



## Red Hat Virtualization 4.3

### Technical Notes

Technical notes for Red Hat Virtualization 4.3 and associated packages



# Red Hat Virtualization 4.3 Technical Notes

---

Technical notes for Red Hat Virtualization 4.3 and associated packages

Red Hat Virtualization Documentation Team

Red Hat Customer Content Services

[rhev-docs@redhat.com](mailto:rhev-docs@redhat.com)

## Legal Notice

Copyright © 2020 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

XFS<sup>®</sup> is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack<sup>®</sup> Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## Abstract

The Technical Notes document provides information about changes made between release 4.2 and release 4.3 of Red Hat Virtualization. This document is intended to supplement the information contained in the text of the relevant errata advisories available through the Content Delivery Network.

---

## Table of Contents

<b>PREFACE</b> .....	<b>3</b>
<b>CHAPTER 1. RHEA-2019:1085 RED HAT VIRTUALIZATION MANAGER (OVIRT-ENGINE) 4.3 GA</b> .....	<b>4</b>
ovirt-engine	4
vdsm	14
<b>CHAPTER 2. RHBA-2019:1077 VDSM 4.3 GA</b> .....	<b>15</b>
vdsm	15



## PREFACE

These Technical Notes provide documentation of the changes made between release 4.2 and release 4.3 of Red Hat Virtualization. They are intended to supplement the information contained in the text of the relevant errata advisories available through the Content Delivery Network. Red Hat Virtualization errata advisories are available at <https://access.redhat.com/errata/>.

A more concise summary of the features added in Red Hat Virtualization 4.3 is available in the Red Hat Virtualization 4.3 Release Notes.

# CHAPTER 1. RHEA-2019:1085 RED HAT VIRTUALIZATION MANAGER (OVIRT-ENGINE) 4.3 GA

The bugs in this chapter are addressed by advisory RHEA-2019:1085. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2019:1085>.

## ovirt-engine

### BZ#1571024

This feature provides the ability to enable live migration for HP VMs (and, in general, to all VM types with pinning settings).

Previously, Red Hat Virtualization 4.2 added a new High-Performance VM profile type. This required configuration settings including pinning the VM to a host based on the host-specific configuration. Due to the pinning settings, the migration option for the HP VM type was automatically forced to be disabled.

Now, Red Hat Virtualization 4.3 provides the ability for live migration of HP VMs (and all other VMs with a pinned configuration like NUMA pinning, CPU pinning, and CPU passthrough enabled). For more details, see the feature page:

<https://ovirt.org/develop/release-management/features/virt/high-performance-vm-migration.html>

### BZ#1619474

This release ensures that if a request occurs to disable I/O threads of a running VM, the I/O threads disable when the VM goes down.

### BZ#1626907

This release prevents VM snapshot creation when the VM is in a non-responding state to preclude database corruption due to an inconsistent image structure.

### BZ#1676461

updated by engine-setup. If an error occurs, engine-setup treats this as a failure and tries to rollback, which is a risky process. To work around this scenario, the package `ovirt-engine-setup-plugin-ovirt-engine` now requires `ovirt-vmconsole 1.0.7-1`. Updating the setup packages with yum should also update `ovirt-vmconsole`. If an error occurs, yum evaluates it as a non-fatal error. See also bug 1665197 for the actual error from `ovirt-vmconsole`.

### BZ#1155676

In this release, users can now export VM templates to OVA files located on shared storage, and import the OVA files from the shared storage into a different data center.

### BZ#1403183

This release ensures that virtual machines with file-based storage created from a template where the Resource Allocation > Storage Allocation > Clone > Format setting is set to Raw results in virtual machines having an Allocation Policy set to "Preallocated."

### BZ#1513398



In the Administration Portal, searching for virtual machines by network label, VM emulated machine, and CPU type are not supported due to the complexity of their implementation.

**BZ#1532969**

Previously, virtual machines could only boot from BIOS. The current release adds support for booting virtual machines via UEFI firmware, a free, newer, more modern way to initialize a system.

