



Red Hat Virtualization 4.1

Technical Notes

Technical Notes for Red Hat Virtualization 4.1 and Associated Packages

Red Hat Virtualization 4.1 Technical Notes

Technical Notes for Red Hat Virtualization 4.1 and Associated Packages

Red Hat Virtualization Documentation Team
Red Hat Customer Content Services
rhev-docs@redhat.com

Legal Notice

Copyright © 2017 Red Hat, Inc..

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). 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 Technical Notes document provides information about changes made between release 4.0 and release 4.1 of Red Hat Virtualization. This document is intended to supplement the information contained in the text of the relevant errata advisories available via Red Hat Network.

Table of Contents

INTRODUCTION	2
CHAPTER 1. RHEA-2017:0997 RED HAT VIRTUALIZATION MANAGER 4.1 GA ENHANCEMENT (OVIRT-ENGINE)	3
CHAPTER 2. RHEA-2017:1114 RED HAT VIRTUALIZATION HOST 4.1 GA ENHANCEMENT	12

INTRODUCTION

These Technical Notes provide documentation of the changes made between release 4.0 and release 4.1 of Red Hat Virtualization. They are intended to supplement the information contained in the text of the relevant errata advisories available via Red Hat Network. Red Hat Virtualization 4.x errata advisories are available at <https://rhn.redhat.com/errata/>.

A more concise summary of the features added in Red Hat Virtualization 4.1 is available in the Red Hat Virtualization 4.1 Release Notes.

CHAPTER 1. RHEA-2017:0997 RED HAT VIRTUALIZATION MANAGER 4.1 GA ENHANCEMENT (OVIRT-ENGINE)

The bugs contained in this chapter are addressed by advisory RHEA-2017:0997. Further information about this advisory is available at <https://rhn.redhat.com/errata/RHEA-2017-0997.html>.

ovirt-engine

BZ#1341106

Previously, when virtual machines were stopped during a shutdown/reboot, the shutdown appeared to the VDSM to have been performed gracefully from within the guest Operating System. The Red Hat Virtualization Manager therefore did not start high availability virtual machines on a different host because it considered the stopping of the virtual machine to be user initiated.

This update ensures that VDSM will now detect that a virtual machine was shutdown from within the system, and can differentiate the unplanned shutdown and reports this information. As a result, virtual machines stopped on a system shutdown, are now restarted on a different system.

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#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#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#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#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#1356767

This update fixes an issue where V2V fails for a VMware external provider which contains special characters in a data center, cluster, or folder name. The characters are now encoded on the V2V engine side before sending them to the VDSM.

1) This solution encodes the VMware data center and VMware cluster parts of the uri. For example, in the URI `vpx://user@vcenter/<DATACENTER>/<CLUSTER>/esxi?no_verify=1` the `<DATACENTER>` and `<CLUSTER>` parts are encoded, including folders if they exist. Other parts of the URI are left decoded.

2) The UI will display the data center and cluster as was originally typed by the user.

3. The encoding is done on the frontend and should be also done while using REST APIs for VMware V2V.

The following special chars are yet not supported:
"%", "/", "\", "&", "+", "=", "?"

BZ#1366786

Previously, the update of the compatibility version of a cluster with many running virtual machines that are installed with the guest-agent caused a deadlock that caused the update to fail. In some cases, these clusters could not be upgraded to a later compatibility version. Now, the deadlock in the database has been prevented so that a cluster with many running virtual machines that are installed with the guest-agent can be upgraded to later compatibility version.

BZ#1372320

Updated messages written to the audit log. Previously backup execution start and finish events were identified only with "Started"/"Finished" values, so it was not obvious which process has started or finished. Events now have the format "Backup started at ..." and "Backup finished at ...".

BZ#1375076

Previously, when merging a snapshot to a base RAW volume, the RAW volume was not refreshed. This meant that merging a snapshot that was extended to a RAW base volume when the virtual machine was not running on the SPM

would fail. This update ensures that the VDSM refreshes the base volume as needed.

