

Oracle Linux Virtualization Manager

Getting Started



F52194-11
February 2024



Oracle Linux Virtualization Manager Getting Started,
F52194-11

Copyright © 2022, 2024, Oracle and/or its affiliates.

Contents

1 About the Docs

Documentation License	1-2
Conventions	1-2
Documentation Accessibility	1-2
Access to Oracle Support for Accessibility	1-2
Diversity and Inclusion	1-2

2 Requirements and Scalability Limits

3 Installation and Configuration

Installing the Engine	3-1
Configuring the Engine	3-3
Engine Configuration Options	3-4
OVN Provider	3-5
WebSocket Proxy	3-5
Data Warehouse	3-5
Keycloak	3-5
VM Console Proxy	3-6
Grafana	3-6
Manager DNS Name	3-6
Automatic Firewall Configuration	3-6
Data Warehouse Database	3-6
Engine Database	3-7
Admin User Password	3-8
Application Mode	3-8
OVN Provider Credentials	3-8
SAN Wipe After Delete	3-9
Web Server Configuration	3-9
Data Warehouse Sampling Scale	3-9
Logging in to the Administration Portal	3-10
Preparing to Log in	3-10

Logging in	3-11
Next Steps	3-11
Logging Out	3-11
Configuring a KVM Host	3-11
Preparing a KVM Host	3-12
Adding a KVM Host	3-14

4 Self-Hosted Engine Deployment

Self-Hosted Engine Prerequisites	4-1
Deploying the Self-Hosted Engine	4-2
Using the Command Line to Deploy	4-4
Using the Cockpit Portal to Deploy	4-9
Enabling High-Availability	4-11
Configuring a Highly Available Host	4-11
Configuring Power Management and Fencing on a Host	4-12
Preventing Host Fencing During Boot	4-14
Checking Fencing Parameters	4-14
Installing Additional Self-Hosted Engine Hosts	4-14
Cleaning up the Deployment	4-15
Upgrading Or Updating the Self-Hosted Engine	4-15

5 Deploying GlusterFS Storage

Deploying GlusterFS Storage Using Cockpit	5-1
Creating a GlusterFS Storage Domain Using the Manager	5-3

1

About the Docs

Oracle Linux Virtualization Manager Release 4.5 is based on [oVirt](#), which is a free, open-source virtualization solution. The product documentation comprises:

- **Release Notes** - A summary of the new features, changes, fixed bugs, and known issues in the Oracle Linux Virtualization Manager. It contains last-minute information, which might not be included in the main body of documentation.
- **Architecture and Planning Guide** - An architectural overview of Oracle Linux Virtualization Manager, prerequisites, and planning information for your environment.
- **Getting Started Guide** - How to install, configure, and get started with the Oracle Linux Virtualization Manager using standard or self-hosted configuration. It also provides information for configuring KVM hosts and deploying GlusterFS storage.
- **Administration Guide** - Provides common administrative tasks for Oracle Linux Virtualization Manager such as:
 - setting up users and groups
 - creating data centers, clusters, and virtual machines
 - using virtual machine templates and snapshots
 - migrating virtual machines
 - configuring logical and virtual networks
 - using local, NFS, iSCSI and FC storage
 - backing up and restoring
 - configuring high-availability, vCPUs, and virtual memory
 - monitoring with event notifications and Grafana dashboards
 - upgrading and updating your environment
 - active-active and active-passive disaster recovery solutions

You can also refer to:

- REST API Guide, which you can access from the Welcome Dashboard or directly through its URL <https://manager-fqdn/ovirt-engine/apidoc>.
- Upstream [oVirt Documentation](#).

If you want to provide feedback about this documentation, please complete the [Oracle Help Center feedback form](#).

To access Oracle Linux Virtualization Manager Release 4.4 documentation, PDFs are available at:

- [Release Notes](#)
- [Getting Started Guide](#)
- [Architecture and Planning Guide](#)
- [Administration Guide](#)

Documentation License

The content in this document is licensed under the [Creative Commons Attribution–Share Alike 4.0](#) (CC-BY-SA) license. In accordance with CC-BY-SA, if you distribute this content or an adaptation of it, you must provide attribution to Oracle and retain the original copyright notices.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the [Oracle Accessibility Program website](#).

For information about the accessibility of the Oracle Help Center, see the [Oracle Accessibility Conformance Report](#).

Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through [Oracle Accessibility Learning and Support](#).

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

2

Requirements and Scalability Limits

Before you begin the tasks in this guide, you should review Oracle Linux Virtualization Manager Release 4.5 concepts, environment requirements, and scalability limitations in the [Oracle Linux Virtualization Manager: Architecture and Planning Guide](#).

3

Installation and Configuration

To deploy Oracle Linux Virtualization Manager, you install and configure the engine on a host with Oracle Linux 8.8 (or later Oracle Linux 8 release), configure KVM hosts, storage, and networks, and create virtual machines. Thoroughly review the [Requirements and Scalability Limits](#) as the requirements for the engine host are different than the KVM hosts.

To review conceptual information and help to plan your installation, see the [Oracle Linux Virtualization Manager: Architecture and Planning Guide](#).

Installing the Engine

To install Oracle Linux Virtualization Manager, you perform a fresh installation of Oracle Linux 8.8 (or later) on the host, install the `ovirt-engine` package, and then run the `engine-setup` command to configure the Manager.



Note:

You can install the Manager in a virtual machine as long as it is not managing that virtual machine, or in a self-hosted engine configuration. For more information, see [Self-Hosted Engine Deployment](#). **Do not configure the same host as a standalone engine and a KVM host.**

You can download the installation ISO for Oracle Linux 8.8 (or later Oracle Linux 8 release) from the Oracle Software Delivery Cloud at <https://edelivery.oracle.com>.

1. Install Oracle Linux 8.8 (or later Oracle Linux 8 release) on the host using the **Minimal Install** base environment.

Follow the instructions in the [Oracle® Linux 8: Installing Oracle Linux](#).



Important:

Do not install any additional packages until after you have installed the Manager packages, because they may cause dependency issues.

2. **(Optional)** If you use a proxy server for Internet access, configure Yum with the proxy server settings. For more information, see the [Oracle® Linux: Managing Software on Oracle Linux](#).
3. Complete one of the following sets of steps:
 - **For ULN registered hosts or using Oracle Linux Manager**
Subscribe the system to the required channels and enable appstream modules.

- a. For ULN registered hosts, log in to <https://linux.oracle.com> with your ULN user name and password. For Oracle Linux Manager registered hosts, access your internal server URL.
- b. On the Systems tab, click the link named for the host in the list of registered machines.
- c. On the System Details page, click **Manage Subscriptions**.
- d. On the System Summary page, select each required channel from the list of available channels and click the right arrow to move the channel to the list of subscribed channels. Subscribe the system to the following channels:

- ol8_x86_64_baseos_latest
- ol8_x86_64_appstream
- ol8_x86_64_kvm_appstream
- ol8_x86_64_ovirt45
- ol8_x86_64_ovirt45_extras
- ol8_x86_64_gluster_appstream
- **(For VDSM)** ol8_x86_64_UEKR7

- e. Click **Save Subscriptions**.
- f. Install the Oracle Linux Virtualization Manager Release 4.5 package, which automatically enables/disables the required repositories.

```
# dnf install oracle-ovirt-release-45-el8
```

- **For Oracle Linux yum server hosts**

Install the Oracle Linux Virtualization Manager Release 4.5 package and enable the required repositories.

- a. Enable the `ol8_baseos_latest` repository.

```
# dnf config-manager --enable ol8_baseos_latest
```

- b. Install the Oracle Linux Virtualization Manager Release 4.5 package, which automatically enables/disables the required repositories.

```
# dnf install oracle-ovirt-release-45-el8
```

- c. Use the `dnf` command to verify that the required repositories are enabled.