**BZ#1550634**

This release removes the Red Hat Virtualization Manager support for clusters levels 3.6 and 4.0. Customers must upgrade their data centers to Red Hat Virtualization Manager 4.1 or later before upgrading to Red Hat Virtualization Manager 4.3.

**BZ#1563271**

An Ansible role, `ovirt-host-deploy-spice-encryption`, has been added to change the cypher string for SPICE consoles. The default cypher string satisfies FIPS requirements ('`TLSv1.2+FIPS:kRSA+FIPS:!eNULL:!aNULL`'). The role can be customized with the Ansible variable `host_deploy_spice_cipher_string`.

**BZ#1574771**

This release enables multiple queues and creates up to four queues per vNIC depending on the number of available vCPUs.

**BZ#1583009**

Previously, during high CPU usage, the balancing process would migrate a single virtual machine that evaluated to a good migration candidate. Now, this enhancement updates the balancing process to migrate multiple virtual machines one-by-one until one of the virtual machine migrations succeeds.

**BZ#1588498**

With this release, the size of the `rhvm` package has been reduced.

**BZ#1590202**

This release adds a feature to control toast notifications. Once any notifications are showing, "Dismiss" and "Do not disturb" buttons will appear that allow the user to silence notifications.

**BZ#1594615**

The release improves upon the fix in BZ#1518253 to allow for a faster abort process and a more easily understood error message.

**BZ#1597085**

A new option has been added to the Administration Portal under Compute > Clusters in the Console configuration screen: Enable VNC Encryption

**BZ#1598391**

Red Hat OpenStack Platform 14's OVN+neutron is now certified as an external network provider for Red Hat Virtualization 4.3.

**BZ#1599321**

There are inconsistencies in the following internal configuration options:

- HotPlugCpuSupported
- HotUnplugCpuSupported
- HotPlugMemorySupported
- HotUnplugMemorySupported
- IsMigrationSupported
- IsMemorySnapshotSupported
- IsSuspendSupported
- ClusterRequiredRngSourcesDefault

Systems that have upgraded from RHV 4.0 to RHV 4.1/4.2 and are experiencing problems with these features should upgrade to RHV 4.2.5 or later.

**BZ#1609139**

Previously, you could only assign one vGPU device type (`mdev_type`) to a virtual machine in the Administration Portal. The current release adds support for assigning multiple Nvidia vGPU device types to a single virtual machine.

**BZ#1619866**

This release ensures that if a request occurs to disable I/O threads of a running VM, the I/O threads setting remains disabled when changing unrelated properties of a running VM.

**BZ#1632055**

This release updates the Red Hat Virtualization Manager power saving policy to allow VM migration from over-utilized hosts to under-utilized hosts to ensure proper balancing.

**BZ#1635405**

This release adds a log entry at the WARN level if an attempt is made to move a disk with a damaged ancestor. A workaround solution is to leverage the REST API to move the disk between storage domains.

**BZ#1638124**

This release enables VM configuration with memory greater than two terabytes.

**BZ#1643512**

This release ensures that red exclamation point appears when a bond is misconfigured.

#### BZ#1646861

This release ensures Red Hat Virtualization Manager sets the recommended options during the creation of a volume from Red Hat Virtualization Manager to distinguish creating volumes from the Cockpit User Interface.

#### BZ#1651225

{enterprise-linux} 8 is fully supported as a guest operating system. Note that GNOME single sign-on functionality, guest application list, and guest-side hooks are not supported.

#### BZ#1655375

Previously, after upgrading to version 4.2 or 4.3, the Compute > Hosts > Network Interfaces page in the Administration Portal did not display host interfaces. Instead, it would throw the following obfuscated exception several times: `webadmin-0.js:formatted:176788 Mon Dec 03 11:46:02 GMT+1000 2018 SEVERE: Uncaught exception com.google.gwt.core.client.JavaScriptException: (TypeError) : Cannot read property 'a' of null`

The current release fixes this issue.

#### BZ#1668720

This release ensures the process to provision of a virtual machine from a template completes correctly.

