



Zenoss Core Installation Guide

Release 5.2.5

Zenoss, Inc.

www.zenoss.com

Zenoss Core Installation Guide

Copyright © 2017 Zenoss, Inc. All rights reserved.

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

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: 1051.17.171

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

Contents

About this guide.....	4
Supported operating systems and browsers.....	7
Supported operating environments.....	8
Chapter 1: Installing Zenoss Core.....	9
Installing the Zenoss Core template.....	9
Deploying Zenoss Core.....	9
Part I: Appliance deployments.....	11
.....	11
Chapter 2: Installing a Control Center master host.....	12
Creating a virtual machine.....	12
Configuring the Control Center master host.....	16
Edit a connection to configure static IPv4 addressing.....	17
Setting the system hostname.....	19
Adding the master host to a resource pool.....	20
Deploying Zenoss Core.....	20
Chapter 3: Installing delegate hosts.....	23
Creating a virtual machine.....	23
Configuring the virtual machine.....	26
Edit a connection to configure static IPv4 addressing.....	27
Setting the system hostname.....	30
Editing the /etc/hosts file.....	31
Chapter 4: Configuring a multi-host Control Center cluster.....	32
Enabling NTP on Microsoft Hyper-V guests.....	32
Adding hosts to the default resource pool.....	35
Chapter 5: Configuring a ZooKeeper ensemble.....	37
ZooKeeper and Control Center.....	37
Understanding the configuration process.....	37
Configuring the master host as a ZooKeeper node.....	39
Configuring delegate host A as a ZooKeeper node.....	40
Configuring delegate host B as a ZooKeeper node.....	41
Starting a ZooKeeper ensemble for the first time.....	42
Updating delegate hosts.....	43

About this guide

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

Related 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 and part number. The part number appears at the end of the list of trademarks, at the front of this guide.

Change history

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

1051.17.171 (5.2.5)

Update release numbers.

1051.17.123 (5.2.4)

Update release numbers.

105117.100 (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 supported operating environments.

Remove procedures for installing Control Center clusters. 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 support for 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 hosts in a Control Center cluster.

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*.

Supported operating systems and browsers

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

Client OS	Supported Browsers
Windows 7 and 8.1	Internet Explorer 11 (Enterprise mode only; compatibility mode is not supported.)
	Internet Explorer 10*
Windows 10	Internet Explorer 11 (Enterprise mode only; compatibility mode is not supported.)
	Internet Explorer 10*
	Firefox 50 and later
	Chrome 54 and later
	Microsoft Edge
Windows Server 2012 R2	Firefox 30
	Chrome 36
Macintosh OS/X 10.9	Firefox 30 and above
	Chrome 36 and above
Ubuntu 14.04 LTS	Firefox 30 and above
	Chrome 37 and above
Red Hat Enterprise Linux 6.5, CentOS 6.5	Firefox 30 and above
	Chrome 37 and above

* Support for Internet Explorer 10 will be withdrawn beginning with the next feature release of Zenoss Core.

Supported operating environments

Zenoss Core, Control Center, and operating systems

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

Zenoss Core Release	Control Center	
	Release	Host OS
5.2.0, 5.2.1, 5.2.2, 5.2.3, 5.2.4, 5.2.5	1.2.0, 1.2.1, 1.2.2, 1.2.3, 1.3.0, 1.3.1, 1.3.2, 1.3.3	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)

1

Installing Zenoss Core

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 Install the Zenoss Core template file.

```
yum --enablerepo=zenoss-stable install -y zenoss-core-service-5.2.5
```

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

Deploying Zenoss Core

Use this procedure to add the Zenoss Core application to Control Center, and to pull application images into 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 pulls Zenoss Core images into the local registry, or just tags the images.

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

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.

Installing a Control Center 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
- Permission to download Zenoss Core software from the [Zenoss Support](#) 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.

