

Oracle Linux Virtualization Manager

Release Notes



F52195-19
April 2026



Oracle Linux Virtualization Manager Release Notes,
F52195-19

Copyright © 2022, 2026, Oracle and/or its affiliates.

Contents

1	About the Docs	
	Documentation License	2
	Conventions	2
	Documentation Accessibility	2
	Access to Oracle Support for Accessibility	2
	Diversity and Inclusion	2
2	What's New	
3	Technology Preview	
	Oracle Linux Virtualization Manager	1
	Compute	1
	Network	2
	Storage	2
	Virtual Machines (Guests)	2
4	Deprecated Features	
5	Known Issues	
	Oracle Linux Virtualization Manager Issues	1
	Global Configuration Issues	2
	Compute Issues	4
	Network Issues	8
	Storage Issues	12
	Virtual Machine Issues	14
6	Feedback and Support	
	Obtaining the Log Files	1

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 consists of:

- **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 the 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
 - Configuring Keycloak authentication
 - 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 the environment
 - Active-active and active-passive disaster recovery solutions

See also:

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

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](#).

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

What's New

Release 4.5 of Oracle Linux Virtualization Manager is based on [oVirt community release version 4.5.5](#).

Release 4.5.5 bug fixes and support for new features

- Federal Information Processing Standard (FIPS) compliance support, which you can enable selectively on specific VMs and bare metal machines, or across environments. For more information, see *FIPS Mode Deployment* in the [Oracle Linux Virtualization Manager: Getting Started](#).
- Intel Emerald Rapids CPU type was added to support the latest generation of Intel x86 servers for Oracle Virtualization compute hosts. The feature is supported with ovirt-engine-4.5.5-1.57 and greater.
- New guest operating system selections were added for users to select for their virtual machines. Oracle Linux 10, Red Hat Enterprise Linux 10, Rocky Linux 10, AlmaLinux 10, Ubuntu 25, and Windows Server 2025 are options starting with ovirt-engine-4.5.5-1.54.
- Windows Server Failover Clustering support for Windows server nodes running as virtual machines. A new user interface selection was added for SCSI-3 persistent reservations to support the feature. The feature is supported with ovirt-engine-4.5.5-1.57 and greater. Additional information on the feature can be found in the [Oracle Linux Virtualization Manager: Administration Guide](#)
- Performance Co-Pilot was enabled on Oracle Linux KVM hosts to gathering more data for troubleshooting issues. The feature is supported with ovirt-engine-4.5.5-1.42 and greater. Additional information on the feature can be found in the [Oracle Linux Virtualization Manager: Administration Guide](#).
- Precheck scripts were delivered for users to check to ensure their hosts are configured correctly before installing the engine. The precheck scripts are delivered starting with the release RPM version oracle-ovirt-release-45-el8-1.0-39. More information on the feature can be found in the [Oracle Linux Virtualization Manager: Getting Started Guide](#)
- Customers can add new Oracle Linux 9 KVM compute hosts while continuing to run Oracle Linux 8 KVM compute hosts. The feature is supported with ovirt-engine-4.5.5-1.65 and greater. Oracle Linux 9 KVM hosts can only be added to a cluster with a compatibility version of 4.7.

Release 4.5.4 bug fixes and support for new features

- For added security, you can enable HTTP Strict Transport Security (HSTS). See *Enabling HTTP Strict Transport Security* in the [Oracle Linux Virtualization Manager: Administration Guide](#).
- Multiple updates to better support backup and restore with snapshots
- Multiple user interface and usability improvements:
 - A new search box for vNIC Profiles was added oVirt Engine WebUI on the main vNIC profiles tab
 - A confirmation prompt was added when suspending a virtual machine

- A new warning event is raised if MAC Address Pool free and available addresses are below the set threshold
- A new warning is raised when the number of hosted-engine hosts exceeds the recommended maximum
- A new progress tracking status was added for Cluster Upgrades
- A new cluster-level 4.7 is now available. The 4.7 cluster level will be the default for new installs. Upgrades will carry forward the 4.4 cluster levels.
- A new parameter was added to the `evenly_distributed` scheduling policy that takes into account the ratio between virtual and physical CPUs on the host.
- Storage domains can now be deactivated and set to maintenance even if there are suspended virtual machines.
- A new user-specified delay parameter, Console Disconnect Action, was added for VNC console connections. This parameter defines what happens when the console is disconnected:
 - **No action** - No action is taken.
 - **Lock screen** (default) - For all Linux machines and for Windows desktops this locks the currently active user session. For Windows servers, this locks the desktop and the currently active user.
 - **Logout user** - For all Linux machines and Windows desktops, this logs out the currently active user session. For Windows servers, the desktop and the currently active user are logged out.
 - **Shutdown virtual machine** - Initiates a graceful virtual machine shutdown.
 - **Reboot virtual machine** - Initiates a graceful virtual machine reboot.

This setting can be changed while the virtual machine is running but will not take effect until a new console connection is established.

The following technology preview features are tested and fully supported with this release:

- **Grafana integration:** You can monitor your environment and display reports using Grafana dashboards that are based on data collected from the Data Warehouse database (`ovirt_engine_history` database). For more information, see *Data Visualization with Grafana* in the [Oracle Linux Virtualization Manager: Architecture and Planning Guide](#).
- **Keycloak integration:** Keycloak provides internal Single Sign-On (SSO) for the Engine and supports modern identity management capabilities such as user federation with Active Directory or LDAP and MFA. For more information, see *Keycloak Integration and Management* in the [Oracle Linux Virtualization Manager: Administrator's Guide](#).
- The `ovirt-engine-rename` command (the oVirt Engine Rename Tool) which changes the fully-qualified DNS name of Oracle Linux Virtualization Manager.

Additional updates:

- Support for the new `kvm_utils3` appstream in the Oracle Linux 8 KVM Appstream Latest channel. This appstream delivers more advanced versions of QEMU and Libvirt.
- Many of the third-party packages and other RPMs

For Requirements and Scalability Limits, see the [Oracle Linux Virtualization Manager: Architecture and Planning Guide](#).

3

Technology Preview

The following features of Oracle Linux Virtualization Manager are currently still under development, but are made available to you for testing and evaluation purposes.

