



# Red Hat Virtualization 4.1

## Release Notes

Release notes for Red Hat Virtualization 4.1



# Red Hat Virtualization 4.1 Release Notes

---

Release notes for Red Hat Virtualization 4.1

Red Hat Virtualization Documentation Team  
Red Hat Customer Content Services  
[rhev-docs@redhat.com](mailto:rhev-docs@redhat.com)

## Legal Notice

Copyright © 2018 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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

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

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

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

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

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

The OpenStack ® 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 Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Virtualization 4.1.

---

## Table of Contents

<b>CHAPTER 1. INTRODUCTION</b> .....	<b>3</b>
1.1. INTRODUCTION TO RED HAT VIRTUALIZATION	3
1.2. SUBSCRIPTIONS	3
<b>CHAPTER 2. RHV FOR IBM POWER</b> .....	<b>7</b>
<b>CHAPTER 3. RELEASE INFORMATION</b> .....	<b>10</b>
3.1. RED HAT VIRTUALIZATION 4.1 GA	10
3.2. RED HAT VIRTUALIZATION BATCH UPDATE 1 (OVIRT-4.1.2)	37
3.3. RED HAT VIRTUALIZATION BATCH UPDATE 2 (OVIRT-4.1.3)	39
3.4. RED HAT VIRTUALIZATION BATCH UPDATE 3 (OVIRT-4.1.4)	43
3.5. RED HAT VIRTUALIZATION BATCH UPDATE 4 (OVIRT-4.1.5)	44
3.6. RED HAT VIRTUALIZATION BATCH UPDATE 5 (OVIRT-4.1.6)	46
3.7. RED HAT VIRTUALIZATION BATCH UPDATE 6 (OVIRT-4.1.7)	51
3.8. RED HAT VIRTUALIZATION BATCH UPDATE 7 (OVIRT-4.1.8)	52
3.9. RED HAT VIRTUALIZATION BATCH UPDATE 8 (OVIRT-4.1.9)	53
3.10. RED HAT VIRTUALIZATION BATCH UPDATE 9 (OVIRT-4.1.10)	55
3.11. RED HAT VIRTUALIZATION BATCH UPDATE 10 (OVIRT-4.1.11)	55



# CHAPTER 1. INTRODUCTION

## 1.1. INTRODUCTION TO RED HAT VIRTUALIZATION

Red Hat Virtualization is an enterprise-grade server and desktop virtualization platform built on Red Hat Enterprise Linux. There are three core components to a Red Hat Virtualization environment:

### Red Hat Virtualization Manager

A server that manages and provides access to the resources in the environment. This server can be a physical machine or a virtual machine. Moreover, the Manager can be set up as a virtual machine on a host in an environment that the Manager itself manages. This setup is known as a self-hosted engine setup.

### Storage

Storage is used to store the data associated with virtual machines in the environment such as the virtual machines themselves and ISO files that can be used to install operating systems or other software on those virtual machines. Storage can be implemented on several storage types, including NFS, iSCSI, GlusterFS, and Fibre Channel.


### Hosts

Hosts are servers that provide the processing capabilities and memory resources used to run virtual machines. In Red Hat Virtualization, servers running Red Hat Enterprise Linux 7 or a minimal operating system specifically designed to operate as a host in Red Hat Virtualization environments, known as the Red Hat Virtualization Host, can be used as hosts.

In addition to these core components, Red Hat Virtualization allows you to set up supporting components such as a scheduler and remote servers for hosting the databases used in the environment, and to take advantage of resources in Red Hat Enterprise Linux OpenStack Platform and VMware environments.

## 1.2. SUBSCRIPTIONS

To install the Red Hat Virtualization Manager and hosts, your systems must be registered with the Content Delivery Network using Red Hat Subscription Management. This section outlines the entitlements and repositories required to set up a Red Hat Virtualization environment.



### IMPORTANT

Red Hat is transitioning the RHN-hosted interface into the Red Hat Subscription Management (RHSM) interfaces by July 31, 2017. If your current systems are registered to RHN Classic, see [Migrating from RHN Classic to Red Hat Subscription Management \(RHSM\) for Red Hat Virtualization](#) for instructions on how to migrate your systems to RHSM.

### 1.2.1. Required Entitlements and Repositories

The packages provided in the following repositories are required to install and configure a functioning Red Hat Virtualization environment. When one of these repositories is required to install a package, the steps required to enable the repository are provided in the appropriate location in the *Installation Guide* or *Self-Hosted Engine Guide*.

**Table 1.1. Red Hat Virtualization Manager**

Subscription Pool	Repository Name	Repository Label	Details
Red Hat Enterprise Linux Server	Red Hat Enterprise Linux Server	rhel-7-server-rpms	Provides the Red Hat Enterprise Linux 7 Server.
Red Hat Enterprise Linux Server	RHEL Server Supplementary	rhel-7-server-supplementary-rpms	Provides the virtio-win package, which provides the Windows VirtIO drivers for use in virtual machines.
Red Hat Virtualization	Red Hat Virtualization	rhel-7-server-rhv-4.1-rpms	Provides the Red Hat Virtualization Manager.
Red Hat Virtualization	Red Hat Virtualization Tools	rhel-7-server-rhv-4-tools-rpms	Provides dependencies for the the Red Hat Virtualization Manager that are common to all Red Hat Virtualization 4 releases.
Red Hat Virtualization	Red Hat JBoss Enterprise Application Platform	jb-eap-7.1-for-rhel-7-server-rpms	Provides the supported release of Red Hat JBoss Enterprise Application Platform on which the Manager runs.
Red Hat Ansible Engine	Red Hat Ansible Engine	rhel-7-server-ansible-2-rpms	Provides Red Hat Ansible Engine.

Table 1.2. Red Hat Virtualization Host

Subscription Pool	Repository Name	Repository Label	Details
Red Hat Virtualization	Red Hat Virtualization Host	rhel-7-server-rhvh-4-rpms	Provides the rhv-hypervisor7-ng-image-update package, which allows you to update the image installed on the host.

Table 1.3. Red Hat Enterprise Linux 7 Hosts

Subscription Pool	Repository Name	Repository Label	Details
Red Hat Enterprise Linux Server	Red Hat Enterprise Linux Server	rhel-7-server-rpms	Provides the Red Hat Enterprise Linux 7 Server.



Subscription Pool	Repository Name	Repository Label	Details
<b>Red Hat Virtualization</b>	<b>Red Hat Virtualization Management Agents (RPMs)</b>	<b>rhel-7-server-rhv-4-mgmt-agent-rpms</b>	Provides the QEMU and KVM packages required for using Red Hat Enterprise Linux 7 servers as virtualization hosts.

## 1.2.2. Optional Entitlements and Repositories

The packages provided in the following repositories are not required to install and configure a functioning Red Hat Virtualization environment. However, they are required to install packages that provide supporting functionality on virtual machines and client systems such as virtual machine resource monitoring. When one of these repositories is required to install a package, the steps required to enable the repository are provided in the appropriate location in the *Installation Guide* or *Self-Hosted Engine Guide*.

**Table 1.4. Optional Entitlements and Repositories**

Subscription Pool	Repository Name	Repository Label	Details
<b>Red Hat Enterprise Linux Server</b>	<b>Red Hat Enterprise Linux 7 Server - RH Common (v.7 Server for x86_64)</b>	<b>rhel-7-server-rh-common-rpms</b>	Provides the <code>rhev-guest-agent-common</code> package for Red Hat Enterprise Linux 7, which allows you to monitor virtual machine resources on Red Hat Enterprise Linux 7 clients.
<b>Red Hat Enterprise Linux Server</b>	<b>Red Hat Enterprise Virt Agent (v.6 Server for x86_64)</b>	<b>rhel-6-server-rhv-4-agent-rpms</b>	Provides the <code>rhev-guest-agent-common</code> package for Red Hat Enterprise Linux 6, which allows you to monitor virtual machine resources on Red Hat Enterprise Linux 6 clients.
<b>Red Hat Enterprise Linux Server</b>	<b>Red Hat Enterprise Virt Agent (v.5 Server for x86_64)</b>	<b>rhel-5-server-rhv-4-agent-rpms</b>	Provides the <code>rhev-guest-agent</code> package for Red Hat Enterprise Linux 5, which allows you to monitor virtual machine resources on Red Hat Enterprise Linux 5 clients.

Subscription Pool	Repository Name	Repository Label	Details
<b>Red Hat Virtualization</b>	<b>Red Hat Virtualization Host Build</b>	<b>rhel-7-server- rhvh-4-build- rpms</b>	Provides packages used to build your own version of the Red Hat Virtualization Host image.

## CHAPTER 2. RHV FOR IBM POWER

This release supports Red Hat Enterprise Linux 7 hosts on IBM POWER8 little endian hardware, and Red Hat Enterprise Linux 7 virtual machines on emulated IBM POWER8 hardware.



### IMPORTANT

Previous releases of RHV for IBM Power required Red Hat Enterprise Linux hosts on POWER8 hardware to be installed from an ISO image. These hosts cannot be updated for use with this release. You must reinstall Red Hat Enterprise Linux 7 hosts using the repositories outlined below.

The packages provided in the following repositories are required to install and configure aspects of a Red Hat Virtualization environment on POWER8 hardware.

**Table 2.1. Required Entitlements and Repositories**

Component	Subscription Pool	Repository Name	Repository Label	Details
Red Hat Virtualization Manager	<b>Red Hat Virtualization for IBM Power</b>	<b>Red Hat Virtualization for IBM Power</b>	<b>rhel-7-server-rhv-4-power-rpms</b>	Provides the Red Hat Virtualization Manager for use with IBM POWER8 hosts. The Manager itself must be installed on x86_64 architecture.
Red Hat Enterprise Linux 7 hosts, little endian	<b>Red Hat Enterprise Linux for Power, little endian</b>	<b>RHV Management Agent for IBM Power, little endian</b>	<b>rhel-7-server-rhv-4-mgmt-agent-for-power-le-rpms</b>	Provides the QEMU and KVM packages required for using Red Hat Enterprise Linux 7 servers on IBM Power (little endian) hardware as virtualization hosts.
	<b>Red Hat Enterprise Linux for Power, little endian</b>	<b>Red Hat Enterprise Linux for IBM Power, little endian</b>	<b>rhel-7-for-power-le-rpms</b>	Provides additional packages required for using Red Hat Enterprise Linux 7 servers on IBM Power (little endian) hardware as virtualization hosts.

Component	Subscription Pool	Repository Name	Repository Label	Details
Red Hat Enterprise Linux 7 virtual machines, big endian	<b>Red Hat Enterprise Linux for Power, big endian</b>	<b>RHV Tools for IBM Power</b>	<b>rhel-7-server-rhv-4-tools-for-power-le-rpms</b>	Provides the rhevm-guest-agent-common package for Red Hat Enterprise Linux 7 virtual machines on emulated IBM Power (big endian) hardware. The guest agents allow you to monitor virtual machine resources on Red Hat Enterprise Linux 7 clients.
Red Hat Enterprise Linux 7 virtual machines, little endian	<b>Red Hat Enterprise Linux for Power, little endian</b>	<b>RHV Tools for IBM Power, little endian</b>	<b>rhel-7-server-rhv-4-tools-for-power-le-rpms</b>	Provides the rhevm-guest-agent-common package for Red Hat Enterprise Linux 7 virtual machines on emulated IBM Power (little endian) hardware. The guest agents allow you to monitor virtual machine resources on Red Hat Enterprise Linux 7 clients.

