Red Hat Enterprise Virtualization 3.6

Technical Notes

Technical Notes for Red Hat Enterprise Virtualization 3.6 and Associated Packages

Red Hat Enterprise Virtualization 3.6 Technical Notes

Technical Notes for Red Hat Enterprise Virtualization 3.6 and Associated Packages

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

Legal Notice

Abstract

These Technical Notes provide documentation of the changes made between release 3.5 and release 3.6 of Red Hat Enterprise Virtualization. Subsequent advisories that provide enhancements, provide bug fixes, or address security flaws are also listed. They are intended to supplement the information contained in the text of the relevant errata advisories available via Red Hat Network.

Table of Contents

INTRODUCTION

These Technical Notes provide documentation of the changes made between release 3.5 and release 3.6 of Red Hat Enterprise Virtualization. They are intended to supplement the information contained in the text of the relevant errata advisories available via Red Hat Network. Red Hat Enterprise Virtualization 3.x errata advisories are available at https://rhn.redhat.com/errata/rhel6-rhev-errata.html.

A more concise summary of the features added in Red Hat Enterprise Virtualization 3.6 is available in the Red Hat Enterprise Virtualization 3.6 Release Notes.

CHAPTER 1. RHEA-2016:0376 RED HAT ENTERPRISE VIRTUALIZATION MANAGER 3.6.0

The bugs contained in this chapter are addressed by advisory RHEA-2016:0376. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2016-0376.html.

ovirt-engine

BZ#1235179

> Previously, if an error occurred during copy of the virtual machine images during the Clone VM operation, the cleanup incorrectly deleted the source virtual machine instead of the destination one. Now, after a failed attempt to clone a virtual machine, the source virtual machine with its disks stays in the system and the incomplete destination virtual machine is deleted.

BZ#1176820

> Prior to Red Hat Enterprise Virtualization 3.5, the <policy>name</policy> tag was used to define the name of the scheduling policy to select. In 3.5, the <name>name</name> tag was introduced. By accident the policy tag was removed, which should be there for backward compatibility. As a result, user REST scripts for setting scheduling policies sometimes did not work. With this release, support for the <policy/> tag has been re-added. The <policy/> tag can now be used again, but <name/> is preferred if it is also present.

BZ#1181672

> Previously, mishandling of an exception to set a host for high availability prevented removing that host from a running setup. With this release, the Manager now logs the error without re-throwing it, allowing it to move the host into maintenance and remove it.

BZ#1149575

> Previously, when running engine-setup with an already generated answer file that included invalid values, the setup process did not return errors but some options were missing from the GUI after the setup. This bug fix adds in error validation so appropriate options are available in the GUI after setup.

BZ#1108153

> A new user role 'VmImporterExporter' is now available. The role allows users to export and import virtual machines in the Administration Portal. The 'DataCenterAdmin' permission is no longer required for performing such actions.

BZ#1278271

Previously, a user couldn't create a virtual machine in the User Portal if they only had the InstanceCreator role permissions. It showed "Not available when no Data Center is up." in the New VM dialog even if there was a data center up. Now, the error message has been changed to a more meaningful "There is no Cluster on which you can create a VM. Please check Data Center status and Cluster permissions." Permission checking has been fixed so that InstanceCreator role permissions are sufficient to create a new virtual machine.

BZ#1275539

Video memory sizes in previous Red Hat Enterprise Virtualization versions were not up to date with the common contemporary screen setups and current state of graphics drivers. As a result, the allocated video memory could be insufficient for high resolutions and/or multihead setups. Now, video memory sizes have been updated to handle typical screen arrangements for all supported screen resolutions in 3.6 clusters. If those settings still do not work well for special requirements, they can be adjusted by advanced users using VDSM hooks, newly including the 'vgamem' attribute of the <video> domain XML element.

BZ#1107512

With this release, virtual machines can be pinned to more than one host. A virtual machine that is pinned to multiple hosts cannot be live migrated, but can be made highly available between the specified hosts. The virtual machine cannot run on any other hosts in the cluster, even if all of the specified hosts are unavailable.

BZ#1185411

Previously, the Manager's JVM heap size was configured to be 1GB by default. As a result, larger setups, such as those with hundred of virtual machines or more, made the Manager run out of heap memory, and required manual configuration to increase the heap size. Now, engine-setup automatically configures the heap size to be the maximum of 1GB and 1/4 of available memory. Larger setups only need to use a machine with enough memory, as is already recommended, and heap size configuration will be done automatically, thus preventing such failures and not requiring manual configuration.

BZ#1133642

Previously, if you created a virtual machine and selected 'Other OS' for the Operating System field, the operation system type defaulted to 32 bit. As a result, when installing a 64 bit operating system not present in the list of operation systems, the virtual machine can use only up to 20 GB of RAM, which is the limit for 32 bit operation systems. With this update, the "Other OS" option now defaults to 64 bit. Any operating system installed in the virtual machine is allowed to use up to 4TB of memory if it is capable for it. 32 bit operating systems are still limited to 20 GB of RAM.

BZ#1094842

> With this release, host network bonds with bond mode 0, 5, or 6 will not support adding a VM network to them. Similarly, a bond with a VM network attached will not be allowed to move into mode 0, 5, or 6.

BZ#1182007

> Previously, the internal monitoring for virtual machines and hosts didn't take into account the latest running virtual machines when calculating the free memory available for running additional virtual machines. As a result, a host would not be able to run virtual machines, or would run virtual machines without having enough memory for them. With this release, the calculation of the host's maximum free memory includes the running virtual machines.

BZ#1208122

> Previously, the --scope option could only be specified once per engine-backup command. If the user wanted to back up both 'files' and 'db', for example, they would have to run the command twice, once for each scope. Now, the engine-backup command accepts more than one --scope option, so a user could back up both 'files' and 'db' in a single command.

BZ#1188147

> This release adds the ability to track backups and alert the administrator if a pre-configurable time has passed since the last successful backup. There is now a clear indication when a new backup is recommended, and minimal Manager downtime in the case of corruption.

BZ#1175313

> Previously, the REST API interface was not aware of a forced logout after 10 hours of being logged in, and would try to access the REST API after it. This caused an authentication popup to appear in the Administration Portal. Now, the REST API interface is aware of the forced logout and stops accessing the REST API after the logout happens, so the popup no longer appears.

BZ#1251893

> Previously, users without VM -> Provisioning Operations -> Create permission were unable to live migrate virtual machines. This has now been fixed and such permission is no longer needed.

BZ#1250556

> Previously, changing the host name of the Manager machine when the machine also provided a local ISO storage domain was prevented by the

rename tool, to prevent the ISO storage domain losing connectivity during the renaming process. Also, changing the host name of the Manager machine when the machine also provided a local data domain could cause running virtual machines to become stuck due to the hosts losing connectivity with their virtual disks. Now, the rename tool checks both of these, and guides the user to eject/shut down/put to maintenance any required virtual machines or storage domains, and only then allows the rename process to continue. To change the host name of the Manager machine, see the product documentation: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Virtualization/3.6/html-single/Administration_Guide/index.html#sect-The_oVirt_Engine_Rename_Tool

BZ#1245801

Previously, missing handling for the Time Zone conversion of (GMT-06:00) Central Standard Time (Mexico) caused the offset used to be 0, putting the time in the guest operating system off by 6 hours. Now that the handling has been added, all selectable timezones are being handled, restoring the correct time to the guest operating system.

BZ#1245230

Previously, incorrect X.509 Subject encoding caused incorrect Unicode encoding. Unicode strings within the subject name were rejected by some applications. Now, X.509 Subject is encoded correctly so that applications accept certificates.

BZ#1176457

Previously, lc_messages was sometimes set to a non-English locale in postgresql.conf. As a result, engine-backup --mode=restore did not filter expected errors, which are in English, and failed. To fix this, engine-backup --mode=restore was changed to require lc_messages to be en_US.UTF-8, and engine-setup was changed to set lc_messages to en_US.UTF-8 on clean installs, and require that on upgrades (and abort otherwise). Now, if lc_messages is not en_US.UTF-8, engine-backup --mode=restore exits with a more helpful error message. engine-setup sets it correctly on clean installs even if the system locale is different, and aborts on upgrades if it is not set correctly.

BZ#1145002

Previously, when a user deleted the default data center, it was not possible to remove the default cluster or associated template afterwards. This has now been fixed.

BZ#1232586

Windows 10 is now supported as a guest operating system (32 and 64 bit).

BZ#1298094

Previously, when a non-administrator created a virtual machine in a virtual machine pool ownership was automatically granted. This caused the virtual machine to be designated as occupied and prevented other users from taking it. Now, the ownership of virtual machines in a virtual machine pool is not granted automatically.

BZ#1239153

Previously, when a virtual machine from a Manual Pool was shut down, its stateless snapshot was removed automatically, while the virtual machine was still assigned to the user. This issue was resolved and now for virtual machines from a Manual Pool the stateless snapshot is removed automatically only when the virtual machine is manually returned to the pool by an administrator. For virtual machines from an Automatic Pool the stateless snapshot is removed automatically on virtual machine shutdown, as usual.

BZ#1253043

Previously, incorrect handling of the virtual machine recovery file caused virtual machine shut down inside the guest. This caused the virtual machine recovery file to be leaked. This may cause incorrect detection of virtual machines, especially in the case of network failures, and cause the Red Hat Enterprise Virtualization Manager to find unknown external virtual machines. Now, VDSM has been fixed to no longer leak recovery files. Recovery files are now correctly handled and this prevents the Red Hat Enterprise Virtualization Manager from incorrectly detecting virtual machines.

BZ#1236267

In the Edit Virtual Machine dialog, if an invalid name was entered, the dialog closed and a server-side error was reported. With this update, an invalid virtual machine name is handled in the standard way. The field is highlighted with a red border.

