Zenoss Core Installation Guide

Release 5.2.3

Zenoss, Inc.

www.zenoss.com

Zenoss Core Installation Guide

Contents

About this guide

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

Related Zenoss Core publications

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

1051.17.100 (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 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	1.2.0, 1.2.1, 1.2.2, 1.2.3, 1.3.0, 1.3.1	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)

Installing Zenoss Core

1

This chapter describes how to install Zenoss Core on a Control Center cluster. For more information about creating a Control Center cluster, refer to the following documents:

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

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-5.2.3
```

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*.

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.

Installing a Control Center master host

2

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

The procedures in this chapter configure a Control Center master host that functions as both master and delegate. Perform the procedures in this chapter whether you are configuring a single-host or a multi-host deployment. (For more information about configuring a multi-host deployment, see *Configuring a multi-host cluster*.)

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 master OVA packages as a virtual machine managed by vSphere Server version 5.5.0, using VMware vSphere Client 5.5.0. The procedure is slightly different with different 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.

Figure 1: vSphere client Home view
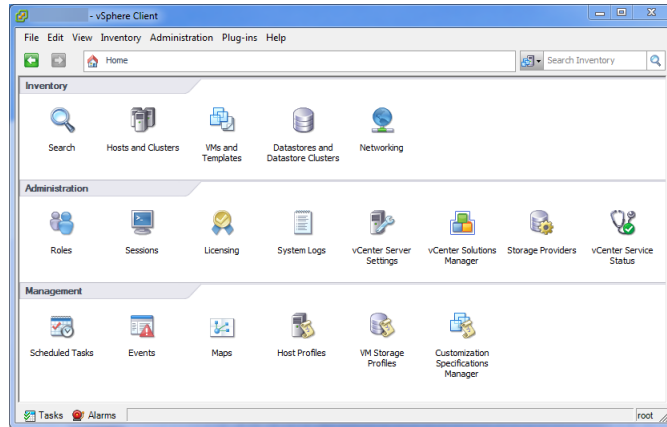
3 From the **File** menu, select **Deploy OVF Template**.

4 In the **Source** panel, specify the path of the Zenoss Core master package, and then click **Next**.

5 In the **OVF Template Details** panel, click **Next**.

6 In the **Name and Location** panel, provide a name and a location for the server.

    a In the **Name** field, enter a new name or use the default.

    b In the **Inventory Location** area, select a data center for the virtual machine.

    c Click **Next**.

7 In the **Host / Cluster** panel, select a host system, and then click **Next**.

8 In the **Storage** panel, select a storage system with sufficient space for your deployment, and then click **Next**.

9 In the **Disk Format** panel, select **Thin Provison**, and then click **Next**.

10 In the **Ready to Complete** panel, review the deployment settings, and then click **Finish**.

Do not check **Power on after deployment**.

11 Navigate to the new virtual machine's **Getting Started** or **Summary** tab, and then click **Edit virtual machine settings**.

12 Optional: Update the memory assigned to the machine.

    a In the **Virtual Machine Properties** dialog box, select **Memory** in the **Hardware** table.

    b In the **Memory Configuration** area, set the **Memory Size** field to 16GB (multi-host deployments) or 32GB (single-host deployments).

       For single-host deployments, you can assign a greater amount of RAM.

13 Optional: Update the number of CPU sockets assigned to the machine, if desired.

Only 4 CPUs are needed for multi-host deployments.

    a In the **Virtual Machine Properties** dialog box, select **CPUs** in the **Hardware** table.

    b Set the **Number of virtual sockets** field to 4 (multi-host deployments), and set the **Number of cores per socket** field to 1.

14 At the bottom of the **Virtual Machine Properties** dialog box, click **OK**.

15 On the new virtual machine's **Getting Started** tab, click **Power on virtual machine**, and then click the **Console** tab.

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 From the **Actions** pane, choose **New** > **Virtual Machine**.