## Unsupported Features

The following Red Hat Virtualization features are not supported:

- Hot-plug CPU
- SPICE display
- SmartCard
- Sound device
- Guest SSO
- Integration with OpenStack Networking (Neutron), OpenStack Image (Glance), and OpenStack Volume (Cinder)
- Self-hosted engine
- Red Hat Virtualization Host (RHVH)

- Disk Block Alignment

For a full list of bugs that affect the RHV for IBM Power release, see [https://bugzilla.redhat.com/show\\_bug.cgi?id=1359843](https://bugzilla.redhat.com/show_bug.cgi?id=1359843).

## CHAPTER 3. RELEASE INFORMATION

These release notes highlight technology preview items, recommended practices, known issues, and deprecated functionality to be taken into consideration when deploying this release of Red Hat Virtualization.

Notes for updates released during the support lifecycle of this Red Hat Virtualization release will appear in the advisory text associated with each update or the *Red Hat Virtualization Technical Notes*. This document is available from the following page:

<https://access.redhat.com/documentation/en/red-hat-virtualization>

### 3.1. RED HAT VIRTUALIZATION 4.1 GA

#### 3.1.1. Enhancements

This release of Red Hat Virtualization features the following enhancements:

##### **BZ#734120**

With this update, the ability to sparsify a thin-provisioned disk has been added to Red Hat Virtualization. When a virtual machine is shut down, the user can sparsify the disk to convert the free space within the disk image back to free space on the host.

##### **BZ#804272**

Previously, if power management fencing was not available, automatic high availability of virtual machines did not work without manual confirmation that the host the virtual machine was running on was rebooted.

In this release, a new option for virtual machine lease on storage domains enables automatic high availability failover of a virtual machine, when the host running the machine is down, without the availability of power management for the host.

##### **BZ#827529**

This release introduces QCOW2 v3, which has a compatibility level of 1.1. This enables the QEMU to use this volume in a more efficient way, with its improved performance capabilities. In addition, it is fully backwards-compatible with the QCOW2 feature set, is easy to upgrade from QCOW2 v2 to QCOW2 v3, and supports extensibility.

##### **BZ#868811**

This enhancement allows live migration of virtual machines that have SR-IOV NICs attached. Before the migration all SR-IOV NICs are hot unplugged, and after successful migration they are hot plugged back.

##### **BZ#1001181**

-

This release provides a cleanup script for completely cleaning the host after a failed attempt to install a self-hosted engine.

**BZ#1024063**

Previously, it was impossible to reboot a host without enabling Power Management. In this release, it is now possible to shut down and reboot a host without using Power Management. From the Management menu, a new option called SSH Management is available, enabling administrators to select either Restart or Stop.

**BZ#1036221**

With this update, if the web console, either noVnc or SPICE HTML5, is unable to connect to the websocket proxy server, a pop-up displays with troubleshooting suggestions. The pop-up contains a link to the default CA certificate.

**BZ#1080379**

Previously, almost all data path operations on the hosts could only be performed on the elected Storage Pool Manager. This could potentially cause bottlenecks.

In this release, a new lightweight host jobs management infrastructure was introduced, which enables data path operations to run on any host. In addition, administrators can monitor the progress of Move Disk operations with the aid of a progress indicator.

**BZ#1081536**

With this release, when creating virtual machine pools using a template that is present in more than one storage domain, virtual machine disks can be distributed to multiple storage domains by selecting "Auto select target" in New Pool -> Resource Allocation -> Disk Allocation.

**BZ#1092907**

Previously, when notification emails were successfully sent to a configured SMTP server, a success message did not appear in the notifier.log file. In this release, when a message is successfully sent to an SMTP server, the following message appears in the notifier.log file: E-mail subject='...' to='...' sent successfully

**BZ#1097589**

This release adds support for CPU hot unplug to Red Hat Virtualization. Note that the guest operating system must also support the feature, and only previously hot plugged CPUs can be hot unplugged.

**BZ#1101554**

■

With this update, the code interfacing with VDSM now uses the VDSM API directly instead of using vdsClient and xmlrpc.

**BZ#1135976**

Previously, CPU pinning information could be silently lost. Now, a pop-up warning, that CPU pinning information will be lost, displays when a user saves a virtual machine. This means that the user is aware that CPU pinning information will be lost and has the choice to cancel the operation.

**BZ#1141422**

With this update, the VDSM thread name is now included in the system monitoring tools. This makes it easier to track the resource usage of the threads.

**BZ#1145240**

Power Management alerts are now disabled when fencing is disabled in a cluster.

**BZ#1161625**

With this update, virtual machines can now be searched for by the user who created them. Using the REST API, the search query is ".../api/vms?search=created\_by\_user\_id%3D[USER\_ID]". The required User ID can be retrieved by using ".../api/users". In addition, the Administration Portal shows the creator's name in the virtual machine general sub-tab. However, it is possible for the user to be removed from the system after the virtual machine is created.

**BZ#1177229**

With this update, image disks are now identifiable from within the guest by engine id, for example, by looking under /dev/disk/by-id. The disk id is now passed to the guest as the disk serial.

**BZ#1177782**

This update provides a link to the gluster volume when creating a gluster storage domain, and enables a single unified flow.

This enables the backup volfile server mount options to be auto-populated, and paves the way for integration features like Disaster Recovery setup using gluster geo-replication.

**BZ#1182369**

Previously, in a hyper-converged cluster environment containing gluster and virt nodes, it was possible to create a replica set containing bricks from the same server. A warning appeared but the action was



enabled even though there was a risk of losing data or service. In this release, it will no longer be possible to create a replica set containing multiple bricks from the same server in a hyper-converged environment.

**BZ#1195140**

The `vioscsi.sys` file is now compatible with Microsoft Cluster Services, which enables the Windows `virtio-scsi` driver to support Windows Server Failover Clustering (WSFC) using shared storage. As a result, `vioscsi.sys` can pass all tests in the "Validate a Configuration" process.

**BZ#1196433**

Previously, in GlusterFS, if a node went down and then returned, GlusterFS would automatically initiate a self-heal operation. During this operation, which could be time-consuming, a subsequent maintenance mode action within the same GlusterFS replica set could result in a split-brain scenario.

In this release, if a Gluster host is performing a self-heal activity, administrators will not be able to move it into maintenance mode. In extreme cases, administrators can use the `force` option to forcefully move a host into maintenance mode.

**BZ#1201020**

Previously, the system performed automatic migrations, when required, without displaying the reason for doing so in the Event log or in the Administrator portal. In this release, after performing an automatic migration, the reason for doing so will be displayed.

**BZ#1213291**

This update introduced a check in the host maintenance flow to ensure glusterFS quorum can be maintained for all glusterFS volumes that have the `'cluster.quorum-type'` option set. Similarly, there is a new check to ensure that the host moving to maintenance does not have a glusterFS brick that is a source of volume self-healing. These checks will be performed by default when moving the host to maintenance.

There is an option in the Manager to skip these checks, but this can result in bringing your system to halt. This option should only be used in extreme cases.

**BZ#1226206**

Previously, when importing a virtual machine from a data storage domain, if the virtual machine had a "bad" MAC address, a MAC collision could occur in the target LAN. A "bad" MAC address is an address that is

already in use or an address that is out of the range in the target cluster. In this release, it is possible to assign a new MAC address when importing the virtual machine from a data storage domain.

**BZ#1235200**

Previously, when restoring a backup of a hosted engine on a different environment, for disaster recovery purposes, administrators were sometimes required to remove the previous hosts from the engine. This was accomplished from within the engine's database, which is a risk-prone procedure.

In this release, a new CLI option can be used during the restore procedure to enable administrators to remove the previous host directly from the engine backup.

**BZ#1240466**

Previously, when restoring a backup of a self-hosted engine on a different environment, for disaster recovery purposes, administrators were sometimes required to remove the previous self-hosted engine's storage domain and virtual machine. This was accomplished from within the engine's database, which was a risk-prone procedure.

With this update, a new CLI option enables administrators to remove the previous self-hosted engine's storage domain and virtual machine directly from the backup of the engine during the restore procedure.

**BZ#1241106**

Previously, discard commands (UNMAP SCSI commands) that were sent from the guest were ignored by QEMU and were not passed on to the underlying storage. This meant that storage that was no longer in use could not be freed up.

In this release, it is now possible to pass on discard commands to the underlying storage. A new property called Pass Discard was added to the Virtual Disk window. When selected, and if all the restrictions are met, discard commands that are sent from the guest will not be ignored by QEMU and will be passed on to the underlying storage. The reported unused blocks in the thinly provisioned LUNs of the underlying storage will be marked as free, and the reported consumed space will be reduced.

**BZ#1246114**

Previously, when the Virtual Machine was powered down, deleting a snapshot could potentially be a very long process. This was due to the need to copy the data from the base snapshot to the top snapshot, where the base snapshot is usually larger than the top snapshot.

Now, when deleting a snapshot when the Virtual Machine is powered down, data is copied from the top snapshot to the base snapshot, which significantly reduces the time required to delete the snapshot.

-

**BZ#1254073**

With this update, support for Gluster arbiter volume creation has been added to Red Hat Virtualization. Arbiter volumes are recommended in place of regular three-way replicated volumes to save storage space.

**BZ#1256500**

Previously, deleting a snapshot while the virtual machine was down took a long time. With this release, snapshot deletion while a virtual machine is down uses 'qemu-img commit' instead of 'qemu-img rebase', improving the performance of the operation.

**BZ#1264085**

With this update, the user can configure the number of memory slots reserved for spare self-hosted engine hosts if the current host crashes. Previously, there was a chance that the self-hosted engine virtual machine would not have a place to start on a loaded cluster and this compromised the high availability feature. Now, the self-hosted engine will have a place to start a backup host so that it is ready to accept the virtual machine if the current host crashes.

**BZ#1270719**

With this update, the option '--accept-defaults' has been added to the engine-setup command. This option causes engine-setup to no longer prompt for answers that have a default. This option saves time for the user as they no longer need to answer the prompts individually, if they are planning to accept the defaults, and also allows other tools to run engine-setup unattended. If the engine-setup command is run using this option and a weak password is provided, the user will be prompted for a stronger password because the default answer to 'Use weak password?' is No. To work around this, add the answer to an answer file.

**BZ#1277675**

This feature allows you to map external VNIC profiles that are defined on an imported VM to the ones that are present in the cluster the VM is going to be imported to. The previous solution exchanged all external VNIC profiles that were not present in the target cluster with an empty profile, which removed the imported VM's network functionality. Now, after importing a VM from a data domain, the VM is configured properly according to the VNIC profiles that are defined in the target cluster.

**BZ#1279378**

With this update, a new option to check for upgrades has been added when installing a host. In the Administration Portal this is available in the host installation menu and it can be triggered by using the hosts upgradecheck endpoint in the REST API.

**BZ#1285499**

Previously, the column control menu in the Administration Portal, which is accessed by right-clicking on a column header in any of the tables, contained arrows for controlling the order of the columns. In this release, the order of the columns in this menu can be defined by dragging and dropping the column to the required position within the menu.

**BZ#1286632**

In this release, a link has been added to the Edit Fence Agent window, which opens the online help and displays information about the parameters that can be set for fence agents.

**BZ#1289155**

With this update, the limit of virtual CPUs has been increased to reflect the capabilities of Red Hat Enterprise Linux 7.3 hosts. It is now possible to configure a virtual machine with up to 288 vCPUs when hosts are in a 4.1 cluster.

**BZ#1289290**

Previously, if a live migration was performed with extreme memory write intensive workloads, the migration would never be able to complete because QEMU could not transfer the memory changes fast enough. In this case, the migration could not reach the non-live finishing phase.

In this release and in these situations, RHV will restrict the amount of CPU given to the guest to reduce the rate at which memory is changed and allow the migration to complete.

