalds.

RabbitMQ is a trademark of Pivotal Software, Inc.

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

Sybase is a registered trademark of Sybase, Inc.

Tomcat is a trademark of the Apache Software Foundation.

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

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

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

Part Number: 1651.17.320

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

Contents

About this guide.....	5
Tested operating environments.....	5
Zenoss Core publications.....	6
Change history.....	6
Part I: Appliance deployments.....	9
Chapter 1: Installing a master host.....	10
Creating a virtual machine.....	10
Configuring the Control Center master host.....	15
Editing a connection to configure static IPv4 addressing.....	15
Setting the system hostname.....	17
Adding the master host to a resource pool.....	18
Deploying Zenoss Core.....	19
Chapter 2: Installing delegate hosts.....	21
Creating a virtual machine.....	21
Configuring the virtual machine.....	25
Editing a connection to configure static IPv4 addressing.....	26
Setting the system hostname.....	28
Editing the /etc/hosts file.....	29
Chapter 3: Configuring a multi-host deployment.....	31
Delegate host authentication.....	31
Chapter 4: Configuring a ZooKeeper ensemble.....	34
ZooKeeper and Control Center.....	34
Understanding the configuration process.....	34
Configuring the master host as a ZooKeeper node.....	36
Configuring delegate host A as a ZooKeeper node.....	37
Configuring delegate host B as a ZooKeeper node.....	38
Starting a ZooKeeper ensemble for the first time.....	39
Updating delegate hosts.....	40
Part II: Non-appliance deployments.....	42
Chapter 5: Adding Zenoss Core to a standard deployment.....	43
Downloading template and image files.....	43
Installing the Zenoss Core template.....	43
Importing Zenoss Core images.....	44
Deploying Zenoss Core.....	44

Appendix A: Enabling NTP on Microsoft Hyper-V guests.....	46
Configuring NTP for public time servers.....	46
Configuring an NTP master server.....	47
Configuring NTP clients.....	48

About this guide

Zenoss Core Installation Guide provides detailed procedures for installing Zenoss Core.

This guide supports the following, mutually-exclusive installation paths:

Part I: Appliance deployments describes how to install a Control Center and Zenoss Core virtual appliance as guest systems on VMWare vSphere or Microsoft Hyper-V hypervisors. For more information about using this part, refer to the *Zenoss Core Planning Guide*.

Part II: Non-appliance deployments describes how to add Zenoss Core to an existing Control Center deployment. For more information about creating Control Center deployments, refer to the *Control Center Planning Guide*.

Tested operating environments

Zenoss Core, Control Center, and operating systems

The following table identifies the tested combinations of Zenoss Core, Control Center, and operating system releases.

Zenoss Core release	Control Center	
	Minimum release	Host OS
6.0.1**	1.5.0	RHEL/CentOS 7.2, 7.3, or 7.4 (64-bit)
5.3.0, 5.3.1, 5.3.2	1.3.0, 1.3.1, 1.3.2, 1.3.3, 1.3.4, 1.4.0, 1.4.1	RHEL/CentOS 7.1, 7.2, or 7.3 (64-bit)
5.2.0, 5.2.1, 5.2.2, 5.2.3, 5.2.4, 5.2.6*	1.2.0, 1.2.1, 1.2.2, 1.2.3, 1.3.0, 1.3.1, 1.3.2, 1.3.3, 1.3.4, 1.4.0, 1.4.1	RHEL/CentOS 7.1, 7.2, or 7.3 (64-bit)
5.1.9, 5.1.10	1.1.9, 1.2.0	RHEL/CentOS 7.1 or 7.2 (64-bit)
5.1.8	1.1.5, 1.1.6, 1.1.7, 1.1.8, 1.1.9	RHEL/CentOS 7.1 or 7.2 (64-bit)
5.1.7	1.1.5, 1.1.6, 1.1.7, 1.1.8	RHEL/CentOS 7.1 or 7.2 (64-bit)
5.1.6 (internal release only)	(none)	(none)
5.1.4, 5.1.5	1.1.5, 1.1.6, 1.1.7	RHEL/CentOS 7.1 or 7.2 (64-bit)
5.1.3	1.1.2, 1.1.3, 1.1.5	RHEL/CentOS 7.1 or 7.2 (64-bit)
5.1.2	1.1.2, 1.1.3	RHEL/CentOS 7.1 or 7.2 (64-bit)
5.1.1	1.1.1, 1.1.2	RHEL/CentOS 7.1 or 7.2 (64-bit)

Supported clients and browsers

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

Client OS	Supported browsers
Windows 7, 10	Internet Explorer 11*

** Version 6.0.0 was withdrawn.

* Version 5.2.5 was withdrawn.

* Enterprise mode only; compatibility mode is not tested.

Client OS	Supported browsers
	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

Zenoss Core publications

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

Additional information and comments

Zenoss welcomes your comments and suggestions regarding our documentation. To share your comments, please send an email to docs@zenoss.com. In the email, include the document title (*Zenoss Core Installation Guide*) and part number (1651.17.320).

Change history

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

1651.17.320 (6.0.1)

Update release numbers.

1651.17. (6.0.0)

Restructure appliance installation procedures.

Update non-appliance procedures for offline installation requirements.

Update release numbers.

1051.17.268 (5.3.2)

Update release numbers.

1051.17.242 (5.3.1)

Update release numbers.

1051.17.229 (5.3.0)

Beginning with this release, you can download and install the appliance artifacts or a converged set of artifacts. This document was reorganized and updated with associated information.

Update release numbers.

1051.17.191

Update release numbers.

About 5.2.5

Version 5.2.5 was withdrawn.

1051.17.123 (5.2.4)

Update release numbers.

1051.17.100 (5.2.3)

Add a chapter for high-availability installations.

New procedures for installing the appliance with Hyper-V.

Update release numbers.

1051.17.58 (5.2.2)

Remove procedure for adding storage for backups (no longer required for appliances).

1051.17.044 (5.2.1)

Update release numbers.

1051.16.335 (5.2.0)

Add a section detailing the tested operating environments.

Remove procedures for installing Control Center. That information is now in the *Control Center Installation Guide*.

1051.16.316

Update release numbers.

1051.16.291

Update release numbers.

1051.16.277

Update release numbers.

1051.16.264

Update release numbers.

1051.16.257

Update release numbers.

1051.16.207

Update release numbers.

1051.16.176

Update release numbers.

1051.16.153

Update release numbers.

1051.16.146

Update release numbers.

1051.16.125

Refine the procedure for creating the application data thin pool.

1051.16.118

Add Zenoss Core 5.1.2.

Add a substep to create the docker override directory.

1051.16.111

Add this document change history.

Add chapters describing how to install the Zenoss Core appliance.

Chapters are organized into parts.

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

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

Rather than editing `/lib/systemd/system/docker.service`, Docker configuration steps now include adding a `systemd` override file.

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

Update the commands for starting and testing a ZooKeeper ensemble.

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

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

Add step to install the `Nmap Ncat` package, which is used to check ZooKeeper ensemble status.

1051.16.060

Planning information is now in the *Zenoss Core Planning Guide*.

Information about how to start and configure Zenoss Core is now in the *Zenoss Core Configuration Guide*.

Part I: Appliance deployments

The chapters in this part describe how to install the Zenoss Core appliance, a pre-configured virtual machine that is ready to deploy to your hypervisor. The instructions include a variety of options for customizing your deployment for your environment.

1

Installing a master host

This chapter describes how to install a Zenoss Core appliance package as a Control Center master host. All Zenoss Core deployments require a Control Center master host.

Creating a virtual machine

You can create a virtual machine for the Zenoss Core appliance with VMware vSphere or Microsoft Hyper-V. Choose one of the procedures in this section.

Creating a virtual machine with vSphere

To perform this task, you need:


- A VMware vSphere client
- An account on the [Zenoss Community](#) site.

This procedure installs Zenoss Core OVA packages as a virtual machine managed by vSphere Server version 6.5.0, using VMware vSphere Web Client 6.5. The procedure might differ with other versions of VMware vSphere Client.