6 In the **Specify Name and Location** panel, provide a name for the virtual machine, and then click **Next**.
7 In the **Specify Generation** panel, select **Generation 1**, and then click **Next**.
8 In the **Assign Memory** panel, specify memory information as follows:

   a In **Startup memory**:

   - For multi-host deployments, enter 16384 (16GB).
   - For single-host deployments, enter 32768 (32GB).

   b Check **Use Dynamic Memory for this virtual machine**.
   c Click **Next**.
9 In the **Configure Networking** panel's **Connection** field, choose a virtual switch, and then click **Next**.
10 In the **Connect Virtual Hard Disk** panel, create a new disk on which to install the guest operating system, as follows:

   a Choose **Create a virtual hard disk**.
   b Specify a name.
   c In the **Size** field, enter 30.
   d Click **Next**.

   In a later step, you will create additional disks.
11 In the **Installation Options** panel, access the ISO file as follows:

   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 ISO image file.

      Choose the master ISO file.
   d Click **Next**.
12 In the **Summary** panel, review the description, and then click **Finish**.

   Hyper-V Manager creates the new virtual machine, and then closes the **New Virtual Machine Wizard**.
13 In the Hyper-V Manager **Virtual Machines** area, choose the new virtual machine, right-click, and then choose **Settings**.
14 In the **Hardware** area of the **Settings** window, 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 the following value for your deployment type:

   - Multi-host deployment: 4
   - Single-host deployment: 8

   d Click **Apply**.
15 Create additional virtual hard disks as follows:

      a   In the **Hardware** area of the **Settings** window, 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 **Hard Drive** area, choose **Virtual hard disk**, and then click **New**.

      d   In the **New Virtual Hard Disk Wizard**:

           1   Navigate to the **Choose Disk Format** panel.

           2   Choose **VHDX**.

           3   Click **Next**.

      e   In the **Choose Disk Type** panel, choose **Dynamically expanding**, and then click **Next**.

      f   In the **Specify Name and Location** panel, enter the disk name, and then click **Next**.

      g   In the **Configure Disk** panel:

           1   Choose **Create a new blank virtual hard disk**.

           2   In the **Size** field, enter 50.

           3   Click **Next**.

      h   In the Summary panel, verify the description, and then click **Finish**.

      i   In the **Settings** window, click **Apply**.

      j   Repeat these substeps to create virtual hard disks sized as follows:

- 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.

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

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

Figure 2: Starting a virtual machine
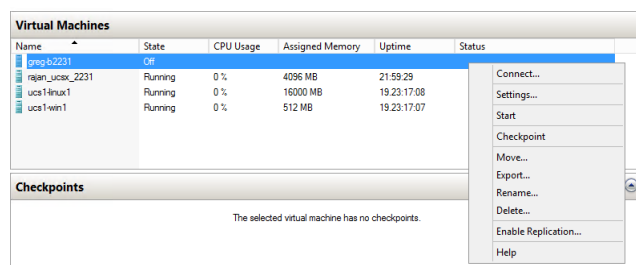
17 In the Hyper-V Manager **Virtual Machines** area, choose the new virtual machine, and then right-click, and then choose **Connect**.

18 In the **Virtual Machine Connection** window, press **Enter**.

The appliance installation process takes about 15 minutes, and should complete with no additional input.

Disregard the Fast TSC calibration failure message.

Configuring the Control Center host

Perform this procedure immediately after creating and starting a Control Center host. All Control Center deployments must include one system 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

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 or press **Tab** or **Shift+Tab**.
- 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*.

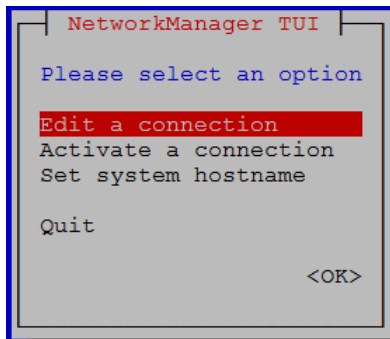**2** Log in as the `root` user.

