IBM VM Recovery Manager HA for Power Systems

Version 1.5

Deployment Guide

Note

Before using this information and the product it supports, read the information in “Notices” on page 95.

This is the latest edition for IBM® VM Recovery Manager HA Version 1.5 for Power Systems until otherwise indicated in a newer edition.

Contents

About this document

The VM Recovery Manager HA solution is a set of software components that together provide a high availability mechanism for virtual machines running on POWER7® processor-based server, or later. This document describes various components, subsystems, and tasks that are associated with the VM Recovery Manager HA solution.

This information provides system administrators with complete information about the following sections:

- Concepts that are used in the VM Recovery Manager HA solution.
- Planning the VM Recovery Manager HA implementation in your production environment and the minimum software requirements.
- Installing the VM Recovery Manager HA filesets.
- Configuring your environment to use the VM Recovery Manager HA solution.
- Troubleshooting any issues associated with the VM Recovery Manager HA solution.
- Using the VM Recovery Manager HA commands.

Highlighting

The following highlighting conventions are used in this document:

Bold	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Bold highlighting also identifies graphical objects, such as buttons, labels, and icons that you select.
Italics	Identifies parameters for actual names or values that you supply.
`Monospace`	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or text that you must type.

Case-sensitivity in VM Recovery Manager HA

The command name and some flags in the VM Recovery Manager HA solution are case-sensitive, which means that it distinguishes between uppercase and lowercase letters. To avoid causing undesirable actions to be performed, always ensure that you use the correct case.

Correct example

```
ksysmgr ADD SITE site_name sitetype=active
```

Incorrect example

```
KSYSMGR ADD Site Site_name SiteType=Active
```

ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

Frequently asked questions

If you have questions about the VM Recovery Manager HA solution, review the following list of answers to some frequently asked questions.

What is VM Recovery Manager HA?

The VM Recovery Manager HA solution implements high availability of virtual machines based on the VM restart technology. Read more about the solution in the "VM Recovery Manager HA overview" on page 5 topic.

What is KSYS?

The KSYS (also known as the controlling system) is a controlling software for the HA operation. The KSYS software is installed in an AIX® Version 7.2.2, or later logical partition. The KSYS LPAR controls the entire cloud environment for the VM Recovery Manager HA solution. If the KSYS subsystem detects failures of hosts, VMs, or critical applications, the KSYS subsystem restarts the virtual machines in another host.

What is a host group?

Hosts are grouped in the VM Recovery Manager HA configuration settings to act as a backup for each other. When failures are detected in any host, VMs in the failed host are relocated and restarted on other healthy hosts within the group. This group of hosts is called a host group. For more information, see "Host group requirements" on page 11.

Which POWER® servers can be included for high availability by using the VM Recovery Manager HA solution?

POWER7+™ processor-based server, or later. For more information, see "Firmware requirements" on page 11.

Which operation systems are supported for virtual machines in hosts?

- AIX Version 6.1, or later
- PowerLinux (Red Hat, SUSE, and Ubuntu Linux® distributions)
- IBM i Version 7.2, or later

For more information, see "Software requirements" on page 10.

What is the minimum version of HMC and VIOS that must be used for this solution?

See "Firmware requirements" on page 11 and "HMC requirements" on page 12.

What are the CPU and memory requirements for the VM Recovery Manager HA solution?

You must have 30 MB of disk space in the `/opt` directory and 200 MB of disk space in the `/var` directory to install the KSYS filesets. For typical cloud environments with approximately 100 LPARs, the KSYS LPAR must have at least one core of CPU and 8 GB of memory. You must also ensure that VIOS contains enough resources for HA management. For typical environments, VIOS resources must be increased by at least one core CPU and at least 2 GB memory beyond the sizing based on the workload requirements. The VM Recovery Manager HA solution also requires two free disks of at least 10 GB that are connected and available across all the Virtual I/O Servers that are included for HA management in the host group. For more information, see "Installation and configuration requirements" on page 11.

How do I install VM Recovery Manager HA solution?

The VM Recovery Manager HA solution consists of few filesets that can be installed by using the AIX installation format on the KSYS LPAR. You can install the GUI server on a separate AIX LPAR or on the KSYS LPAR. Optionally, you can install the VM agents in the guest LPARs for VM and application HA management. For more information, see "Installing the KSYS software" on page 20.

How can I run the operations for configuration and HA management?

You can use both the command-line interface and the GUI to run the HA operations. You can use the **`ksysmgr`** command for all possible operations in the KSYS subsystem. To install the GUI, see "Installing GUI server filesets" on page 21.

Does the GUI depend on open source components?

Yes. The VM Recovery Manager HA solution provides scripts to download the required images and to install them easily. For more information about the GUI installation, see "Installing GUI server filesets" on page 21.

How do I configure my environment by using the VM Recovery Manager HA solution?

After installing the VM Recovery Manager HA software on the AIX LPAR (KSYS LPAR), KSYS can be initialized and managed by using the VM Recovery Manager HA GUI or the **ksysmgr** command. Use the GUI or the CLI to initialize the KSYS LPAR and then create the host groups and register the HMC and hosts. The KSYS subsystem discovers all the resources by working with the HMC and the VIOS. After the discovery operation is complete, you can run the verification operation to ensure that entire environment is free of errors and is ready to move. For configuration steps, see "Setting up the KSYS subsystem" on page 27.

What are the setup requirements for the host group?