#### BZ#1009608

This release allows you to limit east-west traffic of VMs, to enable traffic only between the VM and a gateway. The new filter 'clean-traffic-gateway' has been added to libvirt. With a parameter called GATEWAY\_MAC, a user can specify the MAC address of the gateway that is allowed to communicate with the VM and vice versa. Note that users can specify multiple GATEWAY\_MACs. There are two possible configurations of VM:

- 1) A VM with a static IP. This is the recommended setup. It is also recommended to set the parameter CTRL\_IP\_LEARNING to 'none'. Any other value will result in a leak of initial traffic. This is caused by libvirt's learning mechanism (see <https://libvirt.org/formatnwfilter.html#nwfelemsRulesAdvIPAddrDetection> and [https://bugzilla.redhat.com/show\\_bug.cgi?id=1647944](https://bugzilla.redhat.com/show_bug.cgi?id=1647944) for more details).
- 2) A VM with DHCP. DHCP is working partially. It is not usable in production currently ([https://bugzilla.redhat.com/show\\_bug.cgi?id=1651499](https://bugzilla.redhat.com/show_bug.cgi?id=1651499)).

The filter has a general issue with ARP leak ([https://bugzilla.redhat.com/show\\_bug.cgi?id=1651467](https://bugzilla.redhat.com/show_bug.cgi?id=1651467)). Peer VMs are able to see that the VM using this feature exists (in their arp table), but are not able to contact the VM, as the traffic from peers is still blocked by the filter.

#### BZ#1284775

This release adds support for memory hot-plug for IBM POWER (ppc64le) VMs.

**BZ#1388098**

The current release provides a software hook for the Manager to disable restarting hosts following an outage. For example, this capability would help prevent thermal damage to hardware following an HVAC failure.

**BZ#1441741**

In the current release, the v4 API documentation shows how to retrieve the IP addresses of a virtual machine.

**BZ#1454673**

When renaming a running virtual machine, the new name is now applied immediately, even when the QEMU process is running and is set with the previous name. In this case, the user is provided with a warning that indicates that the running instance of the virtual machine uses the previous name.

**BZ#1497355**

This release ensures the live storage migration process completes properly after creating a snapshot.

**BZ#1520848**

This release updates the VM video RAM settings to ensure enough RAM is present for any Linux guest operating system.

**BZ#1527860**

This release adds USB qemu-xhci controller support to SPICE consoles, for Q35 chipset support. Red Hat Virtualization now expects that when a BIOS type using the Q35 chipset is chosen, and USB is enabled, that the USB controller will be qemu-xhci.

**BZ#1553425**

This release ensures that the number of virtual machines configured to pre-start in a virtual machine pool start after editing an existing virtual machine pool.

**BZ#1559694**

If a VM does not use virtual NUMA nodes, it is better if its whole memory can fit into a single NUMA node on the host. Otherwise, there may be some performance overhead. There are two additions in this RFE:

1. A new warning message is shown in the audit log if a VM is run on a host where its memory cannot fit into a single host NUMA node.
2. A new policy unit is added to the scheduler: 'Fit VM to single host NUMA node'. When starting a VM, this policy prefers hosts where the VM can fit into a single NUMA node. This unit is not active by

default, because it can cause undesired edge cases. For example, the policy unit would cause the following behavior when starting multiple VMs:

In the following setup:

- 9 hosts with 16 GB per NUMA node
- 1 host with 4 GB per NUMA node

When multiple VMs with 6 GB of memory are scheduled, the scheduling unit would prevent them from starting on the host with 4 GB per NUMA node, no matter how overloaded the other hosts are. It would use the last host only when all the others do not have enough free memory to run the VM.

#### **BZ#1580128**

In the Administration Portal, the General subtab in Storage now displays the number of images on the storage domain with the label of "Images." This number corresponds to the number of logical volumes on a block domain.