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#1382807

Previously, attempting to drag icons from elements that were not drag-and-drop enabled caused an exception to be thrown due to a missing empty check on a string pulled from the drag event data. An error would be displayed to the user, requiring a page refresh to clear it. Now, the string pulled from the drag event data is checked for nullness/emptiness, so the exception and error are avoided.

BZ#1387625

This update fixes a bug that was discovered in the sysprep templates for 64-bit platforms where some parts were improperly marked as "for 32-bit platform" and thus ignored by Windows.

BZ#1388963

With this update, an exception displayed in certain cases when editing virtual machine pools has been fixed.

BZ#1393270

After updating autovacuuming options of PostgreSQL, tables should be vacuumed automatically and properly, so no table should grow without any limits.

BZ#1395462

This bug fixes several issues with insufficient synchronization when accessing MAC Pools, which may cause the Manager to allocate the same MAC address to two different virtual machines.

BZ#1406243

With this update, the issue where an incorrect calculation meant that virtual machines with an unsupported number of vCPUs attempted to start and failed has been fixed. The maximum number of allowed vCPUs per

virtual machine formula was adjusted to take into account the limitation of APIC ID. For more information see <https://software.intel.com/en-us/articles/intel-64-architecture-processor-topology-enumeration>

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#1417518

Self-hosted engine always uses an SPM ID of 1 during installation of the first self-hosted engine host, without checking database settings. This release adds options to change the database during the restore process. For disaster recovery, the `--he-remove-hosts` option has been added so that all hosts with `SPM_ID=1` are updated and a different SPM ID assigned.

For bare metal to self-hosted engine migration, a new `engine-migrate-he.py` script is provided. This script should be called before migration, and supplied with the Manager REST API login/password/endpoint and path to CA certificate. Hosts in the selected data center with `SPM_ID=1` will be put into Maintenance mode, so they can accept the new ID safely. Migration can then continue as usual, using the `--he-remove-hosts` option.

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#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#1317584

If a virtual machine was shut down and the user provided a reason, then the reason is displayed as a tooltip in the Status column on the Virtual Machine tab.

BZ#1348405

This update limits the number of snapshots per disk in the Manager due to the maximum path length in QEMU which causes snapshots operations start failing around the 98th image in the chain.

Introduced a new configuration value called `MaxImagesInChain` with a limit number of 90. The limit includes all the image chain, active volume, and the image snapshots. For example if a disk will have 89 snapshots, the next snapshot that will be created will be blocked by a CDA.

Other related operations which uses snapshots like live migrate disk or running a stateless virtual machine will also apply to the same limitation to avoid failure.

BZ#1375379

Previously, Red Hat Virtualization Manager allowed you to rename a virtual machine while it was running. The change was applied immediately in the Manager and the running virtual machine was unchanged. This would cause that you would be unable to start another virtual machine with the original name on the same host.

This update ensures that a running virtual machine name change is applied only when the virtual machine next starts.

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#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#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#1404130

Previously, when a stateless virtual machine was started and then stopped, the network interface associated with the virtual machine was automatically unplugged if there was another virtual machine with a network interface using the same MAC address. This would happen even if

the data center had a MAC pool which allowed duplicate MACs.

This update ensures that if a MAC pool allows duplicates, then a vNIC with a duplicate MAC stays plugged when a stateless virtual machine restarts.

BZ#1408790

Previously, there were frequent `KeyError` tracebacks in the `vdsd.log` after CPU hotplug. Now, these errors are no longer produced by VDSM after CPU hotplug.

BZ#1411460

Previously, the 'VM evenly distributed' policy was not properly taking the pending virtual machines (scheduled, but not yet started) into account. Each scheduling run saw the same situation and selected the same host for the virtual machine it was scheduling. Now, the policy also counts pending virtual machines, so proper balancing is applied.

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#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#1369046

Previously, when checking permissions for a CPU profile, group permissions were not considered. Users that were part of a group could not assign a CPU profile and so could not start a virtual machine. This was fixed by using `PermissionDao` and correct SQL functions when checking permissions, so group permissions are now considered.

BZ#1401077

This update ensures that if an OVF file does not contain the VM Name tag, then the virtual machine is given the OVA file name.