A host group consists of a group of hosts that can protect each other. You must be able to relocate any VM on any of the hosts within the host group. Therefore, network and storage requirements of the VM must be met on each of the hosts. In addition, Live Partition Mobility (LPM) validation of a VM must be successful for each of the host in the host group. For more information, see "Host group requirements" on page 11.

If a host fails, will VM Recovery Manager HA automatically restart the virtual machines?

Yes. By default, the KSYS subsystem automatically restarts the virtual machines on another hosts within the host group and notifies you about the relocation operation. You can customize this setting and set the advisory mode so that the KSYS subsystem only notifies you about the various failures. For more information, see "Restart policy" on page 33.

Will KSYS alert me only for host failures or for VM failures and application failures as well?

By default, the KSYS subsystem provides high-availability feature at host level. You can optionally install the VM agent filesets in the guest virtual machines for high-availability feature at VM and application level. For more information, see "Installing VM agents" on page 22 and "Setting up the VM agent" on page 38.

How do I manage applications for high-availability?

You must install the VM agent in the guest VMs. Currently, the VM agent is supported for AIX and Linux operating systems. After you install the VM agent, you can initialize the VM agent and register the application HA management through a single command that is available in the VM agent component. For more information, see "ksysvmmgr command" on page 65.

Can I register my own custom monitoring methods with VM Recovery Manager HA?

Yes. Use the **ksysvmmgr** command in the VM to register application start, stop, and monitoring methods.

Can I control the start sequences of applications within a VM?

Yes. Use the **ksysvmmgr** command to define the start sequences of applications within the VM. For example, you might want to start the database application first and then start the other applications that use the database application.

Can I start the application after the VM has been relocated automatically?

Yes. You can use the **ksysvmmgr** command to start the application automatically after the VM is restarted. You can also change this setting to start the applications manually.

After the initial configuration, can I modify the environment?

Yes. You can add or delete VMs or hosts and such modifications are discovered automatically by the KSYS subsystem during its scheduled periodic scans. You can also run the discovery operation manually to register the changes in the environment immediately. For more information, see "Modifying the KSYS configuration" on page 37.

Can VM Recovery Manager HA work with LPM?

Yes. You can perform the LPM operation on any of the VMs within a host group that is managed by the VM Recovery Manager HA solution. You can also use the **ksysmgr** command or the GUI to perform the LPM operations. For more information, see "Planned migration of virtual machines" on page 41.

Do I need to perform LPM only through VM Recovery Manager HA or can I do it through the HMC, PowerVC, or LPM tool?

You can use the VM Recovery Manager HA solution to perform planned HA activities by using the LPM operation. However, you can use any other tool to do LPM that works best for you. If you are using another tool, ensure that the VM moves to another host that is within the host group.

Can VM Recovery Manager HA co-exist with other solutions such as PowerVC and PowerHA® SystemMirror?

Yes. For more information, see "Coexistence with other products" on page 13.

Does VM Recovery Manager HA support NovaLink?

Not directly. However, in PowerVM NovaLink-based environments that also have HMCs, you can register scripts that can be plugged into VM Recovery Manager HA. The sample scripts can be customized based on your environment. These scripts change the HMC settings to be in master mode for brief periods so that the KSYS subsystem can work with the HMC to monitor the environment for high availability.

Which storage systems does VM Recovery Manager HA support?

The VM Recovery Manager HA solution can support any storage systems that are certified with the VIOS, except internet Small Computer Systems Interface (iSCSI) storage devices. Storage disks that are related to VMs must be accessible across all the hosts within the host group so that VM can move from a host to any other host within the host group.

Can I use VM Recovery Manager HA with SVC HyperSwap®?

SAN Volume Controller (SVC) and Storwize® HyperSwap technology perform hidden synchronous mirroring across short distances. The VM Recovery Manager HA solution can be used across that distance if the storage system fulfills all the requirements of shared storage and is certified with the VIOS.

Which types of network do VM Recovery Manager HA support?

The VM Recovery Manager HA solution can support any network that is certified with the VIOS that supports the Live Partition Mobility operation.

Which types of VIOS storage configuration do VM Recovery Manager HA support?

The VM Recovery Manager HA solution supports virtual SCSI (vSCSI), N_Port ID virtualization (NPIV), Shared Storage Pool (SSP), and any storage configuration that is supported for the LPM operation.

What all can I manage from the VM Recovery Manager HA GUI?

The VM Recovery Manager HA GUI offers deployment, health monitoring, and administrative experiences.

Does VM Recovery Manager HA relocate a VM if I shut down or stop the VM manually?

No. The VM Recovery Manager HA solution checks the HMC and firmware resources to verify whether the lack of heartbeats from a VM is caused because of an administrator-initiated operation. In those cases, VM will not be relocated.

How is fencing performed during failures?

The VM Recovery Manager HA solution stops the failed VM through the HMC before starting the VM on another host. If the KSYS subsystem cannot connect to the HMC or if the HMC cannot stop the VM successfully, the KSYS subsystem cancels the automated restart operation and instead sends a critical event to the administrator about the VM failure.

Can I configure the alerts in KSYS to be sent as text messages instead of emails?

Yes. Contact your phone company to obtain the phone number that can be represented as an email address, and register the updated email address with the KSYS subsystem. For more information, see Setting contacts for event notification.

Can I register our own method or script to be invoked when specific events occur?