```
┤ Appliance Administration ├

Please select an option to execute:

    Configure Network and DNS
    Configure IPv6 Network CIDR
    Configure Timezone
    Change Docker Hub Credentials
    Install Add-ons
    Change Root Password
    Change ccuser Password
    Update System
    Change SSL settings
    Root Shell
    Reboot / Poweroff System
    Exit

              Run
```
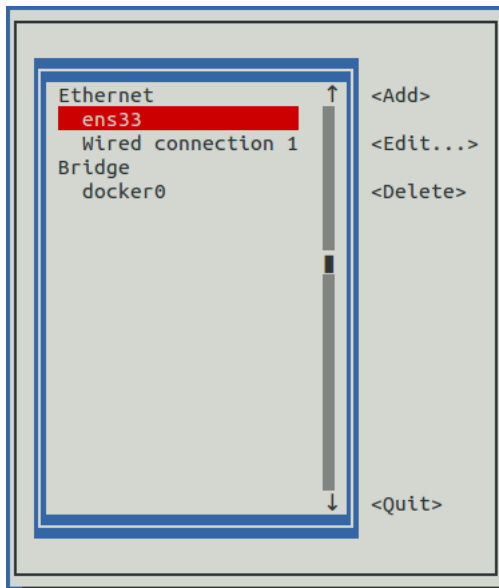
3   Select the **NetworkManager TUI** menu as follows:

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

```
┌┤ NetworkManager TUI ├┐

  Please select an option

  Edit a connection
  Activate a connection
  Set system hostname

  Quit

                <OK>
```

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

Figure 4: Example: Available connections

```
  Ethernet                    ↑     <Add>
    ens33
    Wired connection 1              <Edit...>
  Bridge
    docker0                         <Delete>




                              ↓     <Quit>
```

Note   Do not modify the `docker0` connection.

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

Figure 5: Example: Edit Connection screen

```
                   ┤ Edit Connection ├
      Profile name Wired connection 1
           Device 00:50:56:94:05:4B (ens32)

  ═ ETHERNET                                    <Show>

  ┬ IPv4 CONFIGURATION <Automatic>              <Hide>
         Addresses <Add...>
           Gateway
       DNS servers <Add...>
     Search domains <Add...>

            Routing (No custom routes) <Edit...>
  [ ] Never use this network for default route
  [ ] Ignore automatically obtained routes

  [ ] Require IPv4 addressing for this connection

  ═ IPv6 CONFIGURATION <Automatic>              <Show>

  [X] Automatically connect
  [X] Available to all users

                              <Cancel> <OK>
```

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   Select **<Add>** next to **Addresses**, 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 Select **<OK>** at the bottom of the **Edit Connection** screen, and then press **Enter**.

10 On the **NetworkManager TUI** screen, use the **Tab** key to select the **<Quit>** option, and then press the **Return** key to go back to the **Appliance Administration** menu.
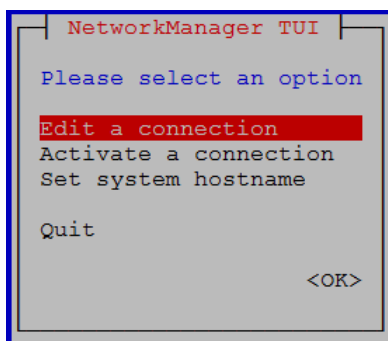
11 Reboot the operating system as follows:

    a   In the **Appliance Administration** menu, select **Reboot / Poweroff System**.

    b   Select **Reboot**.

    c   Select **OK**, and then press **Enter**.

Set system hostname

The default hostname is `core-master` for the Zenoss Core master host and is `core-delegate` for the Zenoss Core delegate host. 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   Press **Tab** to 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 Reboot the operating system as follows:

    a In the **Appliance Administration** menu, select **Reboot / Poweroff System**.

    b Select **Reboot**.

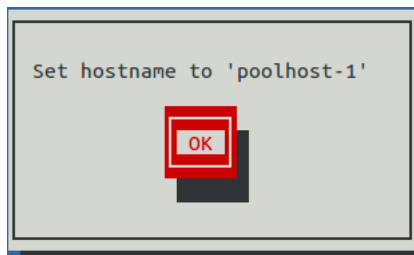   **c** Select **OK**, and then press **Enter**.

Adding the master host to a resource pool

This procedure adds 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 with at least two delegate hosts, perform this step.