- i. Clear the `dnf` cache.

```
# dnf clean all
```

- ii. List the configured repositories and verify that the required repositories are enabled.

```
# dnf repolist
```

The following repositories must be enabled:

- ol8_baseos_latest
- ol8_appstream
- ol8_kvm_appstream

- ovirt-4.5
- ovirt-4.5-extra
- ol8_gluster_appstream
- **(For VDSM)** ol8_UEKR7

- iii. If a required repository is not enabled, use the `dnf config-manager` command to enable it.

```
# dnf config-manager --enable repository
```

4. If your host is running UEK R7:
 - a. Install the *Extra kernel modules* package.


```
# dnf install kernel-uek-modules-extra
```
 - b. Reboot the host.
5. Install the Manager using the `ovirt-engine` command.

```
# dnf install ovirt-engine
```

Proceed to [Configuring the Engine](#).

Configuring the Engine

After you install the Oracle Linux Virtualization Manager, you run the `engine-setup` command (the Setup program) to configure the Manager. You are prompted to answer a series of questions whose values are used to configure the Manager. Some of these questions relate to features that are in technology preview. For more information, see [Technology Preview in the Oracle Linux Virtualization Manager: Release Notes](#).

The Manager uses two PostgreSQL databases: one for the engine and one for the data warehouse. By default, Setup creates and configures the engine database locally on the engine host. Alternatively, you can configure the engine host to use a manually-configured local or remote database. If you choose to use a manually-configured local or remote database, you must set it up **before** running `engine-setup`. Currently, running the engine or data warehouse database on a remote host is a technology preview feature.

To configure the Manager:

1. Run the `engine-setup` command on the host where you installed the Manager.

```
[ INFO ] Stage: Initializing
[ INFO ] Stage: Environment setup
Configuration files: /etc/ovirt-engine-setup.conf.d/10-packaging-jboss.conf, /etc/
ovirt-engine-setup.conf.d/10-packaging.conf
Log file: /var/log/ovirt-engine/setup/ovirt-engine-setup-YYYYMMDDHHMMSS-snz1rn.log
[ INFO ] Stage: Environment packages setup
[ INFO ] Stage: Programs detection
[ INFO ] Stage: Environment setup (late)
[ INFO ] Stage: Environment customization
```

Note:

Run `engine-setup --accept-defaults` to automatically accept all questions that have default answers.

The Setup program prompts you to configure the Manager.

2. Enter Yes if you want to configure Cinderlib integration, which is currently a Tech Preview feature. The default is No.

```
Configure Cinderlib integration (Currently in tech preview) (Yes, No) [No]:
```

3. Enter Yes to configure the Manager.

```
Configure Engine on this host (Yes, No) [Yes]:
```

If you enter No, the configuration stops. To restart, rerun the engine-setup command.

4. For the remaining configuration questions, provide input or accept default values, which are in square brackets after each question. To accept the default value for a given question, press Enter.

Note:

Setup asks you for the fully-qualified DNS name (FQDN) of the Manager host. Although Setup tries to automatically detect the name, you must ensure the FQDN is correct.

For detailed information on the configuration options, see [Engine Configuration Options](#).

Important:

Keycloak integration is a technology preview feature for internal Single-Sign-On (SSO) provider for the Engine and it deprecates AAA. When you get to this configuration option, the default response is Yes; however, since this is a preview feature, enter No.

5. Once you have answered all the questions, Setup displays a list of the values you entered. Review the list carefully and then press Enter to configure the Manager.

Your answers are saved to a file that can be used to reconfigure the Manager using the same values. Setup also displays the location of the log file for the configuration process.

6. When the configuration is complete, details about how to log in to the Administration Portal are displayed. To verify that the configuration was successful, log into the Administration Portal, as described in [Logging in to the Administration Portal](#).

Engine Configuration Options

The information in this section describes the options for configuring Oracle Linux Virtualization Manager when you run the engine-setup command.

! Important:

Some of the configuration options are in technology preview. For more information, see Technology Preview in the [Oracle Linux Virtualization Manager: Release Notes](#).

OVN Provider

```
Configuring ovirt-provider-ovn also sets the Default cluster's default network
provider to ovirt-provider-ovn.
Non-Default clusters may be configured with an OVN after installation.
Configure ovirt-provider-ovn (Yes, No) [Yes]:
```

Install the Open Virtual Network (OVN) provider on the Manager host and add it as an external network provider. The default cluster is automatically configured to use OVN as its network provider.

OVN is an OVS (Open vSwitch) extension which enables you to configure virtual networks.

Using external providers, including the OVN provider, is a technology preview feature.

WebSocket Proxy

```
Configure WebSocket Proxy on this machine? (Yes, No) [Yes]:
```

The WebSocket Proxy enables you to connect to virtual machines using the noVNC or HTML 5 consoles.

For security and performance reasons, you can configure the WebSocket Proxy on a remote host.

Data Warehouse

```
Please note: Data Warehouse is required for the engine.
If you choose to not configure it on this host, you have to configure
it on a remote host, and then configure the engine on this host so that it can
access the database of the remote Data Warehouse host.
Configure Data Warehouse on this host (Yes, No) [Yes]:
```

The Data Warehouse feature can run on the Manager host or on a remote host. Running Data Warehouse on a remote host reduces the load on the Manager host.

Running the Data Warehouse on a remote host is a technology preview feature.

Keycloak

```
* Please note * : Keycloak is now deprecating AAA/JDBC authentication module.
It is highly recommended to install Keycloak based authentication.
Configure Keycloak on this host (Yes, No) [Yes]:No
```

```
Are you really sure not to install internal Keycloak based authentication?
AAA modules are being deprecated
Configure Keycloak on this host (Yes, No) [Yes]:No
```

Keycloak is a technology preview feature for internal Single-Sign-On (SSO) provider for the Engine thus deprecating AAA. In addition, the Provider OVN and the Grafana Portal are reconfigured to use Keycloak SSO as well.

VM Console Proxy

Configure VM Console Proxy on this host (Yes, No) [Yes]:

The VM Console Proxy enables you to access virtual machine serial consoles from a command line. To use this feature, serial consoles must be enabled in the virtual machines.

Grafana

Use Engine admin password as initial Grafana admin password (Yes, No) [Yes]:

Grafana can be configured to use the Engine password to make signing in easier.

Manager DNS Name

Host fully-qualified DNS name of this server [<autodetected-host-name>]:

The fully-qualified DNS name of the Manager host. Check that the automatically detected DNS name is correct.

Automatic Firewall Configuration

Setup can automatically configure the firewall on this system.
Note: automatic configuration of the firewall may overwrite current settings.
Do you want Setup to configure the firewall? (Yes, No) [Yes]:

The following firewall managers were detected on this system: firewalld
Firewall manager to configure (firewalld): firewalld

Configure the firewall on the host to open the ports used for external communication between Oracle Linux Virtualization Manager and the components it manages.

If Setup configures the firewall, and no firewall managers are active, you are prompted to select a firewall manager from a list.

If you enter No, you must manually configure the firewall. When the Manager configuration is complete, Setup displays a list of ports that need to be opened, see for details.

Data Warehouse Database

Where is the DWH database located? (Local, Remote) [Local]:

The Data Warehouse database (the history database) can run on the Manager host or on a remote host. Running the database on a remote host reduces the load on the Manager host.

Running the database on a remote host is a technology preview feature.

▲ Caution:

In this step you configure the name of the database, and the user name and password for connecting to it. Make a note of these details.

Enter Local to connect to a local PostgreSQL server, or Remote to connect to an existing PostgreSQL server running on a remote host.

If you enter Local, you can choose whether to set up a local PostgreSQL server automatically, or to connect to an existing local PostgreSQL server.

```
Setup can configure the local postgresql server automatically for the DWH to run.
This may conflict with existing applications.
Would you like Setup to automatically configure postgresql and create DWH database,
or prefer to perform that manually? (Automatic, Manual) [Automatic]:
```

Enter Automatic to have Setup configure a local database server, or Manual to connect to an existing local database server. If you enter Manual, you are prompted for the details for connecting to the database:

```
DWH database secured connection (Yes, No) [No]:
DWH database name [ovirt_engine_history]:
DWH database user [ovirt_engine_history]:
DWH database password:
```

If you enter Remote to connect to an existing PostgreSQL server running on a remote host, you are prompted for the details for connecting to the database:

```
DWH database host [localhost]:
DWH database port [5432]:
DWH database secured connection (Yes, No) [No]:
DWH database name [ovirt_engine_history]:
DWH database user [ovirt_engine_history]:
DWH database password:
```

Engine Database

```
Where is the Engine database located? (Local, Remote) [Local]:
```

The Oracle Linux Virtualization Manager database (the engine database) can run on the Manager host or on a remote host. Running the database on a remote host reduces the load on the Manager host.