- 1 In a web browser, navigate to the download site, and then log in.
The download site is [Zenoss Community](#).
- 2 Download the Zenoss Core master host OVA file.


```
zenoss-core-6.0.1-1-master.vsphere.ova
```

- 3 Use the VMware vSphere Client to log in to vCenter as `root`, or as a user with superuser privileges, and then display the **Home** view.



- 4 Choose **VMs and Templates**.
- 5 In the top navigation bar, choose **Actions > Deploy OVF Template**.

Figure 1: Deploy OVF Template



- 6 Use panels in the **Deploy OVF Template** wizard to select the OVF package:
 - a To choose the package from a drive on your workstation or network share, browse to the location and choose the OVA file. Click **Next**.
 - b **Select name and location:** Specify a name for the OVF, select a datacenter or folder as the deployment location, and then click **Next**.
 - c **Select a resource:** Select the host, cluster, or other resource on which to run the deployed template, and then click **Next**.

- d **Review details:** Verify the template details, and then click **Next**.
 - e **Select storage:** In **Select virtual disk format**, choose **Thin Provision**, accept defaults for other fields, and then click **Next**.
 - f **Select networks:** Accept defaults and click **Next**.
 - g **Ready to Complete:** Verify the deployment settings, and then click **Finish**.
- The **Recent Tasks** pane displays deployment progress and status information.
- 7 Navigate to the new virtual machine's **Getting Started** tab, and then click **Edit virtual machine settings**.
 - 8 Edit the **Virtual Hardware** settings of the virtual machine.
 - a Change the settings.

For multi-host deployments:

 - Reduce the number of CPUs from 8 to 4.
 - Reduce the amount of memory from 32 to 16.

For single-host deployments, reduce the amount of memory from 32 to 24.
 - b Click **OK**.
 - 9 On the new virtual machine's **Getting Started** tab, click **Power on virtual machine**.
 - 10 In the top navigation bar, choose **Actions > Open console**.

Proceed with [Configuring the Control Center master host](#) on page 15.

Creating a virtual machine with Hyper-V

To perform this task, you need:

- Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- An account on the [Zenoss Community](#) site.

Use this procedure to install the Zenoss Core master host appliance as a virtual machine managed by Microsoft Hyper-V.

- 1 Use a Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
- 2 In a web browser, navigate to the download site, and then log in.
The download site is [Zenoss Community](#).
- 3 Download the Zenoss Core master host ISO file.


```
zenoss-core-6.0.1-1-master.x86_64.iso
```
- 4 Open **Hyper-V Manager**.
- 5 In the left navigation pane, choose a server to host the virtual machine.
- 6 From the **Action** menu, choose **New > Virtual Machine...**
The **New Virtual Machine Wizard** opens.
- 7 In the **New Virtual Machine Wizard** dialog, display the **Specify Name and Location** panel.
If the first panel displayed is the **Before You Begin** panel, click **Next**.
- 8 In the **Specify Name and Location** panel, provide a name for the virtual machine, and then click **Next**.
- 9 In the **Specify Generation** panel, choose **Generation 1**, and then click **Next**.
- 10 In the **Assign Memory** panel, specify the memory for the virtual machine.
 - a In the **Startup memory** field, enter the amount of memory for the host.
 - For multi-host deployments, enter 16384 (16GB).
 - For single-host deployments, enter 24576 (24GB).
 - b Optional: Check **Use Dynamic Memory for this virtual machine**, if desired.

Zenoss Core is tested with dynamic memory enabled.

- c Click **Next**.
- 11 In the **Configure Networking** panel, choose a virtual switch, and then click **Next**.
- 12 In the **Connect Virtual Hard Disk** panel, specify a new disk on which to install the guest operating system.
 - a Choose **Create a virtual hard disk**.
 - b Specify a name.
 - c In the **Size** field, enter 30.
 - d Click **Next**.
- 13 In the **Installation Options** panel, specify the master host ISO file.
 - a Choose **Install an operating system from a bootable CD/DVD-ROM**.
 - b Choose **Image file (.iso)**.
 - c Specify or browse to the location of the master host ISO file.
 - d Click **Next**.
- 14 In the **Completing the New Virtual Machine Wizard** panel, verify the description, and then click **Finish**.
Hyper-V Manager creates the new virtual machine, and then closes the wizard.

Configuring and starting a Hyper-V master host

To perform this task, you need:

- A Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- The master host created in the previous procedure (*Creating a virtual machine with Hyper-V* on page 12)

The Zenoss Core master host virtual machine requires a total of 7 virtual hard disks. The following table identifies the purpose and size of each disk.

Table 1: Master host disks

Purpose	Size
1 Root (/)	30GB
2 Swap	16GB
3 Temporary (/tmp)	16GB
4 Docker data	50GB
5 Control Center internal services data	50GB
6 Application data	200GB
7 Application data backups	150GB

Disk 1 was created when the virtual machine was created, in the previous procedure. Use this procedure to configure hardware resources, create disks 2-6, and start the master host.

- 1 Use a Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
- 2 Open **Hyper-V Manager**.
- 3 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Settings**.
The **Settings** dialog displays.
- 4 In the **Hardware** area, locate the virtual hard disk created previously, and then determine whether it is attached to an IDE controller.


Hyper-V guest machines can only boot from an IDE drive.

- 5 In the **Hardware** area, choose **Processor**, and then change the number of processors assigned to the machine.
 - a In **Number of virtual processors**, enter the value for your deployment.
 - For single-host deployments, enter 8.
 - For multi-host deployments, enter 4.
 - b Click **Apply**.
- 6 In the **Hardware** area on the left, choose **SCSI Controller**, and then create additional virtual hard disks. Repeat the following substeps to create new disks in the following sizes:
 - 50GB
 - 150GB
 - 50GB
 - 200GB
 - 16GB
 - 16GB
 - a In the controller area on the right, choose **Hard Drive**, and then click **Add**.
 - b In the **Location** field, choose an unused location number.
 - c In the **Media** area, choose **Virtual hard disk**, and then click **New**.
 - d Complete panels in the **New Virtual Hard Disk Wizard** as follows:
 - 1 **Choose Disk Format**: Choose **VHDX**, and then click **Next**.
 - 2 **Choose Disk Type**: Choose **Dynamically expanding**, and then click **Next**.
 - 3 **Specify Name and Location**: Enter the disk name, and then click **Next**.
 - 4 **Configure Disk**:
 - a Choose **Create a new blank virtual hard disk**.
 - b **Size**: Enter a disk size from the list at the beginning of this step.
 - c Click **Next**.
 - 5 **Summary/New Virtual Hard Disk Wizard**: Verify the description, and then click **Finish**.
 - e At the bottom of the **Settings** window, click **Apply**.

When all of the disks are created, click **OK**.

- 7 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Start**.

Figure 2: Starting a virtual machine



- 8 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Connect**.
- 9 In the **Virtual Machine Connection** window, press **Enter**.
The appliance installation process takes about 15 minutes, and should complete with no additional input. If received, disregard the Fast TSC calibration failure message.

Configuring the Control Center master host

Perform this procedure immediately after creating and starting a Control Center host. All Control Center deployments must include one system that is configured as the master host.

- 1 Gain access to the console interface of the Control Center host through your hypervisor console interface.

Figure 3: Initial hypervisor console login prompt

```

YOU HAVE NOT ACTIVATED THIS APPLIANCE.
PLEASE LOGIN TO ACTIVATE Zenoss Core

Welcome to Zenoss Core

After initial setup, the Control Center UI can be accessed by
browsing to:

    https://core-master

Ensure that core-master is resolvable to 10.88.121.228, either through your
DNS system or through a HOSTS entry on the browser client. For more
information refer to the installation notes.

You can log in to this console to perform administrative tasks such
as setting up networking and safely rebooting this system.

Control Center UI login credentials are ccuser/core.
To log into the console as root, use the default password 'core'.

Linux Kernel 3.10.0-693.2.2.el7.x86_64 on an x86_64
core-master login: _

```

- 2 Log in as the `root` user.
The initial password is provided in the console.
- 3 The system prompts you to enter a new password for `root`.