```
serviced pool add master
```

4 Add the master host to a resource pool.

- For single-host deployments, add the master host to the `default` resource pool.
- For multi-host deployments with at least two delegate hosts, add the master host to the `master` resource pool.

Replace *Hostname-Or-IP* with the hostname or IP address of the Control Center master host, and replace *Resource-Pool* with `default` or `master`:

```
serviced host add Hostname-Or-IP:4979 Resource-Pool
```

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.

Deploying Zenoss Core

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.

   Replace *Template-ID* with the template identifier returned in the previous step, and 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.

If you are installing a single-host deployment, proceed to the *Zenoss Core Configuration Guide*.

zenoss

Installing delegate hosts

3

This chapter describes how to install appliance-based delegates. You can add as many delegates as you need to a Control Center cluster.

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.

This procedure installs the Zenoss Core delegate OVA package as a virtual machine managed by vSphere Server version 5.5.0, using VMware vSphere Client 5.5.0. The procedure is slightly different with different versions of VMware vSphere Client.

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.

Figure 7: vSphere client Home view
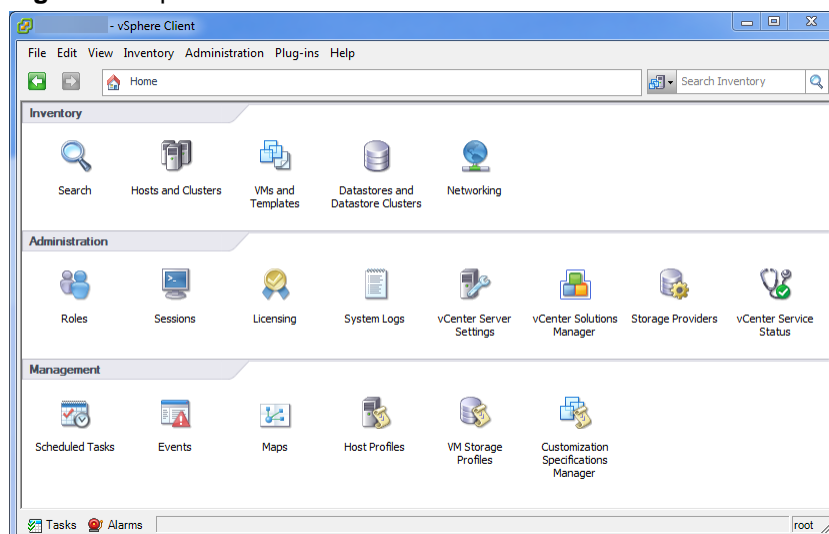
3   From the **File** menu, select **Deploy OVF Template**.

4   In the **Source** panel, specify the path of the Zenoss Core package, and then click **Next**.

5   In the **OVF Template Details** panel, click **Next**.

6   In the **Name and Location** panel, provide a name and a location for the server. as follows, and then click **Next**.

   a   In the **Name** field, enter a new name.

   b   In the **Inventory Location** area, select a data center for the virtual machine.

7   In the **Host / Cluster** panel, select a host system, and then click **Next**.

8   In the **Storage** panel, select a storage system with sufficient space for the virtual machine, and then click **Next**.

9   In the **Disk Format** panel, select **Thin Provison**, and then click **Next**.

10 In the **Ready to Complete** panel, review the deployment settings, and then click **Finish**.

   Do not check **Power on after deployment**.

11 Navigate to the new virtual machine's **Getting Started** or **Summary** tab, and then click **Edit virtual machine settings**.

12 Update the memory assigned to the machine as follows:

   a   In the **Virtual Machine Properties** dialog box, select **Memory** in the **Hardware** table.

   b   In the **Memory Configuration** area, set the **Memory Size** field to 32GB.

   c   At the bottom of the **Virtual Machine Properties** dialog box, click **OK**.

13 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   From the **Actions** pane, choose **New** > **Virtual Machine**.

6   In the **Specify Name and Location** panel, provide a name for the virtual machine, and then click **Next**.

7   In the **Specify Generation** panel, select **Generation 1**, and then click **Next**.

8   In the **Assign Memory** panel, specify memory information as follows:

   a   In **Startup memory**, enter `32768` (32GB).

   b   Check **Use Dynamic Memory for this virtual machine**.

   c   Click **Next**.

9   In the **Configure Networking** panel's **Connection** field, choose a virtual switch, and then click **Next**.

10 In the **Connect Virtual Hard Disk** panel, create a new disk on which to install the guest operating system, as follows:

   a   Choose **Create a virtual hard disk**.

b Specify a name.

c In the **Size** field, enter 30.

In a later step, you will create additional disks.

11 In the **Installation Options** panel, access the ISO file as follows:

    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 ISO image file.

        Choose the delegate ISO file.

    d Click **Next**.

12 In the **Summary** panel, review the description, and then click **Finish**.

Hyper-V Manager creates the new virtual machine, and then closes the **New Virtual Machine Wizard**.

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

14 In the **Hardware** area of the **Settings** window, 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**.

15 Create additional virtual hard disks as follows:

    a In the **Hardware** area of the **Settings** window, choose any controller and location that will accommodate the disk.

    b In the **Hard Drive** area, choose **Virtual hard disk**, and then click **New**.

    c In the **New Virtual Hard Disk Wizard**:

        1 Navigate to the **Choose Disk Format** panel.

        2 Choose **VHDX**.

        3 Click **Next**.

    d In the **Choose Disk Type** panel, choose **Dynamically expanding**, and then click **Next**.

    e In the **Specify Name and Location** panel, enter the disk name, and then click **Next**.

    f In the **Configure Disk** panel:

        1 Choose **Create a new blank virtual hard disk**.

        2 In the **Size** field, enter 50.

        3 Click **Next**.

    g In the **Settings** window, click **Apply**.

    h Repeat these substeps to create two more 16GB virtual hard disks.

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

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

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

Figure 8: Starting a virtual machine
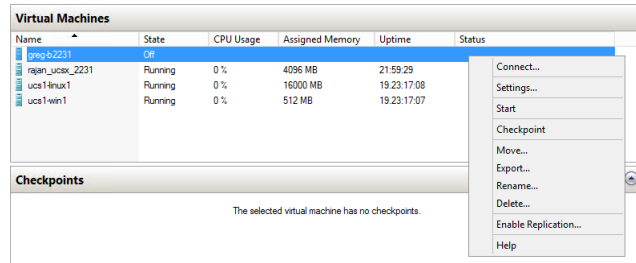
17 In the Hyper-V Manager **Virtual Machines** area, choose the new virtual machine, and then right-click, and then choose **Connect**.

18 In the **Virtual Machine Connection** window, press **Enter**.

The appliance installation process takes about 15 minutes, and should complete with no additional input.

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

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.