BZ#998607

With this update, it is now possible to change the allocation policy of disks imported from an export domain as part of importing a virtual machine. This allows users to specify the format and allocation policy of each disk. Note that `collapse_snapshots` needs to be set to true and the disk IDs should be the IDs of the existing disks that the user is importing.

For example, to force one of the disks to be preallocated:

```
POST /storagedomains/{storagedomain:id}/vms/{vm:id}/import
<action>
  <storage_domain>
    <name>mydata</name>
  </storage_domain>
  <cluster>
```

```
          <name>mycluster</name>
      </cluster>
      <vm>
        <disks>
          <disk id="792f8ca8-3932-429b-9660-006ba144bc72">
            <format>raw</format>
            <sparse>false</sparse>
          </disk>
        </disks>
        <snapshots>
          <collapse_snapshots>true</collapse_snapshots>
        </snapshots>
      </vm>
    </action>
```

BZ#1024010

A virtual machine's buffered or cached memory is now taken into account
in the Extract, Transform and Load (ETL) process. Some applications,
like the Oracle Directory Server, uses cached memory as permanent memory
and never releases it. As a result, users need to take this memory into
account when planning for memory requirements. The data is collected
each minute, hour and day as a default and users can create ad hoc
reports based on the collected data.

BZ#1178480

Previously, less than or equal to (<=) was used in monitoring storage
free space. Alerts for low disk space were triggered when they shouldn't
have. With this update, when monitoring the storage free space, less
than (<) is now used. Alerts for low disk space are now generated
appropriately.

BZ#1252055

Previously, NUMA distances were not always reported properly and this
cause the engine to crash with NPE during host installation or
activation. Now, the engine invents virtual NUMA distances when none are
provided by the hardware so that the engine does not crash and the host
can be activated.

BZ#1048976

With this update, users can now edit the name of virtual disks attached
to a virtual machine in the New VM window when creating a virtual
machine based on a template. This allows users to specify names that are
different to those specified for the disks attached to the template.

BZ#912263

Virtual machine and template names only need to be unique within a data center, the names can be re-used in different data centers in the environment.

BZ#912260

This feature changes the way allocated MAC addresses are tracked. Previously, there was one global MAC address pool, from which all requests for MAC addresses were served. Now you can create several MAC address pools with different settings, and assign one of such pools to a certain data center. This allows each data center to have a different 'mac-space'. Pools can also be shared among multiple data centers.

BZ#1170263

Previously, if a storage domain has less than 1 GB disk space, the available free space was reported as [N/A] in the Administration Portal. This is now altered to display <1.

BZ#877224

With this release, engine-setup notifies the user of the availability of the Reports and Data Warehouse components. These components are not required by all users, so are not installed by default; instead, the installer advertises their availability so that the user can choose whether or not to install them.

BZ#1192465

Previously, distinction was between the Red Hat Gluster Storage versions and the glusterfs versions. This cause versions to be displayed incorrectly. Now, the glusterfs version is always displayed so that the version displays correctly.

BZ#1220690

Previously, the resource allocation tab was not correctly handled for templates. This meant that when creating a virtual machine using a template in the GUI, the 'New Virtual Machine' window would display 'Memory Balloon Device Enabled' regardless of whether or not this was disabled in the template. Now, these template properties are handled properly and creating a virtual machine from a template in the GUI accurately reflects the configuration of the template properties.

BZ#1084117

With this release, virtual machine permissions have been adjusted to be more specific. Previously, VM_BASIC_OPERATIONS allowed a user to run, stop, shut down, pause, and hibernate the virtual machine. Now, REBOOT_VM, STOP_VM, SHUT_DOWN_VM, PAUSE_VM, HIBERNATE_VM, and RUN_VM can each be separately assigned.

BZ#1184164

> With this update websocket proxy configuration values are now validated
> to prevent misconfigureation.

BZ#1264456

> The python-websockify 0.6 dependency package contained changes that
> resulted in novnc and spice-html5 consoles no longer being able to open.
> Fixed ovirt-websocket-proxy.py to work with python-websockify 0.6 so
> that the consoles now open as intended.

BZ#1100201

> When updating the Manager, any package that cannot be rolled back to a
> previous version is now listed in detail so the user can better evaluate
> whether to proceed.

BZ#1184445

> Previously, users could remap the Ctrl+Alt+Del shortcut to Ctrl+Alt+End
> using the Console Options dialog in the Administration or the User
> Portal. With this enhancement, users can now use the engine-config
> utility with the 'RemapCtrlAltDelDefault' option to remap the shortcut
> keys. This global value can still be optionally overridden for each
> virtual machine using the Console Options dialog.

BZ#1267165

> With this release, you can set the threads per core for a virtual
> machine from the Administration Portal. The New Virtual Machine window,
> the Edit Virtual Machine window, and the REST API are updated to include
> the threads per core option along with the number of virtual sockets and
> the cores per virtual sockets options. It is recommended to keep the
> number of threads per core set to 1 for x86 (Intel and AMD) host
> machines, and 1-8 for IBM POWER8 machines. The default value is 1 and
> can be increased to enable simultaneous multi-threading (SMT). Check the
> host architecture before increasing the value otherwise it can lead to
> overutilization of the host.

BZ#1214408

> Previously, when importing an existing, clean storage domain that
> contains OVF_STORE disks from an old setup to an uninitialized data
> center, the OVF_STORE disks did not get registered after the data center
> was initialized and all virtual machine information was lost. With this
> update, when importing clean storage domains to an uninitialized data
> center, the OVF_STORE disks are registered correctly, and new
> unregistered entities are available in the Administration Portal under
> the Storage tab. In addition, storage domains with dirty metadata
> cannot be imported to uninitialized data centers.

BZ#1193643

Previously, if a quota mode was enabled on a data center, editing an existing disk of a running virtual machine always failed. With this update, the default quota value for an existing disk is no longer set to null but set to the current value. A wrong quota change is no longer detected for an existing disk, and editing disks on a running virtual machine works as expected.

BZ#1111253

The global configuration operation has been added to the engine-config tool with the name VmGracefulShutdownTimeout. This configures the length of time, in seconds, the virtual machine will wait before it shuts down to allow for users to close running applications.

BZ#1112332

The Affinity Rules Enforcement Manager has been added to check for broken affinity and anti-affinity policies and migrate virtual machines if necessary.

BZ#1113664

When performing migrations, numbered left and right arrows now indicate incoming and outgoing migrations in the 'Hosts' tab.

BZ#1210633

Previously, the timezone Africa/Algiers was not recognized by Red Hat Enterprise Virtualization Manager 3.5. As a result, when importing virtual machines from previous versions that used the timezone failed. This update adds Africa/Algiers to the list of supported timezones. Virtual machines that used the timezone now imports correctly.

BZ#1152989

Previously, when adding a permission via the REST API the user had to specify the role identifier and not the role name. This meant that the user needed to look up the identifier, issuing an additional call to the top level 'roles' collection. Roles have human readable names, such as UserVmManager, but the REST API did not support specifying these names. Now, the REST API supports both ways to specify a role. This is also supported by the Python SDK, the Java SDK, and the RHEVM Shell.

BZ#1209281

Previously, the pki-pkcs12-extract.sh script relied on the existence of the /dev/fd directory. In Linux, this is normally symbolically linked to the /proc/self/fd directory, allowing processes to access its STDIN, STDOUT, etc as named files. If the /dev/fd directory did not exist, the

script failed. This includes, for example, trying to run engine-setup during installation from a kickstart file. With this update, the script was updated to use the /proc/self/fd directly. Now the script only requires that the /proc direcory is mounted, and does not fail if the /dev/fd directory does not exist.

BZ#1282853

Previously, the custom properties of a virtual machine were not saved for virtual machine pools. Now, custom properties can be set in the Edit Pool dialog, in the same way as in Edit Virtual Machine dialog, and they are correctly applied for all virtual machines in the pool

BZ#1283964

This change allows the user to create NFS v4.1 storage domains (or edit existing domains to use NFS v4.1).

BZ#1206264

This update adjusts the behavior of the Clear All action in the alerts footer tab in the Administration Portal. Previously, Clear All cleared all actions run by the user to hide alerts, and displayed all cleared alerts in the alerts panel. The Clear All action now dismisses all alerts and events in that list. This fix also adds the Display All action to show all alerts and events.

Additionally, an issue has been fixed where actions on the events footer tab affected the alerts tab, and as a result, clearing or displaying all events was simultaneously also applied to the list of alerts. This has been fixed and now operations done on the events tab affects only events.

BZ#1121089

Previously, in rare cases a virtual machine disappeared during migration in the UI and reappeared as an external virtual machine that could no longer be managed. This error has now been fixed.

BZ#1284626

With this release, users can click and drag to expand the width of columns in the Configure dialog box.

BZ#1140569

It is now possible to update the template for an existing virtual machine pool. This enables updated templates to be deployed to all virtual machines in the pool, and the next restarted virtual machine will be based on the new template version.

BZ#1204793

> With this update, the default name for the Red Hat Enterprise
> Virtualization management network has changed from 'rhevm' to
> 'ovirtmgmt'.

BZ#1128426

> With this release, the German translation has been reworked and
> improved.

BZ#1203735

> Previously, engine processes that were not invoked directly by a logged-
> in user but invoked automatically either by an event or by the time
> scheduler had the "null@N/A" user name assigned to them in logs. This
> was confusing for administrators. With this update, the user name
> "SYSTEM" is assigned to all processes which are not invoked directly by
> a logged-in user so the log messages are clearer.

BZ#1132066

> Previously, it was only possible to copy disks attached to a template to
> a different storage domain. With this release, you can now also copy
> virtual machine disks and floating disks to either the same or a
> different storage domain.