Note Passwords must include a minimum of eight characters, with at least one character from three of the following character classes: uppercase letter, lowercase letter, digit, and special.

- 4 The system prompts you to enter a new password for `ccuser`.
The `ccuser` account is the default account for gaining access to the Control Center browser interface.

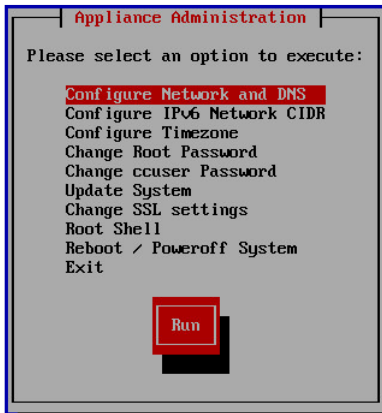
Editing a connection to configure static IPv4 addressing

The default configuration for network connections is DHCP. To configure static IPv4 addressing, perform this procedure.


To navigate in the text user interface (TUI):

- To move forward or backward through options, press the arrow keys.
- To display a menu or choose an option, press **Enter**.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Log in as the `root` user.




- 3 Select the **NetworkManager TUI** menu as follows:
 - a In the **Appliance Administration** menu, select **Configure Network and DNS**, and then press **Enter**.




- 4 On the **NetworkManager TUI** menu, select **Edit a connection**, and then press **Enter**. The TUI displays the connections that are available on the host.

Figure 4: Example: Available connections




Note Do not use this procedure to modify the `docker0` connection.

- 5 Select the virtual connection, and then press **Enter**.

Figure 5: Example: Edit Connection screen

- 6 Optional: If the **IPv4 CONFIGURATION** area is not visible, select its display option (**<Show>**), and then press **Enter**.
- 7 In the **IPv4 CONFIGURATION** area, select **<Automatic>**, and then press **Enter**.


Figure 6: Example: IPv4 Configuration options

- 8 Configure static IPv4 networking as follows:
 - a Select **Manual**, and then press **Enter**.
 - b Beside **Addresses**, select **<Add>**, and then press **Enter**.
 - c In the **Addresses** field, enter an IPv4 address for the virtual machine, and then press **Enter**.
 - d Repeat the preceding two steps for the **Gateway** and **DNS servers** fields.
- 9 Tab to the bottom of the **Edit Connection** screen to select **<OK>**, and then press **Enter**.
- 10 Return to the **Appliance Administration** menu: On the **NetworkManager TUI** screen, select **<Quit>**, and then press **Enter**.
- 11 Reboot the operating system as follows:
 - a In the **Appliance Administration** menu, select **Reboot / Poweroff System**.
 - b Select **Reboot**.
 - c Select **OK**, and then press **Enter**.


Setting the system hostname

The default hostname is `core-master` for the Zenoss Core master host and is `core-delegate` for Zenoss Core delegate hosts. To change the default hostname, perform this procedure.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Select the **NetworkManager TUI** menu as follows:
 - a In the **Appliance Administration** menu, select **Configure Network and DNS**, and then press **Enter**.



- 3 Display the hostname entry field.
 - a In the **NetworkManager TUI** menu, select **Set system hostname**.
 - b Select **OK**, and then press **Enter**.
- 4 In the **Hostname** field, enter the hostname or a fully qualified domain name.
- 5 Press **Tab** twice to select **OK**, and then press **Enter**.



- 6 In the confirmation dialog box, press **Enter**.
- 7 Return to the **Appliance Administration** menu: On the **NetworkManager TUI** screen, select **<Quit>**, and then press **Enter**.
- 8 Reboot the operating system as follows:
 - a In the **Appliance Administration** menu, select **Reboot / Poweroff System**.
 - b Select **Reboot**.
 - c Select **OK**, and then press **Enter**.

Adding the master host to a resource pool

Complete this procedure to add the Control Center master host to the `default` resource pool or to a new resource pool named `master`.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Start a command-line session as `root`.
 - a In the **Appliance Administration** menu, select **Root Shell**.
 - b Select **Run**, and then press **Enter**.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

- 3 Optional: Create a new resource pool, if necessary.
 - For single-host deployments, skip this step.
 - For multi-host deployments, perform the following substeps.

- a Create a new pool named `master`.

```
serviced pool add master
```

- b Assign administrative and distributed file system (DFS) permissions to the new resource pool.

```
serviced pool set-permission --admin --dfs master
```

- 4 Add the master host to a resource pool.

For single-host deployments, add the master host to the `default` resource pool.

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

```
serviced host add Hostname-Or-IP:4979 default
```

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

For multi-host deployments, add the master host to the `master` resource pool and register its authentication token.

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

```
serviced host add --register Hostname-Or-IP:4979 master
```

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

- 5 To exit the command-line session, at the command prompt, enter `exit`.

Deploying Zenoss Core

Use this procedure to add the Zenoss Core application to Control Center and tag application images in the local registry.

- 1 Log in to the Control Center master host as a user with `serviced` CLI privileges.
- 2 Add the Zenoss Core application template to Control Center.

```
serviced template add /opt/serviced/templates/zenoss*.json
```

On success, the `serviced` command returns the template ID.

- 3 Identify the resource pool to which the host belongs.

```
serviced host list
```

- 4 Deploy the application.

- Replace *Template-ID* with the identifier Zenoss Core template
- Replace *Pool* with the name of the resource pool to which the master host belongs (single-host system) or to which the delegate hosts belong (multi-host system)
- Replace *Deployment-ID* with a name for this deployment (for example, `Test` or `Production`)

```
serviced template deploy Template-ID Pool Deployment-ID
```

Control Center tags Zenoss Core images in the local registry.

- If you are creating a single-host deployment, proceed to the *Zenoss Core Configuration Guide*.
- If you are creating a multi-host deployment, proceed to the next chapter.

2

Installing delegate hosts

This chapter describes how to install a Zenoss Core appliance package as a Control Center delegate host. You can add as many delegate hosts as you need to a Control Center deployment.

Perform the procedures in *Installing a master host* on page 10 before performing the procedures in this chapter.

Creating a virtual machine

You can create a virtual machine for the Zenoss Core delegate appliance with VMware vSphere or Microsoft Hyper-V. Choose one of the procedures in this section.

Creating a virtual machine with vSphere

To perform this task, you need:

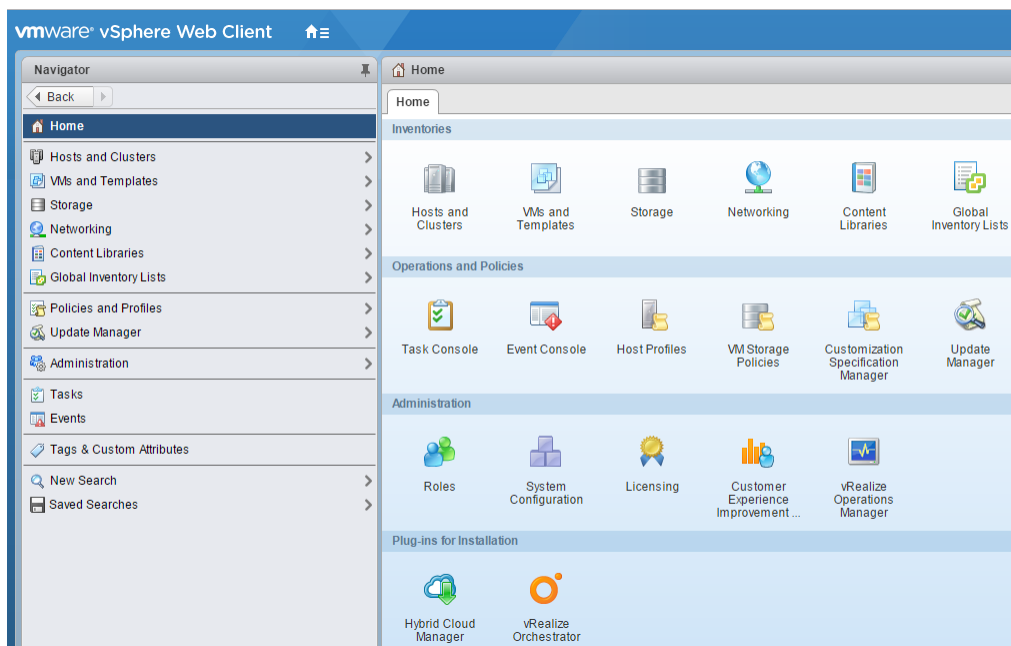
- A VMware vSphere client
- An account on the [Zenoss Community](#) site.