```
┤ Master system IP ├
Enter the IP of the master for this delegate

        IP █
            ┌─────┐          ┌─────────┐
            │  Ok │          │  Cancel │
            └─────┘          └─────────┘
```
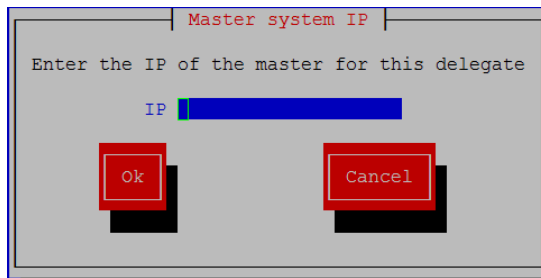
Note   If you enter the hostname or fully qualified domain name of the master host, you need an entry in the /etc/hosts file of the delegate host, or a nameserver on your network, that resolves the name to its IPv4 address.

---

a Press **Tab** to select **Ok**, and then press **Enter**.

The system reboots.

Edit a connection

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 or press **Tab** or **Shift+Tab**.
- 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.

```
┤ Appliance Administration ├
Please select an option to execute:

   Configure Network and DNS
   Configure IPv6 Network CIDR
   Configure Timezone
   Change Docker Hub Credentials
   Install Add-ons
   Change Root Password
   Change ccuser Password
   Update System
   Change SSL settings
   Root Shell
   Reboot / Poweroff System
   Exit

