



Red Hat Virtualization 4.2

Technical Notes

Technical Notes for Red Hat Virtualization 4.2 and Associated Packages

Red Hat Virtualization 4.2 Technical Notes

Technical Notes for Red Hat Virtualization 4.2 and Associated Packages

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

Legal Notice

Copyright © 2018 Red Hat, Inc.

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

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

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

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

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, 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 4.1 and release 4.2 of Red Hat Virtualization. This document is intended to supplement the information contained in the text of the relevant errata advisories available through the Content Delivery Network.

Table of Contents

PREFACE	3
CHAPTER 1. RHEA-2018:1488 RED HAT VIRTUALIZATION MANAGER (OVIRT-ENGINE) 4.2 GA	4
ovirt-engine	4
RFEs	13
vdsm	14
CHAPTER 2. RHEA-2018:1489 VDSM 4.2 GA	15
vdsm	15
CHAPTER 3. RHSA-2018:1524 RED HAT VIRTUALIZATION HOST 4.2 GA	20
imgbased	20
ovirt-node-ng	20
redhat-release-rhev-hypervisor	21
rhev-hypervisor-ng	21
vulnerability	21

PREFACE

These Technical Notes provide documentation of the changes made between release 4.1 and release 4.2 of Red Hat Virtualization. They are intended to supplement the information contained in the text of the relevant errata advisories available through the Content Delivery Network. Red Hat Virtualization errata advisories are available at <https://access.redhat.com/errata/>.

A more concise summary of the features added in Red Hat Virtualization 4.2 is available in the Red Hat Virtualization 4.2 Release Notes.

CHAPTER 1. RHEA-2018:1488 RED HAT VIRTUALIZATION MANAGER (OVIRT-ENGINE) 4.2 GA

The bugs in this chapter are addressed by advisory RHEA-2018:1488. Further information about this advisory is available at <https://access.redhat.com/errata/RHEA-2018:1488>.

ovirt-engine

BZ#1441322

Previously, when the Manager was restored from a backup, it sometimes tried to start a virtual machine that was already running, believing that the virtual machine was down. This resulted in a second instance of the virtual machine being started on a different host. In the current release, the virtual machine will not be restarted automatically after restoration if it is already running somewhere else.

BZ#659847

You can export virtual machines as named OVF/OVA files and save them to a specific path on a host or mounted NFS shared storage.

BZ#988285

The Manager supports the ability to define logical networks, ports, and subnets that are not attached to the physical interfaces of the host. This enables the virtual network interfaces to create an isolated network within Red Hat Virtualization, allowing the virtual machines to communicate among themselves. This feature is configured through the REST API and the UI.

BZ#1319758

If an OVA file is accessible to at least one host in a data center, the OVA file can now be imported into the data center as a virtual machine.

BZ#1372163

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

BZ#1376754

Previously, if a host was placed in maintenance mode and migration was cancelled while at least 3 virtual machines were attempting to migrate to it, the host ended up in an ERROR state. In the current release, the host does not move into an ERROR state in this situation.

BZ#1384321

Previously, if a virtual machine was shut down during a live merge, an illegal snapshot disk was created, the live merge failed, and the virtual machine did not start up. In the current release, the virtual machine can be recovered with a cold merge.

BZ#1408847

Red Hat Virtualization now supports TLSv1.2.

BZ#1414207

Previously, the `ForceRefreshDomainFilesListByDefault` parameter used by `virt-viewer` was set to "True" by default, sometimes causing a flood of ISO domain refresh events in the logs if many console clients were connected. In the current release, setting this parameter to "False" reduces the number of logged refresh events.

BZ#1414970

Previously, the scheduler's Host Device filter policy unit did not filter for the host whose devices were to be used, excluding inappropriate hosts. Consequently, the Manager sometimes tried to run or to migrate a virtual machine with passthrough host devices to an inappropriate host. In the current release, the Host Device filter policy unit correctly filters for the host whose devices are to be used and the scheduler prevents the virtual machine from running on or migrating to an inappropriate host.

BZ#1435485

Previously, when creating large pools of virtual machines, the Manager assigned duplicate MAC addresses, causing network problems. In the current release, the problem has been solved by removing a problematic legacy method.

BZ#1448831

Previously, when creating a new virtual machine from a template with `cloud-init` configured, the virtual machine was created but the stored root password was not copied over. In the current release, the stored root password is copied to the new virtual machine.