BZ#1128763

> The REST API can be used now to retrieve the complete .vv file with
> console connection information, allowing easier scripting of connecting
> to virtual machine graphical consoles.

BZ#1128691

> With this release, the Administration Portal's search function can
> filter disks by the "Wipe After Delete" property.

BZ#1138126

> Previously, when trying to attach an imported storage domain that
> carried existing metadata to an uninitialized data center, an exception
> from VDSM was returned with no proper warning message. With this update,
> the storage domain is checked for existing metadata and an error message
> is provided to advise users to attach a clean storage domain first.

BZ#1196757

> With this release, you can edit the alias of a disk image in the Import
> Image view when importing an image from Glance. Without the ability to
> change the alias, each import of the same image creates a disk with the

same name (which is the default name for that disk). Now, an editable column named "Disk Alias" was added to the table of images which were chosen to be imported. Each image has its own default name, which has not changed and can be left as is. However, one can now edit each image alias before the import.

BZ#1140158

Fixed or removed tooltips for clickable buttons that provided the exact same text as the button text.

BZ#1123391

Red Hat Enterprise Virtualization supports virtual machines with up to 240 vCPUs. The previous version supported a maximum of 160 vCPUs.

BZ#1209185

Previously, the names of virtual machines based on a given template would not be shown in the Virtual Machines subtab of the Templates tab under certain conditions. Now, the names of virtual machines based on the selected template are displayed correctly.

BZ#1209957

Previously, fence agents accepted boolean agent flags as secure or lanplus (ex: lanplus, lanplus=1, or lanplus=true). The options became unsupported, and caused hosts with those fence agents types to stop functioning. With this update, the fencing options for any existing hosts will have the new format. Any implicitly added options will also have the new format. For existing hosts, the fence agents work as normal after the package is updated. If the options field was defined as an encrypted field, redefine the options field manually to align to the new format. For new hosts, the new format (i.e <key>=<value>) should be defined in the options field.

BZ#1154666

Previously, certain form filler single sign on (SSO) solutions (such as Indeed ID) that do not execute 'change' events would not work with the Red Hat Enterprise Virtualization login pages. Form filler SSO solutions that executed 'change' events worked correctly. The login forms have been enhanced to work without 'change' events. All form filler SSOs should now work.

BZ#1157362

This release introduces the Host Upgrade Manager, which notifies the user of available updates to hosts and automates the process of moving the host to maintenance mode, applying updates, and reactivating the host.

BZ#1211231

> With this enhancement, users can connect to Windows 8 and Windows 2012 virtual machines using the SPICE protocol without QXL drivers. Limitations include: no multiple monitors, graphics are not accelerated, etc.

BZ#1157377

> This update adds support for integration with Red Hat Satellite 6.1.x, providing the ability to receive information on errata for hosts, virtual machines, and the Red Hat Enterprise Virtualization Manager in a Red Hat Enterprise Virtualization environment.

BZ#1157999

> Previously, all OpenStack external providers like Neutron and Glance used the same Keystone authentication URL stored in the KeystoneAuthUrl configuration value. Now, a URL field has been added to provider configuration so that different OpenStack providers can use different Keystone authentication URLs.

BZ#1099412

> This release adds the ability to edit the connection details of ISO and Export storage domains.

BZ#1159181

> Fixed a typographical error in an audit log message when adding a cluster policy.

BZ#1221312

> The fence-agents package on Red Hat Enterprise Linux dropped support of boolean flags being supplied to various scripts from the command line without a value. As a result, all existing flags stored in the database have been changed to explicitly set the boolean flag value. This ensures old settings of PM definitions that have the old format now work after upgrade.

BZ#1163105

> Search query filtering has been added to the REST API to match functionality of the Web Admin Portal.

BZ#1223194

> The Administration Portal and the User Portal now support Mozilla Firefox 38 Extended Support Release (ESR) as the client browser.

BZ#1225274

> The "DiskProfileUser" role, which was an administrator role, was assigned to the "Everyone" group by default. As a result, when users logged into the User Portal, they saw the Extended tab by default and were exposed to options that they did not have permissions to operate. With this update, the "DiskProfileUser" role is changed to an end-user type role. Users with end-user type roles now see the Basic tab by default.

BZ#1228257

> It is now possible to receive notifications when a virtual machine enters a paused state. The corresponding entry was added under VM Events in the Users -> Event Notifier -> Manage Events window.

BZ#1169224

> It is now possible to specify a custom name for the disk and/or template when importing a Glance disk through the REST API.

BZ#1169228

> When importing a disk from a Glance storage domain via the REST API, the response will now contain the UUID of the disk created on the target storage domain.

BZ#1172207

> Previously, changing the cluster of a virtual machine failed as the CPU profile from the old cluster did not match the CPU profile of the new cluster. With this update, the CPU profile gets updated automatically to match the newly selected cluster.

BZ#1172629

> Hosts now support more than two power management cards through the use of sequential and concurrent card configurations.

BZ#1231121

> UsbDk has been added to Red Hat Enterprise Virtualization as a preferred alternative to usbclerk. UsbDk is a filter driver that allows Windows guest virtual machines access to the client USB devices through SPICE client (remote-viewer).

BZ#1176555

> This release adds support for fencing policy per cluster to the REST API. Clusters retrieved from REST will now have a <fencing_policy> section.

BZ#1231274

> The status of a storage domain can now be reported by an external system. The status is set using the REST API, and the internal status (as reported by the Manager) and the external status (as reported by a third party) for each storage domain are displayed when a domain is selected in the Administration Portal user interface.

BZ#893927

> With this update, you can configure quotas using the REST API. New entities such as QuotaStorageLimit and QuotaClusterLimit were added to the quotas sub-collection. QuotaStorageLimit allows you to specify a quota policy for individual or all storage domains. QuotaClusterLimit allows you to specify the Memory and CPU limit for individual or all clusters in a data center.

BZ#1259601

> Running engine-setup now checks the expiration date of relevant certificates (including the internal CA certificate and those signed by it) and asks the user whether to renew any certificate that has expired, or is close to expiring. If the user selects 'Yes', those certificates are renewed; 'No' will not alter anything.

BZ#803251

> With this release, the Virtual Machines main tab has been improved to show more detailed data about CPU, memory, and network usage. Instead of a simple bar it now shows historical data and trends.

BZ#1048573

> Enhanced audit log messages with additional information to enable users to easily resolve migration failure issues.

BZ#1200952

> With this release, the Linux Boot Options section of the Run Once window allows users to select kernel and initrd files for Linux virtual machines. Available files are files from the ISO domain of type Unknown (i.e. not ISO, floppy, or disk).

RFEs

BZ#1186812

It is now possible to use the REST API to define specific credentials to each iSCSI target per host by sending a POST request to '/api/hosts/{%host_id%}/storageconnectionextensions'.

BZ#515840

It is now possible to add memory to a virtual machine without having to shut it down first.

BZ#609689

It is now possible to refresh LUN sizes in existing storage domains so that resized LUNs are properly represented in the Manager.

BZ#798601

In the User Portal, the following visual indicators have been added to indicate that the virtual machine console is in use:

1. The virtual machine background of an occupied console is grayed-out when not selected.

2. The text "console in use" is displayed in the virtual machine status description of an occupied virtual machine.

BZ#1136329

This feature introduces new REST API resources for a better way to manage networks. Network related tasks like removing unmanaged networks, attaching new networks, and creating bonds are now easier using the new 'HostSetupNetworks' and 'NetworkAttachment' entities.

BZ#866124

A health status field has been added to host and storage domain entities, which can be set and retrieved in the UI and the REST API.

BZ#867453

QEMU capabilities for auto-convergence and/or Xor Binary Zero Run-Length-Encoding (XBZRLE) can be used to reduce virtual machine downtime and improve convergence during migration. This is supported by hierarchical configuration in 3 levels: global (engine-config), cluster, and virtual machine.

BZ#869804

This release adds SR-IOV enablement to the user interface. In order to connect a vNIC directly to a Virtual Function of an SR-IOV-enabled NIC, the vNIC's profile must be marked as a "passthrough" one. The properties

that must be configured on the Virtual Function are taken from the vNIC's profile/network (vlan tag). Each SR-IOV-enabled host-NIC must have a defined set of networks that it is allowed to service. When starting the virtual machine, its vNIC will be directly connected to one of the free Virtual Functions on the host. Note that not all Physical Functions are equivalent: the vNIC must be connected via a host-NIC that has the vNIC's network as one of its allowed networks.

BZ#878129

Previously, after a query was entered into the search query box the query would disappear when a user switched tabs then returned to the original tab, as there was no infrastructure to store the query. Now, the query in the search query box will be retained after switching tabs.

BZ#902185

This release adds support for Korean to the Administration Portal and User Portal.

BZ#952703

Virtual machines can now be imported from VMware using the Administration Portal graphical user interface. VMware providers can be added in the 'External Providers' dialog, and virtual machines imported using the 'Import' dialog.

BZ#957593

In multiple-monitor setups, it is now possible to specify which monitors are to be used by a virtual guest, and which are to be reserved for the local machine. This is done by editing the ~/.config/virt-viewer/settings file.

BZ#1122990

Previously, when a comment was added to an object, the Administration Portal displayed a yellow paper sheet icon in the "Comment" column. This has now been changed to render the actual comment.

BZ#965289

This feature adds the ability to define a management network at the cluster level. This allows you to create a virtual Data Center with multiple clusters that have different management networks, which in turn allows you to use the same physical network with different VLANs defined for each management network.

BZ#1117489

Previously, creating a new virtual machine would open the 'Guide Me'

window, which offered a link to 'Configure Virtual Disks'. Now, this functionality is integrated in the 'New Virtual Machine' and 'Edit Virtual Machine' windows, and the 'Guide Me' window no longer opens in a virtual machine context.