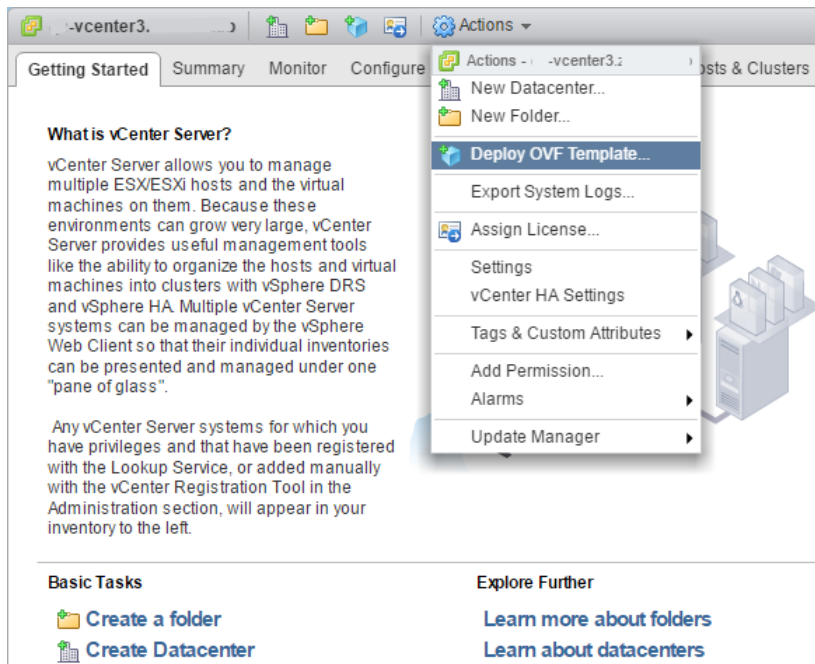
There are two separate OVA packages: one for the master host and one for delegates.

- 1 Download the Zenoss Core *master* host OVA file from the [Zenoss Support](#) site to your workstation.
- 2 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.



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

Figure 1: Deploy OVF Template



- 5 Use panels in the **Deploy OVF Template** wizard to select the OVF package:
 - a **Select template:** To download and install the package from the internet, enter a URL. 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.
- 6 Navigate to the new virtual machine's **Getting Started** tab, and then click **Edit virtual machine settings**.
 - 7 Edit **Virtual Hardware** settings for the virtual machine as follows. Accept default settings for other options.
 - **CPU:** For single-host deployments, specify 8; for multi-host deployments, specify 4.
 - **Cores per socket:** Specify 1.
 - **Memory:** For single-host deployments, specify 64GB; for multi-host deployments, specify 16GB.
- Click **OK**.
- 8 On the new virtual machine's **Getting Started** tab, click **Power on virtual machine**.
 - 9 In the top navigation bar, choose **Actions > Open console**.

Proceed with [Configuring the Control Center master host](#) on page 16.

Creating a virtual machine with Hyper-V

To perform this task, you need:

- A Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- Permission to download Zenoss Core software from the [Zenoss Support](#) site

This procedure installs the Zenoss Core appliance as a virtual machine managed by Microsoft Hyper-V. There are separate ISO files for the master host and delegates.

In this procedure, you download and install the master ISO file. A separate procedure provides instructions for installing delegates.

- 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 Download the Zenoss Core ISO file from the [Zenoss Support](#) site to the Hyper-V host.
Choose the *master* ISO file.
- 3 Open **Hyper-V Manager**.
- 4 In the left navigation pane, choose a server to host the virtual machine.
- 5 In the **Actions** pane, choose **New > Virtual Machine**.
The **New Virtual Machine Wizard** opens.
- 6 Use panels in the wizard to create the virtual machine:
 - a **Specify Name and Location:** Provide a name for the virtual machine, and then click **Next**.
 - b **Specify Generation:** Choose **Generation 1**, and then click **Next**.
 - c **Assign Memory:** Specify memory information as follows:
 - **Startup memory**
 - Multi-host deployment: 16384 (16GB).
 - Single-host deployments: 65536 (64GB).
 - Check **Use Dynamic Memory for this virtual machine**.
 - Click **Next**.
 - d **Configure Networking:** For **Connection**, choose a virtual switch, and then click **Next**.

- e **Connect Virtual Hard Disk:** Create a new disk on which to install the guest operating system as follows:

- 1 Choose **Create a virtual hard disk**.
- 2 Specify a name.
- 3 **Size:** Enter 30.
- 4 Click **Next**.

In a later step, you will create additional disks.