BZ#1449353

Previously, when importing a template into Red Hat Virtualization 4.2 from a storage domain that had previously been attached to an environment earlier than version 4.1, a maximum memory error occurred. In the current release, the maximum memory limitation has been fixed.

BZ#1458501

The REST-API now supports LLDP. See www.ovirt.org/develop/release-management/features/network/lldp/ for details.

BZ#1459908

The precision of the `rx_rate`, `tx_rate`, `rx_drop`, and `tx_drop` of virtual and host network interfaces has been increased, enabling the REST API to generate more accurate network interface statistics.

BZ#1467928

Previously, live merge sometimes failed during "Merge Status", "Destroy Image", or "Destroy Image Check" commands because of network timeout, leaving the top volume in an illegal state. In the current release, the system calls are repeated until they succeed, so that network timeout does not cause live merge to fail.

BZ#1480433

A host must be reinstalled if you turn Kdump integration on or off, or change the kernel command line parameters.

All hosts in a cluster must be reinstalled if you change the firewall type of the cluster.

An exclamation mark icon in Compute > Hosts indicates that a host needs to be reinstalled. Details appear in the host's details view, in the Events tab.

BZ#1490866

iptables has been deprecated in Red Hat Virtualization 4.2 and will be completely removed in version 4.3. Administrators must switch to firewalld, which is introduced in version 4.2. Otherwise, the Manager will review the clusters every 30 days and raise warning events in the audit log. A warning message has been added to engine-config help for all iptables-related settings.

BZ#1509588

Previously, partial import of a virtual machine was only possible using the REST API. In the current release, an "Allow Partial" check box has been added to the Import Virtual Machine(s) window in the Manager.

BZ#1519811

Previously, when upgrading from Red Hat Virtualization 4.0 to 4.1, the maximum memory size allocated for virtual machines was 1TB, which sometimes resulted in a virtual machine's actual memory size exceeding the allocated maximum memory size. In the current release, the allocated maximum memory size adjusts to the actual memory size.

BZ#1537343

Previously, if a virtual machine with a strong positive affinity to a host was down, the affinity rules enforcer tried to migrate it, because it was not running on the specified host. When migration failed, the affinity rules enforcer tried repeatedly to migrate the same virtual machine, ignoring other virtual machines that violated affinity. In the current release, the affinity rules enforcer ignores virtual machines that are down.

BZ#1540289

The default value of JBoss's `jboss.as.management.blocking.timeout` option can be changed by creating `/etc/ovirt-engine/engine.conf.d/99-jboss-blocking-timeout.conf` with `"ENGINE_JBOSS_BLOCKING_TIMEOUT=NNN"`, where "NNN" is the timeout value in seconds.

BZ#1540548

Previously, virtual machines that were paused for too long due to I/O errors were only killed when the engine tried to restart them. The current release adds a default setting, as part of high availability, that kills virtual machines that are paused too long because of I/O error, regardless of whether or when they will be resumed. This allows paused highly available virtual machines to migrate and be restarted.

BZ#1550135

Failed login attempts now appear in the audit log, with details and the user name that failed to log in.

BZ#1554180

This update prevents the user from clearing the "Allow Duplicates" check box in the Edit MAC Address Pool dialog or via a corresponding REST request if duplicate MAC addresses exist.

BZ#1565331

This release restores the ability to search Events.

BZ#1056502

This update allows you to search in Red Hat Virtualization Manager using more than one tag.

BZ#1200963

In the current release, a non-management network can be set as the default network route using the Manager, instead of selecting a custom property.

BZ#1207992

Virtual machines now stay operational when connectivity with CD-ROM images breaks. The error is reported to the guest operating system. Note that the configuration of the storage device may affect the time it takes to detect the error. During this time, the virtual machine is non-operational.

BZ#1285456

In the current release, the Manager's column widths, visibility, and order are persistent.

BZ#1324532

If a user runs 'engine-upgrade-check' without running 'engine-setup', a warning displays, informing the user that the system may not be fully up to date, even if the engine-upgrade-check says no upgrade is available.

BZ#1334982

Previously, in an emergency, users were required to shut down the hosts to preserve the data center. This caused running virtual machines to be killed by the systemd process without performing a graceful shutdown. As a result, the virtual machines' state became undefined, which led to problematic scenarios for virtual machines running databases such as Oracle and SAP.