This procedure installs Zenoss Core OVA packages as a virtual machine managed by vSphere Server version 6.5.0, using VMware vSphere Web Client 6.5. The procedure might differ with other versions of VMware vSphere Client.

- 1 In a web browser, navigate to the download site, and then log in.
The download site is [Zenoss Community](#).
- 2 Download the Zenoss Core delegate host OVA file.


```
zenoss-core-6.0.1-1-delegate.vmware.ova
```

- 3 Use the VMware vSphere Client to log in to vCenter as `root`, or as a user with superuser privileges, and then display the **Home** view.



- 4 In the top navigation bar, choose **Actions > Deploy OVF Template**.

Figure 7: Deploy OVF Template



- 5 Use panels in the **Deploy OVF Template** wizard to select the OVF package:
- To choose the package from a drive on your workstation or network share, browse to the location and choose the OVA file. Click **Next**.
 - Select name and location:** Specify a name for the OVF, select a datacenter or folder as the deployment location, and then click **Next**.
 - Select a resource:** Select the host, cluster, or other resource on which to run the deployed template, and then click **Next**.
 - Review details:** Verify the template details, and then click **Next**.

- e **Select storage:** In **Select virtual disk format**, choose **Thin Provision**, accept defaults for other fields, and then click **Next**.
 - f **Select networks:** Accept defaults and click **Next**.
 - g **Ready to Complete:** Verify the deployment settings, and then click **Finish**.
- The **Recent Tasks** pane displays deployment progress and status information.
- 6 Navigate to the new virtual machine's **Getting Started** tab, and then click **Edit virtual machine settings**.
 - 7 On the new virtual machine's **Getting Started** tab, click **Power on virtual machine**.

Creating a virtual machine with Hyper-V

To perform this task, you need:

- Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- An account on the [Zenoss Community](#) site.

Use this procedure to install the Zenoss Core delegate host appliance as a virtual machine managed by Microsoft Hyper-V.

- 1 Use a Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
- 2 In a web browser, navigate to the download site, and then log in.
The download site is [Zenoss Community](#).
- 3 Download the Zenoss Core delegate host ISO file.

```
zenoss-core-6.0.1-1-delegate.x86_64.iso
```

- 4 Open **Hyper-V Manager**.
- 5 In the left navigation pane, choose a server to host the virtual machine.
- 6 From the **Action** menu, choose **New > Virtual Machine...**
The **New Virtual Machine Wizard** opens.
- 7 In the **New Virtual Machine Wizard** dialog, display the **Specify Name and Location** panel.
If the first panel displayed is the **Before You Begin** panel, click **Next**.
- 8 In the **Specify Name and Location** panel, provide a name for the virtual machine, and then click **Next**.
- 9 In the **Specify Generation** panel, choose **Generation 1**, and then click **Next**.
- 10 In the **Assign Memory** panel, specify the memory for the virtual machine.
 - a In the **Startup memory** field, enter 32768 (32GB).
 - b Optional: Check **Use Dynamic Memory for this virtual machine**, if desired.
Zenoss Core is tested with dynamic memory enabled.
 - c Click **Next**.
- 11 In the **Configure Networking** panel, choose a virtual switch, and then click **Next**.
- 12 In the **Connect Virtual Hard Disk** panel, specify a new disk on which to install the guest operating system.
 - a Choose **Create a virtual hard disk**.
 - b Specify a name.
 - c In the **Size** field, enter 30.
 - d Click **Next**.
- 13 In the **Installation Options** panel, specify the delegate host ISO file.
 - a Choose **Install an operating system from a bootable CD/DVD-ROM**.
 - b Choose **Image file (.iso)**.
 - c Specify or browse to the location of the delegate host ISO file.
 - d Click **Next**.
- 14 In the **Completing the New Virtual Machine Wizard** panel, verify the description, and then click **Finish**.

Hyper-V Manager creates the new virtual machine, and then closes the wizard.

Configuring and starting a Hyper-V delegate host

To perform this task, you need:

- A Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- The delegate host created in the previous procedure (*Creating a virtual machine with Hyper-V* on page 23)

Zenoss Core delegate host virtual machines requires a total of 4 virtual hard disks. The following table identifies the purpose and size of each disk.

Table 2: Delegate host disks


Purpose	Size
1 Root (/)	30GB
2 Swap	16GB
3 Temporary (/tmp)	16GB
4 Docker data	50GB

Disk 1 was created when the virtual machine was created, in the previous procedure. Use this procedure to configure hardware resources, create disks 2-4, and start the master host.

- 1 Use a Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
- 2 Open **Hyper-V Manager**.
- 3 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Settings**.
The **Settings** dialog displays.
- 4 In the **Hardware** area, locate the virtual hard disk created previously, and then determine whether it is attached to an IDE controller.
Hyper-V guest machines can only boot from an IDE drive.
- 5 In the **Hardware** area, choose **Processor**, and then change the number of processors assigned to the machine.
 - a In **Number of virtual processors**, enter 8.
 - b Click **Apply**.
- 6 In the **Hardware** area on the left, choose **SCSI Controller**, and then create additional virtual hard disks. Repeat the following substeps to create new disks in the following sizes:
 - 50GB
 - 16GB
 - 16GB
 - a In the controller area on the right, choose **Hard Drive**, and then click **Add**.
 - b In the **Location** field, choose an unused location number.
 - c In the **Media** area, choose **Virtual hard disk**, and then click **New**.
 - d Complete panels in the **New Virtual Hard Disk Wizard** as follows:
 - 1 **Choose Disk Format:** Choose **VHDX**, and then click **Next**.
 - 2 **Choose Disk Type:** Choose **Dynamically expanding**, and then click **Next**.

- 3 **Specify Name and Location:** Enter the disk name, and then click **Next**.
 - 4 **Configure Disk:**
 - a Choose **Create a new blank virtual hard disk**.
 - b **Size:** Enter a disk size from the list at the beginning of this step.
 - c Click **Next**.
 - 5 **Summary/New Virtual Hard Disk Wizard:** Verify the description, and then click **Finish**.
 - e At the bottom of the **Settings** window, click **Apply**.
- When all of the disks are created, click **OK**.
- 7 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Start**.

Figure 8: Starting a virtual machine



- 8 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Connect**.
 - 9 In the **Virtual Machine Connection** window, press **Enter**.
- The appliance installation process takes about 15 minutes, and should complete with no additional input. If received, disregard the Fast TSC calibration failure message.

Configuring the virtual machine

This procedure configures the new virtual machine as a delegate host.

- 1 Gain access to the console interface of the Control Center host through your hypervisor console interface.

Figure 9: Initial hypervisor console login prompt

```

YOU HAVE NOT ACTIVATED THIS APPLIANCE.
PLEASE LOGIN TO ACTIVATE Zenoss Core

Welcome to Zenoss Core

After initial setup, the Control Center UI can be accessed by
browsing to:

    https://core-delegate

Ensure that core-delegate is resolvable to 10.88.121.217, either through your
DNS system or through a HOSTS entry on the browser client. For more
information refer to the installation notes.

You can log in to this console to perform administrative tasks such
as setting up networking and safely rebooting this system.

Control Center UI login credentials are ccuser/core.
To log into the console as root, use the default password 'core'.

Linux Kernel 3.10.0-693.2.2.el7.x86_64 on an x86_64
core-delegate login:

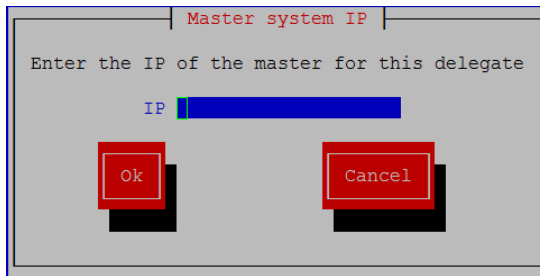
```

- 2 Log in as the `root` user.
The initial password is provided in the console.