- f **Installation Options:** Access the master ISO file as follows:

- 1 Choose **Install an operating system from a bootable CD/DVD-ROM**.
- 2 Choose **Image file (.iso)**.
- 3 Specify or browse to the location of the ISO image file.

- g **Summary/Completing the New Virtual Machine Wizard:** Verify the description, and then click **Finish**.

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

- 7 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Settings**.
- 8 In the **Hardware** area on the left, configure as follows:
 - a Ensure that the virtual hard disk that you created is attached to an IDE controller.
A Hyper-V guest machine can only boot from an IDE drive.
 - b Choose **Processor**.
 - c In the **Processor** area on the right, in **Number of virtual processors**, enter the following value for your deployment type:
 - Multi-host deployment: 4
 - Single-host deployment: 8
 - d Click **Apply**.
- 9 Create additional virtual hard disks as follows:
 - a In the **Hardware** area on the left, choose any controller and location that will accommodate the disk.
 - b In the controller area on the right, choose **Hard Drive**, and then click **Add**.
 - c In the area on the right, 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 50.
 - c Click **Next**.
 - 5 **Summary/New Virtual Hard Disk Wizard:** Verify the description, and then click **Finish**.
- e In the **Settings** window, click **Apply**.
- f Repeat this step to create additional virtual hard disks in the following sizes (for a total of seven disks):
 - 150GB
 - 50GB
 - 200GB
 - 16GB


- 16GB

When viewed from the Linux system, the disks might be reordered; however, the installation is not affected by disk order.

g In the **Settings** dialog box, click **Apply**, and then click **OK**.

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

Figure 2: Starting a virtual machine



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

- 12 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

```

Welcome to Zenoss Core

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

  https://core-master
  (default username/password is ccuser/core)

Ensure that core-master is resolvable to 10.88.121.218, 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. The
root password defaults to 'core'

Linux Kernel 3.10.0-514.2.2.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.

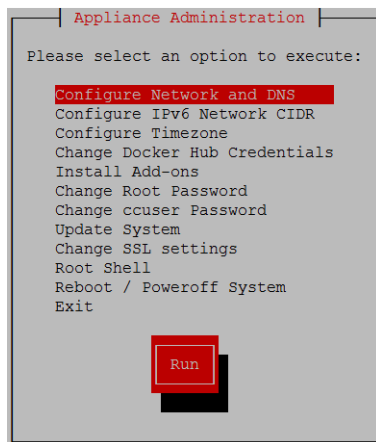
- 4 The system prompts you to enter a new password for `ccuser`.
The `ccuser` account is the default account for gaining access to the Control Center browser interface.

Edit 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



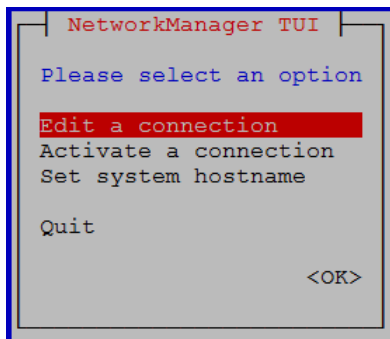
- 8 Configure static IPv4 networking as follows:
 - a Select **Manual**, and then press **Enter**.
 - b Beside **Addresses**, select **<Add>**, and then press **Enter**.

- 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` as follows:
 - 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 cluster 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 cluster 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

You can use the browser interface or the command-line interface to deploy Zenoss Core. Choose one of the procedures in this section.

Using the CLI to deploy

This procedure adds the Zenoss Core application to the list of applications that Control Center manages.

- 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` as follows:
 - 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 the `zenoss.core` application to Control Center.

```
myPath=/opt/serviced/templates
serviced template add $myPath/zenoss-core-*.json
```

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

- 4 Deploy the application as shown in the following example command, replacing variables with your values as shown.
 - Replace *Template-ID* with the template identifier that was returned in the previous step.
 - Replace *Deployment-ID* with a name for this deployment (for example, `Dev` or `Test`).

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

Control Center tags Zenoss Core images in the local registry.

Using the browser interface to deploy

When you log in to the Control Center browser interface for the first time, perform this procedure.


- 1 Display the login page of the Control Center browser interface.