Yes. Use the **ksysmgr** command to register scripts to be called by the KSYS subsystem when specific events occur. For more information, see "Event notification script management" on page 61.

How do I save the KSYS configuration?

Use the **ksysmgr** command to save the KSYS configuration as a snapshot that can be restored later. For more information, see Backing up the configuration data.

How do I update the KSYS LPAR or the KSYS software?

You can update the KSYS software with newer releases of VM Recovery Manager HA and update the KSYS LPAR with newer releases of AIX. During the update process, the production workloads in the VMs are not disrupted. During the KSYS update process, the HA monitoring will be suspended and will resume after the update operation is complete.

Can we predict the impact on the environment because of a VM or host failure?

Use the **ksysmgr report system relocation_plan [vm=vm_list|host=host_list]** command to review the relocation plan for the VMs. When you run this command, the KSYS subsystem reviews the available spare resources and policies to predict where the VMs will be restarted.

How do I know whether the environment is protected from an HA perspective?

You can run the **ksysmgr verify** command to validate the environment. The KSYS subsystem communicates with the HMC and the VIOS to check and report any issues in the environment. Additionally, the KSYS subsystem checks the CPU and memory resources to predict the protection capacity for the HA environment. *Protection capacity* is an attribute of the host group in units of number of hosts. The protection capacity indicates how many host failures can be tolerated by the host group. The KSYS subsystem picks up the largest host, which is hosting the LPARs that sum up to largest amount of CPU resources, and checks whether the LPARs from this host can be recovered by using the spare capacity in the remaining hosts. If that validation passes, then the protection capacity is set as 1. The KSYS subsystem then picks up the next largest host and checks whether the failure of two largest hosts in the host group can be tolerated. If yes, the protection capacity is set as 2. You must deploy the environment such that the environment has enough spare CPU and memory capacity to achieve a protection capacity value of at least 1.

While setting up HG from GUI, can a user access HMC deleted from backend using CLI?

No, user cannot access the HMC deleted from backend using CLI. This is not preferred because unavailability of HMC leads to failure of Host Group deployment.

VM Recovery Manager HA overview

High availability (HA) management is a critical feature of business continuity plans. Any downtime to the software stack can result in loss of revenues and disruption of services. IBM VM Recovery Manager HA for Power Systems is a high availability solution that is easy to deploy and provides an automated solution to recover the virtual machines (VMs), also known as logical partitions (LPARs).

The VM Recovery Manager HA solution implements recovery of the virtual machines based on the VM restart technology. The VM restart technology relies on an out-of-band monitoring and management component that restarts the VMs on another server when the host infrastructure fails. The VM restart technology is different from the conventional cluster-based technology that deploys redundant hardware and software components for a near real-time failover operation when a component fails.

The VM Recovery Manager HA solution is ideal to ensure high availability for many VMs. Additionally, the VM Recovery Manager HA solution is easier to manage because it does not have clustering complexities.

The following figure shows the architecture of the VM Recovery Manager HA solution. A set of hosts is grouped to be backup for each other. When failures are detected, VMs are relocated and restarted on other healthy hosts within the group.

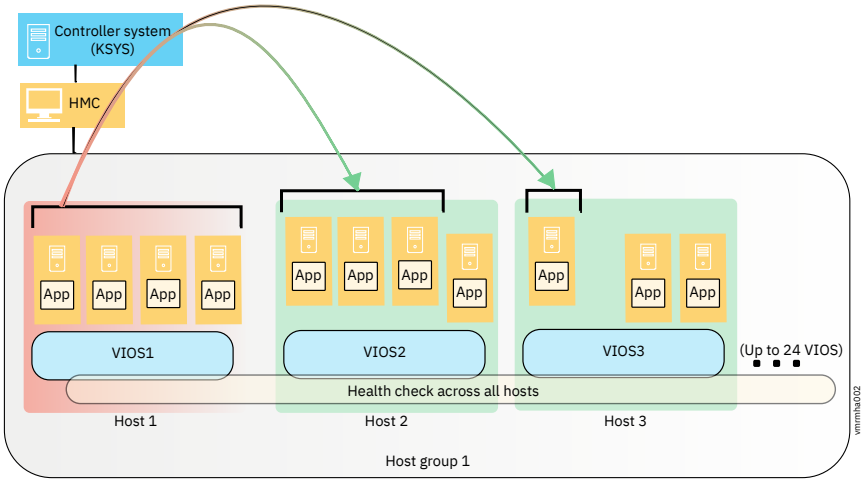

Figure 1. VM Recovery Manager HA solution architecture

The VM Recovery Manager HA solution provides the following capabilities:

Host health monitoring
The VM Recovery Manager HA solution monitors hosts for any failures. If a host fails, the virtual machines in the failed host are automatically restarted on other hosts. The VM Recovery Manager HA solution uses the host monitor module of the VIOS partition in a host to monitor the health of hosts.

VM and application health monitoring
The VM Recovery Manager HA solution monitors the virtual machines, its registered applications, and its hosts, for any failures. If a virtual machine or a critical application fails, the corresponding virtual machines are started automatically on other hosts. The VM Recovery Manager HA solution uses the VM monitor agent that must be installed in each virtual machine to monitor the health of virtual machines and registered applications.

Unplanned HA management
During an unplanned outage, when the VM Recovery Manager HA solution detects a failure in the environment, the virtual machines are restarted automatically on other hosts. You can also change the auto-restart policy to advisory mode. In advisory mode, failed VMs are not relocated automatically, instead email or text messages are sent to the administrator. Administrator can use the interfaces to manually restart the VMs.