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#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 honoured.

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' is almost useless.

If it is necessary to change the default behaviour, 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
```

RFEs

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 fail over of a virtual machine, when the host running the machine is down, without the availability of power management for the host.

BZ#868811

This enhancement allows live migration of virtual machines which 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#1226206

Previously, when importing a virtual machine from a data storage domain, if the virtual machine had a “bad” MAC address, a MAC collision may have occurred 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#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#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#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.

Warning:

- 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#1365937

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

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#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.

CHAPTER 2. RHEA-2017:1114 RED HAT VIRTUALIZATION HOST 4.1 GA ENHANCEMENT

The bugs contained in this chapter are addressed by advisory RHEA-2017:1114. Further information about this advisory is available at <https://rhn.redhat.com/errata/RHEA-2017-1114.html>.

imgbased

BZ#1364040

Previously, the Red Hat Virtualization Host (RHVH) may have repeatedly prompted for upgrades, even when running the most recent version. This was caused by a placeholder package in the image, which was made obsolete in order to perform the upgrade. However, the package that was used to upgrade was not propagated to the RPM database on the new image.

This release now includes the missing update package in the RPM database on the new image, when upgrading the RHVH.

BZ#1432385

Previously, running `imgbase` layout took a very long time to complete on Red Hat Virtualization Hosts (RHVH) that had a very large number of LUNs or other LVM devices attached, due to the fact that it did not filter LVM commands with known tags. In this release, the LVM devices are filtered, which dramatically improves the performance for systems with a large number of LUNs.

ovirt-node-ng

BZ#1380797

Previously, a bug might have enabled services that were disabled when upgrading Red Hat Virtualization Host (RHVH). With this update, the bug was fixed. As a result, when upgrading RHVH, disabled services remain disabled.

BZ#1382833

Previously, the `crond` service was not enabled in the default configuration of Red Hat Virtualization Host (RHVH) 4.x. As a result, the log rotation functionality was not triggered, resulting in very large log files which could potentially fill the disk. In this release, the bug was fixed. As a result, the `crond` service is now enabled in the default configuration and log rotation functions as required.

BZ#1382843

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

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#1417534

Previously, `imgbased` copied `/etc` from old layers into new layers in order to keep configuration changes between upgrades. This behavior differed from that of `RPM`, in that unmodified configuration files would be preserved across `imgbased` upgrades rather than being replaced.

In this release, `imgbased` compares the sums of files to the originals that are kept per-layer in `/usr/share/factory/etc`. As a result, the unmodified configuration files are now handled as required.

BZ#1436172

Previously, `imgbased` did not add groups that were present on the new image but not the old one. As a result, `systemd-tmpfiles-clean.service` failed. In this release, `imgbased` adds groups from new layers and `systemd-tmpfiles-clean.service` is functioning properly.

BZ#1422024

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

BZ#1432359

Previously, Anaconda created a small metadata logical volume that was smaller than `.`. As a result, the logical volume could reach 100% capacity. In this release, during upgrades, if the metadata logical volume is less than 1 GB, it is resized to 1 GB.

redhat-release-rhev-hypervisor**BZ#1379763**

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

rhev-hypervisor-ng**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#1419483

Previously, Red Hat Virtualization Host (RHVH) was shipped with a for version 4.0 certificate and as a result 4.0 appeared in the Subscription page instead of 4.1. In this release, the correct certificate has been shipped, and RHVH now displays 4.1 in the Subscription page.

BZ#1422743

Previously, Red Hat Virtualization Host (RHVH) placed several requisite files (initrd, kernel) in a subdirectory under /boot. As a result, some system utilities, including kdump, were unable to locate the files in the subdirectory, regardless of the BOOT_IMAGE settings. In this release, the boot files are copied to /boot on startup and kdump is functioning properly.

BZ#1429288

Previously, VDSM was not configured after upgrading Red Hat Virtualization Host (RHVH). As a result RHVH could not run together with a Manager that was running an earlier version. In this release, by running `vdsm-tool configure --force` on boot, the VDSM is configured successfully.