



Red Hat Virtualization 4.4

Technical Reference

The technical architecture of Red Hat Virtualization environments

Red Hat Virtualization 4.4 Technical Reference

The technical architecture of Red Hat Virtualization environments

Red Hat Virtualization Documentation Team

Red Hat Customer Content Services

rhev-docs@redhat.com

Legal Notice

Copyright © 2022 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This document describes the concepts, components, and technologies used in a Red Hat Virtualization environment.

Table of Contents

CHAPTER 1. INTRODUCTION	4
1.1. RED HAT VIRTUALIZATION MANAGER	4
1.2. RED HAT VIRTUALIZATION HOST	4
1.3. COMPONENTS THAT SUPPORT THE MANAGER	6
1.4. STORAGE	7
1.5. NETWORK	7
1.6. DATA CENTERS	10
1.7. DATA CENTER AND CLUSTER COMPATIBILITY LEVELS	10
CHAPTER 2. STORAGE	12
2.1. STORAGE DOMAINS OVERVIEW	12
2.2. TYPES OF STORAGE BACKING STORAGE DOMAINS	12
2.3. STORAGE DOMAIN TYPES	13
2.4. STORAGE FORMATS FOR VIRTUAL DISKS	14
2.5. VIRTUAL DISK STORAGE ALLOCATION POLICIES	14
2.6. STORAGE METADATA VERSIONS IN RED HAT VIRTUALIZATION	15
2.7. STORAGE DOMAIN AUTORECOVERY IN RED HAT VIRTUALIZATION	16
2.8. THE STORAGE POOL MANAGER	17
2.9. STORAGE POOL MANAGER SELECTION PROCESS	18
2.10. EXCLUSIVE RESOURCES AND SANLOCK IN RED HAT VIRTUALIZATION	19
2.11. THIN PROVISIONING AND STORAGE OVER-COMMITMENT	20
2.12. LOGICAL VOLUME EXTENSION	20
2.13. THE EFFECT OF STORAGE DOMAIN ACTIONS ON STORAGE CAPACITY	21
CHAPTER 3. NETWORKING	23
3.1. HOST NETWORKING	23
3.2. VIRTUAL MACHINE NETWORKING TYPES	23
3.3. NETWORK ARCHITECTURE	26
3.4. BASIC NETWORKING TERMS	26
3.5. NETWORK INTERFACE CONTROLLER	27
3.6. LINUX BRIDGE	27
3.7. BONDS	27
3.8. BONDING MODES	28
3.9. SWITCH CONFIGURATION FOR BONDING	28
3.10. VIRTUAL NETWORK INTERFACE CARDS	29
3.11. VIRTUAL LAN (VLAN)	30
3.12. NETWORK LABELS	30
3.13. CLUSTER NETWORKING	31
3.14. LOGICAL NETWORKS	33
3.15. REQUIRED NETWORKS, OPTIONAL NETWORKS, AND VIRTUAL MACHINE NETWORKS	33
3.16. PORT MIRRORING	34
3.17. HOST NETWORKING CONFIGURATIONS	34
CHAPTER 4. POWER MANAGEMENT	36
4.1. INTRODUCTION TO POWER MANAGEMENT AND FENCING	36
4.2. POWER MANAGEMENT BY PROXY IN RED HAT VIRTUALIZATION	36
4.3. POWER MANAGEMENT	36
4.4. FENCING	37
4.5. SOFT-FENCING HOSTS	38
4.6. USING MULTIPLE POWER MANAGEMENT FENCING AGENTS	39
CHAPTER 5. LOAD BALANCING, SCHEDULING, AND MIGRATION	40


5.1. LOAD BALANCING, SCHEDULING, AND MIGRATION	40
5.2. LOAD BALANCING POLICY	40
5.3. LOAD BALANCING POLICY: VM_EVENLY_DISTRIBUTED	40
5.4. LOAD BALANCING POLICY: EVENLY_DISTRIBUTED	41
5.5. LOAD BALANCING POLICY: POWER_SAVING	41
5.6. LOAD BALANCING POLICY: NONE	42
5.7. LOAD BALANCING POLICY: CLUSTER_MAINTENANCE	42
5.8. HIGHLY AVAILABLE VIRTUAL MACHINE RESERVATION	42
5.9. SCHEDULING	42
5.10. MIGRATION	43
CHAPTER 6. DIRECTORY SERVICES	44
6.1. DIRECTORY SERVICES	44
6.2. LOCAL AUTHENTICATION: INTERNAL DOMAIN	44
6.3. REMOTE AUTHENTICATION USING GSSAPI	44
CHAPTER 7. TEMPLATES AND POOLS	46
7.1. TEMPLATES AND POOLS	46
7.2. TEMPLATES	46
7.3. POOLS	47
CHAPTER 8. VIRTUAL MACHINE SNAPSHOTS	48
8.1. SNAPSHOTS	48
8.2. LIVE SNAPSHOTS IN RED HAT VIRTUALIZATION	48
8.3. SNAPSHOT CREATION	49
8.4. MONITORING SNAPSHOT HEALTH WITH THE IMAGE DISCREPANCIES TOOL	50
8.5. SNAPSHOT PREVIEWS	51
8.6. SNAPSHOT DELETION	52
CHAPTER 9. HARDWARE DRIVERS AND DEVICES	54
9.1. VIRTUALIZED HARDWARE	54
9.2. STABLE DEVICE ADDRESSES IN RED HAT VIRTUALIZATION	54
9.3. CENTRAL PROCESSING UNIT (CPU)	54
9.4. SYSTEM DEVICES	55
9.5. NETWORK DEVICES	55
9.6. GRAPHICS DEVICES	55
9.7. STORAGE DEVICES	56
9.8. SOUND DEVICES	56
9.9. SERIAL DRIVER	56
9.10. BALLOON DRIVER	56
APPENDIX A. ENUMERATED VALUE TRANSLATION	58
APPENDIX B. EVENT CODES	59
APPENDIX C. TIMEZONES	232
APPENDIX D. LEGAL NOTICE	237

CHAPTER 1. INTRODUCTION

1.1. RED HAT VIRTUALIZATION MANAGER

The Red Hat Virtualization Manager provides centralized management for a virtualized environment. A number of different interfaces can be used to access the Red Hat Virtualization Manager. Each interface facilitates access to the virtualized environment in a different manner.

Figure 1.1. Red Hat Virtualization Manager Architecture



The Red Hat Virtualization Manager provides graphical interfaces and an Application Programming Interface (API). Each interface connects to the Manager, an application delivered by an embedded instance of the Red Hat JBoss Enterprise Application Platform. There are a number of other components which support the Red Hat Virtualization Manager in addition to Red Hat JBoss Enterprise Application Platform.

1.2. RED HAT VIRTUALIZATION HOST


A Red Hat Virtualization environment has one or more hosts attached to it. A host is a server that provides the physical hardware that virtual machines make use of.

Red Hat Virtualization Host (RHVH) runs an optimized operating system installed using a special, customized installation media specifically for creating virtualization hosts.

Red Hat Enterprise Linux hosts are servers running a standard Red Hat Enterprise Linux operating system that has been configured after installation to permit use as a host.

Both methods of host installation result in hosts that interact with the rest of the virtualized environment in the same way, and so, will both be referred to as **hosts**.

Figure 1.2. Host Architecture



Kernel-based Virtual Machine (KVM)

The Kernel-based Virtual Machine (KVM) is a loadable kernel module that provides full virtualization through the use of the Intel VT or AMD-V hardware extensions. Though KVM itself runs in kernel space, the guests running upon it run as individual QEMU processes in user space. KVM allows a host to make its physical hardware available to virtual machines.

QEMU

QEMU is a multi-platform emulator used to provide full system emulation. QEMU emulates a full system, for example a PC, including one or more processors, and peripherals. QEMU can be used to launch different operating systems or to debug system code. QEMU, working in conjunction with KVM and a processor with appropriate virtualization extensions, provides full hardware assisted virtualization.

Red Hat Virtualization Manager Host Agent, VDSM

In Red Hat Virtualization, **VDSM** initiates actions on virtual machines and storage. It also facilitates inter-host communication. VDSM monitors host resources such as memory, storage, and networking. Additionally, VDSM manages tasks such as virtual machine creation, statistics accumulation, and log collection. A VDSM instance runs on each host and receives management commands from the Red Hat Virtualization Manager using the re-configurable port **54321**.

VDSM-REG

VDSM uses **VDSM-REG** to register each host with the Red Hat Virtualization Manager. **VDSM-REG** supplies information about itself and its host using port **80** or port **443**.

libvirt

Libvirt facilitates the management of virtual machines and their associated virtual devices. When Red Hat Virtualization Manager initiates virtual machine life-cycle commands (start, stop, reboot), VDSM invokes libvirt on the relevant host machines to execute them.

Storage Pool Manager, SPM

The Storage Pool Manager (SPM) is a role assigned to one host in a data center. The SPM host has sole authority to make all storage domain structure metadata changes for the data center. This includes creation, deletion, and manipulation of virtual disks, snapshots, and templates. It also

includes allocation of storage for sparse block devices on a Storage Area Network(SAN). The role of SPM can be migrated to any host in a data center. As a result, all hosts in a data center must have access to all the storage domains defined in the data center.

Red Hat Virtualization Manager ensures that the SPM is always available. In case of storage connectivity errors, the Manager re-assigns the SPM role to another host.

Guest Operating System

Guest operating systems do not need to be modified to be installed on virtual machines in a Red Hat Virtualization environment. The guest operating system, and any applications on the guest, are unaware of the virtualized environment and run normally.

Red Hat provides enhanced device drivers that allow faster and more efficient access to virtualized devices. You can also install the Red Hat Virtualization Guest Agent on guests, which provides enhanced guest information to the management console.

1.3. COMPONENTS THAT SUPPORT THE MANAGER

Red Hat JBoss Enterprise Application Platform

Red Hat JBoss Enterprise Application Platform is a Java application server. It provides a framework to support efficient development and delivery of cross-platform Java applications. The Red Hat Virtualization Manager is delivered using Red Hat JBoss Enterprise Application Platform.



IMPORTANT

The version of the Red Hat JBoss Enterprise Application Platform bundled with Red Hat Virtualization Manager is **not** to be used to serve other applications. It has been customized for the specific purpose of serving the Red Hat Virtualization Manager. Using the Red Hat JBoss Enterprise Application Platform that is included with the Manager for additional purposes adversely affects its ability to service the Red Hat Virtualization environment.

Gathering Reports and Historical Data

The Red Hat Virtualization Manager includes a data warehouse that collects monitoring data about hosts, virtual machines, and storage. A number of pre-defined reports are available. Customers can analyze their environments and create reports using any query tools that support SQL.

The Red Hat Virtualization Manager installation process creates two databases. These databases are created on a Postgres instance which is selected during installation.

- The **engine** database is the primary data store used by the Red Hat Virtualization Manager. Information about the virtualization environment like its state, configuration, and performance are stored in this database.
- The **ovirt_engine_history** database contains configuration information and statistical metrics which are collated over time from the **engine** operational database. The configuration data in the **engine** database is examined every minute, and changes are replicated to the **ovirt_engine_history** database. Tracking the changes to the database provides information on the objects in the database. This enables you to analyze and enhance the performance of your Red Hat Virtualization environment and resolve difficulties.

For more information on generating reports based on the **ovirt_engine_history** database see the [History Database](#) in the *Red Hat Virtualization Data Warehouse Guide*.



IMPORTANT

The replication of data to the `ovirt_engine_history` database is performed by the **RHEVM History Service**, `ovirt-engine-dwhd`.

Directory services

Directory services provide centralized network-based storage of user and organizational information. Types of information stored include application settings, user profiles, group data, policies, and access control. The Red Hat Virtualization Manager supports Active Directory, Identity Management (IdM), OpenLDAP, and Red Hat Directory Server 9. There is also a local, internal domain for administration purposes only. This internal domain has only one user: the admin user.

1.4. STORAGE

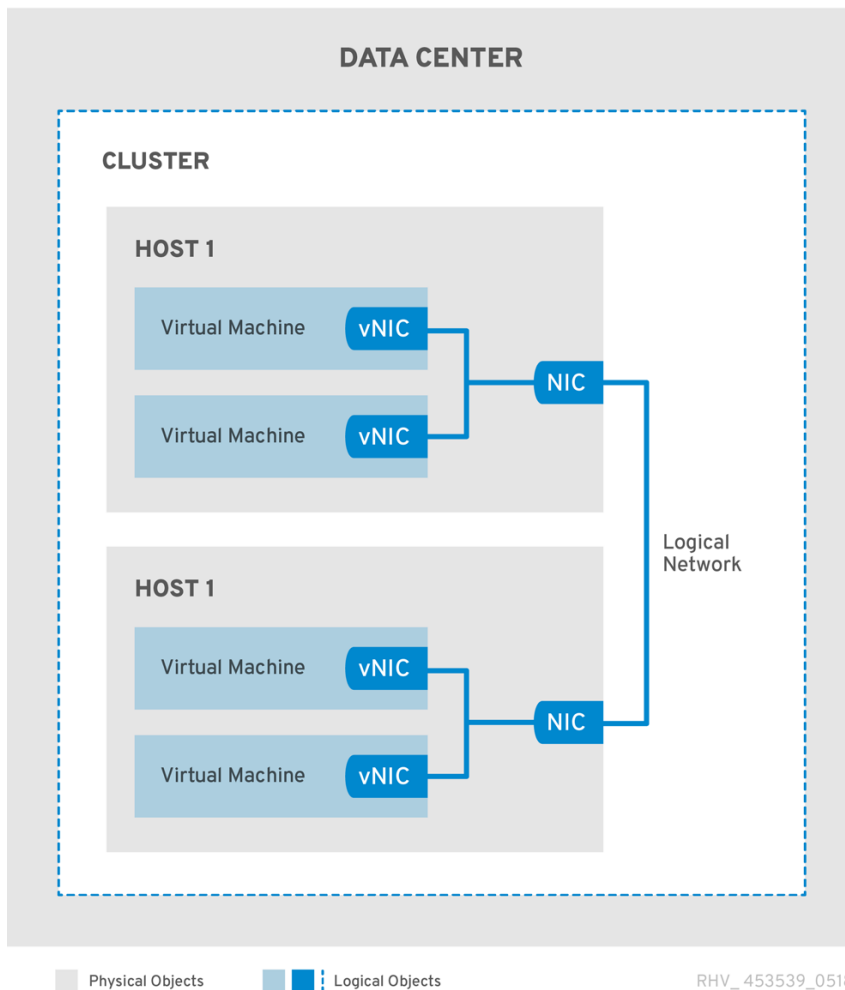
Red Hat Virtualization uses a centralized storage system for virtual disks, templates, snapshots, and ISO files. Storage is logically grouped into storage pools, which are comprised of storage domains. A storage domain is a combination of storage capacity and metadata that describes the internal structure of the storage. See [Storage Domain Types](#)

The data domain is the only one required by each data center. A data storage domain is exclusive to a single data center. Export and ISO domains are optional. Storage domains are shared resources, and must be accessible to all hosts in a data center.

1.5. NETWORK

The Red Hat Virtualization network architecture facilitates connectivity between the different elements of the Red Hat Virtualization environment. The network architecture not only supports network connectivity, it also allows for network segregation.

Figure 1.3. Network Architecture



Networking is defined in Red Hat Virtualization in several layers. The underlying physical networking infrastructure must be in place and configured to allow connectivity between the hardware and the logical components of the Red Hat Virtualization environment.

Networking Infrastructure Layer

The Red Hat Virtualization network architecture relies on some common hardware and software devices:

- Network Interface Controllers (NICs) are physical network interface devices that connect a host to the network.
- Virtual NICs (vNICs) are logical NICs that operate using the host's physical NICs. They provide network connectivity to virtual machines.
- Bonds bind multiple NICs into a single interface.
- Bridges are a packet-forwarding technique for packet-switching networks. They form the basis of virtual machine logical networks.

Logical Networks

Logical networks allow segregation of network traffic based on environment requirements. The types of logical network are:

- logical networks that carry virtual machine network traffic,

- logical networks that do not carry virtual machine network traffic,
- optional logical networks,
- and required networks.

All logical networks can either be required or optional.

A logical network that carries virtual machine network traffic is implemented at the host level as a software bridge device. By default, one logical network is defined during the installation of the Red Hat Virtualization Manager: the **ovirtmgmt** management network.

Other logical networks that can be added by an administrator are: a dedicated storage logical network, and a dedicated display logical network. Logical networks that do not carry virtual machine traffic do not have an associated bridge device on hosts. They are associated with host network interfaces directly.

Red Hat Virtualization segregates management-related network traffic from migration-related network traffic. This makes it possible to use a dedicated network (without routing) for live migration, and ensures that the management network (ovirtmgmt) does not lose its connection to hypervisors during migrations.

Explanation of logical networks on different layers

Logical networks have different implications for each layer of the virtualization environment.

Data Center Layer

Logical networks are defined at the data center level. Each data center has the **ovirtmgmt** management network by default. Further logical networks are optional but recommended. Designation as a **VM Network** and a custom MTU can be set at the data center level. A logical network that is defined for a data center must also be added to the clusters that use the logical network.

Cluster Layer

Logical networks are made available from a data center, and must be added to the clusters that will use them. Each cluster is connected to the management network by default. You can optionally add to a cluster logical networks that have been defined for the cluster's parent data center. When a required logical network has been added to a cluster, it must be implemented for each host in the cluster. Optional logical networks can be added to hosts as needed.

Host Layer

Virtual machine logical networks are implemented for each host in a cluster as a software bridge device associated with a given network interface. Non-virtual machine logical networks do not have associated bridges, and are associated with host network interfaces directly. Each host has the management network implemented as a bridge using one of its network devices as a result of being included in a Red Hat Virtualization environment. Further required logical networks that have been added to a cluster must be associated with network interfaces on each host to become operational for the cluster.

Virtual Machine Layer

Logical networks can be made available to virtual machines in the same way that a network can be made available to a physical machine. A virtual machine can have its virtual NIC connected to any virtual machine logical network that has been implemented on the host that runs it. The virtual machine then gains connectivity to any other devices or destinations that are available on the logical network it is connected to.

Example 1.1. Management Network

The management logical network, named **ovirtmgmt**, is created automatically when the Red Hat Virtualization Manager is installed. The **ovirtmgmt** network is dedicated to management traffic between the Red Hat Virtualization Manager and hosts. If no other specifically purposed bridges are set up, **ovirtmgmt** is the default bridge for all traffic.

1.6. DATA CENTERS

A data center is the highest level of abstraction in Red Hat Virtualization. A data center contains three types of information:

Storage

This includes storage types, storage domains, and connectivity information for storage domains. Storage is defined for a data center, and available to all clusters in the data center. All host clusters within a data center have access to the same storage domains.

Logical networks

This includes details such as network addresses, VLAN tags and STP support. You can define logical networks for a data center and apply them to clusters.

Clusters

Clusters are groups of hosts with compatible processor cores, either AMD or Intel processors. Clusters are migration domains; virtual machines can be live-migrated to any host within a cluster, and not to other clusters. One data center can hold multiple clusters, and each cluster can contain multiple hosts.

1.7. DATA CENTER AND CLUSTER COMPATIBILITY LEVELS

Red Hat Virtualization data centers and clusters have a compatibility version.

The data center compatibility version indicates the version of Red Hat Virtualization that the data center is intended to be compatible with. All clusters in the data center must support the desired compatibility level.

The cluster compatibility version indicates the features of Red Hat Virtualization supported by all of the hosts in the cluster. The cluster compatibility is set according to the version of the least capable host operating system in the cluster.

The table below provides a compatibility matrix of RHV versions and the required data center and cluster compatibility levels.

Table 1.1. Data center and cluster levels supported with Red Hat Virtualization

Compatibility Level	RHV Version	Description
4.7	4.4	Compatibility Level 4.7 was introduced in RHV 4.4 to support new features introduced by RHEL 8.6 hypervisors.

Compatibility Level	RHV Version	Description
4.6	4.4.6	Compatibility Level 4.6 was introduced in RHV 4.4.6 to support new features introduced by RHEL 8.4 hypervisors with Advanced Virtualization 8.4 packages.
4.5	4.4.3	Compatibility Level 4.5 was introduced in RHV 4.4.3 to support new features introduced by RHEL 8.3 hypervisors with Advanced Virtualization 8.3 packages.

Limitations

- Virtio NICs are enumerated as a different device after upgrading the cluster compatibility level to 4.6. Therefore, the NICs might need to be reconfigured. Red Hat recommends that you test the virtual machines before you upgrade the cluster by setting the cluster compatibility level to 4.6 on the virtual machine and verifying the network connection.
If the network connection for the virtual machine fails, configure the virtual machine with a custom emulated machine that matches the current emulated machine, for example `pc-q35-rhel8.3.0` for 4.5 compatibility version, before upgrading the cluster.

CHAPTER 2. STORAGE

2.1. STORAGE DOMAINS OVERVIEW

A storage domain is a collection of images that have a common storage interface. A storage domain contains complete images of templates and virtual machines (including snapshots), ISO files, and metadata about themselves. A storage domain can be made of either block devices (SAN - iSCSI or FCP) or a file system (NAS - NFS, GlusterFS, or other POSIX compliant file systems).



NOTE

GlusterFS Storage is deprecated, and will no longer be supported in future releases.

On NAS, all virtual disks, templates, and snapshots are files.

On SAN (iSCSI/FCP), each virtual disk, template or snapshot is a logical volume. Block devices are aggregated into a logical entity called a volume group, and then divided by LVM (Logical Volume Manager) into logical volumes for use as virtual hard disks. See [Red Hat Enterprise Linux Configuring and managing logical volumes](#) for more information on LVM.

Virtual disks can have one of two formats, either QCOW2 or raw. The type of storage can be either sparse or preallocated. Snapshots are always sparse but can be taken for disks of either format.

Virtual machines that share the same storage domain can be migrated between hosts that belong to the same cluster.

2.2. TYPES OF STORAGE BACKING STORAGE DOMAINS

Storage domains can be implemented using block based and file based storage.

File Based Storage

The file based storage types supported by Red Hat Virtualization are NFS, GlusterFS, other POSIX compliant file systems, and storage local to hosts.



NOTE

GlusterFS Storage is deprecated, and will no longer be supported in future releases.

File based storage is managed externally to the Red Hat Virtualization environment.

NFS storage is managed by a Red Hat Enterprise Linux NFS server, or other third party network attached storage server.

Hosts can manage their own local storage file systems.

Block Based Storage

Block storage uses unformatted block devices. Block devices are aggregated into volume groups by the Logical Volume Manager (LVM). An instance of LVM runs on all hosts, unaware of the instances running on other hosts. VDSM adds clustering logic on top of LVM by scanning volume groups for changes. When changes are detected, VDSM updates individual hosts by telling them to refresh their

volume group information. The hosts divide the volume group into logical volumes, writing logical volume metadata to disk. If more storage capacity is added to an existing storage domain, the Red Hat Virtualization Manager causes VDSM on each host to refresh volume group information. A Logical Unit Number (LUN) is an individual block device. One of the supported block storage protocols, iSCSI or Fibre Channel, is used to connect to a LUN. The Red Hat Virtualization Manager manages software iSCSI connections to the LUNs. All other block storage connections are managed externally to the Red Hat Virtualization environment. Any changes in a block based storage environment, such as the creation of logical volumes, extension or deletion of logical volumes and the addition of a new LUN are handled by LVM on a specially selected host called the Storage Pool Manager. Changes are then synced by VDSM which storage metadata refreshes across all hosts in the cluster.

2.3. STORAGE DOMAIN TYPES

Red Hat Virtualization supports the following types of storage domains, as well as the storage types that each storage domain supports.

- The **Data Domain** stores the hard disk images of all virtual machines in the Red Hat Virtualization environment. Disk images may contain an installed operating system or data stored or generated by a virtual machine. Data storage domains support NFS, iSCSI, FCP, GlusterFS and POSIX compliant storage. A data domain cannot be shared between multiple data centers.



NOTE

GlusterFS Storage is deprecated, and will no longer be supported in future releases.

- The **Export Domain** provides transitory storage for hard disk images and virtual machine templates being transferred between data centers. Additionally, export storage domains store backed up copies of virtual machines. Export storage domains support NFS storage. Multiple data centers can access a single export storage domain but only one data center can use it at a time.



NOTE

The Export domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disks, and templates can then be uploaded from the imported storage domain to the attached data center.

- The **ISO Domain** stores ISO files, also called images. ISO files are representations of physical CDs or DVDs. In the Red Hat Virtualization environment the common types of ISO files are operating system installation disks, application installation disks, and guest agent installation disks. These images can be attached to virtual machines and booted in the same way that physical disks are inserted into a disk drive and booted. ISO storage domains allow all hosts within the data center to share ISOs, eliminating the need for physical optical media.

**NOTE**

The ISO domain is a deprecated storage domain type. The ISO Uploader tool has been deprecated. Red Hat recommends uploading ISO images to the data domain with the Administration Portal or with the REST API.

2.4. STORAGE FORMATS FOR VIRTUAL DISKS

QCOW2 Formatted Virtual Machine Storage

QCOW2 is a storage format for virtual disks. QCOW stands for QEMU copy-on-write. The QCOW2 format decouples the physical storage layer from the virtual layer by adding a mapping between logical and physical blocks. Each logical block is mapped to its physical offset, which enables storage over-commitment and virtual machine snapshots, where each QCOW volume only represents changes made to an underlying virtual disk.

The initial mapping points all logical blocks to the offsets in the backing file or volume. When a virtual machine writes data to a QCOW2 volume after a snapshot, the relevant block is read from the backing volume, modified with the new information and written into a new snapshot QCOW2 volume. Then the map is updated to point to the new place.

Raw


The raw storage format has a performance advantage over QCOW2 in that no formatting is applied to virtual disks stored in the raw format. Virtual machine data operations on virtual disks stored in raw format require no additional work from hosts. When a virtual machine writes data to a given offset in its virtual disk, the I/O is written to the same offset on the backing file or logical volume.

Raw format requires that the entire space of the defined image be preallocated unless using externally managed thin provisioned LUNs from a storage array.

2.5. VIRTUAL DISK STORAGE ALLOCATION POLICIES

Preallocated Storage


All of the storage required for a virtual disk is allocated prior to virtual machine creation. If a 20 GB disk image is created for a virtual machine, the disk image uses 20 GB of storage domain capacity. Preallocated disk images cannot be enlarged. Preallocating storage can mean faster write times because no storage allocation takes place during runtime, at the cost of flexibility. Allocating storage this way reduces the capacity of the Red Hat Virtualization Manager to overcommit storage. Preallocated storage is recommended for virtual machines used for high intensity I/O tasks with less tolerance for latency in storage. Generally, server virtual machines fit this description.

**NOTE**

If thin provisioning functionality provided by your storage back-end is being used, preallocated storage should still be selected from the Administration Portal when provisioning storage for virtual machines.

Sparsely Allocated Storage

The upper size limit for a virtual disk is set at virtual machine creation time. Initially, the disk image does not use any storage domain capacity. Usage grows as the virtual machine writes data to disk, until the upper limit is reached. Capacity is not returned to the storage domain when data in the disk image is removed. Sparsely allocated storage is appropriate for virtual machines with low or medium intensity I/O tasks with some tolerance for latency in storage. Generally, desktop virtual machines fit this description.

**NOTE**

If thin provisioning functionality is provided by your storage back-end, it should be used as the preferred implementation of thin provisioning. Storage should be provisioned from the graphical user interface as preallocated, leaving thin provisioning to the back-end solution.

2.6. STORAGE METADATA VERSIONS IN RED HAT VIRTUALIZATION

Red Hat Virtualization stores information about storage domains as metadata on the storage domains themselves. Each major release of Red Hat Virtualization has seen improved implementations of storage metadata.

V1 metadata (Red Hat Virtualization 2.x series)

- Each storage domain contains metadata describing its own structure, and all of the names of physical volumes that are used to back virtual disks.
- Master domains additionally contain metadata for all the domains and physical volume names in the storage pool. The total size of this metadata is limited to 2 KB, limiting the number of storage domains that can be in a pool.
- Template and virtual machine base images are read only.
- V1 metadata is applicable to NFS, iSCSI, and FC storage domains.

V2 metadata (Red Hat Enterprise Virtualization 3.0)

- All storage domain and pool metadata is stored as logical volume tags rather than written to a logical volume. Metadata about virtual disk volumes is still stored in a logical volume on the domains.
- Physical volume names are no longer included in the metadata.
- Template and virtual machine base images are read only.
- V2 metadata is applicable to iSCSI, and FC storage domains.

V3 metadata (Red Hat Enterprise Virtualization 3.1 and later)

- All storage domain and pool metadata is stored as logical volume tags rather than written to a logical volume. Metadata about virtual disk volumes is still stored in a logical volume on the domains.
- Virtual machine and template base images are no longer read only. This change enables live snapshots, live storage migration, and clone from snapshot.
- Support for unicode metadata is added, for non-English volume names.
- V3 metadata is applicable to NFS, GlusterFS, POSIX, iSCSI, and FC storage domains.

**NOTE**

GlusterFS Storage is deprecated, and will no longer be supported in future releases.

V4 metadata (Red Hat Virtualization 4.1 and later)

- Support for QCOW2 compat levels - the QCOW image format includes a version number to allow introducing new features that change the image format so that it is incompatible with earlier versions. Newer QEMU versions (1.7 and above) support QCOW2 version 3, which is not backwards compatible, but introduces improvements such as zero clusters and improved performance.
- A new xleases volume to support VM leases - this feature adds the ability to acquire a lease per virtual machine on shared storage without attaching the lease to a virtual machine disk. A VM lease offers two important capabilities:
 - Avoiding split-brain.
 - Starting a VM on another host if the original host becomes non-responsive, which improves the availability of HA VMs.

V5 metadata (Red Hat Virtualization 4.3 and later)

- Support for 4K (4096 byte) block storage.
- Support for variable SANLOCK alignments.
- Support for new properties:
 - **BLOCK_SIZE** - stores the block size of the storage domain in bytes.
 - **ALIGNMENT** - determines the formatting and size of the xlease volume. (1MB to 8MB). Determined by the maximum number of host to be supported (value provided by the user) and disk block size.
For example: a 512b block size and support for 2000 hosts results in a 1MB xlease volume.

A 4K block size with 2000 hosts results in a 8MB xlease volume.

The default value of maximum hosts is 250, resulting in an xlease volume of 1MB for 4K disks.
- Deprecated properties:
 - The **LOGBLKSIZE**, **PHYBLKSIZE**, **MTIME**, and **POOL_UUID** fields were removed from the storage domain metadata.
 - The **SIZE** (size in blocks) field was replaced by **CAP** (size in bytes).



NOTE

- You cannot boot from a 4K format disk, as the boot disk always uses a 512 byte emulation.
- The nfs format always uses 512 bytes.

2.7. STORAGE DOMAIN AUTORECOVERY IN RED HAT VIRTUALIZATION

Hosts in a Red Hat Virtualization environment monitor storage domains in their data centers by reading metadata from each domain. A storage domain becomes inactive when all hosts in a data center report that they cannot access the storage domain.

Rather than disconnecting an inactive storage domain, the Manager assumes that the storage domain has become inactive temporarily, because of a temporary network outage for example. Once every 5 minutes, the Manager attempts to re-activate any inactive storage domains.

Administrator intervention may be required to remedy the cause of the storage connectivity interruption, but the Manager handles re-activating storage domains as connectivity is restored.

2.8. THE STORAGE POOL MANAGER

Red Hat Virtualization uses metadata to describe the internal structure of storage domains. Structural metadata is written to a segment of each storage domain. Hosts work with the storage domain metadata based on a single writer, and multiple readers configuration. Storage domain structural metadata tracks image and snapshot creation and deletion, and volume and domain extension.

The host that can make changes to the structure of the data domain is known as the Storage Pool Manager (SPM). The SPM coordinates all metadata changes in the data center, such as creating and deleting disk images, creating and merging snapshots, copying images between storage domains, creating templates and storage allocation for block devices. There is one SPM for every data center. All other hosts can only read storage domain structural metadata.

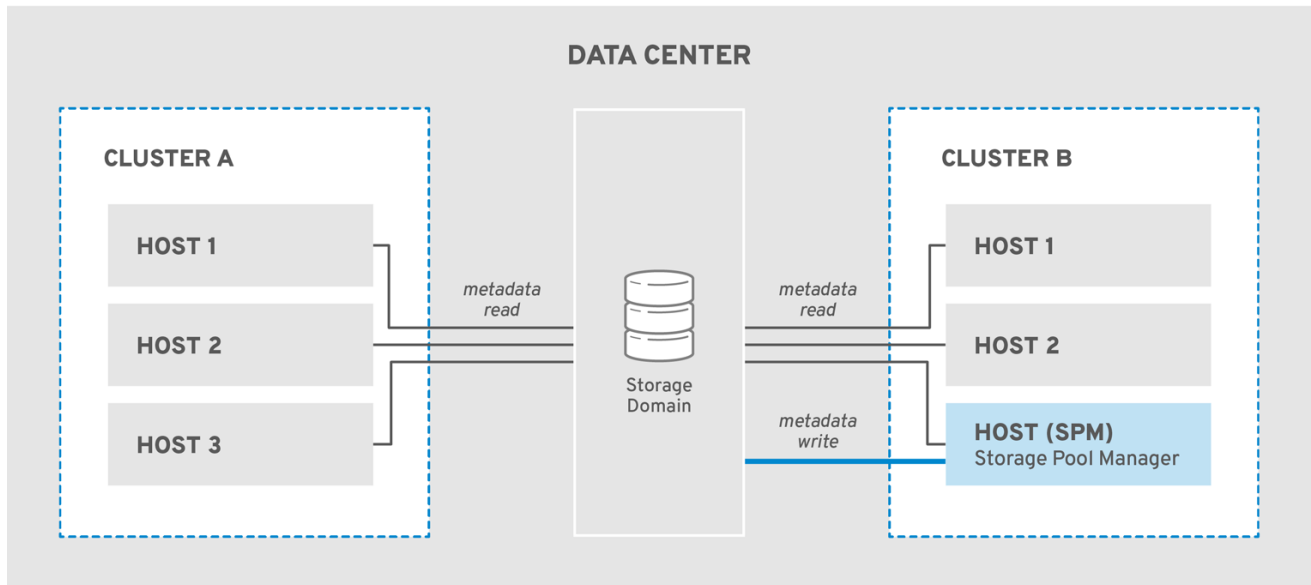
A host can be manually selected as the SPM, or it can be assigned by the Red Hat Virtualization Manager. The Manager assigns the SPM role by causing a potential SPM host to attempt to assume a storage-centric lease. The lease allows the SPM host to write storage metadata. It is storage-centric because it is written to the storage domain rather than being tracked by the Manager or hosts. Storage-centric leases are written to a special logical volume in the **master** storage domain called **leases**. Metadata about the structure of the data domain is written to a special logical volume called **metadata**. The **leases** logical volume protects the **metadata** logical volume from changes.

The Manager uses VDSM to issue the **spmStart** command to a host, causing VDSM on that host to attempt to assume the storage-centric lease. If the host is successful it becomes the SPM and retains the storage-centric lease until the Red Hat Virtualization Manager requests that a new host assume the role of SPM.

The Manager moves the SPM role to another host if:

- The SPM host can not access all storage domains, but can access the **master** storage domain
- The SPM host is unable to renew the lease because of a loss of storage connectivity or the lease volume is full and no write operation can be performed
- The SPM host crashes

Figure 2.1. The Storage Pool Manager Exclusively Writes Structural Metadata.



RHV_453539_0518

2.9. STORAGE POOL MANAGER SELECTION PROCESS

If a host has not been manually assigned the Storage Pool Manager (SPM) role, the SPM selection process is initiated and managed by the Red Hat Virtualization Manager.

First, the Red Hat Virtualization Manager requests that VDSM confirm which host has the storage-centric lease.

The Red Hat Virtualization Manager tracks the history of SPM assignment from the initial creation of a storage domain onward. The availability of the SPM role is confirmed in three ways:

- The "getSPMstatus" command: the Manager uses VDSM to check with the host that had SPM status last and receives one of "SPM", "Contending", or "Free".
- The metadata volume for a storage domain contains the last host with SPM status.
- The metadata volume for a storage domain contains the version of the last host with SPM status.

If an operational, responsive host retains the storage-centric lease, the Red Hat Virtualization Manager marks that host SPM in the administrator portal. No further action is taken.

If the SPM host does not respond, it is considered unreachable. If power management has been configured for the host, it is automatically fenced. If not, it requires manual fencing. The Storage Pool Manager role cannot be assigned to a new host until the previous Storage Pool Manager is fenced.

When the SPM role and storage-centric lease are free, the Red Hat Virtualization Manager assigns them to a randomly selected operational host in the data center.

If the SPM role assignment fails on a new host, the Red Hat Virtualization Manager adds the host to a list containing hosts the operation has failed on, marking these hosts as ineligible for the SPM role. This list is cleared at the beginning of the next SPM selection process so that all hosts are again eligible.

The Red Hat Virtualization Manager continues request that the Storage Pool Manager role and storage-centric lease be assumed by a randomly selected host that is not on the list of failed hosts until the SPM selection succeeds.

Each time the current SPM is unresponsive or unable to fulfill its responsibilities, the Red Hat Virtualization Manager initiates the Storage Pool Manager selection process.

2.10. EXCLUSIVE RESOURCES AND SANLOCK IN RED HAT VIRTUALIZATION

Certain resources in the Red Hat Virtualization environment must be accessed exclusively.

The SPM role is one such resource. If more than one host were to become the SPM, there would be a risk of data corruption as the same data could be changed from two places at once.

Prior to Red Hat Enterprise Virtualization 3.1, SPM exclusivity was maintained and tracked using a VDSM feature called **safelease**. The lease was written to a special area on all of the storage domains in a data center. All of the hosts in an environment could track SPM status in a network-independent way. The VDSM's safe lease only maintained exclusivity of one resource: the SPM role.

Sanlock provides the same functionality, but treats the SPM role as one of the resources that can be locked. Sanlock is more flexible because it allows additional resources to be locked.

Applications that require resource locking can register with Sanlock. Registered applications can then request that Sanlock lock a resource on their behalf, so that no other application can access it. For example, instead of VDSM locking the SPM status, VDSM now requests that Sanlock do so.

Locks are tracked on disk in a **lockspace**. There is one lockspace for every storage domain. In the case of the lock on the SPM resource, each host's liveness is tracked in the lockspace by the host's ability to renew the `hostid` it received from the Manager when it connected to storage, and to write a timestamp to the lockspace at a regular interval. The `ids` logical volume tracks the unique identifiers of each host, and is updated every time a host renews its `hostid`. The SPM resource can only be held by a live host.

Resources are tracked on disk in the **leases** logical volume. A resource is said to be **taken** when its representation on disk has been updated with the unique identifier of the process that has taken it. In the case of the SPM role, the SPM resource is updated with the `hostid` that has taken it.

The Sanlock process on each host only needs to check the resources once to see that they are taken. After an initial check, Sanlock can monitor the lockspaces until timestamp of the host with a locked resource becomes stale.

Sanlock monitors the applications that use resources. For example, VDSM is monitored for SPM status and `hostid`. If the host is unable to renew its `hostid` from the Manager, it loses exclusivity on all resources in the lockspace. Sanlock updates the resource to show that it is no longer taken.

If the SPM host is unable to write a timestamp to the lockspace on the storage domain for a given amount of time, the host's instance of Sanlock requests that the VDSM process release its resources. If the VDSM process responds, its resources are released, and the SPM resource in the lockspace can be taken by another host.

If VDSM on the SPM host does not respond to requests to release resources, Sanlock on the host kills the VDSM process. If the kill command is unsuccessful, Sanlock escalates by attempting to kill VDSM using `sigkill`. If the `sigkill` is unsuccessful, Sanlock depends on the **watchdog daemon** to reboot the host.


Every time VDSM on the host renews its `hostid` and writes a timestamp to the lockspace, the watchdog daemon receives a **pet**. When VDSM is unable to do so, the watchdog daemon is no longer being petted.

After the watchdog daemon has not received a pet for a given amount of time, it reboots the host. This final level of escalation, if reached, guarantees that the SPM resource is released, and can be taken by another host.

2.11. THIN PROVISIONING AND STORAGE OVER-COMMITMENT

The Red Hat Virtualization Manager provides provisioning policies to optimize storage usage within the virtualization environment. A thin provisioning policy allows you to over-commit storage resources, provisioning storage based on the actual storage usage of your virtualization environment.

Storage over-commitment is the allocation of more storage to virtual machines than is physically available in the storage pool. Generally, virtual machines use less storage than what has been allocated to them. Thin provisioning allows a virtual machine to operate as if the storage defined for it has been completely allocated, when in fact only a fraction of the storage has been allocated.



NOTE

While the Red Hat Virtualization Manager provides its own thin provisioning function, you should use the thin provisioning functionality of your storage back-end if it provides one.

To support storage over-commitment, VDSM defines a threshold which compares logical storage allocation with actual storage usage. This threshold is used to make sure that the data written to a disk image is smaller than the logical volume that backs the disk image. QEMU identifies the highest offset written to in a logical volume, which indicates the point of greatest storage use. VDSM monitors the highest offset marked by QEMU to ensure that the usage does not cross the defined threshold. So long as VDSM continues to indicate that the highest offset remains below the threshold, the Red Hat Virtualization Manager knows that the logical volume in question has sufficient storage to continue operations.

When QEMU indicates that usage has risen to exceed the threshold limit, VDSM communicates to the Manager that the disk image will soon reach the size of its logical volume. The Red Hat Virtualization Manager requests that the SPM host extend the logical volume. This process can be repeated as long as the data storage domain for the data center has available space. When the data storage domain runs out of available free space, you must manually add storage capacity to expand it.

2.12. LOGICAL VOLUME EXTENSION

The Red Hat Virtualization Manager uses thin provisioning to overcommit the storage available in a storage pool, and allocates more storage than is physically available. Virtual machines write data as they operate. A virtual machine with a thinly-provisioned disk image will eventually write more data than the logical volume backing its disk image can hold. When this happens, logical volume extension is used to provide additional storage and facilitate the continued operations for the virtual machine.

Red Hat Virtualization provides a thin provisioning mechanism over LVM. When using QCOW2 formatted storage, Red Hat Virtualization relies on the host system process `qemu-kvm` to map storage blocks on disk to logical blocks in a sequential manner. This allows, for example, the definition of a logical 100 GB disk backed by a 1 GB logical volume. When `qemu-kvm` crosses a usage threshold set by VDSM, the local VDSM instance makes a request to the SPM for the logical volume to be extended by another one gigabyte. VDSM on the host running a virtual machine in need of volume extension notifies the SPM VDSM that more space is required. The SPM extends the logical volume and the SPM VDSM instance causes the host VDSM to refresh volume group information and recognize that the extend operation is complete. The host can continue operations.

Logical Volume extension does not require that a host know which other host is the SPM; it could even

be the SPM itself. The storage extension communication is done via a storage mailbox. The storage mailbox is a dedicated logical volume on the data storage domain. A host that needs the SPM to extend a logical volume writes a message in an area designated to that particular host in the storage mailbox. The SPM periodically reads the incoming mail, performs requested logical volume extensions, and writes a reply in the outgoing mail. After sending the request, a host monitors its incoming mail for responses every two seconds. When the host receives a successful reply to its logical volume extension request, it refreshes the logical volume map in device mapper to recognize the newly allocated storage.

When the physical storage available to a storage pool is nearly exhausted, multiple images can run out of usable storage with no means to replenish their resources. A storage pool that exhausts its storage causes QEMU to return an **enospc error**, which indicates that the device no longer has any storage available. At this point, running virtual machines are automatically paused and manual intervention is required to add a new LUN to the volume group.

When a new LUN is added to the volume group, the Storage Pool Manager automatically distributes the additional storage to logical volumes that need it. The automatic allocation of additional resources allows the relevant virtual machines to automatically continue operations uninterrupted or resume operations if stopped.

2.13. THE EFFECT OF STORAGE DOMAIN ACTIONS ON STORAGE CAPACITY

Power on, power off, and reboot a stateless virtual machine

These three processes affect the copy-on-write (COW) layer in a stateless virtual machine. For more information, see the **Stateless** row of the [Virtual Machine General Settings table](#) in the *Virtual Machine Management Guide*.

Create a storage domain

Creating a block storage domain results in files with the same names as the seven LVs shown below, and initially should take less capacity.

ids	64f87b0f-88d6-49e9-b797-60d36c9df497 -wi-ao----	128.00m
inbox	64f87b0f-88d6-49e9-b797-60d36c9df497 -wi-a-----	128.00m
leases	64f87b0f-88d6-49e9-b797-60d36c9df497 -wi-a-----	2.00g
master	64f87b0f-88d6-49e9-b797-60d36c9df497 -wi-ao----	1.00g
metadata	64f87b0f-88d6-49e9-b797-60d36c9df497 -wi-a-----	512.00m
outbox	64f87b0f-88d6-49e9-b797-60d36c9df497 -wi-a-----	128.00m
xleases	64f87b0f-88d6-49e9-b797-60d36c9df497 -wi-a-----	1.00g

Delete a storage domain

Deleting a storage domain frees up capacity on the disk by the same of amount of capacity the process deleted.