Planned HA management

During a planned outage, when you plan to update firmware for a host, you can use the Live Partition Mobility operation of the VM Recovery Manager HA solution to vacate a host by moving all the VMs in the host to the remaining hosts in the group. After the upgrade operation is complete, you can use the VM Recovery Manager HA solution to restore the VM to its original host in a single operation.

Advanced HA policies

The VM Recovery Manager HA solution provides advanced policies to define relationships between VMs such as collocation and anti-collocation of VMs, priority in which the VMs will be restarted, capacity of VMs during failover operations.

GUI and command-line based management

You can use GUI or command-line interface to manage the resources in the VM Recovery Manager HA solution. For GUI, you can install the UI server and then use the web browser to manage the resources. Alternatively, the **ksysmgr** command and the **ksysvmmgr** command on KSYS LPAR provide end-to-end HA management for all resources.

VM Recovery Manager HA concepts

The VM Recovery Manager HA solution provides a highly available environment by identifying a set of resources that are required for processing virtual machines in a server.

The VM Recovery Manager HA solution uses the following concepts:

Controller system (KSYS)
The controlling system, also called KSYS, is a fundamental component that monitors the production environment for any unplanned outage. If an unplanned outage occurs, the KSYS analyzes the situation, notifies the administrator about the failure, and can automatically move the failed virtual machines to another host in the host group. The KSYS interacts with the Hardware Management Console (HMC) to collect configuration information of managed systems. The KSYS subsystem also collects VIOS health information through the HMC.

The KSYS subsystem runs in an AIX logical partition (LPAR). You can customize the security level for the KSYS LPAR according to the AIX security requirements of your organization. In addition, the KSYS LPAR can be protected for failure by using other products such as PowerHA SystemMirror® for AIX. The KSYS subsystem must remain operational even if the site fails. Ensure that you periodically receive KSYS health reports. You can also check the KSYS subsystem health in the VM Recovery Manager HA GUI dashboard.

Host group
Hosts are grouped together to be backup for each other. When failures in any of the hosts are detected, VMs in the failed host are relocated and restarted on other healthy hosts within the group of hosts. This group of hosts is called a host group.

Host monitor
The host monitor daemon is shipped with the Virtual I/O Server (VIOS) and is deployed during the VIOS installation. When you initialize the KSYS subsystem for high-availability feature, the host monitor module becomes active. The KSYS subsystem communicates with the host monitor daemon through the HMC to monitor the hosts for high availability. For information about the VIOS version that contains the host monitor daemon, see the Requirements section.

VM agent
You can optionally install the VM agent filesets, which are shipped along with the KSYS filesets, in the guest virtual machines. The VM agent subsystem provides high-availability feature at the VM and application level. The VM agent monitors the following issues in the production environment:

- **VM failures**: If the operating system of a VM is not working correctly, or if the VM has stopped working because of an error, the VM is restarted on another host within the host group. The KSYS subsystem uses the VM monitor module to monitor the heartbeat from the VM to the host monitor subsystem in a VIOS.
- **Application failures**: Optionally, you can register the applications in the VM agent to enable application monitoring. The VM agent uses the Application HA monitoring framework to monitor the health of the application periodically by running the application-specific monitor scripts, by identifying whether the application has failed, and by identifying whether the VM must be restarted in the same host or another host. This framework can also manage the sequence in which applications are started and stopped within a VM.

Note: The VM agent is supported on AIX and Linux (RHEL and SLES) guest VMs only. Currently, the VM agent subsystem is not supported for the IBM i and Ubuntu VMs. Therefore, IBM i and Ubuntu VMs are relocated from one host to another host within the host group only after a host failure.

The following figure shows the detailed architecture of the VM Recovery Manager HA solution:

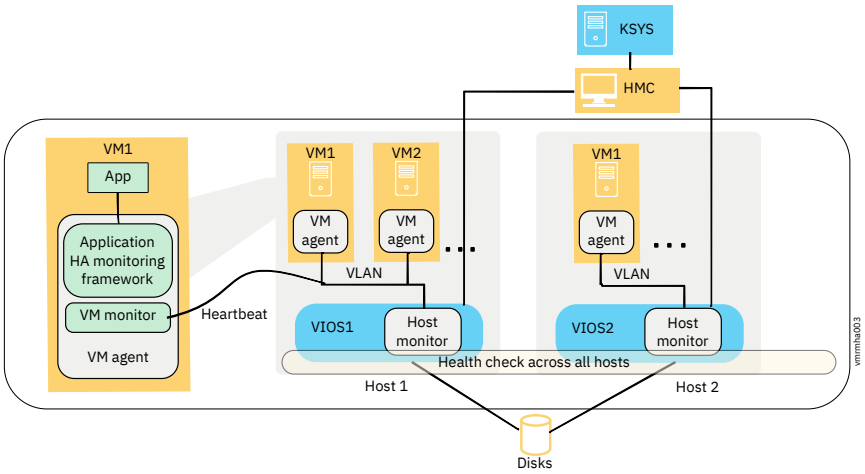

Figure 2. VM Recovery Manager HA solution architecture

Planning VM Recovery Manager HA

To implement the VM Recovery Manager HA solution, you must review your current high availability (HA) recovery plan and consider how the VM Recovery Manager HA solution can be integrated into your current environment.