Running the database on a remote host is a technology preview feature.

▲ Caution:

In this step you configure the name of the database, and the user name and password for connecting to it. Make a note of these details.

Enter Local to connect to a local PostgreSQL server, or Remote to connect to an existing PostgreSQL server running on a remote host.

If you enter **Local**, you can choose whether to set up a local PostgreSQL server automatically, or to connect to an existing local PostgreSQL server.

Setup can configure the local postgresql server automatically for the engine to run.

This may conflict with existing applications.

Would you like Setup to automatically configure postgresql and create Engine database,

or prefer to perform that manually? (Automatic, Manual) [Automatic]:

Enter **Automatic** to have Setup configure a local database server, or **Manual** to connect to an existing local database server. If you enter **Manual**, you are prompted for the details for connecting to the database:

Engine database secured connection (Yes, No) [No]:

Engine database name [engine]:

Engine database user [engine]:

Engine database password:

If you enter **Remote** to connect to an existing PostgreSQL server running on a remote host, you are prompted for the details for connecting to the database:

Engine database host [localhost]:

Engine database port [5432]:

Engine database secured connection (Yes, No) [No]:

Engine database name [engine]:

Engine database user [engine]:

Engine database password:

Admin User Password

Engine admin password:

Confirm engine admin password:

Enter a password for the default administrative user (**admin@internal**). Make a note of the password. If you use a simple password, you might get the following warning:

```
[WARNING] Password is weak: The password fails the dictionary check - it is based on a dictionary word
```

```
Use weak password? (Yes, No) [No]: Yes
```

Application Mode

Application mode (Both, Virt, Gluster) [Both]:

The Manager can be configured to manage virtual machines (**Virt**) or manage Gluster clusters (**Gluster**), or **Both**.

OVN Provider Credentials

Use default credentials (admin@internal) for ovirt-provider-ovn (Yes, No) [Yes]:

oVirt OVN provider user[admin@internal]:

oVirt OVN provider password:

If you installed the OVN provider, configure the credentials for connecting to the OVN (Open vSwitch) databases.

Using external providers, including the OVN provider, is a technology preview feature.

SAN Wipe After Delete

Default SAN wipe after delete (Yes, No) [No]:

Enter Yes to set the default value for the `wipe_after_delete` flag to true, which wipes the blocks of a virtual disk when it is deleted.

Using the wipe after delete functionality is a technology preview feature.

Web Server Configuration

Organization name for certificate [<autodetected-domain-based-name>]:

Provide the organization name to use for the automatically generated self-signed SSL certificate used by the Manager web server.

Setup can configure the default page of the web server to present the application home page. This may conflict with existing applications. Do you wish to set the application as the default web page of the server? (Yes, No) [Yes]:

Enter Yes to make the Oracle Linux Virtualization Manager landing page the default page presented by the web server.

Setup can configure apache to use SSL using a certificate issued from the internal CA. Do you wish Setup to configure that, or prefer to perform that manually? (Automatic, Manual) [Automatic]:

Enter Automatic to generate a self-signed SSL certificate for the web server. Only use self-signed certificates for testing purposes.

Enter Manual to provide the location of the SSL certificate and private key to use the web server.

Note:

For more information, see the following [My Oracle Support](#) articles:

- *How to renew OLVM Hosts Certificate in OLVM Environment/Infrastructure (Doc ID 2885203.1)*
- *VM Migration fails with Error " The server certificate /etc/pki/vdsm/libvirt-vnc/server-cert.pem has expired" (Doc ID 2959537.1)*
- *Moving From Custom 3rd Party CA Certification to Default certification (Doc ID 2963343.1)*

Data Warehouse Sampling Scale

Please choose Data Warehouse sampling scale:

- (1) Basic
 - (2) Full
- (1, 2) [1]:

Set the Data Warehouse sampling scale, either Basic or Full. This step is skipped if the Data Warehouse is not configured to run on the Manager host.

Enter 1 for Basic, which reduces the values of `DWH_TABLES_KEEP_HOURLY` to 720 and `DWH_TABLES_KEEP_DAILY` to 0. Enter 2 for Full.

If the Manager and the Data Warehouse run on the same host, Basic is the recommended sample scale because this reduces the load on the Manager host. Full is recommended only if the Data Warehouse runs on a remote host.

The Full sampling scale is a technology preview feature.

Logging in to the Administration Portal

After you run the engine-setup command to configure Oracle Linux Virtualization Manager, you should log into the Administration Portal to verify that the configuration was successful.

Preparing to Log in

It is recommended that you use the latest version one of the following browsers to access the Administration Portal

- Mozilla Firefox
- Google Chrome
- Microsoft Edge

If Oracle Linux Virtualization Manager was configured to use a self-signed SSL certificate, or an SSL certificate that is signed by a Certificate Authority (CA) that is not trusted by the browser (for example an Intermediate CA), you should install the CA certificate in the browser. Consult your browser's instructions for how to import a CA certificate.

You can download the CA certificate by clicking *Engine CA Certificate* on the Welcome dashboard or by navigating directly to `http://manager-fqdn/ovirt-engine/services/pki-resource?resource=ca-certificate&format=X509-PEM-CA`.

Usually you access the Administration Portal using the fully qualified domain name of the Manager host that you provided during installation. However, you can access the Administration Portal using an alternate host name(s). To do this, you need to add a configuration file to the Manager as follows:

1. Log in to the Manager host as root.
2. Create the file `/etc/ovirt-engine/engine.conf.d/99-custom-ssosetup.conf` with the following content:

```
SSO_ALTERNATE_ENGINE_FQDNS="alias1.example.com alias2.example.com"
```

The list of alternate host names must be separated by spaces.

3. Restart Oracle Linux Virtualization Manager.

```
# systemctl restart ovirt-engine
```

Logging in

You log in to the Administration Portal using a web browser and the default `admin@internal` user.

1. Go to `https://manager-fqdn/ovirt-engine`. The **Welcome** page displays.
2. **(Optional)** Change the preferred language from the drop-down list on the **Welcome** page.
You can view the **Administration Portal** in multiple languages. The default language is based on the locale of your web browser.
3. Click **Administration Portal**. The **Login** page displays.
4. Enter `admin` for the **Username** and the password you specified when you configured the Manager.
5. From the **Profile** list, select `internal` and click **Log In**.

Important:

From the Welcome dashboard, you also have the option of logging into two additional portals:

- The VM Portal
- The Monitoring Portal

For more information, see Access Portals in the [Oracle Linux Virtualization Manager: Architecture and Planning Guide](#)

Next Steps

Now that you have configured and logged into the Manager, the next step is to add Oracle Linux KVM hosts, as described in [Configuring a KVM Host](#).

You will also need to add storage and configure logical networks. See the Storage and Networks tasks in the [Oracle Linux Virtualization Manager: Administration Guide](#).

Logging Out

To log out of the **Administration Portal**, click the person icon in the header bar and click **Sign Out**. You are returned to the **Login** page.

Configuring a KVM Host

To manage an Oracle Linux KVM host using Oracle Linux Virtualization Manager, prepare the KVM host by performing a fresh installation of Oracle Linux 8.8 (or later Oracle Linux 8 release) and enabling the required repositories, and then you add the host to a data center using the Administration Portal.

Before you begin, ensure you have satisfied the *KVM Host Requirements* as detailed in the [Oracle Linux Virtualization Manager: Architecture and Planning Guide](#).

Refer to the [Oracle® Linux: KVM User's Guide](#) for information on the supported guest operating systems.