Migrate a storage domain

Migrating a storage domain does not use additional storage capacity. For more information about migrating storage domains, see [Migrating Storage Domains Between Data Centers in the Same Environment](#) in the *Administration Guide*.

Move a virtual disk to other storage domain

Migrating a virtual disk requires enough free space to be available on the target storage domain. You can see the target domain's approximate free space in the Administration Portal.

The storage types in the move process affect the visible capacity. For example, if you move a preallocated disk from block storage to file storage, the resulting free space may be considerably smaller than the initial free space.

Live migrating a virtual disk to another storage domain also creates a snapshot, which is automatically merged after the migration is complete. To learn more about moving virtual disks, see [Moving a Virtual Disk](#) in the *Administration Guide*.

Pause a storage domain

Pausing a storage domain does not use any additional storage capacity.

Create a snapshot of a virtual machine

Creating a snapshot of a virtual machine can affect the storage domain capacity.

- Creating a live snapshot uses memory snapshots by default and generates two additional volumes per virtual machine. The first volume is the sum of the memory, video memory, and 200 MB of buffer. The second volume contains the virtual machine configuration, which is several MB in size. When using block storage, rounding up occurs to the nearest unit Red Hat Virtualization can provide.
- Creating an offline snapshot initially consumes 1 GB of block storage and is dynamic up to the size of the disk.
- Cloning a snapshot creates a new disk the same size as the original disk.
- Committing a snapshot removes all child volumes, depending on where in the chain the commit occurs.
- Deleting a snapshot eventually removes the child volume for each disk and is only supported with a running virtual machine.
- Previewing a snapshot creates a temporary volume per disk, so sufficient capacity must be available to allow the creation of the preview.
- Undoing a snapshot preview removes the temporary volume created by the preview.

Attach and remove direct LUNs

Attaching and removing direct LUNs does not affect the storage domain since they are not a storage domain component. For more information, see [Overview of Live Storage Migration](#) in the *Administration Guide*.

CHAPTER 3. NETWORKING

3.1. HOST NETWORKING

On the data link layer (layer 2), RHV enables the configuration of Linux bonds to connect to VLANs and define the MTU for network interfaces. These networks can be shared via Linux bridges to virtual machines.

For SR-IOV, you can configure the number of virtual functions and their mapping to logical networks.

FCoE manages its own VLANs. These FCoE managed VLANs are used exclusively for storage access. They are invisible to the Manager and any virtual machines.

iSCSI manages iSCSI bonds. They are not part of RHV's visible host network configuration. You can use iSCSI without iSCSI bonds, which are useful only to improve the reliability of iSCSI storage.



IMPORTANT

All hosts in a cluster must use either IPv4 or IPv6 as the IP stack for their management network. Dual stack is not supported.

You can configure the DNS resolver that the host uses.

It is also possible to manage network roles and QoS.

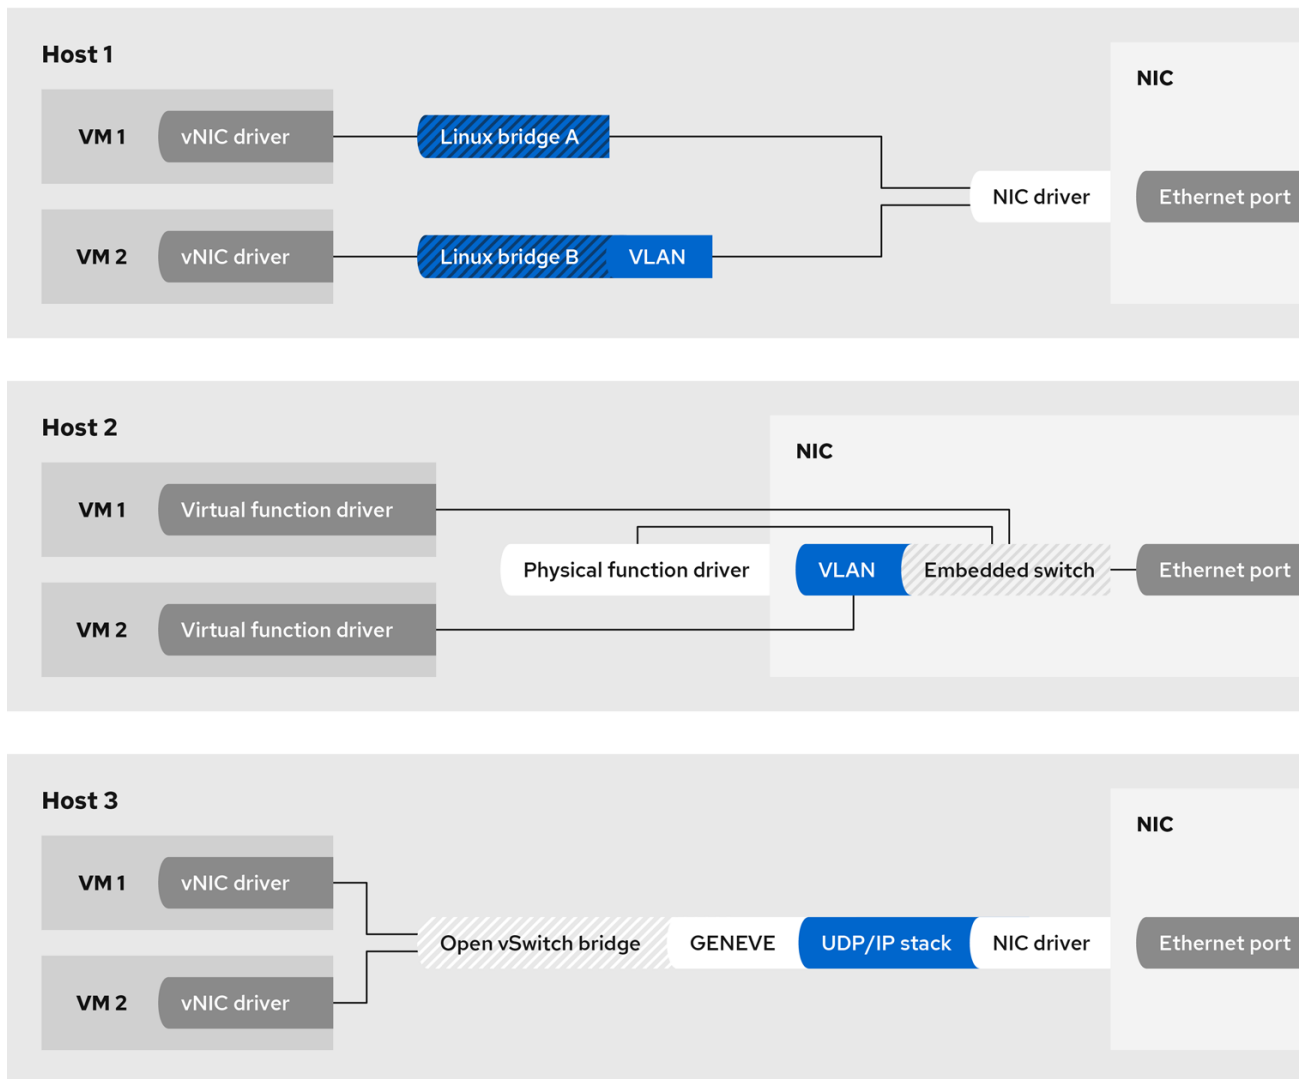
3.2. VIRTUAL MACHINE NETWORKING TYPES

In RHV, virtual NICs of virtual machines can connect to the following types of networks:

- Linux bridges
- SR-IOV NICs
- RHV's internal OVN

The following diagram shows the structure of these three approaches, where:

- **Host 1** represents Linux bridges
- **Host 2** represents SR-IOV NICs
- **Host 3** represents OVN



Configured by RHV host network configuration

103_RHV_0720

Table 3.1. Comparison of network types

	Linux bridge	SR-IOV	RHV internal OVN
Isolation from physical host networks	Layer 3, Separate IP network possible	Layer 2, Separate VLANs possible	Isolated
Live Migration	x	x	x
QoS	x		
Port Mirroring	x		
configuration of plugged vNIC	x	x	
MAC address management	x	x	x

	Linux bridge	SR-IOV	RHV internal OVN
MTU propagation	x		x
VLAN filtering, might require configuration on the physical switch	x	x	Technology Preview
MAC Spoofing Protection	x		x
IP Spoofing Protection	x		x
Predefined Network Filters	x		
Custom Layer 3/4 Filtering			x
NAT			
DHCP/Router Advertisements			x
Layer 3 Router			x
Performance	**	***	*
virtual machine network data encapsulation	flat, VLAN	flat, VLAN	Stable: GENEVE; Technology Preview: flat, VLAN

Networking choices for various scenarios

Linux bridge is the default, and the most proven option. It fits most use cases.

For scenarios that require very low network latency or a large number of Ethernet frames, consider investing in SR-IOV. Keep in mind, however, that SR-IOV requires hardware support and additional configuration steps.

RHV's internal OVN networks enable virtual machines to communicate with each other without any manual network configuration.

The Manager provides only a subset of software-defined networking (SDN) features and user interfaces. To use all SDN features, similar to RHV's internal OVN or a third-party SDN, you need to use an additional client, such as CloudForms.

You can combine all network types on a single host and connect them to the same virtual machine.

3.2.1. Interaction with guest operating system

RHV supports the initial configuration of a virtual machine by providing configuration data via cloud-init. If the qemu-guest-agent runs inside the virtual machine, RHV can report the IP addresses of the virtual machine.

If the virtual machine uses a VirtIO NIC, the MTU of the RHV logical networks are provided to the guest operating system. The guest operating system can pick up the MTU from DHCPv4 or IPv6 router advertisements if the logical network supports these advertisements.

3.2.2. Host and virtual machine networking

Linux bridge networking separates virtual machine and host networking on OSI layer 3. Therefore the networking configuration, including VLAN, bonding, and MTU, is shared between the host and its virtual machines.

To reduce their surface, hosts should not assign IP addresses to VLANs that are connected to virtual machines. By not assigning IP addresses, the hosts can avoid potential confusion caused by virtual machine traffic.

The IP address associated with the Linux bridge is not required to be within the same subnet as the virtual machines that use the bridge for connectivity. If the bridge is assigned an IP address on the same subnet as the virtual machines that use it, the host is addressable within the logical network by virtual machines. As a rule, it is not recommended to run network exposed services on a virtualization host.

3.3. NETWORK ARCHITECTURE

Networking in Red Hat Virtualization includes basic networking, networking within a cluster, and host networking configurations.

Basic networking

The basic hardware and software elements that facilitate networking.

Networking within a cluster

Network interactions among cluster objects such as hosts, logical networks and virtual machines.

Host networking configurations

Supported configurations for networking within a host.

A well designed and built network ensures that high bandwidth tasks receive adequate bandwidth, that latency does not impact user interactions, and that virtual machines can be successfully migrated within a migration domain. A poorly built network can cause unacceptable latency, and migration and cloning failures that result from network flooding.

An alternative method of managing your network is by integrating with Cisco Application Centric Infrastructure (ACI), by configuring Red Hat Virtualization on Cisco's Application Policy Infrastructure Controller (APIC) version 3.1(1) and later according to [Cisco's documentation](#). On the Red Hat Virtualization side, all that is required is connecting the hosts' NICs to the network and the virtual machines' vNICs to the required network. The remaining configuration tasks are managed by Cisco ACI.

3.4. BASIC NETWORKING TERMS

Red Hat Virtualization provides networking functionality between virtual machines, virtualization hosts, and wider networks using:

- Logical networks

- A network interface controller (NIC)
- A Linux bridge
- A Bond
- A virtual network interface controller (vNIC)
- A virtual LAN (VLAN)

NICs, Linux bridges, and vNICs enable network communication between hosts, virtual machines, local area networks, and the Internet. Bonds and VLANs are optionally implemented to enhance security, fault tolerance, and network capacity.

3.5. NETWORK INTERFACE CONTROLLER

The network interface controller (NIC) is a network adapter or LAN adapter that connects a computer to a computer network. The NIC operates on both the physical and data link layers of the machine and enables network connectivity. All virtualization hosts in a Red Hat Virtualization environment have at least one NIC, though it is more common for a host to have two or more NICs.

One physical NIC can have multiple virtual NICs (vNICs) logically connected to it. A virtual NIC acts as a network interface for a virtual machine. To distinguish between a vNIC and the NIC that supports it, the Red Hat Virtualization Manager assigns each vNIC a unique MAC address.

3.6. LINUX BRIDGE


A Linux bridge is a software device that uses packet forwarding in a packet-switched network. Bridging allows multiple network interface devices to share the connectivity of one NIC and appear on a network as separate physical devices. The bridge examines a packet's source addresses to determine relevant target addresses. Once the target address is determined, the bridge adds the location to a table for future reference. This allows a host to redirect network traffic to virtual machine associated vNICs that are members of a bridge.

Custom properties can be defined for both the bridge and the Ethernet connection. VDSM passes the network definition and custom properties to the setup network hook script.

3.7. BONDS

A bond is a collection of multiple network interface cards into a single software-defined device. Because bonded network interfaces combine the transmission capability of the network interface cards included in the bond to act as a single network interface, they can provide greater transmission speed than that of a single network interface card. Also, because all network interface cards in the bond must fail for the bond itself to fail, bonding provides increased fault tolerance. However, one limitation is that the network interface cards that form a bonded network interface must be of the same make and model to ensure that all network interface cards in the bond support the same options and modes.

The packet dispersal algorithm for a bond is determined by the bonding mode used.



IMPORTANT

Modes 1, 2, 3 and 4 support both virtual machine (bridged) and non-virtual machine (bridgeless) network types. Modes 0, 5 and 6 support non-virtual machine (bridgeless) networks only.

3.8. BONDING MODES

Red Hat Virtualization uses Mode 4 by default, but supports the following common bonding modes:

Mode 0 (round-robin policy)

Transmits packets through network interface cards in sequential order. Packets are transmitted in a loop that begins with the first available network interface card in the bond and ends with the last available network interface card in the bond. All subsequent loops then start with the first available network interface card. Mode 0 offers fault tolerance and balances the load across all network interface cards in the bond. However, Mode 0 cannot be used in conjunction with bridges, and is therefore not compatible with virtual machine logical networks.

Mode 1 (active-backup policy)

Sets all network interface cards to a backup state while one network interface card remains active. In the event of failure in the active network interface card, one of the backup network interface cards replaces that network interface card as the only active network interface card in the bond. The MAC address of the bond in Mode 1 is visible on only one port to prevent any confusion that might otherwise be caused if the MAC address of the bond changed to reflect that of the active network interface card. Mode 1 provides fault tolerance and is supported in Red Hat Virtualization.

Mode 2 (XOR policy)

Selects the network interface card through which to transmit packets based on the result of an XOR operation on the source and destination MAC addresses modulo network interface card **slave** count. This calculation ensures that the same network interface card is selected for each destination MAC address used. Mode 2 provides fault tolerance and load balancing and is supported in Red Hat Virtualization.

Mode 3 (broadcast policy)

Transmits all packets to all network interface cards. Mode 3 provides fault tolerance and is supported in Red Hat Virtualization.

Mode 4 (IEEE 802.3ad policy)

Creates aggregation groups in which the interfaces share the same speed and duplex settings. Mode 4 uses all network interface cards in the active aggregation group in accordance with the IEEE 802.3ad specification and is supported in Red Hat Virtualization.

Mode 5 (adaptive transmit load balancing policy)

Ensures the distribution of outgoing traffic accounts for the load on each network interface card in the bond and that the current network interface card receives all incoming traffic. If the network interface card assigned to receive traffic fails, another network interface card is assigned to the role of receiving incoming traffic. Mode 5 cannot be used in conjunction with bridges, therefore it is not compatible with virtual machine logical networks.

Mode 6 (adaptive load balancing policy)

Combines Mode 5 (adaptive transmit load balancing policy) with receive load balancing for IPv4 traffic without any special switch requirements. ARP negotiation is used for balancing the receive load. Mode 6 cannot be used in conjunction with bridges, therefore it is not compatible with virtual machine logical networks.

3.9. SWITCH CONFIGURATION FOR BONDING

Switch configurations vary per the requirements of your hardware. Refer to the deployment and networking configuration guides for your operating system.



IMPORTANT

For every type of switch it is important to set up the switch bonding with the Link Aggregation Control Protocol (LACP) protocol and **not** the Cisco Port Aggregation Protocol (PAgP) protocol.

3.10. VIRTUAL NETWORK INTERFACE CARDS

Virtual network interface cards (vNICs) are virtual network interfaces that are based on the physical NICs of a host. Each host can have multiple NICs, and each NIC can be a base for multiple vNICs.

When you attach a vNIC to a virtual machine, the Red Hat Virtualization Manager creates several associations between the virtual machine to which the vNIC is being attached, the vNIC itself, and the physical host NIC on which the vNIC is based. Specifically, when a vNIC is attached to a virtual machine, a new vNIC and MAC address are created on the physical host NIC on which the vNIC is based. Then, the first time the virtual machine starts after that vNIC is attached, **libvirt** assigns the vNIC a PCI address. The MAC address and PCI address are then used to obtain the name of the vNIC (for example, **eth0**) in the virtual machine.

The process for assigning MAC addresses and associating those MAC addresses with PCI addresses is slightly different when creating virtual machines based on templates or snapshots:

- If PCI addresses have already been created for a template or snapshot, the vNICs on virtual machines created based on that template or snapshot are ordered in accordance with those PCI addresses. MAC addresses are then allocated to the vNICs in that order.
- If PCI addresses have not already been created for a template, the vNICs on virtual machines created based on that template are ordered alphabetically. MAC addresses are then allocated to the vNICs in that order.
- If PCI addresses have not already been created for a snapshot, the Red Hat Virtualization Manager allocates new MAC addresses to the vNICs on virtual machines based on that snapshot.

Once created, vNICs are added to a network bridge device. The network bridge devices connect virtual machines to virtual logical networks.

Running the **ip addr show** command on a virtualization host shows all of the vNICs that are associated with virtual machines on that host. Also visible are any network bridges that have been created to back logical networks, and any NICs used by the host.

```
[root@rhev-host-01 ~]# ip addr show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 00:21:86:a2:85:cd brd ff:ff:ff:ff:ff:ff
    inet6 fe80::221:86ff:fea2:85cd/64 scope link
        valid_lft forever preferred_lft forever
3: wlan0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen
1000
    link/ether 00:21:6b:cc:14:6c brd ff:ff:ff:ff:ff:ff
5: ;vdsmdummy:: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN
    link/ether 4a:d5:52:c2:7f:4b brd ff:ff:ff:ff:ff:ff
```

```

6: bond0: <BROADCAST,MULTICAST,MASTER> mtu 1500 qdisc noop state DOWN
  link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
7: bond4: <BROADCAST,MULTICAST,MASTER> mtu 1500 qdisc noop state DOWN
  link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
8: bond1: <BROADCAST,MULTICAST,MASTER> mtu 1500 qdisc noop state DOWN
  link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
9: bond2: <BROADCAST,MULTICAST,MASTER> mtu 1500 qdisc noop state DOWN
  link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
10: bond3: <BROADCAST,MULTICAST,MASTER> mtu 1500 qdisc noop state DOWN
  link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
11: ovirtmgmt: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state
UNKNOWN
  link/ether 00:21:86:a2:85:cd brd ff:ff:ff:ff:ff:ff
  inet 10.64.32.134/23 brd 10.64.33.255 scope global ovirtmgmt
  inet6 fe80::221:86ff:fea2:85cd/64 scope link
    valid_lft forever preferred_lft forever

```

The console output from the command shows several devices: one loop back device (**lo**), one Ethernet device (**eth0**), one wireless device (**wlan0**), one VDSM dummy device (**vdsmdummy**), five bond devices (**bond0**, **bond4**, **bond1**, **bond2**, **bond3**), and one network bridge (**ovirtmgmt**).

vNICs are all members of a network bridge device and logical network. Bridge membership can be displayed using the **brctl show** command:

```

[root@rhev-host-01 ~]# brctl show
bridge name bridge id STP enabled interfaces
ovirtmgmt 8000.e41f13b7fdd4 no vnet002
  vnet001
  vnet000
  eth0

```

The console output from the **brctl show** command shows that the virtio vNICs are members of the **ovirtmgmt** bridge. All of the virtual machines that the vNICs are associated with are connected to the **ovirtmgmt** logical network. The **eth0** NIC is also a member of the **ovirtmgmt** bridge. The **eth0** device is cabled to a switch that provides connectivity beyond the host.

3.11. VIRTUAL LAN (VLAN)

A VLAN (Virtual LAN) is an attribute that can be applied to network packets. Network packets can be "tagged" into a numbered VLAN. A VLAN is a security feature used to isolate network traffic at the switch level. VLANs are separate and mutually exclusive. The Red Hat Virtualization Manager is VLAN-aware and able to tag and redirect VLAN traffic, however VLAN implementation requires a switch that supports VLANs.

At the switch level, ports are assigned a VLAN designation. A switch applies a VLAN tag to traffic originating from a particular port, marking the traffic as part of a VLAN, and ensures that responses carry the same VLAN tag. A VLAN can extend across multiple switches. VLAN tagged network traffic on a switch is undetectable except by machines connected to a port designated with the correct VLAN. A given port can be tagged into multiple VLANs, which allows traffic from multiple VLANs to be sent to a single port, to be deciphered using software on the machine that receives the traffic.

3.12. NETWORK LABELS

You can use network labels to simplify several administrative tasks associated with creating and administering logical networks and associating those logical networks with physical host network interfaces and bonds.

A network label is a plain text, human readable label that you can attach to a logical network or a physical host network interface. Follow these rules when creating a label:

- There is no limit on the length of a label.
- You must use a combination of lowercase and uppercase letters, underscores and hyphens.
- You cannot use spaces or special characters.

Attaching a label to a logical network or physical host network interface creates an association with other logical networks or physical host network interfaces to which the same label has been attached:

Network Label Associations

- When you attach a label to a logical network, that logical network will be automatically associated with any physical host network interfaces with the given label.
- When you attach a label to a physical host network interface, any logical networks with the given label will be automatically associated with that physical host network interface.
- Changing the label attached to a logical network or physical host network interface is the same as removing a label and adding a new label. The association between related logical networks or physical host network interfaces is updated.

Network Labels and Clusters

- When a labeled logical network is added to a cluster and there is a physical host network interface in that cluster with the same label, the logical network is automatically added to that physical host network interface.
- When a labeled logical network is detached from a cluster and there is a physical host network interface in that cluster with the same label, the logical network is automatically detached from that physical host network interface.

Network Labels and Logical Networks With Roles

- When a labeled logical network is assigned to act as a display network or migration network, that logical network is then configured on the physical host network interface using DHCP so that the logical network can be assigned an IP address.
Setting a label on a role network (for instance, "a migration network" or "a display network") causes a mass deployment of that network on all hosts. Such mass additions of networks are achieved through the use of DHCP. This method of mass deployment was chosen over a method of typing in static addresses, because of the unscalable nature of the task of typing in many static IP addresses.

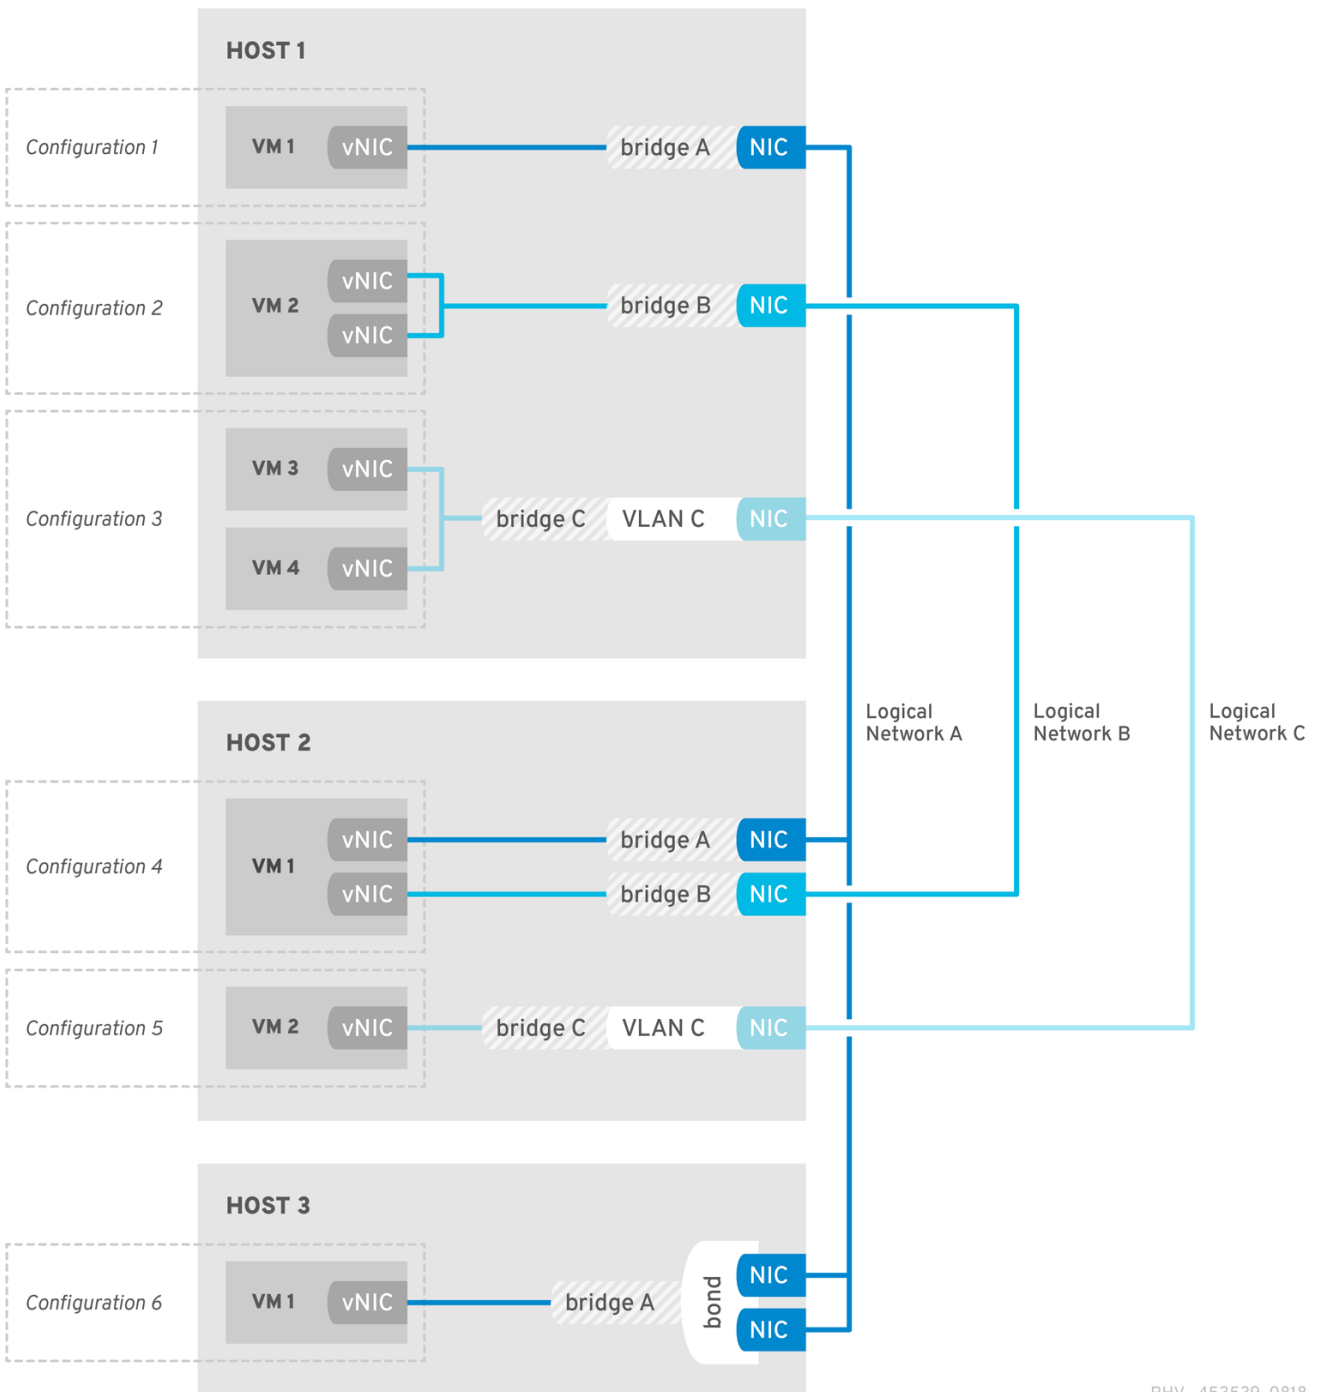
3.13. CLUSTER NETWORKING

Cluster level networking objects include:

- Clusters
- Logical Networks

A data center is a logical grouping of multiple clusters and each cluster is a logical group of multiple hosts. The following diagram depicts the contents of a single cluster.

Figure 3.1. Networking within a cluster



RHV_453539_0818

Hosts in a cluster all have access to the same storage domains. Hosts in a cluster also have logical networks applied to the cluster. For a virtual machine logical network to become operational for use with virtual machines, the network must be defined and implemented for each host in the cluster using the Red Hat Virtualization Manager. Other logical network types can be implemented on only the hosts that use them.

Multi-host network configuration automatically applies any updated network settings to all of the hosts within the data center to which the network is assigned.

3.14. LOGICAL NETWORKS

Logical networking enables the Red Hat Virtualization environment to separate network traffic by type. For example, the **ovirtmgmt** network is created by default during the installation of Red Hat Virtualization to be used for management communication between the Manager and hosts. A typical use for logical networks is to group network traffic with similar requirements and usage together. In many cases, a storage network and a display network are created by an administrator to isolate traffic of each respective type for optimization and troubleshooting.

Red Hat Virtualization supports the following logical network types:

- Logical networks that carry only host network traffic, such as storage or migration traffic
- Logical networks that carry host and virtual machine network traffic
- Logical networks that carry only virtual machine network traffic, such as OVN networks

Logical networks are defined at the data center level.

If necessary, the Red Hat Virtualization Manager automatically instantiates logical networks on the host, depending on the type of virtual machine network. For more information, see [virtual machine networking types](#).

Example 3.1. Example usage of a logical network.

A system administrator wants to use a logical network to test a web server.

There are two hosts called Red Host and White Host in a cluster called Pink Cluster in a data center called Purple Data Center. Both Red Host and White Host have been using the default logical network, **ovirtmgmt**, for all networking functions. The system administrator responsible for Pink Cluster decides to isolate network testing for a web server by placing the web server and some client virtual machines on a separate logical network. She decides to call the new logical network **test_logical_network**.

1. She creates a new logical network, named **test_logical_network**, for the Purple Data Center with VLAN tagging enabled. VLAN tagging is necessary when you have two logical networks connected to the same physical NIC. She applies **test_logical_network** to the Pink Cluster.
2. In Red Host, she attaches **test_logical_network** to a physical NIC that will be included in the bridge that RHV creates. The network is non-operational until she sets up the corresponding bridge in all hosts in the cluster by adding a physical network interface on each host in the Pink cluster to **test_logical_network**. She repeats this step for White Host. When both White Host and Red Host have the **test_logical_network** logical network bridged to a physical network interface, the **test_logical_network** becomes operational and is ready to be used by virtual machines.
3. She associates the virtual machines on the Red Host and White Host with the new network.

3.15. REQUIRED NETWORKS, OPTIONAL NETWORKS, AND VIRTUAL MACHINE NETWORKS

A required network is a logical network that must be available to all hosts in a cluster. When a host's required network becomes non-operational, virtual machines running on that host are migrated to another host; the extent of this migration is dependent upon the chosen scheduling policy. This is

beneficial if you have virtual machines running mission critical workloads.

An optional network is a logical network that has not been explicitly declared as **Required**. Optional networks can be implemented on only the hosts that use them. The presence or absence of optional networks does not affect the **Operational** status of a host. When a non-required network becomes non-operational, the virtual machines running on the network are not migrated to another host. This prevents unnecessary I/O overload caused by mass migrations. Note that when a logical network is created and added to clusters, the **Required** box is selected by default.

To change a network's **Required** designation, from the Administration Portal, select a network, click the **Cluster** tab, and click the **Manage Networks** button.

A virtual machine network, called a **VM network** in the user interface, is a logical network designated to carry only virtual machine network traffic. A virtual machine network can be required or optional. Virtual machines that use an optional virtual machine network only start on hosts with that network.

3.16. PORT MIRRORING

Port mirroring copies layer 3 network traffic on a given logical network and host to a virtual interface on a virtual machine. This virtual machine can be used for network debugging and tuning, intrusion detection, and monitoring the behavior of other virtual machines on the same host and logical network.

The only traffic copied is internal to one logical network on one host. There is no increase in traffic on the network external to the host. However, a virtual machine with port mirroring enabled uses more host CPU and RAM than other virtual machines.

Port mirroring is enabled or disabled in the vNIC profiles of logical networks, and has the following limitations:

- Hot linking vNICs with a profile that has port mirroring enabled is not supported.
- Port mirroring cannot be altered when the vNIC profile is attached to a virtual machine.

Given the above limitations, it is recommended that you enable port mirroring on an additional, dedicated vNIC profile.



IMPORTANT

Enabling port mirroring reduces the privacy of other network users.


3.17. HOST NETWORKING CONFIGURATIONS

[Cluster Networking](#) can be helpful to understand these networking configurations.

Common types of networking configurations for virtualization hosts include:

- **Bridge and NIC configuration.**
This configuration uses a Linux bridge to connect one or more virtual machines to the host's NIC.

An example of this configuration is the automatic creation of the **ovirtmgmt** network when installing the Red Hat Virtualization Manager. Then, during host installation, the Red Hat Virtualization Manager installs **VDSM** on the host. The **VDSM** installation process creates the **ovirtmgmt** bridge which obtains the host's IP address to enable communication with the Manager.



IMPORTANT

All hosts in a cluster must use either IPv4 or IPv6 as the IP stack for their management network. Dual stack is not supported.

- **Bridge, VLAN, and NIC configuration.**

A VLAN can be included in the bridge and NIC configuration to provide a secure channel for data transfer over the network and supports connecting multiple bridges to a single NIC using multiple VLANs.

- **Bridge, Bond, and VLAN configuration.**

A bond creates a logical link that combines the two (or more) physical Ethernet links. The resultant benefits include NIC fault tolerance and potential bandwidth extension, depending on the bonding mode.

- **Multiple Bridge, Multiple VLAN, and NIC configuration.**

This configuration connects a NIC to multiple VLANs.

For example, to connect a single NIC to two VLANs, the network switch can be configured to pass network traffic that has been tagged into one of the two VLANs to one NIC on the host. The host uses two vNICs to separate VLAN traffic, one for each VLAN. Traffic tagged into either VLAN then connects to a separate bridge by having the appropriate vNIC as a bridge member. Each bridge, in turn, connects to multiple virtual machines.



NOTE

You can also bond multiple NICs to facilitate a connection with multiple VLANs. Each VLAN in this configuration is defined over the bond comprising the multiple NICs. Each VLAN connects to an individual bridge and each bridge connects to one or more guests.

CHAPTER 4. POWER MANAGEMENT

4.1. INTRODUCTION TO POWER MANAGEMENT AND FENCING

The Red Hat Virtualization environment is most flexible and resilient when power management and fencing have been configured. Power management allows the Red Hat Virtualization Manager to control host power cycle operations, most importantly to reboot hosts on which problems have been detected. Fencing is used to isolate problem hosts from a functional Red Hat Virtualization environment by rebooting them, in order to prevent performance degradation. Fenced hosts can then be returned to responsive status through administrator action and be reintegrated into the environment.

Power management and fencing make use of special dedicated hardware in order to restart hosts independently of host operating systems. The Red Hat Virtualization Manager connects to a power management devices using a network IP address or hostname. In the context of Red Hat Virtualization, a power management device and a fencing device are the same thing.

4.2. POWER MANAGEMENT BY PROXY IN RED HAT VIRTUALIZATION

The Red Hat Virtualization Manager does not communicate directly with fence agents. Instead, the Manager uses a proxy to send power management commands to a host power management device. The Manager uses VDSM to execute power management device actions, so another host in the environment is used as a fencing proxy.

You can select between:

- Any host in the same cluster as the host requiring fencing.
- Any host in the same data center as the host requiring fencing.

A viable fencing proxy host has a status of either **UP** or **Maintenance**.

4.3. POWER MANAGEMENT

The Red Hat Virtualization Manager is capable of rebooting hosts that have entered a non-operational or non-responsive state, as well as preparing to power off under-utilized hosts to save power. This functionality depends on a properly configured power management device. The Red Hat Virtualization environment supports the following power management devices:

- American Power Conversion (**apc**)
- IBM Bladecenter (**Bladecenter**)
- Cisco Unified Computing System (**cisco_ucs**)
- Dell Remote Access Card 5 (**drac5**)
- Dell Remote Access Card 7 (**drac7**)
- Electronic Power Switch (**eps**)
- HP BladeSystem (**hpblade**)
- Integrated Lights Out (**ilo, ilo2, ilo3, ilo4**)
- Intelligent Platform Management Interface (**ipmilan**)

- Remote Supervisor Adapter (**rsa**)
- Fujitsu–Siemens RSB (**rsb**)
- Western Telematic, Inc (**wti**)

HP servers should use **ilo3** or **ilo4**, Dell servers use **drac5** or Integrated Dell Remote Access Controllers (**idrac**), and IBM servers use **ipmilan**. Integrated Management Module (IMM) uses the IPMI protocol, and therefore IMM users can use **ipmilan**.



NOTE

APC 5.x power management devices are not supported by the **apc** fence agent. Use the **apc_snmp** fence agent instead.

In order to communicate with the listed power management devices, the Red Hat Virtualization Manager makes use of fence agents. The Red Hat Virtualization Manager allows administrators to configure a fence agent for the power management device in their environment with parameters the device will accept and respond to. Basic configuration options can be configured using the graphical user interface. Special configuration options can also be entered, and are passed un-parsed to the fence device. Special configuration options are specific to a given fence device, while basic configuration options are for functionalities provided by all supported power management devices. The basic functionalities provided by all power management devices are:

- **Status:** check the status of the host.
- **Start:** power on the host.
- **Stop:** power down the host.
- **Restart:** restart the host. Actually implemented as stop, wait, status, start, wait, status.

Best practice is to test the power management configuration once when initially configuring it, and occasionally after that to ensure continued functionality.

Resilience is provided by properly configured power management devices in all of the hosts in an environment. Fencing agents allow the Red Hat Virtualization Manager to communicate with host power management devices to bypass the operating system on a problem host, and isolate the host from the rest of its environment by rebooting it. The Manager can then reassign the SPM role, if it was held by the problem host, and safely restart any highly available virtual machines on other hosts.

4.4. FENCING

In the context of the Red Hat Virtualization environment, fencing is a host reboot initiated by the Manager using a fence agent and performed by a power management device. Fencing allows a cluster to react to unexpected host failures as well as enforce power saving, load balancing, and virtual machine availability policies.

Fencing ensures that the role of Storage Pool Manager (SPM) is always assigned to a functional host. If the fenced host was the SPM, the SPM role is relinquished and reassigned to a responsive host. Because the host with the SPM role is the only host that is able to write data domain structure metadata, a non-responsive, un-fenced SPM host causes its environment to lose the ability to create and destroy virtual disks, take snapshots, extend logical volumes, and all other actions that require changes to data domain structure metadata.

When a host becomes non-responsive, all of the virtual machines that are currently running on that host can also become non-responsive. However, the non-responsive host retains the lock on the virtual machine hard disk images for virtual machines it is running. Attempting to start a virtual machine on a second host and assign the second host write privileges for the virtual machine hard disk image can cause data corruption.

Fencing allows the Red Hat Virtualization Manager to assume that the lock on a virtual machine hard disk image has been released; the Manager can use a fence agent to confirm that the problem host has been rebooted. When this confirmation is received, the Red Hat Virtualization Manager can start a virtual machine from the problem host on another host without risking data corruption. Fencing is the basis for highly-available virtual machines. A virtual machine that has been marked highly-available can not be safely started on an alternate host without the certainty that doing so will not cause data corruption.

When a host becomes non-responsive, the Red Hat Virtualization Manager allows a grace period of thirty (30) seconds to pass before any action is taken, to allow the host to recover from any temporary errors. If the host has not become responsive by the time the grace period has passed, the Manager automatically begins to mitigate any negative impact from the non-responsive host. The Manager uses the fencing agent for the power management card on the host to stop the host, confirm it has stopped, start the host, and confirm that the host has been started. When the host finishes booting, it attempts to rejoin the cluster that it was a part of before it was fenced. If the issue that caused the host to become non-responsive has been resolved by the reboot, then the host is automatically set to **Up** status and is once again capable of starting and hosting virtual machines.

4.5. SOFT-FENCING HOSTS

Hosts can sometimes become non-responsive due to an unexpected problem, and though VDSM is unable to respond to requests, the virtual machines that depend upon VDSM remain alive and accessible. In these situations, restarting VDSM returns VDSM to a responsive state and resolves this issue.

"SSH Soft Fencing" is a process where the Manager attempts to restart VDSM via SSH on non-responsive hosts. If the Manager fails to restart VDSM via SSH, the responsibility for fencing falls to the external fencing agent if an external fencing agent has been configured.

Soft-fencing over SSH works as follows. Fencing must be configured and enabled on the host, and a valid proxy host (a second host, in an UP state, in the data center) must exist. When the connection between the Manager and the host times out, the following happens:

1. On the first network failure, the status of the host changes to "connecting".
2. The Manager then makes three attempts to ask VDSM for its status, or it waits for an interval determined by the load on the host. The formula for determining the length of the interval is configured by the configuration values `TimeoutToResetVdsInSeconds` (the default is 60 seconds) + `[DelayResetPerVmInSeconds` (the default is 0.5 seconds)]*(the count of running virtual machines on host) + `[DelayResetForSpmlnSeconds` (the default is 20 seconds)] * 1 (if host runs as SPM) or 0 (if the host does not run as SPM). To give VDSM the maximum amount of time to respond, the Manager chooses the longer of the two options mentioned above (three attempts to retrieve the status of VDSM or the interval determined by the above formula).
3. If the host does not respond when that interval has elapsed, **vdsmd restart** is executed via SSH.
4. If **vdsmd restart** does not succeed in re-establishing the connection between the host and the Manager, the status of the host changes to **Non Responsive** and, if power management is configured, fencing is handed off to the external fencing agent.

**NOTE**

Soft-fencing over SSH can be executed on hosts that have no power management configured. This is distinct from "fencing": fencing can be executed only on hosts that have power management configured.

4.6. USING MULTIPLE POWER MANAGEMENT FENCING AGENTS

Single agents are treated as primary agents. The secondary agent is valid when there are two fencing agents, for example for dual-power hosts in which each power switch has two agents connected to the same power switch. Agents can be of the same or different types.

Having multiple fencing agents on a host increases the reliability of the fencing procedure. For example, when the sole fencing agent on a host fails, the host will remain in a non-operational state until it is manually rebooted. The virtual machines previously running on the host will be suspended, and only fail over to another host in the cluster after the original host is manually fenced. With multiple agents, if the first agent fails, the next agent can be called.

When two fencing agents are defined on a host, they can be configured to use a concurrent or sequential flow:

- **Concurrent:** Both primary and secondary agents have to respond to the Stop command for the host to be stopped. If one agent responds to the Start command, the host will go up.
- **Sequential:** To stop or start a host, the primary agent is used first, and if it fails, the secondary agent is used.

CHAPTER 5. LOAD BALANCING, SCHEDULING, AND MIGRATION

5.1. LOAD BALANCING, SCHEDULING, AND MIGRATION

Individual hosts have finite hardware resources, and are susceptible to failure. To mitigate against failure and resource exhaustion, hosts are grouped into clusters, which are essentially a grouping of shared resources. A Red Hat Virtualization environment responds to changes in demand for host resources using load balancing policy, scheduling, and migration. The Manager is able to ensure that no single host in a cluster is responsible for all of the virtual machines in that cluster. Conversely, the Manager is able to recognize an underutilized host, and migrate all virtual machines off of it, allowing an administrator to shut down that host to save power.

Available resources are checked as a result of three events:

- Virtual machine start - Resources are checked to determine on which host a virtual machine will start.
- Virtual machine migration - Resources are checked in order to determine an appropriate target host.
- Time elapses - Resources are checked at a regular interval to determine whether individual host load is in compliance with cluster load balancing policy.

The Manager responds to changes in available resources by using the load balancing policy for a cluster to schedule the migration of virtual machines from one host in a cluster to another. The relationship between load balancing policy, scheduling, and virtual machine migration are discussed in the following sections.

5.2. LOAD BALANCING POLICY

Load balancing policy is set for a cluster, which includes one or more hosts that may each have different hardware parameters and available memory. The Red Hat Virtualization Manager uses a load balancing policy to determine which host in a cluster to start a virtual machine on. Load balancing policy also allows the Manager determine when to move virtual machines from over-utilized hosts to under-utilized hosts.

The load balancing process runs once every minute for each cluster in a data center. It determines which hosts are over-utilized, which are hosts under-utilized, and which are valid targets for virtual machine migration. The determination is made based on the load balancing policy set by an administrator for a given cluster. The options for load balancing policies are **VM_Evenly_Distributed**, **Evenly_Distributed**, **Power_Saving**, **Cluster_Maintenance**, and **None**.