Oracle Linux Virtualization Manager

- Automate configuring of Oracle Linux Virtualization Manager using the Ansible roles provided in the `ovirt-ansible-roles` package.
- Data warehouse service and history database and engine database installation or migration to a separate host.
- Full sampling scale for data warehouse.
- Dashboard migration to a separate host.
- Additional functionality to the Admin Portal through custom user interface plugins.
- External providers that provide external resources to Oracle Linux Virtualization Manager.
- The `engine-config` command (the Engine Configuration Tool) changes to global configuration settings for Oracle Linux Virtualization Manager.
- USB Filter Editor used to create the `usbfilter.txt` policy file for filtering USB devices on Windows client devices.
- Quotas and service level agreements to control access to resources.
- Custom hooks to extend the functionality of the host agent (VDSM).
- Configuring additional internal local user domains for Oracle Linux Virtualization Manager.
- External directory server apart from Active Directory and OpenLDAP Standard Schema (options 3 and 9 when you use the `ovirt-engine-extension-aaa-ldap-setup` command).
- Users log in automatically to the Administration Portal or VM Portal using the credentials obtained from a Kerberos or LDAP server (single sign-on).
- moVirt Android client for the Oracle Linux Virtualization Manager.

Compute

- CPU quality of service (QoS) entries to control the amount processing capability virtual machines can access on hosts.
- Trusted compute pools to deploy virtual machines on clusters that use Intel Trusted Execution Technology (Intel TXT).
- Existing Gluster storage cluster.
- GPU passthrough that attaches host GPUs directly to virtual machines.

Network

- Host network interfaces configured to use `ethtool` or Fibre Channel over Ethernet (FCoE) custom properties.
- PCI network cards attached directly to virtual machines (PCI passthrough or SR-IOV).
- Network bonds that use bond modes 0 (round-robin policy), 2 (XOR policy), 3 (broadcast policy), and 5 (adaptive transmit load balancing policy).
- Host network QoS (quality of service) entries to control the bandwidth a logical network uses on a physical interface.
- Virtual machine network QoS entries to control to the input/output of VNICs.
- External network providers, including the Open Virtual network (OVN) provider.
- Network security groups provided by the OpenStack Neutron service.
- Networks and interfaces managed with the Cisco Unified Computing System (UCS).

Storage

- Wipe after delete functionality to zero out used blocks in virtual disks.
- Storage quality of service (QoS) entries to control the maximum throughput and I/O operations for virtual disks in a storage domain.
- Disk profiles to define the maximum throughput and I/O for virtual disks in a storage domain.
- POSIX-compliant file system support to allow access to storage not exposed using NFS, iSCSI, or FCP.

Virtual Machines (Guests)

- Seal and automate the initialization of Windows guests using `sysprep`.
- Automatic installs of guest additions in Windows guests using the Application Provisioning Tool (APT).
- User log in automatically to Linux virtual machines using their Oracle Linux Virtualization Manager credentials (single sign-on).
- User log in automatically to virtual machines using the credentials obtained from a Kerberos or LDAP server (single sign-on).
- Host devices directly attached to virtual machines, including SCSI devices such as disks, PCI devices such as NICs and GPUs, and USB devices such as webcams (passthrough or SR-IOV).
- Virtual machine authentication using smart cards on client devices (smart card authentication).
- Multiple displays with virtual machines.
- Remote access to USB devices on Windows client devices using the `usbdk` driver (USB redirection).
- Virtual NUMA nodes configured on virtual machines and pinned them to NUMA nodes on a physical host.

- SAP monitoring in virtual machines enabled through the Administration Portal.

4

Deprecated Features

The following features are marked as deprecated in the upstream release and may be removed in a future release.

- Integration with GlusterFS, an open source scale-out distributed filesystem, to provide a hyperconverged infrastructure (HCI) cluster where both compute and storage are provided from the same hosts.
- The Oracle Enterprise Manager Virtualization Technology (VT) plug-in supporting Oracle Linux Virtualization Manager is deprecated. You can use the Grafana integration as an alternative monitoring and reporting solution. For more information, see Data Visualization with Grafana in the [Oracle Linux Virtualization Manager: Architecture and Planning Guide](#).
- The Log Collector Analyzer Tool (`ovirt-log-collector-analyzer`) is a command-line tool that analyzes and reports on the Oracle Linux Virtualization Manager environment.
- Export domains are temporary storage repositories that are used to copy and move images between data centers and Oracle Linux Virtualization Manager environments. Use data domains instead.
- ISO domains store ISO files which you can attach to virtual machines and use to install and boot operating systems and applications. Use data domains instead.
- The ISO Uploader Tool (`engine-iso-uploader`) is a command-line tool for uploading ISO images to an ISO storage domain. Use the Administration Portal or the REST API to upload ISO images to data domains instead.
- The `ovirt-shell` command line interface has not been updated to support any new features added to the upstream release since version 4.0. For automation purposes, either use the REST API or another tool such as Ansible.
- USB Clerk is a service that is able to install and uninstall USB drivers in Windows virtual machines. Use the `usbdk` driver instead.
- In the event notifications configuration file (`/usr/share/ovirt-engine/services/ovirt-engine-notifier/ovirt-engine-notifier.conf`), the `FAILED_QUERIES_NOTIFICATION_RECIPIENTS` variable is deprecated. Use the `FILTER` variable instead.

5

Known Issues

The following sections provide information about known issues in Oracle Linux Virtualization Manager. In addition, you should also review the following documents:

- Upstream [oVirt 4.5.5 Release Notes](#).
- Applicable Oracle Linux release notes in [Oracle Linux Documentation](#).

Oracle Linux Virtualization Manager Issues

Federated user details not displayed in the Administration Portal

When users are imported through Keycloak (for example, from Active Directory/LDAP user federation) and you navigate to **Administration > Users** in the Administration Portal, the **First Name**, **Last Name**, and **Email** fields might be blank.

Setup Scripts for Active Domain Restrict Users to Verify for Root AD Forest and Not a Specific Subdomain

When attempting to configure authentication using the `ovirt-engine-extension-aaa-ldap` extension, you can add only the forest domain for authentication. When attempting to use the subdomain, you see the following warning and error messages:

```
[ INFO ] Resolving Global Catalog SRV record for
subdomain1.mydomain.domain.local
[WARNING] Cannot resolve Global Catalog SRV record for
subdomain1.mydomain.domain.local.
Please check you have entered correct Active Directory forest name and check
that forest
is resolvable by your system DNS servers
[ ERROR ] Failed to execute stage 'Environment customization': Active
Directory forest is
not resolvable, please make sure you've entered correct forest name. If for
some reason you
can't use forest and you need some special configuration instead, please
refer to examples
directory provided by ovirt-engine-extension-aaa-ldap package.
```

Solution: Issue also occurs in the upstream oVirt release. There is no workaround for this behavior.