BZ#1232559

With this update, Microsoft Windows 10 is now supported as a guest operating system in Red Hat Enterprise Virtualization. The required VirtIO-Win drivers for this operating system are also provided.

BZ#1063343

This feature reports RX/TX byte statistics for host NICs and virtual machine NICS. In the Administration Portal, the Network Interfaces tab of hosts and virtual machines shows statistics for Rx Rate, Tx Rate, Total Rx, Total Tx, and Drops. The REST API (api/hosts/<host id>/nics/<nic id>/statistics) returns statistics for data.current.rx, data.current.tx, data.total.rx, data.total.tx, errors.total.rx, and errors.total.tx. The drops/errors statistics allow you to check the NIC's health and avoid errors associated with NICs experiencing network issues.

BZ#1160731

Previously, it was only possible to import NFS export domains. Now, it is possible to import Gluster and PosixFS export domains, as well as, NFS export domains.

BZ#601863

This feature implements a command line access to a virtual machine's serial console using the SSH client. For more information, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Virtualization/3.6/html/Virtual_Machine_Management _Guide/sect-Starting_the_Virtual_Machine.html#Opening_a_Serial_Console_to_a_Virtual_ Machine.

BZ#1166755

With this update, default options have now been added for the Drac7 fencing agent, allowing it to function without the need for prior manual configuration.

BZ#679070

Two storage-space thresholds have been added for storage domains:
* WarningLowSpaceIndicator - configured as a percentage (0-100) of disk space. A storage domain with disk space below this threshold will display warning to the user and in the log, and
* CriticalSpaceActionBlocker - configured as an integer (in GB) between

> 0 and the MaxInteger, and can exceed the total space of the storage domain. A storage domain with disk space below this threshold will block all operations that consume storage space, even if the operation is temporary, and will display relevant warning to the user and in the log.
>
> These values can be configured when creating, editing, or importing a storage domain.

BZ#758516

> With this release, the request for 'LockScreen' can be configured to be ignored by changing the 'ignore_lock_screen_request value' to True in the guest agent configuration.

BZ#1168576

> When creating a template, users can now set the format of the template's disk volumes on creation. This enables users to define disks in the thin provisioning QCOW2 format.

BZ#840114

> Kernel SamePage Merging (KSM) can now be configured to obey or ignore Non-Uniform Memory Access (NUMA) topology constraints in the Red Hat Enterprise Virtualization Manager. KSM can now either merge pages across the whole machine, ignoring NUMA topology, or merge pages only inside NUMA nodes while obeying the NUMA boundaries.

BZ#1080029

> It is now possible to track the number of days that a user has been attached to a pooled virtual machine by inspecting the 'Creation Date' in the 'Permission' sub-tab.

BZ#1064474

> It is now possible to populate the description of a new LUN disk with all, none, or part of the LUN id. The default setting uses the last four digits of the LUN id and can be configured via engine-config-manager. This value can be set to '-1' for the full LUN id to be used, or '0' for this feature to be ignored. A positive integer will populate the description with the corresponding number of characters of the LUN id.

BZ#1188759

> The Red Hat Enterprise Virtualization Manager now periodically monitors the CA certificate, the Manager certificate, and host certificates, and reports via the event log when a certificate is about to expire.
>
> The Manager also produces an alert for hosts when their certificates

have expired. A new button in the Hosts tab allows the user to enroll a certificate for the host. The same action is exposed in the REST API, by sending a POST request to '/api/hosts/{host:id}/enrollcertificate'.

BZ#1117031

Some operating systems are not supported on older CPU architectures. For example, Windows 8 does not run on Conroe-based hosts. When a cluster is set to a certain CPU model, the unsupported operating systems will no longer be available in the 'New Virtual Machine' and 'Edit Virtual Machine' windows for that cluster.

BZ#1111095

Previously, the DISK_STORAGE_MANIPULATION permission allowed users to perform live storage migration as standard. Now, a new permission, DISK_LIVE_STORAGE_MIGRATION, has been introduced to allow finer control over which users can perform live storage migration.

Upgrading to a version that includes this fix (3.6.0 or 3.5.1) will grant the new permission to all roles that included the DISK_STORAGE_MANIPULATION permission (DataCenterAdmin, StorageAdmin, ClusterAdmin, and relevant custom roles) to maintain functionality.

BZ#1043226

Host Network QoS adds the ability to configure QoS for networks on a host, enabling control of specific network traffic through the physical interfaces.

BZ#869296

This feature provides additional guest system information such as operating system, version, architecture, and the currently configured time zone.

Additionally, it shows a warning in form of an exclamation mark with a tool tip in the Administration Portal's virtual machine overview list when the operating system does not match the configured type, or if a Windows virtual machine timezone configured as a hardware timezone offset does not match the reported value.

This feature requires the Red Hat Enterprise Virtualization Manager Guest Agent to be installed on the guest operating system.

BZ#917026

It is now possible to directly boot a virtual machine from an ISO image directly over HTTPS without the need to have previously imported it on the ISO storage domain. This operation is performed via a VDSM hook. Please note that a stable HTTPS connection is required to avoid EIO errors.

BZ#1117484

> In administration portal it is now possible to open the create virtual machine dialog from the template list populated with the selected template.

BZ#1016839

> This release introduces the Host Upgrade Manager, a set of tools that notifies administrators of available updates to a host, and automates the process of putting the host into maintenance mode, updating its packages, and bringing it back to an Up state.

BZ#1012881

> This release adds prefix format support. Both netmask and prefix format are now supported in the netmask field of static IP configuration.

BZ#1009009

> With this update, the existing notifier filter mechanism has been extended to provide support for severities in an extra, optional field:
>
> ```
>   PATTERN=include|exclude:event[:severity][(subscriber)]
> ```
>
> The parser accepts the previous format for backward compatibility.
>
> An 'include' filter matches events with greater-or-equal severity. An 'exclude' filter matches events with lesser-or-equal severity.

BZ#950248

> A functionality has been added to the Red Hat Enterprise Virtualization Manager to report the SPICE client version that a user connects to a console with. Because of this, it is now possible to block older SPICE clients from connecting to virtual machine graphical consoles. To configure this, use the engine-config tool to set the RemoteViewerSupportedVersions and RemoteViewerNewerVersionUrl fields within vdc_options.

BZ#983019

> Previously, there was no way to disable the file transfer feature supported by the SPICE console. In this release, the feature can be disabled on a per virtual machine basis from the Administration Portal.

BZ#1163365

> With this release, role-bearing network interfaces must have IP configuration, to ensure that they can function as is expected from their role. Red Hat Enterprise Virtualization Manager will no longer

allow configuring role-bearing NICs with no means to actualize the role
function.

BZ#970711

With this release, the downtime during a virtual machine migration is
reported. This is the duration of the handover time needed to transfer
the execution from the source host to the destination host (the last
phase of migration).
Note: as part of this enhancement a more strict clock synchronization is
enforced between the Manager and hosts. Previously, there was an alert
when the host was 5 minutes off the Manager time; now it is 100 ms. The
reason is that for accurate downtime reporting the source and
destination hosts must have the same clock time. This may cause a lot of
new alerts in environments which are not configured properly. The
configuration option (used in engine-config) has changed from
'HostTimeDriftInSec' to 'HostTimeDriftInMS'.

BZ#957608

Normally, the latest client tools are attached to a Windows virtual
machine when it is started. However, if a virtual machine is running for
a long period of time, updates will not be applied automatically. A
notification has been added to the Red Hat Enterprise Virtualization
Manager to indicate the availability of new client tools for Windows
virtual machines. If there are new client tools available, an
exclamation mark will show beside the virtual machine status icon. The
updates take effect at virtual machine reboot.

BZ#1120670

Events can now be dismissed and restored in the Administration Portal.

BZ#976044

This feature allows users to enable multiple console clients (for
example, SPICE and VNC) for a virtual machine, so that users can switch
clients without restarting the virtual machine. The virtual machine can
now be configured with a "SPICE + VNC" console type, which enables this
behavior. The "Console Options" dialog for a virtual machine enables
protocols to be switched on a per-user basis.

BZ#996926

With this update, it is now possible to specify whether encryption is
set for fence agent fields in the Power Management tab of the New Host
and Edit Host windows. This enhancement addresses cases where certain
fence agents require the ability to set whether a field is encrypted.

BZ#1003492

The process to create a template from a virtual machine snapshot has

been simplified. Previously the user needed to create a virtual machine
from a snapshot, and then create a template from the virtual machine.
The new process enables the user to create the template directly from
the snapshot.

BZ#1025796

It is now possible to view all active sessions through the
Administration Portal. 'Guest Information' can be found under 'System'
in tree mode.

BZ#1032793

With this update, a GPU can be directly assigned to a virtual machine by
using a VFIO driver. As a GPU is a PCI-e device, the same constraints
apply as for a generic VFIO assignment. This feature is not supported on
ppc64le hosts due to qemu/host limitations. Additional steps must be
taken to ensure that the hypervisor GPU driver is not bound to the
device. For more information on hardware considerations for device
assignment, see https://access.redhat.com/documentation/en-
US/Red_Hat_Enterprise_Virtualization/3.6/html/Hardware_Considerations_fo
r_Implementing_SR-IOV/Additional_Considerations.html.

Red Hat Enterprise Virtualization only handles the assignment; system
setup must be implemented by a system administrator. For information on
supported GPUs and instructions for system setup, see
https://access.redhat.com/documentation/en-
US/Red_Hat_Enterprise_Linux/7/html/Virtualization_Deployment_and_Adminis
tration_Guide/sect-device-GPU.html.