#### **BZ#1611889**

This feature allows the user to select the cloud-init protocol with which to create a virtual machine's network configuration. The protocol can be selected while creating or editing a VM, or while starting a VM with Run Once. In older versions of cloud-init, backward compatibility needed to be maintained with the ENI protocol, whereas on newer cloud-init versions the OpenStack-Metadata protocol is supported.

#### **BZ#1622068**

Previously, after importing a guest from an ova file, the Import Virtual Machine dialog displayed the network type as "Dual-mode rt8319, VirtIO", when it should have been only "VirtIO". The current release fixes this issue.

#### **BZ#1624069**

This release supports custom Bond Naming in Red Hat Virtualization to include names with up to fifteen printable ASCII characters.

#### **BZ#1625612**

This release adds support for importing VMware virtual machines that include snapshots.

#### **BZ#1630243**

During virtual machine live migration, the migration progress bar is now also shown in the host's Virtual Machine tab.

#### **BZ#1631360**

The previous release changed the system manufacturer of virtual machines from "Red Hat" to "oVirt". This was inconsistent with preceding versions. Some users depended on this field to determine the underlying hypervisor. The current release fixes this issue by setting the SMBIOS manufacturer according to the product being used, which is indicated by the 'OriginType' configuration value. As a result, the manufacturer is set to 'oVirt' when oVirt is being used, and 'Red Hat' when Red Hat Virtualization is being used.

-

#### **BZ#1631392**

Previously, in the Administration Portal, the "New Pool" window uses the "Prestarted" label while the "Edit Pool" window uses the "Prestarted VMs" label. Both of these labels refer to the number of VMs prestarted in the pool. The current release fixes this issue.

#### **BZ#1635942**

Previously, while cloning a virtual machine with a Direct LUN attached, the Administration Portal showed the clone task as red (failed). The current release fixes this issue and displays the clone task as running until it is complete.

#### **BZ#1637015**

With this release, users can now disable pop-up notifications. When a pop-up notification appears in the Administration Portal, the following options are now available for disabling notifications:

- Dismiss All
- Do Not Disturb
- for 10 minutes
- for 1 hour
- for 1 day
- until Next Log In

#### **BZ#1640977**

This release ensures that all values for Quality of Service links are visible.

#### **BZ#1657134**

This release allows a storage domain to be created without a description while using the REST API.

#### **BZ#1657977**

Previously, the "Multi Queues enabled" checkbox was missing from the New- or Edit Instance Types window in the Administration Portal. The current release fixes this issue.

#### **BZ#1663616**

In this release, redirection device types are no longer set to unplugged and can now obtain the proper address from the domain xml when supported or from the host when they are not supported.

#### **BZ#1664342**

The sorting order in the list of Disks in the Storage tab of the Administration Portal was sorted alphabetically by text values in the Creation Date, instead of by time stamp. In this release, the list is now sorted by the time stamp.

#### **BZ#1664490**

This release enhancement preserves a virtual machine's time zone setting of a virtual machine when moving the virtual machine from one cluster to a different cluster.

**BZ#1664540**

A user with a UserRole or a role with a Change CD permit can now change CDs on running VMs in the VM Portal

**BZ#1666958**

This release ensures the SR-IOV vNIC profile does not undergo an invalid update while the vNIC is plugged in and running on the VM during the validation process. To update the SR-IOV vNIC profile, unplug the vNIC from the VM. After the updates are complete, replug the vNIC into the VM.

**BZ#1669367**

This fix ensures that the current propagate\_errors setting does not get reset when changing the disk properties.

**BZ#1676822**

Previously, while testing a RHEL 8 build of the virt-v2v daemon that turns a Red Hat Virtualization Host into a conversion host for CloudForms migration, you could not update the network profile of a running virtual machine guest. The current release fixes this issue.

**BZ#1511697**

Previously, an administrator with the `ClusterAdmin` role was able to modify the self-hosted engine virtual machine, which could cause damage. In the current release, only a `SuperUser` can modify a self-hosted engine and its storage domain.