Preparing a KVM Host

Before you can add an Oracle Linux KVM host, prepare it by performing a fresh installation of Oracle Linux 8.8 (or later Oracle Linux 8 release) and enabling the required repositories. You can download the installation ISO for Oracle Linux 8.8 (or later Oracle Linux 8 release) from the Oracle Software Delivery Cloud at <https://edelivery.oracle.com>.

1. Install Oracle Linux 8.8 (or later Oracle Linux 8 release) on the host.
 - Follow the instructions in the [Oracle® Linux 8: Installing Oracle Linux](#).
 - Select **Minimal Install** as the base environment for the installation.

Caution:

Do **NOT** select any other base environment than **Minimal Install** for the installation or your hosts will have incorrect qemu and libvirt versions, incorrect repositories configured, and no access to virtual machine consoles.

- Do not install any additional packages until after you have added the host to the Manager, because they may cause dependency issues.
2. **(Optional)** If you use a proxy server for Internet access, configure Yum with the proxy server settings. For more information, see the [Oracle® Linux: Managing Software on Oracle Linux](#).
 3. Complete one of the following sets of steps:

- **For ULN registered hosts or using Oracle Linux Manager**

Subscribe the system to the required channels and enable appstream modules.

- a. For ULN registered hosts, log in to <https://linux.oracle.com> with your ULN user name and password. For Oracle Linux Manager registered hosts, access your internal server URL.
- b. On the Systems tab, click the link named for the host in the list of registered machines.
- c. On the System Details page, click **Manage Subscriptions**.
- d. On the System Summary page, select each required channel from the list of available channels and click the right arrow to move the channel to the list of subscribed channels. Subscribe the system to the following channels:

```
– o18_x86_64_baseos_latest
– o18_x86_64_appstream
– o18_x86_64_kvm_appstream
– o18_x86_64_ovirt45
```

- ol8_x86_64_ovirt45_extras
- ol8_x86_64_gluster_appstream
- **(For VDSM)** ol8_x86_64_UEKR7

e. Click **Save Subscriptions**.

f. Install the Oracle Linux Virtualization Manager Release 4.5 package, which automatically enables/disables the required repositories.

```
# dnf install oracle-ovirt-release-45-el8
```

- **For Oracle Linux yum server configured KVM hosts**

Install the Oracle Linux Virtualization Manager Release 4.5 package and enable the required repositories.

 **Note:**

Installing the Oracle Linux Virtualization Manager Release 4.5 package configures an Oracle Linux KVM host; it does not install the Manager.

a. Enable the ol8_baseos_latest repository.

```
# dnf config-manager --enable ol8_baseos_latest
```

b. Install the Oracle Linux Virtualization Manager Release 4.5 package, which automatically enables/disables the required repositories.

```
# dnf install oracle-ovirt-release-45-el8
```

c. Use the dnf command to verify that the required repositories are enabled.

i. Clear the dnf cache.

```
# dnf clean all
```

ii. List the configured repositories and verify that the required repositories are enabled.

```
# dnf repolist
```

The following repositories must be enabled:

- ol8_64_baseos_latest
- ol8_appstream
- ol8_kvm_appstream
- ovirt-4.5
- ovirt-4.5-extra
- ol8_gluster_appstream
- **(For VDSM)** ol8_UEKR7

iii. If a required repository is not enabled, use the dnf config-manager to enable it.

```
# dnf config-manager --enable repository
```

4. If your host is running UEK R7:

a. Install the *Extra kernel modules* package.

```
# dnf install kernel-uek-modules-extra
```

b. Reboot the host.

5. **(Optional)** Activate the Cockpit web console and open the firewall port.

```
# systemctl enable --now cockpit.socket
# firewall-cmd --permanent --zone=public --add-service=cockpit
# firewall-cmd --reload
```

The Cockpit web interface can be used to monitor the host's resources and to perform administrative tasks. You can access the host's Cockpit web interface from the Administration Portal or by connecting directly to the host.

For more information about configuring `firewalld`, see [Configuring a Packet Filtering Firewall](#) in the [Oracle® Linux 8: Configuring the Firewall](#).

6. **(Optional)** Complete the previous steps to prepare additional KVM hosts.

The Oracle Linux KVM host is now ready to be added to the Manager using the Administration Portal.

Adding a KVM Host

Once you have configured an Oracle Linux KVM host, you use the Administration Portal to add the host to a data center so that it can be used to run virtual machines. You can follow the steps below to add KVM hosts installed with other supported guest operating systems.

! Important:

Oracle Linux Virtualization Manager allows you to overallocate a KVM host's memory and CPU resources. As the KVM host itself also needs memory and CPU in order to run, Oracle recommends that you reserve some memory and CPU for the KVM host. To do this, go to **Administration** and set a memory quota and a vCPU quota.

To add an Oracle Linux KVM host:

1. Log in to the Administration Portal.

See [Logging in to the Administration Portal](#) for details.

2. Go to **Compute** and then click **Hosts**.

3. On the **Hosts** pane, click **New**.

The **New Host** dialog box opens with the **General** tab selected on the sidebar.

4. From the **Host Cluster** drop-down list, select the data center and host cluster for the host.

The **Default** data center is auto-selected.

When you install Oracle Linux Virtualization Manager, a data center and cluster named Default is created. You can rename and configure this data center and

cluster, or you can add new data centers and clusters, to meet your needs. See the Data Centers or Clusters tasks in the [Oracle Linux Virtualization Manager: Administration Guide](#).

5. In the **Name** field, enter a name for the host.
6. In the **Hostname** field, enter the fully-qualified domain name or IP address of the host.
7. In the **SSH Port** field, change the standard SSH port 22 if the SSH server on the host uses a different port.
8. Under **Authentication**, select the authentication method to use.

Oracle recommends that you select **SSH PublicKey** authentication. If you select this option, copy the key displayed in the **SSH PublicKey** field to the `/root/.ssh/authorized_keys` file on the host.

Otherwise, enter the root user's password to use password authentication.

9. **(Optional)** Configure other settings for the host from the other tabs on the **New Host** sidebar.

 **Note:**

If you do not want to set any other configuration options now, you can always make changes later by selecting a host from the **Hosts** pane and clicking **Edit**.

10. Click **OK**.

The **Power Management Configuration** screen is displayed.

11. If you do not want to configure power management, click **OK**. Otherwise, click **Configure Power Management**. See [Configuring Power Management and Fencing on a Host](#) for more information.

The host is added to the list of hosts in the Manager. While the Manager is installing the host agent (VDSM) and other required packages on the host, the status of the host is shown as **Installing**. You can view the progress of the installation in the Hosts details pane. When the installation is complete, the host status changes to **Up**.

12. **(Optional)** Complete the previous steps to add more KVM hosts to the Manager.

 **Note:**

After a KVM host is added to a cluster, it is also crucial to avoid any spontaneous changes to the network configuration in `/etc/sysconfig/network-scripts/` or through the NetworkManager (e.g. nmcli). All changes to the network configuration should be carried out through the engine host/manager Administration Portal or REST API.

Now that you have your engine and host(s) configured, see the [Oracle Linux Virtualization Manager: Administration Guide](#) for detailed configuration and administrative tasks.

4

Self-Hosted Engine Deployment

In Oracle Linux Virtualization Manager, a self-hosted engine is a virtualized environment where the engine runs inside a virtual machine on the hosts in the environment. The virtual machine for the engine is created as part of the host configuration process. And, the engine is installed and configured in parallel to the host configuration.

Since the engine runs as a virtual machine and not on physical hardware, a self-hosted engine requires less physical resources. Additionally, since the engine is configured to be highly available, if the host running the Engine virtual machine goes into maintenance mode or fails unexpectedly the virtual machine is migrated automatically to another host in the environment. A minimum of two KVM hosts are required.

To review conceptual information, troubleshooting, and administration tasks, see the oVirt Self-Hosted Engine Guide in [oVirt Documentation](#).

To deploy a self-hosted engine, you perform a fresh installation of Oracle Linux 8.8 (or later Oracle Linux 8 release) on the host, install the Oracle Linux Virtualization Manager Release 4.5 package, and then run the hosted engine deployment tool to complete configuration.