            ┌─────┐
            │ Run │
            └─────┘
```

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 this host.

Figure 10: Example: Available connections
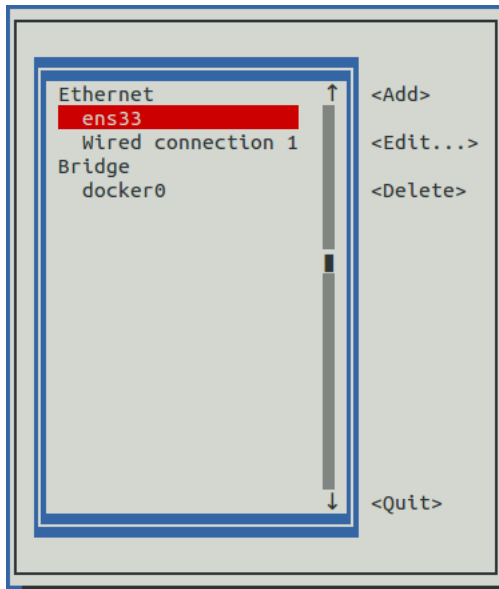
Note   Do not modify the `docker0` connection.

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

Figure 11: Example: Edit Connection screen
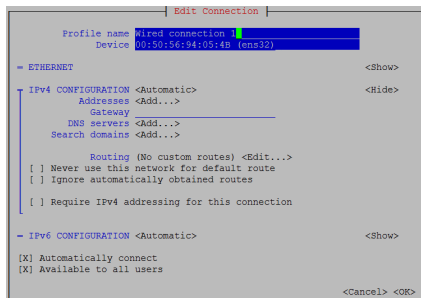
6   Optional: If the **IPv4 CONFIGURATION** area is not visible, select its display option (**<Show>**), and then press **Enter**.

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   Select **<Add>** next to **Addresses**, 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   Select **<OK>** at the bottom of the **Edit Connection** screen, and then press **Enter**.

10   On the **NetworkManager TUI** screen, use the **Tab** key to select the **<Quit>** option, and then press the **Return** key to go back to the **Appliance Administration** menu.

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**.

Set system hostname

The default hostname is `core-master` for the Zenoss Core master host and is `core-delegate` for the Zenoss Core delegate host. 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*.
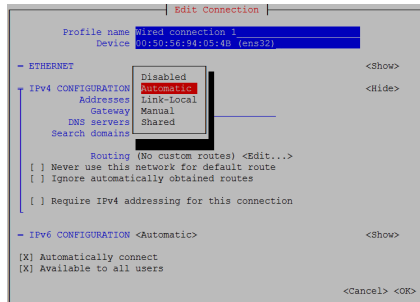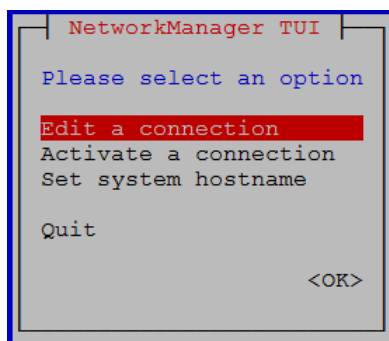
2   Select the **NetworkManager TUI** menu as follows:

    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   Press **Tab** to 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**.

Set hostname to 'poolhost-1'

OK

6 In the confirmation dialog box, press **Enter**.
7 Reboot the operating system as follows:
   a In the **Appliance Administration** menu, select **Reboot / Poweroff System**.
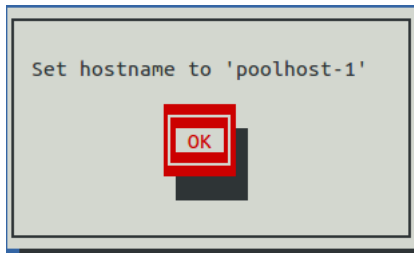   b Select **Reboot**.
   c Select **OK**, and then press **Enter**.

Editing the /etc/hosts file

This procedure is optional. 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 13: Example nano session
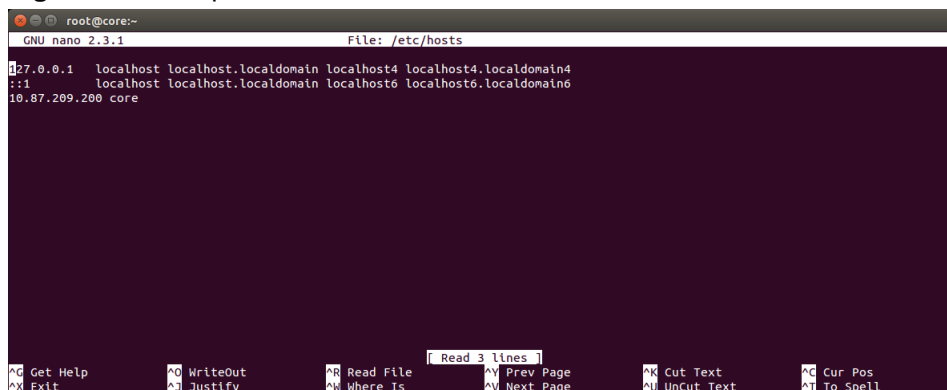
Use the up-arrow and down-arrow keys to select lines, and the right-arrow and left-arrow keys to select characters on a line.

    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 Save the file and exit the editor.

       To save, press **Control-o**. 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

<div align="right">4</div>

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.

Task	Procedure
Install a master host	*Installing a Control Center master host* on page 10
Install delegate hosts	*Installing delegate hosts* on page 19
Add delegate hosts to the cluster	*Adding hosts to the default resource pool* on page 32
Create a ZooKeeper ensemble	*Configuring a ZooKeeper ensemble* on page 33
Create a private NTP server (Hyper-V only)	*Enabling NTP on Microsoft Hyper-V guests* on page 28
Add delegate hosts to the cluster	*Adding hosts to the default resource pool* on page 32

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.

   ```
   ntpd -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.

```
ntpd -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.

Configuring a ZooKeeper ensemble

<div style="text-align: right; font-size: 3em; font-weight: bold;">5</div>

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	*Master*	The IP address or hostname of the master host.
node2	*Delegate-A*	The IP address or hostname of delegate host A.

Variable name	Placeholder value	Actual value
node3	*Delegate-B*	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
SERVICED_ISVCS_START	ZooKeeper ensemble nodes
SERVICED_ISVCS_ZOOKEEPER_ID	ZooKeeper ensemble nodes
SERVICED_ISVCS_ZOOKEEPER_QUORUM	ZooKeeper ensemble nodes
SERVICED_ZK	All Control Center cluster hosts
SERVICED_ZK_SESSION_TIMEOUT	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
```