In this release, virtual machines can be gracefully shut down by delaying the systemd process. After the virtual machines are shut down, the systemd process takes control and continues the shutdown. The VDSM is only shut down after the virtual machines have been gracefully shut down, after passing information to the Manager and waiting 5 seconds for the Manager to acknowledge the virtual machines have been shut down.

BZ#1335642

A virtual machine can now be sealed, in the Manager, in preparation for deployment as a template.

BZ#1335837

This release includes the ability to import Debian/Ubuntu virtual machines from VMware and Xen, which is available as a Technology Preview feature. From RHEL 7.4 virt-v2v can convert Debian and Ubuntu based virtual machines.

Known Issues:

1. virt-v2v cannot change the default kernel in the GRUB2 configuration and the kernel configured on the guest operating system is not changed during the conversion, even if a more optimal version of the kernel is available on the guest.

2. After converting a Debian or Ubuntu VMware guest to KVM, the name of the guest's network interface may change, and will need to be configured manually.

BZ#1364947

In the current release, the storage domain is prevented from going into maintenance if the OVF update fails. An optional check box to force maintenance mode, if desired, has been added.

BZ#1420404

The user can now decide whether a virtual machine should be warm or cold rebooted when started as "Run Once" in the Administration Portal. To facilitate this, the "Trap guest reboots" option has been renamed to "Rollback this configuration during reboots". This enables virtual machines to start on the same host when is it run as "Run Once" and then rebooted.

BZ#1428498

You can now view virtual machines that are pinned to a host, even when they are shut down.

BZ#1434306

A "Force Remove" button has been added to the Administration > Providers screen. Currently, it applies only to volume providers (Cinder). It removes the provider from the database, along with all related entities, such as storage domain, virtual machines, templates, and disks.

BZ#1439332

Previously, engine-backup always used the /tmp directory, which was problematic if the directory was full. In the current release, it is possible to change the engine-backup's temporary directory.

BZ#1439970

Previously, when creating a template from a virtual machine with the REST API, if the disk attachments were customized, the virtual machine could not be referenced by its name. In the current release, it is possible to create a template from a virtual machine referenced by its name and to customize the disk attachments in the same request.

BZ#1441059

The ovirt-engine-notifier tool exposes events that users can subscribe to by using SMTP or SNMP ovirt-engine-notifier providers.

The virtual machine events are:

- VM_PAUSED_EIO: The machine has been paused due to a storage I/O error.
- VM_PAUSED_ENOSPC: The machine has been paused due to lack of storage space.
- VM_PAUSED_EPERM: The machine has been paused due to a storage read/write permissions problem.
- VM_PAUSED_ERROR: The machine has been paused due to an unknown storage error.
- VM_RECOVERED_FROM_PAUSE_ERROR: The machine has recovered from being paused.

BZ#1443989

SPICE HTML5 support has been removed.

BZ#1445681

For REST API clients, the SSO authentication error now appears in the body of the response, for example, "<html><head><title>Error</title></head><body>access_denied: Cannot authenticate user 'admin@internal': The username or password is incorrect.</body></html>".

BZ#1492838

Previously, a bug allowed virtual machines with invalid snapshot IDs to be created in certain scenarios. The virtual machine functioned normally until the user tried to clone or commit the invalid snapshot or to run the virtual machine as 'stateless', resulting in failure or sometimes permanent disk deletion. In the current release, engine-setup checks the database for invalid image IDs. If found, engine-setup displays a warning describing the invalid image ID, the disk alias and virtual machine name assigned to the image, and advises the user to clone the virtual machine to fix the issue.

BZ#1499126

In the current release, ValidNumOfMonitors can be configured with engine-config.

BZ#1506449

In the current release, virtual machine initialization can be cleared with the REST API.

BZ#1513583

A link to the Manager's CA certificate has been added to the Welcome page.

BZ#1524824

Red Hat Virtualization now supports AMD EPYC processors for guest virtual machines.

BZ#1530675

You can now create an OVN network that is connected to a physical host network. This feature enables virtual machines on the external network to be on the same network as the virtual machines within the Data Center.

BZ#1530730

The Manager and the REST API support uploading an ISO image to a data storage domain and attaching it to a virtual machine as a CD-ROM device.

BZ#1555268

Previously, Red Hat Enterprise Linux kernels had kernel address space layout randomization enabled by default. This feature prevented troubleshooting and analysis of the guest's memory dumps. In the current feature, "vmcoreinfo" is enabled for all Linux guests. It allows a compatible kernel to export the debugging information so that the memory image can be analyzed.