- 3 The system prompts you to enter a new password for `root`.

Note Passwords must include a minimum of eight characters, with at least one character from three of the following character classes: uppercase letter, lowercase letter, digit, and special.

- 4 The system prompts you to enter a new password for `ccuser`.
The `ccuser` account is the default account for gaining access to the Control Center browser interface.
- 5 In the **IP** field, enter the hostname, fully qualified domain name, or IPv4 address of the master host.



Note If you enter the hostname or fully qualified domain name of the master host, you need an entry in the `/etc/hosts` file of the delegate host or a nameserver on your network that resolves the name to its IPv4 address.

- a Press **Tab** to select **Ok**, and then press **Enter**.
The system reboots.

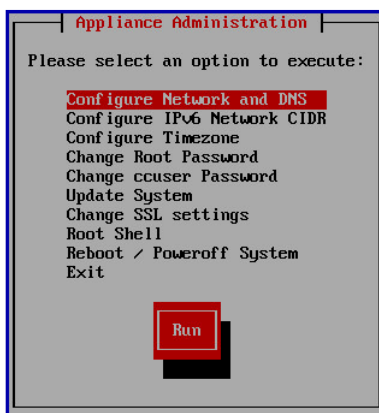
Editing a connection to configure static IPv4 addressing

The default configuration for network connections is DHCP. To configure static IPv4 addressing, perform this procedure.

To navigate in the text user interface (TUI):


- To move forward or backward through options, press the arrow keys.
- To display a menu or choose an option, press **Enter**.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Log in as the `root` user.




- 3 Select the **NetworkManager** TUI menu as follows:

- a In the **Appliance Administration** menu, select **Configure Network and DNS**, and then press **Enter**.



- 4 On the **NetworkManager TUI** menu, select **Edit a connection**, and then press **Enter**. The TUI displays the connections that are available on the host.

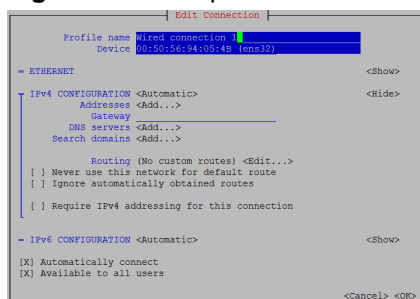
Figure 10: Example: Available connections



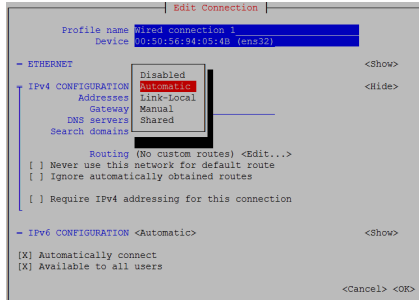
Note Do not use this procedure to modify the `docker0` connection.

- 5 Select the virtual connection, and then press **Enter**.

Figure 11: Example: Edit Connection screen



- 6 Optional: If the **IPv4 CONFIGURATION** area is not visible, select its display option (**<Show>**), and then press **Enter**.
- 7 In the **IPv4 CONFIGURATION** area, select **<Automatic>**, and then press **Enter**.


Figure 12: Example: IPv4 Configuration options

- 8 Configure static IPv4 networking as follows:
 - a Select **Manual**, and then press **Enter**.
 - b Beside **Addresses**, select **<Add>**, and then press **Enter**.
 - c In the **Addresses** field, enter an IPv4 address for the virtual machine, and then press **Enter**.
 - d Repeat the preceding two steps for the **Gateway** and **DNS servers** fields.
- 9 Tab to the bottom of the **Edit Connection** screen to select **<OK>**, and then press **Enter**.
- 10 Return to the **Appliance Administration** menu: On the **NetworkManager TUI** screen, select **<Quit>**, and then press **Enter**.
- 11 Reboot the operating system as follows:
 - a In the **Appliance Administration** menu, select **Reboot / Poweroff System**.
 - b Select **Reboot**.
 - c Select **OK**, and then press **Enter**.


Setting the system hostname

The default hostname is `core-master` for the Zenoss Core master host and is `core-delegate` for Zenoss Core delegate hosts. To change the default hostname, perform this procedure.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Select the **NetworkManager TUI** menu as follows:
 - a In the **Appliance Administration** menu, select **Configure Network and DNS**, and then press **Enter**.



- 3 Display the hostname entry field.
 - a In the **NetworkManager TUI** menu, select **Set system hostname**.
 - b Select **OK**, and then press **Enter**.
- 4 In the **Hostname** field, enter the hostname or a fully qualified domain name.
- 5 Press **Tab** twice to select **OK**, and then press **Enter**.



- 6 In the confirmation dialog box, press **Enter**.
- 7 Return to the **Appliance Administration** menu: On the **NetworkManager TUI** screen, select **<Quit>**, and then press **Enter**.
- 8 Reboot the operating system as follows:
 - a In the **Appliance Administration** menu, select **Reboot / Poweroff System**.
 - b Select **Reboot**.
 - c Select **OK**, and then press **Enter**.

Editing the /etc/hosts file

This procedure is conditional. Perform this procedure only if you use hostnames or fully qualified domain names instead of IPv4 addresses, and only after all delegate hosts are installed and renamed. Perform this procedure on the Control Center master host and on each delegate host.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Start a command-line session as `root`.
 - a In the **Appliance Administration** menu, select **Root Shell**.
 - b Select **Run**, and then press **Enter**.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

- 3 Open the `/etc/hosts` file in a text editor.

The following steps use the *nano* editor.

 - a Start the editor.

```
nano /etc/hosts
```

Figure 13: Example nano session

- b** Optional: On delegate hosts, the file might include two entries with the same IP address. Remove the first of the two entries, which maps the IP address to the `core-master` hostname.
 - c** Add entries for the Control Center master host and for each delegate host.
 - d** To save, press **Control-o**.
 - e** To exit, press **Control-x**.
- 4** Return to the **Appliance Administration** menu.

```
exit
```

- 5** Exit the **Appliance Administration** menu.
- a** Use the down-arrow key to select **Exit**.
 - b** Press **Tab**, and then press **Enter**.

3

Configuring a multi-host deployment

This chapter describes how to configure an appliance-based multi-host deployment, which includes one Control Center master host and one or more delegate hosts. Zenoss recommends deploying at least two delegate hosts, to enable creating a ZooKeeper ensemble.

Note If you are creating a multi-host deployment on a Hyper-V system, you must perform the procedures in [Enabling NTP on Microsoft Hyper-V guests](#) on page 46, in addition to the procedures in this chapter.

Delegate host authentication

Control Center uses RSA key pairs to create the authentication tokens that are required for all delegate communications. When you add a host to a resource pool, the `serviced` instance on the master host creates a private key for the delegate and bundles it with its own public key. The `serviced` instance on the delegate host uses the bundle to sign messages with its unique tokens.

Key bundles are installed by using an SSH connection or a file.

- The command to add a host to a pool can initiate an SSH connection with the delegate and install the key bundle. 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 add a host to a pool creates a file containing the key bundle. You can move the key bundle file to the delegate host with any file transfer method, and then install it on the delegate.

The following procedures demonstrate how to add a host to a resource pool and install its key bundle.

Adding a delegate host through an SSH connection

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 add a delegate host to a resource pool through an SSH connection. Repeat this procedure on each delegate in your Zenoss Core deployment.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Start a command-line session as `root`.
 - a In the **Appliance Administration** menu, select **Root Shell**.
 - b Select **Run**, and then press **Enter**.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

- 3 Add a delegate host to a resource pool.

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.