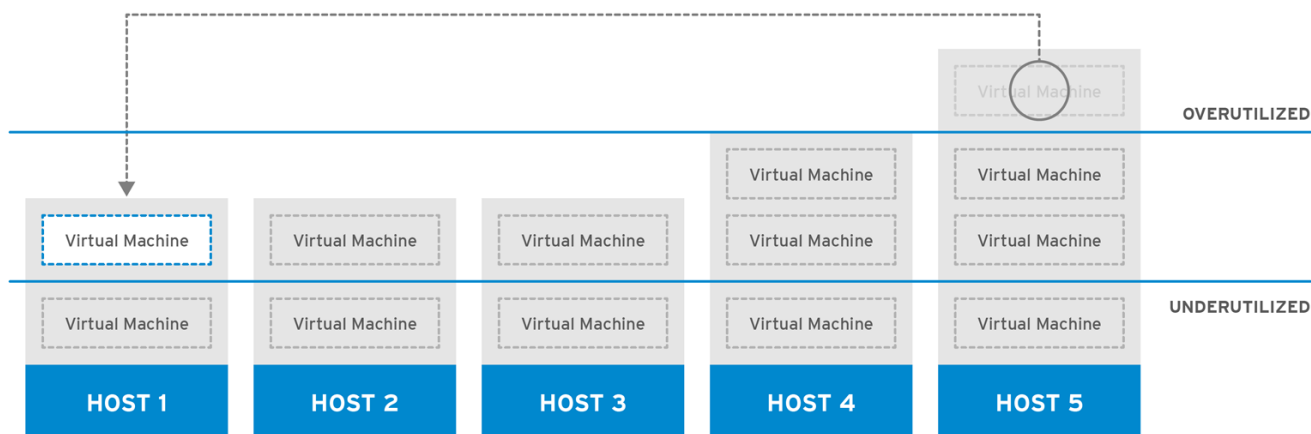
5.3. LOAD BALANCING POLICY: VM_EVENLY_DISTRIBUTED

A virtual machine evenly distributed load balancing policy distributes virtual machines evenly between hosts based on a count of the virtual machines. The high virtual machine count is the maximum number of virtual machines that can run on each host, beyond which qualifies as overloading the host. The **VM_Evenly_Distributed** policy allows an administrator to set a high virtual machine count for hosts. The maximum inclusive difference in virtual machine count between the most highly-utilized host and the least-utilized host is also set by an administrator. The cluster is balanced when every host in the cluster has a virtual machine count that falls inside this migration threshold. The administrator also sets the number of slots for virtual machines to be reserved on SPM hosts. The SPM host will have a lower load than other hosts, so this variable defines how many fewer virtual machines than other hosts it can run. If any host is running more virtual machines than the high virtual machine count and at least one host has a

virtual machine count that falls outside of the migration threshold, virtual machines are migrated one by one to the host in the cluster that has the lowest CPU utilization. One virtual machine is migrated at a time until every host in the cluster has a virtual machine count that falls within the migration threshold.

5.4. LOAD BALANCING POLICY: EVENLY_DISTRIBUTED

Figure 5.1. Evenly Distributed Scheduling Policy

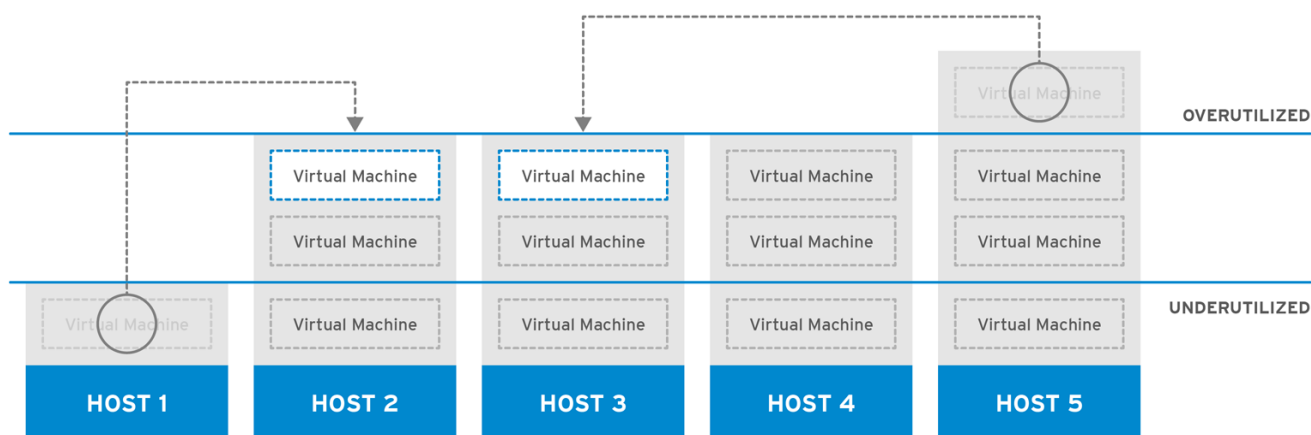


RHV_444396_0417

An evenly distributed load balancing policy selects the host for a new virtual machine according to lowest CPU load or highest available memory. The maximum CPU load and minimum available memory that is allowed for hosts in a cluster for a set amount of time are defined by the evenly distributed scheduling policy's parameters. Beyond these limits the environment's performance will degrade. The evenly distributed policy allows an administrator to set these levels for running virtual machines. If a host has reached the defined maximum CPU load or minimum available memory and the host stays there for more than the set time, virtual machines on that host are migrated one by one to the host in the cluster that has the lowest CPU or highest available memory depending on which parameter is being utilized. Host resources are checked once per minute, and one virtual machine is migrated at a time until the host CPU load is below the defined limit or the host available memory is above the defined limit.

5.5. LOAD BALANCING POLICY: POWER_SAVING

Figure 5.2. Power Saving Scheduling Policy



RHV_444396_0417

A power saving load balancing policy selects the host for a new virtual machine according to lowest CPU or highest available memory. The maximum CPU load and minimum available memory that is allowed for hosts in a cluster for a set amount of time is defined by the power saving scheduling policy's parameters.

Beyond these limits the environment's performance will degrade. The power saving parameters also define the minimum CPU load and maximum available memory allowed for hosts in a cluster for a set amount of time before the continued operation of a host is considered an inefficient use of electricity. If a host has reached the maximum CPU load or minimum available memory and stays there for more than the set time, the virtual machines on that host are migrated one by one to the host that has the lowest CPU or highest available memory depending on which parameter is being utilized. Host resources are checked once per minute, and one virtual machine is migrated at a time until the host CPU load is below the defined limit or the host available memory is above the defined limit. If the host's CPU load falls below the defined minimum level or the host's available memory rises above the defined maximum level the virtual machines on that host are migrated to other hosts in the cluster as long as the other hosts in the cluster remain below maximum CPU load and above minimum available memory. When an under-utilized host is cleared of its remaining virtual machines, the Manager will automatically power down the host machine, and restart it again when load balancing requires or there are not enough free hosts in the cluster.

5.6. LOAD BALANCING POLICY: NONE

If no load balancing policy is selected, virtual machines are started on the host within a cluster with the lowest CPU utilization and available memory. To determine CPU utilization a combined metric is used that takes into account the virtual CPU count and the CPU usage percent. This approach is the least dynamic, as the only host selection point is when a new virtual machine is started. Virtual machines are not automatically migrated to reflect increased demand on a host.

An administrator must decide which host is an appropriate migration target for a given virtual machine. Virtual machines can also be associated with a particular host using pinning. Pinning prevents a virtual machine from being automatically migrated to other hosts. For environments where resources are highly consumed, manual migration is the best approach.

5.7. LOAD BALANCING POLICY: CLUSTER_MAINTENANCE

A cluster maintenance scheduling policy limits activity in a cluster during maintenance tasks. When a cluster maintenance policy is set:

- No new virtual machines may be started, except highly available virtual machines. (Users can create highly available virtual machines and start them manually.)
- In the event of host failure, highly available virtual machines will restart properly and any virtual machine can migrate.

5.8. HIGHLY AVAILABLE VIRTUAL MACHINE RESERVATION

A highly available (HA) virtual machine reservation policy enables the Red Hat Virtualization Manager to monitor cluster capacity for highly available virtual machines. The Manager has the capability to flag individual virtual machines for High Availability, meaning that in the event of a host failure, these virtual machines will be rebooted on an alternative host. This policy balances highly available virtual machines across the hosts in a cluster. If any host in the cluster fails, the remaining hosts can support the migrating load of highly available virtual machines without affecting cluster performance. When highly available virtual machine reservation is enabled, the Manager ensures that appropriate capacity exists within a cluster for HA virtual machines to migrate in the event that their existing host fails unexpectedly.

5.9. SCHEDULING

In Red Hat Virtualization, scheduling refers to the way the Red Hat Virtualization Manager selects a host in a cluster as the target for a new or migrated virtual machine.

For a host to be eligible to start a virtual machine or accept a migrated virtual machine from another host, it must have enough free memory and CPUs to support the requirements of the virtual machine being started on or migrated to it. A virtual machine will not start on a host with an overloaded CPU. By default, a host's CPU is considered overloaded if it has a load of more than 80% for 5 minutes, but these values can be changed using scheduling policies. If multiple hosts are eligible targets, one will be selected based on the load balancing policy for the cluster. For example, if the `Evenly_Distributed` policy is in effect, the Manager chooses the host with the lowest CPU utilization. If the `Power_Saving` policy is in effect, the host with the lowest CPU utilization between the maximum and minimum service levels will be selected. The Storage Pool Manager (SPM) status of a given host also affects eligibility as a target for starting virtual machines or virtual machine migration. A non-SPM host is a preferred target host, for instance, the first virtual machine started in a cluster will not run on the SPM host if the SPM role is held by a host in that cluster.

See [Scheduling Policies](#) in the *Administration Guide* for more information.

5.10. MIGRATION

The Red Hat Virtualization Manager uses migration to enforce load balancing policies for a cluster. Virtual machine migration takes place according to the load balancing policy for a cluster and current demands on hosts within a cluster. Migration can also be configured to automatically occur when a host is fenced or moved to maintenance mode. The Red Hat Virtualization Manager first migrates virtual machines with the lowest CPU utilization. This is calculated as a percentage, and does not take into account RAM usage or I/O operations, except as I/O operations affect CPU utilization. If there are more than one virtual machines with the same CPU usage, the one that will be migrated first is the first virtual machine returned by the database query run by the Red Hat Virtualization Manager to determine virtual machine CPU usage.

Virtual machine migration has the following limitations by default:

- A bandwidth limit of 52 MiBps is imposed on each virtual machine migration.
- A migration will time out after 64 seconds per GB of virtual machine memory.
- A migration will abort if progress is stalled for 240 seconds.
- Concurrent outgoing migrations are limited to one per CPU core per host, or 2, whichever is smaller.

See [Understanding live migration "migration_max_bandwidth" and "max_outgoing_migrations" parameters in vds.conf](#) for details about tuning migration settings.

CHAPTER 6. DIRECTORY SERVICES

6.1. DIRECTORY SERVICES

The Red Hat Virtualization platform relies on directory services for user authentication and authorization. Interactions with all Manager interfaces, including the VM Portal, Administration Portal, and REST API are limited to authenticated, authorized users. Virtual machines within the Red Hat Virtualization environment can use the same directory services to provide authentication and authorization, however they must be configured to do so. The currently supported providers of directory services for use with the Red Hat Virtualization Manager are Identity Management (IdM), Red Hat Directory Server 9 (RHDS), Active Directory (AD), and OpenLDAP. The Red Hat Virtualization Manager interfaces with the directory server for:

- Portal logins (User, Power User, Administrator, REST API).
- Queries to display user information.
- Adding the Manager to a domain.

Authentication is the verification and identification of a party who generated some data, and of the integrity of the generated data. A principal is the party whose identity is verified. The verifier is the party who demands assurance of the principal's identity. In the case of Red Hat Virtualization, the Manager is the verifier and a user is a principal. Data integrity is the assurance that the data received is the same as the data generated by the principal.

Confidentiality and authorization are closely related to authentication. Confidentiality protects data from disclosure to those not intended to receive it. Strong authentication methods can optionally provide confidentiality. Authorization determines whether a principal is allowed to perform an operation. Red Hat Virtualization uses directory services to associate users with roles and provide authorization accordingly. Authorization is usually performed after the principal has been authenticated, and may be based on information local or remote to the verifier.

During installation, a local, internal domain is automatically configured for administration of the Red Hat Virtualization environment. After the installation is complete, more domains can be added.

6.2. LOCAL AUTHENTICATION: INTERNAL DOMAIN

The Red Hat Virtualization Manager creates a limited, internal administration domain during installation. This domain is not the same as an AD or IdM domain, because it exists based on a key in the Red Hat Virtualization PostgreSQL database rather than as a directory service user on a directory server. The internal domain is also different from an external domain because the internal domain will only have one user: the **admin@internal** user. Taking this approach to initial authentication allows Red Hat Virtualization to be evaluated without requiring a complete, functional directory server, and ensures an administrative account is available to troubleshoot any issues with external directory services.

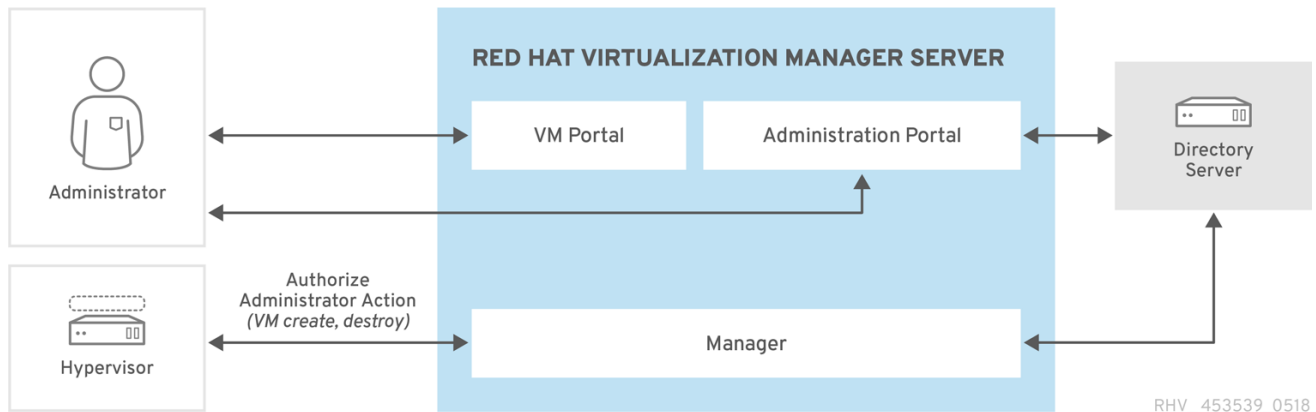
The **admin@internal** user is for the initial configuration of an environment. This includes installing and accepting hosts, adding external AD or IdM authentication domains, and delegating permissions to users from external domains.

6.3. REMOTE AUTHENTICATION USING GSSAPI

In the context of Red Hat Virtualization, remote authentication refers to authentication that is handled by a remote service, not the Red Hat Virtualization Manager. Remote authentication is used for user or API connections coming to the Manager from within an AD, IdM, or RHDS domain. The Red Hat

Virtualization Manager must be configured by an administrator using the **engine-manage-domains** tool to be a part of an RHDS, AD, or IdM domain. This requires that the Manager be provided with credentials for an account from the RHDS, AD, or IdM directory server for the domain with sufficient privileges to join a system to the domain. After domains have been added, domain users can be authenticated by the Red Hat Virtualization Manager against the directory server using a password. The Manager uses a framework called the Simple Authentication and Security Layer (SASL) which in turn uses the Generic Security Services Application Program Interface (GSSAPI) to securely verify the identity of a user, and ascertain the authorization level available to the user.

Figure 6.1. GSSAPI Authentication



CHAPTER 7. TEMPLATES AND POOLS

7.1. TEMPLATES AND POOLS

The Red Hat Virtualization environment provides administrators with tools to simplify the provisioning of virtual machines to users. These are templates and pools. A template is a shortcut that allows an administrator to quickly create a new virtual machine based on an existing, pre-configured virtual machine, bypassing operating system installation and configuration. This is especially helpful for virtual machines that will be used like appliances, for example web server virtual machines. If an organization uses many instances of a particular web server, an administrator can create a virtual machine that will be used as a template, installing an operating system, the web server, any supporting packages, and applying unique configuration changes. The administrator can then create a template based on the working virtual machine that will be used to create new, identical virtual machines as they are required.

Virtual machine pools are groups of virtual machines based on a given template that can be rapidly provisioned to users. Permission to use virtual machines in a pool is granted at the pool level; a user who is granted permission to use the pool will be assigned any virtual machine from the pool. Inherent in a virtual machine pool is the transitory nature of the virtual machines within it. Because users are assigned virtual machines without regard for which virtual machine in the pool they have used in the past, pools are not suited for purposes which require data persistence. Virtual machine pools are best suited for scenarios where either user data is stored in a central location and the virtual machine is a means to accessing and using that data, or data persistence is not important. The creation of a pool results in the creation of the virtual machines that populate the pool, in a stopped state. These are then started on user request.

7.2. TEMPLATES

To create a template, an administrator creates and customizes a virtual machine. Desired packages are installed, customized configurations are applied, the virtual machine is prepared for its intended purpose in order to minimize the changes that must be made to it after deployment. An optional but recommended step before creating a template from a virtual machine is generalization. Generalization is used to remove details like system user names, passwords, and timezone information that will change upon deployment. Generalization does not affect customized configurations. Generalization of Windows and Linux guests in the Red Hat Virtualization environment is discussed further in [Templates](#) in the *Virtual Machine Management Guide*. Red Hat Enterprise Linux guests are generalized using **sys-unconfig**. Windows guests are generalized using **sys-prep**.

When the virtual machine that provides the basis for a template is satisfactorily configured, generalized if desired, and stopped, an administrator can create a template from the virtual machine. Creating a template from a virtual machine causes a read-only copy of the specially configured virtual disk to be created. The read-only image forms the backing image for all subsequently created virtual machines that are based on that template. In other words, a template is essentially a customized read-only virtual disk with an associated virtual hardware configuration. The hardware can be changed in virtual machines created from a template, for instance, provisioning two gigabytes of RAM for a virtual machine created from a template that has one gigabyte of RAM. The template virtual disk, however, cannot be changed as doing so would result in changes for all virtual machines based on the template.

When a template has been created, it can be used as the basis for multiple virtual machines. Virtual machines are created from a given template using a **Thin** provisioning method or a **Clone** provisioning method. Virtual machines that are cloned from templates take a complete writable copy of the template base image, sacrificing the space savings of the thin creation method in exchange for no longer depending on the presence of the template. Virtual machines that are created from a template using the thin method use the read-only image from the template as a base image, requiring that the template and all virtual machines created from it be stored on the same storage domain. Changes to data and newly generated data are stored in a copy-on-write image. Each virtual machine based on a template

uses the same base read-only image, as well as a copy-on-write image that is unique to the virtual machine. This provides storage savings by limiting the number of times identical data is kept in storage. Furthermore, frequent use of the read-only backing image can cause the data being accessed to be cached, resulting in a net performance increase.

7.3. POOLS

Virtual machine pools allow for rapid provisioning of numerous identical virtual machines to users as desktops. Users who have been granted permission to access and use virtual machines from a pool receive an available virtual machine based on their position in a queue of requests. Virtual machines in a pool do not allow data persistence; each time a virtual machine is assigned from a pool, it is allocated in its base state. This is ideally suited to be used in situations where user data is stored centrally.

Virtual machine pools are created from a template. Each virtual machine in a pool uses the same backing read-only image, and uses a temporary copy-on-write image to hold changed and newly generated data. Virtual machines in a pool are different from other virtual machines in that the copy-on-write layer that holds user-generated and -changed data is lost at shutdown. The implication of this is that a virtual machine pool requires no more storage than the template that backs it, plus some space for data generated or changed during use. Virtual machine pools are an efficient way to provide computing power to users for some tasks without the storage cost of providing each user with a dedicated virtual desktop.

Example 7.1. Example Pool Usage

A technical support company employs 10 help desk staff. However, only five are working at any given time. Instead of creating ten virtual machines, one for each help desk employee, a pool of five virtual machines can be created. Help desk employees allocate themselves a virtual machine at the beginning of their shift and return it to the pool at the end.

CHAPTER 8. VIRTUAL MACHINE SNAPSHOTS

8.1. SNAPSHOTS

Snapshots are a storage function that allows an administrator to create a restore point of a virtual machine's operating system, applications, and data at a certain point in time. Snapshots save the data currently present in a virtual machine hard disk image as a COW volume and allow for a recovery to the data as it existed at the time the snapshot was taken. A snapshot causes a new COW layer to be created over the current layer. All write actions performed after a snapshot is taken are written to the new COW layer.

It is important to understand that a virtual machine hard disk image is a chain of one or more volumes. From the perspective of a virtual machine, these volumes appear as a single disk image. A virtual machine is oblivious to the fact that its disk is comprised of multiple volumes.

The term COW volume and COW layer are used interchangeably, however, layer more clearly recognizes the temporal nature of snapshots. Each snapshot is created to allow an administrator to discard unsatisfactory changes made to data **after** the snapshot is taken. Snapshots provide similar functionality to the **Undo** function present in many word processors.



NOTE

Snapshots of virtual machine hard disks marked **shareable** and those that are based on **Direct LUN** connections are not supported, live or otherwise.

The three primary snapshot operations are:

- Creation, which involves the first snapshot created for a virtual machine.
- Previews, which involves previewing a snapshot to determine whether or not to restore the system data to the point in time that the snapshot was taken.
- Deletion, which involves deleting a restoration point that is no longer required.

For task-based information about snapshot operations, see [Snapshots](#) in the *Red Hat Virtualization Virtual Machine Management Guide*.

8.2. LIVE SNAPSHOTS IN RED HAT VIRTUALIZATION

Snapshots of virtual machine hard disks marked **shareable** and those that are based on **Direct LUN** connections are not supported, live or otherwise.

Any other virtual machine that is not being cloned or migrated can have a snapshot taken when running, paused, or stopped.

When a live snapshot of a virtual machine is initiated, the Manager requests that the SPM host create a new volume for the virtual machine to use. When the new volume is ready, the Manager uses VDSM to communicate with libvirt and qemu on the host running the virtual machine that it should begin using the new volume for virtual machine write operations. If the virtual machine is able to write to the new volume, the snapshot operation is considered a success and the virtual machine stops writing to the previous volume. If the virtual machine is unable to write to the new volume, the snapshot operation is considered a failure, and the new volume is deleted.

The virtual machine requires access to both its current volume and the new one from the time when a live snapshot is initiated until after the new volume is ready, so both volumes are opened with read-write access.

Virtual machines with an installed guest agent that supports quiescing can ensure filesystem consistency across snapshots. Registered Red Hat Enterprise Linux guests can install the **qemu-guest-agent** to enable quiescing before snapshots.

If a quiescing compatible guest agent is present on a virtual machine when it a snapshot is taken, VDSM uses libvirt to communicate with the agent to prepare for a snapshot. Outstanding write actions are completed, and then filesystems are frozen before a snapshot is taken. When the snapshot is complete, and libvirt has switched the virtual machine to the new volume for disk write actions, the filesystem is thawed, and writes to disk resume.

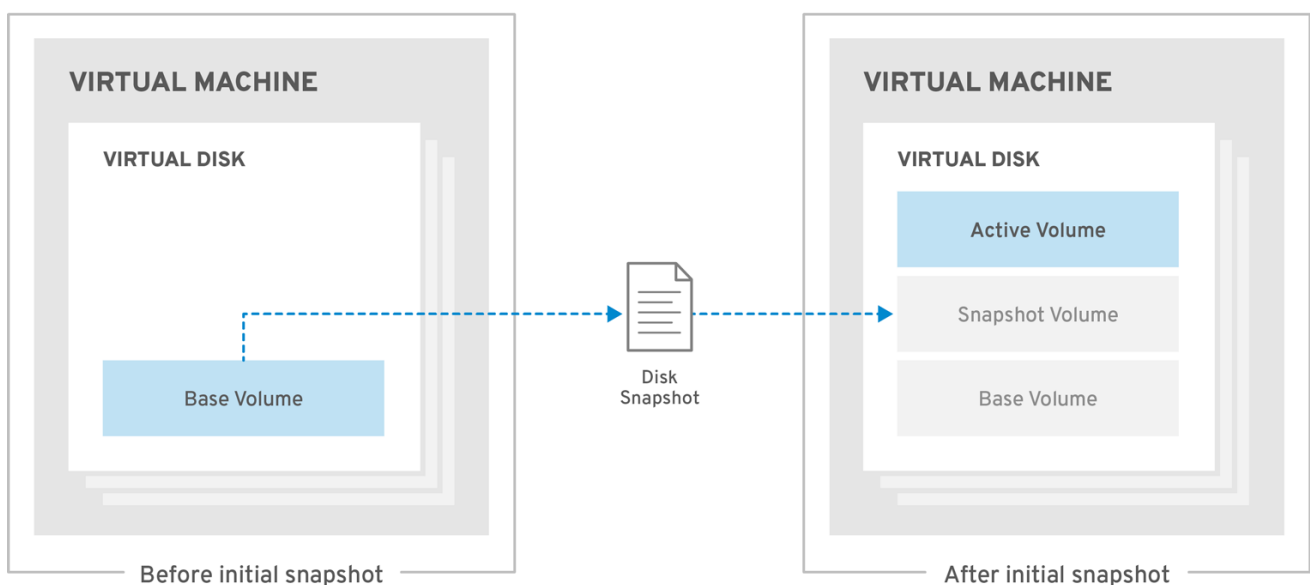
All live snapshots attempted with quiescing enabled. If the snapshot command fails because there is no compatible guest agent present, the live snapshot is re-initiated without the use-quiescing flag. When a virtual machine is reverted to its pre-snapshot state with quiesced filesystems, it boots cleanly with no filesystem check required. Reverting the previous snapshot using an un-quiesced filesystem requires a filesystem check on boot.

8.3. SNAPSHOT CREATION

In Red Hat Virtualization the initial snapshot for a virtual machine is different from subsequent snapshots in that the initial snapshot retains its format, either QCOW2 or raw. The first snapshot for a virtual machine uses existing volumes as a base image. Additional snapshots are additional COW layers tracking the changes made to the data stored in the image since the previous snapshot.

As depicted in [Initial Snapshot Creation](#), the creation of a snapshot causes the volumes that comprise a virtual disk to serve as the base image for all subsequent snapshots.

Figure 8.1. Initial Snapshot Creation

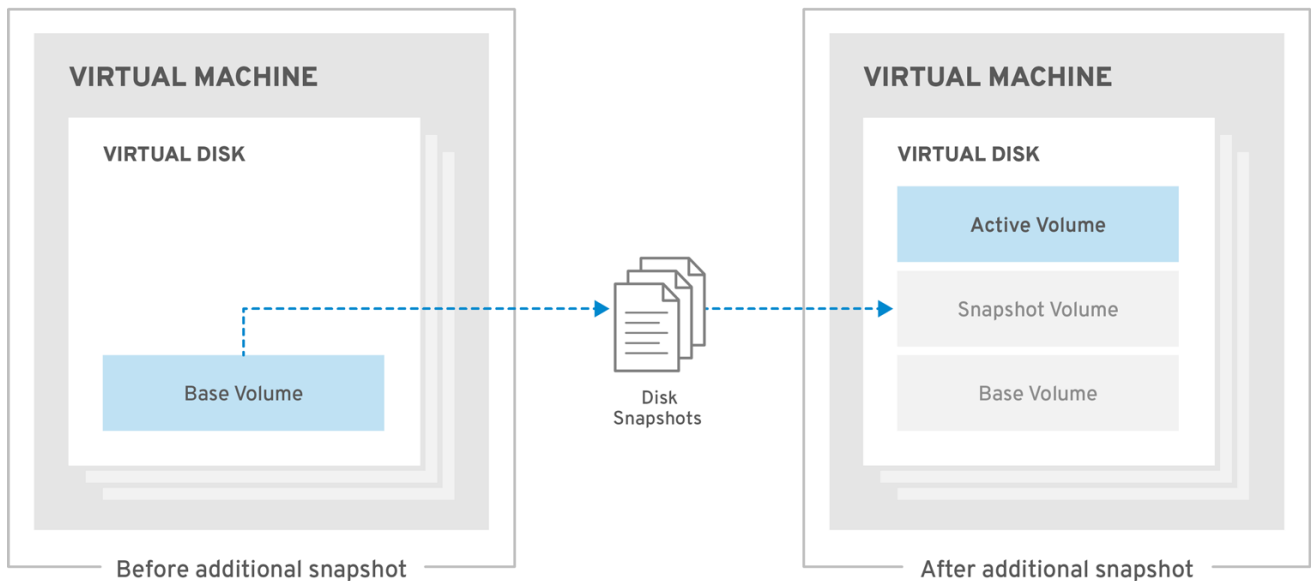


RHV_453539_0718

Snapshots taken after the initial snapshot result in the creation of new COW volumes in which data that is created or changed after the snapshot is taken will be stored. Each newly created COW layer contains only COW metadata. Data that is created by using and operating the virtual machine after a snapshot is taken is written to this new COW layer. When a virtual machine is used to modify data that exists in a

previous COW layer, the data is read from the previous layer, and written into the newest layer. Virtual machines locate data by checking each COW layer from most recent to oldest, transparently to the virtual machine.

Figure 8.2. Additional Snapshot Creation



RHV_453539_0718

8.4. MONITORING SNAPSHOT HEALTH WITH THE IMAGE DISCREPANCIES TOOL

The **RHV Image Discrepancies** tool analyzes image data in the Storage Domain and RHV Database. It alerts you if it finds discrepancies in volumes and volume attributes, but does not fix those discrepancies. Use this tool in a variety of scenarios, such as:

- Before upgrading versions, to avoid carrying over broken volumes or chains to the new version.
- Following a failed storage operation, to detect volumes or attributes in a bad state.
- After restoring the RHV database or storage from backup.
- Periodically, to detect potential problems before they worsen.
- To analyze a snapshot- or live storage migration-related issues, and to verify system health after fixing these types of problems.

Prerequisites

- **Required Versions:** this tool was introduced in RHV version 4.3.8 with **`rhv-log-collector-analyzer-0.2.15-0.el7ev`**.
- Because data collection runs simultaneously at different places and is not atomic, stop all activity in the environment that can modify the storage domains. That is, do not create or remove snapshots, edit, move, create, or remove disks. Otherwise, false detection of inconsistencies may occur. Virtual Machines can remain running normally during the process.

Procedure

1. To run the tool, enter the following command on the RHV Manager:

```
# rhv-image-discrepancies
```

2. If the tool finds discrepancies, rerun it to confirm the results, especially if there is a chance some operations were performed while the tool was running.



NOTE

This tool includes any Export and ISO storage domains and may report discrepancies for them. If so, these can be ignored, as these storage domains do not have entries for images in the RHV database.

Understanding the results

The tool reports the following:

- If there are volumes that appear on the storage but are not in the database, or appear in the database but are not on the storage.
- If some volume attributes differ between the storage and the database.

Sample output:

```
Checking storage domain c277ad93-0973-43d9-a0ca-22199bc8e801
  Looking for missing images...
  No missing images found
  Checking discrepancies between SD/DB attributes...
  image ef325650-4b39-43cf-9e00-62b9f7659020 has a different attribute capacity on
  storage(2696984576) and on DB(2696986624)
  image 852613ce-79ee-4adc-a56a-ea650dcb4cfa has a different attribute capacity on
  storage(5424252928) and on DB(5424254976)

Checking storage domain c64637b4-f0e8-408c-b8af-6a52946113e2
  Looking for missing images...
  No missing images found
  Checking discrepancies between SD/DB attributes...
  No discrepancies found
```

8.5. SNAPSHOT PREVIEWS

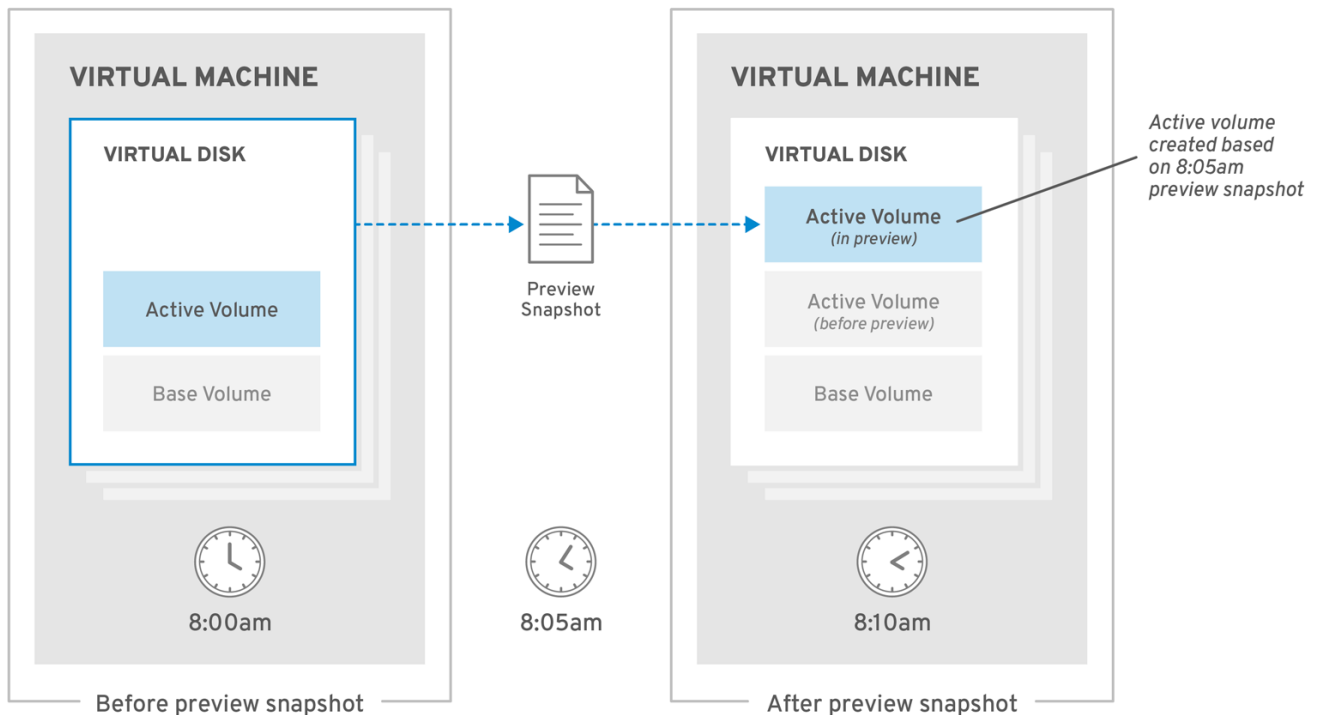
To select which snapshot a virtual disk will be reverted to, the administrator can preview all previously created snapshots.

From the available snapshots per guest, the administrator can select a snapshot volume to preview its contents. As depicted in [Preview Snapshot](#), each snapshot is saved as a COW volume, and when it is previewed, a new preview layer is copied from the snapshot being previewed. The guest interacts with the preview instead of the actual snapshot volume.

After the administrator previews the selected snapshot, the preview can be committed to restore the guest data to the state captured in the snapshot. If the administrator commits the preview, the guest is attached to the preview layer.

After a snapshot is previewed, the administrator can select **Undo** to discard the preview layer of the viewed snapshot. The layer that contains the snapshot itself is preserved despite the preview layer being discarded.

Figure 8.3. Preview Snapshot

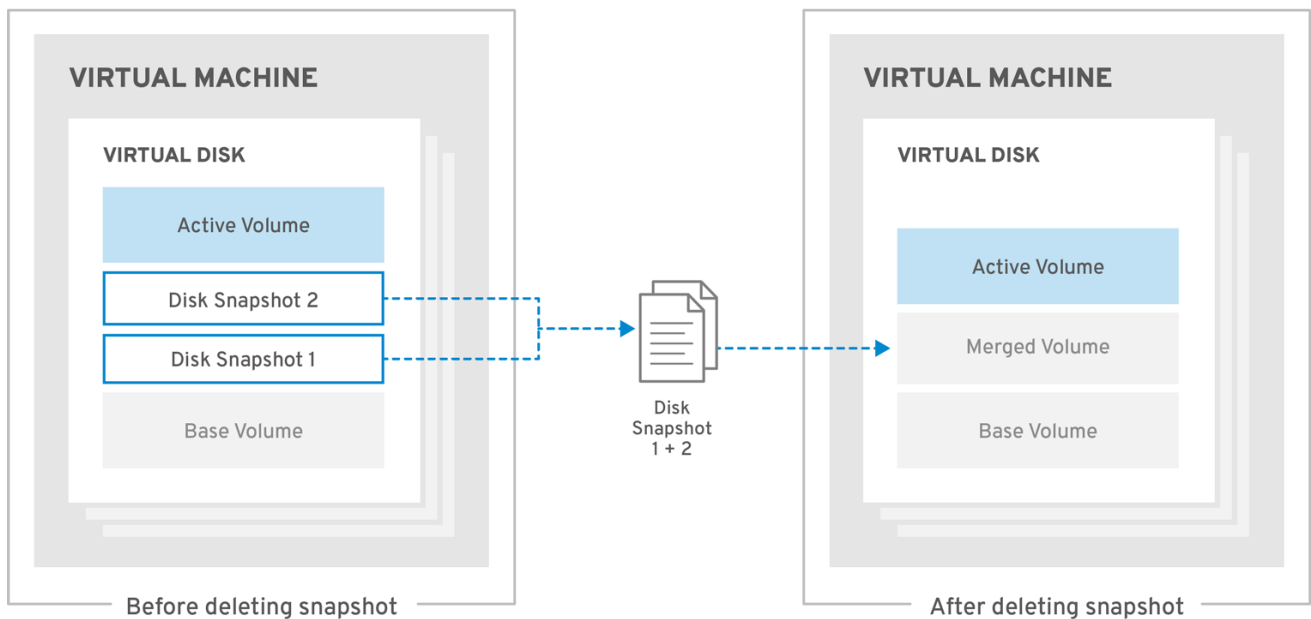


RHV_453539_0718

8.6. SNAPSHOT DELETION


You can delete individual snapshots that are no longer required. Deleting a snapshot removes the ability to restore a virtual disk to that particular restoration point. It does not necessarily reclaim the disk space consumed by the snapshot, nor does it delete the data. The disk space will only be reclaimed if a subsequent snapshot has overwritten the data of the deleted snapshot. For example, if the third snapshot out of five snapshots is deleted, the unchanged data in the third snapshot must be preserved on the disk for the fourth and fifth snapshots to be usable; however, if the fourth or fifth snapshot has overwritten the data of the third, then the third snapshot has been made redundant and the disk space can be reclaimed. Aside from potential disk space reclamation, deleting a snapshot may also improve the performance of the virtual machine.

Figure 8.4. Snapshot Deletion



RHV_453539_0718

Snapshot deletion is handled as an asynchronous block job in which VDSM maintains a record of the operation in the recovery file for the virtual machine so that the job can be tracked even if VDSM is restarted or the virtual machine is shut down during the operation. Once the operation begins, the snapshot being deleted cannot be previewed or used as a restoration point, even if the operation fails or is interrupted. In operations in which the active layer is to be merged with its parent, the operation is split into a two-stage process during which data is copied from the active layer to the parent layer, and disk writes are mirrored to both the active layer and the parent. Finally, the job is considered complete once the data in the snapshot being deleted has been merged with its parent snapshot and VDSM synchronizes the changes throughout the image chain.

**NOTE**

If the deletion fails, fix the underlying problem (for example, a failed host, an inaccessible storage device, or even a temporary network issue) and try again.

CHAPTER 9. HARDWARE DRIVERS AND DEVICES

9.1. VIRTUALIZED HARDWARE

Red Hat Virtualization presents three distinct types of system devices to virtualized guests. These hardware devices all appear as physically attached hardware devices to the virtualized guest but the device drivers work in different ways.

Emulated devices

Emulated devices, sometimes referred to as virtual devices, exist entirely in software. Emulated device drivers are a translation layer between the operating system running on the host (which manages the source device) and the operating systems running on the guests. The device level instructions directed to and from the emulated device are intercepted and translated by the hypervisor. Any device of the same type as that being emulated and recognized by the Linux kernel is able to be used as the backing source device for the emulated drivers.

Para-virtualized Devices

Para-virtualized devices require the installation of device drivers on the guest operating system providing it with an interface to communicate with the hypervisor on the host machine. This interface is used to allow traditionally intensive tasks such as disk I/O to be performed outside of the virtualized environment. Lowering the overhead inherent in virtualization in this manner is intended to allow guest operating system performance closer to that expected when running directly on physical hardware.

Physically shared devices

Certain hardware platforms allow virtualized guests to directly access various hardware devices and components. This process in virtualization is known as passthrough or device assignment. Passthrough allows devices to appear and behave as if they were physically attached to the guest operating system.

9.2. STABLE DEVICE ADDRESSES IN RED HAT VIRTUALIZATION

Virtual hardware PCI address allocations are persisted in the **ovirt-engine** database.

PCI addresses are allocated by **QEMU** at virtual machine creation time, and reported to **VDSM** by **libvirt**. **VDSM** reports them back to the Manager, where they are stored in the **ovirt-engine** database.

When a virtual machine is started, the Manager sends **VDSM** the device address out of the database. **VDSM** passes them to **libvirt** which starts the virtual machine using the PCI device addresses that were allocated when the virtual machine was run for the first time.

When a device is removed from a virtual machine, all references to it, including the stable PCI address, are also removed. If a device is added to replace the removed device, it is allocated a PCI address by **QEMU**, which is unlikely to be the same as the device it replaced.

9.3. CENTRAL PROCESSING UNIT (CPU)

Each host within a cluster has a number of virtual CPUs (vCPUs). The virtual CPUs are in turn exposed to guests running on the hosts. All virtual CPUs exposed by hosts within a cluster are of the type selected when the cluster was initially created via Red Hat Virtualization Manager. Mixing of virtual CPU types within a cluster is not possible.

Each available virtual CPU type has characteristics based on physical CPUs of the same name. The virtual CPU is indistinguishable from the physical CPU to the guest operating system.

**NOTE**

Support for x2APIC:

All virtual CPU models provided by Red Hat Enterprise Linux 7 hosts include support for x2APIC. This provides an Advanced Programmable Interrupt Controller (APIC) to better handle hardware interrupts.

9.4. SYSTEM DEVICES

System devices are critical for the guest to run and cannot be removed. Each system device attached to a guest also takes up an available PCI slot. The default system devices are:

- Host bridge
- ISA bridge and USB bridge (The USB and ISA bridges are the same device)
- Graphics card using the VGA or qxl driver
- Memory balloon device

For information about how to use PCI Express and conventional PCI devices with Intel Q35-based virtual machines, see [Using PCI Express and Conventional PCI Devices with the Q35 Virtual Machine](#) .

9.5. NETWORK DEVICES

Red Hat Virtualization is able to expose three different types of network interface controller to guests. The type of network interface controller to expose to a guest is chosen when the guest is created but is changeable from the Red Hat Virtualization Manager.

- The **e1000** network interface controller exposes a virtualized Intel PRO/1000 (e1000) to guests.
- The **virtio** network interface controller exposes a para-virtualized network device to guests.
- The **rtl8139** network interface controller exposes a virtualized **Realtek Semiconductor Corp RTL8139** to guests.

Multiple network interface controllers are permitted per guest. Each controller added takes up an available PCI slot on the guest.

9.6. GRAPHICS DEVICES

The SPICE or VNC graphics protocols can be used to connect to the emulated graphics devices.

You can select a **Video Type** in the Administration Portal:

- **QXL**: Emulates a para-virtualized video card that works best with QXL guest drivers
- **VGA**: Emulates a dummy VGA card with **Bochs** VESA extensions
- **BOCHS**: Emulates a dummy VGA card without legacy emulation for guest machines that that run with UEFI. This is the default display video card emulator for UEFI servers.



NOTE

For a virtual machine of type **server** that is set with UEFI and uses compatibility level 4.6 or above, **BOCHS** is the default value of **Video Type**.

In Red Hat Virtualization 4.4.5, you must do the following to enable this feature:

1. Run the following command:

```
engine-config --set EnableBochsDisplay=true --cver=<version>
```

where **<version>** is the compatibility version.

2. Restart the engine.
3. Set **Video Type** to **BOCHS** manually.

9.7. STORAGE DEVICES

Storage devices and storage pools can use the block device drivers to attach storage devices to virtualized guests. Note that the storage drivers are not storage devices. The drivers are used to attach a backing storage device, file or storage pool volume to a virtualized guest. The backing storage device can be any supported type of storage device, file, or storage pool volume.

- The **IDE** driver exposes an emulated block device to guests. The emulated **IDE** driver can be used to attach any combination of up to four virtualized **IDE** hard disks or virtualized **IDE** CD-ROM drives to each virtualized guest. The emulated **IDE** driver is also used to provide virtualized DVD-ROM drives.
- The **VirtIO** driver exposes a para-virtualized block device to guests. The para-virtualized block driver is a driver for all storage devices supported by the hypervisor attached to the virtualized guest (except for floppy disk drives, which must be emulated).