BZ#1430865

Previously, if a snapshot of a disk attached to a virtual machine was deleted and the user tried to attach the storage domain containing this virtual machine before the OVF_STORE had been updated with the change, the attachment operation would fail. Because the OVF indicated the presence of a disk with a snapshot, this disk was fetched as a potential disk to register, even though it was already part of a virtual machine. In the current release, the disks are counted only once and the storage domain can be attached.

BZ#1459320

The table styling in the UI has been updated to use PatternFly to make text easier to read in selected table rows.

BZ#1484392

Previously, a template could only be imported as a clone, and a virtual machine based on such a template could only be imported with the "Collapse Snapshot" option set. In the current release, cloning and "Collapse Snapshot" are not enforced when importing a template or its virtual machine.

BZ#1486006

SUSE Linux Enterprise Server 11+ has been added to the list of guest operating systems in the Manager.

BZ#[1491771](#)

The Red Hat Virtualization Manager allows storage domain creation with a single Gluster brick.

BZ#[1506217](#)

Previously, unresponsive hosts with power management enabled had to be fenced manually. In the current release, the Manager, upon start-up, will automatically attempt to fence the hosts after a configurable period (5 minutes, by default) of inactivity has elapsed.

BZ#[1523614](#)

Previously, when a user attempted to move a disk with a snapshot that had been created before the disk was extended, the operation failed in storage domains whose data center was 4.0 or earlier. This occurred because "qemu-img convert" with compat=0.10 images interprets the space after the backing file as zeroes, sometimes causing the output disk to be larger than the logical volume created for it. In the current release, an attempt to move such a disk is blocked with an error message stating that the disk's snapshot must be deleted before moving the disk.

BZ#[1535256](#)

In the current release, MaxBlockSize is checked when extending a disk, to ensure that its size does not exceed the MaxBlockSize value.

BZ#[1539777](#)

Previously, the migration summary message showed the same value for 'total migration time' and 'actual migration time'. This value was calculated as the period of time from the start of execution of the migration command until the end of the entire migration process. In the current release, 'actual migration time' is calculated from the first migration progress event to the end of the entire migration process. If the migration command is run several times, 'actual migration time' reflects only the last run, while the 'total migration time' reflects the total time for all runs.

BZ#[1543218](#)

The descriptions of the self-hosted engine disks have been improved in the Manager to reduce the risk of their being accidentally deleted.

BZ#[1545193](#)

Previously, if a user tried to perform a live merge of a snapshot that

included unattached disks, the live merge did not finish and the snapshot remained locked. In the current release, live merge is blocked if the snapshot includes unattached disks.

BZ#[1551722](#)

Previously, the `VM32BitMaxMemorySizeInMB` and `VM64BitMaxMemorySizeInMB` parameters were defined incorrectly as both global and version-specific parameters. In the current release, they are defined per version and behave correctly.

RFEs**BZ#[723931](#)**

The Manager now displays multipath device alerts, to enable users to troubleshoot storage issues caused by faulty paths.

BZ#[856337](#)

Link Layer Discovery Protocol (LLDP) support is integrated into the Manager, so that LLDP information gathered by each host is available through the REST API and the UI. This feature improves and simplifies the detection of network configuration issues in large enterprise environments.

BZ#[872530](#)

The default zeroing method has been changed from `"dd"` to `"blkdiscard"`, which allows storage offloading, if supported by the storage array, and consumes much less network bandwidth. The zeroing method can be reverted by adding `"zero_method = dd"` to `/etc/vdsm/vdsm.conf`.

BZ#[1263602](#)

Additional mount options for the self-hosted engine storage domain, such as NFS mount, are available in the Manager.

BZ#[1066114](#)

collectd's swap plugin has been expanded to provide the following swap statistics: used, free, cached, swap_io - in and swap_io out.

BZ#[1388595](#)

The current default size of the pool of available MAC addresses, 133, is too small and has been increased to 1024.