   For high-availability deployments, use the virtual IP address or hostname of the high-availability cluster as the value for `node1`.
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 Clean up the Control Center configuration file.

   a Open `/etc/default/serviced` in a text editor.
   b Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK* variable declaration at that location.
   c Locate the original *SERVICED_ZK* variable declaration, and then paste the cut line immediately below it.
   d Navigate to the end of the file, and cut the line that contains the *SERVICED_ISVCS_ZOOKEEPER_ID* variable declaration at that location.

> **e** Locate the original *SERVICED_ISVCS_ZOOKEEPER_ID* variable declaration, and then paste the cut line immediately below it.
>
> **f** Navigate to the end of the file, and cut the line that contains the *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable declaration at that location.
>
> **g** Locate the original *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable declaration, and then paste the cut line immediately below it.
>
> **h** Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK_SESSION_TIMEOUT* variable declaration at that location.
>
> **i** Locate the original *SERVICED_ZK_SESSION_TIMEOUT* variable declaration, and then paste the cut line immediately below it.
>
> **j** Save the file, and then close the editor.

8 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
```

For high-availability deployments, use the virtual IP address or hostname of the high-availability cluster as the value for `node1`.

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

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=2" >> /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 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
```

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   Set the *SERVICED_ISVCS_START* variable, and clean up the Control Center configuration file.

   a   Open /etc/default/serviced in a text editor.
   b   Locate the line for the *SERVICED_ISVCS_START* variable, and then make a copy of the line, immediately below the original.
   c   Remove the number sign character (#) from the beginning of the line.
   d   Delete all but zookeeper from its list of values.
   e   Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK* variable declaration at that location.
   f   Locate the original *SERVICED_ZK* variable declaration, and then paste the cut line immediately below it.
   g   Comment the original *SERVICED_ZK* declaration, which references only the master host.

         Insert the number sign character (#) immediately in front of the original *SERVICED_ZK* variable.

   h   Navigate to the end of the file, and cut the line that contains the *SERVICED_ISVCS_ZOOKEEPER_ID* variable declaration at that location.
   i   Locate the original *SERVICED_ISVCS_ZOOKEEPER_ID* variable declaration, and then paste the cut line immediately below it.
   j   Navigate to the end of the file, and cut the line that contains the *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable declaration at that location.
   k   Locate the original *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable declaration, and then paste the cut line immediately below it.
   l   Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK_SESSION_TIMEOUT* variable declaration at that location.
   m   Locate the original *SERVICED_ZK_SESSION_TIMEOUT* variable declaration, and then paste the cut line immediately below it.
   n   Save the file, and then close the editor.