9.8. SOUND DEVICES

Two emulated sound devices are available:

- The **ac97** emulates an **Intel 82801AA AC97 Audio** compatible sound card.
- The **es1370** emulates an **ENSONIQ AudioPCI ES1370** sound card.

9.9. SERIAL DRIVER

The para-virtualized serial driver (**virtio-serial**) is a bytestream-oriented, character stream driver. The para-virtualized serial driver provides a simple communication interface between the host's user space and the guest's user space where networking is not be available or unusable.

9.10. BALLOON DRIVER

The balloon driver allows guests to express to the hypervisor how much memory they require. The balloon driver allows the host to efficiently allocate and memory to the guest and allow free memory to be allocated to other guests and processes.

Guests using the balloon driver can mark sections of the guest's RAM as not in use (balloon inflation).

The hypervisor can free the memory and use the memory for other host processes or other guests on that host. When the guest requires the freed memory again, the hypervisor can reallocate RAM to the guest (balloon deflation).

APPENDIX A. ENUMERATED VALUE TRANSLATION

The API uses Red Hat Virtualization Query Language to perform search queries. For more information, see [Searches](#) in the *Introduction to the Administration Portal*.

Note that certain enumerated values in the API require a different search query when using the Query Language. The following tables provides a translation for these key enumerated values according to resource type.

Table A.1. Enumerated Value Translations

API Enumerable Type	API Enumerable Value	Query Language Value
data_center_states	not_operational	notoperational
host_states	non_responsive	nonresponsive
	install_failed	installfailed
	preparing_for_maintenance	preparingformaintenance
	non_operational	nonoperational
	pending_approval	pendingapproval
vm_states	powering_up	poweringup
	powering_down	poweringdown
	migrating	migratingfrom
	migrating	migratingto
	not_responding	notresponding
	wait_for_launch	waitforlaunch
	reboot_in_progress	rebootinprogress
	saving_state	savingstate
	restoring_state	restoringstate

APPENDIX B. EVENT CODES

This table lists all event codes.

Table B.1. Event codes

Code	Name	Severity	Message
0	UNASSIGNED	Info	
1	VDC_START	Info	Starting oVirt Engine.
2	VDC_STOP	Info	Stopping oVirt Engine.
12	VDS_FAILURE	Error	Host \${VdsName} is non responsive.
13	VDS_DETECTED	Info	Status of host \${VdsName} was set to \${HostStatus}.
14	VDS_RECOVER	Info	Host \${VdsName} is rebooting.
15	VDS_MAINTENANCE	Normal	Host \${VdsName} was switched to Maintenance Mode.
16	VDS_ACTIVATE	Info	Activation of host \${VdsName} initiated by \${UserName}.
17	VDS_MAINTENANCE_FAILED	Error	Failed to switch Host \${VdsName} to Maintenance mode.
18	VDS_ACTIVATE_FAILED	Error	Failed to activate Host \${VdsName}. (User: \${UserName}).
19	VDS_RECOVER_FAILED	Error	Host \${VdsName} failed to recover.
20	USER_VDS_START	Info	Host \${VdsName} was started by \${UserName}.

Code	Name	Severity	Message
21	USER_VDS_STOP	Info	Host <code>{VdsName}</code> was stopped by <code>{UserName}</code> .
22	IRS_FAILURE	Error	Failed to access Storage on Host <code>{VdsName}</code> .
23	VDS_LOW_DISK_SPACE	Warning	Warning, Low disk space. Host <code>{VdsName}</code> has less than <code>{DiskSpace}</code> MB of free space left on: <code>{Disks}</code> .
24	VDS_LOW_DISK_SPACE_ERROR	Error	Critical, Low disk space. Host <code>{VdsName}</code> has less than <code>{DiskSpace}</code> MB of free space left on: <code>{Disks}</code> . Low disk space might cause an issue upgrading this host.
25	VDS_NO_SELINUX_ENFORCEMENT	Warning	Host <code>{VdsName}</code> does not enforce SELinux. Current status: <code>{Mode}</code>
26	IRS_DISK_SPACE_LOW	Warning	Warning, Low disk space. <code>{StorageDomainName}</code> domain has <code>{DiskSpace}</code> GB of free space.
27	VDS_STATUS_CHANGE_FAILED_DUE_TO_STOP_SPM_FAILURE	Warning	Failed to change status of host <code>{VdsName}</code> due to a failure to stop the spm.
28	VDS_PROVISION	Warning	Installing OS on Host <code>{VdsName}</code> using Hostgroup <code>{HostGroupName}</code> .

Code	Name	Severity	Message
29	USER_ADD_VM_TEMPLATE_SUCCESS	Info	Template <code>{VmTemplateName}</code> was created successfully.
31	USER_VDC_LOGOUT	Info	User <code>{UserName}</code> connected from <code>{SourceIP}</code> using session <code>{SessionID}</code> logged out.
32	USER_RUN_VM	Info	VM <code>{VmName}</code> started on Host <code>{VdsName}</code>
33	USER_STOP_VM	Info	VM <code>{VmName}</code> powered off by <code>{UserName}</code> (Host: <code>{VdsName}</code>) <code>{OptionalReason}</code> .
34	USER_ADD_VM	Info	VM <code>{VmName}</code> was created by <code>{UserName}</code> .
35	USER_UPDATE_VM	Info	VM <code>{VmName}</code> configuration was updated by <code>{UserName}</code> .
36	USER_ADD_VM_TEMPLATE_FAILURE	Error	Failed creating Template <code>{VmTemplateName}</code> .
37	USER_ADD_VM_STARTED	Info	VM <code>{VmName}</code> creation was initiated by <code>{UserName}</code> .

Code	Name	Severity	Message
38	USER_CHANGE_DISK_VM	Info	CD <code>{DiskName}</code> was inserted to VM <code>{VmName}</code> by <code>{UserName}</code> .
39	USER_PAUSE_VM	Info	VM <code>{VmName}</code> was suspended by <code>{UserName}</code> (Host: <code>{VdsName}</code>).
40	USER_RESUME_VM	Info	VM <code>{VmName}</code> was resumed by <code>{UserName}</code> (Host: <code>{VdsName}</code>).
41	USER_VDS_RESTART	Info	Host <code>{VdsName}</code> was restarted by <code>{UserName}</code> .
42	USER_ADD_VDS	Info	Host <code>{VdsName}</code> was added by <code>{UserName}</code> .
43	USER_UPDATE_VDS	Info	Host <code>{VdsName}</code> configuration was updated by <code>{UserName}</code> .
44	USER_REMOVE_VDS	Info	Host <code>{VdsName}</code> was removed by <code>{UserName}</code> .
45	USER_CREATE_SNAPSHOT	Info	Snapshot <code>{SnapshotName}</code> creation for VM <code>{VmName}</code> was initiated by <code>{UserName}</code> .
46	USER_TRY_BACK_TO_SNAPSHOT	Info	Snapshot-Preview <code>{SnapshotName}</code> for VM <code>{VmName}</code> was initiated by <code>{UserName}</code> .

Code	Name	Severity	Message
47	USER_RESTORE_FROM_SNAPSHOT	Info	VM \${VmName} restored from Snapshot by \${UserName}.
48	USER_ADD_VM_TEMPLATE	Info	Creation of Template \${VmTemplateName} from VM \${VmName} was initiated by \${UserName}.
49	USER_UPDATE_VM_TEMPLATE	Info	Template \${VmTemplateName} configuration was updated by \${UserName}.
50	USER_REMOVE_VM_TEMPLATE	Info	Removal of Template \${VmTemplateName} was initiated by \${UserName}.
51	USER_ADD_VM_TEMPLATE_FINISHED_SUCCESS	Info	Creation of Template \${VmTemplateName} from VM \${VmName} has been completed.
52	USER_ADD_VM_TEMPLATE_FINISHED_FAILURE	Error	Failed to complete creation of Template \${VmTemplateName} from VM \${VmName}.
53	USER_ADD_VM_FINISHED_SUCCESS	Info	VM \${VmName} creation has been completed.
54	USER_FAILED_RUN_VM	Error	Failed to run VM \${VmName}\${DueToError} (User: \${UserName}).
55	USER_FAILED_PAUSE_VM	Error	Failed to suspend VM \${VmName} (Host: \${VdsName}, User: \${UserName}).

Code	Name	Severity	Message
56	USER_FAILED_STOP_VM	Error	Failed to power off VM <code>{VmName}</code> (Host: <code>{VdsName}</code> , User: <code>{UserName}</code>).
57	USER_FAILED_ADD_VM	Error	Failed to create VM <code>{VmName}</code> (User: <code>{UserName}</code>).
58	USER_FAILED_UPDATE_VM	Error	Failed to update VM <code>{VmName}</code> (User: <code>{UserName}</code>).
59	USER_FAILED_REMOVE_VM	Error	
60	USER_ADD_VM_FINISHED_FAILURE	Error	Failed to complete VM <code>{VmName}</code> creation.
61	VM_DOWN	Info	VM <code>{VmName}</code> is down. <code>{ExitMessage}</code>
62	VM_MIGRATION_START	Info	Migration started (VM: <code>{VmName}</code> , Source: <code>{VdsName}</code> , Destination: <code>{DestinationVdsName}</code> , User: <code>{UserName}</code>). <code>{OptionalReason}</code>
63	VM_MIGRATION_DONE	Info	Migration completed (VM: <code>{VmName}</code> , Source: <code>{VdsName}</code> , Destination: <code>{DestinationVdsName}</code> , Duration: <code>{Duration}</code> , Total: <code>{TotalDuration}</code> , Actual downtime: <code>{ActualDowntime}</code>)

Code	Name	Severity	Message
64	VM_MIGRATION_ABORT	Error	Migration failed: \${MigrationError} (VM: \${VmName} , Source: \${VdsName}).
65	VM_MIGRATION_FAILED	Error	Migration failed \${DueToMigrationError} (VM: \${VmName} , Source: \${VdsName}).
66	VM_FAILURE	Error	VM \${VmName} cannot be found on Host \${VdsName} .
67	VM_MIGRATION_START_SYSTEM_INITIATED	Info	Migration initiated by system (VM: \${VmName} , Source: \${VdsName} , Destination: \${DestinationVdsName} , Reason: \${OptionalReason}).
68	USER_CREATE_SNAPSHOT_FINISHED_SUCCESS	Info	Snapshot '\${SnapshotName}' creation for VM '\${VmName}' has been completed.
69	USER_CREATE_SNAPSHOT_FINISHED_FAILURE	Error	Failed to complete snapshot '\${SnapshotName}' creation for VM '\${VmName}' .
70	USER_RUN_VM_AS_STATELESS_FINISHED_FAILURE	Error	Failed to complete starting of VM \${VmName} .
71	USER_TRY_BACK_TO_SNAPSHOT_FINISHED_SUCCESS	Info	Snapshot-Preview \${SnapshotName} for VM \${VmName} has been completed.

Code	Name	Severity	Message
72	MERGE_SNAPSHOT_S_ON_HOST	Info	Merging snapshots ({SourceSnapshot}) into {DestinationSnapshot} of disk {Disk} on host {VDS}
73	USER_INITIATED_SHUTDOWN_VM	Info	VM shutdown initiated by {UserName} on VM {VmName} (Host: {VdsName}) {OptionalReason} .
74	USER_FAILED_SHUTDOWN_VM	Error	Failed to initiate shutdown on VM {VmName} (Host: {VdsName} , User: {UserName}).
75	VDS_SOFT_RECOVER	Info	Soft fencing on host {VdsName} was successful.
76	USER_STOPPED_VM_INSTEAD_OF_SHUTDOWN	Info	VM {VmName} was powered off ungracefully by {UserName} (Host: {VdsName}) {OptionalReason} .
77	USER_FAILED_STOPPING_VM_INSTEAD_OF_SHUTDOWN	Error	Failed to power off VM {VmName} (Host: {VdsName} , User: {UserName}).
78	USER_ADD_DISK_TO_VM	Info	Add-Disk operation of {DiskAlias} was initiated on VM {VmName} by {UserName} .
79	USER_FAILED_ADD_DISK_TO_VM	Error	Add-Disk operation failed on VM {VmName} (User: {UserName}).
80	USER_REMOVE_DISK_FROM_VM	Info	Disk was removed from VM {VmName} by {UserName} .

Code	Name	Severity	Message
81	USER_FAILED_REMOVE_DISK_FROM_VM	Error	Failed to remove Disk from VM {VmName} (User: {UserName}).
88	USER_UPDATE_VM_DISK	Info	VM {VmName} {DiskAlias} disk was updated by {UserName} .
89	USER_FAILED_UPDATE_VM_DISK	Error	Failed to update VM {VmName} disk {DiskAlias} (User: {UserName}).
90	VDS_FAILED_TO_GET_HOST_HARDWARE_INFO	Warning	Could not get hardware information for host {VdsName}
94	USER_COMMIT_RESTORE_FROM_SNAPSHOT_START	Info	Committing a Snapshot-Preview for VM {VmName} was initialized by {UserName} .
95	USER_COMMIT_RESTORE_FROM_SNAPSHOT_FINISH_SUCCESS	Info	Committing a Snapshot-Preview for VM {VmName} has been completed.
96	USER_COMMIT_RESTORE_FROM_SNAPSHOT_FINISH_FAILURE	Error	Failed to commit Snapshot-Preview for VM {VmName} .
97	USER_ADD_DISK_TO_VM_FINISHED_SUCCESS	Info	The disk {DiskAlias} was successfully added to VM {VmName} .
98	USER_ADD_DISK_TO_VM_FINISHED_FAILURE	Error	Add-Disk operation failed to complete on VM {VmName} .
99	USER_TRY_BACK_TO_SNAPSHOT_FINISH_FAILURE	Error	Failed to complete Snapshot-Preview {SnapshotName} for VM {VmName} .

Code	Name	Severity	Message
100	USER_RESTORE_FROM_SNAPSHOT_FINISH_SUCCESS	Info	VM <code>{VmName}</code> restoring from Snapshot has been completed.
101	USER_RESTORE_FROM_SNAPSHOT_FINISH_FAILURE	Error	Failed to complete restoring from Snapshot of VM <code>{VmName}</code> .
102	USER_FAILED_CHANGE_DISK_VM	Error	Failed to change disk in VM <code>{VmName}</code> (Host: <code>{VdsName}</code> , User: <code>{UserName}</code>).
103	USER_FAILED_RESUME_VM	Error	Failed to resume VM <code>{VmName}</code> (Host: <code>{VdsName}</code> , User: <code>{UserName}</code>).
104	USER_FAILED_ADD_VDS	Error	Failed to add Host <code>{VdsName}</code> (User: <code>{UserName}</code>).
105	USER_FAILED_UPDATE_VDS	Error	Failed to update Host <code>{VdsName}</code> (User: <code>{UserName}</code>).
106	USER_FAILED_REMOVE_VDS	Error	Failed to remove Host <code>{VdsName}</code> (User: <code>{UserName}</code>).
107	USER_FAILED_VDS_RESTART	Error	Failed to restart Host <code>{VdsName}</code> , (User: <code>{UserName}</code>).
108	USER_FAILED_ADD_VM_TEMPLATE	Error	Failed to initiate creation of Template <code>{VmTemplateName}</code> from VM <code>{VmName}</code> (User: <code>{UserName}</code>).

Code	Name	Severity	Message
109	USER_FAILED_UPDATE_VM_TEMPLATE	Error	Failed to update Template <code>{VmTemplateName}</code> (User: <code>{UserName}</code>).
110	USER_FAILED_REMOVE_VM_TEMPLATE	Error	Failed to initiate removal of Template <code>{VmTemplateName}</code> (User: <code>{UserName}</code>).
111	USER_STOP_SUSPENDED_VM	Info	Suspended VM <code>{VmName}</code> has had its save state cleared by <code>{UserName}</code> <code>{OptionalReason}</code> .
112	USER_STOP_SUSPENDED_VM_FAILED	Error	Failed to power off suspended VM <code>{VmName}</code> (User: <code>{UserName}</code>).
113	USER_REMOVE_VM_FINISHED	Info	VM <code>{VmName}</code> was successfully removed.
115	USER_FAILED_TRY_BACK_TO_SNAPSHOT	Error	Failed to preview Snapshot <code>{SnapshotName}</code> for VM <code>{VmName}</code> (User: <code>{UserName}</code>).
116	USER_FAILED_RESTORE_FROM_SNAPSHOT	Error	Failed to restore VM <code>{VmName}</code> from Snapshot (User: <code>{UserName}</code>).
117	USER_FAILED_CREATE_SNAPSHOT	Error	Failed to create Snapshot <code>{SnapshotName}</code> for VM <code>{VmName}</code> (User: <code>{UserName}</code>).
118	USER_FAILED_VDS_START	Error	Failed to start Host <code>{VdsName}</code> , (User: <code>{UserName}</code>).

Code	Name	Severity	Message
119	VM_DOWN_ERROR	Error	VM {VmName} is down with error. {ExitMessage} .
120	VM_MIGRATION_TO_SERVER_FAILED	Error	Migration failed {DueToMigrationError} (VM: {VmName} , Source: {VdsName} , Destination: {DestinationVdsName}).
121	SYSTEM_VDS_RESTART	Info	Host {VdsName} was restarted by the engine.
122	SYSTEM_FAILED_VDS_RESTART	Error	A restart initiated by the engine to Host {VdsName} has failed.
123	VDS_SLOW_STORAGE_RESPONSE_TIME	Warning	Slow storage response time on Host {VdsName} .
124	VM_IMPORT	Info	Started VM import of {ImportedVmName} (User: {UserName})
125	VM_IMPORT_FAILED	Error	Failed to import VM {ImportedVmName} (User: {UserName})
126	VM_NOT_RESPONDING	Warning	VM {VmName} is not responding.
127	VDS_RUN_IN_NO_KVM_MODE	Error	Host {VdsName} running without virtualization hardware acceleration

Code	Name	Severity	Message
128	VM_MIGRATION_TRYING_RERUN	Warning	Failed to migrate VM \${VmName} to Host \${DestinationVdsName}\${DueToMigrationError}. Trying to migrate to another Host.
129	VM_CLEARED	Info	Unused
130	USER_SUSPEND_VM_FINISH_FAILURE_WILL_TRY_AGAIN	Error	Failed to complete suspending of VM \${VmName}, will try again.
131	USER_EXPORT_VM	Info	VM \${VmName} exported to \${ExportPath} by \${UserName}
132	USER_EXPORT_VM_FAILED	Error	Failed to export VM \${VmName} to \${ExportPath} (User: \${UserName})
133	USER_EXPORT_TEMPLATE	Info	Template \${VmTemplateName} exported to \${ExportPath} by \${UserName}
134	USER_EXPORT_TEMPLATE_FAILED	Error	Failed to export Template \${VmTemplateName} to \${ExportPath} (User: \${UserName})
135	TEMPLATE_IMPORT	Info	Started Template import of \${ImportedVmTemplateName} (User: \${UserName})

Code	Name	Severity	Message
136	TEMPLATE_IMPORT_FAILED	Error	Failed to import Template <code>\${ImportedVmTemplateName}</code> (User: <code>\${UserName}</code>)
137	USER_FAILED_VDS_STOP	Error	Failed to stop Host <code>\${VdsName}</code> , (User: <code>\${UserName}</code>).
138	VM_PAUSED_ENOSPC	Error	VM <code>\${VmName}</code> has been paused due to no Storage space error.
139	VM_PAUSED_ERROR	Error	VM <code>\${VmName}</code> has been paused due to unknown storage error.
140	VM_MIGRATION_FAILED_DURING_MOVE_TO_MAINTENANCE	Error	Migration failed <code>\${DueToMigrationError}</code> while Host is in 'preparing for maintenance' state. Consider manual intervention: stopping/migrating Vms as Host's state will not turn to maintenance while VMs are still running on it.(VM: <code>\${VmName}</code> , Source: <code>\${VdsName}</code> , Destination: <code>\${DestinationVdsName}</code>).

Code	Name	Severity	Message
141	VDS_VERSION_NOT_SUPPORTED_FOR_CLUSTER	Error	Host {VdsName} is installed with VDSM version {VdsSupportedVersions} and cannot join cluster {ClusterName} which is compatible with VDSM versions {CompatibilityVersion} .
142	VM_SET_TO_UNKNOWN_STATUS	Warning	VM {VmName} was set to the Unknown status.
143	VM_WAS_SET_DOWN_DUE_TO_HOST_REBOOT_OR_MANUAL_FENCE	Info	Vm {VmName} was shut down due to {VdsName} host reboot or manual fence
144	VM_IMPORT_INFO	Info	Value of field {FieldName} of imported VM {VmName} is {FieldValue} . The field is reset to the default value
145	VM_PAUSED_EIO	Error	VM {VmName} has been paused due to storage I/O problem.
146	VM_PAUSED_PERMISSION	Error	VM {VmName} has been paused due to storage permissions problem.
147	VM_POWER_DOWN_FAILED	Warning	Shutdown of VM {VmName} failed.
148	VM_MEMORY_UNGUARANTEED_VALUE	Error	VM {VmName} on host {VdsName} was guaranteed {MemGuaranteed} MB but currently has {MemActual} MB

Code	Name	Severity	Message
149	USER_ADD	Info	User '\${NewUserName}' was added successfully to the system.
150	USER_INITIATED_RUN_VM	Info	Starting VM \${VmName} was initiated by \${UserName} .
151	USER_INITIATED_RUN_VM_FAILED	Warning	Failed to run VM \${VmName} on Host \${VdsName} .
152	USER_RUN_VM_ON_NON_DEFAULT_VDS	Warning	Guest \${VmName} started on Host \${VdsName} . (Default Host parameter was ignored - assigned Host was not available).
153	USER_STARTED_VM	Info	VM \${VmName} was started by \${UserName} (Host: \${VdsName}).
154	VDS_CLUSTER_VERSION_NOT_SUPPORTED	Error	Host \${VdsName} is compatible with versions (\${VdsSupportedVersions}) and cannot join Cluster \${ClusterName} which is set to version \${CompatibilityVersion} .

Code	Name	Severity	Message
155	VDS_ARCHITECTURE_NOT_SUPPORTED_FOR_CLUSTER	Error	Host \${VdsName} has architecture \${VdsArchitecture} and cannot join Cluster \${ClusterName} which has architecture \${ClusterArchitecture}.
156	CPU_TYPE_UNSUPPORTED_IN_THIS_CLUSTER_VERSION	Error	Host \${VdsName} moved to Non-Operational state as host CPU type is not supported in this cluster compatibility version or is not supported at all
157	USER_REBOOT_VM	Info	User \${UserName} initiated reboot of VM \${VmName}.
158	USER_FAILED_REBOOT_VM	Error	Failed to reboot VM \${VmName} (User: \${UserName}).
159	USER_FORCE_SELECTED_SPM	Info	Host \${VdsName} was force selected by \${UserName}
160	USER_ACCOUNT_DISABLED_OR_LOCKED	Error	User \${UserName} cannot login, as it got disabled or locked. Please contact the system administrator.
161	VM_CANCEL_MIGRATION	Info	Migration cancelled (VM: \${VmName}, Source: \${VdsName}, User: \${UserName}).
162	VM_CANCEL_MIGRATION_FAILED	Error	Failed to cancel migration for VM: \${VmName}

Code	Name	Severity	Message
163	VM_STATUS_RESTORED	Info	VM <code>{VmName}</code> status was restored to <code>{VmStatus}</code> .
164	VM_SET_TICKET	Info	User <code>{UserName}</code> initiated console session for VM <code>{VmName}</code>
165	VM_SET_TICKET_FAILED	Error	User <code>{UserName}</code> failed to initiate a console session for VM <code>{VmName}</code>
166	VM_MIGRATION_NO_VDS_TO_MIGRATE_TO	Warning	No available host was found to migrate VM <code>{VmName}</code> to.
167	VM_CONSOLE_CONNECTED	Info	User <code>{UserName}</code> is connected to VM <code>{VmName}</code> .
168	VM_CONSOLE_DISCONNECTED	Info	User <code>{UserName}</code> got disconnected from VM <code>{VmName}</code> .
169	VM_FAILED_TO_PRESTART_IN_POOL	Warning	Cannot pre-start VM in pool <code>'{VmPoolName}'</code> . The system will continue trying.
170	USER_CREATE_LIVE_SNAPSHOT_FINISHED_FAILURE	Warning	Failed to create live snapshot <code>'{SnapshotName}'</code> for VM <code>'{VmName}'</code> . VM restart is recommended. Note that using the created snapshot might cause data inconsistency.

Code	Name	Severity	Message
171	USER_RUN_VM_AS_STATELESS_WITH_DISKS_NOT_ALLOWING_SNAPSHOT	Warning	VM {VmName} was run as stateless with one or more of disks that do not allow snapshots (User: {UserName}).
172	USER_REMOVE_VM_FINISHED_WITH_ILLEGAL_DISKS	Warning	VM {VmName} has been removed, but the following disks could not be removed: {DisksNames} . These disks will appear in the main disks tab in illegal state, please remove manually when possible.
173	USER_CREATE_LIVE_SNAPSHOT_NO_MEMORY_FAILURE	Error	Failed to save memory as part of Snapshot {SnapshotName} for VM {VmName} (User: {UserName}).
174	VM_IMPORT_FROM_CONFIGURATION_EXECUTED_SUCCESSFULLY	Info	VM {VmName} has been successfully imported from the given configuration.
175	VM_IMPORT_FROM_CONFIGURATION_ATTACH_DISKS_FAILED	Warning	VM {VmName} has been imported from the given configuration but the following disk(s) failed to attach: {DiskAliases} .
176	VM_BALLOON_DRIVER_ERROR	Error	The Balloon driver on VM {VmName} on host {VdsName} is requested but unavailable.

Code	Name	Severity	Message
177	VM_BALLOON_DRIVER_UNCONTROLLED	Error	The Balloon device on VM <code>{VmName}</code> on host <code>{VdsName}</code> is inflated but the device cannot be controlled (guest agent is down).
178	VM_MEMORY_NOT_IN_RECOMMENDED_RANGE	Warning	VM <code>{VmName}</code> was configured with <code>{VmMemInMb}</code> MiB of memory while the recommended value range is <code>{VmMinMemInMb}</code> MiB - <code>{VmMaxMemInMb}</code> MiB
179	USER_INITIATED_RUN_VM_AND_PAUSED	Info	Starting in paused mode VM <code>{VmName}</code> was initiated by <code>{UserName}</code> .
180	TEMPLATE_IMPORT_FROM_CONFIGURATION_SUCCESS	Info	Template <code>{VmTemplateName}</code> has been successfully imported from the given configuration.
181	TEMPLATE_IMPORT_FROM_CONFIGURATION_FAILED	Error	Failed to import Template <code>{VmTemplateName}</code> from the given configuration.
182	USER_FAILED_ATTACH_USER_TO_VM	Error	Failed to attach User <code>{AdUserName}</code> to VM <code>{VmName}</code> (User: <code>{UserName}</code>).
183	USER_ATTACH_TAG_TO_TEMPLATE	Info	Tag <code>{TagName}</code> was attached to Template(s) <code>{TemplatesNames}</code> by <code>{UserName}</code> .

Code	Name	Severity	Message
184	USER_ATTACH_TAG_TO_TEMPLATE_FAILED	Error	Failed to attach Tag \${TagName} to Template(s) \${TemplatesNames} (User: \${UserName}).
185	USER_DETACH_TEMPLATE_FROM_TAG	Info	Tag \${TagName} was detached from Template(s) \${TemplatesNames} by \${UserName} .
186	USER_DETACH_TEMPLATE_FROM_TAG_FAILED	Error	Failed to detach Tag \${TagName} from TEMPLATE(s) \${TemplatesNames} (User: \${UserName}).
187	VDS_STORAGE_CONNECTION_FAILED_BUT_LAST_VDS	Error	Failed to connect Host \${VdsName} to Data Center, due to connectivity errors with the Storage. Host \${VdsName} will remain in Up state (but inactive), as it is the last Host in the Data Center, to enable manual intervention by the Administrator.
188	VDS_STORAGE_CONNECTION_FAILED	Error	Failed to connect Host \${VdsName} to the Storage Domains \${failedStorageDomains} .
189	VDS_STORAGE_VDS_STATS_FAILED	Error	Host \${VdsName} reports about one of the Active Storage Domains as Problematic.

Code	Name	Severity	Message
190	UPDATE_OVF_FOR_STORAGE_DOMAIN_FAILED	Warning	Failed to update VMs/Templates OVF data for Storage Domain <code>\${StorageDomainName}</code> in Data Center <code>\${StoragePoolName}</code> .
191	CREATE_OVF_STORE_FOR_STORAGE_DOMAIN_FAILED	Warning	Failed to create OVF store disk for Storage Domain <code>\${StorageDomainName}</code> . The Disk with the id <code>\${DiskId}</code> might be removed manually for automatic attempt to create new one. OVF updates won't be attempted on the created disk.
192	CREATE_OVF_STORE_FOR_STORAGE_DOMAIN_INITIATE_FAILED	Warning	Failed to create OVF store disk for Storage Domain <code>\${StorageDomainName}</code> . OVF data won't be updated meanwhile for that domain.
193	DELETE_OVF_STORE_FOR_STORAGE_DOMAIN_FAILED	Warning	Failed to delete the OVF store disk for Storage Domain <code>\${StorageDomainName}</code> . In order to detach the domain please remove it manually or try to detach the domain again for another attempt.
194	VM_CANCEL_CONVERSION	Info	Conversion cancelled (VM: <code>\${VmName}</code> , Source: <code>\${VdsName}</code> , User: <code>\${UserName}</code>).

Code	Name	Severity	Message
195	VM_CANCEL_CONVERSION_FAILED	Error	Failed to cancel conversion for VM: \${VmName}
196	VM_RECOVERED_FROM_PAUSE_ERROR	Normal	VM \${VmName} has recovered from paused back to up.
197	SYSTEM_SSH_HOST_RESTART	Info	Host \${VdsName} was restarted using SSH by the engine.
198	SYSTEM_FAILED_SSH_HOST_RESTART	Error	A restart using SSH initiated by the engine to Host \${VdsName} has failed.
199	USER_UPDATE_OVF_STORE	Info	OVF_STORE for domain \${StorageDomainName} was updated by \${UserName}.
200	IMPORTEXPORTE_GET_VMS_INFO_FAILED	Error	Failed to retrieve VM/Templates information from export domain \${StorageDomainName}
201	IRS_DISK_SPACE_LOW_ERROR	Error	Critical, Low disk space. \${StorageDomainName} domain has \${DiskSpace} GB of free space.
202	IMPORTEXPORTE_GET_EXTERNAL_VMS_INFO_FAILED	Error	Failed to retrieve VMs information from external server \${URL}

Code	Name	Severity	Message
204	IRS_HOSTED_ON_VDS	Info	Storage Pool Manager runs on Host <code>\${VdsName}</code> (Address: <code>\${ServerIp}</code>), Data Center <code>\${StoragePoolName}</code> .
205	PROVIDER_ADDED	Info	Provider <code>\${ProviderName}</code> was added. (User: <code>\${UserName}</code>)
206	PROVIDER_ADDITION_FAILED	Error	Failed to add provider <code>\${ProviderName}</code> . (User: <code>\${UserName}</code>)
207	PROVIDER_UPDATED	Info	Provider <code>\${ProviderName}</code> was updated. (User: <code>\${UserName}</code>)
208	PROVIDER_UPDATE_FAILED	Error	Failed to update provider <code>\${ProviderName}</code> . (User: <code>\${UserName}</code>)
209	PROVIDER_REMOVED	Info	Provider <code>\${ProviderName}</code> was removed. (User: <code>\${UserName}</code>)
210	PROVIDER_REMOVAL_FAILED	Error	Failed to remove provider <code>\${ProviderName}</code> . (User: <code>\${UserName}</code>)
213	PROVIDER_CERTIFICATE_IMPORTED	Info	Certificate for provider <code>\${ProviderName}</code> was imported. (User: <code>\${UserName}</code>)
214	PROVIDER_CERTIFICATE_IMPORT_FAILED	Error	Failed importing Certificate for provider <code>\${ProviderName}</code> . (User: <code>\${UserName}</code>)

Code	Name	Severity	Message
215	PROVIDER_SYNCHRONIZED	Info	
216	PROVIDER_SYNCHRONIZED_FAILED	Error	Failed to synchronize networks of Provider \${ProviderName}.
217	PROVIDER_SYNCHRONIZED_PERFORMED	Info	Networks of Provider \${ProviderName} were successfully synchronized.
218	PROVIDER_SYNCHRONIZED_PERFORMED_FAILED	Error	Networks of Provider \${ProviderName} were incompletely synchronized.
219	PROVIDER_SYNCHRONIZED_DISABLED	Error	Failed to synchronize networks of Provider \${ProviderName}, because the authentication information of the provider is invalid. Automatic synchronization is deactivated for this Provider.
250	USER_UPDATE_VM_CLUSTER_DEFAULT_HOST_CLEARED	Info	\${VmName} cluster was updated by \${UserName}, Default host was reset to auto assign.
251	USER_REMOVE_VM_TEMPLATE_FINISHED	Info	Removal of Template \${VmTemplateName} has been completed.
252	SYSTEM_FAILED_UPDATE_VM	Error	Failed to Update VM \${VmName} that was initiated by system.
253	SYSTEM_UPDATE_VM	Info	VM \${VmName} configuration was updated by system.

Code	Name	Severity	Message
254	VM_ALREADY_IN_REQUESTED_STATUS	Info	VM \${VmName} is already \${VmStatus} , \${Action} was skipped. User: \${UserName} .
302	USER_ADD_VM_POOL_WITH_VMS	Info	VM Pool \${VmPoolName} (containing \${VmsCount} VMs) was created by \${UserName} .
303	USER_ADD_VM_POOL_WITH_VMS_FAILED	Error	Failed to create VM Pool \${VmPoolName} (User: \${UserName}).
304	USER_REMOVE_VM_POOL	Info	VM Pool \${VmPoolName} was removed by \${UserName} .
305	USER_REMOVE_VM_POOL_FAILED	Error	Failed to remove VM Pool \${VmPoolName} (User: \${UserName}).
306	USER_ADD_VM_TO_POOL	Info	VM \${VmName} was added to VM Pool \${VmPoolName} by \${UserName} .
307	USER_ADD_VM_TO_POOL_FAILED	Error	Failed to add VM \${VmName} to VM Pool \${VmPoolName} (User: \${UserName}).
308	USER_REMOVE_VM_FROM_POOL	Info	VM \${VmName} was removed from VM Pool \${VmPoolName} by \${UserName} .

Code	Name	Severity	Message
309	USER_REMOVE_VM_FROM_POOL_FAILED	Error	Failed to remove VM \${VmName} from VM Pool \${VmPoolName} (User: \${UserName}).
310	USER_ATTACH_USER_TO_POOL	Info	User \${AdUserName} was attached to VM Pool \${VmPoolName} by \${UserName} .
311	USER_ATTACH_USER_TO_POOL_FAILED	Error	Failed to attach User \${AdUserName} to VM Pool \${VmPoolName} (User: \${UserName}).
312	USER_DETACH_USER_FROM_POOL	Info	User \${AdUserName} was detached from VM Pool \${VmPoolName} by \${UserName} .
313	USER_DETACH_USER_FROM_POOL_FAILED	Error	Failed to detach User \${AdUserName} from VM Pool \${VmPoolName} (User: \${UserName}).
314	USER_UPDATE_VM_POOL	Info	VM Pool \${VmPoolName} configuration was updated by \${UserName} .
315	USER_UPDATE_VM_POOL_FAILED	Error	Failed to update VM Pool \${VmPoolName} configuration (User: \${UserName}).

Code	Name	Severity	Message
316	USER_ATTACH_USER_TO_VM_FROM_POOL	Info	Attaching User <code>{AdUserName}</code> to VM <code>{VmName}</code> in VM Pool <code>{VmPoolName}</code> was initiated by <code>{UserName}</code> .
317	USER_ATTACH_USER_TO_VM_FROM_POOL_FAILED	Error	Failed to attach User <code>{AdUserName}</code> to VM from VM Pool <code>{VmPoolName}</code> (User: <code>{UserName}</code>).
318	USER_ATTACH_USER_TO_VM_FROM_POOL_FINISHED_SUCCESS	Info	User <code>{AdUserName}</code> successfully attached to VM <code>{VmName}</code> in VM Pool <code>{VmPoolName}</code> .
319	USER_ATTACH_USER_TO_VM_FROM_POOL_FINISHED_FAILURE	Error	Failed to attach user <code>{AdUserName}</code> to VM <code>{VmName}</code> in VM Pool <code>{VmPoolName}</code> .
320	USER_ADD_VM_POOL_WITH_VMS_ADD_VDS_FAILED	Error	Pool <code>{VmPoolName}</code> Created, but some Vms failed to create (User: <code>{UserName}</code>).
321	USER_REMOVE_VM_POOL_INITIATED	Info	VM Pool <code>{VmPoolName}</code> removal was initiated by <code>{UserName}</code> .
325	USER_REMOVE_AD_USER	Info	User <code>{AdUserName}</code> was removed by <code>{UserName}</code> .

Code	Name	Severity	Message
326	USER_FAILED_REMOVE_ADUSER	Error	Failed to remove User <code>\${AdUserName}</code> (User: <code>\${UserName}</code>).
327	USER_FAILED_ADD_ADUSER	Warning	Failed to add User <code>'\${NewUserName}'</code> to the system.
342	USER_REMOVE_SNAPSHOT	Info	Snapshot <code>'\${SnapshotName}'</code> deletion for VM <code>'\${VmName}'</code> was initiated by <code>\${UserName}</code> .
343	USER_FAILED_REMOVE_SNAPSHOT	Error	Failed to remove Snapshot <code>\${SnapshotName}</code> for VM <code>\${VmName}</code> (User: <code>\${UserName}</code>).
344	USER_UPDATE_VM_POOL_WITH_VMS	Info	VM Pool <code>\${VmPoolName}</code> was updated by <code>\${UserName}</code> , <code>\${VmsCount}</code> VMs were added.
345	USER_UPDATE_VM_POOL_WITH_VMS_FAILED	Error	Failed to update VM Pool <code>\${VmPoolName}</code> (User: <code>\${UserName}</code>).
346	USER_PASSWORD_CHANGED	Info	Password changed successfully for <code>\${UserName}</code>
347	USER_PASSWORD_CHANGE_FAILED	Error	Failed to change password. (User: <code>\${UserName}</code>)

Code	Name	Severity	Message
348	USER_CLEAR_UNKNOWN_VMS	Info	All VMs' status on Non Responsive Host \${VdsName} were changed to 'Down' by \${UserName}
349	USER_FAILED_CLEAR_UNKNOWN_VMS	Error	Failed to clear VMs' status on Non Responsive Host \${VdsName}. (User: \${UserName}).
350	USER_ADD_BOOKMARK	Info	Bookmark \${BookmarkName} was added by \${UserName}.
351	USER_ADD_BOOKMARK_FAILED	Error	Failed to add bookmark: \${BookmarkName} (User: \${UserName}).
352	USER_UPDATE_BOOKMARK	Info	Bookmark \${BookmarkName} was updated by \${UserName}.
353	USER_UPDATE_BOOKMARK_FAILED	Error	Failed to update bookmark: \${BookmarkName} (User: \${UserName})
354	USER_REMOVE_BOOKMARK	Info	Bookmark \${BookmarkName} was removed by \${UserName}.
355	USER_REMOVE_BOOKMARK_FAILED	Error	Failed to remove bookmark \${BookmarkName} (User: \${UserName})
356	USER_REMOVE_SNAPSHOT_FINISHED_SUCCESS	Info	Snapshot '\${SnapshotName}' deletion for VM '\${VmName}' has been completed.

Code	Name	Severity	Message
357	USER_REMOVE_SNAPSHOT_FINISHED_FAILURE	Error	Failed to delete snapshot '{SnapshotName}' for VM '{VmName}'.
358	USER_VM_POOL_MAXIMUM_SUBSEQUENT_FAILURES_REACHED	Warning	Not all VMs were successfully created in VM Pool '{VmPoolName}'.
359	USER_REMOVE_SNAPSHOT_FINISHED_FAILURE_PARTIAL_SNAPSHOT	Warning	Due to partial snapshot removal, Snapshot '{SnapshotName}' of VM '{VmName}' now contains only the following disks: '{DiskAliases}'.
360	USER_DETACH_USER_FROM_VM	Info	User '{AdUserName}' was detached from VM '{VmName}' by '{UserName}'.
361	USER_FAILED_DETACH_USER_FROM_VM	Error	Failed to detach User '{AdUserName}' from VM '{VmName}' (User: '{UserName}').
362	USER_REMOVE_SNAPSHOT_FINISHED_FAILURE_BASE_IMAGE_NOT_FOUND	Error	Failed to merge images of snapshot '{SnapshotName}': base volume '{BaseVolumeId}' is missing. This may have been caused by a failed attempt to remove the parent snapshot; if this is the case, please retry deletion of the parent snapshot before deleting this one.

Code	Name	Severity	Message
370	USER_EXTEND_DISK_SIZE_FAILURE	Error	Failed to extend size of the disk <code>'\${DiskAlias}'</code> to <code>\${NewSize}</code> GB, User: <code>\${UserName}</code> .
371	USER_EXTEND_DISK_SIZE_SUCCESS	Info	Size of the disk <code>'\${DiskAlias}'</code> was successfully updated to <code>\${NewSize}</code> GB by <code>\${UserName}</code> .
372	USER_EXTEND_DISK_SIZE_UPDATE_VM_FAILURE	Warning	Failed to update VM <code>'\${VmName}'</code> with the new volume size. VM restart is recommended.
373	USER_REMOVE_DISK_SNAPSHOT	Info	Disk <code>'\${DiskAlias}'</code> from Snapshot(s) <code>'\${Snapshots}'</code> of VM <code>'\${VmName}'</code> deletion was initiated by <code>\${UserName}</code> .
374	USER_FAILED_REMOVE_DISK_SNAPSHOT	Error	Failed to delete Disk <code>'\${DiskAlias}'</code> from Snapshot(s) <code>\${Snapshots}</code> of VM <code>\${VmName}</code> (User: <code>\${UserName}</code>).
375	USER_REMOVE_DISK_SNAPSHOT_FINISHED_SUCCESS	Info	Disk <code>'\${DiskAlias}'</code> from Snapshot(s) <code>'\${Snapshots}'</code> of VM <code>'\${VmName}'</code> deletion has been completed (User: <code>\${UserName}</code>).
376	USER_REMOVE_DISK_SNAPSHOT_FINISHED_FAILURE	Error	Failed to complete deletion of Disk <code>'\${DiskAlias}'</code> from snapshot(s) <code>'\${Snapshots}'</code> of VM <code>'\${VmName}'</code> (User: <code>\${UserName}</code>).

Code	Name	Severity	Message
377	USER_EXTENDED_DISK_SIZE	Info	Extending disk '{DiskAlias}' to \${NewSize} GB was initiated by \${UserName}.
378	USER_REGISTER_DISK_FINISHED_SUCCESS	Info	Disk '{DiskAlias}' has been successfully registered as a floating disk.
379	USER_REGISTER_DISK_FINISHED_FAILURE	Error	Failed to register Disk '{DiskAlias}'.
380	USER_EXTEND_DISK_SIZE_UPDATE_HOST_FAILURE	Warning	Failed to refresh volume size on host '{VdsName}'. Please try the operation again.
381	USER_REGISTER_DISK_INITIATED	Info	Registering Disk '{DiskAlias}' has been initiated.
382	USER_REDUCE_DISK_FINISHED_SUCCESS	Info	Disk '{DiskAlias}' has been successfully reduced.
383	USER_REDUCE_DISK_FINISHED_FAILURE	Error	Failed to reduce Disk '{DiskAlias}'.
400	USER_ATTACH_VM_TO_AD_GROUP	Info	Group \${GroupName} was attached to VM \${VmName} by \${UserName}.
401	USER_ATTACH_VM_TO_AD_GROUP_FAILED	Error	Failed to attach Group \${GroupName} to VM \${VmName} (User: \${UserName}).

Code	Name	Severity	Message
402	USER_DETACH_VM_TO_AD_GROUP	Info	Group <code>\${GroupName}</code> was detached from VM <code>\${VmName}</code> by <code>\${UserName}</code> .
403	USER_DETACH_VM_TO_AD_GROUP_FAILED	Error	Failed to detach Group <code>\${GroupName}</code> from VM <code>\${VmName}</code> (User: <code>\${UserName}</code>).
404	USER_ATTACH_VM_POOL_TO_AD_GROUP	Info	Group <code>\${GroupName}</code> was attached to VM Pool <code>\${VmPoolName}</code> by <code>\${UserName}</code> .
405	USER_ATTACH_VM_POOL_TO_AD_GROUP_FAILED	Error	Failed to attach Group <code>\${GroupName}</code> to VM Pool <code>\${VmPoolName}</code> (User: <code>\${UserName}</code>).
406	USER_DETACH_VM_POOL_TO_AD_GROUP	Info	Group <code>\${GroupName}</code> was detached from VM Pool <code>\${VmPoolName}</code> by <code>\${UserName}</code> .
407	USER_DETACH_VM_POOL_TO_AD_GROUP_FAILED	Error	Failed to detach Group <code>\${GroupName}</code> from VM Pool <code>\${VmPoolName}</code> (User: <code>\${UserName}</code>).
408	USER_REMOVE_AD_GROUP	Info	Group <code>\${GroupName}</code> was removed by <code>\${UserName}</code> .