BZ#[957788](#)

Within the VM, direct LUNs are automatically assigned specific file systems and mount points so that they can be identified in `/dev/disk/by-id` by the `'lvm-pv-uuid'` prefix.

vdsm

BZ#[1358717](#)

Previously, when a virtual machine with a thin provisioned disk was exported from an NFS data domain and imported into a block data domain, the disk's virtual and actual size appeared to be the same. Consequently, the disk's virtual size was used to allocate space for the disk. In the current release, `qemu-img` map is used to detect sparseness, allowing space allocation based on the actual size of the disk. This fix is only for NFS v4.2. Earlier NFS versions will continue to use pre-allocation, based on the virtual size of the disk.

CHAPTER 2. RHEA-2018:1489 VDSM 4.2 GA

The bugs in this chapter are addressed by advisory RHEA-2018:1489. Further information about this advisory is available at <https://access.redhat.com/errata/RHSA-2018:1489>.

vdsm

BZ#1446492

Previously, an incorrect storage domain procedure could create invalid storage domain LVM metadata. When detected, the system would fail to activate the storage domain. Now, the system logs a warning when invalid storage domain metadata is detected, without failing the activation.

BZ#1205739

A previously imported storage domain that was destroyed or detached can now be imported into an uninitialized Data Center. In the past, this operation failed because the storage domain retained its old metadata.

BZ#1432039

Previously, VDSM was refreshing active logical volumes that did not change (or never change) and do not need refresh, increasing the load on the storage server, delaying other LVM operations, and adding noise to the logs. Now, VDSM only refreshes logical volumes that have been changed, so there are no more useless refresh operations.

BZ#1449968

Currently, LVM scans and activates raw volumes during boot. Then it scans and activates guest logical volumes created inside a guest on top of the raw volumes. It also scans and activates guest logical volumes inside LUNs which are not part of a Red Hat Virtualization storage domain. As a result, there may be thousands of active logical volumes on a host, which should not be active. This leads to very slow boot time and may lead to data corruption later if a logical volume active on the host was extended on another host.

To avoid this, you can configure an LVM filter using the "vdsm-tool config-lvm-filter" command. The LVM filter prevents scanning and activation of logical volumes not required by the host, which improves boot time.

BZ#1470435

Previously, when copying disks using qcow2 compressed format, the destination disk size was not calculated correctly, because it was incorrectly assumed that the disk was not compressed. Copying an uploaded disk using qcow2 compressed format, or cloning a virtual machine using such a disk, would fail. Now, the system estimates the destination disk size based on the qcow2 actual image format, so it is possible to copy compressed disks and clone virtual machines that use them.

BZ#1474566

Previously, incorrect LVM configuration resulted in incorrect LVM output. The LVM configuration has now been fixed so that the correct LVM output is generated. The names of the generated files are as follows:

```
lvm_lvs_-v_-o_tags_--  
config_global_locking_type_0_use_lvmetad_0_devices_preferred_names_.dev.ma  
pper._ignore_suspended_devices_1_write_cache_state_0_disable_after_error_c  
ount_3_filter_a_.dev.mapper.._r
```

```
lvm_pvs_-v_-o_all_--  
config_global_locking_type_0_use_lvmetad_0_devices_preferred_names_.dev.ma  
pper._ignore_suspended_devices_1_write_cache_state_0_disable_after_error_c  
ount_3_filter_a_.dev.mapper.._r
```

```
lvm_vgs_-v_-o_tags_--  
config_global_locking_type_0_use_lvmetad_0_devices_preferred_names_.dev.ma  
pper._ignore_suspended_devices_1_write_cache_state_0_disable_after_error_c  
ount_3_filter_a_.dev.mapper.._r
```

BZ#1486543

Previously, when a VM was migrating and the source host became non-operational, the VM could end up running on two hosts simultaneously. This has now been fixed.

BZ#1494921

Previously, VDSM expected the optional object identifier (OID) field in LLDP. If this was absent, parsing of LLDP failed. Now, VDSM no longer expects the OID field in LLDP. LLDP info from the host is presented via the REST API without an OID.

BZ#1511962

Previously, TLSv12 support was backported into Red Hat Virtualization 4.1.5 (BZ#1412552) but it was turned off by default and enabling TLSv12 required manual configuration. Now, TLSv12 support is enabled by default and no manual configuration is required.

BZ#1518676

Previously, VDSM was accessing NFS storage directly when performing lease operations. A bug in Python could cause the entire VDSM process to hang if the NFS storage was not responsive. VDSM could become unkillable (D state) for many hours, until the host was rebooted. Now, VDSM uses an external process to access NFS storage, so it can continue to function correctly even if the NFS storage becomes non-responsive, and can be restarted if needed.