Bug: 29182338

Active Directory Profile Missing After Restoring the Engine

When restoring from a backup to a fresh installation of the Oracle Linux Virtualization Manager, the Active Directory profile is missing.

Solution: Perform the following steps:

1. Reinstall the `ovirt-engine-extension-aaa-ldap-setup` package:

```
dnf install ovirt-engine-extension-aaa-ldap-setup
```

2. Restart the Manager.

```
service ovirt-engine restart
```

3. Add the IP address of the Active Directory to the `/etc/host` and `/etc/resolve.conf` files, respectively.

Bug: 29410228

Removed Storage Domains Are Still Shown in the Dashboard

After you remove a storage domain, it is still shown in the Dashboard.

Solution: There is no workaround for this behavior.

Bug: 29494264

Unable to Set up a Connection to an OpenLDAP Directory

When you use `ovirt-engine-extension-aaa-ldap` extension to set up a connection to an OpenLDAP server (option 9 - OpenLDAP Standard Schema), the setup fails because you are unable to authenticate to the directory server.

Solution: There is no workaround for this behavior. Issue also occurs in the upstream oVirt release.

Bug: 29525988

Global Configuration Issues

Virtual Machine Started on KVM Host Whose Virtual Machine Count Exceeds the HighVirtualMachineCount Property Set for an Evenly_Distributed Scheduling Policy

A virtual machine started on a KVM host whose virtual machine count exceeded the number of virtual machines set in the **HighVMCount** property for an `Evenly_Distributed` scheduling policy. Based on the scheduling policy configured in this scenario, load balancing should have been triggered and this virtual machine should have started on another KVM host in the cluster.

Solution: There is no workaround for this behavior.

Bug: 29168788

Virtual Machine Can Be Started on a KVM Host Whose CPU Utilization Exceeds the HighUtilization Property Set for an Evenly_Distributed Scheduling Policy

In a 3-host cluster where only one host is active (the other two hosts are in `Maintenance` mode), 5 virtual machines are created by importing OVA files. An `Evenly_Distributed` scheduling policy is configured with the **HighUtilization** property set to **50**. When the CPU utilization exceed 50% on the KVM host and a virtual machine is started, the virtual machine should fail to startup; however, the virtual machine is starting up on the KVM host in this scenario.

Solution: There is no workaround for this behavior.

Bug: 29171712

CPU Load Not Evenly Load Balanced for an Evenly_Distributed Scheduling Policy

In a 3-host cluster where an Evenly_Distributed scheduling policy is configured with the **HighUtilization** property set to **50** and the **CPUOverCommitDuration** set to **1**, CPU load did not evenly distribute across the KVM hosts in the cluster. In this scenario, virtual machines did not migrate that should have been migrated due to load balancing based on the configured scheduling policy.

Solution: There is no workaround for this behavior.

Bug: 29172270

Power_Saving Scheduling Policy Not Shutting Down Any of the KVM Hosts in a Cluster with CPU Utilization Less Than 20%

In a cluster with 3 running KVM hosts and with 4 running virtual machines, a Power_Savings scheduling policy is configured with the **EnableAutomaticHostPowerManagement** property set to **true** when the CPU and memory is found to be low on the KVM hosts. After this policy is set, the KVM hosts are not being shutdown and the virtual machines are not being migrated, even though the CPU utilization is less than 20%. Given the configured Power_Savings scheduling policy for this scenario, some of the hosts should have been shutdown.

Solution: There is no workaround for this behavior.

Bug: 29418541

Virtual Machine Not Migrating After Exceeding the MaxFreeMemoryForOverUtilized Property Value for a Power_Savings Scheduling Policy

A virtual machine is observed not migrating to another KVM host in a cluster that has enough free memory when the virtual machine exceeds the value set for the **MaxFreeMemoryForOverUtilized** property of a Power_Savings scheduling policy.

Solution: There is no workaround for this behavior.

Bug: 29419399

MinFreeMemoryForUnderUtilized Property Not Working for Evenly_Distributed and Power_Savings Scheduling Policies

In a cluster with 3 active hosts where there are 4 running virtual machines (3 virtual machines running on one of the hosts and 1 virtual machine running on another one of the hosts), an Evenly_Distributed policy is configured with a value set for the **MinFreeMemoryForUnderUtilized** property. The virtual machines in this environment then exceed the **MinFreeMemoryForUnderUtilized** property value set for the policy, but neither the KVM hosts are shutdown nor are the virtual machines migrated in this scenario.

The policy is then changed to a Power_Savings scheduling policy and the **MinFreeMemoryForUnderUtilized** property is changed to the same value as previously set for the Evenly_Distributed scheduling policy, and again it is observed that neither the KVM hosts are shutdown nor the virtual machines are migrated when this property value is exceeded.

Solution: There is no workaround for this behavior.

Bug: 29425062

MacPoolAdmin Role Is Available Only for System-Level Users

Although the **MacPoolAdmin** role can be assigned to users of different levels (for example, **System**, **Data Center**, **Cluster**, and so on), only users who are given this role at the **System** level are actually able to perform **MacPoolAdmin** tasks on the Oracle Linux Virtualization Manager, such as creating, editing, or deleting MAC address pools.

Solution: If a user requires **MacPoolAdmin** privileges, ensure that the user is assigned the **MacPoolAdmin** role at the **System** level on the Manager.

Bug: 29534106

Compute Issues

Host Becomes Nonresponsive After Hardening pam.d

A KVM host becomes nonresponsive after appending the following line in the `/etc/pam.d/system-auth` file:

```
session required pam_lastlog.so showfailed
```

Solution: Append the following line in the `/etc/pam.d/system-auth` file:

```
session required pam_lastlog.so silent
```

The `silent` option does not inform you about any previous logins; however it does update the `/var/log/lastlog` file.

For additional information, refer to the [My Oracle Support \(MOS\)](#) article *Host becomes nonresponsive after hardening pam.d (Doc ID 2796734.1)*.