Code	Name	Severity	Message
409	USER_REMOVE_AD_GROUP_FAILED	Error	Failed to remove group \${GroupName} (User: \${UserName}).
430	USER_UPDATE_TAG	Info	Tag \${TagName} configuration was updated by \${UserName} .
431	USER_UPDATE_TAG_FAILED	Error	Failed to update Tag \${TagName} (User: \${UserName}).
432	USER_ADD_TAG	Info	New Tag \${TagName} was created by \${UserName} .
433	USER_ADD_TAG_FAILED	Error	Failed to create Tag named \${TagName} (User: \${UserName}).
434	USER_REMOVE_TAG	Info	Tag \${TagName} was removed by \${UserName} .
435	USER_REMOVE_TAG_FAILED	Error	Failed to remove Tag \${TagName} (User: \${UserName}).
436	USER_ATTACH_TAG_TO_USER	Info	Tag \${TagName} was attached to User(s) \${AttachUserNames} by \${UserName} .
437	USER_ATTACH_TAG_TO_USER_FAILED	Error	Failed to attach Tag \${TagName} to User(s) \${AttachUserNames} (User: \${UserName}).

Code	Name	Severity	Message
438	USER_ATTACH_TAG_TO_USER_GROUP	Info	Tag <code>\${TagName}</code> was attached to Group(s) <code>\${AttachGroupsNames}</code> by <code>\${UserName}</code> .
439	USER_ATTACH_TAG_TO_USER_GROUP_FAILED	Error	Failed to attach Group(s) <code>\${AttachGroupsNames}</code> to Tag <code>\${TagName}</code> (User: <code>\${UserName}</code>).
440	USER_ATTACH_TAG_TO_VM	Info	Tag <code>\${TagName}</code> was attached to VM(s) <code>\${VmsNames}</code> by <code>\${UserName}</code> .
441	USER_ATTACH_TAG_TO_VM_FAILED	Error	Failed to attach Tag <code>\${TagName}</code> to VM(s) <code>\${VmsNames}</code> (User: <code>\${UserName}</code>).
442	USER_ATTACH_TAG_TO_VDS	Info	Tag <code>\${TagName}</code> was attached to Host(s) <code>\${VdsNames}</code> by <code>\${UserName}</code> .
443	USER_ATTACH_TAG_TO_VDS_FAILED	Error	Failed to attach Tag <code>\${TagName}</code> to Host(s) <code>\${VdsNames}</code> (User: <code>\${UserName}</code>).
444	USER_DETACH_VDS_FROM_TAG	Info	Tag <code>\${TagName}</code> was detached from Host(s) <code>\${VdsNames}</code> by <code>\${UserName}</code> .
445	USER_DETACH_VDS_FROM_TAG_FAILED	Error	Failed to detach Tag <code>\${TagName}</code> from Host(s) <code>\${VdsNames}</code> (User: <code>\${UserName}</code>).

Code	Name	Severity	Message
446	USER_DETACH_VM_FROM_TAG	Info	Tag \${TagName} was detached from VM(s) \${VmsNames} by \${UserName}.
447	USER_DETACH_VM_FROM_TAG_FAILED	Error	Failed to detach Tag \${TagName} from VM(s) \${VmsNames} (User: \${UserName}).
448	USER_DETACH_USER_FROM_TAG	Info	Tag \${TagName} detached from User(s) \${DetachUserNames} by \${UserName}.
449	USER_DETACH_USER_FROM_TAG_FAILED	Error	Failed to detach Tag \${TagName} from User(s) \${DetachUserNames} (User: \${UserName}).
450	USER_DETACH_USER_GROUP_FROM_TAG	Info	Tag \${TagName} was detached from Group(s) \${DetachGroupNames} by \${UserName}.
451	USER_DETACH_USER_GROUP_FROM_TAG_FAILED	Error	Failed to detach Tag \${TagName} from Group(s) \${DetachGroupNames} (User: \${UserName}).
452	USER_ATTACH_TAG_TO_USER_EXISTS	Warning	Tag \${TagName} already attached to User(s) \${AttachUserNamesExists}.
453	USER_ATTACH_TAG_TO_USER_GROUP_EXISTS	Warning	Tag \${TagName} already attached to Group(s) \${AttachGroupNamesExists}.

Code	Name	Severity	Message
454	USER_ATTACH_TAG_TO_VM_EXISTS	Warning	Tag \${TagName} already attached to VM(s) \${VmsNamesExists}.
455	USER_ATTACH_TAG_TO_VDS_EXISTS	Warning	Tag \${TagName} already attached to Host(s) \${VdsNamesExists}.
456	USER_LOGGED_IN_VM	Info	User \${GuestUser} logged in to VM \${VmName}.
457	USER_LOGGED_OUT_VM	Info	User \${GuestUser} logged out from VM \${VmName}.
458	USER_LOCKED_VM	Info	User \${GuestUser} locked VM \${VmName}.
459	USER_UNLOCKED_VM	Info	User \${GuestUser} unlocked VM \${VmName}.
460	USER_ATTACH_TAG_TO_TEMPLATE_EXISTS	Warning	Tag \${TagName} already attached to Template(s) \${TemplatesNamesExists}.
467	UPDATE_TAGS_VM_DEFAULT_DISPLAY_TYPE	Info	Vm \${VmName} tag default display type was updated
468	UPDATE_TAGS_VM_DEFAULT_DISPLAY_TYPE_FAILED	Info	Failed to update Vm \${VmName} tag default display type
470	USER_ATTACH_VM_POOL_TO_AD_GROUP_INTERNAL	Info	Group \${GroupName} was attached to VM Pool \${VmPoolName}.

Code	Name	Severity	Message
471	USER_ATTACH_VM_POOL_TO_AD_GROUP_FAILED_INTERNAL	Error	Failed to attach Group <code>{GroupName}</code> to VM Pool <code>{VmPoolName}</code> .
472	USER_ATTACH_USER_TO_POOL_INTERNAL	Info	User <code>{AdUserName}</code> was attached to VM Pool <code>{VmPoolName}</code> .
473	USER_ATTACH_USER_TO_POOL_FAILED_INTERNAL	Error	Failed to attach User <code>{AdUserName}</code> to VM Pool <code>{VmPoolName}</code> (User: <code>{UserName}</code>).
493	VDS_ALREADY_IN_REQUESTED_STATUS	Warning	Host <code>{HostName}</code> is already <code>{AgentStatus}</code> , Power Management <code>{Operation}</code> operation skipped.
494	VDS_MANUAL_FENCE_STATUS	Info	Manual fence for host <code>{VdsName}</code> was started.
495	VDS_MANUAL_FENCE_STATUS_FAILED	Error	Manual fence for host <code>{VdsName}</code> failed.
496	VDS_FENCE_STATUS	Info	Host <code>{VdsName}</code> power management was verified successfully.
497	VDS_FENCE_STATUS_FAILED	Error	Failed to verify Host <code>{VdsName}</code> power management.
498	VDS_APPROVE	Info	Host <code>{VdsName}</code> was successfully approved by user <code>{UserName}</code> .
499	VDS_APPROVE_FAILED	Error	Failed to approve Host <code>{VdsName}</code> .

Code	Name	Severity	Message
500	VDS_FAILED_TO_RUN_VMS	Error	Host \${VdsName} will be switched to Error status for \${Time} minutes because it failed to run a VM.
501	USER_SUSPEND_VM	Info	Suspending VM \${VmName} was initiated by User \${UserName} (Host: \${VdsName}).
502	USER_FAILED_SUSPEND_VM	Error	Failed to suspend VM \${VmName} (Host: \${VdsName}).
503	USER_SUSPEND_VM_OK	Info	VM \${VmName} on Host \${VdsName} is suspended.
504	VDS_INSTALL	Info	Host \${VdsName} installed
505	VDS_INSTALL_FAILED	Error	Host \${VdsName} installation failed. \${FailedInstallMessage}.
506	VDS_INITIATED_RUN_VM	Info	Trying to restart VM \${VmName} on Host \${VdsName}
509	VDS_INSTALL_IN_PROGRESS	Info	Installing Host \${VdsName}. \${Message}.
510	VDS_INSTALL_IN_PROGRESS_WARNING	Warning	Host \${VdsName} installation in progress . \${Message}.
511	VDS_INSTALL_IN_PROGRESS_ERROR	Error	An error has occurred during installation of Host \${VdsName}: \${Message}.

Code	Name	Severity	Message
512	USER_SUSPEND_VM_FINISH_SUCCESS	Info	Suspending VM {VmName} has been completed.
513	VDS_RECOVER_FAILED_VMS_UNKNOWN	Error	Host {VdsName} cannot be reached, VMs state on this host are marked as Unknown.
514	VDS_INITIALIZING	Warning	Host {VdsName} is initializing. Message: {ErrorMessage}
515	VDS_CPU_LOWER_THAN_CLUSTER	Warning	Host {VdsName} moved to Non-Operational state as host does not meet the cluster's minimum CPU level. Missing CPU features : {CpuFlags}
516	VDS_CPU_RETRIEVE_FAILED	Warning	Failed to determine Host {VdsName} CPU level - could not retrieve CPU flags.
517	VDS_SET_NONOPERATIONAL	Info	Host {VdsName} moved to Non-Operational state.
518	VDS_SET_NONOPERATIONAL_FAILED	Error	Failed to move Host {VdsName} to Non-Operational state.
519	VDS_SET_NONOPERATIONAL_NETWORK	Warning	Host {VdsName} does not comply with the cluster {ClusterName} networks, the following networks are missing on host: '{Networks}'

Code	Name	Severity	Message
520	USER_ATTACH_USER_TO_VM	Info	User <code>\${AdUserName}</code> was attached to VM <code>\${VmName}</code> by <code>\${UserName}</code> .
521	USER_SUSPEND_VM_FINISH_FAILURE	Error	Failed to complete suspending of VM <code>\${VmName}</code> .
522	VDS_SET_NONOPERATIONAL_DOMAIN	Warning	Host <code>\${VdsName}</code> cannot access the Storage Domain(s) <code>\${StorageDomainNames}</code> attached to the Data Center <code>\${StoragePoolName}</code> . Setting Host state to Non-Operational.
523	VDS_SET_NONOPERATIONAL_DOMAIN_FAILED	Error	Host <code>\${VdsName}</code> cannot access the Storage Domain(s) <code>\${StorageDomainNames}</code> attached to the Data Center <code>\${StoragePoolName}</code> . Failed to set Host state to Non-Operational.
524	VDS_DOMAIN_DELAY_INTERVAL	Warning	Storage domain <code>\${StorageDomainName}</code> experienced a high latency of <code>\${Delay}</code> seconds from host <code>\${VdsName}</code> . This may cause performance and functional issues. Please consult your Storage Administrator.
525	VDS_INITIATED_RUN_AS_STATELESS_VM_NOT_YET_RUNNING	Info	Starting VM <code>\${VmName}</code> as stateless was initiated.

Code	Name	Severity	Message
528	USER_EJECT_VM_DISK	Info	CD was ejected from VM \${VmName} by \${UserName}.
530	VDS_MANUAL_FENCE_FAILED_CALL_FENCE_SPM	Warning	Manual fence did not revoke the selected SPM (\${VdsName}) since the master storage domain\n was not active or could not use another host for the fence operation.
531	VDS_LOW_MEM	Warning	Available memory of host \${HostName} in cluster \${Cluster} [\${AvailableMemory} MB] is under defined threshold [\${Threshold} MB].
532	VDS_HIGH_MEM_USAGE	Warning	Used memory of host \${HostName} in cluster \${Cluster} [\${UsedMemory}%] exceeded defined threshold [\${Threshold}%].
533	VDS_HIGH_NETWORK_USE	Warning	
534	VDS_HIGH_CPU_USAGE	Warning	Used CPU of host \${HostName} [\${UsedCpu}%] exceeded defined threshold [\${Threshold}%].
535	VDS_HIGH_SWAP_USAGE	Warning	Used swap memory of host \${HostName} [\${UsedSwap}%] exceeded defined threshold [\${Threshold}%].

Code	Name	Severity	Message
536	VDS_LOW_SWAP	Warning	Available swap memory of host <code>{HostName}</code> [<code>{AvailableSwapMemory}</code> MB] is under defined threshold [<code>{Threshold}</code> MB].
537	VDS_INITIATED_RUN_VM_AS_STATELESS	Info	VM <code>{VmName}</code> was restarted on Host <code>{VdsName}</code> as stateless
538	USER_RUN_VM_AS_STATELESS	Info	VM <code>{VmName}</code> started on Host <code>{VdsName}</code> as stateless
539	VDS_AUTO_FENCE_STATUS	Info	Auto fence for host <code>{VdsName}</code> was started.
540	VDS_AUTO_FENCE_STATUS_FAILED	Error	Auto fence for host <code>{VdsName}</code> failed.
541	VDS_AUTO_FENCE_FAILED_CALL_FENCE_SPM	Warning	Auto fence did not revoke the selected SPM (<code>{VdsName}</code>) since the master storage domain was not active or could not use another host for the fence operation.
550	VDS_PACKAGES_IN_PROGRESS	Info	Package update Host <code>{VdsName}</code> . <code>{Message}</code> .
551	VDS_PACKAGES_IN_PROGRESS_WARNING	Warning	Host <code>{VdsName}</code> update packages in progress . <code>{Message}</code> .
552	VDS_PACKAGES_IN_PROGRESS_ERROR	Error	Failed to update packages Host <code>{VdsName}</code> . <code>{Message}</code> .

Code	Name	Severity	Message
555	USER_MOVE_TAG	Info	Tag \${TagName} was moved from \${OldParentTagName} to \${NewParentTagName} by \${UserName} .
556	USER_MOVE_TAG_FAILED	Error	Failed to move Tag \${TagName} from \${OldParentTagName} to \${NewParentTagName} (User: \${UserName}).
560	VDS_ANSIBLE_INSTALL_STARTED	Info	Ansible host-deploy playbook execution has started on host \${VdsName} .
561	VDS_ANSIBLE_INSTALL_FINISHED	Info	Ansible host-deploy playbook execution has successfully finished on host \${VdsName} .
562	VDS_ANSIBLE_HOST_REMOVE_STARTED	Info	Ansible host-remove playbook execution started on host \${VdsName} .
563	VDS_ANSIBLE_HOST_REMOVE_FINISHED	Info	Ansible host-remove playbook execution has successfully finished on host \${VdsName} . For more details check log \${LogFile}
564	VDS_ANSIBLE_HOST_REMOVE_FAILED	Warning	Ansible host-remove playbook execution failed on host \${VdsName} . For more details please check log \${LogFile}

Code	Name	Severity	Message
565	VDS_ANSIBLE_HOST_REMOVE_EXECUTION_FAILED	Info	Ansible host-remove playbook execution failed on host <code>{VdsName}</code> with message: <code>{Message}</code>
600	USER_VDS_MAINTENANCE	Info	Host <code>{VdsName}</code> was switched to Maintenance mode by <code>{UserName}</code> (Reason: <code>{Reason}</code>).
601	CPU_FLAGS_NX_IS_MISSING	Warning	Host <code>{VdsName}</code> is missing the NX cpu flag. This flag can be enabled via the host BIOS. Please set Disable Execute (XD) for an Intel host, or No Execute (NX) for AMD. Please make sure to completely power off the host for this change to take effect.
602	USER_VDS_MAINTENANCE_MIGRATION_FAILED	Warning	Host <code>{VdsName}</code> cannot change into maintenance mode - not all Vms have been migrated successfully. Consider manual intervention: stopping/migrating Vms: <code>{failedVms}</code> (User: <code>{UserName}</code>).

Code	Name	Severity	Message
603	VDS_SET_NONOPERATIONAL_IFACE_DOWN	Warning	Host \${VdsName} moved to Non-Operational state because interfaces which are down are needed by required networks in the current cluster: '\${NicsWithNetworks}'.
604	VDS_TIME_DRIFT_ALERT	Warning	Host \${VdsName} has time-drift of \${Actual} seconds while maximum configured value is \${Max} seconds.
605	PROXY_HOST_SELECTION	Info	Host \${Proxy} from \${Origin} was chosen as a proxy to execute fencing on Host \${VdsName}.
606	HOST_REFRESHED_CAPABILITIES	Info	Successfully refreshed the capabilities of host \${VdsName}.
607	HOST_REFRESH_CAPABILITIES_FAILED	Error	Failed to refresh the capabilities of host \${VdsName}.
608	HOST_INTERFACE_HIGH_NETWORK_USAGE	Warning	Host \${HostName} has network interface which exceeded the defined threshold [\${Threshold}%] (\${InterfaceName}: transmit rate[\${TransmitRate}%], receive rate [\${ReceiveRate}%])
609	HOST_INTERFACE_STATE_UP	Normal	Interface \${InterfaceName} on host \${VdsName}, changed state to up

Code	Name	Severity	Message
610	HOST_INTERFACE_STATE_DOWN	Warning	Interface <code>{InterfaceName}</code> on host <code>{VdsName}</code> , changed state to down
611	HOST_BOND_SLAVE_STATE_UP	Normal	Slave <code>{SlaveName}</code> of bond <code>{BondName}</code> on host <code>{VdsName}</code> , changed state to up
612	HOST_BOND_SLAVE_STATE_DOWN	Warning	Slave <code>{SlaveName}</code> of bond <code>{BondName}</code> on host <code>{VdsName}</code> , changed state to down
613	FENCE_KDUMP_LISTENER_IS_NOT_ALIVE	Error	Unable to determine if Kdump is in progress on host <code>{VdsName}</code> , because <code>fence_kdump</code> listener is not running.
614	KDUMP_FLOW_DETECTED_ON_VDS	Info	Kdump flow is in progress on host <code>{VdsName}</code> .
615	KDUMP_FLOW_NOT_DETECTED_ON_VDS	Info	Kdump flow is not in progress on host <code>{VdsName}</code> .
616	KDUMP_FLOW_FINISHED_ON_VDS	Info	Kdump flow finished on host <code>{VdsName}</code> .
617	KDUMP_DETECTION_NOT_CONFIGURED_ON_VDS	Warning	Kdump integration is enabled for host <code>{VdsName}</code> , but <code>kdump</code> is not configured properly on host.

Code	Name	Severity	Message
618	HOST_REGISTRATION_FAILED_INVALID_CLUSTER	Info	No default or valid cluster was found, Host <code>{VdsName}</code> registration failed
619	HOST_PROTOCOL_INCOMPATIBLE_WITH_CLUSTER	Warning	Host <code>{VdsName}</code> uses not compatible protocol during activation (xmlrpc instead of jsonrpc). Please examine installation logs and VDSM logs for failures and reinstall the host.
620	USER_VDS_MAINTENANCE_WITHOUT_REASON	Info	Host <code>{VdsName}</code> was switched to Maintenance mode by <code>{UserName}</code> .
650	USER_UNDO_RESTORE_FROM_SNAPSHOT_START	Info	Undoing a Snapshot-Preview for VM <code>{VmName}</code> was initialized by <code>{UserName}</code> .
651	USER_UNDO_RESTORE_FROM_SNAPSHOT_FINISH_SUCCESS	Info	Undoing a Snapshot-Preview for VM <code>{VmName}</code> has been completed.
652	USER_UNDO_RESTORE_FROM_SNAPSHOT_FINISH_FAILURE	Error	Failed to undo Snapshot-Preview for VM <code>{VmName}</code> .
700	DISK_ALIGNMENT_SCAN_START	Info	Starting alignment scan of disk <code>'{DiskAlias}'</code> .
701	DISK_ALIGNMENT_SCAN_FAILURE	Warning	Alignment scan of disk <code>'{DiskAlias}'</code> failed.
702	DISK_ALIGNMENT_SCAN_SUCCESS	Info	Alignment scan of disk <code>'{DiskAlias}'</code> is complete.

Code	Name	Severity	Message
809	USER_ADD_CLUSTER	Info	Cluster \${ClusterName} was added by \${UserName}
810	USER_ADD_CLUSTER_FAILED	Error	Failed to add Host cluster (User: \${UserName})
811	USER_UPDATE_CLUSTER	Info	Host cluster \${ClusterName} was updated by \${UserName}
812	USER_UPDATE_CLUSTER_FAILED	Error	Failed to update Host cluster (User: \${UserName})
813	USER_REMOVE_CLUSTER	Info	Host cluster \${ClusterName} was removed by \${UserName}
814	USER_REMOVE_CLUSTER_FAILED	Error	Failed to remove Host cluster (User: \${UserName})
815	USER_VDC_LOGOUT_FAILED	Error	Failed to log out user \${UserName} connected from ' \${SourceIP} ' using session ' \${SessionID} '.
816	MAC_POOL_EMPTY	Warning	No MAC addresses left in the MAC Address Pool.
817	CERTIFICATE_FILE_NOT_FOUND	Error	Could not find oVirt Engine Certificate file.
818	RUN_VM_FAILED	Error	Cannot run VM \${VmName} on Host \${VdsName} . Error: \${ErrMsg}

Code	Name	Severity	Message
819	VDS_REGISTER_ERROR_UPDATING_HOST	Error	Host registration failed - cannot update Host Name for Host \${VdsName2}. (Host: \${VdsName1})
820	VDS_REGISTER_ERROR_UPDATING_HOST_ALL_TAKEN	Error	Host registration failed - all available Host Names are taken. (Host: \${VdsName1})
821	VDS_REGISTER_HOST_IS_ACTIVE	Error	Host registration failed - cannot change Host Name of active Host \${VdsName2}. (Host: \${VdsName1})
822	VDS_REGISTER_ERROR_UPDATING_NAME	Error	Host registration failed - cannot update Host Name for Host \${VdsName2}. (Host: \${VdsName1})
823	VDS_REGISTER_ERROR_UPDATING_NAMES_ALL_TAKEN	Error	Host registration failed - all available Host Names are taken. (Host: \${VdsName1})
824	VDS_REGISTER_NAME_IS_ACTIVE	Error	Host registration failed - cannot change Host Name of active Host \${VdsName2}. (Host: \${VdsName1})
825	VDS_REGISTER_AUTO_APPROVE_PATTERN	Error	Host registration failed - auto approve pattern error. (Host: \${VdsName1})
826	VDS_REGISTER_FAILED	Error	Host registration failed. (Host: \${VdsName1})

Code	Name	Severity	Message
827	VDS_REGISTER_EXISTING_VDS_UPDATE_FAILED	Error	Host registration failed - cannot update existing Host. (Host: \${VdsName1})
828	VDS_REGISTER_SUCCEEDED	Info	Host \${VdsName1} registered.
829	VM_MIGRATION_ON_CONNECT_CHECK_FAILED	Error	VM migration logic failed. (VM name: \${VmName})
830	VM_MIGRATION_ON_CONNECT_CHECK_SUCCEEDED	Info	Migration check failed to execute.
831	USER_VDC_SESSION_TERMINATED	Info	User \${UserName} forcibly logged out user \${TerminatedSessionUsername} connected from '\${SourceIP}' using session '\${SessionID}'.
832	USER_VDC_SESSION_TERMINATION_FAILED	Error	User \${UserName} failed to forcibly log out user \${TerminatedSessionUsername} connected from '\${SourceIP}' using session '\${SessionID}'.
833	MAC_ADDRESS_IS_IN_USE	Warning	Network Interface \${IfaceName} has MAC address \${MACAddr} which is in use.
834	VDS_REGISTER_EMPTY_ID	Warning	Host registration failed, empty host id (Host: \${VdsHostName})

Code	Name	Severity	Message
835	SYSTEM_UPDATE_CLUSTER	Info	Host cluster \${ClusterName} was updated by system
836	SYSTEM_UPDATE_CLUSTER_FAILED	Info	Failed to update Host cluster by system
837	MAC_ADDRESSES_POOL_NOT_INITIALIZED	Warning	Mac Address Pool is not initialized. \${Message}
838	MAC_ADDRESS_IS_IN_USE_UNPLUG	Warning	Network Interface \${IfaceName} has MAC address \${MACAddr} which is in use, therefore it is being unplugged from VM \${VmName} .
839	HOST_AVAILABLE_UPDATES_FAILED	Error	Failed to check for available updates on host \${VdsName} with message '\${Message}' .
840	HOST_UPGRADE_STARTED	Info	Host \${VdsName} upgrade was started (User: \${UserName}).
841	HOST_UPGRADE_FAILED	Error	Failed to upgrade Host \${VdsName} (User: \${UserName}).
842	HOST_UPGRADE_FINISHED	Info	Host \${VdsName} upgrade was completed successfully.
845	HOST_CERTIFICATION_IS_ABOUT_TO_EXPIRE	Warning	Host \${VdsName} certification is about to expire at \${ExpirationDate} . Please renew the host's certification.

Code	Name	Severity	Message
846	ENGINE_CERTIFICATION_HAS_EXPIRED	Info	Engine's certification has expired at <code>{ExpirationDate}</code> . Please renew the engine's certification.
847	ENGINE_CERTIFICATION_IS_ABOUT_TO_EXPIRE	Warning	Engine's certification is about to expire at <code>{ExpirationDate}</code> . Please renew the engine's certification.
848	ENGINE_CA_CERTIFICATION_HAS_EXPIRED	Info	Engine's CA certification has expired at <code>{ExpirationDate}</code> .
849	ENGINE_CA_CERTIFICATION_IS_ABOUT_TO_EXPIRE	Warning	Engine's CA certification is about to expire at <code>{ExpirationDate}</code> .
850	USER_ADD_PERMISSION	Info	User/Group <code>{SubjectName}</code> , Namespace <code>{Namespace}</code> , Authorization provider: <code>{Authz}</code> was granted permission for Role <code>{RoleName}</code> on <code>{VdcObjectType}</code> <code>{VdcObjectName}</code> , by <code>{UserName}</code> .
851	USER_ADD_PERMISSION_FAILED	Error	User <code>{UserName}</code> failed to grant permission for Role <code>{RoleName}</code> on <code>{VdcObjectType}</code> <code>{VdcObjectName}</code> to User/Group <code>{SubjectName}</code> .

Code	Name	Severity	Message
852	USER_REMOVE_PERMISSION	Info	User/Group \${SubjectName} Role \${RoleName} permission was removed from \${VdcObjectType} \${VdcObjectName} by \${UserName}
853	USER_REMOVE_PERMISSION_FAILED	Error	User \${UserName} failed to remove permission for Role \${RoleName} from \${VdcObjectType} \${VdcObjectName} to User/Group \${SubjectName}
854	USER_ADD_ROLE	Info	Role \${RoleName} granted to \${UserName}
855	USER_ADD_ROLE_FAILED	Error	Failed to grant role \${RoleName} (User \${UserName})
856	USER_UPDATE_ROLE	Info	\${UserName} Role was updated to the \${RoleName} Role
857	USER_UPDATE_ROLE_FAILED	Error	Failed to update role \${RoleName} to \${UserName}
858	USER_REMOVE_ROLE	Info	Role \${RoleName} removed from \${UserName}
859	USER_REMOVE_ROLE_FAILED	Error	Failed to remove role \${RoleName} (User \${UserName})
860	USER_ATTACHED_ACTION_GROUP_TO_ROLE	Info	Action group \${ActionGroup} was attached to Role \${RoleName} by \${UserName}

Code	Name	Severity	Message
861	USER_ATTACHED_ACTION_GROUP_TO_ROLE_FAILED	Error	Failed to attach Action group <code>\${ActionGroup}</code> to Role <code>\${RoleName}</code> (User: <code>\${UserName}</code>)
862	USER_DETACHED_ACTION_GROUP_FROM_ROLE	Info	Action group <code>\${ActionGroup}</code> was detached from Role <code>\${RoleName}</code> by <code>\${UserName}</code>
863	USER_DETACHED_ACTION_GROUP_FROM_ROLE_FAILED	Error	Failed to attach Action group <code>\${ActionGroup}</code> to Role <code>\${RoleName}</code> by <code>\${UserName}</code>
864	USER_ADD_ROLE_WITH_ACTION_GROUP	Info	Role <code>\${RoleName}</code> was added by <code>\${UserName}</code>
865	USER_ADD_ROLE_WITH_ACTION_GROUP_FAILED	Error	Failed to add role <code>\${RoleName}</code>
866	USER_ADD_SYSTEM_PERMISSION	Info	User/Group <code>\${SubjectName}</code> was granted permission for Role <code>\${RoleName}</code> on <code>\${VdcObjectType}</code> by <code>\${UserName}</code> .
867	USER_ADD_SYSTEM_PERMISSION_FAILED	Error	User <code>\${UserName}</code> failed to grant permission for Role <code>\${RoleName}</code> on <code>\${VdcObjectType}</code> to User/Group <code>\${SubjectName}</code> .

Code	Name	Severity	Message
868	USER_REMOVE_SYSTEM_PERMISSION	Info	User/Group \${SubjectName} Role \${RoleName} permission was removed from \${VdcObjectType} by \${UserName}
869	USER_REMOVE_SYSTEM_PERMISSION_FAILED	Error	User \${UserName} failed to remove permission for Role \${RoleName} from \${VdcObjectType} to User/Group \${SubjectName}
870	USER_ADD_PROFILE	Info	Profile created for \${UserName}
871	USER_ADD_PROFILE_FAILED	Error	Failed to create profile for \${UserName}
872	USER_UPDATE_PROFILE	Info	Updated profile for \${UserName}
873	USER_UPDATE_PROFILE_FAILED	Error	Failed to update profile for \${UserName}
874	USER_REMOVE_PROFILE	Info	Removed profile for \${UserName}
875	USER_REMOVE_PROFILE_FAILED	Error	Failed to remove profile for \${UserName}
876	HOST_CERTIFICATION_IS_INVALID	Error	Host \${VdsName} certification is invalid. The certification has no peer certificates.

Code	Name	Severity	Message
877	HOST_CERTIFICATI ON_HAS_EXPIRED	Info	Host \${VdsName} certification has expired at \${ExpirationDate}. Please renew the host's certification.
878	ENGINE_CERTIFICA TION_IS_ABOUT_TO _EXPIRE_ALERT	Info	Engine's certification is about to expire at \${ExpirationDate}. Please renew the engine's certification.
879	HOST_CERTIFICATI ON_IS_ABOUT_TO_ EXPIRE_ALERT	Info	Host \${VdsName} certification is about to expire at \${ExpirationDate}. Please renew the host's certification.
880	HOST_CERTIFICATI ON_ENROLLMENT_ STARTED	Normal	Enrolling certificate for host \${VdsName} was started (User: \${UserName}).
881	HOST_CERTIFICATI ON_ENROLLMENT_ FINISHED	Normal	Enrolling certificate for host \${VdsName} was completed successfully (User: \${UserName}).
882	HOST_CERTIFICATI ON_ENROLLMENT_ FAILED	Error	Failed to enroll certificate for host \${VdsName} (User: \${UserName}).
883	ENGINE_CA_CERTIF ICATION_IS_ABOUT _TO_EXPIRE_ALERT	Info	Engine's CA certification is about to expire at \${ExpirationDate}.
884	HOST_AVAILABLE_ UPDATES_STARTED	Info	Started to check for available updates on host \${VdsName}.

Code	Name	Severity	Message
885	HOST_AVAILABLE_UPDATES_FINISHED	Info	Check for available updates on host <code>{VdsName}</code> was completed successfully with message <code>{Message}</code> .
886	HOST_AVAILABLE_UPDATES_PROCESS_IS_ALREADY_RUNNING	Warning	Failed to check for available updates on host <code>{VdsName}</code> : Another process is already running.
887	HOST_AVAILABLE_UPDATES_SKIPPED_UNSUPPORTED_STATUS	Warning	Failed to check for available updates on host <code>{VdsName}</code> : Unsupported host status.
890	HOST_UPGRADE_FINISHED_MANUAL_HA	Warning	Host <code>{VdsName}</code> upgrade was completed successfully, but the Hosted Engine HA service may still be in maintenance mode. If necessary, please correct this manually.
900	AD_COMPUTER_ACCOUNT_SUCCESSFUL	Info	Account creation successful.
901	AD_COMPUTER_ACCOUNT_FAILED	Error	Account creation failed.
918	USER_FORCE_REMOVE_STORAGE_POOL	Info	Data Center <code>{StoragePoolName}</code> was forcibly removed by <code>{UserName}</code>
919	USER_FORCE_REMOVE_STORAGE_POOL_FAILED	Error	Failed to forcibly remove Data Center <code>{StoragePoolName}</code> . (User: <code>{UserName}</code>)

Code	Name	Severity	Message
925	MAC_ADDRESS_IS_EXTERNAL	Warning	VM {VmName} has MAC address(es) {MACAddr} , which is/are out of its MAC pool definitions.
926	NETWORK_REMOVE_BOND	Info	Remove bond: {BondName} for Host: {VdsName} (User: {UserName}).
927	NETWORK_REMOVE_BOND_FAILED	Error	Failed to remove bond: {BondName} for Host: {VdsName} (User: {UserName}).
928	NETWORK_VDS_NETWORK_MATCH_CLUSTER	Info	Vds {VdsName} network match to cluster {ClusterName}
929	NETWORK_VDS_NETWORK_NOT_MATCH_CLUSTER	Error	Vds {VdsName} network does not match to cluster {ClusterName}
930	NETWORK_REMOVE_VM_INTERFACE	Info	Interface {InterfaceName} ({InterfaceType}) was removed from VM {VmName} . (User: {UserName})
931	NETWORK_REMOVE_VM_INTERFACE_FAILED	Error	Failed to remove Interface {InterfaceName} ({InterfaceType}) from VM {VmName} . (User: {UserName})
932	NETWORK_ADD_VM_INTERFACE	Info	Interface {InterfaceName} ({InterfaceType}) was added to VM {VmName} . (User: {UserName})

Code	Name	Severity	Message
933	NETWORK_ADD_VM_INTERFACE_FAILED	Error	Failed to add Interface <code>{InterfaceName}</code> (<code>{InterfaceType}</code>) to VM <code>{VmName}</code> . (User: <code>{UserName}</code>)
934	NETWORK_UPDATE_VM_INTERFACE	Info	Interface <code>{InterfaceName}</code> (<code>{InterfaceType}</code>) was updated for VM <code>{VmName}</code> . <code>{LinkState}</code> (User: <code>{UserName}</code>)
935	NETWORK_UPDATE_VM_INTERFACE_FAILED	Error	Failed to update Interface <code>{InterfaceName}</code> (<code>{InterfaceType}</code>) for VM <code>{VmName}</code> . (User: <code>{UserName}</code>)
936	NETWORK_ADD_TEMPLATE_INTERFACE	Info	Interface <code>{InterfaceName}</code> (<code>{InterfaceType}</code>) was added to Template <code>{VmTemplateName}</code> . (User: <code>{UserName}</code>)
937	NETWORK_ADD_TEMPLATE_INTERFACE_FAILED	Error	Failed to add Interface <code>{InterfaceName}</code> (<code>{InterfaceType}</code>) to Template <code>{VmTemplateName}</code> . (User: <code>{UserName}</code>)
938	NETWORK_REMOVE_TEMPLATE_INTERFACE	Info	Interface <code>{InterfaceName}</code> (<code>{InterfaceType}</code>) was removed from Template <code>{VmTemplateName}</code> . (User: <code>{UserName}</code>)

Code	Name	Severity	Message
939	NETWORK_REMOVE_TEMPLATE_INTERFACE_FAILED	Error	Failed to remove Interface <code>\${InterfaceName}</code> (<code>\${InterfaceType}</code>) from Template <code>\${VmTemplateName}</code> . (User: <code>\${UserName}</code>)
940	NETWORK_UPDATE_TEMPLATE_INTERFACE	Info	Interface <code>\${InterfaceName}</code> (<code>\${InterfaceType}</code>) was updated for Template <code>\${VmTemplateName}</code> . (User: <code>\${UserName}</code>)
941	NETWORK_UPDATE_TEMPLATE_INTERFACE_FAILED	Error	Failed to update Interface <code>\${InterfaceName}</code> (<code>\${InterfaceType}</code>) for Template <code>\${VmTemplateName}</code> . (User: <code>\${UserName}</code>)
942	NETWORK_ADD_NETWORK	Info	Network <code>\${NetworkName}</code> was added to Data Center: <code>\${StoragePoolName}</code>
943	NETWORK_ADD_NETWORK_FAILED	Error	Failed to add Network <code>\${NetworkName}</code> to Data Center: <code>\${StoragePoolName}</code>
944	NETWORK_REMOVE_NETWORK	Info	Network <code>\${NetworkName}</code> was removed from Data Center: <code>\${StoragePoolName}</code>
945	NETWORK_REMOVE_NETWORK_FAILED	Error	Failed to remove Network <code>\${NetworkName}</code> from Data Center: <code>\${StoragePoolName}</code>

Code	Name	Severity	Message
946	NETWORK_ATTACH_NETWORK_TO_CLUSTER	Info	Network \${NetworkName} attached to Cluster \${ClusterName}
947	NETWORK_ATTACH_NETWORK_TO_CLUSTER_FAILED	Error	Failed to attach Network \${NetworkName} to Cluster \${ClusterName}
948	NETWORK_DETACH_NETWORK_TO_CLUSTER	Info	Network \${NetworkName} detached from Cluster \${ClusterName}
949	NETWORK_DETACH_NETWORK_TO_CLUSTER_FAILED	Error	Failed to detach Network \${NetworkName} from Cluster \${ClusterName}
950	USER_ADD_STORAGE_POOL	Info	Data Center \${StoragePoolName} , Compatibility Version \${CompatibilityVersion} and Quota Type \${QuotaEnforcementType} was added by \${UserName}
951	USER_ADD_STORAGE_POOL_FAILED	Error	Failed to add Data Center \${StoragePoolName} . (User: \${UserName})
952	USER_UPDATE_STORAGE_POOL	Info	Data Center \${StoragePoolName} was updated by \${UserName}
953	USER_UPDATE_STORAGE_POOL_FAILED	Error	Failed to update Data Center \${StoragePoolName} . (User: \${UserName})

Code	Name	Severity	Message
954	USER_REMOVE_STORAGE_POOL	Info	Data Center <code>\${StoragePoolName}</code> was removed by <code>\${UserName}</code>
955	USER_REMOVE_STORAGE_POOL_FAILED	Error	Failed to remove Data Center <code>\${StoragePoolName}</code> . (User: <code>\${UserName}</code>)
956	USER_ADD_STORAGE_DOMAIN	Info	Storage Domain <code>\${StorageDomainName}</code> was added by <code>\${UserName}</code>
957	USER_ADD_STORAGE_DOMAIN_FAILED	Error	Failed to add Storage Domain <code>\${StorageDomainName}</code> . (User: <code>\${UserName}</code>)
958	USER_UPDATE_STORAGE_DOMAIN	Info	Storage Domain <code>\${StorageDomainName}</code> was updated by <code>\${UserName}</code>
959	USER_UPDATE_STORAGE_DOMAIN_FAILED	Error	Failed to update Storage Domain <code>\${StorageDomainName}</code> . (User: <code>\${UserName}</code>)
960	USER_REMOVE_STORAGE_DOMAIN	Info	Storage Domain <code>\${StorageDomainName}</code> was removed by <code>\${UserName}</code>
961	USER_REMOVE_STORAGE_DOMAIN_FAILED	Error	Failed to remove Storage Domain <code>\${StorageDomainName}</code> . (User: <code>\${UserName}</code>)

Code	Name	Severity	Message
962	USER_ATTACH_STORAGE_DOMAIN_TO_POOL	Info	Storage Domain <code>\${StorageDomainName}</code> was attached to Data Center <code>\${StoragePoolName}</code> by <code>\${UserName}</code>
963	USER_ATTACH_STORAGE_DOMAIN_TO_POOL_FAILED	Error	Failed to attach Storage Domain <code>\${StorageDomainName}</code> to Data Center <code>\${StoragePoolName}</code> . (User: <code>\${UserName}</code>)
964	USER_DETACH_STORAGE_DOMAIN_FROM_POOL	Info	Storage Domain <code>\${StorageDomainName}</code> was detached from Data Center <code>\${StoragePoolName}</code> by <code>\${UserName}</code>
965	USER_DETACH_STORAGE_DOMAIN_FROM_POOL_FAILED	Error	Failed to detach Storage Domain <code>\${StorageDomainName}</code> from Data Center <code>\${StoragePoolName}</code> . (User: <code>\${UserName}</code>)
966	USER_ACTIVATED_STORAGE_DOMAIN	Info	Storage Domain <code>\${StorageDomainName}</code> (Data Center <code>\${StoragePoolName}</code>) was activated by <code>\${UserName}</code>
967	USER_ACTIVATE_STORAGE_DOMAIN_FAILED	Error	Failed to activate Storage Domain <code>\${StorageDomainName}</code> (Data Center <code>\${StoragePoolName}</code>) by <code>\${UserName}</code>

Code	Name	Severity	Message
968	USER_DEACTIVATE_STORAGE_DOMAIN	Info	Storage Domain <code>\${StorageDomainName}</code> (Data Center <code>\${StoragePoolName}</code>) was deactivated and has moved to 'Preparing for maintenance' until it will no longer be accessed by any Host of the Data Center.
969	USER_DEACTIVATE_STORAGE_DOMAIN_FAILED	Error	Failed to deactivate Storage Domain <code>\${StorageDomainName}</code> (Data Center <code>\${StoragePoolName}</code>).
970	SYSTEM_DEACTIVATED_STORAGE_DOMAIN	Warning	Storage Domain <code>\${StorageDomainName}</code> (Data Center <code>\${StoragePoolName}</code>) was deactivated by system because it's not visible by any of the hosts.
971	SYSTEM_DEACTIVATE_STORAGE_DOMAIN_FAILED	Error	Failed to deactivate Storage Domain <code>\${StorageDomainName}</code> (Data Center <code>\${StoragePoolName}</code>).
972	USER_EXTENDED_STORAGE_DOMAIN	Info	Storage <code>\${StorageDomainName}</code> has been extended by <code>\${UserName}</code> . Please wait for refresh.
973	USER_EXTENDED_STORAGE_DOMAIN_FAILED	Error	Failed to extend Storage Domain <code>\${StorageDomainName}</code> . (User: <code>\${UserName}</code>)

Code	Name	Severity	Message
974	USER_REMOVE_VG	Info	Volume group \${VgId} was removed by \${UserName} .
975	USER_REMOVE_VG_FAILED	Error	Failed to remove Volume group \${VgId} . (User: UserName)
976	USER_ACTIVATE_STORAGE_POOL	Info	Data Center \${StoragePoolName} was activated. (User: \${UserName})
977	USER_ACTIVATE_STORAGE_POOL_FAILED	Error	Failed to activate Data Center \${StoragePoolName} . (User: \${UserName})
978	SYSTEM_FAILED_CHANGE_STORAGE_POOL_STATUS	Error	Failed to change Data Center \${StoragePoolName} status.
979	SYSTEM_CHANGE_STORAGE_POOL_STATUS_NO_HOST_FOR_SPM	Error	Fencing failed on Storage Pool Manager \${VdsName} for Data Center \${StoragePoolName} . Setting status to Non-Operational.
980	SYSTEM_CHANGE_STORAGE_POOL_STATUS_PROBLEMATIC	Warning	Invalid status on Data Center \${StoragePoolName} . Setting status to Non Responsive.
981	USER_FORCE_REMOVE_STORAGE_DOMAIN	Info	Storage Domain \${StorageDomainName} was forcibly removed by \${UserName}

Code	Name	Severity	Message
982	USER_FORCE_REMOVE_STORAGE_DOMAIN_FAILED	Error	Failed to forcibly remove Storage Domain <code>{StorageDomainName}</code> . (User: <code>{UserName}</code>)
983	RECONSTRUCT_MASTER_FAILED_NO_MASTER	Warning	No valid Data Storage Domains are available in Data Center <code>{StoragePoolName}</code> (please check your storage infrastructure).
984	RECONSTRUCT_MASTER_DONE	Info	Reconstruct Master Domain for Data Center <code>{StoragePoolName}</code> completed.
985	RECONSTRUCT_MASTER_FAILED	Error	Failed to Reconstruct Master Domain for Data Center <code>{StoragePoolName}</code> .
986	SYSTEM_CHANGE_STORAGE_POOL_STATUS_PROBLEMATIC_SEARCHING_NEW_SPM	Warning	Data Center is being initialized, please wait for initialization to complete.
987	SYSTEM_CHANGE_STORAGE_POOL_STATUS_PROBLEMATIC_WITH_ERROR	Warning	Invalid status on Data Center <code>{StoragePoolName}</code> . Setting Data Center status to Non Responsive (On host <code>{VdsName}</code> , Error: <code>{Error}</code>).
988	USER_CONNECT_HOSTS_TO_LUN_FAILED	Error	Failed to connect Host <code>{VdsName}</code> to device. (User: <code>{UserName}</code>)

Code	Name	Severity	Message
989	SYSTEM_CHANGE_STORAGE_POOL_STATUS_PROBLEMATIC_FROM_NON_OPERATIONAL	Info	Try to recover Data Center \${StoragePoolName} . Setting status to Non Responsive.
990	SYSTEM_MASTER_DOMAIN_NOT_IN_SYNC	Warning	Sync Error on Master Domain between Host \${VdsName} and oVirt Engine. Domain: \${StorageDomainName} is marked as Master in oVirt Engine database but not on the Storage side. Please consult with Support on how to fix this issue.
991	RECOVERY_STORAGE_POOL	Info	Data Center \${StoragePoolName} was recovered by \${UserName}
992	RECOVERY_STORAGE_POOL_FAILED	Error	Failed to recover Data Center \${StoragePoolName} (User: \${UserName})
993	SYSTEM_CHANGE_STORAGE_POOL_STATUS_RESET_IRS	Info	Data Center \${StoragePoolName} was reset. Setting status to Non Responsive (Elect new Storage Pool Manager).
994	CONNECT_STORAGE_SERVERS_FAILED	Warning	Failed to connect Host \${VdsName} to Storage Servers
995	CONNECT_STORAGE_POOL_FAILED	Warning	Failed to connect Host \${VdsName} to Storage Pool \${StoragePoolName}

Code	Name	Severity	Message
996	STORAGE_DOMAIN_ERROR	Error	The error message for connection <code>{Connection}</code> returned by VDSM was: <code>{ErrorMessage}</code>
997	REFRESH_REPOSITORY_IMAGE_LIST_FAILED	Error	Refresh image list failed for domain(s): <code>{imageDomains}</code> . Please check domain activity.
998	REFRESH_REPOSITORY_IMAGE_LIST_SUCCEEDED	Info	Refresh image list succeeded for domain(s): <code>{imageDomains}</code>
999	STORAGE_ALERT_VG_METADATA_CRITICALLY_FULL	Error	The system has reached the 80% watermark on the VG metadata area size on <code>{StorageDomainName}</code> . This is due to a high number of Vdisks or large Vdisks size allocated on this specific VG.
1000	STORAGE_ALERT_SMALL_VG_METADATA	Warning	The allocated VG metadata area size is smaller than 50MB on <code>{StorageDomainName}</code> , which might limit its capacity (the number of Vdisks and/or their size).
1001	USER_RUN_VM_FAILURE_STATELESS_SNAPSHOT_LEFT	Error	Failed to start VM <code>{VmName}</code> , because exist snapshot for stateless state. Snapshot will be deleted.