BZ#1543103

Previously, the `after_vm_pause` VDSM hook was not executed after I/O errors. This has now been fixed.

BZ#1130527

LVM scans and activates raw volumes during boot. Then it scans and activates guest logical volumes created inside a guest on top of the raw volumes. It also scans and activates guest logical volumes inside LUNs which are not part of a Red Hat Virtualization storage domain. As a result, it may find logical volumes with the same volume name or volume group name as groups or volumes on the host, causing errors. To avoid this, you can configure an LVM filter using the `"vdsmd-tool config-lvm-filter"` command. The LVM filter prevents scanning and activation of logical volumes not required by the host, thereby avoiding naming collisions.

BZ#1421746

In this release, a new VDSM hook that configures nested virtualization, has been introduced as a Technology Preview. Support for nested virtualization was introduced in Red Hat Enterprise Linux 7 and it enables a virtual machine to serve as a host. VDSM hooks are a means to insert code, commands, or scripts into a point in the lifecycle of a virtual machine or the VDSM daemon.

BZ#1430799

`vdsmd-tool` now provides commands for VDSM network cleanup, such as ``vdsmd-tool clear-nets`` and ``vdsmd-tool dummybr-remove``. You can remove networks configured by VDSM following the steps below. Note that the VDSM service does not need to be running:

1. To prevent loss of connectivity, it might be necessary to exclude the default route network from the cleanup. Look for a network providing the default route (`ovirtmgmt` by default):

```
# vdsmd-tool list-nets
...
ovirtmgmt (default route)
...
```

2. Remove all networks configured by VDSM except for the default network:

```
# vdsmd-tool clear-nets --exclude-net ovirtmgmt
```

3. Remove the libvirt dummy bridge `;vdsmdummy;`:

```
# vdsmd-tool dummybr-remove
```

4. Now that the host is clean, you can remove VDSM.

BZ#1447300

`Sparsify` and `sysprep` can now be run on POWER hosts.

BZ#1449869

Less additional space is now required in `/var/tmp` during VMware OVA import.

BZ#[1532133](#)

Red Hat Virtualization uses the `qemu-img` tool to copy disks during live storage migration, instead of `dd`. This tool converts unused space in the image to holes, making the destination disk sparse. Raw preallocated disks copied during live storage migration were converted to raw sparse disks. Now, you can use the `qemu-img` preallocation option when copying raw preallocated disks to file-based storage domains, so that the disks are kept preallocated after the migration.

BZ#[1202595](#)

LVM scans and activates LUNs and raw volumes, and any logical volume inside them, like those created within a guest. It then displays the unexpected guest logical volumes, or confusing errors about them. You can now use `"vdsmd-tool config-lvm-filter"` to configure an LVM filter so that LVM cannot scan or activate guest logical volumes, preventing the unexpected output.

BZ#[1335176](#)

Previously, VDSM could miss short storage outages that caused a VM to pause, so the VM paused during the short outage was not resumed. Now, Multipath queues I/O for more time and fails I/O only if there was no access to storage after the timeout, and VDSM uses a shorter timeout to detect inaccessible storage. If the outage is very short, a VM will not pause, and does not need to be resumed. If the outage is longer and a VM did pause, VDSM is more likely to detect the outage and resume the VM when storage becomes available again.

BZ#[1422670](#)

In this release, VDSM will no longer detect libvirt block jobs that have successfully completed, which will allow live merge operations to complete successfully in these cases.

BZ#[1427184](#)

Previously, the libvirt API would report live merges as complete before they were completed, resulting in errors. With this release, live merge progress is now detected using the libvirt xml, resulting in correct reporting of live merge completion status.

BZ#[1502352](#)

Previously, when a block job was manually aborted during a live merge, the merge operation would fail. In this release, `vdsmd` will now detect the failure and correct it.

BZ#1502488

Leaked clusters on an image are now correctly identified and handled, allowing cold merges to succeed when they are present.

BZ#1512547

The host monitoring task will now correctly report host statistics after the vdsd task execution queue becomes full under extreme load.

BZ#1523314

Libvirt no longer logs an incorrect error when a virtual machine is shut down correctly.