The VM Recovery Manager HA package consists of filesets for the installation of KSYS, GUI, and VM agent. The following table describes the key components of the VM Recovery Manager HA solution:

Table 1. VM Recovery Manager HA filesets

Subsystem name	Description	Filesets
KSYS	Base product software that must be installed on an AIX logical partition. It provides the technical foundation of VM Recovery Manager HA and command line-based administration by using the **ksysmgr** command.	• `ksys.hautils.rte` • `ksys.ha.license` • `ksys.main.cmds` • `ksys.main.msg.en_US.cmds` • `ksys.main.rte`
GUI	Optional fileset that can be installed on an AIX logical partition for accessing the VM Recovery Manager HA solution by using GUI. You can install the server on the KSYS logical partition also.	• `ksys.ui.agent`: GUI agent fileset that must be installed on the KSYS nodes. • `ksys.ui.server`: GUI server fileset that must be installed on the system that manages the KSYS nodes. This fileset can be installed on one of the KSYS nodes. • `ksys.ui.common`: GUI common fileset that must be installed along with both the `ksys.ui.server` (GUI server) fileset and the `ksys.ui.agent` (GUI agent) fileset.
VM agent	Fileset that can be installed on the guest virtual machines that are running either AIX or Linux operating systems. If you install the VM agent and enable HA monitoring feature at the VM level , you can monitor the individual virtual machines and the applications that are running in the virtual machine. Otherwise, only host level monitoring is supported. VM agents are not supported for Ubuntu Linux distribution and IBM® i operating system. Therefore, the virtual machines that run Ubuntu Linux distribution or IBM® i operating system are monitored only for host failures.	• AIX: `ksys.vmmon.rte` • Red Hat Enterprise Linux (RHEL): `vmagent-1.5.0-1.0.el7.ppc64le` • SUSE Linux Enterprise Server (SLES): `vmagent-1.5.0-1.0.suse123.ppc64le`

VM Recovery Manager HA requirements

Before you plan the implementation of the VM Recovery Manager HA solution, you must understand the other entities and resources that the VM Recovery Manager HA solution requires for disaster recovery operations.

The following requirements must be met before you can install the VM Recovery Manager HA solution:

Software requirements

- The KSYS logical partition must be running IBM AIX 7.2 with Technology Level 2
- You must install the OpenSSL software version 1.0.2.800, or later for the AIX operating system. The latest version of the OpenSSL software is also included on the AIX base media.
- Each LPAR in the host must have one of the following operating systems:
  - AIX Version 6.1, or later
  - PowerLinux
    - Red Hat Enterprise Linux (little endian) Version 7.4, or later (kernel version: 3.10.0-693)
    - SUSE Linux Enterprise Server (little endian) Version 12.3, or later (kernel version - 4.4.126-94.22)
    - Ubuntu Linux distribution Version 16.04
  - IBM i Version 7.1, or later
- You can install the VM agent to monitor the virtual machine and applications on the LPARs that run only the following operating systems:
  - AIX Version 6.1, or later
  - PowerLinux
    - Red Hat Enterprise Linux (little endian) Version 7.4, or later (kernel version: 3.10.0-693)
    - SUSE Linux Enterprise Server Version 12.3, or later (kernel version - 4.4.126-94.22)
- This release requires the `IJ29125m0a.201110.epkg.Z` KSYS efix. You can download and install the efix from the following location:
  https://aix.software.ibm.com/aix/efixes/IJ29125m/IJ29125m0a.201110.epkg.Z
- >|For VIOS version 3.1.1.0, the following VIOS efix is required. You must download the efix and install it before installing the VM Recovery Manager HA for Power Systems Version 1.5:
  https://aix.software.ibm.com/aix/ifixes/IJ21043/IJ21043m1b.200218.epkg.Z |<
- >|For VIOS versions 3.1.1.21 and 3.1.1.25, the following VIOS efix is required. You must download the efix and install it before installing the VM Recovery Manager HA for Power Systems Version 1.5:
  https://aix.software.ibm.com/aix/efixes/ij25165/IJ25165m2c.200727.epkg.Z |<
- For VIOS version 3.1.2.10, the following VIOS efix is required. You must download the efix and install it before installing the VM Recovery Manager HA for Power Systems Version 1.5:
  https://aix.software.ibm.com/aix/efixes/IJ28933/IJ28933m1a.201106.epkg.Z

- PowerVM® Enterprise Edition must be deployed on all hosts to use the high-availability feature.
- VIOS Version 3.1.1.10, or later, with all the subsequent patches must be installed in VIOS partitions. Also, your production environment must have two Virtual I/O Servers per host. You can have a maximum of 24 Virtual I/O Servers in a single host group. If more than two Virtual I/O Servers exist in a host, you can exclude it from the KSYS configuration settings. For more information about setting up dual VIOS in your environment, see Setting up a dual VIOS by using the HMC.
- As a best practice, you must deploy AIX rules in the VIOS. The VIOS must have enough free space in the /(root), `/var`, and `/usr` file system. Additional CPU and memory resources are needed in each VIOS for VM Recovery Manager HA management. You must add at least 0.5 core CPU and 2 GB memory apart from the VIOS sizing that you are planning to deploy for your production environment.