Code	Name	Severity	Message
1002	USER_ATTACH_STORAGE_DOMAINS_TO_POOL	Info	Storage Domains were attached to Data Center \${StoragePoolName} by \${UserName}
1003	USER_ATTACH_STORAGE_DOMAINS_TO_POOL_FAILED	Error	Failed to attach Storage Domains to Data Center \${StoragePoolName} . (User: \${UserName})
1004	STORAGE_DOMAIN_TASKS_ERROR	Warning	Storage Domain \${StorageDomainName} is down while there are tasks running on it. These tasks may fail.
1005	UPDATE_OVF_FOR_STORAGE_POOL_FAILED	Warning	Failed to update VMs/Templates OVF data in Data Center \${StoragePoolName} .
1006	UPGRADE_STORAGE_POOL_ENCOUNTED_PROBLEMS	Warning	Data Center \${StoragePoolName} has encountered problems during upgrade process.
1007	REFRESH_REPOSITORY_IMAGE_LIST_INCOMPLETE	Warning	Refresh image list probably incomplete for domain \${imageDomain}, only \${imageListSize} images discovered.
1008	NUMBER_OF_LVS_ON_STORAGE_DOMAIN_EXCEEDED_THRESHOLD	Warning	The number of LVs on the domain \${storageDomainName} exceeded \${maxNumOfLVs}, you are approaching the limit where performance may degrade.

Code	Name	Severity	Message
1009	USER_DEACTIVATE_STORAGE_DOMAIN_OVF_UPDATE_INCOMPLETE	Warning	Failed to deactivate Storage Domain <code>\${StorageDomainName}</code> as the engine was restarted during the operation, please retry. (Data Center <code>\${StoragePoolName}</code>).
1010	RELOAD_CONFIGURATIONS_SUCCESS	Info	System Configurations reloaded successfully.
1011	RELOAD_CONFIGURATIONS_FAILURE	Error	System Configurations failed to reload.
1012	NETWORK_ACTIVATE_VM_INTERFACE_SUCCESS	Info	Network Interface <code>\${InterfaceName}</code> (<code>\${InterfaceType}</code>) was plugged to VM <code>\${VmName}</code> . (User: <code>\${UserName}</code>)
1013	NETWORK_ACTIVATE_VM_INTERFACE_FAILURE	Error	Failed to plug Network Interface <code>\${InterfaceName}</code> (<code>\${InterfaceType}</code>) to VM <code>\${VmName}</code> . (User: <code>\${UserName}</code>)
1014	NETWORK_DEACTIVATE_VM_INTERFACE_SUCCESS	Info	Network Interface <code>\${InterfaceName}</code> (<code>\${InterfaceType}</code>) was unplugged from VM <code>\${VmName}</code> . (User: <code>\${UserName}</code>)
1015	NETWORK_DEACTIVATE_VM_INTERFACE_FAILURE	Error	Failed to unplug Network Interface <code>\${InterfaceName}</code> (<code>\${InterfaceType}</code>) from VM <code>\${VmName}</code> . (User: <code>\${UserName}</code>)

Code	Name	Severity	Message
1016	UPDATE_FOR_OVF_STORES_FAILED	Warning	Failed to update OVF disks \${DiskIds}, OVF data isn't updated on those OVF stores (Data Center \${DataCenterName}, Storage Domain \${StorageDomainName}).
1017	RETRIEVE_OVF_STORE_FAILED	Warning	Failed to retrieve VMs and Templates from the OVF disk of Storage Domain \${StorageDomainName}.
1018	OVF_STORE_DOES_NOT_EXISTS	Warning	This Data center compatibility version does not support importing a data domain with its entities (VMs and Templates). The imported domain will be imported without them.
1019	UPDATE_DESCRIPTION_FOR_DISK_FAILED	Error	Failed to update the meta data description of disk \${DiskName} (Data Center \${DataCenterName}, Storage Domain \${StorageDomainName}).
1020	UPDATE_DESCRIPTION_FOR_DISK_SKIPPED_SINCE_STORAGE_DOMAIN_NOT_ACTIVE	Warning	Not updating the metadata of Disk \${DiskName} (Data Center \${DataCenterName}). Since the Storage Domain \${StorageDomainName} is not in active.

Code	Name	Severity	Message
1022	USER_REFRESH_LUN_STORAGE_DOMAIN	Info	Resize LUNs operation succeeded.
1023	USER_REFRESH_LUN_STORAGE_DOMAIN_FAILED	Error	Failed to resize LUNs.
1024	USER_REFRESH_LUN_STORAGE_DIFFERENT_SIZE_DOMAIN_FAILED	Error	Failed to resize LUNs. Not all the hosts are seeing the same LUN size.
1025	VM_PAUSED	Info	VM <code>{VmName}</code> has been paused.
1026	FAILED_TO_STORE_ENTIRE_DISK_FIELD_IN_DISK_DESCRIPTION_METADATA	Warning	Failed to store field <code>{DiskFieldName}</code> as a part of <code>{DiskAlias}</code> 's description metadata due to storage space limitations. The field <code>{DiskFieldName}</code> will be truncated.
1027	FAILED_TO_STORE_ENTIRE_DISK_FIELD_AND_REST_OF_FIELDS_IN_DISK_DESCRIPTION_METADATA	Warning	Failed to store field <code>{DiskFieldName}</code> as a part of <code>{DiskAlias}</code> 's description metadata due to storage space limitations. The value will be truncated and the following fields will not be stored at all: <code>{DiskFieldsNames}</code> .
1028	FAILED_TO_STORE_DISK_FIELDS_IN_DISK_DESCRIPTION_METADATA	Warning	Failed to store the following fields in the description metadata of disk <code>{DiskAlias}</code> due to storage space limitations: <code>{DiskFieldsNames}</code> .

Code	Name	Severity	Message
1029	STORAGE_DOMAIN_MOVED_TO_MAINTENANCE	Info	Storage Domain <code>\${StorageDomainName}</code> (Data Center <code>\${StoragePoolName}</code>) successfully moved to Maintenance as it's no longer accessed by any Host of the Data Center.
1030	USER_DEACTIVATE_D_LAST_MASTER_STORAGE_DOMAIN	Info	Storage Domain <code>\${StorageDomainName}</code> (Data Center <code>\${StoragePoolName}</code>) was deactivated.
1031	TRANSFER_IMAGE_INITIATED	Info	Image <code>\${TransferType}</code> with disk <code>\${DiskAlias}</code> was initiated by <code>\${UserName}</code> .
1032	TRANSFER_IMAGE_SUCCEEDED	Info	Image <code>\${TransferType}</code> with disk <code>\${DiskAlias}</code> succeeded.
1033	TRANSFER_IMAGE_CANCELLED	Info	Image <code>\${TransferType}</code> with disk <code>\${DiskAlias}</code> was cancelled.
1034	TRANSFER_IMAGE_FAILED	Error	Image <code>\${TransferType}</code> with disk <code>\${DiskAlias}</code> failed.
1035	TRANSFER_IMAGE_TEARDOWN_FAILED	Info	Failed to tear down image <code>\${DiskAlias}</code> after image transfer session.
1036	USER_SCAN_STORAGE_DOMAIN_FOR_UNREGISTERED_DISKS	Info	Storage Domain <code>\${StorageDomainName}</code> has finished to scan for unregistered disks by <code>\${UserName}</code> .

Code	Name	Severity	Message
1037	USER_SCAN_STORAGE_DOMAIN_FOR_UNREGISTERED_DISKS_FAILED	Error	Storage Domain <code>\${StorageDomainName}</code> failed to scan for unregistered disks by <code>\${UserName}</code> .
1039	LUNS_BROKE_STORAGE_DOMAIN_SUPPORT	Warning	Luns with IDs: <code>[\${LunIds}]</code> were updated in the DB but caused the storage domain <code>\${StorageDomainName}</code> (ID <code>\${storageDomainId}</code>) to stop supporting passing discard from the guest to the underlying storage. Please configure these luns' discard support in the underlying storage or disable 'Enable Discard' for vm disks on this storage domain.
1040	DISKS_WITH_ILLEGAL_PASS_DISCARD_EXIST	Warning	Disks with IDs: <code>[\${DiskIds}]</code> have their 'Enable Discard' on even though the underlying storage does not support it. Please configure the underlying storage to support discard or disable 'Enable Discard' for these disks.
1041	USER_REMOVE_DEVICE_FROM_STORAGE_DOMAIN_FAILED	Error	Failed to remove <code>\${LunId}</code> from Storage Domain <code>\${StorageDomainName}</code> . (User: <code>\${UserName}</code>)

Code	Name	Severity	Message
1042	USER_REMOVE_DE VICE_FROM_STORA GE_DOMAIN	Info	`\${LunId}` was removed from Storage Domain `\${StorageDomainName}`. (User: `\${UserName}`)
1043	USER_REMOVE_DE VICE_FROM_STORA GE_DOMAIN_START ED	Info	Started to remove `\${LunId}` from Storage Domain `\${StorageDomainName}`. (User: `\${UserName}`)
1044	ILLEGAL_STORAGE _DOMAIN_DISCARD _AFTER_DELETE	Warning	The storage domain with id `\${storageDomainId}` has its 'Discard After Delete' enabled even though the underlying storage does not support discard. Therefore, disks and snapshots on this storage domain will not be discarded before they are removed.
1045	LUNS_BROKE_SD_ DISCARD_AFTER_D ELETE_SUPPORT	Warning	Luns with IDs: `\${LunsIds}` were updated in the DB but caused the storage domain `\${StorageDomainName}` (ID `\${storageDomainId}`) to stop supporting discard after delete. Please configure these luns' discard support in the underlying storage or disable 'Discard After Delete' for this storage domain.

Code	Name	Severity	Message
1046	STORAGE_DOMAINS_COULD_NOT_BE_SYNCED	Info	Storage domains with IDs [\${StorageDomainsIds}] could not be synchronized. To synchronize them, please move them to maintenance and then activate.
1048	DIRECT_LUNS_COULD_NOT_BE_SYNCED	Info	Direct LUN disks with IDs [\${DirectLunDiskIds}] could not be synchronized because there was no active host in the data center. Please synchronize them to get their latest information from the storage.
1052	OVF_STORES_UPDATED_IGNORED	Normal	OVFs update was ignored - nothing to update for storage domain '\${StorageDomainName}'
1060	UPLOAD_IMAGE_CLIENT_ERROR	Error	Unable to upload image to disk \${DiskId} due to a client error. Make sure the selected file is readable.
1061	UPLOAD_IMAGE_XHR_TIMEOUT_ERROR	Error	Unable to upload image to disk \${DiskId} due to a request timeout error. The upload bandwidth might be too slow. Please try to reduce the chunk size: 'engine-config - s UploadImageChunk SizeKB

Code	Name	Severity	Message
1062	UPLOAD_IMAGE_NETWORK_ERROR	Error	Unable to upload image to disk <code>{DiskId}</code> due to a network error. Ensure that <code>ovirt-imageio</code> service is installed and configured and that <code>ovirt-engine</code> 's CA certificate is registered as a trusted CA in the browser. The certificate can be fetched from <code>{EngineUrl}/ovirt-engine/services/pki-resource?resource</code>
1063	DOWNLOAD_IMAGE_NETWORK_ERROR	Error	Unable to download disk <code>{DiskId}</code> due to a network error. Make sure <code>ovirt-imageio</code> service is installed and configured, and <code>ovirt-engine</code> 's certificate is registered as a valid CA in the browser. The certificate can be fetched from <code>https://<engine_url>/ovirt-engine/services/pki-resource?resource</code>
1064	TRANSFER_IMAGE_STOPPED_BY_SYSTEM_TICKET_RENEW_FAILURE	Error	Transfer was stopped by system. Reason: failure in transfer image ticket renewal.
1065	TRANSFER_IMAGE_STOPPED_BY_SYSTEM_MISSING_TICKET	Error	Transfer was stopped by system. Reason: missing transfer image ticket.

Code	Name	Severity	Message
1067	TRANSFER_IMAGE_STOPPED_BY_SYSTEM_MISSING_HOST	Error	Transfer was stopped by system. Reason: Could not find a suitable host for image data transfer.
1068	TRANSFER_IMAGE_STOPPED_BY_SYSTEM_FAILED_TO_CREATE_TICKET	Error	Transfer was stopped by system. Reason: failed to create a signed image ticket.
1069	TRANSFER_IMAGE_STOPPED_BY_SYSTEM_FAILED_TO_ADD_TICKET_TO_DAEMON	Error	Transfer was stopped by system. Reason: failed to add image ticket to ovirt-imageio-daemon.
1070	TRANSFER_IMAGE_STOPPED_BY_SYSTEM_FAILED_TO_ADD_TICKET_TO_PROXY	Error	Transfer was stopped by system. Reason: failed to add image ticket to ovirt-imageio.
1071	UPLOAD_IMAGE_PAUSED_BY_SYSTEM_TIMEOUT	Error	Upload was paused by system. Reason: timeout due to transfer inactivity.
1072	DOWNLOAD_IMAGE_CANCELED_TIMEOUT	Error	Download was canceled by system. Reason: timeout due to transfer inactivity.
1073	TRANSFER_IMAGE_PAUSED_BY_USER	Normal	Image transfer was paused by user ({UserName}).
1074	TRANSFER_IMAGE_RESUMED_BY_USER	Normal	Image transfer was resumed by user ({UserName}).

Code	Name	Severity	Message
1098	NETWORK_UPDATE_DISPLAY_FOR_HOST_WITH_ACTIVE_VM	Warning	Display Network was updated on Host \${VdsName} with active VMs attached. The change will be applied to those VMs after their next reboot. Running VMs might loose display connectivity until then.
1099	NETWORK_UPDATE_DISPLAY_FOR_CLUSTER_WITH_ACTIVE_VM	Warning	Display Network (\${NetworkName}) was updated for Cluster \${ClusterName} with active VMs attached. The change will be applied to those VMs after their next reboot.
1100	NETWORK_UPDATE_DISPLAY_TO_CLUSTER	Info	Update Display Network (\${NetworkName}) for Cluster \${ClusterName}. (User: \${UserName})
1101	NETWORK_UPDATE_DISPLAY_TO_CLUSTER_FAILED	Error	Failed to update Display Network (\${NetworkName}) for Cluster \${ClusterName}. (User: \${UserName})
1102	NETWORK_UPDATE_NETWORK_TO_VDS_INTERFACE	Info	Update Network \${NetworkName} in Host \${VdsName}. (User: \${UserName})
1103	NETWORK_UPDATE_NETWORK_TO_VDS_INTERFACE_FAILED	Error	Failed to update Network \${NetworkName} in Host \${VdsName}. (User: \${UserName})

Code	Name	Severity	Message
1104	NETWORK_COMMIT_NETWORK_CHANGES	Info	Network changes were saved on host <code>\${VdsName}</code>
1105	NETWORK_COMMIT_NETWORK_CHANGES_FAILED	Error	Failed to commit network changes on <code>\${VdsName}</code>
1106	NETWORK_HOST_USING_WRONG_CLUSTER_VLAN	Warning	<code>\${VdsName}</code> is having wrong vlan id: <code>\${VlanIdHost}</code> , expected vlan id: <code>\${VlanIdCluster}</code>
1107	NETWORK_HOST_MISSING_CLUSTER_VLAN	Warning	<code>\${VdsName}</code> is missing vlan id: <code>\${VlanIdCluster}</code> that is expected by the cluster
1108	VDS_NETWORK_MTU_DIFFER_FROM_LOGICAL_NETWORK	Info	
1109	BRIDGED_NETWORK_OVER_MULTIPLE_INTERFACES	Warning	Bridged network <code>\${NetworkName}</code> is attached to multiple interfaces: <code>\${Interfaces}</code> on Host <code>\${VdsName}</code> .
1110	VDS_NETWORKS_OUT_OF_SYNC	Warning	Host <code>\${VdsName}</code> 's following network(s) are not synchronized with their Logical Network configuration: <code>\${Networks}</code> .

Code	Name	Severity	Message
1111	VM_MIGRATION_FAILED_DURING_MOVE_TO_MAINTENANCE_NO_DESTINATION_VDS	Error	Migration failed\${DueToMigrationError} while Source Host is in 'preparing for maintenance' state.\n Consider manual intervention\: stopping/migrating Vms as Host's state will not\n turn to maintenance while VMs are still running on it.(VM: \${VmName}, Source: \${VdsName}).
1112	NETWORK_UPDATED_NETWORK_ON_CLUSTER	Info	Network \${NetworkName} on Cluster \${ClusterName} updated.
1113	NETWORK_UPDATED_NETWORK_ON_CLUSTER_FAILED	Error	Failed to update Network \${NetworkName} on Cluster \${ClusterName}.
1114	NETWORK_UPDATE_NETWORK	Info	Network \${NetworkName} was updated on Data Center: \${StoragePoolName}
1115	NETWORK_UPDATE_NETWORK_FAILED	Error	Failed to update Network \${NetworkName} on Data Center: \${StoragePoolName}
1116	NETWORK_UPDATE_VM_INTERFACE_LINK_UP	Info	Link State is UP.
1117	NETWORK_UPDATE_VM_INTERFACE_LINK_DOWN	Info	Link State is DOWN.

Code	Name	Severity	Message
1118	INVALID_BOND_INTERFACE_FOR_MANAGEMENT_NETWORK_CONFIGURATION	Error	Failed to configure management network on host <code>\${VdsName}</code> . Host <code>\${VdsName}</code> has an invalid bond interface (<code>\${InterfaceName}</code> contains less than 2 active slaves) for the management network configuration.
1119	VLAN_ID_MISMATCH_FOR_MANAGEMENT_NETWORK_CONFIGURATION	Error	Failed to configure management network on host <code>\${VdsName}</code> . Host <code>\${VdsName}</code> has an interface <code>\${InterfaceName}</code> for the management network configuration with VLAN-ID (<code>\${VlanId}</code>), which is different from data-center definition (<code>\${MgmtVlanId}</code>).
1120	SETUP_NETWORK_FAILED_FOR_MANAGEMENT_NETWORK_CONFIGURATION	Error	Failed to configure management network on host <code>\${VdsName}</code> due to setup networks failure.
1121	PERSIST_NETWORK_FAILED_FOR_MANAGEMENT_NETWORK	Warning	Failed to configure management network on host <code>\${VdsName}</code> due to failure in persisting the management network configuration.

Code	Name	Severity	Message
1122	ADD_VNIC_PROFILE	Info	VM network interface profile <code>{VnicProfileName}</code> was added to network <code>{NetworkName}</code> in Data Center: <code>{DataCenterName}</code> . (User: <code>{UserName}</code>)
1123	ADD_VNIC_PROFILE_FAILED	Error	Failed to add VM network interface profile <code>{VnicProfileName}</code> to network <code>{NetworkName}</code> in Data Center: <code>{DataCenterName}</code> (User: <code>{UserName}</code>)
1124	UPDATE_VNIC_PROFILE	Info	VM network interface profile <code>{VnicProfileName}</code> was updated for network <code>{NetworkName}</code> in Data Center: <code>{DataCenterName}</code> . (User: <code>{UserName}</code>)
1125	UPDATE_VNIC_PROFILE_FAILED	Error	Failed to update VM network interface profile <code>{VnicProfileName}</code> for network <code>{NetworkName}</code> in Data Center: <code>{DataCenterName}</code> . (User: <code>{UserName}</code>)
1126	REMOVE_VNIC_PROFILE	Info	VM network interface profile <code>{VnicProfileName}</code> was removed from network <code>{NetworkName}</code> in Data Center: <code>{DataCenterName}</code> . (User: <code>{UserName}</code>)

Code	Name	Severity	Message
1127	REMOVE_VNIC_PROFILE_FAILED	Error	Failed to remove VM network interface profile <code>{VnicProfileName}</code> from network <code>{NetworkName}</code> in Data Center: <code>{DataCenterName}</code> . (User: <code>{UserName}</code>)
1128	NETWORK_WITHOUT_INTERFACES	Warning	Network <code>{NetworkName}</code> is not attached to any interface on host <code>{VdsName}</code> .
1129	VNIC_PROFILE_UNSUPPORTED_FEATURES	Warning	VM <code>{VmName}</code> has network interface <code>{NicName}</code> which is using profile <code>{VnicProfile}</code> with unsupported feature(s) <code>'{UnsupportedFeatures}'</code> by VM cluster <code>{ClusterName}</code> (version <code>{CompatibilityVersion}</code>).
1131	REMOVE_NETWORK_BY_LABEL_FAILED	Error	Network <code>{Network}</code> cannot be removed from the following hosts: <code>{HostNames}</code> in data-center <code>{StoragePoolName}</code> .
1132	LABEL_NETWORK	Info	Network <code>{NetworkName}</code> was labeled <code>{Label}</code> in data-center <code>{StoragePoolName}</code> .

Code	Name	Severity	Message
1133	LABEL_NETWORK_FAILED	Error	Failed to label network <code>{NetworkName}</code> with label <code>{Label}</code> in data-center <code>{StoragePoolName}</code> . .
1134	UNLABEL_NETWORK	Info	Network <code>{NetworkName}</code> was unlabeled in data-center <code>{StoragePoolName}</code> . .
1135	UNLABEL_NETWORK_FAILED	Error	Failed to unlabeled network <code>{NetworkName}</code> in data-center <code>{StoragePoolName}</code> . .
1136	LABEL_NIC	Info	Network interface card <code>{NicName}</code> was labeled <code>{Label}</code> on host <code>{VdsName}</code> .
1137	LABEL_NIC_FAILED	Error	Failed to label network interface card <code>{NicName}</code> with label <code>{Label}</code> on host <code>{VdsName}</code> .
1138	UNLABEL_NIC	Info	Label <code>{Label}</code> was removed from network interface card <code>{NicName}</code> on host <code>{VdsName}</code> .
1139	UNLABEL_NIC_FAILED	Error	Failed to remove label <code>{Label}</code> from network interface card <code>{NicName}</code> on host <code>{VdsName}</code> .

Code	Name	Severity	Message
1140	SUBNET_REMOVED	Info	Subnet <code>\${SubnetName}</code> was removed from provider <code>\${ProviderName}</code> . (User: <code>\${UserName}</code>)
1141	SUBNET_REMOVAL_FAILED	Error	Failed to remove subnet <code>\${SubnetName}</code> from provider <code>\${ProviderName}</code> . (User: <code>\${UserName}</code>)
1142	SUBNET_ADDED	Info	Subnet <code>\${SubnetName}</code> was added on provider <code>\${ProviderName}</code> . (User: <code>\${UserName}</code>)
1143	SUBNET_ADDITION_FAILED	Error	Failed to add subnet <code>\${SubnetName}</code> on provider <code>\${ProviderName}</code> . (User: <code>\${UserName}</code>)
1144	CONFIGURE_NETWORK_BY_LABELS_WHEN_CHANGING_CLUSTER_FAILED	Error	Failed to configure networks on host <code>\${VdsName}</code> while changing its cluster.
1145	PERSIST_NETWORK_ON_HOST	Info	(<code>\${Sequence}</code> / <code>\${Total}</code>): Applying changes for network(s) <code>\${NetworkNames}</code> on host <code>\${VdsName}</code> . (User: <code>\${UserName}</code>)
1146	PERSIST_NETWORK_ON_HOST_FINISHED	Info	(<code>\${Sequence}</code> / <code>\${Total}</code>): Successfully applied changes for network(s) <code>\${NetworkNames}</code> on host <code>\${VdsName}</code> . (User: <code>\${UserName}</code>)

Code	Name	Severity	Message
1147	PERSIST_NETWORK_ON_HOST_FAILED	Error	{Sequence}/{Total}: Failed to apply changes for network(s) {NetworkNames} on host {VdsName}. (User: {UserName})
1148	MULTI_UPDATE_NETWORK_NOT_POSSIBLE	Warning	Cannot apply network {NetworkName} changes to hosts on unsupported data center {StoragePoolName}. (User: {UserName})
1149	REMOVE_PORT_FROM_EXTERNAL_PROVIDER_FAILED	Warning	Failed to remove vNIC {NicName} from external network provider {ProviderName}. The vNIC can be identified on the provider by device id {NicId}.
1150	IMPORTEXPONENT_EXPORT_VM	Info	Vm {VmName} was exported successfully to {StorageDomainName}
1151	IMPORTEXPONENT_EXPORT_VM_FAILED	Error	Failed to export Vm {VmName} to {StorageDomainName}
1152	IMPORTEXPONENT_IMPORT_VM	Info	Vm {VmName} was imported successfully to Data Center {StoragePoolName}, Cluster {ClusterName}

Code	Name	Severity	Message
1153	IMPORTEXPOR_T_IM PORT_VM_FAILED	Error	Failed to import Vm \${VmName} to Data Center \${StoragePoolName} , Cluster \${ClusterName}
1154	IMPORTEXPOR_T_RE MOVE_TEMPLATE	Info	Template \${VmTemplateName} was removed from \${StorageDomainNa me}
1155	IMPORTEXPOR_T_RE MOVE_TEMPLATE_F AILED	Error	Failed to remove Template \${VmTemplateName} from \${StorageDomainNa me}
1156	IMPORTEXPOR_T_EX PORT_TEMPLATE	Info	Template \${VmTemplateName} was exported successfully to \${StorageDomainNa me}
1157	IMPORTEXPOR_T_EX PORT_TEMPLATE_F AILED	Error	Failed to export Template \${VmTemplateName} to \${StorageDomainNa me}
1158	IMPORTEXPOR_T_IM PORT_TEMPLATE	Info	Template \${VmTemplateName} was imported successfully to Data Center \${StoragePoolName} , Cluster \${ClusterName}
1159	IMPORTEXPOR_T_IM PORT_TEMPLATE_F AILED	Error	Failed to import Template \${VmTemplateName} to Data Center \${StoragePoolName} , Cluster \${ClusterName}

Code	Name	Severity	Message
1160	IMPORTEXPOR_T_RE MOVE_VM	Info	Vm \${VmName} was removed from \${StorageDomainName}
1161	IMPORTEXPOR_T_RE MOVE_VM_FAILED	Error	Failed to remove Vm \${VmName} remove from \${StorageDomainName}
1162	IMPORTEXPOR_T_ST ARTING_EXPORT_VM	Info	Starting export Vm \${VmName} to \${StorageDomainName}
1163	IMPORTEXPOR_T_ST ARTING_IMPORT_TEMPLATE	Info	Starting to import Template \${VmTemplateName} to Data Center \${StoragePoolName}, Cluster \${ClusterName}
1164	IMPORTEXPOR_T_ST ARTING_EXPORT_TEMPLATE	Info	Starting to export Template \${VmTemplateName} to \${StorageDomainName}
1165	IMPORTEXPOR_T_ST ARTING_IMPORT_VM	Info	Starting to import Vm \${VmName} to Data Center \${StoragePoolName}, Cluster \${ClusterName}
1166	IMPORTEXPOR_T_ST ARTING_REMOVE_TEMPLATE	Info	Starting to remove Template \${VmTemplateName} remove \${StorageDomainName}
1167	IMPORTEXPOR_T_ST ARTING_REMOVE_VM	Info	Starting to remove Vm \${VmName} remove from \${StorageDomainName}

Code	Name	Severity	Message
1168	IMPORTEXPOR_FAILED_TO_IMPORT_VM	Warning	Failed to read VM '{ImportedVmName}' OVF, it may be corrupted. Underlying error message: {ErrorMessage}
1169	IMPORTEXPOR_FAILED_TO_IMPORT_TEMPLATE	Warning	Failed to read Template '{Template}' OVF, it may be corrupted. Underlying error message: {ErrorMessage}
1170	IMPORTEXPOR_IMPORT_TEMPLATE_INVALID_INTERFACE	Normal	While importing Template {EntityName}, the Network/s {Networks} were found to be Non-VM Networks or do not exist in Cluster. Network Name was not set in the Interface/s {Interfaces}.
1171	USER_ACCOUNT_PASSWORD_EXPIRED	Error	User {UserName} cannot login, as the user account password has expired. Please contact the system administrator.
1172	AUTH_FAILED_INVALID_CREDENTIALS	Error	User {UserName} cannot login, please verify the username and password.
1173	AUTH_FAILED_CLOCK_SKEW_TOO_LARGE	Error	User {UserName} cannot login, the engine clock is not synchronized with directory services. Please contact the system administrator.

Code	Name	Severity	Message
1174	AUTH_FAILED_NO_KDCS_FOUND	Error	User \${UserName} cannot login, authentication domain cannot be found. Please contact the system administrator.
1175	AUTH_FAILED_DNS_ERROR	Error	User \${UserName} cannot login, there's an error in DNS configuration. Please contact the system administrator.
1176	AUTH_FAILED_OTHER	Error	User \${UserName} cannot login, unknown kerberos error. Please contact the system administrator.
1177	AUTH_FAILED_DNS_COMMUNICATION_ERROR	Error	User \${UserName} cannot login, cannot lookup DNS for SRV records. Please contact the system administrator.
1178	AUTH_FAILED_CONNECTION_TIMED_OUT	Error	User \${UserName} cannot login, connection to LDAP server has timed out. Please contact the system administrator.
1179	AUTH_FAILED_WRONG_REALM	Error	User \${UserName} cannot login, please verify your domain name.

Code	Name	Severity	Message
1180	AUTH_FAILED_CONNECTION_ERROR	Error	User \${UserName} cannot login, connection refused or some configuration problems exist. Possible DNS error. Please contact the system administrator.
1181	AUTH_FAILED_CANNOT_FIND_LDAP_SERVER_FOR_DOMAIN	Error	User \${UserName} cannot login, cannot find valid LDAP server for domain. Please contact the system administrator.
1182	AUTH_FAILED_NO_USER_INFORMATION_WAS_FOUND	Error	User \${UserName} cannot login, no user information was found. Please contact the system administrator.
1183	AUTH_FAILED_CLIENT_NOT_FOUND_IN_KERBEROS_DATABASE	Error	User \${UserName} cannot login, user was not found in domain. Please contact the system administrator.
1184	AUTH_FAILED_INTERNAL_KERBEROS_ERROR	Error	User \${UserName} cannot login, an internal error has occurred in the Kerberos implementation of the JVM. Please contact the system administrator.
1185	USER_ACCOUNT_EXPIRED	Error	The account for \${UserName} got expired. Please contact the system administrator.

Code	Name	Severity	Message
1186	IMPORTEXPONENT_PROXY_HOST_AVAILABLE_IN_DC	Error	No Host in Data Center '\${StoragePoolName}' can serve as a proxy to retrieve remote VMs information (User: \${UserName}).
1187	IMPORTEXPONENT_CANNOT_SERVE_AS_PROXY	Error	Host \${VdsName} cannot be used as a proxy to retrieve remote VMs information since it is not up (User: \${UserName}).
1188	IMPORTEXPONENT_PARTIAL_VM_MISSING_ENTITIES	Warning	The following entities could not be verified and will not be part of the imported VM \${VmName}: '\${MissingEntities}' (User: \${UserName}).
1189	IMPORTEXPONENT_IMPORT_VM_FAILED_UPDATING_OVF	Error	Failed to import Vm \${VmName} to Data Center \${StoragePoolName}, Cluster \${ClusterName}, could not update VM data in export.
1190	USER_RESTORE_FROM_SNAPSHOT_START	Info	Restoring VM \${VmName} from snapshot started by user \${UserName}.
1191	VM_DISK_ALREADY_CHANGED	Info	CD \${DiskName} is already inserted to VM \${VmName}, disk change action was skipped. User: \${UserName}.

Code	Name	Severity	Message
1192	VM_DISK_ALREADY_EJECTED	Info	CD is already ejected from VM <code>{VmName}</code> , disk change action was skipped. User: <code>{UserName}</code> .
1193	IMPORTEXPOR_T_STARTING_CONVERT_VM	Info	Starting to convert Vm <code>{VmName}</code>
1194	IMPORTEXPOR_T_CONVERT_FAILED	Info	Failed to convert Vm <code>{VmName}</code>
1195	IMPORTEXPOR_T_CANNOT_GET_OVF	Info	Failed to get the configuration of converted Vm <code>{VmName}</code>
1196	IMPORTEXPOR_T_INVALID_OVF	Info	Failed to process the configuration of converted Vm <code>{VmName}</code>
1197	IMPORTEXPOR_T_PARTIAL_TEMPLATE_MISSING_ENTITIES	Warning	The following entities could not be verified and will not be part of the imported Template <code>{VmTemplateName}</code> : <code>'{MissingEntities}'</code> (User: <code>{UserName}</code>).
1200	ENTITY_RENAMED	Info	<code>{EntityType}</code> <code>{OldEntityName}</code> was renamed from <code>{OldEntityName}</code> to <code>{NewEntityName}</code> by <code>{UserName}</code> .
1201	UPDATE_HOST_NIC_VFS_CONFIG	Info	The VFs configuration of network interface card <code>{NicName}</code> on host <code>{VdsName}</code> was updated.

Code	Name	Severity	Message
1202	UPDATE_HOST_NIC_VFS_CONFIG_FAILED	Error	Failed to update the VFs configuration of network interface card \${NicName} on host \${VdsName}.
1203	ADD_VFS_CONFIG_NETWORK	Info	Network \${NetworkName} was added to the VFs configuration of network interface card \${NicName} on host \${VdsName}.
1204	ADD_VFS_CONFIG_NETWORK_FAILED	Info	Failed to add \${NetworkName} to the VFs configuration of network interface card \${NicName} on host \${VdsName}.
1205	REMOVE_VFS_CONFIG_NETWORK	Info	Network \${NetworkName} was removed from the VFs configuration of network interface card \${NicName} on host \${VdsName}.
1206	REMOVE_VFS_CONFIG_NETWORK_FAILED	Info	Failed to remove \${NetworkName} from the VFs configuration of network interface card \${NicName} on host \${VdsName}.
1207	ADD_VFS_CONFIG_LABEL	Info	Label \${Label} was added to the VFs configuration of network interface card \${NicName} on host \${VdsName}.

Code	Name	Severity	Message
1208	ADD_VFS_CONFIG_LABEL_FAILED	Info	Failed to add <code>{Label}</code> to the VFs configuration of network interface card <code>{NicName}</code> on host <code>{VdsName}</code> .
1209	REMOVE_VFS_CONFIG_LABEL	Info	Label <code>{Label}</code> was removed from the VFs configuration of network interface card <code>{NicName}</code> on host <code>{VdsName}</code> .
1210	REMOVE_VFS_CONFIG_LABEL_FAILED	Info	Failed to remove <code>{Label}</code> from the VFs configuration of network interface card <code>{NicName}</code> on host <code>{VdsName}</code> .
1211	USER_REDUCE_MAIN_DEVICES_STARTED	Info	Started to reduce Storage <code>{StorageDomainName}</code> devices. (User: <code>{UserName}</code>).
1212	USER_REDUCE_MAIN_DEVICES_FAILED_METADATA_DEVICES	Error	Failed to reduce Storage <code>{StorageDomainName}</code> . The following devices contains the domain metadata <code>{deviceIds}</code> and can't be reduced from the domain. (User: <code>{UserName}</code>).
1213	USER_REDUCE_MAIN_DEVICES_FAILED	Error	Failed to reduce Storage <code>{StorageDomainName}</code> . (User: <code>{UserName}</code>).
1214	USER_REDUCE_MAIN_DEVICES_SUCCEEDED	Info	Storage <code>{StorageDomainName}</code> has been reduced. (User: <code>{UserName}</code>).

Code	Name	Severity	Message
1215	USER_REDUCE_DO MAIN_DEVICES_FAI LED_NO_FREE_SPA CE	Error	Can't reduce Storage \${StorageDomainNa me}. There is not enough space on the destination devices of the storage domain. (User: \${UserName}).
1216	USER_REDUCE_DO MAIN_DEVICES_FAI LED_TO_GET_DOM AIN_INFO	Error	Can't reduce Storage \${StorageDomainNa me}. Failed to get the domain info. (User: \${UserName}).
1217	CANNOT_IMPORT_V M_WITH_LEASE_CO MPAT_VERSION	Warning	The VM \${VmName} has a VM lease defined yet will be imported without it as the VM compatibility version does not support VM leases.
1218	CANNOT_IMPORT_V M_WITH_LEASE_ST ORAGE_DOMAIN	Warning	The VM \${VmName} has a VM lease defined yet will be imported without it as the Storage Domain for the lease does not exist or is not active.
1219	FAILED_DETERMINE _STORAGE_DOMAI N_METADATA_DEVI CES	Error	Failed to determine the metadata devices of Storage Domain \${StorageDomainNa me}.
1220	HOT_PLUG_LEASE_ FAILED	Error	Failed to hot plug lease to the VM \${VmName}. The VM is running without a VM lease.
1221	HOT_UNPLUG_LEA SE_FAILED	Error	Failed to hot unplug lease to the VM \${VmName}.

Code	Name	Severity	Message
1222	DETACH_DOMAIN_WITH_VMS_AND_TEMPLATES_LEASES	Warning	The deactivated domain <code>{storageDomainName}</code> contained leases for the following VMs/Templates: <code>{entitiesNames}</code> , a part of those VMs will not run and need manual removal of the VM leases.
1223	IMPORTEXPARTING_EXPORT_VM_TO_OVA	Info	Starting to export Vm <code>{VmName}</code> as a Virtual Appliance
1224	IMPORTEXPART_VM_TO_OVA	Info	Vm <code>{VmName}</code> was exported successfully as a Virtual Appliance to path <code>{OvaPath}</code> on Host <code>{VdsName}</code>
1225	IMPORTEXPART_VM_TO_OVA_FAILED	Error	Failed to export Vm <code>{VmName}</code> as a Virtual Appliance to path <code>{OvaPath}</code> on Host <code>{VdsName}</code>
1226	IMPORTEXPARTING_EXPORT_TEMPLATE_TO_OVA	Info	Starting to export Template <code>{VmTemplateName}</code> as a Virtual Appliance
1227	IMPORTEXPART_TEMPLATE_TO_OVA	Info	Template <code>{VmTemplateName}</code> was exported successfully as a Virtual Appliance to path <code>{OvaPath}</code> on Host <code>{VdsName}</code>

Code	Name	Severity	Message
1228	IMPORTEXPOR_TEMPLA_T O_OVA_FAILED	Error	Failed to export Template {VmTemplateName} as a Virtual Appliance to path {OvaPath} on Host {VdsName}
1300	NUMA_ADD_VM_NUMA_NODE_SUCCESS	Info	Add VM NUMA node successfully.
1301	NUMA_ADD_VM_NUMA_NODE_FAILED	Error	Add VM NUMA node failed.
1310	NUMA_UPDATE_VM_NUMA_NODE_SUCCESS	Info	Update VM NUMA node successfully.
1311	NUMA_UPDATE_VM_NUMA_NODE_FAILED	Error	Update VM NUMA node failed.
1320	NUMA_REMOVE_VM_NUMA_NODE_SUCCESS	Info	Remove VM NUMA node successfully.
1321	NUMA_REMOVE_VM_NUMA_NODE_FAILED	Error	Remove VM NUMA node failed.
1322	USER_ADD_VM_TEMPLATE_CREATE_TEMPLATE_FAILURE	Error	Failed to create Template {VmTemplateName} or its disks from VM {VmName}.
1323	USER_ADD_VM_TEMPLATE_ASSIGN_ILLEGAL_FAILURE	Error	Failed preparing Template {VmTemplateName} for sealing (VM: {VmName}).
1324	USER_ADD_VM_TEMPLATE_SEAL_FAILURE	Error	Failed to seal Template {VmTemplateName} (VM: {VmName}).

Code	Name	Severity	Message
1325	USER_SPARSIFY_IMAGE_START	Info	Started to sparsify \${DiskAlias}
1326	USER_SPARSIFY_IMAGE_FINISH_SUCCESS	Info	\${DiskAlias} sparsified successfully.
1327	USER_SPARSIFY_IMAGE_FINISH_FAILURE	Error	Failed to sparsify \${DiskAlias}.
1328	USER_AMEND_IMAGE_START	Info	Started to amend \${DiskAlias}
1329	USER_AMEND_IMAGE_FINISH_SUCCESS	Info	\${DiskAlias} has been amended successfully.
1330	USER_AMEND_IMAGE_FINISH_FAILURE	Error	Failed to amend \${DiskAlias}.
1340	VM_DOES_NOT_FIT_TO_SINGLE_NUMA_NODE	Warning	VM \${VmName} does not fit to a single NUMA node on host \${HostName}. This may negatively impact its performance. Consider using vNUMA and NUMA pinning for this VM.
1400	ENTITY_RENAMED INTERNALLY	Info	\${EntityType} \${OldEntityName} was renamed from \${OldEntityName} to \${NewEntityName}.
1402	USER_LOGIN_ON_BEHALF_FAILED	Error	Failed to execute login on behalf - \${LoginOnBehalfLogInfo}.

Code	Name	Severity	Message
1403	IRS_CONFIRMED_DISK_SPACE_LOW	Warning	Warning, low confirmed disk space. \${StorageDomainName} domain has \${DiskSpace} GB of confirmed free space.
2000	USER_HOTPLUG_DISK	Info	VM \${VmName} disk \${DiskAlias} was plugged by \${UserName}.
2001	USER_FAILED_HOTPLUG_DISK	Error	Failed to plug disk \${DiskAlias} to VM \${VmName} (User: \${UserName}).
2002	USER_HOTUNPLUG_DISK	Info	VM \${VmName} disk \${DiskAlias} was unplugged by \${UserName}.
2003	USER_FAILED_HOTUNPLUG_DISK	Error	Failed to unplug disk \${DiskAlias} from VM \${VmName} (User: \${UserName}).
2004	USER_COPIED_DISK	Info	User \${UserName} is copying disk \${DiskAlias} to domain \${StorageDomainName}.
2005	USER_FAILED_COPY_DISK	Error	User \${UserName} failed to copy disk \${DiskAlias} to domain \${StorageDomainName}.

Code	Name	Severity	Message
2006	USER_COPIED_DISK_FINISHED_SUCCESS	Info	User \${UserName} finished copying disk \${DiskAlias} to domain \${StorageDomainName}.
2007	USER_COPIED_DISK_FINISHED_FAILURE	Error	User \${UserName} finished with error copying disk \${DiskAlias} to domain \${StorageDomainName}.
2008	USER_MOVED_DISK	Info	User \${UserName} moving disk \${DiskAlias} to domain \${StorageDomainName}.
2009	USER_FAILED_MOVED_VM_DISK	Error	User \${UserName} failed to move disk \${DiskAlias} to domain \${StorageDomainName}.
2010	USER_MOVED_DISK_FINISHED_SUCCESS	Info	User \${UserName} finished moving disk \${DiskAlias} to domain \${StorageDomainName}.
2011	USER_MOVED_DISK_FINISHED_FAILURE	Error	User \${UserName} have failed to move disk \${DiskAlias} to domain \${StorageDomainName}.
2012	USER_FINISHED_REMOVE_DISK_NO_DOMAIN	Info	Disk \${DiskAlias} was successfully removed (User \${UserName}).