**BZ#1294629**

With this update, the loading performance of external virtual machines from an external server has been improved for VMware, KVM, and Xen. Previously, when displaying a list of virtual machines, libvirt was asked for the full information for each virtual machine when only the virtual machine names were required. Now, libvirt is only asked for the virtual machine names at the first import dialog and only imports the full virtual machine data list after the user has selected the required virtual machines.

**BZ#1299677**

With this update, CPU hotplug is now supported on POWER for selected guest operating systems. CPU hotplug is now supported on both x86\_64 architecture and on selected guest operating systems for ppc64le.

**BZ#1300591**

With this update, the user can now customize the virtual machine disk size when using the engine-appliance.

**BZ#1300959**

The latest virtio-win release, which includes Windows 10 drivers, is now required by the Manager to ensure that suitable drivers can be supplied to virtual machines during installation of Windows 10.

**BZ#1301681**

This feature allows you to edit the configuration stored in the self-hosted engine's shared storage.

**BZ#1302185**

Users can now change an initialized Data Center type to Shared or Local. The following updates will be available:

1. Shared to Local - Only for a data center that does not contain more than one host and more than one cluster, since a local data center does not support it. The Manager will validate and block this operation with the following messages:

```
CLUSTER_CANNOT_ADD_MORE_THEN_ONE_HOST_TO_LOCAL_STORAGE  
VDS_CANNOT_ADD_MORE_THEN_ONE_HOST_TO_LOCAL_STORAGE
```

2. Local to Shared - Only for a data center that does not contain a local storage domain. The Manager will validate and block this operation with the following message:

```
ERROR_CANNOT_CHANGE_STORAGE_POOL_TYPE_WITH_LOCAL.
```

This update is useful for data domains that are used to move virtual machines or templates around.

**BZ#1306263**

The weighting for virtual machine scheduling has been updated. The best host for the virtual machine is now selected using a weighted rank algorithm instead of the pure sum of weights. A rank is calculated for the policy unit and host, and the weight multiplier is then used to multiply the ranks for the given policy unit. The host with the highest number is selected.

The reason for the change is that current weight policy units do not use a common result value range. Each unit reports numbers as needed, and this causes issues with user configured preferences. For example, memory (which has high numbers) always wins over CPU (0-100).

This update ensures that the impact of the policy unit multiplier for

the scheduling policy configuration is more predictable. However, users who are using it should check their configuration for sanity when upgrading.

**BZ#1309930**

This update ensures that the Manager signs certificates using the SHA-256 algorithm instead of SHA-1 because SHA-256 is more secure and is expected to have a longer life expectancy.

Only the default for new certificates was changed. To change certificates for existing hosts, they need to be reinstalled, or to have their certificates enrolled. Other certificates require a completely new setup, using engine-cleanup and engine-setup, including the one for httpd.

**BZ#1314387**

This update adds the ability to download or upload Red Hat Virtualization images (for example, virtual machine images) using the Red Hat Virtualization API.

**BZ#1316692**

High availability could previously not be enabled for virtual machines in Hyper-converged mode. Previous fencing policies ignored Gluster processes. But in Hyper-converged mode, fencing policies are required to ensure that a host is not fenced if there is a brick process running, or to ensure no loss of quorum when shutting down the host with an active brick.

The following fencing policies have been added to Hyper-converged clusters:

- SkipFencingIfGlusterBricksUp: Fencing will be skipped if bricks are running and can be reached from other peers.
- SkipFencingIfGlusterQuorumNotMet: Fencing will be skipped if bricks are running and shutting down the host will cause loss of quorum

Virtual machine high availability can be tested by enabling power management on hyper-converged nodes.

**BZ#1317429**

This update adds the ability to acquire a lease per virtual machine on shared storage, without attaching the lease to a disk. This adds the capability to avoid split-brain, and to avoid starting a virtual machine on another host if the original host becomes non-responsive, thereby improving virtual machine high availability.

**BZ#1317436**

MAC address pools are now bound to the cluster instead of the data center because certain environments require MAC address pools on the

cluster level.

On the REST layer, the `mac_pool` attribute was added to the cluster, and can be set or queried. The `StoragePool` resource (represents the data center) was also altered. When updating its `mac_pool_id`, all clusters in a given data center will be updated to use this MAC address pool. When `StoragePool` is queried using the REST GET method, the ID of the MAC address pool will be reported only when all clusters in the given data center are using the same MAC address pool. Otherwise, the user needs to use the `Cluster` resource to get the MAC address pool of each individual cluster.

#### **BZ#1317447**

This feature allows you to request a new MAC address when importing a virtual machine from a data storage domain. This allows you to avoid importing a virtual machine with a bad MAC address, which might cause a MAC address collision in the target LAN. A MAC address would be considered "bad" if it is already in use in the target cluster or is out of the range of the MAC pool of the target cluster.

#### **BZ#1318350**

The NTP configuration is automatically set when deploying the self-hosted engine.

#### **BZ#1325864**

This update adds support for deploying Gluster storage during the self-hosted engine deployment through the Cockpit UI. Previously, the user needed to first deploy the Gluster storage using `gdeploy`, then deploy the self-hosted engine using the Cockpit UI, and configuration files had to be manually updated.

#### **BZ#1330138**

This update ensures that a self-hosted engine deployment works without needing to disable `NetworkManager`.

#### **BZ#1331858**

This update allows you to enable SSH access for the Manager virtual machine when deploying the self-hosted engine. You can choose between `yes`, `no`, and `without-password`. You can also pass a public SSH key for the root user during deployment.

#### **BZ#1333045**

A new `'original_template'` property has been introduced for the `'vm'` REST API resource. This enables the user to get information about the template the virtual machine was based on before cloning.

**BZ#1333436**

Previously, support for Legacy USB was deprecated and the UI displayed three options: Native, Legacy (Deprecated), and Disabled. In this release, the Legacy option has been completely removed and the UI now displays two options: Enabled and Disabled.

**BZ#1337101**

With this release, `/dev/random` is now the default random number generator in clusters with a cluster compatibility level of 4.0 and earlier, and `/dev/urandom` is now the default random number generator in clusters with a cluster compatibility level of 4.1 and later. Because these random number generators are enabled by default, the option to enable them has now been removed from the New Cluster and Edit Cluster windows. However, you can select the random number generator source for individual virtual machines from the New Virtual Machine and Edit Virtual Machine windows.

**BZ#1341153**

This update includes the subversion name and subversion number of the template being removed from the "remove template(s)" screen.

The two "remove template(s)" screens now display the following:

Are you sure you want to remove the following items?

- template-name (Version: subversion-name(subversion-number))

**BZ#1342919**

With this update, it is now possible to configure discard after delete per block storage domain. Previously, a user could get similar functionality by configuring the `discard_enable` parameter in the VDSM configuration file. This caused each logical volume (disk or snapshot) that was about to be removed by this specific host to be discarded first. Now, discard after delete can be enabled for a block storage domain, rather than a host. This means that if discard after delete is enabled, it no longer matters which host removes the logical volume, as each logical volume under this domain will be discarded before it is removed.

**BZ#1343562**

This update ensures that only hosts that have the status Up or NonOperational are checked for upgrades. Previously hosts with the status Maintenance were also checked, but often they were not reachable, which caused errors in Events.

**BZ#1344521**

Previously, if the guest agent was not running or was out of date, the hover text message that appeared next to the explanation mark for the



problematic virtual machine informed the user that the operating system did not match or that the timezone configuration was incorrect. In this release, the hover text will correctly display a message informing the user that the guest agent needs to be installed and running in the guest.

**BZ#1347631**

With this update, a log file located in `/var/log/httpd/ovirt-requests-log` now logs all requests made to the Red Hat Virtualization Manager via HTTPS, including how long each request took. There is a 'Correlation-Id' header included to allow for easier comparison of requests with the `engine.log`. CorrelationIds are now generated for every request automatically and can be passed to the REST API per Correlation-Id header or per `correlation_id` query parameter.

**BZ#1348107**

With this update, a user can now save a provider for external libvirt connection in the Providers tree section. When a user tries to import a virtual machine from libvirt+kvm to the Red Hat Virtualization environment the saved provider is available instead of having to re-enter the address.

**BZ#1349301**

The self-hosted engine only supports deployment using the RHV-M Appliance. With this release, the deployment script allows you to download and install the Appliance RPM directly, instead of having to install it before deployment.

**BZ#1349321**

When users import virtual machines from Xen on RHEL to Red Hat Virtualization it will access the saved provider address instead of the user having to re-enter the address.

**BZ#1349907**

Previously, `after_hibernation` hooks were never executed. With this release, `before_hibernation` and `after_hibernation` hooks are always executed on the guest operating system (with the guest agent installed) when suspending and resuming a virtual machine.

**BZ#1350465**

Previously, when importing a virtual machine, if the import failed, the output of the `virt-v2v` tool was not available for investigating the reason for the failure, and the import had to be reproduced manually. In this release, the output of `virt-v2v` is now stored in the `/var/log/vdsm/import` directory. All logs older than 30 days are automatically removed.

**BZ#1353556**

Previously, a Dashboard tab was introduced to the Administration Portal. However, when loading the Administration Portal the user landed at the Virtual Machines tab followed by an immediate switch to the Dashboard tab. Now, the UI plugin has been improved to allow pre-loading of UI plugins, such as `ovirt-engine-dashboard`. This means that the user lands directly at the Dashboard tab.

**BZ#1353750**

With this update, the debug logging for `ovirt-engine-extension-aaa-ldap` has been updated. When `ovirt-engine-extension-aaa-ldap` is enabled the following messages will show in the logs. The LDAP server that authenticated a user is shown as "User 'myuser1' is performing bind request to: ldap.example.com" and the LDAP server that performed a search request is shown as "Performing SearchRequest '...' request on server ldap.example.com."

**BZ#1354343**

This update includes the Post-copy migration policy, which is available as a Technology Preview feature. The policy is similar to the Minimal Downtime policy, and enables the virtual machine to start running on the destination host as soon as possible. During the final phase of the migration (post-copy phase), the missing parts of the memory content is transferred between the hosts on demand. This guarantees that the migration will eventually converge with very little downtime. The disadvantage of this policy is that in the post-copy phase, the virtual machine may slow down significantly as the missing parts of memory are transferred between the hosts. If anything goes wrong during the post-copy phase, such as a network failure between the hosts, the running virtual machine instance will be lost. It is therefore not possible to abort a migration during the post-copy phase.

**BZ#1356161**

With this release, if you do not specify any NUMA mapping, Red Hat Virtualization defaults to a NUMA node that contains the host device's memory-mapped I/O (MMIO). This configuration is only preferred, rather than strictly required.

**BZ#1358716**

The self-hosted engine setup wizard now warns users if the host is already registered to Red Hat Virtualization Manager. Previously, a host that was registered to the Manager but not running a self-hosted engine would present the option to set up a self-hosted engine, which ran the risk of unregistering the host. Now, hosts that are registered to the Manager present a "Redeploy" button in the Hosted Engine wizard in Cockpit, which must be selected in order to continue.

**BZ#1361115**

This update adds Gluster-related fencing policies for hyper-converged clusters. Previous fencing policies ignored Gluster processes. But in Hyper-converged mode, fencing policies are required to ensure that a host is not fenced if there is a brick process running, or to ensure no loss of quorum when shutting down the host with an active brick.

The following fencing policies have been added to Hyper-converged clusters:

- SkipFencingIfGlusterBricksUp: Fencing will be skipped if bricks are running and can be reached from other peers.
- SkipFencingIfGlusterQuorumNotMet: Fencing will be skipped if bricks are running and shutting down the host will cause loss of quorum

#### **BZ#1361230**

Red Hat Virtualization Host (RHVH) 4.0 allows users to install RPMs. However, installed RPMs are lost after upgrading RHVH.

RHVH 4.1 now includes a yum plugin that saves and reinstalls RPM packages after upgrading, to ensure that installed RPMs are no longer lost after upgrading.