Bug: 33128719

KVM Host Under the Control of a Engine Host Can Be Accidentally Added to Another Manager Host Without Validation of Its Current State

When an KVM host is already deployed on a Manager host, you can add this KVM host to another Manager host, causing this Manager host to take the KVM host away from the original owning Manager host. Adding a KVM host that is already under the control of a Manager host is highly not recommended, but it can be done by accident. In this event, the KVM host that resides on the original owning Manager host changes to a status of `Nonresponsive` and all virtual machines running on it change to a status of `Unknown`. When you attempt to put the nonresponsive host into `Maintenance` mode, the following error message is generated:

```
Error while executing action: Cannot switch Host to Maintenance mode.  
Host still has running VMs on it and is in Non Responsive state.
```

Solution: For information about the workaround for this issue, refer to the related issue [Removing a Stolen, Defunct, Nonresponsive, or Destroyed KVM Host](#).

Bug: 29127349

Duplicate KVM Host Cannot Be Moved into Maintenance Mode After Host Addition Fails

If you add a KVM host that is running virtual machines to the Manager host and you attempt to add the same KVM host using a different host name (either IP address or FQDN), the Manager host fails to add this KVM host. When you try moving the failed KVM host into `Maintenance` mode to remove it from the Manager, it causes the KVM host to be stuck in `Preparing for Maintenance` mode.

Solution: Migrate all virtual machines that are `UP` on the running KVM host to a different KVM host in the cluster. Doing that changes the status on the failed KVM host to `Maintenance` mode. You can then remove the KVM host.

Bug: 29127707

SSH Connection Hangs When Adding a New KVM Host with Firewalld Disabled to the Engine

SSH Connection hangs when adding a new KVM host with `firewalld` disabled to the Manager.

Solution: There is no workaround for this behavior. The SSH connection is eventually restored after some delay.

Bug: 29135914

Bridge Is Not Cleaned When Undeploying and Removing a KVM Host

After a KVM host is undeployed and removed from the Manager host, the previous bridges are not cleaned on the KVM host.

Solution: Networking resources should be manually removed or cleaned after removing the KVM host from the Manager host, and the default management network, `ovirtmgmt`, should not be modified. To revert any other network configurations, run the following command:

```
vdsm-tool clear-nets --exclude-net ovirtmgmt
```

Bug: 29167000

Incomplete Error Message Stopping a KVM Host

This issue is seen with a KVM host whose **Status** is `Unassigned`. The steps to reproduce the issue are as follows:

1. Go to **Compute** and then click **Hosts**.
2. On the Hosts pane, click **Management** and from the drop-down list select **Stop**.

The following incomplete **Operation Canceled** error message is generated:

```
Error while executing action
```

Although this error message is incomplete, it is accurate; however, the message does not provide a workaround.

Solution: The **Status** of the KVM host must be `UP` before you can stop the KVM host. Try setting the KVM host to `Maintenance` mode and then stopping it. If this does not work, you must log in to the KVM host, resolve the issue, and bring it back up.

Bug: 29298704

Host Console Indicates KVM Host Is Registered on the Engine Host After Its Removal

The Host Console indicates that an KVM host is still registered on the Manager host after it has been removed from the Manager host. This issue is observed in the following scenario.

1. Remove the KVM host by using the Manager as follows:
 - a. On the **Hosts** pane, put the KVM host in `Maintenance` mode by clicking **Management** and then selecting `Maintenance` from the drop-down list.
 - b. Click **Installation** and then select **Reinstall** from the drop-down list.
 - c. Click **Hosted Engine** and then select **Undeploy** from the drop-down list.
 - d. Click **OK**.
 - e. Then set the KVM host to `Maintenance` mode again and click **Remove** to open the **Remove Host(s)** confirmation window and click **OK**.
2. Open the Host Console.
3. Go to **Virtualization** and then click **Hosted Engine**.

The **Host Engine Setup** screen indicates that the system is already registered to the removed KVM host.

Solution: There is no workaround for this behavior.

Bug: 29444179

Removing a Stolen, Defunct, Nonresponsive, or Destroyed KVM Host

A KVM host that is under the control of a Manager host can be taken over by another Manager host. See also [KVM Host Under the Control of a Engine Host Can Be Accidentally Added to Another Manager Host Without Validation of Its Current State](#).

For example, a KVM host that is already in use by a Manager host can be added as a **New Host** by a different Manager host. In this scenario, the original Manager is unable to communicate with its KVM host. Because the KVM host has a running a virtual machine, attempts to put the host into `Maintenance` mode are rejected with the following error message:

```
Host still has running VMs on it, and is Non responsive state
```

Solution: There are two available methods for handling this situation:

- Method to resolve the issue of a host being taken over by another Manager host.
- Method to remove a KVM host that has become permanently nonfunctional.

Method to Resolve the Issue of a Host Being Taking Over by Another Engine Host

When the new Manager host takes over the KVM host, the following symptoms are observed:

- On the original owning Manager host, the KVM host shows a status of `Connecting` and the virtual machines show up as running.
- On the new Manager host, copies of the virtual machines show up as `external_vm-name` and they are running.

Perform the following steps:

1. Power off the KVM host.

After powering off the KVM host, the following events occur:

- The KVM host eventually goes into a status of `NonResponsive` on both Manager hosts .
- The virtual machines go into a status of `Unknown` on both Manager hosts.
- The following events are triggered for this KVM host on both Manager hosts: `Handling non responsive Host host-name`.

2. Wait for these events to fail.

On both Manager hosts, perform the following steps:

1. Click **More Actions** and select **Confirm Host has been Rebooted** from the drop-down list.
2. Select the **Confirm Operation** checkbox and click **OK**.

This action causes the virtual machines to be marked as **Down** on the Managers and the host to display a value of 0 under the **Virtual Machines** column. The virtual machines may now be migrated by the old manager to another host if that is permitted by the **Migration Mode** settings of the virtual machines.

3. Set the KVM host to `Maintenance` mode.

Then, on the new Manager host, perform the following steps:

1. Remove the virtual machine copies.
2. Remove the KVM host.

When the KVM host comes back up, SSH to it and remove all authorized keys.

```
rm /root/.ssh/authorized_keys*
```

Then, on the original owning Manager host, reinstall the KVM host using password authentication (for simplicity).