Firmware requirements

- Minimum required levels of IBM Power Systems servers follow:
  - POWER7+™ Systems that have one of the following firmware levels:
    - FW770.90, or later
    - FW780.70, or later except MMB systems (9117-MMB models)
    - FW783.50, or later
  - POWER8® Systems that have one of the following firmware levels:
    - FW840.60, or later
    - FW860.30, or later
  - POWER9™ Systems that have the following firmware levels:
    - FW910, or later

Installation and configuration requirements

- You must have root authority to perform any installation tasks.
- The KSYS LPAR must have at least 1 core CPU and 8 GB memory. These requirements can be higher if you have a large environment of more than 100 LPARs in the data center.
- Ensure that you have enough space in the LPAR so that KSYS filesets can be installed successfully. You must have 30 MB of disk space in the `/opt` directory and 200 MB of disk space in the `/var` directory.
- You must check whether a KSYS installation is already in progress by using the **`ksysmgr q cl`** command. If the KSYS software is installed previously, you must uninstall the KSYS software.
- For the installation of VM agent, ensure each virtual machine meets the following disk space requirements:
  - At least 100 MB disk space in the `/usr` directory to install the VM agent filesets.
  - At least 1 GB disk space in the `/var` directory for log files.

Host group requirements

- Host group can have a logical name with a maximum of 64 characters.
- A single KSYS LPAR can manage up to 4 host groups. A host group can consist of maximum 12 hosts.
- Network and storage must be configured on all hosts in the host group such that any VM from any host can be migrated to any other host within the host group.
- For each host group, the KSYS subsystem requires two disks for health cluster management. A disk of at least 10 GB is required for health monitoring of the hosts, called as repository disk, and another disk of at least 10 GB is required for health data tracking, called as HA disk, for each host group. All these disks must be accessible to all the Virtual I/O Servers on each of the hosts in the host group. >|The hostname of the Virtual I/O Servers must be in the Fully Qualified Domain Name (FQDN) format.|<

HMC requirements

- The VM Recovery Manager HA solution requires HMC Version 9 Release 9.1.0, or later. It is recommended to have each host managed by two HMCs. The HMC must have enough free space.
- Ensure that you can perform Live Partition Mobility (LPM) operation on the VMs among the hosts that you want to be a part of VM Recovery Manager HA management. You can use HMC-based LPM validation to ensure that the virtual machines can move from a host to any other host in the host group.
  - To display the ongoing logical partition movement progress (LPM), HMC Version 9 Release 9.3.0, or later is required.
- For POWER7+™ Systems and POWER8 Systems in which the **Simplified Remote Restart** attribute is not set for a VM in the HMC, when the VM is relocated from a source host to a destination host, the time-of-day change might not be updated correctly. To retain the correct time-of-day clock for a VM across the hosts, you must set the VIOS partitions profile of both the source host and the destination host as the time reference partition by enabling the **Time Reference Partition** option in the HMC. For more information about how to set the Time Reference Partition, see Synchronizing the hypervisor and Service Processor time-of-day clocks to Time Reference Partition.
- To migrate an IBM i virtual machine from the source host to the destination host, verify that the **Restricted I/O Partition** check box for the IBM i logical partition is selected in the HMC. For more information about how to verify the restricted I/O mode, see Verifying that the IBM i mobile partition is in the restricted I/O mode.
- Ensure that the **automatic reboot** attribute is not set for any VM in the HMC. The KSYS subsystem validates this attribute and notifies you to disable this attribute. If you set this attribute, it can lead to unpredictable results such as the VM can restart on two hosts simultaneously.
- When you add a host or manage a new VM that is co-managed by the HMC and PowerVM NovaLink, set the HMC to be in the master mode. Otherwise, the discovery operation fails and the virtual machines in the host are not monitored for the high-availability function.

Network requirements

- All virtual machines (VMs) that are managed by the VM Recovery Manager HA solution must use virtual I/O resources through VIOS. The VMs must not be connected to a physical network adapter or any dedicated devices.
- Storage area network (SAN) connectivity and zoning must be configured so that VIOS can access the disks that are relevant to the hosts.
- Ensure independent redundant SAN and network connections are established across the VIOS in each hosts in the host group.
- Ensure that the KSYS LPAR has HTTP Secure (HTTPS) connectivity to all the HMCs that can manage the hosts in the host group.
- The same virtual LAN (VLAN) must be configured across the site.
- Ensure redundant connections are established from the KSYS LPAR to HMC and from HMC to VIOS logical partitions. Any connectivity issues between KSYS, HMC, and VIOS logical partitions can lead to disruption in the regular data collection activity and disaster recovery operations.
- Ensure proper RMC connection between the VMs and HMC. If RMC connection between VM and HMC has issues, the Partition Load Manager (PLM) cannot work and hence the VMs cannot be recovered.

GUI requirements

- The logical partition (LPAR), in which you want to install the GUI filesets, must be running IBM AIX 7.2 with Technology Level 2 Service Pack 1 (7200-02-01), or later. You can choose to install the GUI server fileset on one of the KSYS nodes.
- The LPAR in which you are installing the GUI server must run in Enhanced Korn shell that uses the `/usr/bin/ksh93` shell script.

- The LPAR in which you are installing the GUI server fileset must have at least 1 core CPU and 8 GB memory.
- Google Chrome and Mozilla Firefox web browsers are supported to access the GUI for the VM Recovery Manager HA solution.