For Windows 8 or 10 virtual machines, install the virtual machine using
only a SPICE or VNC console. After the installation is complete, install
the correct drivers for the device and restart the virtual machine. If
the GPU output is to be used, use the Windows display settings to
configure the physical screen connected to the GPU as either a
extension, mirror, or single screen of the operating system.

BZ#1201511

With this update the search capability of Red Hat Enterprise
Virtualization has been improved so that virtual machines can be
searched for based upon their virtual machine cluster compatibility
level, custom virtual machine emulated-machine type, and CPU model.

BZ#1201508

With this update virtual machines in the cluster level of Red Hat
Enterprise Virtualization 3.6 run with the pc-i440fx-rhel7.2.0 machine
type on x86 architecture and pseries-rhel7.2.0 on ppc64 architecture by
default. This means that virtual machines can use all of the
capabilities provided by these new machine types.

BZ#1225368

The communication between the Red Hat Enterprise Virtualization Manager and hosts has been improved. An event-based mechanism now communicates changes in virtual machine statuses instead of polling. This significantly reduces latency and network bandwidth. The improvement is especially noticeable in large-scale environments where hosts or data centers are remote, or connected over a slow connection.

BZ#1200067

Using the remote-viewer tool to connect to an ovirt:// URI now displays a menu that allows the user to change the CD image inserted in the virtual machine. This makes it possible to change the inserted CD while the virtual machine is running, without the need to access the Administration Portal or the User Portal.

BZ#1109154

QEMU capabilities for auto-convergence and/or Xor Binary Zero Run-Length-Encoding (XBZRLE) can be used to reduce virtual machine downtime and improve convergence during migration. This is supported by hierarchical configuration in 3 levels: global (engine-config), cluster, and virtual machine.

BZ#825045

Red Hat Enterprise Linux 7 virtualization hosts now support Virtual Function I/O (VFIO), a new UIO-like kernel driver that allows for a cleaner PCI device assignment architecture in KVM.

Note that IOMMU support needs to be enabled manually on the host's kernel command line. For GPU passthrough, additional kernel configuration is needed: for example, in the case of nvidia, the default driver needs to be avoided in both the host and the guest (via blacklisting of the OSS driver, and by binding the PCI device to the pci-stub driver on the host).

BZ#1110919

Users can now specify custom locations for database backups using an answer file generated by engine-setup that includes the following default key/value.

engine:
OVESETUP_CONFIG/engineDbBackupDir=str:/var/lib/ovirt-engine/backups
dwh:
OVESETUP_DWH_CONFIG/dwhDbBackupDir=str:/var/lib/ovirt-engine-dwh/backups
reports:
OVESETUP_REPORTS_CONFIG/reportsDbBackupDir=str:/var/lib/ovirt-engine-reports/backups

Users can edit these values to specify alternate file paths, which will be used by engine-setup when the '--config-append' option is added.

BZ#1101577

Red Hat Enterprise Virtualization 3.6 supports virtio-blk data plane, a QEMU feature that increases performance in heavy I/O scenarios. This allows virtual machines to be configured with additional QEMU threads on the host to offload disk I/O processing from the main emulation thread. The recommended number of threads varies, but is a trade-off between additional threads overhead (increasing the load on the hypervisor in general) and offloading benefits on per-disk basis (heavy I/O virtual machines would benefit most). Note, virtio-blk data plane must be configured per virtual machine for all disks.

This feature is not supported for the virtio-scsi interface.

BZ#678977

With this release, you can give a reason when moving a host to maintenance mode. In some cases, it can be of benefit to see why hosts are in maintenance. On the cluster level, you can set whether you want to be asked for the maintenance reason. If you set it to true (select the check box), you will see a Reason field when moving a host to maintenance.

BZ#1190853

HPE Superdome X servers support using the HPE BladeSystem power management agent. To enable power management on a Superdome X or a HPE C7000 blade, in the Administration Portal, use the add hosts or edit hosts dialog, and select 'hpblade' as the fence agent type. In the SSH Port field enter the blade bay number for a C7000 blade or, for Superdome X, the partition number you want to control. To use SSH, add secure=1 in the Options field. This is not required for telnet. Also add the command prompt to the Options field; for use with SSH the Options field would contain "secure=1,cmd_prompt=[prompt]".

BZ#1106562

With this update the Red Hat Enterprise Virtualization Manager provides a disk snapshot (volume) ID. This allows users to view the mapping of each Red Hat Enterprise Virtualization virtual machine's virtual disk to one or more data storage LUNs. This can be useful for troubleshooting storage performance issues. The disk snapshot (volume) IDs can be viewed in two places: Virtual Machines -> Snapshots -> Disks or Storage -> Disk Snapshots.

BZ#1083523

Previously, Red Hat Enterprise Virtualization provided three ways to access a virtual machine graphical console from web portals: using an ActiveX/Firefox SPICE plugin, .vv file MIME type association to remote-viewer (called Native in the user interface), and web-based clients (technology preview). In this version of Red Hat Enterprise Virtualization, the browser-specific plugins have been deprecated in

favor of more universal generic .vv file association, which allows for more flexibility (for example, using custom VNC viewers to view the console) and does not suffer from frequent browser changes and instabilities.

As a result, the default SPICE console mode has changed in new installations, which now default to the Native setting in the user interface. For existing installations that are upgraded this value is not changed, but can be configured by the global ClientModeSpiceDefault parameter with the engine-config tool.

BZ#664030

With this release, it is now possible to customize the system UUID of virtual machines.

CHAPTER 2. RHEA-2016:0375 OVIRT-HOSTED-ENGINE-SETUP

The bugs contained in this chapter are addressed by advisory RHEA-2016:0375. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2016-0375.html.

ovirt-hosted-engine-setup

BZ#1167074

> Previously, iptables-services was not a required package during hosted-engine deployment, so hosted-engine was unable to configure iptables. Now, iptables-services is required, so hosted-engine can configure iptables correctly.

BZ#1216172

> Previously, HostId was treated as an integer on the first host and as a string on additional hosts due to bad parsing of the answerfile, causing setup to fail. Now, this failure has been fixed by treating HostId as an integer on all hosts.

BZ#1251912

> Previously, if more than one data center was defined in Red Hat Enterprise Virtualization Manager while using tagged VLANs, 'hosted-engine --deploy' failed updating the VLAN property on the management network. Now, multiple data centers are handled correctly.

BZ#1077850

> With this release, you can deploy the Self-Hosted Engine with a virtual appliance image instead of an operating system ISO image.

BZ#1116469

> Previously, after a power outage or another disruptive event, the hosted-engine.lockspace file refused to accept new connections, and virtual machines failed to start. The --reinitialize-lockspace command line option has been added to the hosted-engine command, which reinitializes the sanlock lockspace file and wipes all locks. This option is available only in clusters in global maintenance mode with all high availability agents shut down. Additionally, the --force option can be used with the --reinitialize-lockspace option to override the safety checks, but should be only used with caution.

BZ#1198138

> With this release, users can attach a cloud-init ISO image to automatically configure the RHEVM Appliance from hosted-engine setup.

BZ#1196038

> With this update, hosted-engine-setup is able to deploy the Red Hat
> Enterprise Virtualization Manager bridge using a bonded VLAN interface.

BZ#1191074

> Previously, user defined iptables rules were overwritten when the host
> was added to the Manager and the automatic configuration of the firewall
> was turned off. Now, when the automatic configuration of the firewall is
> turned off the iptables rules will not be rewritten.

BZ#1169290

> Previously, the self-hosted engine was talking with the VDSM-wrapping
> vdsClient utility. This could cause SSL timeout errors for long sync
> commands. With this release, the self-hosted engine uses the vdscli
> library for storage operations instead of vdsClient, reducing SSL
> timeouts.

BZ#1155637

> Additional VDSM flags have been added to improve handling of pre-
> existing storage domains in hosted engine deployments.

BZ#1241426

> Previously, a value for OVEHOSTED_CORE/nodeSetup was set later in the
> hosted-engine setup script, which caused an error if the user decided to
> abort setup at the first question, when the script tried to access
> OVEHOSTED_CORE/nodeSetup and found nothing. This was fixed by initially
> setting a default value for OVEHOSTED_CORE/nodeSetup, so the user can
> abort setup at the first question without errors.

BZ#1153278

> With this update, users can now deploy self-hosted engine environments
> using fibre channel storage domains.

BZ#1139793

> Previously, every host in a self-hosted engine environment stored a
> configuration file for the self-hosted engine virtual machine. As a
> result, any change to the configuration required users to update the
> file on each host. Now, a single self-hosted engine configuration file
> is saved on shared storage, and can be accessed by all hosts.

BZ#1131918

> Previously, hosted-engine-setup was letting the user choose an invalid
> interface to create the management bridge on. Now, the list of allowed

responses is better filtered.

BZ#1205663

With this release, users can override the RHEVM Appliance's default memory value during hosted-engine deployment.

BZ#1120780

Users can now add an answerfile entry to enable automatic shutdown of a virtual machine at the end of the hosted engine deployment process. This allows for full automation of the process.

BZ#1108341

Previously, when manually deploying a hosted engine, deployment would fail if the admin password provided in the hosted-engine setup on the host did not match the admin password provided in the engine-setup on the virtual machine. Now, the user will be prompted again until the passwords match.

BZ#1101297

With this release, Self-Hosted Engine setup can be fully automated by using the RHEV-M Virtual Appliance and dynamically configuring it with cloud-init.

BZ#1215437

Previously, hosted-engine-setup had the capability to redeploy a host reusing the same host ID it was using before, but this was not allowed for host 1. With this release, it is now possible to redeploy host 1 reusing the same host ID.

BZ#1083025

With this release, Gluster storage can be used as a data domain on the Self-Hosted Engine.

BZ#1234915

This release rebases package(s) to version 1.3.0.