Note:

If you are deploying a self-hosted engine as a hyperconverged infrastructure with GlusterFS storage, you must deploy GlusterFS *BEFORE* you deploy the self-hosted engine. See [Deploying GlusterFS Storage](#).

You can also deploy a self-hosted engine using the command line or Cockpit portal. If you want to use the command line, proceed to [Using the Command Line to Deploy](#). If you want to use the Cockpit portal, proceed to [Using the Cockpit Portal to Deploy](#).



Note:

If you are behind a proxy, you must use the command line option to deploy.

Self-Hosted Engine Prerequisites

In addition to the [Requirements and Scalability Limits](#), you must satisfy the following prerequisites before deploying a self-hosted engine.

- A minimum of two KVM hosts.
- A fully-qualified domain name for your engine and host with forward and reverse lookup records set in the DNS.
- A directory of at least 5 GB on the host for the oVirt Engine Appliance. During the deployment process the `/var/tmp` directory is checked to see if it has enough space to

extract the appliance files. If the `/var/tmp` directory does not have enough space, you can specify a different directory or mount external storage.

 **Note:**

The VDSM user and KVM group must have read, write, and execute permissions on the directory.

- Prepared storage of at least 74 GB to be used as a data storage domain dedicated to the engine virtual machine. The data storage domain is created during the self-hosted engine deployment.

If you are using iSCSI storage, do not use the same iSCSI target for the self-hosted engine storage domain and any additional storage domains.

NOT_SUPPORTED:

When you have a data center with only one active data storage domain and that domain gets corrupted, you are unable to add new data storage domains or remove the corrupted data storage domain. If you have deployed your self-hosted engine in such a data center and its data storage domain gets corrupted, you must redeploy your self-hosted engine.

- The host you are using to deploy a self-hosted engine, must be able to access yum.oracle.com.

Deploying the Self-Hosted Engine

You must perform a fresh installation of Oracle Linux 8.8 (or later Oracle Linux 8 release) on an Oracle Linux Virtualization Manager host before deploying a self-hosted engine. You can download the installation ISO for from the Oracle Software Delivery Cloud at <https://edelivery.oracle.com>.

1. Install Oracle Linux 8.8 (or later Oracle Linux 8 release) on the host using the **Minimal Install** base environment.

 **Caution:**

Do **NOT** select any other base environment than **Minimal Install** for the installation or your hosts will have incorrect `qemu` and `libvirt` versions, incorrect repositories configured, and no access to virtual machine consoles.

Do not install any additional packages until after you have installed the Manager packages, because they may cause dependency issues.

Follow the instructions in the [Oracle® Linux 8: Installing Oracle Linux](#).

2. Ensure that the `firewalld` service is enabled and started.

For more information about configuring `firewalld`, see [Configuring a Packet Filtering Firewall in the Oracle® Linux 8: Configuring the Firewall](#).

3. Complete one of the following sets of steps:

- **For ULN registered hosts or using Oracle Linux Manager**

Subscribe the system to the required channels.

- a. For ULN registered hosts, log in to <https://linux.oracle.com> with your ULN user name and password. For Oracle Linux Manager registered hosts, access your internal server URL.
- b. On the Systems tab, click the link named for the host in the list of registered machines.
- c. On the System Details page, click **Manage Subscriptions**.
- d. On the System Summary page, select each required channel from the list of available channels and click the right arrow to move the channel to the list of subscribed channels. Subscribe the system to the following channels:

- `ol8_x86_64_baseos_latest`
- `ol8_x86_64_appstream`
- `ol8_x86_64_kvm_appstream`
- `ol8_x86_64_ovirt45`
- `ol8_x86_64_ovirt45_extras`
- `ol8_x86_64_gluster_appstream`
- **(For VDSM)** `ol8_x86_64_UEKR7`

- e. Click **Save Subscriptions**.
- f. Install the Oracle Linux Virtualization Manager Release 4.5 package, which automatically enables/disables the required repositories.

```
# dnf install oracle-ovirt-release-45-el8
```

- **For Oracle Linux yum server hosts**

Install the Oracle Linux Virtualization Manager Release 4.5 package and enable the required repositories.

- a. Enable the `ol8_baseos_latest` yum repository.

```
# dnf config-manager --enable ol8_baseos_latest
```

- b. Install the Oracle Linux Virtualization Manager Release 4.5 package, which automatically enables/disables the required repositories.

```
# dnf install oracle-ovirt-release-45-el8
```

- c. Use the `dnf` command to verify that the required repositories are enabled.

- i. Clear the yum cache.

```
# dnf clean all
```

- ii. List the configured repositories and verify that the required repositories are enabled.

```
# dnf repolist
```

The following repositories must be enabled:

```

- ol8_baseos_latest
- ol8_appstream
- ol8_kvm_appstream
- ovirt-4.5
- ovirt-4.5-extra
- ol8_gluster_appstream
- (For VDSM) ol8_UEKR7

```

- iii. If a required repository is not enabled, use the config-manager to enable it.

```

# dnf config-manager --enable
repository

```

4. If your host is running UEK R7:
 - a. Install the *Extra kernel modules* package.

```

# dnf install kernel-uek-modules-extra

```
 - b. Reboot the host.
5. Install the hosted engine deployment tool and engine appliance.

```

# dnf install ovirt-hosted-engine-setup -y

```

Using the Command Line to Deploy

You can deploy the self-hosted engine from the command line. A script collects the details of your environment and uses them to configure the host and the engine.

1. Start the deployment. IPv6 is used by default. To use IPv4, specify the `--4` option:

```

# hosted-engine --deploy --4

```

Optionally, use the `--ansible-extra-vars` option to define variables for the deployment. For example:

```

# hosted-engine --deploy --4 --ansible-extra-vars="@/root/extra-vars.yml"

```

```

cat /root/extra-vars.yml
---
he_pause_host: true
he_proxy: "http://<host>:<port>"
he_enable_keycloak: false

```

See the [oVirt Documentation](#) for more information.

2. Enter Yes to begin deployment.

```

Continuing will configure this host for serving as hypervisor and will
create a local VM
with a running engine. The locally running engine will be used to configure
a new storage
domain and create a VM there. At the end the disk of the local VM will be
moved to the
shared storage.
Are you sure you want to continue? (Yes, No) [Yes]:

```

 **Note:**

The hosted-engine script creates a virtual machine and uses cloud-init to configure it. The script also runs engine-setup and reboots the system so that the virtual machine can be managed by the high availability agent.

3. Enter the name of the data center or accept the default.

```
Please enter the name of the data center where you want to deploy this hosted-
engine
host. Data center [Default]:
```

4. Enter a name for the cluster or accept the default.

```
Please enter the name of the cluster where you want to deploy this hosted-engine
host.
Cluster [Default]:
```

5. Keycloak integration is a technology preview feature for internal Single-Sign-On (SSO) provider for the Engine and it deprecates AAA. The default response is Yes; however, since this is a preview feature, enter No.

```
Configure Keycloak integration on the engine(Yes, No) [Yes]:No
```

6. Configure the network.

- a.** If the gateway that displays is correct, press Enter to configure the network.
- b.** Enter a pingable address on the same subnet so the script can check the host's connectivity.

```
Please indicate a pingable gateway IP address [X.X.X.X]:
```

- c.** The script detects possible NICs to use as a management bridge for the environment. Select the default.

```
Please indicate a nic to set ovirtmgmt bridge on: (eth1, eth0) [eth1]:
```

7. Enter the path to an OVA archive if you want to use a custom appliance for the virtual machine installation. Otherwise, leave this field empty to use the oVirt Engine Appliance.

```
If you want to deploy with a custom engine appliance image, please specify the
path to
the OVA archive you would like to use.
Entering no value will use the image from the ovirt-engine-appliance rpm,
installing it if needed.
Appliance image path []:
```

8. Specify the fully-qualified domain name for the engine virtual machine.

```
Please provide the FQDN you would like to use for the engine appliance.
Note: This will be the FQDN of the engine VM you are now going to launch,
it should not point to the base host or to any other existing machine.
Engine VM FQDN: manager.example.com
Please provide the domain name you would like to use for the engine appliance.
Engine VM domain: [example.com]
```