Note

If the installation fails on the setup due to network issues causing the KVM host to go into a `NonResponsive` status on the Manager host, set the KVM host back to `Maintenance` mode and reinstall it with an SSH key.

Method to Remove a KVM Host That Has Become Permanently Nonfunctional

Assuming the Manager host shows the KVM host status as `Nonresponsive`, the following symptoms are observed:

- The virtual machines that were running on the KVM host show a status of `Unknown`.
- The following events are generated on the Manager host for this KVM host:

```
Handling non responsive Host host-name
```

These events eventually fail.

Perform the following steps:

1. Click **More Actions** and select **Confirm Host has been Rebooted** from the drop-down list.

2. Select the **Confirm Operation** checkbox and click **OK**.

This action causes the virtual machines to be shutdown and the host to display a value of 0 under the **Virtual Machines** column.

3. Set the KVM host to `Maintenance mode`.
4. Remove the KVM host.

For additional information, refer to the [My Oracle Support \(MOS\)](#) article *How to Remove a Defunct Host From oVirt Engine (Doc ID 2540819.1)*.

Bug: 29685904

Network Issues

MAC Address Displayed on the Engine for the Non-Primary Subordinate Port of a Bond Not Synchronized With KVM Host

After a bond is created on an KVM host with 2 subordinate ports, the MAC address of the non-primary subordinate port is changed to be the same as the primary subordinate port on the KVM host, but the Manager continues to display the previous MAC address of the subordinate port.

Solution: There is no workaround for this behavior.

Bug: 29049447

Engine Does Not Update the IP Address of a Virtual Machine When Switching the Virtual Machine Network for a Running Virtual Machine

In a scenario where a virtual machine is running with one VNIC interface connected to a VM network, it is observed that the Manager allows the user to change the VNIC interface for the running virtual machine to a different VM network. This action causes the virtual machine to lose its VM network and to become inaccessible (that is, its IP address is no longer pingable). The Manager, however, still displays its old IP address and does not generate any error messages.

Solution: There is no workaround for this behavior.

Bug: 29060999

Default Route Still Set to Yes After Default Route Role Is Removed from Connected Network

Changing the default route on a network that is attached to a host and that has virtual machines connected to that network is illegal and is rejected by VDSM. The Manager host gets notified about the request failure and logs the error, but the Manager is unaware of the failure; that is, the **Manage Networks** pane shows that default route change was successful.

Solution: When you modify the network configuration, make sure the interface configuration accurately reflects the most recent changes.

Bug: 29133844

Exception Generated When VLAN Network Is Added with Interface Name Longer Than 15 Characters

Adding an VLAN interface whose interface exceeded 15 characters generates the following failure message:

```
VDSErrorException: Failed to HostSetupNetworksVDS, error = Internal JSON-RPC error:
{'reason': '[Errno 8] Input data out of range'}, code = -32603
```

The VLAN ID ranges from 0 to 4096, which is a maximum of four characters and the total interface name has a limit of 15 characters. This is a Linux kernel limitation. The maximum interface name length is defined by the kernel headers and is a global limit, affecting all applications.

Solution: Limit interface names to 15 characters (including periods and VLAN IDs).

Bug: 29221397

Adding KVM Host Connected to VLAN Network Fails When VLAN Network Is Not Configured on Management Network

Adding a KVM host that is connected to a VLAN network fails when the `ovirtmgmt` management network is not configured on the VLAN network.

Solution: All servers that are part of a single cluster must be on the same VLAN management network.

Perform the following steps:

1. Go to **Network** and then click **Networks**.
2. Select the `ovirtmgmt` network under the **Name** column.
3. Click **Edit**.
4. Select the **Enable VLAN tagging** checkbox and enter the VLAN tag ID in the text entry field.

Make sure that the VLAN tag ID is the same as the VLAN ID set on the host interface to be connected to the `ovirtmgmt` management network.

5. Go to **Compute** and then click **Hosts**.
6. Click **New**.
7. On the **Hosts** pane, click on the KVM host.
8. Click on the **Network Interfaces** tab.
9. Click **Setup Host Networks** and add the `ovirtmgmt` management network to the appropriate interface.

Bug: 29245869

Network Synchronization Fails but Is Reported as Succeeding

After creating a VLAN network with a MTU value and adding the network to a host interface, the network is not synchronized because the host retains the default MTU value of 1500. When

clicking **Sync All Networks**, the Manager reports the following message, indicating that the network synchronization is completed:

```
Finished Synchronizing networks on host
```

The network, however, is still not synchronized.

Solution: Try clicking **Sync All Networks** again.

Bug: 29311422

KVM Host Becomes Non-Operational When a Bondport Connected to Virtual Machine Network Changed from an Unsupported Mode to a Supported Mode

When a bondport is created with a mode that is not supported for a VM network, the network can be connected to a non-VM network successfully; however, if the non-VM network is changed to VM network, the network becomes out of sync. Furthermore, if the user follows the instructions provided in the error message to fix the network issue, the KVM host becomes non-operational. This issue is caused by the network change (non-VM to VM network) combined with the bond mode change.

Note

No issues occur when performing these actions in reverse; that is, changing a VM network to a non-VM network and then changing the bond mode accordingly.

The steps to reproduce this issue are as follows:

1. Create a non-VM network and connect the network to a bondport mode in 0, 5, or 6.
2. Change the network mode to a VM Network by editing the network in the Manager and selecting the **VM Network** checkbox.

This action causes the network to become out of sync.

3. Click **Sync All Networks**.

The following error is generated: Network *name* is attached to bond *number* . VM networks cannot be attached to bonds in mode 0, 5 or 6.

4. Change the bond mode to a mode available for a VM network (1, 2, or 4) and click **OK**.

After the network setup window closes, the **Sync All Network** button appears again, indicating that the network is out of sync, and the host becomes non-operational.

Solution: You can click **Sync All Networks** again to bring the KVM host up. Alternatively, to bring the KVM host up, you can put this KVM host in `Maintenance` mode and activate it after the network is back in sync.

Bug: 29312752

Engine Does Not Check Connectivity When Removing an Interface

When you connect a host interface to a network, the Manager validates their connectivity before establishing the connection. However, when removing an interface, the Manager does not check their connectivity. This behavior is observed in a network configuration where the management network is already connected to a bondport with 3 or more subordinate ports and only 1 of the ports is in UP state and is able to communicate with the management network.