Coexistence with other products

The VM Recovery Manager HA solution and other products can co-exist in the same environment with the following considerations:

IBM Power® Virtualization Center (PowerVC)

- If you use the PowerVC solution to perform the live partition mobility (LPM) operation of virtual machines, ensure that VM remains within the host group that is defined in the KSYS configuration settings. Otherwise, the virtual machine might be disconnected from the VM Recovery Manager HA configuration and cannot be monitored for failures.
- Do not use the HA management function in the PowerVC solution. If both the PowerVC and VM Recovery Manager HA solutions are used to handle unplanned events, the recovery operations might fail.

IBM PowerHA SystemMirror software

Even if you are using the PowerHA SystemMirror solution for application high-availability, you can use the VM Recovery Manager HA solution for hardware failures and automated VM restart operations. This type of deployment can bring the entire cluster back online by restarting the failed PowerHA node or VMs on other hosts.

If you have deployed both the PowerHA SystemMirror and VM Recovery Manager HA solutions, consider the following guidelines:

- If the primary node of the PowerHA configuration fails, the PowerHA solution starts the workload on secondary node of the PowerHA configuration. The VM Recovery Manager HA solution restarts the failed VM, which is the old primary LPAR of PowerHA, on another host. The restarted VM can rejoin the cluster and continue to provide higher level of high-availability.
- If you deploy both the PowerHA solution and the VM Recovery Manager HA solution, you must not configure the VM agent in the PowerHA VMs because PowerHA application management is sufficient for application high-availability.
- Ensure that you define the anti-collocation policies for nodes or LPARs of the PowerHA cluster.

The VM Recovery Manager HA solution can co-exist with other cluster technologies also, such as Oracle RAC and Veritas Cluster Server.

PowerVM® NovaLink

Currently, the VM Recovery Manager HA solution works with the PowerVM environment that is configured with the HMC. When NovaLink and HMCs are deployed together, the VM Recovery Manager HA solution can work only if HMCs are set to be in the master mode. The VM Recovery Manager HA solution communicates with the HMC for continuous monitoring and relocation operations.

Licensing considerations

- With the VM Recovery Manager HA solution, the virtual machines (VMs) that must be replicated or managed by the solution are hosted by processor cores. These managed VMs do not determine the licensing count, but the count of processor cores that are hosting the managed VMs determines the licensing cost. This means that the count of whole number of processor cores that are hosting the VMs that are being replicated or managed by the VM Recovery Manager HA solution determines the license count.
- The VM Recovery Manager HA licenses are installed on an AIX partition that is designated to be the partition that is hosting the KSYS orchestrator. The VM Recovery Manager HA license enables the KSYS orchestrator

- The AIX partition that is hosting the VM Recovery Manager HA solution can be located anywhere in proximity to the HMCs and storage servers that the VM Recovery Manager HA solution is managing. For instance, if you are implementing the VM Recovery Manager HA solution, you can install the KSYS subsystem on an AIX partition of a system that is located outside of the managed host group.
- The VM Recovery Manager HA solution conforms to the active/inactive technology, which includes Live Partition Mobility (LPM) in your high availability configuration. Therefore, the entire VM is restarted on an alternate logical partition. You do not need to install the production licenses in the target system because the VM is restated along with the corresponding workload.

VM Recovery Manager HA limitations

Consider the following restrictions for the VM Recovery Manager HA solution.

KSYS limitations

- If a user-defined Shared Storage Pool (SSP) cluster name is same as the KSYS subsytem defined cluster name, or is in the KSYS_<KSYS_CLUSTER_NAME>_1_<siteID> format, the user-defined SSP cluster name and the KSYS subsystem defined cluster are considered to be the same. For example, if both the cluster names are the same, the KSYS subsystem removes the user-defined SSP cluster automatically while removing the KSYS defined cluster.
- The following commands can run without considering the policies:

```
ksysmgr verify host_group
ksysmgr lpm host action=validation
```

  Therefore, successful completion of the verification and validation operations does not mean that the virtual machines can be relocated successfully.
- In user-defined cluster, the following commands might remove all the data that you added:

```
ksysmgr modify host_group
ksysmgr remove host
```

- The KSYS subsystem follows the KSYS_*peer domain_HG_ID* format to name an HA SSP. The KSYS subsystem uses this format to differentiate between an HA SSP and a user-defined SSP. Therefore, you must not use this format for all user-defined SSPs.
- If access to the repository disk and the pool disk across the cluster or across some of the Virtual I/O Servers in the SSP cluster is lost, the status reporting operations and failover operations might be delayed. The discovery operation might also fail. Contact IBM Support to check whether any fixes are available for these issues.
- When the database network (DBN) node lose network connective or lose access to the pool of disks for a long time, all Virtual I/O Servers operate in local mode.
- When you remove the KSYS cluster, the KSYS subsystem fails to delete HA-specific VM and VIOS adapters if the cleanup operation continues for a long time. You must delete the VIOS adapters manually to avoid inconsistencies across the Virtual I/O Servers. If you create the KSYS cluster again, the KSYS subsystem can reuse the previous HA-specific adapters.
- The KSYS subsystem supports the Shared Storage Pool (SSP) cluster's high availability disk only when the Shared Storage Pool (SSP) is created from the KSYS subsystem.
- The KSYS subsystem does not display the high availability disk in any query when you use a user-defined SSP cluster.
- You cannot modify a KSYS subsystem's high availability disk after creating the SSP cluster from the KSYS node.
- After configuring a KSYS cluster and after configuring applications on that cluster, if you shut down a virtual machine (VM) or a logical partition (LPAR) of the cluster, the KSYS subsystem does not change the status of the VM or LPAR to red. The status remains green. However, if you shut down the same VM or LPAR from the HMC, the KSYS subsystem changes the status of the VM or LPAR to red.