9. Enter and confirm a root password for the engine.

```
Enter root password that will be used for the engine appliance:
Confirm appliance root password:
```

10. Optionally, enter an SSH public key to enable you to log in to the engine as the root user and specify whether to enable SSH access for the root user.

```

Enter ssh public key for the root user that will be used for the engine
appliance (leave it empty to skip):
Do you want to enable ssh access for the root user (yes, no, without-
password)
[yes]:
You may provide an SSH public key, that will be added by the deployment
script to the
authorized_keys file of the root user in the engine appliance.
This should allow you passwordless login to the engine machine after
deployment.
If you provide no key, authorized_keys will not be touched.
SSH public key []:
[WARNING] Skipping appliance root ssh public key
Do you want to enable ssh access for the root user? (yes, no, without-
password) [yes]:

```

11. Enter the virtual machine's CPU and memory configuration.

```

Please specify the number of virtual CPUs for the VM (Defaults to appliance
OVF value): [4]:
Please specify the memory size of the VM in MB. The default is the appliance
OVF value [16384]:

```

12. Enter a MAC address for the engine virtual machine or accept a randomly generated MAC address.

```

You may specify a unicast MAC address for the VM or accept a randomly
generated default [00:16:3e:3d:34:47]:

```

 **Note:**

If you want to provide the engine virtual machine with an IP address using DHCP, ensure that you have a valid DHCP reservation for this MAC address. The deployment script does not configure the DHCP server for you.

13. Enter the virtual machine's networking details.

```

How should the engine VM network be configured (DHCP, Static)[DHCP]?

```

 **Note:**

If you specified Static, enter the IP address of the Engine. The static IP address must belong to the same subnet as the host. For example, if the host is in 10.1.1.0/24, the Engine virtual machine's IP must be in the same subnet range (10.1.1.1-254/24).

```

Please enter the IP address to be used for the engine VM [x.x.x.x]:
Please provide a comma-separated list (max 3) of IP addresses of
domain
name servers for the engine VM
Engine VM DNS (leave it empty to skip):

```

14. Specify whether to add entries in the virtual machine's /etc/hosts file for the engine virtual machine and the base host. Ensure that the host names are resolvable.

Add lines for the appliance itself and for this host to /etc/hosts on the engine VM?

Note: ensuring that this host could resolve the engine VM hostname is still up to you.

Add lines to /etc/hosts? (Yes, No)[Yes]:

15. Provide the name and TCP port number of the SMTP server, the email address used to send email notifications, and a comma-separated list of email addresses to receive these notifications. Or, press Enter to accept the defaults.

Please provide the name of the SMTP server through which we will send notifications [localhost]:

Please provide the TCP port number of the SMTP server [25]:

Please provide the email address from which notifications will be sent [root@localhost]:

Please provide a comma-separated list of email addresses which will get notifications [root@localhost]:

16. Enter and confirm a password for the admin@internal user to access the Administration Portal.

Enter engine admin password:

Confirm engine admin password:

The script creates the virtual machine which can take time if it needs to install the oVirt Engine Appliance. After creating the virtual machine, the script continues gathering information.

17. Select the type of storage to use.

Please specify the storage you would like to use (glusterfs, iscsi, fc, nfs)[nfs]:

- If you selected NFS, enter the version, full address and path to the storage, and any mount options.

Please specify the nfs version you would like to use (auto, v3, v4, v4_1)[auto]:

Please specify the full shared storage connection path to use (example: host:/path):

storage.example.com:/hosted_engine/nfs

If needed, specify additional mount options for the connection to the hosted-engine storage domain []:

- If you selected iSCSI, enter the portal details and select a target and LUN from the auto-detected lists. You can only select one iSCSI target during the deployment, but multipathing is supported to connect all portals of the same portal group.

 **Note:**

To specify more than one iSCSI target, you must enable multipathing before deploying the self-hosted engine. There is also a Multipath Helper tool that generates a script to install and configure multipath with different options.

Please specify the iSCSI portal IP address:

Please specify the iSCSI portal port [3260]:

Please specify the iSCSI discover user:

Please specify the iSCSI discover password:

Please specify the iSCSI portal login user:

Please specify the iSCSI portal login password:

The following targets have been found:

```
[1]   iqn.2017-10.com.redhat.example:he
      TPGT: 1, portals:
          192.168.1.xxx:3260
          192.168.2.xxx:3260
          192.168.3.xxx:3260
```

Please select a target (1) [1]: 1

The following luns have been found on the requested target:

```
[1] 360003ff44dc75adcb5046390a16b4beb 199GiB MSFT Virtual HD
      status: free, paths: 1 active
```

Please select the destination LUN (1) [1]:

- If you selected GlusterFS, enter the full address and path to the storage, and any mount options. Only replica 3 Gluster storage is supported.

* Configure the volume as follows as per [Gluster Volume Options for Virtual

```
Machine Image Store]
(documentation/admin-guide/chap-Working_with_Gluster_Storage#Options
set on Gluster Storage Volumes to Store Virtual Machine Images)
```

Please specify the full shared storage connection path to use
(example: host:/path):

```
storage.example.com:/hosted_engine/gluster_volume
```

If needed, specify additional mount options for the connection to the
hosted-engine storage domain []:

- If you selected Fibre Channel, select a LUN from the auto-detected list. The host bus adapters must be configured and connected. The deployment script auto-detects the available LUNs, and the LUN must not contain any existing data.

The following luns have been found on the requested target:

```
[1] 3514f0c5447600351 30GiB XtremIO XtremApp
      status: used, paths: 2 active
```

```
[2] 3514f0c5447600352 30GiB XtremIO XtremApp
      status: used, paths: 2 active
```

Please select the destination LUN (1, 2) [1]:

18. Enter the engine disk size:

Please specify the size of the VM disk in GB: [50]:

If successful, one data center, cluster, host, storage domain, and the engine virtual machine are already running.

19. Optionally, log into the Oracle Linux Virtualization Manager Administration Portal to add any other resources.

In the Administration Portal, the engine virtual machine, the host running it, and the self-hosted engine storage domain are flagged with a gold crown.

20. Enable the required repositories on the Engine virtual machine.

21. Optionally, add a directory server using the `ovirt-engine-extension-aaa-ldap-setup` interactive setup script so you can add additional users to the environment.

Using the Cockpit Portal to Deploy



Note:

If you are behind a proxy, you must use the command line option to deploy your self-hosted engine.

To deploy the self-hosted engine using the Cockpit portal, complete the following steps.

1. Install the Cockpit dashboard.

```
# dnf install cockpit-ovirt-dashboard -y
```

2. Open the Cockpit port 9090 on firewalld.

```
# firewall-cmd --permanent --zone=public --add-port=9090/tcp
```

```
# firewall-cmd --reload
```

3. Enable and start the Cockpit service

```
# systemctl enable --now cockpit.socket
```

4. Log into the Cockpit portal from the following URL:

```
https://host_IP_or_FQDN:9090
```

5. To start the self-hosted engine deployment, click **Virtualization** and select **Hosted Manager**.

6. Click **Start** under **Hosted Manager**.

7. Provide the following details for the Engine virtual machine.

- a. In the **Engine VM FQDN** field, enter the Engine virtual machine FQDN. Do not use the FQDN of the host.
- b. In the **MAC Address** field, enter a MAC address for the Engine virtual machine or leave blank and the system provides a randomly-generated address.
- c. From the **Network Configuration** drop-down list, select **DHCP** or **Static**.
 - To use **DHCP**, you must have a DHCP reservation (a pre-set IP address on the DHCP server) for the Engine virtual machine. In the **MAC Address** field, enter the MAC address.
 - To use **Static**, enter the virtual machine IP, the gateway address, and the DNS servers. The IP address must belong to the same subnet as the host.

- d. Select the **Bridge Interface** from the drop-down list.

- e. Enter and confirm the virtual machine's **Root Password**.