**BZ#1560132**

In the Administration Portal, it is possible to set a threshold for cluster level monitoring as a percentage or an absolute value, for example, 95% or 2048 MB. When usage exceeds 95% or free memory falls below 2048 MB, a "high memory usage" or "low memory available" event is logged. This reduces log clutter for clusters with large (1.5 TB) amounts of memory.

**BZ#1561539**

A new option, Activate Host After Install, has been added to the Administration Portal under Compute > Hosts, in the New Host or Edit Host screen. This option is selected by default.

**BZ#1570040**

This release adds support for external OpenID Connect authentication using Keycloak in both the user interface and the REST API.

**BZ#1586126**

This release ensures that hosts can be set to maintenance mode after upgrading Red Hat Virtualization from 4.1 to 4.2.3.

**BZ#1595489**

This release ensures that VMs existing in Red Hat Virtualization Manager version 4.2.3 or earlier do not lose their CD-ROM device if the VMs are restarted in 4.2.3 or later versions.

**BZ#1635304**

This release allows users in Red Hat Virtualization Manager to view the full path of the host group in the host group drop-down list to facilitate host group configuration.

**BZ#1636331**

Previously, trying to update a disk attribute using the `/api/disks/{disk_id}` API failed without an error. The current release fixes this issue.

**BZ#1643486**

This release ensures the value of the migration bandwidth limit is correct.

**BZ#1645007**

Previously, making an API call to the foreman (hosts, hostgroups, compute resources) returned only 20 entries. The current release fixes this issue and displays all of the entries.

**BZ#1647928**

This release ensures the ``setupnetworks`` REST API can remove or modify an existing network attachment.

**BZ#1651140**

Red Hat Virtualization Manager now requires JBoss Enterprise Application Platform.

**BZ#1651255**

You can now set the number of IO threads in the new/edit VM dialog in the Administration Portal, instead of just the REST API.

**BZ#1654442**

There was a bug in the REST API for non-administrator users related to VNIC Profiles. Consequently, an error message appeared saying "GET\_ALL\_VNIC\_PROFILES failed query execution failed due to insufficient permissions." The code was fixed and the error no longer occurs.

**BZ#1655659**



This release ensures the correct parsing of the `rhv-toossetup_x.x.x.iso` file.

#### BZ#1655911

In this release, the following changes have been made in the view filters for VMs in the Administration Portal under Compute > Hosts > selected host:

New view filter names:

- From "Running on host" to "Running on current host" (default view)
- From "Pinned to host" to "Pinned to current host"
- From "All" to "Both" - when "Both" is selected, a new column named "Attachment to current host" is displayed to indicate that the VM is: "Running on current host", "Pinned to current host", or "Pinned and Running on current host".

#### BZ#1667842

Previously, the background process to migrate virtual machines considered affinity groups. This release updates the background process to migrate virtual machines to consider both affinity groups and affinity labels.

#### BZ#1536397

Previously, CloudInit passed the `dns_search` value incorrectly as the `dns_nameserver` value. For example, after configuring a the Networks settings of a virtual machine and running it, the `dns_search` value showed up in the `resolv.conf` file as the `dns_nameserver` value. The current release fixes this issue.

#### BZ#1540921

Conroe and Penryn CPU types are no longer supported. They will not appear as options for Compatibility Version 4.3, and a warning is displayed for older versions.

#### BZ#1600788

This release provides a check to evaluate self-hosted engine volumes prior to deleting the self-hosted engine volumes.

#### BZ#1602968

Previously, "Power Off" was missing from the virtual machine context menu in the Administration Portal; although it was present in previous versions, it was removed as part of the new user interface in 4.2. Now, "Power Off" is once again present when a running virtual machine is right-clicked.

#### BZ#1635845

This release ensures the clearing of the VM uptime during a guest operating system reboot, and the uptime that does display corresponds to the guest operating system.

#### BZ#1636256

In previous versions, it was not possible to limit the number of simultaneous sessions for each user,

so active sessions could significantly grow up until they expired. Now, Red Hat Virtualization Manager 4.3 introduces the `ENGINE_MAX_USER_SESSIONS` option, which can limit simultaneous sessions per user. The default value is `-1` and allows unlimited sessions per user.