Use the hostname or IP address to identify a Control Center host. If you use a hostname, all Control Center hosts must be able to resolve it, either through an entry in `/etc/hosts` or through a nameserver on the network. In the following example, replace *Hostname-Or-IP* with the hostname or IP address of a delegate host, and replace *Resource-Pool* with the name of a resource pool.

If the host is behind a router or firewall for network address translation (NAT), include the option `--nat-address` to specify the NAT device's hostname or IP address and port of the delegate host.

```
serviced host add --register Hostname-Or-IP:4979 Resource-Pool \
  --nat-address==NAT-Hostname-Or-IP:NAT-Port
```

Adding a delegate host using a file

Use this procedure to add a delegate host to a resource pool by using a key bundle file. Repeat this procedure on each delegate in your Zenoss Core deployment.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Start a command-line session as `root`.
 - a In the **Appliance Administration** menu, select **Root Shell**.
 - b Select **Run**, and then press **Enter**.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

- 3 Add a delegate host to a resource pool.

Use the hostname or IP address to identify a Control Center host. If you use a hostname, all Control Center hosts must be able to resolve it, either through an entry in `/etc/hosts` or through a nameserver on the network. In the following example, replace *Hostname-Or-IP* with the hostname or IP address of a delegate host, and replace *Resource-Pool* with the name of a resource pool.

If the host is behind a router or firewall for network address translation (NAT), include the option `--nat-address` to specify the NAT device's hostname or IP address and port of the delegate host.

```
serviced host add Hostname-Or-IP:4979 Resource-Pool \
  --nat-address==NAT-Hostname-Or-IP:NAT-Port
```

The command creates a unique key bundle file in the local directory.

- 4 Use a file transfer utility such as `scp` to copy the key bundle file to the delegate host. Once copied to the delegate host, the key bundle file is not needed on the master host and can be deleted.

- 5 Log in to the Control Center delegate host as a user with `serviced` CLI privileges.
- 6 Install the key bundle.
Replace *Key-Bundle-Path* with the pathname of the key bundle file:

```
serviced host register Key-Bundle-Path
```

- 7 Delete the key bundle file.
The file is no longer needed on the delegate host.
Replace *Key-Bundle-Path* with the pathname of the key bundle file:

```
rm Key-Bundle-Path
```

4

Configuring a ZooKeeper ensemble

This chapter describes how to create a ZooKeeper ensemble (cluster) for a multi-host Control Center deployment that includes a minimum of three hosts. If your deployment includes just one host or two hosts, skip this chapter.

ZooKeeper and Control Center

Control Center relies on *Apache ZooKeeper* to distribute and manage application services. ZooKeeper maintains the definitions of each service and the list of services assigned to each host. The scheduler, which runs on the master host, determines assignments and sends them to the ZooKeeper node that is serving as the ensemble leader. The leader replicates the assignments to the other ensemble nodes, so that the other nodes can assume the role of leader if the leader node fails.

All Control Center hosts retrieve assignments and service definitions from the ZooKeeper ensemble leader and then start services in Docker containers as required. So, the Control Center configuration files of all Control Center hosts must include a definition for the `SERVICED_ZK` variable, which specifies the ZooKeeper endpoints of the ensemble nodes. Additional variables are required on ensemble nodes.

A ZooKeeper ensemble requires a minimum of three nodes, which is sufficient for most environments. An odd number of nodes is recommended and an even number of nodes is strongly discouraged. A five-node ensemble improves failover protection during maintenance windows but larger ensembles yield no benefits.

The Control Center master host is always an ensemble node. All ensemble nodes should be on the same subnet.

Understanding the configuration process

The procedures in this chapter instruct you to create temporary variables that are used as building blocks, to construct Control Center configuration variables accurately. You append the Control Center variables to `/etc/default/serviced`, and then edit the file to move the variables to more appropriate locations.

The most important temporary variables specify the IP address or hostname of each host in the ZooKeeper ensemble. The following table identifies these important variables, the names and values of which must be identical on every Control Center host.

Variable name	Placeholder value	Actual value
node1	<i>Master</i>	The IP address or hostname of the master host.
node2	<i>Delegate-A</i>	The IP address or hostname of delegate host A.

Variable name	Placeholder value	Actual value
node3	<i>Delegate-B</i>	The IP address or hostname of delegate host B.

Note All ensemble hosts should be on the same subnet.

ZooKeeper variables

The variables in the following table are set only on ZooKeeper ensemble nodes, except *SERVICED_ZK*, which must be identical on all Control Center hosts.

Variable	Where to set
<i>SERVICED_ISVCS_START</i>	ZooKeeper ensemble nodes
<i>SERVICED_ISVCS_ZOOKEEPER_ID</i>	ZooKeeper ensemble nodes
<i>SERVICED_ISVCS_ZOOKEEPER_QUORUM</i>	ZooKeeper ensemble nodes
<i>SERVICED_ZK</i>	All Control Center hosts
<i>SERVICED_ZK_SESSION_TIMEOUT</i>	ZooKeeper ensemble nodes

Example multi-host ZooKeeper configuration

This example shows the ZooKeeper variables in the `/etc/default/serviced` configuration file of each host in a 4-node Control Center deployment. For convenience, the relevant settings for each node or host are also included in subsequent procedures.

Note The value of the *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable is formatted to fit the available space. In the configuration file, the variable and value are on the same line.

Master host and ZooKeeper ensemble node, 198.51.100.135:

```
SERVICED_ISVCS_ZOOKEEPER_ID=1
SERVICED_ZK=198.51.100.135:2181,198.51.100.136:2181,198.51.100.137:2181
SERVICED_ISVCS_ZOOKEEPER_QUORUM=1@0.0.0.0:2888:3888,\
 2@198.51.100.136:2888:3888,3@198.51.100.137:2888:3888
SERVICED_ZK_SESSION_TIMEOUT=15
```

Delegate host and ZooKeeper ensemble node, 198.51.100.136:

```
SERVICED_ISVCS_START=zookeeper
SERVICED_ISVCS_ZOOKEEPER_ID=2
SERVICED_ZK=198.51.100.135:2181,198.51.100.136:2181,198.51.100.137:2181
SERVICED_ISVCS_ZOOKEEPER_QUORUM=1@198.51.100.135:2888:3888,\
 2@0.0.0.0:2888:3888,3@198.51.100.137:2888:3888
SERVICED_ZK_SESSION_TIMEOUT=15
```

Delegate host and ZooKeeper ensemble node, 198.51.100.137:

```
SERVICED_ISVCS_START=zookeeper
SERVICED_ISVCS_ZOOKEEPER_ID=3
SERVICED_ZK=198.51.100.135:2181,198.51.100.136:2181,198.51.100.137:2181
SERVICED_ISVCS_ZOOKEEPER_QUORUM=1@198.51.100.135:2888:3888,\
```

```
2@198.51.100.136:2888:3888,3@0.0.0.0:2888:3888
SERVICED_ZK_SESSION_TIMEOUT=15
```

Delegate host, 198.51.100.138:

```
SERVICED_ZK=198.51.100.135:2181,198.51.100.136:2181,198.51.100.137:2181
```

Configuring the master host as a ZooKeeper node

This procedure configures the Control Center master host as a node in a ZooKeeper ensemble.

- 1 Log in to the master host as `root`, or as a user with superuser privileges.
- 2 Define the IP address variables for each node in the ZooKeeper ensemble.
Replace *Master* with the IP address or hostname of the Control Center master host, and replace *Delegate-A* and *Delegate-B* with the IP addresses or hostnames of the delegate hosts to include in the ensemble:

```
node1=Master
node2=Delegate-A
node3=Delegate-B
```

- 3 Set the ZooKeeper node ID to 1.

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

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

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

- 5 Specify the nodes in the ZooKeeper quorum.

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

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

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

- 6 Specify the timeout for inactive connections.
You can copy the following text and paste it in your console:

```
echo "SERVICED_ZK_SESSION_TIMEOUT=15" >> /etc/default/serviced
```

- 7 Verify the ZooKeeper environment variables.

```
grep -E '^\\b*SERVICED' /etc/default/serviced | grep -E '_Z(OO|K)'
```

The following example shows the environment variables for a master host with IP address 198.51.100.135.