BZ#1211532

With this release, hosted engine setup auto-detects available Appliance images and suggests them for use during installation.

BZ#1227366

–

Fully automated setup is now the default option in hosted engine setup, using cloud-init and the RHEV-M Appliance.
Setup also recommends to install the appliance if it is missing.

BZ#1217106

The self-hosted engine setup work flow has been improved. By using the RHEV-M Virtual Appliance instead of the traditional operating system installation for the Manager virtual machine, the setup time now takes less than 10 minutes (excluding the RPM download time).

BZ#1305885

This release rebases package(s) to version 1.3.3.3.

BZ#1177800

Previously, using multipath IO, each device was represented as physical devices (/dev/sdx) mapped together. Now, instead of showing each physical device, the LUN list provided during deployment shows the GUID of the multipath device.

BZ#1169792

Previously, hosted-engine-setup was checking for the ISO image to be readable by the VDSM user, but it was the KVM user that needed to read it to work properly. If the ISO image was readable by the VDSM user without being readable by the KVM user, the check passed but the virtual machine was still unable to boot from the ISO. With this release, hosted-engine-setup has a coherent permission check. Now if the ISO image passes the check, it works.

BZ#1254838

Previously, hosted-engine setup always used the Manager CA certificate to trust the signature of the REST API (apache) certificate, but some users replaced that with an externally signed one. In that case the validation, and so the deployment of additional hosts, failed. With this release, hosted-engine setup lets the user specify the local path of an external CA file, or proceed in insecure mode, if validation with the internal CA certificate fails.

BZ#1253939

Previously, VDSM used a heuristic to configure the default route property based on the bridge name. This was removed in Red Hat Enterprise Virtualization 3.6 as the name was not mandatory. However, this caused the default bridge property not to be configured for a hosted engine. Now, VDSM explicitly sets the default bridge property.

BZ#1254104

Previously, hosted engine deployment over iSCSI using the RHEVM-appliance failed with an endless loop if the destination storage had insufficient space. This has now been fixed.

BZ#1249118

Previously, ovirt-hosted-engine-setup assumed that if an answer file was passed to the command line on an additional host, it was generated by ovirt-hosted-engine-setup in a previous execution on the first host. When the host is RHEV-H an answer file is always passed to ovirt-hosted-engine-setup. On additional hosts based on RHEV-H, the ovirt-hosted-engine-setup failed. With this release, an additional question has been added to setup to confirm whether the answer file is the one generated on the first host. Now setup works on additional hosts.

BZ#1032069

Previously, the correctness of an ISO file type was not enforced at the customization level, so the user could pass an invalid file, causing the virtual machine to refuse to boot. This error has now been fixed.

BZ#1101553

The tool previously used vdsClient for interacting with VDSM. Now it uses vdscli API, resulting in better error handling and more control.

BZ#1210763

Previously, the length of a username wasn't checked, even though the database field was limited to store only 50 characters. A user could enter a username longer than 50 characters, causing an error while storing the username into the database during the setup. With this release, username length is now validated to ensure it can fit in the database. The user can no longer enter usernames which may cause the database insert to fail due to their length. Please note: work is in progress to modify the database to allow longer usernames, so in future this limitation will be lifted.

BZ#1215436

It is now possible to connect to the serial console of the engine virtual machine using a text-only connection when deploying the self-hosted engine. Instructions for connecting to the serial console are provided during the deployment process.

CHAPTER 3. RHEA-2016:0377 RHEVM-SPICE-CLIENT

The bugs contained in this chapter are addressed by advisory RHEA-2016:0377. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2016-0377.html.

mingw-virt-viewer

BZ#1065936

> This release adds remote-viewer support for proxy authentication. With this update, SPICE can connect via a proxy that requires authentication (such as user name and password) in order to allow only authenticated users to connect via the proxy. As a result, when remote-viewer is configured to use a proxy requiring user authentication, the SPICE client passes authentication information to the proxy, and users connecting with the proxy are authenticated.

BZ#1155764

> Previously, the Red Hat Enterprise Virtualization Manager always opened the virt-viewer client in full screen mode, even if the client was not configured to open in full screen. This occurred because the CONTROLLER_AUTO_DISPLAY_RES flag was automatically set. This has been fixed and the Red Hat Enterprise Virtualization Manager now ignores CONTROLLER_AUTO_DISPLAY_RES flag, and the client now opens in full screen mode only when the CONTROLLER_SET_FULL_SCREEN flag is set.

BZ#1146016

> Previously, virt-viewer was installed per user (with the ALLUSERS property unspecified); when virt-viewer was updated, the Windows installer could not find the remote-viewer.exe client binary associated with the appropriate MIME type. To work around this, users needed to set the path to the remote-viewer.exe executable manually to access virtual machines managed by Red Hat Enterprise Virtualization Manager that used the Native Client console option. This was problematic for users without administrator access, so virt-viewer is now installed per-machine (ALLUSERS is now explicitly set to "1"). As a result, after an upgrade, users without administrator access can run the newly installed executable, and the Native Client console option works for all users.
>
> Note: Red Hat recommends uninstalling previous virt-viewer packages and cleaning up the registry for all users before installing this version.

BZ#1201387

> Isochronous devices (for example, a webcam and headset) previously caused remote-viewer to freeze temporarily due to high memory consumption, which was a result of poor buffer management in the virt-viewer client. This issue has now been improved by adding a buffer limit in the client, while still allowing enough data to be sent to update the remote device. Data can now be sent over a low bandwidth network without high memory consumption in the client.

BZ#1244880

Previously, Windows did not use the API to provide coordinates when using multiple monitors, which caused the virtual monitors to change position. Now, Windows uses the API for managing monitor position, so monitors stay in the expected position.

BZ#1169286

Previously, the shift keys behaved incorrectly on Windows machines accessed with virt-viewer. When both shift keys were pressed at the same time and then released, Windows sent only one release key event, causing the SPICE client to think the other shift key is still pressed, and the guest to input capital letters. This has been fixed so GDK sends an extra release event so that both shift keys are released, and guest input is as expected.

BZ#1136846

Previously, users could not control a Windows 7 virtual machine's sound volume with the guest volume controls. This occurred because the Windows client's audio back-end (GStreamer) did not have mute and volume control capabilities for playback and record. This was fixed by updating GStreamer to version 1.0, and implementing volume and mute capabilities in the directsoundsrc element. Volume control on Windows virtual machines now works as expected.

BZ#1275343

Previously, remote-viewer enabled more displays than necessary when running in full screen mode. This was fixed, and remote-viewer now only enables the same number of displays as the user's monitors when started in full screen mode.

BZ#1099327

Previously, Red Hat Enterprise Linux 7 hosts failed to correctly redirect USB devices to virtual machines. This occurred because of the channel creation order, where USBREDIR channel(s) get added to the session after the display channel. Since virt-viewer checked for a USBREDIR channel immediately as the display channel was added to the session, this returned a FALSE result for USBREDIR-capable virtual machines. This fix refactors the code so as to not depend on channel creation order, and virt-viewer no longer fails to redirect USB devices to virtual machines.

BZ#1181289

With this release, the virt-viewer package has been updated and is now based on virt-viewer 2.0. A list of bug fixes and enhancements can be found at: https://git.fedorahosted.org/cgit/virt-viewer.git/plain/NEWS.

BZ#1158007

The mingw-openssl package has been rebased to upstream version 1.0.2a-1 to fix potential security vulnerabilities in CVE-2014-3566, CVE-2014-3567, CVE-2014-3513, and CVE-2014-3568.

BZ#1161081

Previously, clipboard images were truncated when copying from the SPICE client to guest, as clipboard data size was incorrectly calculated as a text string. Data size is now calculated according to its type so that copying and pasting images between client and guest works as expected.

BZ#1135105

Previously, the SPICE client did not send a data message when transferring zero-sized files by drag and drop. Because the guest agent expects to receive a message, it would keep a file open, causing a leak. This has been fixed by adding a transfer message for zero-sized files to prevent guests from leaking file descriptors in the guest agent during the copy.

BZ#975835

Using the remote-viewer tool to connect to an ovirt:// URI now displays a menu that allows the user to change the CD image inserted in the virtual machine (VM). This makes it possible to change the inserted CD while the VM is running without the need to use Red Hat Enterprise Virtualization or the oVirt portal.

BZ#1273977

Improved memory handling so that applications using the libgovirt library no longer crash when multiple ISO domains with the same name are found in REST requests. Applications using libgovirt no longer crash due to memory corruption.

BZ#857087

With this release, the menu File->USB device selection is disabled on Windows client machines if usbclerk is not running and usbdk is not installed. For Spice usbredir to work, it is required that either usbclerk is running or usbdk is installed (or both). When it is known that usbredir will fail, it is better and clearer to the user to disable it.

BZ#1247927

Previously, when the first display on a Windows guest was disabled, it could not be enabled in a new session, because virt-viewer did not check events on the first display channel. With this fix, virt-viewer checks

for events in all display channels. It is now possible to disable a Windows guest's first display in one session, and enable it in a different session through the View menu in virt-viewer.

BZ#1200103

The virt-viewer tool can now report the SPICE client version being used to connect to a console. This allows a minimum SPICE client version to be set using the engine-config tool, and older SPICE clients to be blocked from connecting to virtual machine graphical consoles.

BZ#1055600

Previously, shortcuts were missing from the "Send key" menu when virt-viewer was started using the ActiveX plug-in. This occurred because enabling hotkeys required a rebuild of the menu containing the new hotkeys, which was done before parsing the string containing the new hotkeys. This has been fixed and the new shortcuts now properly show in the menu when connecting using the plug-in.