To limit the number of simultaneous sessions per user, create the `99-limit-user-sessions.conf` file in `/etc/ovirt-engine/engine.conf.d` and add `ENGINE_MAX_USER_SESSIONS=NNN`, where `NNN` is the maximum number of allowed simultaneous sessions per user. Save and restart using: `systemctl restart ovirt-engine`.

#### **BZ#1650422**

Red Hat Virtualization Manager no longer logs messages regarding non-preferred host penalizations if the VM is not configured to have a preferred host.

#### **vdsm**

#### **BZ#1625543**

When Importing KVM VMs and Sparseness is specified, the actual Disk Size should be preserved to improve the performance of the Import as well as to conserve disk space on the Destination Storage Domain. Previously, when you set thin provisioning for importing a KVM-based VM into a Red Hat Virtualization environment, the disk size of the VM within the Red Hat Virtualization storage domain was inflated to the volume size or larger, even when the original KVM-based VM was much smaller. KVM Sparseness is now supported so that when you import a virtual machine with thin provisioning enabled into a Red Hat Virtualization environment, the disk size of the original virtual machine image is preserved. However, KVM Sparseness is not supported for Block Storage Domains.

## CHAPTER 2. RHBA-2019:1077 VDSM 4.3 GA

The bugs in this chapter are addressed by advisory RHBA-2019:1077. Further information about this advisory is available at <https://access.redhat.com/errata/RHBA-2019:1077>.

### vdsm

#### BZ#1593568

Previously, if a CD-ROM was ejected from a virtual machine and VDSM was fenced or restarted, the virtual machine became unresponsive and/or the Manager reported its status as "Unknown." In the current release, a virtual machine with an ejected CD-ROM recovers after restarting VDSM.

#### BZ#1111783

In the current release, Windows clustering is supported for directly attached LUNs and shared disks.

#### BZ#1111784

The current release supports Windows clustering for directly attached LUNs and shared disks.

#### BZ#1561033

The current release adds AMD SMT-awareness to VDSM and RHV-M. This change helps meet the constraints of schedulers and software that are licensed per-core. It also improves cache coherency for VMs by presenting a more accurate view of the CPU topology. As a result, SMT works as expected on AMD CPUs.

#### BZ#1614430

Vdsm-gluster tries to run heal operations on all volumes. Previously, if the gluster commands got stuck, VDSM started waiting indefinitely for them, exhausting threads, until it timed-out. Then it stopped communicating with the Manager and went offline. The current release adds a timeout to the gluster heal info command so the command terminates within a set timeout and threads do not become exhausted. On timeout, the system issues a GlusterCommandTimeoutException, which causes the command to exit and notifies the Manager. As a result, VDSM threads are not stuck, and VDSM does not go offline.

#### BZ#1617745

Previously, when a migrating virtual machine was not properly set up on the destination host, it could still start there under certain circumstances, then run unnoticed and without VDSM supervision. This situation sometimes resulted in split-brain. Now migration is always prevented from starting if the virtual machine set up fails on the destination host.

#### BZ#1654417

Previously, if an xlease volume was corrupted, VDSM could not acquire leases and features like high-availability virtual machines did not work. The current release adds rebuild-xleases and format-xleases commands to the VDSM tool. Administrators can use these commands to rebuild or format corrupted xlease volumes.

**BZ#1655115**

The current release removes the VDSM daemon's support for cluster levels 3.6/4.0 and Red Hat Virtualization Manager 3.6/4.0. This means that VDSM from RHV 4.3 cannot be used with the Manager from RHV 3.6/4.0. To use the new version of VDSM, upgrade the Manager to version 4.1 or later.

**BZ#1662449**

If a user with an invalid sudo configuration uses sudo to run commands, sudo appends a "last login" message to the command output. When this happens, VDSM fails to run lvm commands. Previously, the VDSM log did not contain helpful information about what caused those failures.