In this situation, when you attempt to remove the UP port (primary port) in the **Setup Host Networks** dialog box, the Manager generates the following **Operation Canceled** error message:

```
Error while executing action HostSetupNetworks: Could not connect to peer host
```

This error message, which is generated during the removal of the subordinate port, is referring to the new connection that cannot be established, even though the removal of the primary port has already occurred.

Solution: There is no workaround for this behavior. In this situation, you need to make sure the remaining communication between the bondport and the management network are not interrupted after the subordinate port is removed, otherwise the host loses connection with the Manager and becomes non-responsive.

Bug: 29338703

Network Scale Limitation When Adding Networks to a KVM Host

In an Oracle Linux environment you can add approximately 389 networks to a KVM host, after which network communication timeout errors start appearing in the Manager log file and no further networks can be added to the host.

Solution: There is no workaround for this behavior.

Bug: 29383782

Network Label Failure Generates Illegal Network Parameters Error When Performing Other Network Operations on the KVM Host

In a scenario where two VLANs are erroneously created with duplicate VLAN numbers, a warning in the form of an orange exclamation mark appears as expected in the Manager next to the duplicate VLAN on the **Setup Host** dialog box with the following error message:

```
Cannot have networks with duplicate vlan id on same interface.
```

After this network configuration error, attempting to perform other network operations on the KVM host generates the following **Operation Canceled** error message:

```
Error while executing action HostSetupNetworks: Illegal Network parameters
```

Solution: You must resolve the network label failure before other network operations can be performed.

Bug: 29424399

Cannot Modify a VNIC in a Running Virtual Machine if an MTU Value is Set

If a virtual machine is running, editing a VNIC that has an MTU value set fails with the following message:

```
Error while executing action Edit VM Interface properties: General Exception
```

The following message is displayed in the `/var/log/vdsm/vdsm.log` file:

```
libvirtError: Operation not supported: cannot modify MTU
```

Solution: Shut down the virtual machine before editing the VNIC. A known issue in the upstream oVirt release.

Bug: 29456945

VM Network Bridge Name Does Not Match VM Network Name

When you create a new VM network, VDSM creates a bridge named after the VM network and its associated configuration file (`/etc/sysconfig/network-scripts/ifcfg-vm_network`) on the Oracle Linux KVM hosts.

If the VM network name is longer than 15 characters, or it contains special characters or a space character, VDSM generates a name for the bridge and the configuration file in the format `"onXXXXXXXXXXXXX "` where `XXXXXXXXXXXXX` is the first 13 hexadecimal characters extracted from the UUID of the network.

This is a Linux kernel limitation, bridge names must be 15 characters or less, and must not contain any special characters.

Solution: Limit VM network names to 15 characters and avoid any special characters.

Bug: 29409851

Storage Issues

QEMU Direct LUN Disk Naming Changed in Virtual Machine After Upgrade

The QEMU naming convention for the direct LUNs changes after upgrading from 4.3.x to 4.4. This means that any applications or databases using those disks might not start up or will throw errors.

On 4.3.x, the disk names of the direct LUNs before the upgrade:

```
0QEMU_QEMU_HARDDISK_drive-ua_<UUID>
```

On 4.4.10, the disk names of the direct LUNs after the upgrade to 4.4.10:

```
0QEMU_QEMU_HARDDISK_<UUID>
```

Solution: Manually change the udev rule, for example:

```
KERNEL=="sd*", SUBSYSTEM=="block", ENV{DEVTYPE}=="disk", /
ENV{ID_SERIAL}=="0QEMU_QEMU_HARDDISK_9f02f5ed-2787-41d6-899c-#####", /
SYMLINK+="mapper/ORDATA01", OWNER="grid", GROUP="asmadmin", MODE="0660",
ENV{ASM_DISK}="RDATA01"
```

Bug: 35376968

Storage Must be on the Same Subnet as the KVM Hosts

In order to avoid issues with routing, storage must be located on the same subnet as the Oracle Linux KVM hosts that will use the storage.

Bug: 29220930

Decreasing the Storage Size Used in Storage Domains Is Not Recommended

Although the size of NFS Ext 3 and 4 filesystems can be decreased and Block Storage (iSCSI and Fibre Channel) provides the capability to decrease storage size by resizing the underlying LUN, decreasing the storage size that is used for a storage domain is potentially a dangerous operation and is highly not recommended.

Solution: Do not decrease the storage size used in storage domains in your virtualization environment.

Bug: 29285337

Manager Not Updating the Size of an iSCSI-Based Virtual Disk After ZFS Storage Resized

The Oracle Linux Virtualization Manager is not updating the size of an iSCSI-based virtual disk after the ZFS storage is resized. This issue is seen after creating an LUN, attaching the virtual disk to a running virtual machine, resizing the LUN, and restarting VDSM. After the refreshing the **Disks** pane, the iSCSI-based virtual disk size does not update accordingly. However, when checking the iSCSI disk on the SPM host, the size of the disk is correctly updated.

Solution: There is no workaround for this behavior.

Bug: 29370809

Virtual Machine Templates Can Be Imported with MAC Addresses Not in Range for the MAC Address Pool

virtual machine templates can be imported with MAC addresses that are out of range for the MAC address pool.

Solution: If there is a MAC address range restriction for your virtualization environment, the workaround is to manually change the MAC address to an address within the range of the MAC address pool before starting the virtual machine.

Bug: 29449334

Oracle Linux Virtualization Manager Does Not Prevent Adding a Direct LUN Disk Being Used by an Active Storage Domain

The Manager prevents creating a storage data domain on LUNs that are being used by direct LUN disks. The Manager, however, does not prevent adding a direct LUN disk that is being used by an active storage domain.

Solution: There is no workaround for this behavior. If you attempt to add a direct LUN disk that is being used by an active storage domain, the Manager prompts you to approve the operation with the following warning message: This message might be unrecoverable and destructive!

Bug: 29465439

Microsoft Windows 8 (and Later) Virtual Machines Disabling Third-Party Mini-Port Drivers

Microsoft Windows 8 (and later) virtual machines are disabling third-party mini-port drivers.