8   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
```

For high-availability deployments, use the virtual IP address or hostname of the high-availability cluster as the value for `node1`.

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

```
echo "SERVICED_ISVCS_ZOOKEEPER_ID=3" >> /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 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
```

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 Set the *SERVICED_ISVCS_START* variable, and clean up the Control Center configuration file.

a Open `/etc/default/serviced` in a text editor.

b Locate the line for the *SERVICED_ISVCS_START* variable, and then make a copy of the line, immediately below the original.

c Remove the number sign character (#) from the beginning of the line.

d Delete all but `zookeeper` from its list of values.

e Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK* variable declaration at that location.

f Locate the original *SERVICED_ZK* variable declaration, and then paste the cut line immediately below it.

g Comment the original *SERVICED_ZK* declaration, which references only the master host.

Insert the number sign character (#) immediately in front of the original *SERVICED_ZK* variable.

h Navigate to the end of the file, and cut the line that contains the *SERVICED_ISVCS_ZOOKEEPER_ID* variable declaration at that location.

     i   Locate the original *SERVICED_ISVCS_ZOOKEEPER_ID* variable declaration, and then paste the cut line immediately below it.

     j   Navigate to the end of the file, and cut the line that contains the *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable declaration at that location.

     k   Locate the original *SERVICED_ISVCS_ZOOKEEPER_QUORUM* variable declaration, and then paste the cut line immediately below it.

     l   Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK_SESSION_TIMEOUT* variable declaration at that location.

     m Locate the original *SERVICED_ZK_SESSION_TIMEOUT* variable declaration, and then paste the cut line immediately below it.

     n   Save the file, and then close the editor.

8 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
```

Installing the ZooKeeper image for Docker Engine

Use one of the procedures in this section to pull or import the Docker Engine image for the ZooKeeper service on each ZooKeeper ensemble node.

Pulling the ZooKeeper image for Docker Engine

To perform this procedure, ZooKeeper ensemble nodes must have internet access.

Use this procedure to pull the Docker Engine image for the ZooKeeper service from Docker Hub to a ZooKeeper ensemble node. Repeat this procedure on each node in the ZooKeeper ensemble.

1 Log in to the ZooKeeper ensemble node as root, or as a user with superuser privileges.

2 Identify the images associated with the installed version of serviced.

```
serviced version | grep Images
```

Example result:

```
IsvcsImages: [zenoss/serviced-isvcs:v54 zenoss/isvcs-zookeeper:v8]
```

In this example, the tag of the ZooKeeper image is isvcs-zookeeper:v8.

3 Pull the required image from Docker Hub to the local repository.

Replace *ZooKeeper-Tag* with the tag displayed in the previous step:

```
docker pull zenoss/ZooKeeper-Tag
```

Importing the ZooKeeper image for Docker Engine

Use this procedure to import the Docker Engine image for the ZooKeeper service from a ZooKeeper image file. Repeat this procedure on each node in the ZooKeeper ensemble.

1   Log in to the ZooKeeper ensemble node as `root`, or as a user with superuser privileges.
2   Change directory to `/root`.

```
cd /root
```

3   Extract the ZooKeeper image.

```
./install-zenoss-isvcs-zookeeper:v*.run
```

Image extraction begins when you press the **y** key.
4   Optional: Delete the archive file, if desired.

```
rm -i ./install-zenoss-isvcs-zookeeper:v*.run
```

Starting a ZooKeeper ensemble

Use this procedure to start a ZooKeeper ensemble.

This procedures uses the `nc` utility to query ensemble hosts. If `nc` is not available, you can use `telnet` with *interactive ZooKeeper commands*.

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

1   Log in to the Control Center master host as `root`, or as a user with superuser privileges.
2   In a separate window, log in to the second node of the ZooKeeper ensemble (*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 and start `serviced`.

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

5   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
```

For high-availability deployments, use the virtual IP address or hostname of the high-availability cluster as the value for `node1`.

3   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
```

4   Update the variable.

   a   Open `/etc/default/serviced` in a text editor.
   b   Navigate to the end of the file, and cut the line that contains the *SERVICED_ZK* variable declaration at that location.

     The value of this declaration specifies three endpoints.

   c   Locate the *SERVICED_ZK* variable near the beginning of the file, and then delete the line it is on.

     The value is just the master host endpoint.

   d   Paste the *SERVICED_ZK* variable declaration from the end of the file in the location of the just-deleted declaration.
   e   Save the file, and then close the editor.

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
```