BZ#1205804

When viewing a remote SPICE session in windowed mode and resizing the viewer window be wider than the virtual screen, the viewer window was automatically resized to fit the screen width, and the resolution of the remote screen was scaled down. This unintentional behavior has been corrected, and in the described scenario, the both the viewer window and the resolution of its contents stay unchanged.

BZ#1233422

It is now possible to configure the position in which the guest displays in multi-monitor setups. To do so, edit the ~/.config/virt-viewer/settings file. For more information about this feature, refer to the CONFIGURATION section of the remote-viewer(1) manual page.

BZ#1242600

Previously, connecting to a virtual machine with remote-viewer when the resize-guest option was disabled changed the guest's display resolution. The code has been edited so that when resize-guest is disabled, virt-viewer no longer requests a display update. This also ensures that incorrect monitor configuration will not be sent, and the guest's display resolution will not change upon connection.

BZ#1244869

Prior to this update, Microsoft Windows virtual machines with multi-monitor setup in some cases waited for an image that had already been deleted, which caused the client to become unresponsive. With this update, reference counting has been added to channel images in the

client and thread safety has been improved for image caching in the server. As a result, Windows virtual machines no longer hang due to the mentioned reasons.

spice-activex-win

BZ#1171717

Previously, the SPICE ActiveX plug-in sometimes crashed if a dynamic menu string was set larger than 4096 bytes. This happened because the menu string property in the SPICE ActiveX plug-in was stored in a 4096-byte array, and when a new menu string was set, its size was not checked. This was fixed by making the menu string property dynamically allocated according to the size of the new string. Additionally, to limit the memory allocated, the SPICE ActiveX plug-in also checks and rejects all strings that are too large. Strings are now sent to the client as UTF-8, and log messages size is limited. As a result, the plug-in no longer crashes when given a large menu string.

spice-usbdk-win

BZ#1144043

Redirecting or releasing a USB device can take 3-5 seconds in some cases, and since this was performed synchronously in spice-gtk, it caused virt-viewer to freeze for several seconds. In this release, these operations have been made asynchronous on the spice-gtk level, and users will no longer experience virt-viewer freezing when redirecting or releasing a USB device.

BZ#1202371

UsbDk has been added to Red Hat Enterprise Virtualization as a preferred alternative to usbclerk. UsbDk is a filter driver that is able to allow and disallow access to USB devices from spice-client on Windows machines. This allows virtual machines access to the client USB devices through SPICE client (remote-viewer) running on Windows machines.

CHAPTER 4. RHBA-2016:0362 VDSM

The bugs contained in this chapter are addressed by advisory RHBA-2016:0362. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2016-0362.html.

vdsm

BZ#1188251

> Previously, VDSM did not consume pre-defined ifcfg interfaces and did not consider them as belonging to it. When a setupNetworks command was issued and failed, VDSM failed to restore the original ifcfg file. Note that this bug only occurs when failing to set up network on top of a pre-existing ifcfg file. Now, VDSM stores pre-defined ifcfg files before they are modified, even with unified persistence.

BZ#1182094

> Previously, NUMA statistics were collected every time VDSM was queried for host statistics. This resulted in a higher load and unnecessary delays as collecting the data was time consuming as an external process was executed. Now, NUMA statistic collection has been moved to the statistics threads and the host statistic query reports the last collected result.

BZ#1207610

> The Memory Overcommitment Manager (MOM) policy formula for CPU limits previously used fixed constants and divided those by the amount of CPUs. The result was too low on hosts with more than 100 CPUs and the value was refused by libvirt, which caused performance degradation in virtual machines. The CPU limit formulas have been improved and as a result, the CPU limits can now handle any number of CPUs.

BZ#1269424

> Previously, VDSM memory consumption continually increased on some environments, caused by a memory leak in VDSM. The code has been updated to eliminate the VDSM memory leak, and there is no longer a memory usage increase when running VDSM.

BZ#1215967

> With this update prepareImage is now called by VDSM to mount the required NFS storage to deploy additional hosts using NFS.

BZ#1155583

> When live merging snapshots on a block storage domain, the merge target volume is pro-actively extended to accommodate active writing of the source volume. This may cause some over-extension of the target volume.

BZ#1256949

Previously, a Memory Overcommitment Manager (MOM) policy rule computed KSM's sleep_millisecs value using a division with the amount of host memory being part of the divider. As a result, the sleep_millisecs value dropped below 10ms on hosts with more than 16GiB of RAM. That value was invalid and too aggressive, causing a huge CPU load on the host. In this release, the sleep_millisecs value was bounded to never drop below 10ms, thus improving the CPU load on affected machines.

BZ#1247075

Previously, excessive thread usage in VDSM and Python runtime architecture caused poor VDSM performance on multicore hosts. VDSM now supports affinity in order to pin its processes to specific cores. CPU usage is reduced as a result of pinning VDSM threads to a smaller number of cores.

BZ#1247058

NUMA nodes can exist without memory (for example, when hotswapping memory modules). This was not considered in VDSM, causing the statistics reporting mechanism (getVdsStats) to break. Now, this error has been fixed by explicitly checking for NUMA nodes with zero memory, and returning a memory usage of 100%.

BZ#1112861

Previously, a host became non-operational when a network interface was blocked in a multipath environment where the host contained two network devices both configured for the same subnet and configured with an iSCSI bond. This occurred because by the time the failing path is active or ready again, it is not considered by multipath anymore until a new host changes state to Up. Now, when configuring the iSCSI bond network interfaces, VDSM configures the multipath with the correct interface passed by the engine. As a result, when one of the network interfaces on the same subnet becomes non-responsive, path 2 will be used to reach the iSCSI target, and hosts will continue to operate normally.

BZ#1226911

In some cases after migration fails, many error messages flooded the VDSM log, which caused one of the VDSM threads to consume 100% CPU. This was caused by incorrect usage of epoll, which has been fixed in this update.

BZ#1261007

Previously, when using a separate display network, the VDSM service ignored the specific listening IP address and listened to all connections via the management network. With this update, the VDSM service uses the display network settings as expected.

BZ#1156194

Previously, when a virtual machine was resumed from a suspended state its time was not updated, resulting in incorrect time in the guest operating system. This happened because the functionality to update time on resume was missing in VDSM. Now, the functionality has been added and the time is now updated after a virtual machine is resumed from a suspended state. As long as the guest operating system supports this feature and qemu-guest-agent is running. It is still recommended to have NTP services running in the guest operating system to provide precise time for the guest.

BZ#1267444

Some host deployments or upgrades failed previously because VDSM did not handle empty supplementary groups in the sanlock process. These groups were left unconfigured when a race occurred in the supplementary groups configuration during sanlock startup. In this release, VDSM handles empty supplementary groups, and host deployment or upgrade will not fail if VDSM checks the sanlock configuration before supplementary groups are configured.

BZ#1205058

The previous virtual machine payload ISO used only the Rock Ridge extension, and as a result Windows virtual machines could not use payloads with long filenames. This fix adds Microsoft's Joliet filesystem extension to the generated ISO, so that payloads with long filenames are now displayed correctly on both Windows and Linux systems.

BZ#1203891

On the public internet, there are many random attempts to expose well-known ports on servers, in order for VNC to gain access to the machine behind the ports. These remote connection attempts triggered disconnect events which locked the console screen, even if the connection had not established a valid VNC or SPICE session. Now, only the client IP and port of the current known session can disconnect the console.

BZ#974510

Previously, when a user was connected to a SPICE console via a SPICE proxy, the console connection would drop during virtual machine migration. This happened because the client machine was not able to connect to the display on the destination host machine. Now, for both SPICE and VNC, console access is not interrupted.

Note the following limitations:
1. On virtual machines where both displays (SPICE and VNC) are configured, the console connection persists only when using SPICE, and will otherwise fail.

2. Uninterrupted console access only works with remote-viewer and plugins. It does not work with integrated web-based clients (noVNC and SPICE HTML5) or with third-party VNC viewers.

BZ#1296936

With this update, the MOM component no longer fails to enforce QoS policies, KSM, and memory ballooning.

BZ#1182247

Red Hat Enterprise Virtualization supports virtual machines with up to 240 vCPUs. The previous version supported a maximum of 160 vCPUs.

BZ#922744

With this change, the glusterfs-cli package must be installed. Please note that in the VDSM spec file, there is no dependency to glusterfs-cli. This means that VDSM installation will succeed even if glustefs-cli is not installed, and that the glusterfs-cli package must be installed manually.

BZ#880738

Previously, when accessing a device that was no longer available various commands used to block for several minutes because of various issues in the underlying platform. This caused long delays in VDSM threads waiting on blocked commands, leading to long delays and timeouts in various flows. Now, there have been several fixes in the kernel and device-mapper-multipath and VDSM multipath has been improved. VDSM now requires a fixed version of these components. When accessing a device that is unavailable minimal delays are expected when the device is first accessed. After the first access the device is now considered faulty by multipath and no further delays are expected.

BZ#1229177

This change disables the ability to use VDSM in clusters from 3.0 to 3.3. Red Hat Enterprise Virtualization 3.6 drops support for cluster levels between 3.0 and 3.3 and Manager versions 3.3 and below. Support for Red Hat Enterprise Virtualization Manager 3.4 is still available as tech-preview.

BZ#1286997

With this update, the vdsm-hook-vmfex-dev package is included with VDSM. Users can now connect their virtual machine network profiles to Cisco UCS-defined port profiles.

BZ#1004101

Previous versions of VDSM used an inadequate algorithm to calculate the downtime of virtual machines in certain scenarios, causing migration to fail when virtual machines were running heavy loads. This version of VDSM implements a new algorithm to estimate a virtual machine's downtime, and migrations in these scenarios converge more easily.

BZ#1128881

Previously, VDSM reported all channel devices as 'Unknown' device types with a warning. This was not correct and has now been fixed.