Solution: To bring back the third-party mini-port drivers, perform the following steps:

1. Start a virtual machine using the IDE interface.
2. Open the console for the virtual machine.
3. Open the Command Prompt in the Microsoft Windows virtual machine and set the machine to boot in Safe Mode.

```
bcdedit /set {current} safeboot minimal
```

4. Shutdown the virtual machine.
5. Change the storage type.
The virtual machine now starts normally.
6. Reset the `bcdedit` setting to boot the virtual machine in Normal Mode.

```
bcdedit /deletvalue {current} safeboot
```

7. Reboot the virtual machine.

Bug: 29472477

Manager Does Not Prevent Attaching VirtIO-SCSI Disks to OL6U10_X86 Virtual Machines

Manager does not prevent attaching VirtIO-SCSI disks to OL6U10_X86 virtual machines. The VirtIO-SCSI option in the Manager should be disabled for OL6U10_X86 virtual machines in the **Virtual Machine: Resource Allocation Settings**.

Solution: There is no workaround for this behavior. Do not attach VirtIO-SCSI disks to OL6U10_X86 virtual machines.

Bug: 29499061

Virtual Machine Issues

Snapshots of KVM Guests Using TPM Fail if Save Memory Option is Enabled

If you have virtual machines with UEFI and TPM (Trusted Platform Module) enabled **and** attempt to create snapshots using the **Save Memory** option enabled, you get an error:

```
Cannot create Snapshot. VM has a TPM device attached.
```

Solution: If you don't need the TPM functionality for the snapshot, consider temporarily disabling the TPM device before creating the snapshot or do not use the **Save Memory** option.

Bug: 37603737

Virtual Machine Does Not Boot with *Allow Privileged SCSI I/O* Option Enabled

If you have configured a virtual machine on Oracle Linux 7.7 with a physical LUN and the *Allow Privileged SCSI I/O* option enabled, the virtual machine does not boot and reports an error:

```
libvirtError: Requested operation is not valid: unpriv_sgio is not supported  
by this kernel
```

Solution: Do not use the *Allow Privileged SCSI I/O* option if your virtual machine has access to a physical LUN.

Bug: 31021044

Virtual Machines Not Migrated As Expected Through Virtual Machine Affinity and Host Affinity Rules Set to Active and Enforcing

In a 2-host cluster where 4 virtual machines are running on one of the hosts, a **Vm Affinity Rule** and a **Host Affinity Rule** are set on the Manager to **Active** and **Enforcing**. These rules, however, are not applied correctly because the virtual machines defined in the affinity group should have migrated to the other host; instead, the other KVM host is filtered out in this scenario.

Solution: Enabling the VM Affinity Rule or the Host Affinity Rule **Enforcing** option imposes strict limits on whether virtual machines can start or be migrated. With VM Affinity Rule **Enforcing** enabled (shown as **Hard** in the list of Affinity Groups), the system does not migrate a virtual machine to a host different from where the other virtual machines in its affinity group are running, even if attempting to satisfy a Host Affinity Rule that only permits the virtual machines to run on a host other than where they are presently running.

The workaround for both VM and Host affinity rules is to **not** enable enforcing unless all scenarios including migration scenarios are understood. The rules are not disabled if you do not select the **Enforcing** option; it changes the enforcement from **Hard** to **Soft**, permitting the system to temporarily break the rule as adjustments in virtual machine placement are made.

Bug: 29190112

Changing Watchdog Actions on a Running Virtual Machine Requires a Virtual Machine Restart for New Attribute to Take Effect

After enabling a watchdog device and setting the watchdog `action` attribute to `reset`, the virtual machine is restarted. Changing the watchdog `action` to a different attribute, such as `poweroff`, causes the **Pending Virtual Machine Changes** window to open with the following message:

```
Changes that require Virtual Machine restart: watchdog. The VM is then  
restarted.
```

The issue is that even though watchdog `action` is changed to `poweroff`, the attribute value did not take effect immediately. The virtual machine must first be restarted before the new attribute takes effect.

Solution: There is no workaround for this behavior. If a watchdog action is changed on a running virtual machine, you must restart the virtual machine for the attribute changes to take effect.

Bug: 29213956

Guest Agent Hooks Do Not Work If a Migration Policy Is Set to Legacy

The guest agent hooks do not work for migration policies that are set to **Legacy**.

Solution: This behavior is expected, as the guest hooks are enabled depending on the migration policy that is configured, and the guest agent hook mechanism is disabled for migration policies set to **Legacy**. To use guest agent hooks, you must use one of the other available migration policy types.

Bug: 29261746

Importing Several OVA Files Resulted in Only One Virtual Machine Being Imported

When importing virtual machines, multiple OVA files can be selected, but only one virtual machine is actually imported. If there is a limitation, the Manager should block attempts to import multiple OVA files.

Solution: There is no workaround for this behavior.

Bug: 29373795

High-Availability Virtual Machines Fail to Restart After Restarting KVM Host

High-availability virtual machines fail to restart after the KVM host is restarted on the Manager.

Solution: There is no workaround for this behavior. The KVM host eventually starts up several minutes later.

Bug: 29393228

Migration Mode Greyed Out When Setting Pinning-to-Host Setting for a Virtual Machine

This issue is observed in a 2-host cluster. The steps to reproduce the issue are as follows:

1. Edit a virtual machine by selecting the **Highly Available** checkbox on the **High Availability** tab, and click **OK**.
2. Edit the virtual machine and update the pinning-to-host settings by selecting **Specific Hosts** and adding the two hosts from the drop-down list on the **Host** tab, and click **OK**.

After doing completing these steps, the **Migration Mode** drop-down list is greyed out. This option should be available in this scenario.

Solution: Try changing the pinning-to-host option to **Any Host in Cluster** and then change the setting back to **Specific Hosts** again.

Bug: 29374044

Cannot Assign Virtual Machine from a Virtual Machine Pool: Manager Claims That the Virtual Machine Is Attached When It Is Not

After creating a virtual machine pool with 10 virtual machines, setting the **Maximum number of VMs per user** field to 10, and granting permissions to a user, the user correctly attaches some virtual machines. Later the user tries to attach more virtual machines, but the Manager generates the following error, indicating that virtual machines are already attached when they are not:

```
PoolName:  
Cannot attach VM to VM-Pool. VM-Pool is already attached to a User.
```

Cannot allocate and run VM from VM-Pool. There are no available VMs in the VM-Pool.

Solution: A known issue in the upstream oVirt release. There is no workaround for this behavior.