Note The value of the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable is formatted to fit the available space. The result of the `grep` command shows the variable and value on the same line.

```
SERVICED_ZK=198.51.100.135:2181,198.51.100.136:2181,198.51.100.137:2181
SERVICED_ISVCS_ZOOKEEPER_ID=1
SERVICED_ISVCS_ZOOKEEPER_QUORUM=1@0.0.0.0:2888:3888,\
 2@198.51.100.136:2888:3888,3@198.51.100.137:2888:3888
SERVICED_ZK_SESSION_TIMEOUT=15
```

Configuring delegate host A as a ZooKeeper node

Use this procedure to configure the delegate host designated as *Delegate-A* as a ZooKeeper node.

- 1 Log in to the delegate host as `root`, or as a user with superuser privileges.
- 2 Define the IP address variables for each node in the ZooKeeper ensemble. Replace *Master* with the IP address or hostname of the Control Center master host, and replace *Delegate-A* and *Delegate-B* with the IP addresses or hostnames of the delegate hosts to include in the ensemble:

```
node1=Master
node2=Delegate-A
node3=Delegate-B
```

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

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

- 4 Remove the existing definition of the `SERVICED_ZK` variable, which specifies only the Control Center master host.

```
sed -i.bak '/SERVICED_ZK=/d' /etc/default/serviced
```

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

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

- 6 Specify the nodes in the ZooKeeper quorum.

ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address or hostname of delegate host A with `0.0.0.0`.

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

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

- 7 Specify the timeout for inactive connections. You can copy the following text and paste it in your console:

```
echo "SERVICED_ZK_SESSION_TIMEOUT=15" >> /etc/default/serviced
```

- 8 Configure Control Center to start the ZooKeeper service.

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

```
echo "SERVICED_ISVCS_START=zookeeper" >> /etc/default/serviced
```

9 Verify the ZooKeeper environment variables.

```
grep -E '^\\b*SERVICED' /etc/default/serviced \
| grep -E '(CS_ZO|_ZK|CS_ST)'
```

The following example shows the environment variables for a delegate host with IP address 198.51.100.136.

Note The value of the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable is formatted to fit the available space. The result of the `grep` command shows the variable and value on the same line.

```
SERVICED_ZK=198.51.100.135:2181,198.51.100.136:2181,198.51.100.137:2181
SERVICED_ISVCS_START=zookeeper
SERVICED_ISVCS_ZOOKEEPER_ID=2
SERVICED_ISVCS_ZOOKEEPER_QUORUM=1@198.51.100.135:2888:3888,\
2@0.0.0.0:2888:3888,3@198.51.100.137:2888:3888
SERVICED_ZK_SESSION_TIMEOUT=15
```

Configuring delegate host B as a ZooKeeper node

Use this procedure to configure the delegate host designated as *Delegate-B* as a ZooKeeper node.

- 1 Log in to the delegate host as `root`, or as a user with superuser privileges.
- 2 Define the IP address variables for each node in the ZooKeeper ensemble.
Replace *Master* with the IP address or hostname of the Control Center master host, and replace *Delegate-A* and *Delegate-B* with the IP addresses or hostnames of the delegate hosts to include in the ensemble:

```
node1=Master
node2=Delegate-A
node3=Delegate-B
```

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

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

- 4 Remove the existing definition of the `SERVICED_ZK` variable, which specifies only the Control Center master host.

```
sed -i.bak '/SERVICED_ZK=/d' /etc/default/serviced
```

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

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

- 6 Specify the nodes in the ZooKeeper quorum.
ZooKeeper requires a unique quorum definition for each node in its ensemble. To achieve this, replace the IP address or hostname of delegate host B with `0.0.0.0`.
You can copy the following text and paste it in your console:

```
q1="1@${node1}:2888:3888"
```

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

- 7 Specify the timeout for inactive connections.

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

```
echo "SERVICED_ZK_SESSION_TIMEOUT=15" >> /etc/default/serviced
```

- 8 Configure Control Center to start the ZooKeeper service.

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

```
echo "SERVICED_ISVCS_START=zookeeper" >> /etc/default/serviced
```

- 9 Verify the ZooKeeper environment variables.

```
grep -E '^\\b*SERVICED' /etc/default/serviced \
  | grep -E '(CS_ZO|_ZK|CS_ST)'
```

The following example shows the environment variables for a delegate host with IP address 198.51.100.137.

Note The value of the `SERVICED_ISVCS_ZOOKEEPER_QUORUM` variable is formatted to fit the available space. The result of the `grep` command shows the variable and value on the same line.

```
SERVICED_ZK=198.51.100.135:2181,198.51.100.136:2181,198.51.100.137:2181
SERVICED_ISVCS_START=zookeeper
SERVICED_ISVCS_ZOOKEEPER_ID=3
SERVICED_ISVCS_ZOOKEEPER_QUORUM=1@198.51.100.135:2888:3888,\
  2@198.51.100.136:2888:3888,3@0.0.0.0:2888:3888
SERVICED_ZK_SESSION_TIMEOUT=15
```

Starting a ZooKeeper ensemble for the first time

Use this procedure to start a ZooKeeper ensemble.

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

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

```
systemctl stop serviced
```

- 5 When `serviced` is stopped on all ensemble hosts, start `serviced` on all ensemble hosts at the same time.

```
systemctl start serviced
```

- 6 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 /bin/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*.

- c Query delegate host A and identify its role in the ensemble.
Replace *Delegate-A* with the hostname or IP address of delegate host A:

```
{ echo stats; sleep 1; } | nc Delegate-A 2181 | grep Mode
```

- d Query delegate host B and identify its role in the ensemble.
Replace *Delegate-B* with the hostname or IP address of delegate host B:

```
{ echo stats; sleep 1; } | nc Delegate-B 2181 | grep Mode
```

- e Detach from the container of the ZooKeeper service.

```
exit
```

If none of the hosts reports that it is the ensemble leader within a few minutes of starting *serviced*, reboot the hosts.

Updating delegate hosts

The default configuration of delegate hosts sets the value of the *SERVICED_ZK* variable to the master host only. Use this procedure to update the setting to include all of the hosts in the ZooKeeper ensemble. Perform this procedure on each delegate host that is not an ensemble node.

- 1 Log in to the delegate host as *root*, or as a user with superuser privileges.
- 2 Define the IP address variables for each node in the ZooKeeper ensemble.
Replace *Master* with the IP address or hostname of the Control Center master host, and replace *Delegate-A* and *Delegate-B* with the IP addresses or hostnames of the delegate hosts to include in the ensemble:

```
node1=Master
node2=Delegate-A
node3=Delegate-B
```

- 3 Remove the existing definition of the *SERVICED_ZK* variable, which specifies only the Control Center master host.

```
sed -i.bak '/SERVICED_ZK=/d' /etc/default/serviced
```

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

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


5 Verify the setting.

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

The following example shows the environment variable for a delegate host that is not a node in the ZooKeeper ensemble:

```
SERVICED_ZK=198.51.100.135:2181,198.51.100.136:2181,198.51.100.137:2181
```

6 Restart Control Center.

```
systemctl restart serviced
```

Part II: Non-appliance deployments

The chapters in this part describe how to add the Zenoss Core application to a Control Center deployment. You can add Zenoss Core to a standard deployment of Control Center or to a high-availability deployment.

Adding Zenoss Core to a standard deployment

5

This chapter describes how to add the Zenoss Core application to a standard deployment of Control Center. For more information about creating a standard Control Center deployment, refer to the following documents:

- *Control Center Planning Guide*
- *Zenoss Core Planning Guide*
- *Control Center Installation Guide*

Downloading template and image files

To perform this procedure, you need:

- A workstation with internet access
- An account on the [Zenoss Community](#) site.
- A secure network copy program

Use this procedure to

- download required files to a workstation
 - copy the files to a Control Center master host
- 1 In a web browser, navigate to the download site, and then log in.
The download site is [Zenoss Community](#).
 - 2 Download the self-installing Docker image files for Zenoss Core.
 - `install-zenoss-hbase-24.0.8.run`
 - `install-zenoss-opentsdb-24.0.8.run`
 - `install-zenoss-core_6.0-6.0.1_1.run`
 - 3 Download the Zenoss Core service definition, which is distributed as an RPM file.

```
zenoss-core-service-6.0.1-1.noarch.rpm
```

- 4 Use a secure copy program to copy the files to the Control Center master host.