BZ#1142776

Simultaneous migration of many virtual machines could create a deadlock between the threads that monitor the hosts. As a result, the hosts were not monitored, thus the status of their virtual machines were not updated by the Red Hat Enterprise Virtualization Manager. The virtual machines are now monitored to prevent the deadlock.

BZ#1219364

With this update, the CPU usage of the VDSM management process has been reduced. This increases the performance and scalability of each hypervisor. As a consequence some of the functionality of VDSM was separated out to a standalone "MOM" process. These functions are KSM, ballooning and QoS policy enforcing.

BZ#1219903

Previously, libvirt reported a "metadata not found" error in vdsm.log when a query was made to get the missing metadata element. This was not actually an error, but a misleading message issued by VDSM. An empty metadata element has been added to the code so that this message will no longer appear in the log.

BZ#1217401

Previously, adding a direct LUN to a virtual machine sometimes timed out. This occurred because a physical volume (PV) create test is performed for each device when calling getDeviceList. Since a PV create test requires significant resources, it affects the response time of GetDeviceList in scale setups, sometimes causing timeouts on the Red Hat Enterprise Virtualization Manager. This has been fixed and the PV create test can now be skipped using a flag. If the PV test is needed to know the usage state of the device, it can be run on specific devices, therefore minimizing the impact and decreasing the user waiting time on the following operations:
1. UI - Add direct LUN disk
2. UI/REST - Add ISCSI/FC storage domain
3. UI/REST - Edit ISCSI/FC storage domain

4. UI/REST - Extend ISCSI/FC storage domain
5. REST - Add direct LUN disk (if <host> parameter is provided)
6. UI/REST - Import an iSCSI/FCP domain

BZ#1215610

Previously, Red Hat Enterprise Virtualization incorrectly configured the hypervisor for certain Windows versions, resulting in significant time drift on Windows virtual machines running high CPU loads. Code has been added to VDSM to inject periodic RTC interrupts, to prevent lost interrupts which caused time drift in Windows guests. The recommended hypervisor settings are now configured for Windows versions and there is no longer time drift in Windows virtual machines.

BZ#1215387

With the release of Red Hat Enterprise Virtualization 3.6, VDSM will no longer support engines older than Red Hat Enterprise Virtualization 3.3.0. It will, however, continue to support clusters and data centers with compatibility versions lower than 3.3.

BZ#1126206

File-type storage domains now use separate IOProcess instances. This improves performance, and prevents one slow or unreachable storage domain from affecting other storage domains.

BZ#965929

When Red Hat Enterprise Virtualization configures a network on a host, it generates several ifcfg files with specific content. If a user wants to tweak the content by adding or removing any initscripts, they must deploy a hook script to do so whenever the ifcfg file is rewritten.

The example hook script will add the following entries to the ifcfg file of nic 'ens11' when ifcfg is modified:
  USERCTL=yes
  ETHTOOL_OPTS="autoneg on speed 1000 duplex full"

To add the hook script to VDSM hooks, place the file in /usr/libexec/vdsm/hooks/before_ifcfg_write, and ensure the VDSM has permissions to the file.

The VDSM checks this directory every time ifcfg configuration is changed, and executes each script in this directory.
As input to the script, the VDSM will pass the path to a json file containing the ifcfg file data, for example:

```
{
  "config":
"DEVICE=ens13\nHWADDR=52:54:00:d1:3d:c8\nBRIDGE=z\nONBOOT=yes\nMTU=1500\
nNM_CONTROLLED=no\nIPV6INIT=no\n",
  "ifcfg_file": "/etc/sysconfig/network-scripts/ifcfg-ens7"
```

```
}
```

Modified ifcfg file contents (under the "config" entry) can be written
to a json file, and will be used by VDSM as the new ifcfg file content.
If no file is given, VDSM will use the unmodified content.

The following is a description of the example hook script.

Reading in the data from the json file:

```
  hook_data = hooking.read_json()
```

Getting the value of the new ifcfg file content:

```
  config_data = hook_data['config']
```

Getting the name of the ifcfg file which will be modified:

```
  ifcfg_file = hook_data['ifcfg_file']
```

Modify and write the content of the ifcfg file:

```
  config_data += "USERCTL=yes\nETHTOOL_OPTS=\"autoneg on speed 1000
duplex full\"\n"
  hook_data['config'] = config_data
  hooking.write_json(hook_data)
```

The file also dumps the data read from the json file to file
(/tmp/hook_data), to show the format of the input json file:

```
  with open("/tmp/hook_data",mode='w') as file:
      file.write( json.dumps(hook_data))
```

BZ#1123052

Red Hat recommends using glusterfs volumes that are replica type and
replica count 1 or 3. Previously, VDSM did not validate glusterfs
volumes when adding glusterfs storage domains. As a result,
administrators could configure glusterfs volumes with an unsupported
replica type and count. VDSM now validates the required parameters
before using a glusterfs volume as a storage domain, and if the
glusterfs volume configuration is not supported, there is a warning.

BZ#1206231

VDSM in Red Hat Enterprise Virtualization 3.6 no longer supports Red Hat
Enterprise Linux 6. In this bug, el6 support was removed from the
vdsm.spec file to reduce the maintenance required.

CHAPTER 5. RHBA-2016:0378 OVIRT-NODE

The bugs contained in this chapter are addressed by advisory RHBA-2016:0378. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2016-0378.html.

ovirt-node

BZ#1237212

> Previously, FCoE storage did not work by default due to FCoE services not being enabled. Now, FCoE services are enabled by default in the Red Hat Enterprise Virtualization Hypervisor.

BZ#1203176

> The Red Hat Enterprise Virtualization Hypervisor (RHEV-H) could not be installed on systems with more than 26 disks in /dev (including multiple paths to the same disk) due to the local disk being filtered out of the installer. The filtering now operates on single disks instead of lists of disks ensuring that RHEV-H can be installed.

BZ#1264195

> Fixed an error in the code that caused device discovery to take longer than it should when installing Red Hat Enterprise Virtualization.

BZ#1266530

> Previously, the Red Hat Enterprise Virtualization Hypervisor was not updated with the VNC and SPICE ports that VDSM sets in QEMU. Updated the port range in the firewall to 5900:6923 for VNC and Spice usage, to enable customers to see the graphical interface.

BZ#1177308

> Free space validation is now executed when enabling or disabling the "Fill all space" option when installing Red Hat Enterprise Virtualization Hypervisor. Previously, the amount of free space in the storage volume page of the installer was only validated when entering the page, which made it possible to continue an installation with negative free space or partition sizes.

BZ#1187034

> Added a confirmation message to indicate whether a file has been successfully or unsuccessfully persisted.

BZ#1269752

> Fixed the Red Hat Subscription Management (RHSM) daemon to be enabled only if RHSM is configured. This prevents the RHSM daemon from writing to the log files when RHSM is not configured.

to the log files when RHSM is not configured.

BZ#1255277

RHN Classic is being replaced with Red Hat Subscription Management (RHSM) interfaces. Updated the subscription page text to refer to RHSM instead of RHN.

BZ#1203837

With this update XFS based storage for local storage domains is now supported on the Red Hat Enterprise Virtualization Hypervisor. Red Hat Enterprise Virtualization 3.6 is not required to use XFS based storage domains.

BZ#1233035

Previously, creating a new bond used the balance-rr mode by default and this could lead to connection issues. Now, active-backup mode is used by default for newly created bonds as active-backup mode does not increase the available bandwidth but is safe to use in any setup.

BZ#1212902

Fixed the 'redhat-release'file to include the 'release' keyword to ensure that the Red Hat Enterprise Virtualization Hypervisor version is displayed when using the command:

'cat /etc/redhat-release'

BZ#1215016

Fixed a python file issue which prevented configuration migration from occuring when updating Red Hat Enterprise Virtualization Hypervisor from 6.5 to 6.6.

BZ#1238775

With this release, modifications to a persisted multipath.conf file are applied at boot time when previously they were not.

BZ#1061081

With this update the NTP service can now be configured to not start automatically after a manual or autoinstall as the ntpd package provides ntp.conf with pre-configured servers which are invalid for Red Hat Enterprise Virtualization Hypervisor deployments without internet access. This update allows users to use ntp=off in grub or autoinstall to disable the NTP daemon. To re-enable the NTP daemon users must use the the text user interface (TUI) to set a non-default NTP server, such as an internal NTP server.

BZ#1182048

> On some Red Hat Enterprise Virtualization Hypervisor hosts, a device-mapper error (device-mapper: table: 253:6: multipath: error getting device) is displayed on the login screen. This error does not have a functional impact on those hosts.

BZ#1148080

> With this update the perf tool has been added to the Red Hat Enterprise Virtualization Hypervisor to allow for better performance monitoring.

BZ#1218941

> Fixed an issue where iSCSI information is missing for devices when re-installing Red Hat Enterprise Virtualization Hypervisor.

BZ#1153417

> Red Hat Enterprise Virtualization Hypervisor users can now configure the logrotate interval and maximum size via automated installation options. The parameters can be used together or separately. If used separately, the default value for 'logrotate_max_size' is '1024' and the default value for 'logrotate_interval' is 'daily'.

ovirt-node-plugin-hosted-engine

BZ#1208780

> Users can now perform maintenance tasks on a RHEV-H self-hosted engine through the text user interface (TUI) without entering the rescue shell.

ovirt-node-plugin-vdsm

BZ#1231379

> The 'vdsm-reg' tool has been deprecated and replaced with 'vdsm-tool register' to register supported distributions with Red Hat Enterprise Virtualization.

APPENDIX A. REVISION HISTORY

Revision 3.6-1   **Wed 09 Mar 2016**   **Red Hat Enterprise Virtualization Documentation Team**

Initial creation for the Red Hat Enterprise Virtualization 3.6 release.