Bug: 29379502

Manager Does Not Generate a Warning Message to Prevent OVA Export from a Virtual Machine Whose Guest OS Is Installed on a Direct LUN Disk

The Manager does not generate an warning message that prevents OVA export from a virtual machine whose guest OS is installed on a direct LUN disk.

Solution: There is no workaround for this behavior.

Bug: 29432214

Manager Allows OVA Export from a Virtual Machine Whose Guest OS Is Installed on a Direct LUN Disk

The Manager allows OVA export from a virtual machine whose guest OS is installed on a direct LUN disk. The Manager should prevent attempts to export an OVA from a virtual machine whose guest OS is installed on a direct LUN disk image.

Solution: There is no workaround for this behavior.

Bug: 29432323

Manager Does Not Prevent Virtual Machine Shutdown When Creating a Snapshot

The Manager does not prevent a virtual machine from being shutdown when creating a snapshot, thereby causing the snapshot to hang. In this situation, the virtual machine cannot be started or removed and the creation of the snapshot cannot be stopped.

When inspecting the Manager log, the follow message is seen repeatedly:

```
[org.ovirt.engine.core.bl1.SerialChildCommandsExecutionCallback]
(EE-ManagedThreadFactory-engineScheduled-Thread-34)
[7363ae20-1f20-451c-8318-215f122fcc15] Command 'CreateSnapshotForVm' (id:
'b50fee33-204b-4fc0-a8e6-d40c7f58d485') waiting on child command id:
'3c843771-8b8f-4c07-9230-822bd2892481' type:'AddDisk' to complete
2019-03-08 03:55:22,992Z INFO
```

Solution: There is no workaround for this behavior.

Bug: 29457750

Engine Exception When Unplugging a NIC Connected to a Virtual Machine Network on a Running Virtual Machine

Unplugging a NIC that is connected to a VM network on a running virtual machine may fail on the Manager with the following **Operation Canceled** message:

```
Error while executing action Edit VM Interface properties: Failed to
deactivate
VM Network Interface.
```

When inspecting the Manager log, you may see the following EngineException:

```
EngineException: org.ovirt.engine.core.vdsbroker.vdsbroker.VDSErrorException:  
VDSGenericException: VDSErrorException: Failed to HotUnplugNicVDS, error =  
Timeout detaching <Interface name=vnet3, type=bridge, mac=00:21:f6:00:04:6b  
at 0x7f6d30377868>, code = 50 (Failed with error DEACTIVATE_NIC_FAILED and  
code 50)
```

Solution: Try unplugging the NIC again.

Bug: 29460927

Highly Available Virtual Machine Not Started on Other Host When Source Host Dies Unexpectedly

Highly available virtual machines are not started on the other host when the source host is powered off or dies unexpectedly.

Solution: If a host that has running highly available virtual machines is powered off or dies unexpectedly and the virtual machines do not have additional split-brain prevention mechanisms in place (such as Sanlock, VM leases, or VM resume), you may be required to use the **Confirm 'Host has been Rebooted'** option from the **More Actions** drop-down list on the **Hosts** pane before the Manager allows the highly available virtual machines on the host to be restarted on the other host.

Bug: 29491043

VNC Console Session for a Running Virtual Machine Closes During Virtual Machine Migration

If you have an VNC console open for an active virtual machine during virtual machine migration, the session closes. This behavior occurs when using either the native client or noVNC.

Solution: There is no workaround for this behavior.

Bug: 29491251

Migrate VM Dialog Box Cannot Be Closed While a Virtual Machine Is Being Started

The **Migrate VM** dialog box cannot be closed while a virtual machine is being started in the Manager.

Solution: Refresh the browser page.

Bug: 29498386

Unable to Perform CPU Hot Unplug with Default Machine Type (pc-i440fx-2.6) for Oracle Linux Virtual Machine

The CPU hot unplug feature on the Manager is not working with the default emulated machine type for Oracle Linux virtual machines (pc-i440fx-2.6).

Solution: To enable CPU hot unplug feature for a particular Oracle Linux virtual machine, you must perform the following steps:

1. Shutdown the virtual machine.
2. Change the Custom Emulated Machine to pc-i440fx-2.7 by performing the following steps in the Administration Portal:

- a. Go to **Compute** and click **Virtual Machines**.
 - b. Select the Oracle Linux virtual machine and click **Edit Virtual Machine**.
 - c. Click the **System** tab on the sidebar.
 - d. Click **Advanced Parameters** to expand the menu.
 - e. On the Custom Emulated Machine drop-down menu, select `pc-i440fx-2.7`.
 - f. Click **OK** to save your changes.
3. Restart the virtual machine.

Bug: 29517731

Cannot Import OVA Files That Contain More Than 1 Virtual Machine

If an Open Virtualization Appliance file contains more than one virtual machine, then the virtual machines in the file cannot be imported. The import begins and appears to import the first virtual machine, but then the import fails and no virtual machine are installed successfully.

Solution: Restrict Open Virtualization Appliance files to one virtual machine.

Bug: 30771759

6

Feedback and Support

If you have an Oracle Linux Premier support subscription or Oracle Premier Support for Systems, you can open a case with [My Oracle Support](#).

Provide the following information:

- Description of the problem, including the situation where the problem occurs, and its impact on your operation.
- Machine type, operating system release, browser type and version, locale and product release, including any patches you have applied, and other software that might be affecting the problem.
- Detailed steps on the method you have used, to reproduce the problem.
- Any error logs or core dumps.
- Any metrics from [Performance Co-Pilot logs](#).

Obtaining the Log Files

When requesting support, use the log collector tool to collect logs from across the virtualization environment. The tool generates an archive file you should include when submitting a Service Request (SR).

1. Install the log collector tool:

```
dnf install ovirt-log-collector
```

2. Log in as the `root` user and provide the administration credentials for the Manager.
3. Use `ovirt-log-collector -h` to displays usage information all valid options.
4. Perform the log collection on the Manager host:

```
ovirt-log-collector
```

If you don't specify parameters, the tool collect all logs from the Manager and its attached hosts, including all database logs unless you specify `--no-postgresql`.

5. Locate the archive file in `/tmp/logcollector`.

Important

Performance Co-Pilot (PCP) collects OS and network metrics that you can use to diagnose performance issues and that can help the support team to resolve issues. See [Performance Co-Pilot logs](#).