Code	Name	Severity	Message
2013	USER_FINISHED_FAILED_REMOVE_DISK_NO_DOMAIN	Warning	Failed to remove disk <code>{DiskAlias}</code> (User <code>{UserName}</code>).
2014	USER_FINISHED_REMOVE_DISK	Info	Disk <code>{DiskAlias}</code> was successfully removed from domain <code>{StorageDomainName}</code> (User <code>{UserName}</code>).
2015	USER_FINISHED_FAILED_REMOVE_DISK	Warning	Failed to remove disk <code>{DiskAlias}</code> from storage domain <code>{StorageDomainName}</code> (User: <code>{UserName}</code>).
2016	USER_ATTACH_DISK_TO_VM	Info	Disk <code>{DiskAlias}</code> was successfully attached to VM <code>{VmName}</code> by <code>{UserName}</code> .
2017	USER_FAILED_ATTACH_DISK_TO_VM	Error	Failed to attach Disk <code>{DiskAlias}</code> to VM <code>{VmName}</code> (User: <code>{UserName}</code>).
2018	USER_DETACH_DISK_FROM_VM	Info	Disk <code>{DiskAlias}</code> was successfully detached from VM <code>{VmName}</code> by <code>{UserName}</code> .
2019	USER_FAILED_DETACH_DISK_FROM_VM	Error	Failed to detach Disk <code>{DiskAlias}</code> from VM <code>{VmName}</code> (User: <code>{UserName}</code>).
2020	USER_ADD_DISK	Info	Add-Disk operation of <code>{DiskAlias}</code> was initiated by <code>{UserName}</code> .
2021	USER_ADD_DISK_FINISHED_SUCCESS	Info	The disk <code>{DiskAlias}</code> was successfully added.

Code	Name	Severity	Message
2022	USER_ADD_DISK_FINISHED_FAILURE	Error	Add-Disk operation failed to complete.
2023	USER_FAILED_ADD_DISK	Error	Add-Disk operation failed (User: \${UserName}).
2024	USER_RUN_UNLOCK_ENTITY_SCRIPT	Info	
2025	USER_MOVE_IMAGE_GROUP_FAILED_TO_DELETE_SRC_IMAGE	Warning	Possible failure while deleting \${DiskAlias} from the source Storage Domain \${StorageDomainName} during the move operation. The Storage Domain may be manually cleaned-up from possible leftovers (User:\${UserName}).
2026	USER_MOVE_IMAGE_GROUP_FAILED_TO_DELETE_DST_IMAGE	Warning	Possible failure while clearing possible leftovers of \${DiskAlias} from the target Storage Domain \${StorageDomainName} after the move operation failed to copy the image to it properly. The Storage Domain may be manually cleaned-up from possible leftovers (User:\${UserName}).
2027	USER_IMPORT_IMAGE	Info	User \${UserName} importing image \${RepolImageName} to domain \${StorageDomainName}.

Code	Name	Severity	Message
2028	USER_IMPORT_IMAGE_FINISHED_SUCCESS	Info	User \${UserName} successfully imported image \${RepolImageName} to domain \${StorageDomainName}.
2029	USER_IMPORT_IMAGE_FINISHED_FAILURE	Error	User \${UserName} failed to import image \${RepolImageName} to domain \${StorageDomainName}.
2030	USER_EXPORT_IMAGE	Info	User \${UserName} exporting image \${RepolImageName} to domain \${DestinationStorageDomainName}.
2031	USER_EXPORT_IMAGE_FINISHED_SUCCESS	Info	User \${UserName} successfully exported image \${RepolImageName} to domain \${DestinationStorageDomainName}.
2032	USER_EXPORT_IMAGE_FINISHED_FAILURE	Error	User \${UserName} failed to export image \${RepolImageName} to domain \${DestinationStorageDomainName}.
2033	HOT_SET_NUMBER_OF_CPUS	Info	Hotplug CPU: changed the number of CPUs on VM \${vmName} from \${previousNumberOfCpus} to \${numberOfCpus}

Code	Name	Severity	Message
2034	FAILED_HOT_SET_NUMBER_OF_CPUS	Error	Failed to hot set number of CPUs to VM <code>{vmName}</code> . Underlying error message: <code>{ErrorMessage}</code>
2035	USER_ISCSI_BOND_HOST_RESTART_WARNING	Warning	The following Networks has been removed from the iSCSI bond <code>{IscsiBondName}</code> : <code>{NetworkNames}</code> . for those changes to take affect, the hosts must be moved to maintenance and activated again.
2036	ADD_DISK_INTERNAL	Info	Add-Disk operation of <code>{DiskAlias}</code> was initiated by the system.
2037	ADD_DISK_INTERNAL_FAILURE	Info	Add-Disk operation of <code>{DiskAlias}</code> failed to complete.
2038	USER_REMOVE_DISK_INITIATED	Info	Removal of Disk <code>{DiskAlias}</code> from domain <code>{StorageDomainName}</code> was initiated by <code>{UserName}</code> .
2039	HOT_SET_MEMORY	Info	Hotset memory: changed the amount of memory on VM <code>{vmName}</code> from <code>{previousMem}</code> to <code>{newMem}</code>
2040	FAILED_HOT_SET_MEMORY	Error	Failed to hot set memory to VM <code>{vmName}</code> . Underlying error message: <code>{ErrorMessage}</code>

Code	Name	Severity	Message
2041	DISK_PREALLOCAT ION_FAILED	Error	
2042	USER_FINISHED_RE MOVE_DISK_ATTAC HED_TO_VMS	Info	Disk \${DiskAlias} associated to the VMs \${VmNames} was successfully removed from domain \${StorageDomainName} (User \${UserName}).
2043	USER_FINISHED_RE MOVE_DISK_ATTAC HED_TO_VMS_NO_ DOMAIN	Info	Disk \${DiskAlias} associated to the VMs \${VmNames} was successfully removed (User \${UserName}).
2044	USER_REMOVE_DIS K_ATTACHED_TO_V MS_INITIATED	Info	Removal of Disk \${DiskAlias} associated to the VMs \${VmNames} from domain \${StorageDomainName} was initiated by \${UserName} .
2045	USER_COPY_IMAGE _GROUP_FAILED_T O_DELETE_DST_IM AGE	Warning	Possible failure while clearing possible leftovers of \${DiskAlias} from the target Storage Domain \${StorageDomainName} after the operation failed. The Storage Domain may be manually cleaned-up from possible leftovers (User: \${UserName}).

Code	Name	Severity	Message
2046	MEMORY_HOT_UNPLUG_SUCCESSFULLY_REQUESTED	Info	Hot unplug of memory device (<code>{deviceId}</code>) of size <code>{memoryDeviceSize}MB</code> was successfully requested on VM ' <code>{vmName}</code> '. Physical memory guaranteed updated from <code>{oldMinMemoryMb}MB</code> to <code>{newMinMemoryMb}MB</code> .
2047	MEMORY_HOT_UNPLUG_FAILED	Error	Failed to hot unplug memory device (<code>{deviceId}</code>) of size <code>{memoryDeviceSize}MiB</code> out of VM ' <code>{vmName}</code> ': <code>{errorMessage}</code>
2048	FAILED_HOT_SET_MEMORY_NOT_DIVISIBLE	Error	Failed to hot plug memory to VM <code>{vmName}</code> . Amount of added memory (<code>{memoryAdded}MiB</code>) is not dividable by <code>{requiredFactor}MiB</code> .

Code	Name	Severity	Message
2049	MEMORY_HOT_UNPLUG_SUCCESSFULLY_REQUESTED_PLUS_MEMORY_INFO	Info	Hot unplug of memory device (<code>{deviceId}</code>) of size <code>{memoryDeviceSize}MiB</code> was successfully requested on VM ' <code>{vmName}</code> '. Defined Memory updated from <code>{oldMemoryMb}MiB</code> to <code>{newMemoryMb}MiB</code> . Physical memory guaranteed updated from <code>{oldMinMemoryMb}MiB</code> to <code>{newMinMemoryMb}MiB</code> .
2050	NO_MEMORY_DEVICE_TO_HOT_UNPLUG	Info	Defined memory can't be decreased. There are no hot plugged memory devices on VM <code>{vmName}</code> .
2051	NO_SUITABLE_MEMORY_DEVICE_TO_HOT_UNPLUG	Info	There is no memory device to hot unplug to satisfy request to decrement memory from <code>{oldMemoryMb}MiB</code> to <code>{newMemoryMB}MiB</code> on VM <code>{vmName}</code> . Available memory devices (decremented memory sizes): <code>{memoryHotUnplugOptions}</code> .
3000	USER_ADD_QUOTA	Info	Quota <code>{QuotaName}</code> has been added by <code>{UserName}</code> .

Code	Name	Severity	Message
3001	USER_FAILED_ADD_QUOTA	Error	Failed to add Quota <code>\${QuotaName}</code> . The operation was initiated by <code>\${UserName}</code> .
3002	USER_UPDATE_QUOTA	Info	Quota <code>\${QuotaName}</code> has been updated by <code>\${UserName}</code> .
3003	USER_FAILED_UPDATE_QUOTA	Error	Failed to update Quota <code>\${QuotaName}</code> . The operation was initiated by <code>\${UserName}</code> ..
3004	USER_DELETE_QUOTA	Info	Quota <code>\${QuotaName}</code> has been deleted by <code>\${UserName}</code> .
3005	USER_FAILED_DELETE_QUOTA	Error	Failed to delete Quota <code>\${QuotaName}</code> . The operation was initiated by <code>\${UserName}</code> ..
3006	USER_EXCEEDED_QUOTA_CLUSTER_GRACE_LIMIT	Error	Cluster-Quota <code>\${QuotaName}</code> limit exceeded and operation was blocked. Utilization: <code>\${Utilization}</code> , Requested: <code>\${Requested}</code> - Please select a different quota or contact your administrator to extend the quota.

Code	Name	Severity	Message
3007	USER_EXCEEDED_Q UOTA_CLUSTER_LI MIT	Warning	Cluster-Quota \${QuotaName} limit exceeded and entered the grace zone. Utilization: \${Utilization} (It is advised to select a different quota or contact your administrator to extend the quota).
3008	USER_EXCEEDED_Q UOTA_CLUSTER_TH RESHOLD	Warning	Cluster-Quota \${QuotaName} is about to exceed. Utilization: \${Utilization}
3009	USER_EXCEEDED_Q UOTA_STORAGE_G RACE_LIMIT	Error	Storage-Quota \${QuotaName} limit exceeded and operation was blocked. Utilization(used/requ ested): \${CurrentStorage}%/ \${Requested}% - Please select a different quota or contact your administrator to extend the quota.
3010	USER_EXCEEDED_Q UOTA_STORAGE_LI MIT	Warning	Storage-Quota \${QuotaName} limit exceeded and entered the grace zone. Utilization: \${CurrentStorage}% (It is advised to select a different quota or contact your administrator to extend the quota).

Code	Name	Severity	Message
3011	USER_EXCEEDED_QUOTA_STORAGE_THRESHOLD	Warning	Storage-Quota \${QuotaName} is about to exceed. Utilization: \${CurrentStorage}%
3012	QUOTA_STORAGE_RESIZE_LOWER_THAN_CONSUMPTION	Warning	Storage-Quota \${QuotaName} : the new size set for this quota is less than current disk utilization.
3013	MISSING_QUOTA_STORAGE_PARAMETERS_PERMISSIVE_MODE	Warning	Missing Quota for Disk, proceeding since in Permissive (Audit) mode.
3014	MISSING_QUOTA_CLUSTER_PARAMETERS_PERMISSIVE_MODE	Warning	Missing Quota for VM \${VmName} , proceeding since in Permissive (Audit) mode.
3015	USER_EXCEEDED_QUOTA_CLUSTER_GRACE_LIMIT_PERMISSIVE_MODE	Warning	Cluster-Quota \${QuotaName} limit exceeded, proceeding since in Permissive (Audit) mode. Utilization: \${Utilization} , Requested: \${Requested} - Please select a different quota or contact your administrator to extend the quota.

Code	Name	Severity	Message
3016	USER_EXCEEDED_QUOTA_STORAGE_GRACE_LIMIT_PERMISSIVE_MODE	Warning	Storage-Quota \${QuotaName} limit exceeded, proceeding since in Permissive (Audit) mode. Utilization(used/requested): \${CurrentStorage}%/ \${Requested}% - Please select a different quota or contact your administrator to extend the quota.
3017	USER_IMPORT_IMAGE_AS_TEMPLATE	Info	User \${UserName} importing image \${RepolImageName} as template \${TemplateName} to domain \${StorageDomainName} .
3018	USER_IMPORT_IMAGE_AS_TEMPLATE_FINISHED_SUCCESS	Info	User \${UserName} successfully imported image \${RepolImageName} as template \${TemplateName} to domain \${StorageDomainName} .
3019	USER_IMPORT_IMAGE_AS_TEMPLATE_FINISHED_FAILURE	Error	User \${UserName} failed to import image \${RepolImageName} as template \${TemplateName} to domain \${StorageDomainName} .

Code	Name	Severity	Message
4000	GLUSTER_VOLUME_CREATE	Info	Gluster Volume <code>\${glusterVolumeName}</code> created on cluster <code>\${clusterName}</code> .
4001	GLUSTER_VOLUME_CREATE_FAILED	Error	Creation of Gluster Volume <code>\${glusterVolumeName}</code> failed on cluster <code>\${clusterName}</code> .
4002	GLUSTER_VOLUME_OPTION_ADDED	Info	Volume Option <code>\${Key}</code>
4003	GLUSTER_VOLUME_OPTION_SET_FAILED	Error	Volume Option <code>\${Key}</code>
4004	GLUSTER_VOLUME_START	Info	Gluster Volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> started.
4005	GLUSTER_VOLUME_START_FAILED	Error	Could not start Gluster Volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4006	GLUSTER_VOLUME_STOP	Info	Gluster Volume <code>\${glusterVolumeName}</code> stopped on cluster <code>\${clusterName}</code> .
4007	GLUSTER_VOLUME_STOP_FAILED	Error	Could not stop Gluster Volume <code>\${glusterVolumeName}</code> on cluster <code>\${clusterName}</code> .
4008	GLUSTER_VOLUME_OPTIONS_RESET	Info	Volume Option <code>\${Key}</code>

Code	Name	Severity	Message
4009	GLUSTER_VOLUME_OPTIONS_RESET_FAILED	Error	Could not reset Gluster Volume <code>{glusterVolumeName}</code> Options on cluster <code>{clusterName}</code> .
4010	GLUSTER_VOLUME_DELETE	Info	Gluster Volume <code>{glusterVolumeName}</code> deleted on cluster <code>{clusterName}</code> .
4011	GLUSTER_VOLUME_DELETE_FAILED	Error	Could not delete Gluster Volume <code>{glusterVolumeName}</code> on cluster <code>{clusterName}</code> .
4012	GLUSTER_VOLUME_REBALANCE_START	Info	Gluster Volume <code>{glusterVolumeName}</code> rebalance started on cluster <code>{clusterName}</code> .
4013	GLUSTER_VOLUME_REBALANCE_START_FAILED	Error	Could not start Gluster Volume <code>{glusterVolumeName}</code> rebalance on cluster <code>{clusterName}</code> .
4014	GLUSTER_VOLUME_REMOVE_BRICKS	Info	Bricks removed from Gluster Volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> .
4015	GLUSTER_VOLUME_REMOVE_BRICKS_FAILED	Error	Could not remove bricks from Gluster Volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> .

Code	Name	Severity	Message
4016	GLUSTER_VOLUME_REPLACE_BRICK_FAILED	Error	Replace Gluster Volume <code>\${glusterVolumeName}</code> Brick failed on cluster <code>\${clusterName}</code>
4017	GLUSTER_VOLUME_REPLACE_BRICK_START	Info	Gluster Volume <code>\${glusterVolumeName}</code> Replace Brick started on cluster <code>\${clusterName}</code> .
4018	GLUSTER_VOLUME_REPLACE_BRICK_START_FAILED	Error	Could not start Gluster Volume <code>\${glusterVolumeName}</code> Replace Brick on cluster <code>\${clusterName}</code> .
4019	GLUSTER_VOLUME_ADD_BRICK	Info	<code>\${NoOfBricks}</code> brick(s) added to volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4020	GLUSTER_VOLUME_ADD_BRICK_FAILED	Error	Failed to add bricks to the Gluster Volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4021	GLUSTER_SERVER_REMOVE_FAILED	Error	Failed to remove host <code>\${VdsName}</code> from Cluster <code>\${ClusterName}</code> .
4022	GLUSTER_VOLUME_PROFILE_START	Info	Gluster Volume <code>\${glusterVolumeName}</code> profiling started on cluster <code>\${clusterName}</code> .

Code	Name	Severity	Message
4023	GLUSTER_VOLUME_PROFILE_START_FAILED	Error	Could not start profiling on gluster volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code>
4024	GLUSTER_VOLUME_PROFILE_STOP	Info	Gluster Volume <code>{glusterVolumeName}</code> profiling stopped on cluster <code>{clusterName}</code> .
4025	GLUSTER_VOLUME_PROFILE_STOP_FAILED	Error	Could not stop Profiling on gluster volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> .
4026	GLUSTER_VOLUME_CREATED_FROM_CLI	Warning	Detected new volume <code>{glusterVolumeName}</code> on cluster <code>{ClusterName}</code> , and added it to engine DB.
4027	GLUSTER_VOLUME_DELETED_FROM_CLI	Info	Detected deletion of volume <code>{glusterVolumeName}</code> on cluster <code>{ClusterName}</code> , and deleted it from engine DB.
4028	GLUSTER_VOLUME_OPTION_SET_FROM_CLI	Warning	Detected new option <code>{key}</code>
4029	GLUSTER_VOLUME_OPTION_RESET_FROM_CLI	Warning	Detected option <code>{key}</code>

Code	Name	Severity	Message
4030	GLUSTER_VOLUME_PROPERTIES_CHANGED_FROM_CLI	Warning	Detected changes in properties of volume <code>{glusterVolumeName}</code> of cluster <code>{ClusterName}</code> , and updated the same in engine DB.
4031	GLUSTER_VOLUME_BRICK_ADDED_FROM_CLI	Warning	Detected new brick <code>{brick}</code> on volume <code>{glusterVolumeName}</code> of cluster <code>{ClusterName}</code> , and added it to engine DB.
4032	GLUSTER_VOLUME_BRICK_REMOVED_FROM_CLI	Info	Detected brick <code>{brick}</code> removed from Volume <code>{glusterVolumeName}</code> of cluster <code>{ClusterName}</code> , and removed it from engine DB.
4033	GLUSTER_SERVER_REMOVED_FROM_CLI	Info	Detected server <code>{VdsName}</code> removed from Cluster <code>{ClusterName}</code> , and removed it from engine DB.
4034	GLUSTER_VOLUME_INFO_FAILED	Error	Failed to fetch gluster volume list from server <code>{VdsName}</code> .
4035	GLUSTER_COMMAND_FAILED	Error	Gluster command <code>{Command}</code> failed on server <code>{Server}</code> .
4038	GLUSTER_SERVER_REMOVE	Info	Host <code>{VdsName}</code> removed from Cluster <code>{ClusterName}</code> .

Code	Name	Severity	Message
4039	GLUSTER_VOLUME_STARTED_FROM_CLI	Warning	Detected that Volume <code>\${glusterVolumeName}</code> of Cluster <code>\${ClusterName}</code> was started, and updated engine DB with it's new status.
4040	GLUSTER_VOLUME_STOPPED_FROM_CLI	Warning	Detected that Volume <code>\${glusterVolumeName}</code> of Cluster <code>\${ClusterName}</code> was stopped, and updated engine DB with it's new status.
4041	GLUSTER_VOLUME_OPTION_CHANGED_FROM_CLI	Info	Detected change in value of option <code>\${key}</code> from <code>\${oldValue}</code> to <code>\${newValue}</code> on volume <code>\${glusterVolumeName}</code> of cluster <code>\${ClusterName}</code> , and updated it to engine DB.
4042	GLUSTER_HOOK_ENABLED	Info	Gluster Hook <code>\${GlusterHookName}</code> enabled on cluster <code>\${ClusterName}</code> .
4043	GLUSTER_HOOK_ENABLED_FAILED	Error	Failed to enable Gluster Hook <code>\${GlusterHookName}</code> on cluster <code>\${ClusterName}</code> . <code>\${FailureMessage}</code>
4044	GLUSTER_HOOK_ENABLED_PARTIAL	Warning	Gluster Hook <code>\${GlusterHookName}</code> enabled on some of the servers on cluster <code>\${ClusterName}</code> . <code>\${FailureMessage}</code>

Code	Name	Severity	Message
4045	GLUSTER_HOOK_DISABLE	Info	Gluster Hook <code>{GlusterHookName}</code> disabled on cluster <code>{ClusterName}</code> .
4046	GLUSTER_HOOK_DISABLE_FAILED	Error	Failed to disable Gluster Hook <code>{GlusterHookName}</code> on cluster <code>{ClusterName}</code> . <code>{FailureMessage}</code>
4047	GLUSTER_HOOK_DISABLE_PARTIAL	Warning	Gluster Hook <code>{GlusterHookName}</code> disabled on some of the servers on cluster <code>{ClusterName}</code> . <code>{FailureMessage}</code>
4048	GLUSTER_HOOK_LIST_FAILED	Error	Failed to retrieve hook list from <code>{VdsName}</code> of Cluster <code>{ClusterName}</code> .
4049	GLUSTER_HOOK_CONFLICT_DETECTED	Warning	Detected conflict in hook <code>{HookName}</code> of Cluster <code>{ClusterName}</code> .
4050	GLUSTER_HOOK_DETECTED_NEW	Info	Detected new hook <code>{HookName}</code> in Cluster <code>{ClusterName}</code> .
4051	GLUSTER_HOOK_DETECTED_DELETE	Info	Detected removal of hook <code>{HookName}</code> in Cluster <code>{ClusterName}</code> .
4052	GLUSTER_VOLUME_OPTION_MODIFIED	Info	Volume Option <code>{Key}</code> changed to <code>{Value}</code> from <code>{oldvalue}</code> on <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> .

Code	Name	Severity	Message
4053	GLUSTER_HOOK_GETCONTENT_FAILED	Error	Failed to read content of hook <code>{HookName}</code> in Cluster <code>{ClusterName}</code> .
4054	GLUSTER_SERVICE_LIST_FAILED	Error	Could not fetch statuses of services from server <code>{VdsName}</code> . Updating statuses of all services on this server to UNKNOWN.
4055	GLUSTER_SERVICE_TYPE_ADDED_TO_CLUSTER	Info	Service type <code>{ServiceType}</code> was not mapped to cluster <code>{ClusterName}</code> . Mapped it now.
4056	GLUSTER_CLUSTER_SERVICE_STATUSES_CHANGED	Info	Status of service type <code>{ServiceType}</code> changed from <code>{OldStatus}</code> to <code>{NewStatus}</code> on cluster <code>{ClusterName}</code>
4057	GLUSTER_SERVICE_ADDED_TO_SERVER	Info	Service <code>{ServiceName}</code> was not mapped to server <code>{VdsName}</code> . Mapped it now.
4058	GLUSTER_SERVER_SERVICE_STATUS_CHANGED	Info	Status of service <code>{ServiceName}</code> on server <code>{VdsName}</code> changed from <code>{OldStatus}</code> to <code>{NewStatus}</code> . Updating in engine now.

Code	Name	Severity	Message
4059	GLUSTER_HOOK_UPDATED	Info	Gluster Hook <code>{GlusterHookName}</code> updated on conflicting servers.
4060	GLUSTER_HOOK_UPDATE_FAILED	Error	Failed to update Gluster Hook <code>{GlusterHookName}</code> on conflicting servers. <code>{FailureMessage}</code>
4061	GLUSTER_HOOK_ADDED	Info	Gluster Hook <code>{GlusterHookName}</code> added on conflicting servers.
4062	GLUSTER_HOOK_ADD_FAILED	Error	Failed to add Gluster Hook <code>{GlusterHookName}</code> on conflicting servers. <code>{FailureMessage}</code>
4063	GLUSTER_HOOK_REMOVED	Info	Gluster Hook <code>{GlusterHookName}</code> removed from all servers in cluster <code>{ClusterName}</code> .
4064	GLUSTER_HOOK_REMOVE_FAILED	Error	Failed to remove Gluster Hook <code>{GlusterHookName}</code> from cluster <code>{ClusterName}</code> . <code>{FailureMessage}</code>
4065	GLUSTER_HOOK_REFRESH	Info	Refreshed gluster hooks in Cluster <code>{ClusterName}</code> .
4066	GLUSTER_HOOK_REFRESH_FAILED	Error	Failed to refresh gluster hooks in Cluster <code>{ClusterName}</code> .

Code	Name	Severity	Message
4067	GLUSTER_SERVICE_STARTED	Info	`\${servicetype}` service started on host `\${VdsName}` of cluster `\${ClusterName}`.
4068	GLUSTER_SERVICE_START_FAILED	Error	Could not start `\${servicetype}` service on host `\${VdsName}` of cluster `\${ClusterName}`.
4069	GLUSTER_SERVICE_STOPPED	Info	`\${servicetype}` services stopped on host `\${VdsName}` of cluster `\${ClusterName}`.
4070	GLUSTER_SERVICE_STOP_FAILED	Error	Could not stop `\${servicetype}` service on host `\${VdsName}` of cluster `\${ClusterName}`.
4071	GLUSTER_SERVICE_S_LIST_NOT_FETCHED	Info	Could not fetch list of services from `\${ServiceGroupType}` named `\${ServiceGroupName}`.
4072	GLUSTER_SERVICE_RESTARTED	Info	`\${servicetype}` service re-started on host `\${VdsName}` of cluster `\${ClusterName}`.
4073	GLUSTER_SERVICE_RESTART_FAILED	Error	Could not re-start `\${servicetype}` service on host `\${VdsName}` of cluster `\${ClusterName}`.

Code	Name	Severity	Message
4074	GLUSTER_VOLUME_OPTIONS_RESET_ALL	Info	All Volume Options reset on <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4075	GLUSTER_HOST_UID_NOT_FOUND	Error	Could not find gluster uuid of server <code>\${VdsName}</code> on Cluster <code>\${ClusterName}</code> .
4076	GLUSTER_VOLUME_BRICK_ADDED	Info	Brick <code>[\${brickpath}]</code> on host <code>[\${servername}]</code> added to volume <code>[\${glusterVolumeName}]</code> of cluster <code>\${clusterName}</code>
4077	GLUSTER_CLUSTER_SERVICE_STATUSES_ADDED	Info	Status of service type <code>\${ServiceType}</code> set to <code>\${NewStatus}</code> on cluster <code>\${ClusterName}</code>
4078	GLUSTER_VOLUME_REBALANCE_STOP	Info	Gluster Volume <code>\${glusterVolumeName}</code> rebalance stopped of cluster <code>\${clusterName}</code> .
4079	GLUSTER_VOLUME_REBALANCE_STOP_FAILED	Error	Could not stop rebalance of gluster volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4080	START_REMOVING_GLUSTER_VOLUME_BRICKS	Info	Started removing bricks from Volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>

Code	Name	Severity	Message
4081	START_REMOVING_GLUSTER_VOLUME_BRICKS_FAILED	Error	Could not start remove bricks from Volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4082	GLUSTER_VOLUME_REMOVE_BRICKS_STOP	Info	Stopped removing bricks from Volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4083	GLUSTER_VOLUME_REMOVE_BRICKS_STOP_FAILED	Error	Failed to stop remove bricks from Volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4084	GLUSTER_VOLUME_REMOVE_BRICKS_COMMIT	Info	Gluster volume <code>\${glusterVolumeName}</code> remove bricks committed on cluster <code>\${clusterName}</code> . <code>\${NoOfBricks}</code> brick(s) removed from volume <code>\${glusterVolumeName}</code> .
4085	GLUSTER_VOLUME_REMOVE_BRICKS_COMMIT_FAILED	Error	Gluster volume <code>\${glusterVolumeName}</code> remove bricks could not be committed on cluster <code>\${clusterName}</code>

Code	Name	Severity	Message
4086	GLUSTER_BRICK_STATUS_CHANGED	Warning	Detected change in status of brick <code>{brickpath}</code> of volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> from <code>{oldValue}</code> to <code>{newValue}</code> via <code>{source}</code> .
4087	GLUSTER_VOLUME_REBALANCE_FINISHED	Info	<code>{action}</code> <code>{status}</code> on volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> .
4088	GLUSTER_VOLUME_MIGRATE_BRICK_DATA_FINISHED	Info	<code>{action}</code> <code>{status}</code> for brick(s) on volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> . Please review to abort or commit.
4089	GLUSTER_VOLUME_REBALANCE_START_DETECTED_FROM_CLI	Info	Detected start of rebalance on volume <code>{glusterVolumeName}</code> of Cluster <code>{ClusterName}</code> from CLI.
4090	START_REMOVING_GLUSTER_VOLUME_BRICKS_DETECTED_FROM_CLI	Info	Detected start of brick removal for bricks <code>{brick}</code> on volume <code>{glusterVolumeName}</code> of Cluster <code>{ClusterName}</code> from CLI.

Code	Name	Severity	Message
4091	GLUSTER_VOLUME_REBALANCE_NOT_FOUND_FROM_CLI	Warning	Could not find information for rebalance on volume <code>\${glusterVolumeName}</code> of Cluster <code>\${ClusterName}</code> from CLI. Marking it as unknown.
4092	REMOVE_GLUSTER_VOLUME_BRICKS_NOT_FOUND_FROM_CLI	Warning	Could not find information for remove brick on volume <code>\${glusterVolumeName}</code> of Cluster <code>\${ClusterName}</code> from CLI. Marking it as unknown.
4093	GLUSTER_VOLUME_DETAILS_REFRESH	Info	Refreshed details of the volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4094	GLUSTER_VOLUME_DETAILS_REFRESH_FAILED	Error	Failed to refresh the details of volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4095	GLUSTER_HOST_UUID_ALREADY_EXISTS	Error	Gluster UUID of host <code>\${VdsName}</code> on Cluster <code>\${ClusterName}</code> already exists.
4096	USER_FORCE_SELECTED_SPM_STOP_FAILED	Error	Failed to force select <code>\${VdsName}</code> as the SPM due to a failure to stop the current SPM.

Code	Name	Severity	Message
4097	GLUSTER_GEOREP_SESSION_DELETE_FROM_CLI	Warning	Detected deletion of geo-replication session <code>\${geoRepSessionKey}</code> from volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4098	GLUSTER_GEOREP_SESSION_DETECTED_FROM_CLI	Warning	Detected new geo-replication session <code>\${geoRepSessionKey}</code> for volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> . Adding it to engine.
4099	GLUSTER_GEOREP_SESSION_REFRESH	Info	Refreshed geo-replication sessions for volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4100	GLUSTER_GEOREP_SESSION_REFRESH_FAILED	Error	Failed to refresh geo-replication sessions for volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4101	GEOREP_SESSION_STOP	Info	Geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> has been stopped.
4102	GEOREP_SESSION_STOP_FAILED	Error	Failed to stop geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>

Code	Name	Severity	Message
4103	GEOREP_SESSION_DELETED	Info	Geo-replication session deleted on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4104	GEOREP_SESSION_DELETE_FAILED	Error	Failed to delete geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4105	GLUSTER_GEOREP_CONFIG_SET	Info	Configuration <code>\${key}</code> has been set to <code>\${value}</code> on the geo-rep session <code>\${geoRepSessionKey}</code> .
4106	GLUSTER_GEOREP_CONFIG_SET_FAILED	Error	Failed to set the configuration <code>\${key}</code> to <code>\${value}</code> on geo-rep session <code>\${geoRepSessionKey}</code> .
4107	GLUSTER_GEOREP_CONFIG_LIST	Info	Refreshed configuration options for geo-replication session <code>\${geoRepSessionKey}</code>
4108	GLUSTER_GEOREP_CONFIG_LIST_FAILED	Error	Failed to refresh configuration options for geo-replication session <code>\${geoRepSessionKey}</code>
4109	GLUSTER_GEOREP_CONFIG_SET_DEFAULT	Info	Configuration of <code>\${key}</code> of session <code>\${geoRepSessionKey}</code> reset to its default value .

Code	Name	Severity	Message
4110	GLUSTER_GEOREP_CONFIG_SET_DEFAULT_FAILED	Error	Failed to set \${key} of session \${geoRepSessionKey} to its default value.
4111	GLUSTER_VOLUME_SNAPSHOT_DELETED	Info	Gluster volume snapshot \${snapname} deleted.
4112	GLUSTER_VOLUME_SNAPSHOT_DELETE_FAILED	Error	Failed to delete gluster volume snapshot \${snapname}.
4113	GLUSTER_VOLUME_ALL_SNAPSHOTS_DELETED	Info	Deleted all the gluster volume snapshots for the volume \${glusterVolumeName} of cluster \${clusterName}.
4114	GLUSTER_VOLUME_ALL_SNAPSHOTS_DELETE_FAILED	Error	Failed to delete all the gluster volume snapshots for the volume \${glusterVolumeName} of cluster \${clusterName}.
4115	GLUSTER_VOLUME_SNAPSHOT_ACTIVATED	Info	Activated the gluster volume snapshot \${snapname} on volume \${glusterVolumeName} of cluster \${clusterName}.
4116	GLUSTER_VOLUME_SNAPSHOT_ACTIVATION_FAILED	Error	Failed to activate the gluster volume snapshot \${snapname} on volume \${glusterVolumeName} of cluster \${clusterName}.

Code	Name	Severity	Message
4117	GLUSTER_VOLUME_SNAPSHOT_DEACTIVATED	Info	De-activated the gluster volume snapshot <code>{snapname}</code> on volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> .
4118	GLUSTER_VOLUME_SNAPSHOT_DEACTIVATE_FAILED	Error	Failed to de-activate gluster volume snapshot <code>{snapname}</code> on volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> .
4119	GLUSTER_VOLUME_SNAPSHOT_RESTORED	Info	Restored the volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> to the state of gluster volume snapshot <code>{snapname}</code> .
4120	GLUSTER_VOLUME_SNAPSHOT_RESTORE_FAILED	Error	Failed to restore the volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> to the state of gluster volume snapshot <code>{snapname}</code> .
4121	GLUSTER_VOLUME_SNAPSHOT_CONFIG_UPDATED	Info	Updated Gluster volume snapshot configuration(s).
4122	GLUSTER_VOLUME_SNAPSHOT_CONFIG_UPDATE_FAILED	Error	Failed to update gluster volume snapshot configuration(s).
4123	GLUSTER_VOLUME_SNAPSHOT_CONFIG_UPDATE_FAILED_PARTIALLY	Error	Failed to update gluster volume snapshot configuration(s) <code>{failedSnapshotConfigs}</code> .

Code	Name	Severity	Message
4124	NEW_STORAGE_DEVICE_DETECTED	Info	Found new storage device <code>{storageDevice}</code> on host <code>{VdsName}</code> , and added it to engine DB."
4125	STORAGE_DEVICE_REMOVED_FROM_THE_HOST	Info	Detected deletion of storage device <code>{storageDevice}</code> on host <code>{VdsName}</code> , and deleting it from engine DB."
4126	SYNC_STORAGE_DEVICES_IN_HOST	Info	Manually synced the storage devices from host <code>{VdsName}</code>
4127	SYNC_STORAGE_DEVICES_IN_HOST_FAILED	Error	Failed to sync storage devices from host <code>{VdsName}</code>
4128	GEOREP_OPTION_SET_FROM_CLI	Warning	Detected new option <code>{key}</code>
4129	GEOREP_OPTION_CHANGED_FROM_CLI	Warning	Detected change in value of option <code>{key}</code> from <code>{oldValue}</code> to <code>{value}</code> for geo-replication session on volume <code>{glusterVolumeName}</code> of cluster <code>{ClusterName}</code> , and updated it to engine.
4130	GLUSTER_MASTER_VOLUME_STOP_FAILED_DURING_SNAPSHOT_RESTORE	Error	Could not stop master volume <code>{glusterVolumeName}</code> of cluster <code>{clusterName}</code> during snapshot restore.

Code	Name	Severity	Message
4131	GLUSTER_MASTER_VOLUME_SNAPSHOT_RESTORE_FAILED	Error	Could not restore master volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4132	GLUSTER_VOLUME_SNAPSHOT_CREATED	Info	Snapshot <code>\${snapname}</code> created for volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4133	GLUSTER_VOLUME_SNAPSHOT_CREATE_FAILED	Error	Could not create snapshot for volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4134	GLUSTER_VOLUME_SNAPSHOT_SCHEDULED	Info	Snapshots scheduled on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4135	GLUSTER_VOLUME_SNAPSHOT_SCHEDULE_FAILED	Error	Failed to schedule snapshots on the volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4136	GLUSTER_VOLUME_SNAPSHOT_RESCHEDULED	Info	Rescheduled snapshots on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .
4137	GLUSTER_VOLUME_SNAPSHOT_RESCHEDULE_FAILED	Error	Failed to reschedule snapshots on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> .

Code	Name	Severity	Message
4138	CREATE_GLUSTER_BRICK	Info	Brick <code>\${brickName}</code> created successfully on host <code>\${vdsName}</code> of cluster <code>\${clusterName}</code> .
4139	CREATE_GLUSTER_BRICK_FAILED	Error	Failed to create brick <code>\${brickName}</code> on host <code>\${vdsName}</code> of cluster <code>\${clusterName}</code> .
4140	GLUSTER_GEO_REP_PUB_KEY_FETCH_FAILED	Error	Failed to fetch public keys.
4141	GLUSTER_GET_PUB_KEY	Info	Public key fetched.
4142	GLUSTER_GEOREP_PUBLIC_KEY_WRITE_FAILED	Error	Failed to write public keys to <code>\${VdsName}</code>
4143	GLUSTER_WRITE_PUB_KEYS	Info	Public keys written to <code>\${VdsName}</code>
4144	GLUSTER_GEOREP_SETUP_MOUNT_BROKER_FAILED	Error	Failed to setup geo-replication mount broker for user <code>\${geoRepUserName}</code> on the slave volume <code>\${geoRepSlaveVolumeName}</code> .
4145	GLUSTER_SETUP_GEOREP_MOUNT_BROKER	Info	Geo-replication mount broker has been setup for user <code>\${geoRepUserName}</code> on the slave volume <code>\${geoRepSlaveVolumeName}</code> .

Code	Name	Severity	Message
4146	GLUSTER_GEOREP_SESSION_CREATE_FAILED	Error	Failed to create geo-replication session between master volume : \${glusterVolumeName} of cluster \${clusterName} and slave volume : \${geoRepSlaveVolumeName} for the user \${geoRepUserName} .
4147	CREATE_GLUSTER_VOLUME_GEOREP_SESSION	Info	Created geo-replication session between master volume : \${glusterVolumeName} of cluster \${clusterName} and slave volume : \${geoRepSlaveVolumeName} for the user \${geoRepUserName} .
4148	GLUSTER_VOLUME_SNAPSHOT_SOFT_LIMIT_REACHED	Info	Gluster Volume Snapshot soft limit reached for the volume \${glusterVolumeName} on cluster \${clusterName}.
4149	HOST_FEATURES_INCOMPATIBLE_WITH_CLUSTER	Error	Host \${VdsName} does not comply with the list of features supported by cluster \${ClusterName}. \${UnSupportedFeature} is not supported by the Host
4150	GLUSTER_VOLUME_SNAPSHOT_SCHEDULE_DELETED	Info	Snapshot schedule deleted for volume \${glusterVolumeName} of \${clusterName}.

Code	Name	Severity	Message
4151	GLUSTER_BRICK_STATUS_DOWN	Info	Status of brick <code>{brickpath}</code> of volume <code>{glusterVolumeName}</code> on cluster <code>{ClusterName}</code> is down.
4152	GLUSTER_VOLUME_SNAPSHOT_DETECTED_NEW	Info	Found new gluster volume snapshot <code>{snapname}</code> for volume <code>{glusterVolumeName}</code> on cluster <code>{ClusterName}</code> , and added it to engine DB."
4153	GLUSTER_VOLUME_SNAPSHOT_DELETED_FROM_CLI	Info	Detected deletion of gluster volume snapshot <code>{snapname}</code> for volume <code>{glusterVolumeName}</code> on cluster <code>{ClusterName}</code> , and deleting it from engine DB."
4154	GLUSTER_VOLUME_SNAPSHOT_CLUSTER_CONFIG_DETECTED_NEW	Info	Found new gluster volume snapshot configuration <code>{snapConfigName}</code> with value <code>{snapConfigValue}</code> on cluster <code>{ClusterName}</code> , and added it to engine DB."

Code	Name	Severity	Message
4155	GLUSTER_VOLUME_SNAPSHOT_VOLUME_CONFIG_DETECTED_NEW	Info	Found new gluster volume snapshot configuration <code>\${snapConfigName}</code> with value <code>\${snapConfigValue}</code> for volume <code>\${glusterVolumeName}</code> on cluster <code>\${ClusterName}</code> , and added it to engine DB."
4156	GLUSTER_VOLUME_SNAPSHOT_HARD_LIMIT_REACHED	Info	Gluster Volume Snapshot hard limit reached for the volume <code>\${glusterVolumeName}</code> on cluster <code>\${clusterName}</code> .
4157	GLUSTER_CLI_SNAPSHOT_SCHEDULE_DISABLE_FAILED	Error	Failed to disable gluster CLI based snapshot schedule on cluster <code>\${clusterName}</code> .
4158	GLUSTER_CLI_SNAPSHOT_SCHEDULE_DISABLED	Info	Disabled gluster CLI based scheduling successfully on cluster <code>\${clusterName}</code> .
4159	SET_UP_PASSWORDLESS_SSH	Info	Password-less SSH has been setup for user <code>\${geoRepUserName}</code> on the nodes of remote volume <code>\${geoRepSlaveVolumeName}</code> from the nodes of the volume <code>\${glusterVolumeName}</code> .

Code	Name	Severity	Message
4160	SET_UP_PASSWORDLESS_SSH_FAILED	Error	Failed to setup Passwordless ssh for user <code>\${geoRepUserName}</code> on the nodes of remote volume <code>\${geoRepSlaveVolumeName}</code> from the nodes of the volume <code>\${glusterVolumeName}</code> .
4161	GLUSTER_VOLUME_TYPE_UNSUPPORTED	Warning	Detected a volume <code>\${glusterVolumeName}</code> with type <code>\${glusterVolumeType}</code> on cluster <code>\${Cluster}</code> and it is not fully supported by engine.
4162	GLUSTER_VOLUME_BRICK_REPLACED	Info	Replaced brick <code>'\${brick}'</code> with new brick <code>'\${newBrick}'</code> of Gluster Volume <code>\${glusterVolumeName}</code> on cluster <code>\${clusterName}</code>
4163	GLUSTER_SERVER_STATUS_DISCONNECTED	Info	Gluster server <code>\${vdsName}</code> set to DISCONNECTED on cluster <code>\${clusterName}</code> .
4164	GLUSTER_STORAGE_DOMAIN_SYNC_FAILED	Info	Failed to synchronize data from storage domain <code>\${storageDomainName}</code> to remote location.
4165	GLUSTER_STORAGE_DOMAIN_SYNCED	Info	Successfully synchronized data from storage domain <code>\${storageDomainName}</code> to remote location.

Code	Name	Severity	Message
4166	GLUSTER_STORAGE_DOMAIN_SYNC_STARTED	Info	Successfully started data synchronization data from storage domain <code>{storageDomainName}</code> to remote location.
4167	STORAGE_DOMAIN_DR_DELETED	Error	Deleted the data synchronization schedule for storage domain <code>{storageDomainName}</code> as the underlying geo-replication session <code>{geoRepSessionKey}</code> has been deleted.
4168	GLUSTER_WEBHOOK_ADDED	Info	Added webhook on <code>{clusterName}</code>
4169	GLUSTER_WEBHOOK_ADD_FAILED	Error	Failed to add webhook on <code>{clusterName}</code>
4170	GLUSTER_VOLUME_RESET_BRICK_FAILED	Error	
4171	GLUSTER_VOLUME_BRICK_RESETED	Info	
4172	GLUSTER_VOLUME_CONFIRMED_SPACE_LOW	Warning	Warning! Low confirmed free space on gluster volume <code>{glusterVolumeName}</code>
4436	GLUSTER_SERVER_ADD_FAILED	Error	Failed to add host <code>{VdsName}</code> into Cluster <code>{ClusterName}</code> . <code>{ErrorMessage}</code>

Code	Name	Severity	Message
4437	GLUSTER_SERVERS_LIST_FAILED	Error	Failed to fetch gluster peer list from server <code>\${VdsName}</code> on Cluster <code>\${ClusterName}</code> . <code>\${ErrorMessage}</code>
4595	GLUSTER_VOLUME_GEO_REP_START_FAILED_EXCEPTION	Error	Failed to start geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4596	GLUSTER_VOLUME_GEO_REP_START	Info	Geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> has been started.
4597	GLUSTER_VOLUME_GEO_REP_PAUSE_FAILED	Error	Failed to pause geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4598	GLUSTER_VOLUME_GEO_REP_RESUME_FAILED	Error	Failed to resume geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code>
4599	GLUSTER_VOLUME_GEO_REP_RESUME	Info	Geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> has been resumed.
4600	GLUSTER_VOLUME_GEO_REP_PAUSE	Info	Geo-replication session on volume <code>\${glusterVolumeName}</code> of cluster <code>\${clusterName}</code> has been paused.