CHAPTER 3. RHSA-2018:1524 RED HAT VIRTUALIZATION HOST 4.2 GA

The bugs in this chapter are addressed by advisory RHSA-2018:1524. Further information about this advisory is available at <https://access.redhat.com/errata/RHSA-2018:1524>.

imgbased

BZ#1454536

Currently, the Red Hat Virtualization Host generates VDSM certificates at the time of the first boot. This means that if the system clock was not set correctly at install time then chrony or ntpd may resynchronize the clock after the VDSM certificate was generated, leading to a certificate that is not valid yet if the appropriate timezone is behind UTC. A workaround for this is to set the system clock appropriately at install time. Now, imgbased-configure-vdsm starts after chronyd or ntpd and waits two seconds for the clock to synchronize, but this is not a guarantee.

BZ#1535791

Red Hat Virtualization Host now ensures that deleted configuration files are no longer restored from previous layers when upgrading.

ovirt-node-ng

BZ#1502920

This update ensures that Red Hat Virtualization Host (RHVH) synchronizes system-specific data in /usr. This ensures that systems registered to Satellite are still able to receive updates after upgrading.

BZ#1515698

In order to allow for faster remediation of kernel CVEs and for testing of fixes from newer kernels, RHVH now supports installation of new kernels without a full image update. New kernel installations properly update the bootloader configuration.

BZ#1549630

Previously, Red Hat Virtualization Host did not warn users if a local storage domain was placed on the same file system as / (root). As a result, local storage domains on the same file system as / were not migrated when the host was updated, leading to a potential loss of local virtual machines.

In this release, the host now fails to upgrade if storage domains are located on the same file system as /, and a message appears instructing the user how to resolved the issue and local storage domains on / are no longer at risk.

BZ#1489567

Red Hat Virtualization Manager now displays the Red Hat Virtualization Host version installed.

BZ#1561258

To properly set boot flags if grub2-mkconfig is run, or a new kernel is installed, Red Hat Virtualization Host (RHVH) ships with a custom grub generator. This generator uses a list of all local LVM volume groups to create boot arguments, but systems with local storage domains on separate volume groups are currently present in the list, which could result in incorrect flags being in the grub.conf file. This update enables RHVH grub generators to explicitly look for the RHVH volume group, and ignore others.

redhat-release-rhev-hypervisor**BZ#1516123**

This update ensures that tuned.service is enabled by default to enable tuned-adm to set the active profile.

rhev-hypervisor-ng**BZ#1374007**

In this release, a new version of Anaconda now includes storage constraint checks and default settings for Red Hat Virtualization Hosts (RHVH), which require a special partitioning layout. When custom partitioning is selected, LVM-thin is the default for RHVH.

BZ#1420068

In this release, Red Hat Virtualization Host supports NIST SP 800-53 partitioning requirements to improve the security. Environments upgrading to Red Hat Virtualization 4.2 will also be configured to match NIST SP 800-53 partitioning requirements.

BZ#1503148

NTP is deprecated in favor of chrony in RHV 4.2. The updated default configuration allows users upgrading to RHV 4.2 from RHV-H to seamlessly transition from NTP to chrony without intervention.

vulnerability**BZ#1567974**

A command injection vulnerability was found in the 11-dhclient script provided by dhcp-client located in /etc/NetworkManager/dispatcher.d/11-

`dhclient`. A malicious DHCP server, or an attacker on the local network able to spoof DHCP responses, could use this flaw to execute arbitrary commands with root privileges on systems using NetworkManager and configured to obtain network configuration using the DHCP protocol.

BZ#[1558721](#)

A privilege escalation flaw was found in gluster snapshot scheduler. Any gluster client allowed to mount gluster volumes could also mount shared gluster storage volume and escalate privileges by scheduling malicious cronjob via symlink.

BZ#[1566837](#)

A flaw was found in the way the Linux kernel's KVM hypervisor handled exceptions delivered after a stack switch operation via `Mov SS` or `Pop SS` instructions. During the stack switch operation, the processor did not deliver interrupts and exceptions, rather they are delivered once the first instruction after the stack switch is executed. An unprivileged KVM guest user could use this flaw to crash the guest or, potentially, escalate their privileges in the guest.

BZ#[1567074](#)

A flaw was found in the way the Linux kernel handled exceptions delivered after a stack switch operation via `Mov SS` or `Pop SS` instructions. During the stack switch operation, the processor did not deliver interrupts and exceptions, rather they are delivered once the first instruction after the stack switch is executed. An unprivileged system user could use this flaw to crash the system kernel resulting in the denial of service.