This will not work when upgrading from RHVH 4.0 to RHVH 4.1.

#### **BZ#1362604**

The rng-tools package has been added to oVirt Node NG / RHV-H. This tool is required for the TPM module to be able to work with the Random Number Generator.

#### **BZ#1365021**

This enhancement is a rebase on the jsonrpc Dispatcher APIs to provide better performance and make the code more robust.

#### **BZ#1365937**

A mobile client for Red Hat Enterprise Virtualization, which is compatible with Red Hat Enterprise Virtualization 3.5 onwards, is available for Android devices.

#### **BZ#1366118**

oVirt release now provides repository configuration files for enabling GlusterFS 3.8 repositories on Red Hat Enterprise Linux, CentOS Linux, and similar operating systems.

#### **BZ#1366183**

Since Red Hat Virtualization now has the capability to deploy additional self-hosted engine hosts from the Manager with host-deploy, the capability to deploy additional self-hosted engine hosts from hosted-

engine setup is not required anymore. It has now been removed.

Similarly, the RHV-M Appliance has proved to be the easiest flow to have a working self-hosted engine environment; all other bootstrap flows have now been removed.

**BZ#1366899**

This release adds support for overlay networks using Open Virtual Network (OVN) as a Technology Preview. This feature allows you to add OVN as an external network provider, and import or create networks from it in the Red Hat Virtualization Manager. You can then provision virtual machines with network interfaces connected using these logical overlays (OVN networks).

**BZ#1367826**

Previously, Python SDK was configured to communicate with the server using uncompressed responses. This caused long response times. In this release, the default configuration is to send compressed responses.

**BZ#1367924**

Multiple updates were made to the UI for Self-Hosted Engine.

New icons have been added:

- To virtual machines to indicate whether it is the Manager virtual machine.
- To hosts to indicate whether it can run the Manager virtual machine.
- To storage domains to indicate whether it contains the Manager virtual machine.

Buttons to enable and disable global maintenance mode have been moved to the context menu of a host that can run the Manager virtual machine. Depending on the current status of global maintenance mode, either the enable or disable option will be enabled.

**BZ#1369175**

The "Enable USB Auto-Share" option in the "Console options" dialog is now only available if "USB Support" is enabled on the virtual machine.

**BZ#1370485**

Previously, Java SDK was configured to communicate with the server using uncompressed responses. This caused long response times. In this release, the default configuration is to send compressed responses.

**BZ#1371530**

In this release, when installing or reinstalling hosts, the collectd and fluentd packages are now installed, including the relevant plugins. These can be used to send statistics to a central metrics store.

■

**BZ#1373223**

Previously, if SPICE USB redirection was disabled, libvirt created a default USB controller. With this update, if SPICE USB redirection is disabled then the virtual machine has a new USB controller, which is configurable per guest operating system and cluster version. This is defined in the `osinfo-defaults.properties` configuration file.

**BZ#1375568**

The `tcpdump` package is now included with Red Hat Virtualization Host.

**BZ#1376454**

A reinstalling and restoring workflow was tested and confirmed for moving from version 3.6 Red Hat Enterprise Virtualization Hypervisor hosts to the new implementation, Red Hat Virtualization Host, in 4.0 or 4.1.

**BZ#1379000**

With this update, IBM Security (Tivoli) Directory Server has been added to supported LDAP servers in `ovirt-engine-extension-aaa-ldap`. This allows customers to attach Red Hat Virtualization 4.1 to their IBM Security (Tivoli) Directory Server setup and to use users and groups from this setup in Red Hat Virtualization.

**BZ#1379074**

Previously, the `ExportVmCommand` appeared in the Engine log without the ID of the virtual machine being exported. This information has now been added to the log.

Note: After this change, users must have export permissions for the virtual machine and its disks to export a virtual machine. Previously, permissions to export virtual machine disks were sufficient.

**BZ#1379154**

Previously, users who wanted to use Cockpit for system configuration needed to log in to the system and retrieve IP address information manually. Now, Red Hat Virtualization Host provides a message on login informing users of the URL to the Cockpit user interface.

**BZ#1379763**

The `"screen"` package is now available as part of the base RHVH image.

**BZ#1379771**

This release introduces a 'force' flag, which can be used to update a storage server connection regardless of the associated storage domain status (allowing updates even when the storage domain is not in Maintenance mode).

For example: `PUT /ovirt-engine/api/storageconnections/123;force=true`

#### **BZ#1380365**

This update adds the ability to import partial virtual machines using the REST API.

The Hyper Converged Infrastructure (HCI) Disaster Recovery (DR) solution is based on the concept that only data disks are replicated and system disks are not. Previously, if some of the virtual machine's disks were not replicated, the virtual machine import would fail. Since disks have snapshots, they cannot be imported as floating disks. To allow the DR to work, a virtual machine is forced to import from a storage domain, even if some of its disks are not accessible.

The following is a REST request for importing a partial unregistered virtual machine.

```
POST /api/storagedomains/xxxxxxx-xxxx-xxxx-xxxxxxx/vms/xxxxxxx-xxxx-xxxx-xxxxxxx/register HTTP/1.1
```

```
Accept: application/xml
```

```
Content-type: application/xml
```

```
<action>
  <cluster id='bf5a9e9e-5b52-4b0d-aeba-4ee4493f1072'></cluster>
  <allow_partial_import>true</allow_partial_import>
</action>
```

#### **BZ#1381184**

Red Hat Virtualization now supports headless virtual machines that run without a graphical console and display device. Headless mode is also supported for templates, pools, and instance types. This feature supports running a headless virtual machine from start, or after the initial setup (after "Run Once"). Headless mode can be enabled or disabled for a new or existing virtual machine at any time.

#### **BZ#1382843**

With this update, the Red Hat Virtualization Host includes sysstat as part of the base image.

#### **BZ#1383342**

This feature allows you to request a console ticket for a specific graphics device by means of the REST API. The existing endpoint, `/api/vms/{vmId}/ticket`, defaulted to SPICE in scenarios when SPICE+VNC was configured as the graphics protocol, making it impossible to request

a VNC ticket. Now, a ticket action has been added to the `/api/vms/{vmId}/graphicsconsoles/{consoleId}` resource, making it possible to request a ticket for a specific console. This specific endpoint is now preferred, and the pre-existing per-VM endpoint is considered deprecated.

**BZ#1387254**

Previously when integrating the Manager with an LDAP server using the `ovirt-engine-extension-aaa-ldap-setup` tool, the root of the LDAP tree (base DN) was selected automatically based on the LDAP server defaults. However, sometimes the defaults are incorrect for Manager integrations, and administrators are required to edit configuration files manually after the setup job completes.

Now the `ovirt-engine-extension-aaa-ldap-setup` tool offers to override the default base DN retrieved from LDAP server, so manual changes are no longer necessary.

**BZ#1388245**

This release adds the ability to specify a Maximum Memory value in all VM-like dialogs (Virtual Machine, Template, Pool, and Instance Type). It is accessible in the `{vm, template, instance_type}/memory_policy/max` tag in the REST API. The value defines the upper limit to which memory hot plug can be performed. The default value is 4x memory size.

**BZ#1388430**

This release adds a maintenance tool to run vacuum actions on the engine database (or specific tables). This tool optimizes table stats and compacts the internals of tables, resulting in less disk space usage, more efficient future maintenance work, and updated table stats for better query planning. Also provided is an engine-setup dialog that offers to perform vacuum during upgrades. This can be automated by the answer file.

**BZ#1388724**

Previously, it was not possible to install Windows Server 2016 on a virtual machine. In this release, it is now possible to install Windows Server 2016 on a virtual machine. When adding a virtual machine, Windows Server 2016 appears in the list of available operating systems.

**BZ#1389324**

Previously, the Networking tab in Cockpit was disabled in Red Hat Virtualization Host (RHVH) images. This is now enabled, meaning that system networking can be configured through Cockpit in RHVH.

**BZ#1392393**

Support for virtual machine to host affinity has been added. This

enables users to create affinity groups for virtual machines to be associated with designated hosts. Virtual machine host affinity can be disabled or enabled on request.

Virtual machine to host affinity is useful in the following scenarios:

- Hosts with specific hardware are required by certain virtual machines.
- Virtual machines that form a logical management unit can be run on a certain set of hosts for SLA or management, for example, a separate rack for each customer.
- Virtual machines with licensed software must run on specific physical machines to avoid scheduling virtual machines to hosts that need to be decommissioned or upgraded.

#### **BZ#1392418**

The user experience for HA global maintenance has been improved in the UI by moving the options to a more logical location, and providing a visual indication about the current state of HA global maintenance for a given host.

The "Enable HA Global Maintenance" and "Disable HA Global Maintenance" buttons are now displayed on the right-click menu for hosts instead of virtual machines, and reflect the global maintenance state of the host by disabling the button matching the host's current HA global maintenance state.

The previous method of displaying the options for virtual machines was unintuitive. Additionally, both the enable and disable options remained available regardless of whether or not the host was in HA global maintenance mode.

#### **BZ#1392872**

With this release, Intel Skylake family CPUs are now supported.

#### **BZ#1393257**

This update adds the ability to import partial templates through REST. You can register a template even if some of the storage domains are missing.

The following is a REST request for importing a partial unregistered template:

```
POST /api/storagedomains/xxxxxxx-xxxx-xxxx-xxxxxx/templates/xxxxxxx-xxxx-xxxx-xxxxxx/register HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

```
<action>
  <cluster id='bf5a9e9e-5b52-4b0d-aeba-4ee4493f1072'></cluster>
  <allow_partial_import>true</allow_partial_import>
</action>
```



**BZ#1393407**

During the authorization stage of the login flow, the user's group memberships, including nested groups, are retrieved. Nested group memberships are resolved using recursive LDAP searches, which could take significant amount of time.

This update uses a special Active Directory feature called `LDAP_MATCHING_RULE_IN_CHAIN`, which allows you to fetch complete group memberships, including nested groups, in one LDAP search.

**BZ#1393918**

With this update, some ancillary self-hosted engine commands that were still based on `xmlrpc` have been moved to `jsonrpc`.

**BZ#1396672**

Since Red Hat Enterprise Virtualization 3.6, `ovirt-ha-agent` read its configuration and the Manager virtual machine specification from shared storage. Previously, they were local files replicated on each involved host. This enhancement modifies the output of `hosted-engine --vm-status` to show whether the configuration and the Manager virtual machine specification have been, on each reported host, correctly read from the shared storage.

**BZ#1397833**

Previously, the Java heap size for Data Warehouse was not explicitly set. This resulted in the Java virtual machine using the default size, which could have been as large as a quarter of the machine's total memory. With this release, Data Warehouse's configuration was updated to allocate 1 GB of RAM, with the addition of two new parameters:

```
DWH_HEAP_MIN=1g
DWH_HEAP_MAX=1g
```

The size can be set to a higher value for larger environments using these new parameters.

**BZ#1398593**

This feature integrates the setup for data sync to a remote location using geo-replication for Gluster-based storage domains, to improve disaster recovery. A user is able to schedule data sync to a remote location from the Red Hat Virtualization UI.

**BZ#1399142**

This release changes the default disk interface type from `virtio-blk` to `virtio-scsi`. `virtio-blk` is still supported, but users are encouraged to use the more modern `virtio-scsi`. When creating or attaching a disk to a virtual machine, the `virtio-scsi` interface type will now be selected by default.

**BZ#1400366**

This update allows you to change the default network used by the host from the management network (ovirtmgmt) to a non-management network.

**BZ#1402435**

The self-hosted engine's machine type has now been upgraded for Red Hat Enterprise Linux 7 compatibility.

**BZ#1403839**

With this update, the ability to remove LUNs from a block data domain has been added. This means that LUNs can be removed from a block data domain provided that there is enough free space on the other domain devices to contain the data stored on the LUNs being removed.

**BZ#1404232**

Previously, when a USB hub containing a redirected device was unplugged, spice-usbd-win failed to clean up the redirected device. When the USB hub and its attached device were replugged, the device could not be redirected.

In this release, the issue has been fixed. spice-usbd-win will now clean up the redirected device as required. When the USB hub and the USB device are re-plugged, the device can be redirected to the guest.

**BZ#1404660**

This feature adds rule enforcement support for VM to host affinity. VM to host affinity groups require the affinity rule enforcer to handle them in addition to the existing enforcement of VM to VM affinity. The rule enforcer will now be able to find VM to host affinity violations and choose a VM to migrate according to these violations.

**BZ#1405032**

This release adds the VirtIO-RNG driver installer to the guest tools ISO for supported Windows versions.

**BZ#1405813**

A script is now supplied to configure collectd and fluentd on hosts to send statistics to a central store.

**BZ#1405912**

In this release, the following optional troubleshooting packages have been added to the RHV-H repository:

- tcpdump
- vim-enhanced

```
screen
strace/ltrace
wget
sysstat
dropwatch
systemtap
```

These packages can be installed on Red Hat Virtualization Host.

#### BZ#1405975

With this update, support for Red Hat Virtualization and oVirt has been added to Ansible. For more information about oVirt Ansible modules see [http://docs.ansible.com/ansible/list\\_of\\_cloud\\_modules.html#ovirt](http://docs.ansible.com/ansible/list_of_cloud_modules.html#ovirt)

#### BZ#1406814

This fix allows administrators to set the engine-config option "HostPackagesUpdateTimeInHours" to 0, which disables automatic periodical checks for host upgrades. Automatic periodical checks are not always needed, for example when managing hosts using Satellite.

#### BZ#1408193

From now on, all timestamp records for the engine and engine tools logs will contain a time zone to ease correlation between logs on the Manager and hosts. Previously engine.log contained a timestamp without a time zone, for example:

```
2017-02-27 13:35:06,720 INFO
[org.ovirt.engine.core.dal.dbbroker.DbFacade] (ServerService Thread Pool
-- 51) [] Initializing the DbFacade
```

From now on there will always be a timezone identifier at the end of the timestamp part, for example:

```
2017-02-27 13:35:06,720+01 INFO
[org.ovirt.engine.core.dal.dbbroker.DbFacade] (ServerService Thread Pool
-- 51) [] Initializing the DbFacade
```

#### BZ#1408876

This release enables Virtual Machines to lease areas on the storage domain. If a Virtual Machine has a lease on a storage domain, it will not be possible to move this storage domain into Maintenance mode. If the user attempts to do so, an error message will appear explaining that a virtual machine currently has a lease on this storage.

#### BZ#1411654

Previously, rhvm-appliance was not available via subscriptions on the RHV-H repositories. In this release, rhvm-appliance is now used as the preferred deployment mechanism by `ovirt-hosted-engine-setup`, and is now

available in the RHV-H repositories.

**BZ#1412095**

Previously, the Networking tab was not available in the Cockpit in Red Hat Virtualization Host, even though the NetworkManager was enabled. With this release, the Networking tab is now available in the Cockpit, and administrators can use it to configure the network.

**BZ#1412547**

Previously, when the Manager attempted to connect to VDSM it tried to negotiate the highest available version of TLS but due to previous issues there was a limitation to try TLSv1.0 as the highest version and to not try any higher version. Now, the limit has been removed so that TLSv1.1 and TLSv1.2 can be negotiated if they are available on the VDSM side. Removing this limit will allow TLSv1.0 to be dropped from future versions of VDSM.

**BZ#1413150**

The Red Hat Virtualization Manager now provides warnings for all data centers and clusters that have not been upgraded to latest installed version. The compatibility version of all data centers is checked once a week and on Manager startup. If it is not the latest version, an alert is raised and stored in the audit log. The Data Centers and Clusters main tabs now also show an exclamation mark icon for each data center or cluster that is not at the latest version. Hovering over this icon displays a recommendation to upgrade the compatibility version.

**BZ#1413181**

In RHV 4.1 a new tools repository containing packages required by the Red Hat Virtualization Manager has been added. See the Release Notes or Installation Guide for repository details.

**BZ#1418145**

The 'localdisk' hook adds the ability to use fast local storage instead of shared storage, while using shared storage for managing virtual machine templates. Currently, a user has to choose between fast local storage, where nothing is shared with other hosts, or shared storage where everything is shared between the hosts and fast local storage cannot be used. This update mixes local and shared storage.

The 'localdisk' hook works as follows:

- 1) A user will create a virtual machine normally on shared storage of any type. To use the virtual machine with local storage, the user will need to pin the virtual machine to a certain host and enable the localdisk hook.

- 2) When starting the virtual machine on the pinned host, the localdisk

hook will copy the virtual machine disks from shared storage into the host's local storage and modify the disk path to use the local copy of the disk.

3) The original disk may be a single volume or a chain of volumes based on a template. The local copy is a raw preallocated volume using a LVM logical volume on the special "ovirt-local" volume group.

To change storage on a virtual machine using local storage, the localdisk hook must be disabled.

Warnings:

- Virtual machines using local disk must be pinned to a specific host and cannot be migrated between hosts.
- No storage operations on a virtual machines using local disks are allowed, for example, creating/deleting snapshots, moving disks, creating templates from the virtual machine.
- The virtual machine disks on the shared storage should not be deleted and the storage domain needs to be active and accessible.

#### **BZ#1422024**

With this release, the net-snmp package is part of the Red Hat Virtualization Host image by default.

#### **BZ#1424821**

It is now possible to create NFS storage domains with NFS version 4.2 via the REST API.

#### **BZ#1427987**

The API supports the 'filter' parameter to indicate if results should be filtered according to the permissions of the user. Due to the way this is implemented, non-admin users need to set this parameter for almost all operations because the default value is 'false'. To simplify things for non-admin users, a configuration option `ENGINE_API_FILTER_BY_DEFAULT` has been added, which allows you to change the default value to 'true', but only for non-admin users. If the value is explicitly given in a request, it will be honored.

If you change the value of `ENGINE_API_FILTER_BY_DEFAULT` to 'true', be aware that this is not backwards compatible, as clients that used non-admin users and did not explicitly provide the 'filter' parameter will start to behave differently. However, this is unlikely because calls from non-admin users without 'filter=true' are almost useless.

If it is necessary to change the default behavior, it can be achieved by changing this parameter in a configuration file inside the '/etc/ovirt-engine/engine.conf.d' directory. For example:

```
# echo 'ENGINE_API_FILTER_BY_DEFAULT="true"' > \
```

```
/etc/ovirt-engine/engine.conf.d/99-filter-by-default.conf
```

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

### 3.1.2. Technology Preview

The items listed in this section are provided as Technology Previews. For further information on the scope of Technology Preview status, and the associated support implications, refer to <https://access.redhat.com/support/offerings/techpreview/>.

#### **BZ#1400569**

In this release, a new user interface for the User Portal has been introduced as a Technology Preview. The new user interface offers improved performance. The new User Portal can be accessed from the following link: [https://\[ENGINE\\_HOST\]/ovirt-engine/web-ui](https://[ENGINE_HOST]/ovirt-engine/web-ui)

### 3.1.3. Release Notes

This section outlines important details about the release, including recommended practices and notable changes to Red Hat Virtualization. You must take this information into account to ensure the best possible outcomes for your deployment.

#### **BZ#1364978**

Previously, Red Hat Virtualization Host (RHVH) was shipped without an End User License Agreement (EULA). In this release, this bug has been fixed and RHVH now includes an EULA.

#### **BZ#1375563**

Currently, virtual machines with block devices cannot be imported from Xen to Red Hat Virtualization using the Administration Portal. This release adds a workaround to manually import Xen virtual machines with block devices:

1. Create an export domain.
2. Run the following command to copy the files locally:  

```
$ virt-v2v-copy-to-local -ic xen+ssh://root@xenserver.com vmname
```
3. Run the following command to convert the virtual machine to libvirt XML:  

```
$ virt-v2v -i libvirtxml vmname.xml -o rhev -of raw -os servername:/path/to/export/domain
```
4. The virtual machine should now appear in your export domain, and can be imported to a data domain.

#### **BZ#1394831**

RHEV Agent has been renamed to oVirt Agent on Windows.

#### **BZ#1405612**

-

rhevmspice-client packages were renamed to spice-client-msi.

#### BZ#1414085

Previously the name of the “rhevmspice-client” RPM contained only a timestamp, without versioning information. In this release, the Red Hat Virtualization release will now be included in the name of the “rhevmspice-client” RPM and will be visible from the node channel.

### 3.1.4. Known Issues

These known issues exist in Red Hat Virtualization at this time:

#### BZ#1326709

When hosted-engine --deploy is run on additional hosts that have multiple FQDNs associated with it, the script will pick the host address that is returned by default.

Ensure that a host's hostname resolves to the required FQDN/em1 before deploying the self-hosted engine.

#### BZ#1416153

Previously, after deleting a snapshot in a data center, the original volume's allocation policy and size differed from the pre-snapshot state. In this release, if a snapshot is created from a preallocated volume, when the snapshot is deleted, qemu-img is called to copy data from the top volume to the base volume. As a result, the original volume's allocation policy and size are identical.

#### BZ#1422442

Due to an unstable slave order in NetworkManager, DHCP over a bond created by NetworkManager may receive a different IP address after adding it to Red Hat Virtualization (RHV) or after rebooting. The workaround is to avoid using DHCP on a NetworkManager-controlled bond.

NetworkManager may also remove a DHCP-provided host name after a host is added to RHV. To avoid this, persist the host name explicitly via Cockpit or hostnamectl.

#### BZ#1433434

In Red Hat Virtualization 4.1, when the Manager deploys a host, collectd is always installed; however, host deployment will fail if you are attempting to deploy a new or reinstalled version 3.y host (in a cluster with 3.6 compatibility level), because collectd is not shipped in the 3.y repositories.

To avoid this, ensure that you install and deploy any version 3.y hosts prior to upgrading the Manager to 4.1.

Note that after the Manager upgrade, these hosts will continue to work, but you will not be able to reinstall them without first upgrading them to version 4.1.

**BZ#1441711**

Cockpit is currently available only for x86\_64 architecture. As a result, in Red Hat Virtualization, Cockpit is supported only for x86\_64 hosts, and is not supported for ppc64le (IBM POWER8) hosts.

**BZ#1442254**

File conflicts, as a result of package renaming, caused direct upgrades of debuginfo from version 1.0.12 and earlier to a later version to fail due to duplicate files in the same location.

The workaround for this problem is to manually uninstall the previous version of the debuginfo package before an installation or upgrade of the newer ovirt-guest-agent packages, such as:

```
# yum remove rhvm-guest-agent-debuginfo
```

This workaround does not introduce any limitations and is simple to execute for users of the debuginfo packages.

### 3.1.5. Deprecated Functionality

The items in this section are either no longer supported or will no longer be supported in a future release

**BZ#1219545**

This release removes the ability to export Gluster volume profile statistics as a PDF file (a feature that was not widely used) as part of removing the dependency on the avalon-framework package.

**BZ#1350334**

The rhvm-guest-agent packages for Red Hat Enterprise Linux have now been renamed to ovirt-guest-agent, to align with upstream.

**BZ#1372237**

This release removes a no-longer-needed workaround for the vds-m-jsonrpc deprecation warning.

**BZ#1429809**

IFCFG persistence mode has been declared deprecated. The Unified persistence mode has been the default from version 3.5 and should now be used in all systems.



## 3.2. RED HAT VIRTUALIZATION BATCH UPDATE 1 (OVIRT-4.1.2)

These release notes highlight Technology Preview items, recommended practices, known issues, and deprecated functionality to be taken into consideration when deploying this release of Red Hat Enterprise Virtualization.

Notes for updates released during the support lifecycle of this Red Hat Enterprise Virtualization release will appear in the advisory text associated with each update or the *Red Hat Enterprise Virtualization Technical Notes*. This document is available from the following page:

[https://access.redhat.com/site/documentation/en-US/Red\\_Hat\\_Enterprise\\_Virtualization](https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Virtualization)

### 3.2.1. Enhancements

This release of Red Hat Enterprise Virtualization features the following enhancements:

#### BZ#1360983

With this update, the host name is automatically set to the virtual machine name in the RunOnce dialog. If required, this name can be changed.

#### BZ#1388433

Previously, the Manager's frequently updated PostgreSQL tables would fill up with obsolete data, creating the risk of disk flooding and transaction ID wraparound issues.

This release has introduced a more aggressive vacuum daemon configuration for collecting obsolete rows. This enables the Manager's tables to remain healthy, and for disk space usage to be better correlated with the actual amount of kept data.

#### BZ#1406398

Red Hat Virtualization Manager now supports NFS version 4.2 connections (when also supported by the storage).

#### BZ#1421533

In this release, the vdsclient has been added as a dependency, and it replaces the now deprecated vdsClient.

#### BZ#1427790

Previously, it was possible to limit the highest SSL/TLS protocol version that was negotiated when establishing a connection between the Manager and VDSM. This was required for older clients.

In this release, this option has been removed from engine-config as it was verified that it is no longer required for VDSM 3.6 and later. VDSM 3.6 can successfully negotiate the highest available version.

**BZ#1443508**

Previously, `ovirt-host-deploy` failed if the `collectd/fluentd` packages and their plug-ins were not available.

In this release, if these packages are missing, `ovirt-host-deploy` emits a warning but does not fail. This enables a 3.6 host to be added to a 4.1 Manager, even though the 3.6 repositories do not include these packages.

**BZ#1446056**

Previously, after running `engine-upgrade-check`, the user was not informed that the system may not be up-to-date if `engine-setup` was not run after running `yum update`, for example. This is despite `engine-upgrade-check` stating that no upgrade is available. This has now been fixed so that a warning message is displayed to the user if `engine-setup` was not run.

### 3.2.2. Release Notes

This section outlines important details about the release, including recommended practices and notable changes to Red Hat Enterprise Virtualization. You must take this information into account to ensure the best possible outcomes for your deployment.

**BZ#1367107**

Previously, the API and Administration Portal differed in their firewall configuration when installing new hosts or reinstalling existing ones:

- Administration Portal - The firewall configuration was enabled by default.
- API - The firewall configuration was disabled by default and customers were required to specify the `override_iptables=true` parameter to complete the installation successfully. Alternatively, they had to manually configure the firewall prior to adding it to the Manager.

Now, both the API and the Administration Portal host installations enable firewall configuration by default.

Note that for customers upgrading to this release who are using API scripts to add hosts, if these scripts rely on the pre-4.1.2 API behavior to not configure the firewall, they must add `override_iptables=false` to their scripts.

**BZ#1433961**

Previously, it was possible to configure memory overcommit without setting memory ballooning or KSM control. This configuration affected the scheduling, but the memory was not freed. This has now been fixed by disabling ballooning and KSM by default, and setting memory optimization to "None" (100%).

**BZ#1436161**

The recommended number of logical volumes per storage domain was increased from 300 to 1000. The alert that is issued when exceeding this number was changed to reflect that.

Note that the limit will only change to 1000 if it was set to 300. If the customer manually configured the limit, the original configuration will not be overridden.

### 3.3. RED HAT VIRTUALIZATION BATCH UPDATE 2 (OVIRT-4.1.3)

#### 3.3.1. Enhancements

This release of Red Hat Virtualization features the following enhancements:

**BZ#1408825**

With this update, the ability to unplug a lease from and plug another lease into a running virtual machine has been added using new APIs. This provides the ability to move a virtual machine lease from one storage domain to another so that the original storage domain can be placed into maintenance.

**BZ#1418659**

With this update, engine.log has been added to the logs collection. The records are now parsed and sent to the central metrics store. This allows the administrator to analyse the logs in a more simple and comfortable way. This means that the engine.log records are available for analysis in Kibana.

**BZ#1421204**

Previous versions of Red Hat Virtualization were able to negotiate encryption protocol up to TLSv1 for external provider communication. This update adds the the ability to negotiate encryption protocol up to TLSv1.2.

Note: The exact version used for communication depends on highest version available on the external provider target.

**BZ#1429861**

With this update, ovirt-engine-metrics can now configure collectd to connect to the Red Hat Virtualization Manager's PostgreSQL database and get information from it. Previously, making collectd connect to the PostgreSQL database did not work and was disabled by BZ#1436001.

**BZ#1431545**

With this update, GWT symbol maps, used to clarify obfuscated client-side (GWT UI) stack traces (`/var/log/ovirt-engine/ui.log`), are now part of the core Red Hat Virtualization Manager UI packages. The packages are installed by default when the Red Hat Virtualization Manager is installed or upgraded.

This update means that the user now has meaningful client-side stack traces in `/var/log/ovirt-engine/ui.log` without having to install additional packages, which saves time when analyzing UI-related issues. This update also reduces the symbol maps' size by zipping them. The Administration Portal UI uncompressed GWT symbol maps previously took ~1 GB of disk space. The zipped version consumes only 50 MB. The content of the zipped GWT symbol map files is streamed by the Manager as required and the content is not physically extracted to disk. After the Manager install or upgrade, the GWT symbol maps files use the following location:

```
/usr/share/ovirt-engine/gwt-symbols/webadmin/symbolMaps.zip
```

NOTE: The `webadmin-portal-debuginfo` package, which previously provided the GWT symbol map files, must now be manually removed prior to engine upgrade, or the upgrade will fail. (Because this package is optional, you only need to remove it if you installed it manually.)

#### **BZ#1436981**

The Python SDK now supports asynchronous requests and HTTP pipe-lining. Users can send requests asynchronously and wait for the response later in code. This makes it possible to send multiple requests using multiple connections or pipelined connections, and wait for the response later, which improves performance when fetching multiple objects from the API.

#### **BZ#1439281**

The self-hosted engine Manager virtual machine upgrade procedure from Red Hat Enterprise Virtualization 3.6 on Red Hat Enterprise Linux 6 to Red Hat Virtualization 4.0 on Red Hat Enterprise Linux 7 now always tries to connect to the Manager API as `admin@internal`, making it more flexible.

#### **BZ#1444992**

This update adds a new VDSM hook called `vdsm-hook-vfio-mdev`, which makes the host capable of working with mdev capable devices such as GPUs. The hook is automatically installed. The following prerequisites apply for a virtual machine to use an mdev instance:

- 1) The host must have a kernel that supports the mediated device, and it must have the device and the correct drivers.
- 2) The predefined property, `'mdev_type'`, must be set and correspond to one of the mdev types supported by the device.
- 3) The virtual machine must be pinned to the host(s) with the device(s).

The supported `mdev_type` values can be determined by checking the host on

which the device is present (not visible in the engine) by querying vdsm (vdsm-client Host hostdevListByCaps) or running the following bash script:

```
for device in /sys/class/mdev_bus/*; do
  for mdev_type in $device/mdev_supported_types/*; do
    MDEV_TYPE=$(basename $mdev_type)
    DESCRIPTION=$(cat $mdev_type/description)
    echo "mdev_type: $MDEV_TYPE --- description: $DESCRIPTION";
  done;
done | sort | uniq
```

**BZ#1445631**

With this update, the user can now control the location of the temporary directory created by the engine-backup command, using the command line option `--tmpdir=DIR` and the variable `$TMPDIR`. Previously, the engine-backup command always used `/tmp`. If `/tmp` was full, or close to being full, it was not possible to tell the engine-backup command to use a different directory.

**BZ#1447887**

Red Hat Virtualization Host (RHVH) now supports NIST 800-53 partitioning requirements to improve security. RHVH uses a NIST 800-53 partition layout by default, and existing configurations will be changed on update.

**BZ#1450646**

This update makes it easier to debug a disk snapshot's live deletion failures because the initial state of the volume chain is logged.

**BZ#1455771**

When collecting SOS reports from hosts, `chrony` and `systemd` SOS plugins can collect information about time synchronization. In addition, a new option `--time-only` has been added to `ovirt-log-collector` allowing information about time differences to be gathered from the hosts without gathering full SOS reports, saving a considerable amount of time for the operation.

**BZ#1456877**

Previously, the self-hosted engine setup failed if the system clock was set to a time zone that was behind UTC during the installation. This is due to the fact that the Red Hat Virtualization Host generates VDSM certificates at the first boot and if the clock is incorrect, the `chronyd` or `ntpd` processes resynchronized the clock. This led to an invalid certificate if the time zone was behind UTC.

Now, Red Hat Virtualization Host generates the certificates after the `chronyd` or `ntpd` processes and waits two seconds for the clock to

synchronize.

Note that if the Red Hat Virtualization Host is configured after the installation, or if the NTP server is too slow, the self-hosted engine setup may fail.

### 3.3.2. Technology Preview

The items listed in this section are provided as Technology Previews. For further information on the scope of Technology Preview status, and the associated support implications, refer to <https://access.redhat.com/support/offerings/techpreview/>.

#### **BZ#1456568**

This release introduces a new VM Portal as a Technology Preview. A "VM Portal" link is now available on the Red Hat Virtualization Welcome Page. The VM Portal provides the same functionality that is available in the Basic tab of the current User Portal (now deprecated).

### 3.3.3. Known Issues

These known issues exist in Red Hat Virtualization at this time:

#### **BZ#1455441**

Previously, the self-hosted engine setup failed if the system clock was set to a time zone that was behind UTC during the installation. This is due to the fact that the Red Hat Virtualization Host generates VDSM certificates at the first boot and if the clock is incorrect, the chronyd or ntpd processes resynchronized the clock. This led to an invalid certificate if the time zone was behind UTC.

Now, Red Hat Virtualization Host generates the certificates after the chronyd or ntpd processes and waits two seconds for the clock to synchronize.

Note that if the Red Hat Virtualization Host is configured after the installation, or if the NTP server is too slow, the self-hosted engine setup may fail.

### 3.3.4. Deprecated Functionality

The items in this section are either no longer supported or will no longer be supported in a future release

#### **BZ#1441632**

With this update, the migration of a virtual machine to a different cluster can no longer be invoked from the UI. Regular migration within a cluster remains unchanged.

#### **BZ#1464113**

The User Portal is now deprecated. From Red Hat Virtualization 4.2, it will no longer be available. A new VM Portal will be added, but certain features may not be present at the time of release. A technology preview of the VM Portal is available with the 4.1.3 release. Customers utilizing the Red Hat CloudForms Self Service user interface to manage virtual machines will not be affected.

## 3.4. RED HAT VIRTUALIZATION BATCH UPDATE 3 (OVIRT-4.1.4)

### 3.4.1. Enhancements

This release of Red Hat Virtualization features the following enhancements:

#### BZ#917062

ABRT is now integrated with Red Hat Virtualization when initializing hypervisors. ABRT is configured by VDSM and saves meaningful debug information, instead of enormous core-dump reports.

#### BZ#1438408

The ClusterInMaintenance scheduling policy was added to limit activity in a cluster during maintenance tasks. When set, no new VMs may be started, except highly available VMs. Highly available VMs are still restarted properly upon host failure and any VM can be migrated. Users can create new highly available VMs and start them manually.

#### BZ#1450293

Newer browsers require the subjectAltName extension in certificates for HTTPS. Previously (version 4.1.2 and later), engine-setup created certificates with subjectAltName for new setups, but did not update existing certificates during upgrade.

Now engine-setup can add subjectAltName to existing internal certificates, so that newer browsers will accept them.

#### BZ#1452668

Previously, administrators had to enter an unencrypted password when invoking 'ovirt-aaa-jdbc-tool user password-reset'. The password was then encrypted inside ovirt-aaa-jdbc-tool and stored in the database.

This update enables administrators to use the new --encrypted option to enter an already encrypted password when invoking 'ovirt-aaa-jdbc-tool user password-reset'.

However there are some caveats when providing encrypted passwords:

1. Entering an encrypted password means that password validity tests cannot be performed, so they are skipped and the password is accepted even if it does not comply with the password validation policy.

2. A password has to be encrypted using the same configured algorithm. To encrypt passwords, administrators can use the `'/usr/share/ovirt-engine/bin/ovirt-engine-crypto-tool.sh'` tool, which provides the `'pbe-encode'` command to encrypt passwords using the default PBKDF2WithHmacSHA1 algorithm.

**BZ#1461813**

RHV-H now supports vGPU for nVidia GPUs, by including `vdsm-hook-vfio-mdev`.

**BZ#1465862**

Previously, users could use the REST API to add/edit affinity labels for virtual machines and hosts, but they could not view this information in the Administration Portal.

In this release, affinity label information is now available in the Administration Portal. It can be seen in the Add/Edit windows for virtual machines and hosts. The window displays all available affinity labels in the system, and for entities with assigned labels, those labels are selected.

**BZ#1467058**

Previously, when creating a new VLAN-tagged network, the Manager did not allow the same VLAN ID to be duplicated across various entities, such as networks and data centers.

In this release, it is now possible to use the same VLAN ID across multiple entities, as long as they are not used on the same interface on the hardware node.

### 3.4.2. Release Notes

This section outlines important details about the release, including recommended practices and notable changes to Red Hat Virtualization. You must take this information into account to ensure the best possible outcomes for your deployment.

**BZ#1430285**

In order to run VMs on hosts with limited memory resources, the cluster scheduling policy should be set to `"evenly_distributed"`, with `"maxFreeMemoryForOverUtilized = 99"`. This configuration enables the engine to schedule VMs on hosts with more free memory, for better distribution.

## 3.5. RED HAT VIRTUALIZATION BATCH UPDATE 4 (OVIRT-4.1.5)



### 3.5.1. Enhancements

This release of Red Hat Virtualization features the following enhancements:

#### BZ#1022961

This release adds libgfapi support to the Manager and VDSM. libgfapi provides virtual machines with faster access to their images, stored on a Gluster volume, compared to a fuse interface. With the 'LibgfApi' data center feature enabled, or 'lubgfapi\_supported' cluster-level feature enabled, virtual machines will access their images, stored on Gluster volumes, directly via libgfapi.

#### BZ#1364083

With this update, oVirt Host related packages are now also available for ppc64le architecture.

#### BZ#1412552

VDSM now supports TLSv1.1 and TLSv1.2, enabling all clients to leverage these new encryption protocols, depending on their capabilities.

#### BZ#1451490

This release adds the list of required collectd and fluentd packages to the metrics script, and installs them if missing. Previously, there were cases where the packages were available in the repository and required for the metrics setup script, but were not installed, causing the configuration to fail. Now, the missing packages are installed.

#### BZ#1454827

Previously, Red Hat Virtualization supported TLS version 1.0. From this release, all components can use TLS version 1.2 for encrypted communication, providing that it is used by all components.

#### BZ#1457357

You can now change more self-hosted engine configuration in the hosted-engine.conf file.

For example, to set the value:

```
hosted-engine --set-shared-config gateway 10.35.1.254
```

To get the value:

```
hosted-engine --get-shared-config gateway
```

The commands can be used with the additional parameter "type", if there is concern for duplicate keys in different files:

```
hosted-engine --set-shared-config gateway 10.35.1.254 --type=he_conf
```

```
hosted-engine --get-shared-config gateway --type=he_conf
```

**BZ#1462500**

This release adds checks to the `ovirt_env_name`, to make sure it is a valid OpenShift namespace identifier. If it is not, the metrics script will now fail. The name must be valid in order for the records to be kept by the Elasticsearch running in OpenShift.

### 3.5.2. Release Notes

This section outlines important details about the release, including recommended practices and notable changes to Red Hat Virtualization. You must take this information into account to ensure the best possible outcomes for your deployment.

**BZ#1472047**

The `batik` package is no longer required by the Red Hat Virtualization Manager. Red Hat recommends removing it after upgrading to Red Hat Virtualization 4.1.

**BZ#1472048**

The `fop` package is no longer required by the Red Hat Virtualization Manager. Red Hat recommends removing it after upgrading to Red Hat Virtualization 4.1.

### 3.5.3. Known Issues

These known issues exist in Red Hat Virtualization at this time:

**BZ#1351211**

When removing a Gluster host after moving it to maintenance mode, the host is not removed from the peer list of other hosts. To work around this issue, do not stop the Gluster services while moving the host to maintenance mode if the host is going to be removed from cluster.

### 3.5.4. Deprecated Functionality

The items in this section are either no longer supported or will no longer be supported in a future release

**BZ#1473179**

The `ovirt-engine-setup-plugin-dockerc` package is now deprecated, and will be removed in version 4.2.

## 3.6. RED HAT VIRTUALIZATION BATCH UPDATE 5 (OVIRT-4.1.6)

### 3.6.1. Enhancements

This release of Red Hat Virtualization features the following enhancements:

**BZ#1472254**

When configuring Active Directory (AD) with the `ovirt-engine-extension-aaa-ldap-setup` tool, regardless of whether you are defining a multiple or single domain forest, you can only configure the name of the forest - you cannot define the name of a specific domain or a specific server.

This release provides examples for common advanced AD configuration which users can copy to their local environment and adapt as required. These examples are bundled within the `ovirt-engine-extension-aaa-ldap` package and after installing the package, a description of these examples can be found in `/usr/share/ovirt-engine-extension-aaa-ldap/examples/README.md`

In addition, the following improvements have been made to the `ovirt-engine-extension-aaa-ldap-setup` tool:

1. A more detailed error reporting for various AD forest configuration steps.
2. A mandatory login test that checks the configuration (previously this test was optional).

**BZ#1473535**

The `ovirt-ansible-roles` package contains Ansible roles, which can help administrators with common Red Hat Virtualization administration tasks. All roles can be executed from the command line using Ansible, but some of those roles can also be executed directly from the Red Hat Virtualization Manager. More details about the roles can be found in `README.md` included in the package (`/usr/share/doc/ovirt-ansible-roles/README.md`), or directly in the source code repository (<https://github.com/ovirt/ovirt-ansible>).

**BZ#1477053**

This update lists the additional features available for a cluster level, and allows you to enable or disable additional cluster features using an API.

**BZ#1479677**

This update adds support for Link Layer Discovery Protocol (LLDP) in Red Hat Virtualization.

To support the integration of Red Hat Virtualization with other software, the information gathered by LLDP on each host will be provided by the Manager's REST API. Information such as the switch port number the host's network interface is connected to, will be shown in the GUI.

Known Issue BZ#1487930:

Cisco switches might send VLAN tagged 802.1Q LLDP frames. These LLDP frames are ignored if the NIC of the Red Hat Virtualization host is not attached to the corresponding Red Hat Virtualization logical network.

**BZ#1481095**

The user now has the option of specifying a path to the appliance OVF, as an alternative to installing the appliance RPM.

**BZ#1482104**

The `ovirt-ansible-roles` package contains Ansible roles, which can help administrators with common Red Hat Virtualization administration tasks. All roles can be executed from the command line using Ansible, but some of those roles can also be executed directly from the Red Hat Virtualization Manager. More details about the roles can be found in `README.md` included in the package (`/usr/share/doc/ovirt-ansible-roles/README.md`), or directly in the source code repository (<https://github.com/ovirt/ovirt-ansible>).

Please note that currently Ansible has the following bug:

[https://bugzilla.redhat.com/show\\_bug.cgi?id=1487113](https://bugzilla.redhat.com/show_bug.cgi?id=1487113)

Because of this bug you should avoid calling multiple oVirt Ansible roles in a single playbook, or calling oVirt Ansible modules prior to calling oVirt Ansible roles in a single playbook (the working solution is to call the modules from the `pre_tasks`).

**BZ#1486573**

This Ansible role can be used to automate upgrading the whole cluster to a higher version. More information can be found at <https://github.com/oVirt/ovirt-ansible/blob/ovirt-ansible-1.0/roles/ovirt-cluster-upgrade/README.md>

Please note that currently Ansible has the following bug:

[https://bugzilla.redhat.com/show\\_bug.cgi?id=1487113](https://bugzilla.redhat.com/show_bug.cgi?id=1487113)

Because of this bug you should avoid calling multiple oVirt Ansible roles in a single playbook, or calling oVirt Ansible modules prior to calling oVirt Ansible roles in a single playbook (the working solution is to call the modules from the `pre_tasks`).

**BZ#1486574**

This Ansible role can be used to automate data center setup. More information can be found at <https://github.com/oVirt/ovirt-ansible/blob/ovirt-ansible-1.0/roles/ovirt-infra/README.md>

Please note that currently Ansible has the following bug:

[https://bugzilla.redhat.com/show\\_bug.cgi?id=1487113](https://bugzilla.redhat.com/show_bug.cgi?id=1487113)

Because of this bug you should avoid calling multiple oVirt Ansible roles in a single playbook, or calling oVirt Ansible modules prior to calling oVirt Ansible roles in a single playbook (the working solution is to call the modules from the `pre_tasks`).

**BZ#1486575**

This Ansible role can be used to automate creating a VM template from an external image. More information can be found at <https://github.com/oVirt/ovirt-ansible/blob/ovirt-ansible-1.0/roles/ovirt-image-template/README.md>

Please note that currently Ansible has the following bug: [https://bugzilla.redhat.com/show\\_bug.cgi?id=1487113](https://bugzilla.redhat.com/show_bug.cgi?id=1487113)  
Because of this bug you should avoid calling multiple oVirt Ansible roles in a single playbook, or calling oVirt Ansible modules prior to calling oVirt Ansible roles in a single playbook (the working solution is to call the modules from the `pre_tasks`).

**BZ#1486576**

This Ansible role can be used to download a ManageIQ qcow2 image, create a VM from it, and register a provider within the ManageIQ installation. More information can be found at <https://github.com/oVirt/ovirt-ansible/blob/ovirt-ansible-1.0/roles/ovirt-manageiq/README.md>

Please note that currently Ansible has the following bug: [https://bugzilla.redhat.com/show\\_bug.cgi?id=1487113](https://bugzilla.redhat.com/show_bug.cgi?id=1487113)  
Because of this bug you should avoid calling multiple oVirt Ansible roles in a single playbook, or calling oVirt Ansible modules prior to calling oVirt Ansible roles in a single playbook (the working solution is to call the modules from the `pre_tasks`).

**BZ#1486577**

This Ansible role can be used to automate creating a virtual machine. More information can be found at <https://github.com/oVirt/ovirt-ansible/blob/ovirt-ansible-1.0/roles/ovirt-vm-infra/README.md>

Please note that currently Ansible has the following bug: [https://bugzilla.redhat.com/show\\_bug.cgi?id=1487113](https://bugzilla.redhat.com/show_bug.cgi?id=1487113)  
Because of this bug you should avoid calling multiple oVirt Ansible roles in a single playbook, or calling oVirt Ansible modules prior to calling oVirt Ansible roles in a single playbook (the working solution is to call the modules from the `pre_tasks`).

### 3.6.2. Release Notes

This section outlines important details about the release, including recommended practices and notable changes to Red Hat Virtualization. You must take this information into account to ensure the best possible outcomes for your deployment.

**BZ#1471615**

This release includes a new Spice QXL WDDM DOD driver (Display-only driver for QXL virtual devices) for Windows 10 guests.

**BZ#1472047**

The batik package is no longer required by the Red Hat Virtualization Manager. Red Hat recommends removing it after upgrading to Red Hat Virtualization 4.1.

**BZ#1472048**

The fop package is no longer required by the Red Hat Virtualization Manager. Red Hat recommends removing it after upgrading to Red Hat Virtualization 4.1.

**BZ#1481150**

A new VDSM hook that configures the nested virtualization feature in Red Hat Enterprise Linux 7 is now available as a tech preview. Nested virtualization allows a virtual machine to serve as a hypervisor. VDSM hooks are a means to insert code, commands, or scripts into a point in the lifecycle of a virtual machine, as well as the lifecycle of the VDSM daemon.

Packages and features released under tech preview are not fully supported under Red Hat Subscription Level Agreements, and may not be functionally complete, and are not intended for production use.

**BZ#1481212**

Added the following changes to the GUI in this update:

- Ability to change the virtual machine's CD-ROM ISO.
- Infinite scroller for VMS list added to decrease response time.
- Added a hyperlink on "Runs on Host" to the host the running the particular virtual machine.
- Added "Loading ..." indicator during background activity.
- "Pending Changes" tag rendered when NEXT\_RUN configuration exists.
- "VM Type" field added.
- Improved error messages.
- Added a "usb-filter" in the .vv file for remote-viewer.
- The setting of the Virtual Machine icon has been fixed.
- Added multiple additional functional and user interface fixes.
- Application supports branding.
- Basic support for internationalization, but translation not yet provided.
- Built as a "noarch" RPM.

### 3.6.3. Known Issues

These known issues exist in Red Hat Virtualization at this time:

**BZ#1478865**

On Red Hat Enterprise Linux, consistent network device naming [1] is used to generate consistent names for network devices. On some systems, the generated name is too long, limiting automated tail additions to be

applied (the classic example is to create a VLAN interface over such an interface).

In cases where the long-name interface is used as a base of a VLAN, the total VLAN interface name may exceed the kernel name limit (15). As a result, the creation of the vlan fails. For example,  $\text{len}(\langle\text{base-name}\rangle + 1(\text{for the dot}) + \text{len}\langle\text{vlan-id}\rangle + 1) \leq 15$

The simplest workaround is to disable the new name mechanism. Alternatively, add a static udev rule to manipulate the name of the interfaces:

```
ACTION=="add", SUBSYSTEM=="net", DRIVERS=="?*"  
ATTR{address}=="00:50:56:8e:12:34", NAME="eth123"
```

[1] [https://access.redhat.com/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/7/html/Networking\\_Guide/ch-Consistent\\_Network\\_Device\\_Naming.html](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Networking_Guide/ch-Consistent_Network_Device_Naming.html)

## 3.7. RED HAT VIRTUALIZATION BATCH UPDATE 6 (OVIRT-4.1.7)

These release notes highlight technology preview items, recommended practices, known issues, and deprecated functionality to be taken into consideration when deploying this release of Red Hat Enterprise Virtualization.

Notes for updates released during the support lifecycle of this Red Hat Enterprise Virtualization release will appear in the advisory text associated with each update or the *Red Hat Enterprise Virtualization Technical Notes*. This document is available from the following page:

[https://access.redhat.com/site/documentation/en-US/Red\\_Hat\\_Enterprise\\_Virtualization](https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Virtualization)

### 3.7.1. Enhancements

This release of Red Hat Enterprise Virtualization features the following enhancements:

#### BZ#1490810

Previously, while VDSM performed periodic monitoring and maintenance tasks, operations would occasionally become too slow or even become blocked. In this case, the internal queue became full and the periodic operations were no longer performed. A "TooManyTasks" warning appeared in the log files at a maximum rate of once every 10 seconds.

In this release, if VDSM cannot perform its periodic operations, in addition to issuing a warning in the log files, VDSM also dumps the contents of the queue into the logs.

#### BZ#1502510

The precision of rx\_rate, tx\_rate, rx\_drop, and tx\_drop of virtual and host network interfaces have been increased. Network traffic 100 times smaller can now be detected on network interface statistics.

If traffic on the network interface is below the precision of the network interface statistics, it is not reflected in the statistics.

**BZ#1502653**

This release adds caching of the OVF storage location. As the OVF storage location rarely changes, it does not need to be searched for on every monitoring loop iteration. Instead it can be saved and reused, and expired only in the case of an error. As a result, the monitoring loop execution time is decreased significantly.

### 3.7.2. Known Issues

These known issues exist in Red Hat Enterprise Virtualization at this time:

**BZ#1399766**

If no self-hosted engine host is available other than the host that is currently running the self-hosted engine, you will not be able to move the host that runs the Manager virtual machine to maintenance mode. Furthermore, even after another host moves to a status of "Up", it could take a few minutes for it to receive a score that will enable it to run the Manager virtual machine.

The host will remain in a "preparing for maintenance" state until the Manager virtual machine can migrate to another host.

## 3.8. RED HAT VIRTUALIZATION BATCH UPDATE 7 (OVIRT-4.1.8)

### 3.8.1. Enhancements

This release of Red Hat Virtualization features the following enhancements:

**BZ#1501793**

There are several cluster and host settings which require reinstallation of the host if changed. The requirement to reinstall was always mentioned in documentation and a WARNING event was raised. With this release, the Administration Portal now also shows an exclamation mark icon for each host that needs to be reinstalled. When an exclamation mark icon is shown, you can find the details about it in the Action Items section of the host's details view.

**BZ#1504691**

In this release, a refresh button has been added to the Dashboard tab in the Administration Portal to enable users to view up-to-date system summary information.

**BZ#1516194**



The ipa-client package is now installed on hosts, and is included in the Red Hat Virtualization Host image. This enables Cockpit certificate signing and SSO with Red Hat IdM, and adding the host to an IdM realm.

### 3.8.2. Release Notes

This section outlines important details about the release, including recommended practices and notable changes to Red Hat Virtualization. You must take this information into account to ensure the best possible outcomes for your deployment.

#### BZ#1513886

This update ensures that TLSv12 support is enabled by default and no manual configuration is required.

### 3.8.3. Known Issues

These known issues exist in Red Hat Virtualization at this time:

#### BZ#1455441

Previously, the self-hosted engine setup failed if the system clock was set to a time zone that was behind UTC during the installation. This is due to the fact that the Red Hat Virtualization Host generates VDSM certificates at the first boot and if the clock is incorrect, the chronyd or ntpd processes resynchronized the clock. This led to an invalid certificate if the time zone was behind UTC.

Now, Red Hat Virtualization Host generates the certificates after the chronyd or ntpd processes, and waits two seconds for the clock to synchronize.

Note that if the Red Hat Virtualization Host is configured after the installation, or if the NTP server is too slow, the self-hosted engine setup may fail.

## 3.9. RED HAT VIRTUALIZATION BATCH UPDATE 8 (OVIRT-4.1.9)

### 3.9.1. Enhancements

This release of Red Hat Virtualization features the following enhancements:

#### BZ#1484404

You can now use key=key=value format (where the value of a fence option key is itself in a key=value format) when setting fence agent options, in order to support old drac5 devices that require "ssh\_options=-oCiphers="+3des-cbc" in the Options field to enable no-longer-enabled-insecure ciphers for SSH connection.

**BZ#1518545**

otopi now optionally logs the list of network connections on the machine after failures. Enable this by installing the optional `otopi-debug-plugins` package. This can help debug failures to start network services due to "Address already in use" errors.

**BZ#1525473**

Proton is a high performance, lightweight messaging library. It can be used in the widest range of messaging applications including brokers, client libraries, routers, bridges, proxies, and more. Proton is based on the AMQP 1.0 messaging standard. Proton enables integration with the AMQP 1.0 ecosystem from any platform, environment, or language.

**BZ#1525933**

Katello Agent is included in RHV and RHVH, and installed on the hosts during deployment. Katello Agent provides a list of the installed RPMs, for better integration with Satellite.

**BZ#1528290**

otopi's yum plugin is now more verbose, also logging all RPM scriptlets' output. This can help debug failures caused by errors from such scriptlets.

**BZ#1531016**

This update adds IBRS-enabled CPU Models to clusters, which can be used to mitigate Spectre vulnerability in guests. See <https://access.redhat.com/solutions/3307851> for more information.

**BZ#1532180**

The Software subtab in the Host General pane now contains a new "Kernel Features" field showing the PTI, IBPB, and IBRS kernel features of the host system, if the information is available. The Virtual Machine General pane now contains a new "Guest CPU Type" field showing the guest machine CPU type. These can be used to easily recognize if meltdown fixes are applied. For more info on meltdown patches detection, see <https://access.redhat.com/solutions/3317841>.

**BZ#1533061**

When live or cold merge fails, snapshot disks may be left in an illegal state. If virtual machines with illegal snapshot disks are shut down, they will not re-start. Virtual machines with illegal snapshot disks are now marked with an exclamation mark and a warning message not to shut them down.

### 3.9.2. Known Issues

These known issues exist in Red Hat Virtualization at this time:

#### BZ#1516494

Previously, having an empty VM directory in the export domain caused an error when trying to read the VMs stored in it, resulting in endless SPM re-elections. Now, this issue no longer causes SPM re-election. To fix the VM list loading error, you must remove the empty directory manually.

### 3.9.3. Deprecated Functionality

The items in this section are either no longer supported or will no longer be supported in a future release

#### BZ#1529305

The kernel no longer supports the `/sys/block/dm-X/queue/discard zeroes data` file in sysfs. It is therefore no longer possible to determine whether discarded blocks from a block device returns zeros or the actual data. Therefore, the virtual machine disk properties "Wipe After Delete" and "Enable Discard" are no longer supported at the same time.

## 3.10. RED HAT VIRTUALIZATION BATCH UPDATE 9 (OVIRT-4.1.10)

### 3.10.1. Enhancements

This release of Red Hat Virtualization features the following enhancements:

#### BZ#1534421

Previously, the self-hosted engine did not support IBRS-compatible CPUs. In this release, IBRS-compatible CPUs are now supported by the self-hosted engine.

#### BZ#1547031

In this release, the `ENGINE_JBOSS_BLOCKING_TIMEOUT` option, which changes the default value of `jboss.as.management.blocking.timeout`, is exposed. To change the default value, create `/etc/ovirt-engine/engine.conf.d/99-jboss-blocking-timeout.conf` with the line, `"ENGINE_JBOSS_BLOCKING_TIMEOUT=NNN"`, where "NNN" is the number of seconds for the JBoss timeout.

## 3.11. RED HAT VIRTUALIZATION BATCH UPDATE 10 (OVIRT-4.1.11)

### 3.11.1. Release Notes

This section outlines important details about the release, including recommended practices and notable changes to Red Hat Virtualization. You must take this information into account to ensure the best possible outcomes for your deployment.

**BZ#1472047**

The batik package is no longer required by the Red Hat Virtualization Manager. Red Hat recommends removing it after upgrading to Red Hat Virtualization 4.1.

**BZ#1472048**

The fop package is no longer required by the Red Hat Virtualization Manager. Red Hat recommends removing it after upgrading to Red Hat Virtualization 4.1.

**BZ#1562074**

Do not upgrade to RHEL 7.5 if you have virtual machines using Direct LUN disks with the "Use Privileged SCSI I/O" option checked, to avoid unexpected behavior.