Code	Name	Severity	Message
9000	VDS_ALERT_FENCE_IS_NOT_CONFIGURED	Info	Failed to verify Power Management configuration for Host \${VdsName}.
9001	VDS_ALERT_FENCE_TEST_FAILED	Info	Power Management test failed for Host \${VdsName}. \${Reason}
9002	VDS_ALERT_FENCE_OPERATION_FAILED	Info	Failed to power fence host \${VdsName}. Please check the host status and it's power management settings, and then manually reboot it and click "Confirm Host Has Been Rebooted"
9003	VDS_ALERT_FENCE_OPERATION_SKIPPED	Info	Host \${VdsName} became non responsive. Fence operation skipped as the system is still initializing and this is not a host where hosted engine was running on previously.
9004	VDS_ALERT_FENCE_NO_PROXY_HOST	Info	There is no other host in the data center that can be used to test the power management settings.
9005	VDS_ALERT_FENCE_STATUS_VERIFICATION_FAILED	Info	Failed to verify Host \${Host} \${Status} status, Please \${Status} Host \${Host} manually.

Code	Name	Severity	Message
9006	CANNOT_HIBERNATE_RUNNING_VMS_AFTER_CLUSTER_CPU_UPGRADE	Warning	Hibernation of VMs after CPU upgrade of Cluster <code>{Cluster}</code> is not supported. Please stop and restart those VMs in case you wish to hibernate them
9007	VDS_ALERT_SECONDARY_AGENT_USED_FOR_FENCE_OPERATION	Info	Secondary fence agent was used to <code>{Operation}</code> Host <code>{VdsName}</code>
9008	VDS_HOST_NOT_RESPONDING_CONNECTING	Warning	Host <code>{VdsName}</code> is not responding. It will stay in Connecting state for a grace period of <code>{Seconds}</code> seconds and after that an attempt to fence the host will be issued.
9009	VDS_ALERT_PM_HEALTH_CHECK_FENCE_AGENT_NON_RESPONSIVE	Info	Health check on Host <code>{VdsName}</code> indicates that Fence-Agent <code>{AgentId}</code> is non-responsive.
9010	VDS_ALERT_PM_HEALTH_CHECK_START_MIGHT_FAIL	Info	Health check on Host <code>{VdsName}</code> indicates that future attempts to Start this host using Power-Management are expected to fail.
9011	VDS_ALERT_PM_HEALTH_CHECK_STOP_MIGHT_FAIL	Info	Health check on Host <code>{VdsName}</code> indicates that future attempts to Stop this host using Power-Management are expected to fail.

Code	Name	Severity	Message
9012	VDS_ALERT_PM_HEALTH_CHECK_RESTART_MIGHT_FAIL	Info	Health check on Host \${VdsName} indicates that future attempts to Restart this host using Power-Management are expected to fail.
9013	VDS_ALERT_FENCE_OPERATION_SKIPPED_BROKEN_CONNECTIVITY	Info	Host \${VdsName} became non responsive and was not restarted due to Fencing Policy: \${Percents} percents of the Hosts in the Cluster have connectivity issues.
9014	VDS_ALERT_NOT_RESTARTED_DUE_TO_POLICY	Info	Host \${VdsName} became non responsive and was not restarted due to the Cluster Fencing Policy.
9015	VDS_ALERT_FENCE_DISABLED_BY_CLUSTER_POLICY	Info	Host \${VdsName} became Non Responsive and was not restarted due to disabled fencing in the Cluster Fencing Policy.
9016	FENCE_DISABLED_IN_CLUSTER_POLICY	Info	Fencing is disabled in Fencing Policy of the Cluster \${ClusterName}, so HA VMs running on a non-responsive host will not be restarted elsewhere.
9017	FENCE_OPERATION_STARTED	Info	Power management \${Action} of Host \${VdsName} initiated.

Code	Name	Severity	Message
9018	FENCE_OPERATION_SUCCEEDED	Info	Power management \${Action} of Host \${VdsName} succeeded.
9019	FENCE_OPERATION_FAILED	Error	Power management \${Action} of Host \${VdsName} failed.
9020	FENCE_OPERATION_USING_AGENT_AND_PROXY_STARTED	Info	Executing power management \${Action} on Host \${Host} using Proxy Host \${ProxyHost} and Fence Agent \${AgentType}:\${AgentType} .
9021	FENCE_OPERATION_USING_AGENT_AND_PROXY_FAILED	Warning	Execution of power management \${Action} on Host \${Host} using Proxy Host \${ProxyHost} and Fence Agent \${AgentType}:\${AgentType} failed.
9022	ENGINE_NO_FULL_BACKUP	Info	There is no full backup available, please run engine-backup to prevent data loss in case of corruption.
9023	ENGINE_NO_WARM_BACKUP	Info	Full backup was created on \${Date} and it's too old. Please run engine-backup to prevent data loss in case of corruption.
9024	ENGINE_BACKUP_STARTED	Normal	Engine backup started.
9025	ENGINE_BACKUP_COMPLETED	Normal	Engine backup completed successfully.

Code	Name	Severity	Message
9026	ENGINE_BACKUP_FAILED	Error	Engine backup failed.
9028	VDS_ALERT_NO_PM_CONFIG_FENCE_OPERATION_SKIPPED	Info	Host \${VdsName} became non responsive. It has no power management configured. Please check the host status, manually reboot it, and click "Confirm Host Has Been Rebooted"
9500	TASK_STOPPING_ASYNC_TASK	Info	Stopping async task \${CommandName} that started at \${Date}
9501	TASK_CLEARING_ASYNC_TASK	Info	Clearing asynchronous task \${CommandName} that started at \${Date}
9506	USER_ACTIVATE_STORAGE_DOMAIN_FAILED_ASYNC	Warning	Failed to autorecover Storage Domain \${StorageDomainName} (Data Center \${StoragePoolName}).
9600	IMPORTEXPOR_IMPORT_VM_INVALID_INTERFACES	Warning	While importing VM \${EntityName}, the Network/s \${Networks} were found to be Non-VM Networks or do not exist in Cluster or are missing a suitable VM network interface profile. Network Name was not set in the Interface/s \${Interfaces}.

Code	Name	Severity	Message
9601	VDS_SET_NON_OPERATIONAL_VM_NETWORK_IS_BRIDGELESS	Warning	Host <code>{VdsName}</code> does not comply with the cluster <code>{ClusterName}</code> networks, the following VM networks are non-VM networks: <code>{Networks}</code> . The host will become NonOperational.
9602	HA_VM_FAILED	Error	Highly Available VM <code>{VmName}</code> failed. It will be restarted automatically.
9603	HA_VM_RESTART_FAILED	Error	Restart of the Highly Available VM <code>{VmName}</code> failed.
9604	EMULATED_MACHINES_INCOMPATIBLE_WITH_CLUSTER	Warning	Host <code>{VdsName}</code> does not comply with the cluster <code>{ClusterName}</code> emulated machine. The cluster emulated machine is <code>{clusterEmulatedMachines}</code> and the host emulated machines are <code>{hostSupportedEmulatedMachines}</code> .
9605	EXCEEDED_MAXIMUM_NUM_OF_RESTART_HA_VM_ATTEMPTS	Error	Highly Available VM <code>{VmName}</code> could not be restarted automatically, exceeded the maximum number of attempts.

Code	Name	Severity	Message
9606	IMPORTEXPORT_SNAPSHOT_VM_INVALID_INTERFACES	Warning	While previewing a snapshot of VM \${EntityName} , the Network/s \${Networks} were found to be Non-VM Networks or do not exist in Cluster. Network Name was not set in the Interface/s \${Interfaces} .
9607	ADD_VM_FROM_SNAPSHOT_INVALID_INTERFACES	Warning	While adding vm \${EntityName} from snapshot, the Network/s \${Networks} were found to be Non-VM Networks or do not exist in Cluster. Network Name was not set in the Interface/s \${Interfaces} .
9608	RNG_SOURCES_INCOMPATIBLE_WITH_CLUSTER	Warning	Host \${VdsName} does not comply with the cluster \${ClusterName} Random Number Generator sources. The Hosts supported sources are: \${hostSupportedRngSources} ; and the cluster requirements are: \${clusterRequiredRngSources} .

Code	Name	Severity	Message
9609	EMULATED_MACHINES_INCOMPATIBLE_WITH_CLUSTER_LEVEL	Warning	Host <code>{VdsName}</code> does not comply with the cluster <code>{ClusterName}</code> emulated machines. The current cluster compatibility level supports <code>{clusterEmulatedMachines}</code> and the host emulated machines are <code>{hostSupportedEmulatedMachines}</code> .
9610	MIXING_RHEL_VERSIONS_IN_CLUSTER	Warning	Not possible to mix RHEL 6.x and 7.x hosts in one cluster. Tried adding <code>{addingRhel}</code> host to a cluster with <code>{previousRhel}</code> hosts.
9611	COLD_REBOOT_VM_DOWN	Info	VM <code>{VmName}</code> is down as a part of cold reboot process
9612	COLD_REBOOT_FAILED	Error	Cold reboot of VM <code>{VmName}</code> failed
9613	EXCEEDED_MAXIMUM_NUM_OF_COLD_REBOOT_VM_ATTEMPTS	Error	VM <code>{VmName}</code> could not be rebooted, exceeded the maximum number of attempts.
9700	DWH_STARTED	Info	ETL Service started.
9701	DWH_STOPPED	Info	ETL Service stopped.
9704	DWH_ERROR	Error	Error in ETL Service.

Code	Name	Severity	Message
9801	EXTERNAL_EVENT_NORMAL	Info	An external event with NORMAL severity has been added.
9802	EXTERNAL_EVENT_WARNING	Warning	An external event with WARNING severity has been added.
9803	EXTERNAL_EVENT_ERROR	Error	An external event with ERROR severity has been added.
9804	EXTERNAL_ALERT	Info	An external event with ALERT severity has been added.
9901	WATCHDOG_EVENT	Warning	Watchdog event ({wdaction}) triggered on {VmName} at {wdevent} (host time).
9910	USER_ADD_CLUSTER_POLICY	Info	Scheduling Policy {ClusterPolicy} was added. (User: {UserName})
9911	USER_FAILED_TO_ADD_CLUSTER_POLICY	Error	Failed to add Scheduling Policy: {ClusterPolicy}. (User: {UserName})
9912	USER_UPDATE_CLUSTER_POLICY	Info	Scheduling Policy {ClusterPolicy} was updated. (User: {UserName})
9913	USER_FAILED_TO_UPDATE_CLUSTER_POLICY	Error	Failed to update Scheduling Policy: {ClusterPolicy}. (User: {UserName})
9914	USER_REMOVE_CLUSTER_POLICY	Info	Scheduling Policy {ClusterPolicy} was removed. (User: {UserName})

Code	Name	Severity	Message
9915	USER_FAILED_TO_REMOVE_CLUSTER_POLICY	Error	Failed to remove Scheduling Policy: <code>\${ClusterPolicy}</code> . (User: <code>\${UserName}</code>)
9920	FAILED_TO_CONNECT_TO_SCHEDULER_PROXY	Error	Failed to connect to external scheduler proxy. External filters, scoring functions and load balancing will not be performed.
10000	VDS_UNTRUSTED	Error	Host <code>\${VdsName}</code> was set to non-operational. Host is not trusted by the attestation service.
10001	USER_UPDATE_VM_FROM_TRUSTED_TO_UNTRUSTED	Warning	The VM <code>\${VmName}</code> was updated from trusted cluster to non-trusted cluster.
10002	USER_UPDATE_VM_FROM_UNTRUSTED_TO_TRUSTED	Warning	The VM <code>\${VmName}</code> was updated from non-trusted cluster to trusted cluster.
10003	IMPORTEXPOR_IMPORT_VM_FROM_TRUSTED_TO_UNTRUSTED	Warning	The VM <code>\${VmName}</code> was created in trusted cluster and imported into a non-trusted cluster
10004	IMPORTEXPOR_IMPORT_VM_FROM_UNTRUSTED_TO_TRUSTED	Warning	The VM <code>\${VmName}</code> was created in non-trusted cluster and imported into a trusted cluster

Code	Name	Severity	Message
10005	USER_ADD_VM_FROM_TRUSTED_TO_UNTRUSTED	Warning	The VM <code>{VmName}</code> was created in an untrusted cluster. It was originated from the Template <code>{VmTemplateName}</code> which was created in a trusted cluster.
10006	USER_ADD_VM_FROM_UNTRUSTED_TO_TRUSTED	Warning	The VM <code>{VmName}</code> was created in a trusted cluster. It was originated from the Template <code>{VmTemplateName}</code> which was created in an untrusted cluster.
10007	IMPORTEXPOR_IMPORT_TEMPLATE_FROM_TRUSTED_TO_UNTRUSTED	Warning	The Template <code>{VmTemplateName}</code> was created in trusted cluster and imported into a non-trusted cluster
10008	IMPORTEXPOR_IMPORT_TEMPLATE_FROM_UNTRUSTED_TO_TRUSTED	Warning	The Template <code>{VmTemplateName}</code> was created in non-trusted cluster and imported into a trusted cluster
10009	USER_ADD_VM_TEMPLATE_FROM_TRUSTED_TO_UNTRUSTED	Warning	The non-trusted Template <code>{VmTemplateName}</code> was created from trusted Vm <code>{VmName}</code> .
10010	USER_ADD_VM_TEMPLATE_FROM_UNTRUSTED_TO_TRUSTED	Warning	The trusted template <code>{VmTemplateName}</code> was created from non-trusted Vm <code>{VmName}</code> .
10011	USER_UPDATE_VM_TEMPLATE_FROM_TRUSTED_TO_UNTRUSTED	Warning	The Template <code>{VmTemplateName}</code> was updated from trusted cluster to non-trusted cluster.

Code	Name	Severity	Message
10012	USER_UPDATE_VM_TEMPLATE_FROM_UNTRUSTED_TO_TRUSTED	Warning	The Template <code>{VmTemplateName}</code> was updated from non-trusted cluster to trusted cluster.
10013	IMPORTEXPOR_T_EXTERNAL_VMS_NOT_IN_DOWN_STATUS	Warning	The following VMs retrieved from external server <code>{URL}</code> are not in down status: <code>{Vms}</code> .
10100	USER_ADDED_NETWORK_QOS	Info	Network QoS <code>{QosName}</code> was added. (User: <code>{UserName}</code>)
10101	USER_FAILED_TO_ADD_NETWORK_QOS	Error	Failed to add Network QoS <code>{QosName}</code> . (User: <code>{UserName}</code>)
10102	USER_REMOVED_NETWORK_QOS	Info	Network QoS <code>{QosName}</code> was removed. (User: <code>{UserName}</code>)
10103	USER_FAILED_TO_REMOVE_NETWORK_QOS	Error	Failed to remove Network QoS <code>{QosName}</code> . (User: <code>{UserName}</code>)
10104	USER_UPDATED_NETWORK_QOS	Info	Network QoS <code>{QosName}</code> was updated. (User: <code>{UserName}</code>)
10105	USER_FAILED_TO_UPDATE_NETWORK_QOS	Error	Failed to update Network QoS <code>{QosName}</code> . (User: <code>{UserName}</code>)
10110	USER_ADDED_QOS	Info	QoS <code>{QoSName}</code> was added. (User: <code>{UserName}</code>)

Code	Name	Severity	Message
10111	USER_FAILED_TO_ADD_QOS	Error	Failed to add QoS \${QoSName}. (User: \${UserName})
10112	USER_REMOVED_QOS	Info	QoS \${QoSName} was removed. (User: \${UserName})
10113	USER_FAILED_TO_REMOVE_QOS	Error	Failed to remove QoS \${QoSName}. (User: \${UserName})
10114	USER_UPDATED_QOS	Info	QoS \${QoSName} was updated. (User: \${UserName})
10115	USER_FAILED_TO_UPDATE_QOS	Error	Failed to update QoS \${QoSName}. (User: \${UserName})
10120	USER_ADDED_DISK_PROFILE	Info	Disk Profile \${ProfileName} was successfully added (User: \${UserName}).
10121	USER_FAILED_TO_ADD_DISK_PROFILE	Error	Failed to add Disk Profile (User: \${UserName}).
10122	USER_REMOVED_DISK_PROFILE	Info	Disk Profile \${ProfileName} was successfully removed (User: \${UserName}).
10123	USER_FAILED_TO_REMOVE_DISK_PROFILE	Error	Failed to remove Disk Profile \${ProfileName} (User: \${UserName}).
10124	USER_UPDATED_DISK_PROFILE	Info	Disk Profile \${ProfileName} was successfully updated (User: \${UserName}).

Code	Name	Severity	Message
10125	USER_FAILED_TO_UPDATE_DISK_PROFILE	Error	Failed to update Disk Profile <code>{ProfileName}</code> (User: <code>{UserName}</code>).
10130	USER_ADDED_CPU_PROFILE	Info	CPU Profile <code>{ProfileName}</code> was successfully added (User: <code>{UserName}</code>).
10131	USER_FAILED_TO_ADD_CPU_PROFILE	Error	Failed to add CPU Profile (User: <code>{UserName}</code>).
10132	USER_REMOVED_CPU_PROFILE	Info	CPU Profile <code>{ProfileName}</code> was successfully removed (User: <code>{UserName}</code>).
10133	USER_FAILED_TO_REMOVE_CPU_PROFILE	Error	Failed to remove CPU Profile <code>{ProfileName}</code> (User: <code>{UserName}</code>).
10134	USER_UPDATED_CPU_PROFILE	Info	CPU Profile <code>{ProfileName}</code> was successfully updated (User: <code>{UserName}</code>).
10135	USER_FAILED_TO_UPDATE_CPU_PROFILE	Error	Failed to update CPU Profile <code>{ProfileName}</code> (User: <code>{UserName}</code>).
10200	USER_UPDATED_MOM_POLICIES	Info	Mom policy was updated on host <code>{VdsName}</code> .
10201	USER_FAILED_TO_UPDATE_MOM_POLICIES	Warning	Mom policy could not be updated on host <code>{VdsName}</code> .

Code	Name	Severity	Message
10250	PM_POLICY_UP_TO_MAINTENANCE	Info	Host \${Host} is not currently needed, activating maintenance mode in preparation for shutdown.
10251	PM_POLICY_MAINTENANCE_TO_DOWN	Info	Host \${Host} is not currently needed, shutting down.
10252	PM_POLICY_TO_UP	Info	Reactivating host \${Host} according to the current power management policy.
10300	CLUSTER_ALERT_HA_RESERVATION	Info	Cluster \${ClusterName} failed the HA Reservation check, HA VMs on host(s): \${Hosts} will fail to migrate in case of a failover, consider adding resources or shutting down unused VMs.
10301	CLUSTER_ALERT_HA_RESERVATION_DOWN	Info	Cluster \${ClusterName} passed the HA Reservation check.
10350	USER_ADDED_AFFINITY_GROUP	Info	Affinity Group \${affinityGroupName} was added. (User: \${UserName})
10351	USER_FAILED_TO_ADD_AFFINITY_GROUP	Error	Failed to add Affinity Group \${affinityGroupName}. (User: \${UserName})
10352	USER_UPDATED_AFFINITY_GROUP	Info	Affinity Group \${affinityGroupName} was updated. (User: \${UserName})

Code	Name	Severity	Message
10353	USER_FAILED_TO_UPDATE_AFFINITY_GROUP	Error	Failed to update Affinity Group <code>\${affinityGroupName}</code> . (User: <code>\${UserName}</code>)
10354	USER_REMOVED_AFFINITY_GROUP	Info	Affinity Group <code>\${affinityGroupName}</code> was removed. (User: <code>\${UserName}</code>)
10355	USER_FAILED_TO_REMOVE_AFFINITY_GROUP	Error	Failed to remove Affinity Group <code>\${affinityGroupName}</code> . (User: <code>\${UserName}</code>)
10356	VM_TO_HOST_CONFLICT_IN_ENFORCING_POSITIVE_AND_NEGATIVE_AFFINITY	Error	The affinity groups: <code>\${AffinityGroups}</code> , with hosts: <code>:\${Hosts}</code> and VMs: <code>:\${Vms}</code> , have VM to host conflicts between positive and negative enforcing affinity groups.
10357	VM_TO_HOST_CONFLICT_IN_POSITIVE_AND_NEGATIVE_AFFINITY	Warning	The affinity groups: <code>\${AffinityGroups}</code> , with hosts: <code>:\${Hosts}</code> and VMs: <code>:\${Vms}</code> , have VM to host conflicts between positive and negative affinity groups.
10358	VM_TO_HOST_CONFLICTS_POSITIVE_VM_TO_VM_AFFINITY	Warning	The affinity groups: <code>\${AffinityGroups}</code> , with hosts: <code>:\${Hosts}</code> and VMs: <code>:\${Vms}</code> , have conflicts between VM to host affinity and VM to VM positive affinity.

Code	Name	Severity	Message
10359	VM_TO_HOST_CONFLICTS_NEGATIVE_VM_TO_VM_AFFINITY	Warning	The affinity groups: \${AffinityGroups}, with hosts : \${Hosts} and VMs: \${Vms}, have conflicts between VM to host affinity and VM to VM negative affinity.
10360	NON_INTERSECTING_POSITIVE_HOSTS_AFFINITY_CONFLICTS	Warning	The affinity groups: \${AffinityGroups}, with hosts : \${Hosts} and VMs : \${Vms} , have non intersecting positive hosts conflicts.
10361	VM_TO_VM_AFFINITY_CONFLICTS	Error	
10380	USER_ADDED_AFFINITY_LABEL	Info	Affinity Label \${labelName} was added. (User: \${UserName})
10381	USER_FAILED_TO_ADD_AFFINITY_LABEL	Error	Failed to add Affinity Label \${labelName}. (User: \${UserName})
10382	USER_UPDATED_AFFINITY_LABEL	Info	Affinity Label \${labelName} was updated. (User: \${UserName})
10383	USER_FAILED_TO_UPDATE_AFFINITY_LABEL	Error	Failed to update Affinity Label \${labelName}. (User: \${UserName})
10384	USER_REMOVED_AFFINITY_LABEL	Info	Affinity Label \${labelName} was removed. (User: \${UserName})
10385	USER_FAILED_TO_REMOVE_AFFINITY_LABEL	Error	Failed to remove Affinity Label \${labelName}. (User: \${UserName})

Code	Name	Severity	Message
10400	ISCSI_BOND_ADD_SUCCESS	Info	iSCSI bond <code>'\${IscsiBondName}'</code> was successfully created in Data Center <code>'\${StoragePoolName}'</code> .
10401	ISCSI_BOND_ADD_FAILED	Error	Failed to create iSCSI bond <code>'\${IscsiBondName}'</code> in Data Center <code>'\${StoragePoolName}'</code> .
10402	ISCSI_BOND_EDIT_SUCCESS	Info	iSCSI bond <code>'\${IscsiBondName}'</code> was successfully updated.
10403	ISCSI_BOND_EDIT_FAILED	Error	Failed to update iSCSI bond <code>'\${IscsiBondName}'</code> .
10404	ISCSI_BOND_REMOVE_SUCCESS	Info	iSCSI bond <code>'\${IscsiBondName}'</code> was removed from Data Center <code>'\${StoragePoolName}'</code> .
10405	ISCSI_BOND_REMOVE_FAILED	Error	Failed to remove iSCSI bond <code>'\${IscsiBondName}'</code> from Data Center <code>'\${StoragePoolName}'</code> .
10406	ISCSI_BOND_EDIT_SUCCESS_WITH_WARNING	Warning	iSCSI bond <code>'\${IscsiBondName}'</code> was successfully updated but some of the hosts encountered connection issues.

Code	Name	Severity	Message
10407	ISCSI_BOND_ADD_SUCCESS_WITH_WARNING	Warning	iSCSI bond '{IscsiBondName}' was successfully created in Data Center '{StoragePoolName}' but some of the hosts encountered connection issues.
10450	USER_SET_HOSTED_ENGINE_MAINTENANCE	Info	Hosted Engine HA maintenance mode was updated on host \${VdsName}.
10451	USER_FAILED_TO_SET_HOSTED_ENGINE_MAINTENANCE	Error	Hosted Engine HA maintenance mode could not be updated on host \${VdsName}.
10452	VDS_MAINTENANCE_MANUAL_HA	Warning	Host \${VdsName} was switched to Maintenance mode, but Hosted Engine HA maintenance could not be enabled. Please enable it manually.
10453	USER_VDS_MAINTENANCE_MANUAL_HA	Warning	Host \${VdsName} was switched to Maintenance mode by \${UserName}, but Hosted Engine HA maintenance could not be enabled. Please enable it manually.

Code	Name	Severity	Message
10454	VDS_ACTIVATE_MA NUAL_HA	Warning	Host <code>{VdsName}</code> was activated by <code>{UserName}</code> , but the Hosted Engine HA service may still be in maintenance mode. If necessary, please correct this manually.
10455	VDS_ACTIVATE_MA NUAL_HA_ASYNC	Warning	Host <code>{VdsName}</code> was autorecovered, but the Hosted Engine HA service may still be in maintenance mode. If necessary, please correct this manually.
10456	HOSTED_ENGINE_V M_IMPORT_SUCCEE DED	Normal	Hosted Engine VM was imported successfully
10460	HOSTED_ENGINE_D OMAIN_IMPORT_SU CCEDED	Normal	Hosted Engine Storage Domain imported successfully
10461	HOSTED_ENGINE_D OMAIN_IMPORT_FAI LED	Error	Failed to import the Hosted Engine Storage Domain
10500	EXTERNAL_SCHED ULER_PLUGIN_ERR OR	Error	Running the external scheduler plugin <code>'{PluginName}'</code> failed: <code>'{ErrorMessage}'</code>
10501	EXTERNAL_SCHED ULER_ERROR	Error	Running the external scheduler failed: <code>'{ErrorMessage}'</code>
10550	VM_SLA_POLICY_C PU	Info	VM <code>{VmName}</code> SLA Policy was set. CPU limit is set to <code>{cpuLimit}</code>

Code	Name	Severity	Message
10551	VM_SLA_POLICY_STORAGE	Info	VM \${VmName} SLA Policy was set. Storage policy changed for disks: [\${diskList}]
10552	VM_SLA_POLICY_CPU_STORAGE	Info	VM \${VmName} SLA Policy was set. CPU limit is set to \${cpuLimit}. Storage policy changed for disks: [\${diskList}]
10553	FAILED_VM_SLA_POLICY	Error	Failed to set SLA Policy to VM \${VmName}. Underlying error message: \${ErrorMessage}
10600	USER_REMOVE_AUDIT_LOG	Info	Event list message \${AuditLogId} was removed by User \${UserName}.
10601	USER_REMOVE_AUDIT_LOG_FAILED	Error	User \${UserName} failed to remove event list message \${AuditLogId}.
10602	USER_CLEAR_ALL_AUDIT_LOG_EVENTS	Info	All events were removed. (User: \${UserName})
10603	USER_CLEAR_ALL_AUDIT_LOG_EVENTS_FAILED	Error	Failed to remove all events. (User: \${UserName})
10604	USER_DISPLAY_ALL_AUDIT_LOG	Info	All events were displayed. (User: \${UserName})
10605	USER_DISPLAY_ALL_AUDIT_LOG_FAILED	Error	Failed to display all events. (User: \${UserName})
10606	USER_CLEAR_ALL_AUDIT_LOG_ALERTS	Info	All alerts were removed. (User: \${UserName})

Code	Name	Severity	Message
10607	USER_CLEAR_ALL_AUDIT_LOG_ALERTS_FAILED	Error	Failed to remove all alerts. (User: \${UserName})
10700	MAC_POOL_ADD_SUCCESS	Info	MAC Pool '\${MacPoolName}' (id
10701	MAC_POOL_ADD_FAILED	Error	Failed to create MAC Pool '\${MacPoolName}'. (User: \${UserName})
10702	MAC_POOL_EDIT_SUCCESS	Info	MAC Pool '\${MacPoolName}' (id
10703	MAC_POOL_EDIT_FAILED	Error	Failed to update MAC Pool '\${MacPoolName}' (id
10704	MAC_POOL_REMOVE_SUCCESS	Info	MAC Pool '\${MacPoolName}' (id
10705	MAC_POOL_REMOVE_FAILED	Error	Failed to remove MAC Pool '\${MacPoolName}' (id
10750	CINDER_PROVIDER_ERROR	Error	An error occurred on Cinder provider: '\${CinderException}'
10751	CINDER_DISK_CONNECTION_FAILURE	Error	Failed to retrieve connection information for Cinder Disk '\${DiskAlias}'.
10752	CINDER_DISK_CONNECTION_VOLUME_DRIVER_UNSUPPORTED	Error	Unsupported volume driver for Cinder Disk '\${DiskAlias}'.

Code	Name	Severity	Message
10753	USER_FINISHED_FAILED_REMOVE_CINDER_DISK	Error	Failed to remove disk <code>\${DiskAlias}</code> from storage domain <code>\${StorageDomainName}</code> . The following entity id could not be deleted from the Cinder provider <code>'\${imageId}'</code> . (User: <code>\${UserName}</code>).
10754	USER_ADDED_LIBVIRT_SECRET	Info	Authentication Key <code>\${LibvirtSecretUUID}</code> was added. (User: <code>\${UserName}</code>).
10755	USER_FAILED_TO_ADD_LIBVIRT_SECRET	Error	Failed to add Authentication Key <code>\${LibvirtSecretUUID}</code> . (User: <code>\${UserName}</code>).
10756	USER_UPDATE_LIBVIRT_SECRET	Info	Authentication Key <code>\${LibvirtSecretUUID}</code> was updated. (User: <code>\${UserName}</code>).
10757	USER_FAILED_TO_UPDATE_LIBVIRT_SECRET	Error	Failed to update Authentication Key <code>\${LibvirtSecretUUID}</code> . (User: <code>\${UserName}</code>).
10758	USER_REMOVED_LIBVIRT_SECRET	Info	Authentication Key <code>\${LibvirtSecretUUID}</code> was removed. (User: <code>\${UserName}</code>).
10759	USER_FAILED_TO_REMOVE_LIBVIRT_SECRET	Error	Failed to remove Authentication Key <code>\${LibvirtSecretUUID}</code> . (User: <code>\${UserName}</code>).

Code	Name	Severity	Message
10760	FAILED_TO_REGISTER_LIBVIRT_SECRET	Error	Failed to register Authentication Keys for storage domain <code>{StorageDomainName}</code> on host <code>{VdsName}</code> .
10761	FAILED_TO_UNREGISTER_LIBVIRT_SECRET	Error	Failed to unregister Authentication Keys for storage domain <code>{StorageDomainName}</code> on host <code>{VdsName}</code> .
10762	FAILED_TO_REGISTER_LIBVIRT_SECRET_ON_VDS	Error	Failed to register Authentication Keys on host <code>{VdsName}</code> .
10763	NO_LIBRBD_PACKAGE_AVAILABLE_ON_VDS	Error	Librbd1 package is not available on host <code>{VdsName}</code> , which is mandatory for using Cinder storage domains.
10764	FAILED_TO_FREEZE_VM	Warning	Failed to freeze guest filesystems on VM <code>{VmName}</code> . Note that using the created snapshot might cause data inconsistency.
10765	FAILED_TO_THAW_VM	Warning	Failed to thaw guest filesystems on VM <code>{VmName}</code> . The filesystems might be unresponsive until the VM is restarted.
10766	FREEZE_VM_INITIATED	Normal	Freeze of guest filesystems on VM <code>{VmName}</code> was initiated.
10767	FREEZE_VM_SUCCESS	Normal	Guest filesystems on VM <code>{VmName}</code> have been frozen successfully.

Code	Name	Severity	Message
10768	THAW_VM_SUCCESS	Normal	Guest filesystems on VM {VmName} have been thawed successfully.
10769	USER_FAILED_TO_FREEZE_VM	Warning	Failed to freeze guest filesystems on {VmName} (Host: {VdsName} , User: {UserName}).
10770	USER_FAILED_TO_THAW_VM	Warning	Failed to thaw guest filesystems on {VmName} (Host: {VdsName} , User: {UserName}).
10771	VDS_CANNOT_CONNECT_TO_GLUSTERFS	Error	Host {VdsName} cannot connect to Glusterfs. Verify that glusterfs-cli package is installed on the host.
10780	AFFINITY_RULES_ENFORCEMENT_MANAGER_START	Normal	Affinity Rules Enforcement Manager started.
10781	AFFINITY_RULES_ENFORCEMENT_MANAGER_INTERVAL_REACHED	Normal	
10800	VM_ADD_HOST_DEVICES	Info	Host devices {NamesAdded} were attached to Vm {VmName} by User {UserName} .
10801	VM_REMOVE_HOST_DEVICES	Info	Host devices {NamesRemoved} were detached from Vm {VmName} by User {UserName} .
10802	VDS_BROKER_COMMAND_FAILURE	Error	VDSM {VdsName} command {CommandName} failed: {message}

Code	Name	Severity	Message
10803	IRS_BROKER_COMMAND_FAILURE	Error	VDSM command <code>{CommandName}</code> failed: <code>{message}</code>
10804	VDS_UNKNOWN_HOST	Error	The address of host <code>{VdsName}</code> could not be determined
10810	SYSTEM_CHANGE_STORAGE_POOL_STATUS_UP_REPORTING_HOSTS	Normal	Data Center <code>{StoragePoolName}</code> status was changed to UP as some of its hosts are in status UP.
10811	SYSTEM_CHANGE_STORAGE_POOL_STATUS_NON_RESPONSIVE_NO_REPORTING_HOSTS	Info	Data Center <code>{StoragePoolName}</code> status was changed to Non Responsive as none of its hosts are in status UP.
10812	STORAGE_POOL_LOWER_THAN_ENGINE_HIGHEST_CLUSTER_LEVEL	Info	Data Center <code>{StoragePoolName}</code> compatibility version is <code>{dcVersion}</code> , which is lower than latest engine version <code>{engineVersion}</code> . Please upgrade your Data Center to latest version to successfully finish upgrade of your setup.
10900	HOST_SYNC_ALL_NETWORKS_FAILED	Error	Failed to sync all host <code>{VdsName}</code> networks
10901	HOST_SYNC_ALL_NETWORKS_FINISHED	Info	Managed to sync all host <code>{VdsName}</code> networks.

Code	Name	Severity	Message
10902	PERSIST_HOST_SETUP_NETWORK_ON_HOST	Info	(\${Sequence}/\${Total}): Applying network's changes on host \${VdsName}. (User: \${UserName})
10903	PERSIST_SETUP_NETWORK_ON_HOST_FINISHED	Info	(\${Sequence}/\${Total}): Successfully applied changes on host \${VdsName}. (User: \${UserName})
10904	PERSIST_SETUP_NETWORK_ON_HOST_FAILED	Error	(\${Sequence}/\${Total}): Failed to apply changes on host \${VdsName}. (User: \${UserName})
10905	CLUSTER_SYNC_ALL_NETWORKS_FAILED	Error	Failed to sync all cluster \${ClusterName} networks
10906	CLUSTER_SYNC_ALL_NETWORKS_STARTED	Info	Started sync of all cluster \${ClusterName} networks.
10910	NETWORK_REMOVE_NIC_FILTER_PARAMETER	Info	Network interface filter parameter (id \${VmNicFilterParameterId}) was successfully removed by \${UserName}.
10911	NETWORK_REMOVE_NIC_FILTER_PARAMETER_FAILED	Error	Failed to remove network interface filter parameter ((id \${VmNicFilterParameterId}) by \${UserName}.

Code	Name	Severity	Message
10912	NETWORK_ADD_NIC_FILTER_PARAMETER	Info	Network interface filter parameter <code>{VmNicFilterParameterName}</code> (id <code>{VmNicFilterParameterId}</code>) was successfully added to Interface with id <code>{VmInterfaceId}</code> on VM <code>{VmName}</code> by <code>{UserName}</code> .
10913	NETWORK_ADD_NIC_FILTER_PARAMETER_FAILED	Error	Failed to add network interface filter parameter <code>{VmNicFilterParameterName}</code> (id <code>{VmNicFilterParameterId}</code>) to Interface with id <code>{VmInterfaceId}</code> on VM <code>{VmName}</code> by <code>{UserName}</code> by <code>{UserName}</code> .
10914	NETWORK_UPDATE_NIC_FILTER_PARAMETER	Info	Network interface filter parameter <code>{VmNicFilterParameterName}</code> (id <code>{VmNicFilterParameterId}</code>) on Interface with id <code>{VmInterfaceId}</code> on VM <code>{VmName}</code> was successfully updated by <code>{UserName}</code> .
10915	NETWORK_UPDATE_NIC_FILTER_PARAMETER_FAILED	Error	Failed to update network interface filter parameter <code>{VmNicFilterParameterName}</code> (id <code>{VmNicFilterParameterId}</code>) on Interface with id <code>{VmInterfaceId}</code> on VM <code>{VmName}</code> by <code>{UserName}</code> .

Code	Name	Severity	Message
10916	MAC_ADDRESS_HAD_TO_BE_REALLOCATED	Warning	Some MAC addresses had to be reallocated because they are duplicate.
10917	MAC_ADDRESS_VIOLATES_NO_DUPLICATES_SETTING	Error	Duplicate MAC addresses had to be introduced into mac pool violating no duplicates setting.
10918	MAC_ADDRESS_COULDNT_BE_REALLOCATED	Error	Some MAC addresses had to be reallocated, but operation failed because of insufficient amount of free MACs.
10920	NETWORK_IMPORT_EXTERNAL_NETWORK	Info	Successfully initiated import of external network <code>\${NetworkName}</code> from provider <code>\${ProviderName}</code> .
10921	NETWORK_IMPORT_EXTERNAL_NETWORK_FAILED	Error	Failed to initiate external network <code>\${NetworkName}</code> from provider <code>\${ProviderName}</code> .
10922	NETWORK_IMPORT_EXTERNAL_NETWORK_INTERNAL	Info	
10923	NETWORK_IMPORT_EXTERNAL_NETWORK_INTERNAL_FAILED	Error	

Code	Name	Severity	Message
10924	NETWORK_AUTO_DEFINE_NO_DEFAULT_EXTERNAL_PROVIDER	Warning	Cannot create auto-defined network connected to \${NetworkName}. Cluster \${ClusterName} does not have default external network provider.
11000	USER_ADD_EXTERNAL_JOB	Info	New external Job \${description} was added by user \${UserName}
11001	USER_ADD_EXTERNAL_JOB_FAILED	Error	Failed to add new external Job \${description}
11500	FAULTY_MULTIPATHS_ON_HOST	Warning	Faulty multipath paths on host \${VdsName} on devices: [\${MpathGuids}]
11501	NO_FAULTY_MULTIPATHS_ON_HOST	Normal	No faulty multipath paths on host \${VdsName}
11502	MULTIPATH_DEVICES_WITHOUT_VALID_PATHS_ON_HOST	Warning	Multipath devices without valid paths on host \${VdsName} : [\${MpathGuids}]
12000	MIGRATION_REASON_AFFINITY_ENFORCEMENT	Info	Affinity rules enforcement
12001	MIGRATION_REASON_LOAD_BALANCING	Info	Load balancing
12002	MIGRATION_REASON_HOST_IN_MAINTENANCE	Info	Host preparing for maintenance

Code	Name	Severity	Message
12003	VM_MIGRATION_NOT_ALL_VM_NICS_WERE_PLUGGED_BACK	Error	After migration of \${VmName} , following vm nics failed to be plugged back: \${NamesOfNotRepluggedNics} .
12004	VM_MIGRATION_PLUGGING_VM_NICS_FAILED	Error	After migration of \${VmName} vm nics failed to be plugged back.
12005	CLUSTER_CANNOT_UPDATE_VM_COMPATIBILITY_VERSION	Error	Cannot update compatibility version of Vm/Template: [\${VmName}] , Message: \${Message}
13000	DEPRECATED_API	Warning	Client from address "\${ClientAddress}" is using version \${ApiVersion} of the API, which has been \
13001	DEPRECATED_IPTABLES_FIREWALL	Warning	Cluster \${ClusterName} uses IPTables firewall, which has been deprecated in \

APPENDIX C. TIMEZONES

The API maps Windows Standard Format timezone names to tz database format when specifying a timezone for a virtual machine or VM template. This means the API only accepts certain tz database codes, which the following table lists:

Table C.1. Accepted tz database codes

tz database Format	Windows Standard Format
Africa/Cairo	Egypt Standard Time
Africa/Casablanca	Morocco Standard Time
Africa/Johannesburg	South Africa Standard Time
Africa/Lagos	W. Central Africa Standard Time
Africa/Nairobi	E. Africa Standard Time
Africa/Reykjavik	Greenwich Standard Time
Africa/Windhoek	Namibia Standard Time
America/Anchorage	Alaskan Standard Time
America/Bogota	SA Pacific Standard Time
America/Buenos_Aires	Argentina Standard Time
America/Caracas	Venezuela Standard Time
America/Chicago	Central Standard Time
America/Chihuahua	Mexico Standard Time
America/Chihuahua	Mountain Standard Time
America/Denver	Mountain Standard Time
America/Godthab	Greenland Standard Time
America/Guatemala	Central America Standard Time
America/Halifax	Atlantic Standard Time
America/La_Paz	SA Western Standard Time

tz database Format	Windows Standard Format
America/Los_Angeles	Pacific Standard Time
America/Manaus	Central Brazilian Standard Time
America/Mexico_City	Central Standard Time
America/Mexico_City	Mexico Standard Time
America/Montevideo	Montevideo Standard Time
America/New_York	Eastern Standard Time
America/Phoenix	US Mountain Standard Time
America/Regina	Canada Central Standard Time
America/Santiago	Pacific SA Standard Time
America/Sao_Paulo	E. South America Standard Time
America/St_Johns	Newfoundland Standard Time
America/Tijuana	Pacific Standard Time
Asia/Amman	Jordan Standard Time
Asia/Baghdad	Arabic Standard Time
Asia/Baku	Azerbaijan Standard Time
Asia/Bangkok	SE Asia Standard Time
Asia/Beirut	Middle East Standard Time
Asia/Calcutta	India Standard Time
Asia/Colombo	Sri Lanka Standard Time
Asia/Dhaka	Central Asia Standard Time
Asia/Dubai	Arabian Standard Time
Asia/Irkutsk	North Asia East Standard Time

tz database Format	Windows Standard Format
Asia/Jerusalem	Israel Standard Time
Asia/Kabul	Afghanistan Standard Time
Asia/Karachi	Pakistan Standard Time
Asia/Katmandu	Nepal Standard Time
Asia/Krasnoyarsk	North Asia Standard Time
Asia/Novosibirsk	N. Central Asia Standard Time
Asia/Rangoon	Myanmar Standard Time
Asia/Riyadh	Arab Standard Time
Asia/Seoul	Korea Standard Time
Asia/Shanghai	China Standard Time
Asia/Singapore	Singapore Standard Time
Asia/Taipei	Taipei Standard Time
Asia/Tashkent	West Asia Standard Time
Asia/Tehran	Iran Standard Time
Asia/Tokyo	Tokyo Standard Time
Asia/Vladivostok	Vladivostok Standard Time
Asia/Yakutsk	Yakutsk Standard Time
Asia/Yekaterinburg	Ekaterinburg Standard Time
Asia/Yerevan	Armenian Standard Time
Asia/Yerevan	Caucasus Standard Time
Atlantic/Azores	Azores Standard Time
Atlantic/Cape_Verde	Cape Verde Standard Time
Atlantic/South_Georgia	Mid-Atlantic Standard Time

tz database Format	Windows Standard Format
Australia/Adelaide	Cen. Australia Standard Time
Australia/Brisbane	E. Australia Standard Time
Australia/Darwin	AUS Central Standard Time
Australia/Hobart	Tasmania Standard Time
Australia/Perth	W. Australia Standard Time
Australia/Sydney	AUS Eastern Standard Time
Etc/GMT-3	Georgian Standard Time
Etc/GMT+12	Dateline Standard Time
Etc/GMT+3	SA Eastern Standard Time
Etc/GMT+5	US Eastern Standard Time
Europe/Berlin	W. Europe Standard Time
Europe/Budapest	Central Europe Standard Time
Europe/Istanbul	GTB Standard Time
Europe/Kiev	FLE Standard Time
Europe/London	GMT Standard Time
Europe/Minsk	E. Europe Standard Time
Europe/Moscow	Russian Standard Time
Europe/Paris	Romance Standard Time
Europe/Warsaw	Central European Standard Time
Indian/Mauritius	Mauritius Standard Time
Pacific/Apia	Samoa Standard Time
Pacific/Auckland	New Zealand Standard Time
Pacific/Fiji	Fiji Standard Time

tz database Format	Windows Standard Format
Pacific/Guadalcanal	Central Pacific Standard Time
Pacific/Honolulu	Hawaiian Standard Time
Pacific/Port_Moresby	West Pacific Standard Time
Pacific/Tongatapu	Tonga Standard Time

APPENDIX D. LEGAL NOTICE

Copyright © 2022 Red Hat, Inc.

Licensed under the ([Creative Commons Attribution–ShareAlike 4.0 International License](#)). Derived from documentation for the ([oVirt Project](#)). If you distribute this document or an adaptation of it, you must provide the URL for the original version.

Modified versions must remove all Red Hat trademarks.

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.