```
https://IP-Address:50443
```

Replace *IP-Address* with the IP address or, if name resolution of the hostname is configured in your environment, the hostname of the Control Center virtual machine. The Control Center web server listens at the hostname or IP address of the Control Center master host and port 50443.

The web server that listens for Control Center interface requests typically takes a few minutes to start after the Control Center master host is started.

Figure 7: Control Center login page



- 2 At the login page, enter `ccuser` and its password.
If the **Deployment Wizard** does not display automatically, click **Application**, located at the right side of the page.
 - 3 In the **Select Applications** step, select the application name (**Zenoss.core**), and then click **Next**.
 - 4 In the **Select Resource Pool** step, select `default`, and then click **Next**.
 - 5 In the **Deployment Applications** step, enter an identifier for this deployment of Zenoss Core, and then click **Deploy**.
 - 6 At the top of the page, click **Logout**.
- 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.

3

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 cluster.

Perform the procedures in *Installing a Control Center master host* on page 12 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
- Permission to download Zenoss Core software from the [Zenoss Support](#) 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.

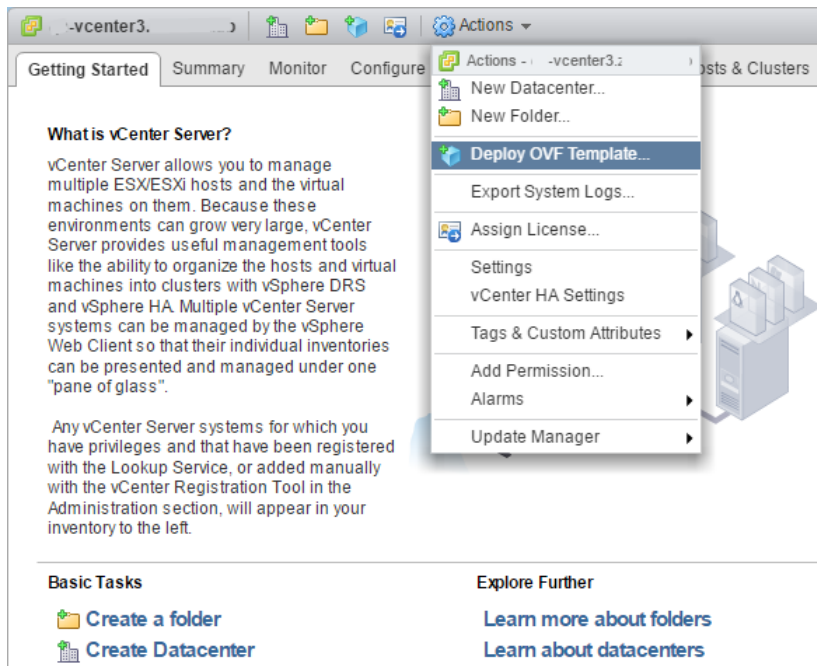
There are two separate OVA packages: one for the master host and one for delegates.

- 1 Download the Zenoss Core delegate OVA file from the [Zenoss Support](#) site to your workstation, if necessary.
- 2 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.



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

Figure 8: Deploy OVF Template



- 4 Use panels in the **Deploy OVF Template** wizard to select the OVF package:
 - a **Select template:** To download and install the package from the internet, enter a URL. 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.
- 5 Navigate to the new virtual machine's **Getting Started** tab, and then click **Edit virtual machine settings**.
- 6 Update the memory assigned to the machine to 32GB, and then click **OK**.
- 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:

- A Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- Permission to download Zenoss Core software from the [Zenoss Support](#) site

This procedure installs the Zenoss Core appliance as a virtual machine managed by Microsoft Hyper-V. There are separate ISO files for the master host and delegates.

In this procedure, you download and install the delegate ISO file. A separate procedure provides instructions for installing the master.


- 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 Download the Zenoss Core ISO file from the [Zenoss Support](#) site to the Hyper-V host.
Choose the *delegate* ISO file.
- 3 Open **Hyper-V Manager**.
- 4 In the left navigation pane, choose a server to host the virtual machine.
- 5 In the **Actions** pane, choose **New > Virtual Machine**.
The **New Virtual Machine Wizard** opens.
- 6 Use panels in the wizard to create the virtual machine:
 - a **Specify Name and Location:** Provide a name for the virtual machine, and then click **Next**.
 - b **Specify Generation:** Choose **Generation 1**, and then click **Next**.
 - c **Assign Memory:** Specify memory information as follows:
 - **Startup memory:** Enter 32768 (32GB).
 - Check **Use Dynamic Memory for this virtual machine**.
 - Click **Next**.
 - d **Configure Networking:** For **Connection**, choose a virtual switch, and then click **Next**.
 - e **Connect Virtual Hard Disk:** Create a new disk on which to install the guest operating system as follows:
 - 1 Choose **Create a virtual hard disk**.
 - 2 Specify a name.
 - 3 **Size:** Enter 30.
 - 4 Click **Next**.

In a later step, you will create additional disks.

- f **Installation Options:** Access the delegate ISO file as follows:
 - 1 Choose **Install an operating system from a bootable CD/DVD-ROM**.
 - 2 Choose **Image file (.iso)**.

- 3 Specify or browse to the location of the ISO image file.
- g **Summary/Completing the New Virtual Machine Wizard:** Verify the description, and then click **Finish**.
- 7 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Settings**.
- 8 In the **Hardware** area on the left, configure as follows:
 - a Ensure that the virtual hard disk that you created is attached to an IDE controller.
A Hyper-V guest machine can only boot from an IDE drive.
 - b Choose **Processor**.
 - c In **Number of virtual processors**, enter 8, and then click **Apply**.
- 9 Create additional virtual hard disks as follows:
 - a In the **Hardware** area on the left, choose any controller and location that will accommodate the disk.
 - b In the controller area on the right, choose **Hard Drive**, and then click **Add**.
 - c In the area on the right, 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 50.
 - c Click **Next**.
 - e In the **Settings** window, click **Apply**.
 - f Repeat this step to create two 16GB virtual hard disks (for a total of four disks).
When viewed from the Linux system, the disks might be reordered; however, the installation is not affected by disk order.
 - g In the **Settings** dialog box, click **Apply**, and then click **OK**.
- 10 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Start**.

Figure 9: Starting a virtual machine



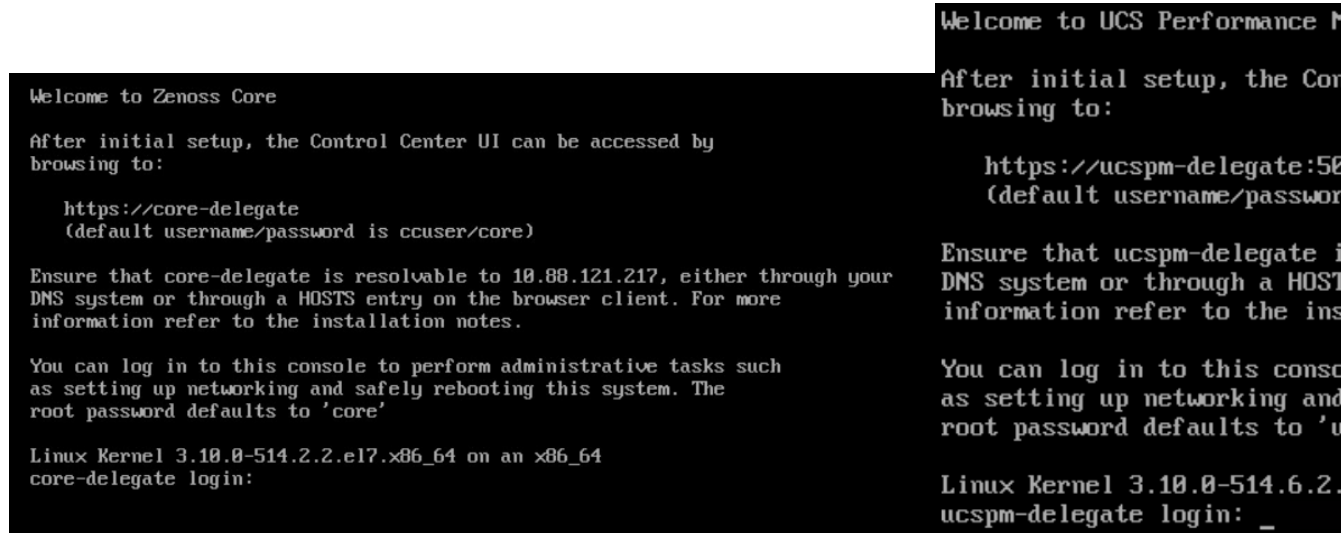
- 11 In the Hyper-V Manager **Virtual Machines** area, right-click the new virtual machine, and then choose **Connect**.
- 12 In the **Virtual Machine Connection** window, press **Enter**.
The appliance installation process takes about 15 minutes, and should complete with no additional input.
If received, 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 10: Initial hypervisor console login prompt



- 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.

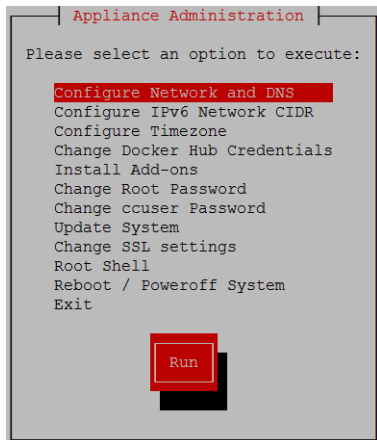
- a Press **Tab** to select **Ok**, and then press **Enter**.
The system reboots.

Edit 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.

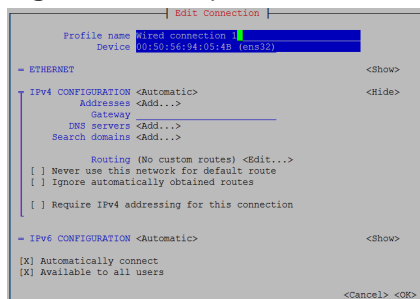
Figure 11: Example: Available connections



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

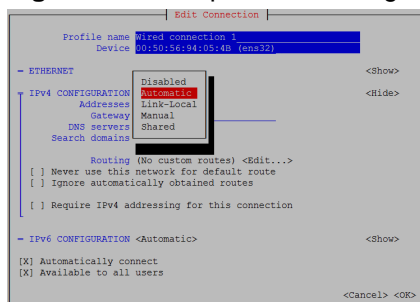
- 5 Select the virtual connection, and then press **Enter**.