- f. Specify whether to allow **Root SSH Access**.

- g. Enter the **Number of Virtual CPUs** for the virtual machine.

- h. Enter the **Memory Size (MiB)**. The available memory is displayed next to the field.

8. Optionally, click **Advanced** to provide any of the following information.

- Enter a **Root SSH Public Key** to use for root access to the Engine virtual machine.

- Select the **Edit Hosts File** check box if you want to add entries for the Engine virtual machine and the base host to the virtual machine's `/etc/hosts` file. You must ensure that the host names are resolvable.
 - Change the management **Bridge Name**, or accept the default of `ovirtmgmt`.
 - Enter the **Gateway Address** for the management bridge.
 - Enter the **Host FQDN** of the first host to add to the Engine. This is the FQDN of the host you are using for the deployment.
9. Click **Next**.
 10. Enter and confirm the **Admin Portal Password** for the `admin@internal` user.
 11. Optionally, configure event notifications.
 - Enter the **Server Name** and **Server Port Number** of the SMTP server.
 - Enter a **Sender E-Mail Address**.
 - Enter **Recipient E-Mail Addresses**.
 12. Click **Next**.
 13. Review the configuration of the Engine and its virtual machine. If the details are correct, click **Prepare VM**.
 14. When the virtual machine installation is complete, click **Next**.
 15. Select the **Storage Type** from the drop-down list and enter the details for the self-hosted engine storage domain.
 - For NFS:
 - a. In the **Storage Connection** field, enter the full address and path to the storage.
 - b. If required, enter any **Mount Options**.
 - c. Enter the **Disk Size (GiB)**.
 - d. Select the **NFS Version** from the drop-down list.
 - e. Enter the **Storage Domain Name**.
 - For iSCSI:
 - a. Enter the **Portal IP Address**, **Portal Port**, **Portal Username**, and **Portal Password**.
 - b. Click **Retrieve Target List** and select a target. You can only select one iSCSI target during the deployment, but multipathing is supported to connect all portals of the same portal group.

 **Note:**

To specify more than one iSCSI target, you must enable multipathing before deploying the self-hosted engine. There is also a Multipath Helper tool that generates a script to install and configure multipath with different options.

- c. Enter the **Disk Size (GiB)**.
- d. Enter the **Discovery Username** and **Discovery Password**.

- For FibreChannel:
 - a. Enter the **LUN ID**. The host bus adapters must be configured and connected and the LUN must not contain any existing data.
 - b. Enter the **Disk Size (GiB)**.
 - For Gluster Storage:
 - a. In the **Storage Connection** field, enter the full address and path to the storage.
 - b. If required, enter any **Mount Options**.
 - c. Enter the **Disk Size (GiB)**.
16. Click **Next**.
 17. Review the storage configuration. If the details are correct, click **Finish Deployment**.
 18. When the deployment is complete, click **Close**.

If successful, one data center, cluster, host, storage domain, and the engine virtual machine are already running.
 19. Optionally, log into the Oracle Linux Virtualization Manager Administration Portal to add any other resources.

In the Administration Portal, the engine virtual machine, the host running it, and the self-hosted engine storage domain are flagged with a gold crown.
 20. Enable the required repositories on the Engine virtual machine.
 21. Optionally, add a directory server using the `ovirt-engine-extension-aaa-ldap-setup` interactive setup script so you can add additional users to the environment.
 22. To view the self-hosted engine's status in Cockpit, under **Virtualization** click **Hosted Engine**.

Enabling High-Availability

The host that houses the self-hosted engine is not highly available by default. Since the self-hosted engine runs inside a virtual machine on a host, if you do not configure high-availability for the host, then virtual machine recovery after a host crash is not possible.

Configuring a Highly Available Host

If you want the hosts in a cluster to be responsive and available when unexpected failures happen, you should use fencing. Fencing allows a cluster to react to unexpected host failures and enforce power saving, load balancing, and virtual machine availability policies. You should configure the fencing parameters for your host's power management device and test their correctness from time to time.

A *Non Operational* host is different from a *Non Responsive* host. A *Non Operational* host can communicate with the Manager, but has incorrect configuration, for example a missing logical network. A *Non Responsive* host cannot communicate with the Manager.

In a fencing operation, a non-responsive host is rebooted, and if the host does not return to an active status within a prescribed time, it remains non-responsive pending manual intervention and troubleshooting.

The Manager can perform management operations after it reboots, by a proxy host, or manually in the **Administration Portal**. All the virtual machines running on the non-

responsive host are stopped, and highly available virtual machines are restarted on a different host. At least two hosts are required for power management operations.

! Important:

If a host runs virtual machines that are highly available, power management must be enabled and configured.

Configuring Power Management and Fencing on a Host

The Manager uses a proxy to send power management commands to a host power management device because the engine does not communicate directly with fence agents. The host agent (VDSM) executes power management device actions and another host in the environment is used as a fencing proxy. This means that you must have at least two hosts for power management operations.

When you configure a fencing proxy host, make sure the host is in:

- the same cluster as the host requiring fencing.
- the same data center as the host requiring fencing.
- `UP` or `Maintenance` status to remain viable.

Power management operations can be performed in three ways:

- by the Manager after it reboots
- by a proxy host
- manually in the **Administration Portal**

To configure power management and fencing on a host:

1. Click **Compute** and select **Hosts**.
2. Select a host and click **Edit**.
3. Click the **Power Management** tab.
4. Check **Enable Power Management** to enable the rest of the fields.
5. Check **Kdump integration** to prevent the host from fencing while performing a kernel crash dump. Kdump integration is enabled by default.

! Important:

If you enable or disable Kdump integration on an existing host, you must reinstall the host.

6. **(Optional)** Check **Disable policy control of power management** if you do not want your host's power management to be controlled by the scheduling policy of the host's cluster.
7. To configure a fence agent, click the plus sign (+) next to **Add Fence Agent**. The **Edit fence agent** pane opens.

8. Enter the **Address** (IP Address or FQDN) to access the host's power management device.
9. Enter the **User Name** and **Password** of the of the account used to access the power management device.
10. Select the power management device **Type** from the drop-down list.
11. Enter the **Port** (SSH) number used by the power management device to communicate with the host.
12. Enter the **Slot** number used to identify the blade of the power management device.
13. Enter the **Options** for the power management device. Use a comma-separated list of key-value pairs.
 - If you leave the **Options** field blank, you are able to use both IPv4 and IPv6 addresses
 - To use only IPv4 addresses, enter `inet4_only=1`
 - To use only IPv6 addresses, enter `inet6_only=1`
14. Check **Secure** to enable the power management device to connect securely to the host.
You can use ssh, ssl, or any other authentication protocol your power management device supports.
15. Click **Test** to ensure the settings are correct and then click **OK**.
Test Succeeded, Host Status is: on displays if successful.

NOT_SUPPORTED:

Power management parameters (userid, password, options, etc.) are tested by the Manager only during setup and manually after that. If you choose to ignore alerts about incorrect parameters, or if the parameters are changed on the power management hardware without changing in the Manager as well, fencing is likely to fail when most needed.

16. Fence agents are sequential by default. To change the sequence in which the fence agents are used:
 - a. Review your fence agent order in the **Agents by Sequential Order** field.
 - b. To make two fence agents concurrent, next to one fence agent click the **Concurrent with** drop-down list and select the other fence agent.
You can add additional fence agents to this concurrent fence agent group.
17. Expand the **Advanced Parameters** and use the up and down buttons to specify the order in which the Manager searches the host's **cluster** and **dc** (data center) for a power management proxy.
18. To add an additional power management proxy:
 - a. Click the plus sign (+) next to **Add Power Management Proxy**.
The **Select fence proxy preference type to add** pane opens.
 - b. Select a power management proxy from the drop-down list and then click **OK**.
Your new proxy displays in the **Power Management Proxy Preference** list.

 **Note:**

By default, the Manager searches for a fencing proxy within the same cluster as the host. If The Manager cannot find a fencing proxy within the cluster, it searches the data center.