The current release improves error handling in the VDSM code running lvm commands. Now, if VDSM fails, an error message clearly states that there was invalid output from the lvm commands, and shows the output added by sudo. Although this change does not fix the root cause, an invalid sudo configuration, it makes it easier to understand the issue.

**BZ#1510336**

This release adds the ability to manage the MTU of VM networks in a centralized way, enabling oVirt to manage MTU all the way from the host network to the guest in the VM. This feature allows for the consistent use of MTUs in logical networks with small MTU (e.g., tunneled networks) and large MTU (e.g., jumbo frames) in VMs, even without DHCP.

**BZ#1510856**

Making large snapshots and other abnormal events can pause virtual machines, impacting their system time, and other functions, such as timestamps. The current release provides Guest Time Synchronization, which, after a snapshot is created and the virtual machine is un-paused, uses VDSM and the guest agent to synchronize the system time of the virtual machine with that of the host. The `time_sync_snapshot_enable` option enables synchronization for snapshots. The `time_sync_cont_enable` option enables synchronization for abnormal events that may pause virtual machines. By default, these features are disabled for backward compatibility.

**BZ#1511891**

Previously, copying volumes to preallocated disks was slower than necessary and did not make optimal use of available network resources. In the current release, qemu-img uses out-of-order writing to improve the speed of write operations by up to six times. These operations include importing, moving, and copying large disks to preallocated storage.

**BZ#1560460**

Previously, VDSM used `stat()` to implement `islink()` checks when using `ioprocess` to run commands. As a result, if a user or storage system created a recursive symbolic link inside the ISO storage domain, VDSM failed to report file information. In the current release, VDSM uses `lstat()` to implement `islink()` so it can report file information from recursive symbolic links.

**BZ#1575777**

Previously, a floppy drive in a virtual machine could prevent the virtual machine from being imported. In the current release, floppy drives are ignored during import.

**BZ#1625591**

Previously, after importing and removing a Kernel-based Virtual Machine (KVM), trying to re-import the same virtual machine failed with a "Job ID already exists" error. The current release deletes completed import jobs from the VDSM. You can re-import a virtual machine without encountering the same error.

**BZ#1636254**

VDSM uses Ildpad. Due to a bug, Ildpad confuses NetXtreme II BCM57810 FCoE-enabled cards. When the VDSM configuration enables Ildpad to read Ildp data from the card, it renders the card unusable. To work around this issue, set `enable_ildp=false` in `vdsm.conf.d` and restart VDSM. Check that Ildpad is disabled on all relevant interfaces by entering the command, `"lldptool get-ildp -i $ifname adminStatus"`. If Ildp is enabled, disable it by entering `"lldptool set-ildp -i $ifname adminStatus=disabled"`. After ensuring that Ildp support is disabled in VDSM, networking should be unaffected.

**BZ#1514004**

The TLSv1 and TLSv1.1 protocols are no longer secure. In the current release, they have been forcefully disabled in the VDSM configuration and cannot be enabled. Only TLSv1.2 and higher versions of the protocol are enabled. The exact version enabled depends on the underlying OpenSSL version.

**BZ#1585008**

The current release adds a new `'ssl_ciphers'` option to VDSM, which enables you to configure available ciphers for encrypted connections (for example, between the Manager and VDSM, or between VDSM and VDSM). The values this option uses conform to the OpenSSL standard. For more information, see <https://access.redhat.com/articles/4056301>

**BZ#1589612**

When a virtual machine starts, VDSM uses the domain metadata section to store data which is required to configure a virtual machine but which is not adequately represented by the standard libvirt domain.

Previously, VDSM stored drive IO tune settings in this metadata that were redundant because they already had proper representation in the libvirt domain. Furthermore, if IO tune settings were enabled, a bug in storing the IO tune settings prevented the virtual machine from starting. The current release removes the redundant information from the domain metadata and fixes the bug that prevented virtual machines from starting.

**BZ#1647607**

Previously, an incorrectly named USB3 controller, `"qemu_xhci,"` prevented virtual machines from booting if they used a host passthrough with this controller. The current release corrects the controller name to `"qemu-xhci,"` which resolves the booting issue.