Figure 12: 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 13: 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` as follows:
 - 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 14: 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**.

Configuring a multi-host Control Center cluster

4

This chapter describes how to configure an appliance-based multi-host Control Center cluster. A multi-host Zenoss Core deployment 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.

Enabling NTP on Microsoft Hyper-V guests

Note The procedures in this section 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.

Like most distributed applications, Control Center requires a common time source. The procedures in this section enable *NTP* to synchronize the system clocks of all hosts in your Control Center cluster.

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

- 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

(Hyper-V only) This procedure uses the default configuration of NTP to synchronize system clocks with public time servers. If all hosts in the Control Center cluster can access the internet, repeat this procedure on each host in the cluster, 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` as follows:
 - 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` as follows:

- 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 as follows:

- a Open `/etc/ntp.conf` with a text editor.

- b Replace all 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 the hosts in a Control Center cluster:

```
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, the 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 delegate in your Control Center cluster.

- 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` as follows:
 - 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 as follows:

- a Open `/etc/ntp.conf` with a text editor.
- b Replace all 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
```

Adding hosts to the default resource pool

This procedure adds one or more hosts to the default resource pool.

- 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` as follows:
 - 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 and register a delegate host.

Replace *Hostname-Or-IP* with the hostname or IP address of the delegate host to add and include the register flag to authenticate the delegate host:

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

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

- 4 Repeat the preceding command for each delegate in your Control Center cluster.

5

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 of the hosts in a Control Center cluster 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 cluster 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 cluster 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 cluster 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 cluster hosts
<i>SERVICED_ZK_SESSION_TIMEOUT</i>	ZooKeeper ensemble nodes

Example multi-host ZooKeeper configuration

This example shows the ZooKeeper variables in the `/etc/defaults/serviced` configuration file of each host in a 4-node Control Center cluster. 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 in a Control Center cluster 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
```