19. Click OK.

From the list of hosts, the exclamation mark next to the host's name disappeared, signifying that you have successfully configured power management and fencing.

Preventing Host Fencing During Boot

After you configure power management and fencing, when you start the Manager it automatically attempts to fence non-responsive hosts that have power management enabled *after* the quiet time (5 minutes by default) has elapsed. You can opt to extend the quiet time to prevent, for example, a scenario where the Manager attempts to fence hosts while they boot up. This can happen after a data center outage because a host's boot process is normally longer than the Manager boot process.

You can configure quiet time using the `engine-config` command option

`DisableFenceAtStartupInSec`:

```
# engine-config -s DisableFenceAtStartupInSec=number
```

Checking Fencing Parameters

To automatically check the fencing parameters, you can configure the `PMHealthCheckEnabled` (false by default) and `PMHealthCheckIntervalInSec` (3600 sec by default) `engine-config` options.

```
# engine-config -s PMHealthCheckEnabled=True
```

```
# engine-config -s PMHealthCheckIntervalInSec=number
```

When set to true, `PMHealthCheckEnabled` checks all host agents at the interval specified by `PMHealthCheckIntervalInSec` and raises warnings if it detects issues.

Installing Additional Self-Hosted Engine Hosts

You add self-hosted engine hosts the same way as a regular host, with an additional step to deploy the host as a self-hosted engine host. The shared storage domain is automatically detected and the host can be used as a failover host to host the Engine virtual machine when required. You can also add regular hosts to a self-hosted engine environment, but they cannot be used to host the Engine virtual machine.

 **Important:**

Before you begin, refer to [Preparing a KVM Host](#).

To install an additional self-hosted engine host, complete the following steps.

1. In the **Administration Portal**, go to **Compute** and click **Hosts**.
2. Click **New**.
For information on additional host settings, see the Admin Guide in the latest upstream [oVirt Documentation](#).
3. Use the drop-down list to select the **Data Center** and **Host Cluster** for the new host.
4. Enter the **Name** and the **Address** of the new host. The standard SSH port, port 22, is auto-filled in the **SSH Port** field.
5. Select an authentication method to use for the engine to access the host.
 - Enter the root user's password to use password authentication.
 - Alternatively, copy the key displayed in the **SSH PublicKey** field to `/root/.ssh/authorized_keys` on the host to use public key authentication.
6. Optionally, configure power management, where the host has a supported power management card. For information, see [Configuring Power Management and Fencing on a Host](#).
7. Click the **Hosted Engine** sub-tab.
8. Select the **Deploy** radio button.
9. Click **OK**.

Cleaning up the Deployment

If your self-hosted engine deployment fails, you must perform a few cleanup tasks before retrying.

1. Run the hosted engine cleanup command:

```
# /usr/sbin/ovirt-hosted-engine-cleanup
```
2. Remove the storage:

```
# rm -rf <storage_repo>/*
```
3. If the deployment failed after the local, temporary hosted engine virtual machine is created, you might need to clean up the local virtual machine repository:

```
# rm -rf /var/tmp/localvm*
```

Upgrading Or Updating the Self-Hosted Engine

See [Upgrading Your Environment to 4.5](#) or [Updating the Self-Hosted Engine](#) in the [Oracle Linux Virtualization Manager: Administration Guide](#).

5

Deploying GlusterFS Storage

Oracle Linux Virtualization Manager has been integrated with GlusterFS, an open source scale-out distributed filesystem, to provide a hyperconverged solution where both compute and storage are provided from the same hosts. Gluster volumes residing on the hosts are used as storage domains in the Manager to store the virtual machine images. In this scenario, the Manager is run as a self-hosted engine within a virtual machine on these hosts; although, you can deploy GlusterFS within a standalone environment.

For instructions on creating a GlusterFS storage domain, refer to the [My Oracle Support \(MOS\)](#) article *How to Create Glusterfs Storage Domain (Doc ID 2679824.1)*.

Note:

If you are deploying a self-hosted engine as hyperconverged infrastructure with GlusterFS storage, you must deploy GlusterFS *before* you deploy the self-hosted engine. For more information about using GlusterFS, including prerequisites, see the [Oracle Linux GlusterFS documentation](#).

Deploying GlusterFS Storage Using Cockpit

To deploy GlusterFS storage using the Cockpit web interface, complete the following steps.

Important:

To deploy GlusterFS, you must have at least three (3) KVM hosts. If you want more than three KVM hosts, they must be added in factors of three.

1. Ensure that on all KVM hosts you have installed the following packages:
 - `cockpit-ovirt-dashboard` to provide a UI for installation
 - `vdsm-gluster` to manage gluster services
 - `ovirt-host` on the KVM host used for cockpit deployment
2. Go to **Compute**, and then click **Hosts**.
3. Under the **Name** column, click the host to be used as the designated server.
4. Click **Host Console**.
5. Enter your login credentials (the user name and password of the root account.).
6. Go to **Virtualization** and then click **Hosted Engine**.
7. Click **Redeploy** under **Hosted Engine Setup**.
8. Click **Start** under **Hyperconverged**.

9. On the **Hosts** screen, enter 3 (or more) KVM hosts that are in the data center to be used for GlusterFS, with the main designated KVM host entered first and click **Next** when finished.
10. On the **FQDNs** screen, enter the FQDN (or IP address) for the hosts to be managed by the Hosted Engine and click **Next** when finished.

 **Note:**

The FQDN of the designated server is input during the Hosted Engine deployment process and is not asked for here.

11. Click **Next** on the **Packages** screen.
12. On the **Volumes** screen, create the minimum storage domains that are required: `engine` or `data`. Click **Next** when finished.

For example:

engine

- **Name:** `engine`
- **Volume Type:** `Replicate` (default)
- **Arbiter:** Ensure the check box is selected.
- **Brick Dirs:** `/gluster_bricks/engine/engine` (default)

data

- **Name:** `data`
- **Volume Type:** `Replicate` (default)
- **Arbiter:** Ensure the check box is selected.
- **Brick Dirs:** `/gluster_bricks/data/data` (default)

13. On the **Brick Locations** screen, specify the brick locations for your volumes and click **Next** when finished.

For this step, you specify the brick locations for your volumes (`engine`, `data`, `export`, and `iso`).

14. Review the screen and click **Deploy**.
 - If you are using an internal disk as the Gluster disk, no edits are required and you can simply click **Deploy** to continue with the deployment.
 - If you are using an external iSCSI ZFS external drive as the Gluster disk, click **Edit** to edit the `gdeployConfig.conf` file and specify the block device on each server that is being used for storage. Click **Save** and then click **Deploy** to continue with the deployment.

This process takes some time to complete, as the `gdeploy` tool installs required packages and configures Gluster volumes and their underlying storage.

A message display on the screen when the deployment completes successfully.

Creating a GlusterFS Storage Domain Using the Manager

To add a GlusterFS storage volume as a storage domain:

1. Go to **Storage** and then click **Domains**.
2. On the **Storage Domains** pane, click the **New Domain** button.
3. For the **Name** field, enter a name for the data domain.
4. From the **Data Center** drop-down list, select the data center where the GlusterFS volume is deployed. By default, the **Default** option is selected in the drop-down list.
5. From the **Domain Function** drop-down list, select the domain function. By default, the **Data** option is selected in the drop-down list.

For this step, leave **Data** as the domain function because a data domain is being created in this example.

6. From the **Storage Type** drop-down list, select **GlusterFS**.
7. For the **Host to Use** drop-down list, select the host for which to attach the data domain.
8. When **GlusterFS** is selected for the **Storage Type**, the **New Domain** dialog box updates to display additional configuration fields associated with GlusterFS storage domains.
9. Ensure the **Use managed gluster volume** check box is not selected.
10. From the **Gluster** drop-down list, select the path to which domain function you are creating.
11. For the **Mount Options** option, specify additional mount options in a comma-separated list, as you would using the mount -o command.
12. **(Optional)** Configure the advanced parameters.
13. Click **OK** to mount the volume as a storage domain.

You can click **Tasks** to monitor the various processing steps that are completed to add the GlusterFS storage domain to the data center.