Installing the Zenoss Core template

Use this procedure to install the Zenoss Core service definition template on the Control Center master host.

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

- 2 Move the RPM file to /tmp.

```
mv zenoss-core-service-6.0.1-1.noarch.rpm /tmp
```

- 3 Install the Zenoss Core template file.

```
yum install /tmp/zenoss-core-service-6.0.1-1.noarch.rpm
```

The template file is stored in /opt/serviced/templates.

Importing Zenoss Core images

Use this procedure to import the Zenoss Core images into the local registry.

- 1 Log in to the Control Center master host as root, or as a user with superuser privileges.
- 2 Move the Docker image files to /root.

```
mv install-zenoss-*.run /root
```

- 3 Add execute permission to the image files.

```
chmod +x /root/install-zenoss-*.run
```

- 4 Change directory to /root.

```
cd /root
```

- 5 Import the images.

The images are contained in self-extracting archive files.

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

- 6 List the images in the registry.

```
docker images
```

The result should include one image for each archive file.

- 7 Optional: Delete the archive files.

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

Deploying Zenoss Core

Use this procedure to add the Zenoss Core application to Control Center, and to tag application images in the local registry.

- 1 Log in to the Control Center master host as a user with serviced CLI privileges.
- 2 Add the Zenoss Core application template to Control Center.

```
serviced template add /opt/serviced/templates/zenoss*.json
```

On success, the `serviced` command returns the template ID.

- 3 Identify the resource pool to which the host belongs.

```
serviced host list
```

- 4 Deploy the application.

- Replace *Template-ID* with the identifier Zenoss Core template
- Replace *Pool* with the name of the resource pool to which the master host belongs (single-host system) or to which the delegate hosts belong (multi-host system)
- Replace *Deployment* with a name for this deployment (for example, `Test` or `Production`)

```
serviced template deploy Template-ID Pool Deployment
```

Control Center tags Zenoss Core images in the local registry.

Zenoss Core is ready to be configured for your environment. For more information, refer to the *Zenoss Core Configuration Guide*.

A

Enabling NTP on Microsoft Hyper-V guests

Control Center requires a common time source. The procedures in this appendix configure *NTP* to synchronize the system clocks of Control Center hosts.

Note The procedures in this appendix are required only for multi-host deployments running as Microsoft Hyper-V guests. VMware vSphere guests use an hourly `cron` job to synchronize their system clocks with the host.

Configure NTP to rely on a time source as follows, depending on internet access of the Control Center hosts:

- If all hosts can access the internet, configure NTP to rely on public time servers.
- If no host can access the internet, configure NTP to rely on a private master server.

Configuring NTP for public time servers

This procedure uses the default configuration of NTP to synchronize system clocks with public time servers. If all Control Center hosts can access the internet, repeat this procedure on each host, starting with the Control Center master host.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Start a command-line session as `root`.
 - a In the **Appliance Administration** menu, select **Root Shell**.
 - b Select **Run**, and then press **Enter**.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

- 3 Stop Control Center.

```
systemctl stop serviced
```

- 4 Synchronize the system clock and enable the NTP daemon.
 - a Set the system time.

```
ntpdate -gq
```

- b Enable the `ntpd` daemon.

```
systemctl enable ntpd
```

- c Configure `ntpd` to start when the system starts.

Currently, an unresolved issue associated with NTP prevents `ntpd` from restarting correctly after a reboot, and the following commands provide a workaround to ensure that it does.

```
echo "systemctl start ntpd" >> /etc/rc.d/rc.local
chmod +x /etc/rc.d/rc.local
```

- d Start `ntpd`.

```
systemctl start ntpd
```

- 5 Start Control Center.

```
systemctl start serviced
```

Configuring an NTP master server

(*Hyper-V only*) This procedure configures an NTP master server on the Control Center master host. Perform this procedure only if the host does not have internet access.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Start a command-line session as `root`.
 - a In the **Appliance Administration** menu, select **Root Shell**.
 - b Select **Run**, and then press **Enter**.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

- 3 Create a backup of the NTP configuration file.

```
cp -p /etc/ntp.conf /etc/ntp.conf.orig
```

- 4 Edit the NTP configuration file.

- a Open `/etc/ntp.conf` with a text editor.
- b Replace all of the lines in the file with the following lines:

```
# Use the local clock
server 127.127.1.0 prefer
fudge 127.127.1.0 stratum 10
driftfile /var/lib/ntp/drift
broadcastdelay 0.008

# Give localhost full access rights
restrict 127.0.0.1

# Grant access to client hosts
restrict Address-Range mask Netmask nomodify notrap
```

- c Replace *Address-Range* with the range of IPv4 network addresses that are allowed to query this NTP server.

For example, the following IP addresses are assigned to Control Center hosts:

```
203.0.113.10
203.0.113.11
203.0.113.12
203.0.113.13
```

For the preceding addresses, the value for *Address-Range* is `203.0.113.0`.

- d Replace *Netmask* with the IPv4 network mask that corresponds with the address range.
For example, a valid network mask for `203.0.113.0` is `255.255.255.0`.
- e Save the file and exit the editor.

5 Stop Control Center.

```
systemctl stop serviced
```

6 Enable and start the NTP daemon as follows:

- a Enable the `ntpd` daemon.

```
systemctl enable ntpd
```

- b Configure `ntpd` to start when the system starts.

Currently, an unresolved issue associated with NTP prevents `ntpd` from restarting correctly after a reboot, and the following commands provide a workaround to ensure that it does.

```
echo "systemctl start ntpd" >> /etc/rc.d/rc.local
chmod +x /etc/rc.d/rc.local
```

- c Start `ntpd`.

```
systemctl start ntpd
```

7 Start Control Center.

```
systemctl start serviced
```

Configuring NTP clients

(*Hyper-V only*) This procedure configures delegates to synchronize their clocks with the NTP server on the Control Center master host. Perform this procedure only if the delegates do not have internet access. Repeat this procedure on each Control Center delegate host.

- 1 Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as *PuTTY*.
- 2 Start a command-line session as `root`.
 - a In the **Appliance Administration** menu, select **Root Shell**.
 - b Select **Run**, and then press **Enter**.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```


- 3 Create a backup of the NTP configuration file.

```
cp -p /etc/ntp.conf /etc/ntp.conf.orig
```

- 4 Edit the NTP configuration file./

- a Open `/etc/ntp.conf` with a text editor.
b Replace all of the lines in the file with the following lines:

```
# Point to the master time server
server Master-Address

restrict default ignore
restrict 127.0.0.1
restrict Master-Address mask 255.255.255.255 nomodify notrap noquery

driftfile /var/lib/ntp/drift
```

- c Replace both instances of *Master-Address* with the IPv4 address of the host where the NTP server is running (the Control Center master host).
d Save the file and exit the editor.
- 5 Stop Control Center.

```
systemctl stop serviced
```

- 6 Synchronize the clock with the master server.

```
ntpdate -gq
```

- 7 Enable and start the NTP daemon as follows:

- a Enable the `ntpd` daemon.

```
systemctl enable ntpd
```

- b Configure `ntpd` to start when the system starts.

Currently, an unresolved issue associated with NTP prevents `ntpd` from restarting correctly after a reboot, and the following commands provide a workaround to ensure that it does.

```
echo "systemctl start ntpd" >> /etc/rc.d/rc.local
chmod +x /etc/rc.d/rc.local
```

- c Start `ntpd`.

```
systemctl start ntpd
```

- 8 Start Control Center.

```
systemctl start serviced
```