- A maximum of 10 scripts can be added in the KSYS sub system for `add notify` command.
- The ksysmgr command fails to start or stop applications that are part of a dependency setup if Virtual I/O Servers are in local database mode.
  - Workaround: You must run the resume or suspend command for the VM.
- On VIOS nodes, if the disks of a shared storage pool (SSP) are not accessible after the system is re-activated due to shutdown or reboot, the disk state continues to be down. This impacts the start of pool and requires a quorum to come back online. As a workaround, choose one of the following options. If you do not want to reboot your VIOS, follow the workaround option 1.
  - Workaround option 1: Complete the following procedure:
    1. Restore the disk connectivity.
    2. Run the `cfgmgr` command as a root user to make the system aware of the disks.
    3. Run the command `padmin: clstartstop -stop -m <node>`.
    4. Run the command `padmin: clstartstop -start -m <node>`.
  - Workaround option 2: Complete the following procedure:
    1. Restore the disk connectivity.
    2. Reboot the VIOS node.
- For a VM with vSCSI disk, the cleanup operation fails in the local database mode.
  - Workaround: You must bring the SSP cluster back to the global mode.
- The KSYS subsystem does not handle the application dependency, if the VM has been shut down manually and the dependent application is part of the VM.
- VM Recovery Manager HA does not work if the Live Partition Mobility (LPM) features is disabled at firmware level.
- >| If a current repository disk is down, automatic replacement does not occur on previously used repository disk that has the same cluster signature. In this case, a free backup repository disk might not be available, hence the automatic replacement operation fails.
  - Workaround: Run the following command to clear the previous cluster signatures:

    ```
    cleandisk -r <diskname>
    ```

  |<
- >| The `ksysmgr -t remove cec` command is not supported on a user-defined KSYS cluster.
  - Workaround: Reconfigure the KSYS cluster. Otherwise, use the KSYS controlled VIOS cluster.

  |<
- >| In the scalability environment where the VMs are spread across the hosts of a host group, and the LPM verification operation is run on the host group, based on the type of configuration, at some point of time, many requests might go to one host and if the number of requests are more than the maximum requests that the host can handle, the verification operation might fail with following error:

```
HSCLB401 The maximum number of partition migration commands allowed are already in progress.
```

|<

Migration limitations

- Disable the quick discovery feature before running the Live Partition Mobility (LPM) and the restart operations on a virtual machines.
- You cannot run the Live Partition Mobility (LPM) operation simultaneously on multiple hosts by using the **`ksysmgr`** command. You must specify multiple virtual machines, in a comma-separated list, in the

ksysmgr command. Also, you can perform the LPM operation on a list of virtual machines simultaneously only if all the virtual machines are present in the same host.

- The flexible capacity policy is applicable only for VM failover operations. The flexible capacity function is not supported for virtual machines that are migrated by using the LPM operation.
- The flexible capacity reduction rules are applied even if 100% of resources are available on the target host.
- The flexible capacity policy is applicable only on CPU and memory resources. It is not applied on I/O resources. You must ensure enough I/O resources are available in the target host.
- If a VM migrates from host1 to host2, and applications in the VM become stable. At a later point of time, if the VM from the host2 needs to be migrated due to an application failure, the host1 will not be considered as a backup for application failure migration, because the VM had previously failed on host1. If host1 needs to be considered as a backup for future application failure, use the following workaround.
  - Workaround: After the VM is stable on the host2, clear the **FailedHostList** list of the VM. Run the command `chrsrc -s 'Name="VMName"' IBM.VMR_LPAR VmRestartFailedCecs='{""}'` to clear the **FailedHostList** list for the VM.
- The discovery operation or the KSYS restart operation automatically starts the dependency applications that were stopped by the user before the discovery or the restart of the KSYS subsystem.
  - Workaround: Complete the following procedure:
    1. Do not perform the discovery operation after stopping the dependency application.
    2. Disable the auto discover and the quick discovery features.
    3. Do not perform the KSYS subsystem restart.

VM agent limitations

- The **ksysvmmgr start|stop app** command supports only one application at a time.
- The **ksysvmmgr suspend|resume** command does not support the application dependencies.
- For all applications that are installed on the non-rootvg disks, you must enable the **automatic varyon** option for volume groups and the **auto mount** option for file systems after the virtual machine is restarted on the AIX operating system.
- If the application is in any of the failure states, for example, `NOT_STOPPABLE`, `NOT_STARTABLE`, `ABNORMAL`, or `FAILURE`, you must fix the failure issue, and then use the **ksysvmmgr start|resume application** command to start and monitor the application.
- If the KSYS cluster is deleted, or if a virtual machine is not included for the HA management, the VM agent daemon becomes inoperative. You must manually re-start the VM agent daemon in the virtual machine to bring the VM agent daemon to operative state.
- On a critical application failure, the KSYS subsystem continues to relocate the virtual machine from one host to another host even if the virtual machine relocated to the home host. For example, if a host group contains two hosts (host1 and host2) and a if the registered critical application in the vm1_host1 virtual machine fails, the KSYS subsystem relocates the vm1_host1 virtual machine to host2. If the application does not start in the NORMAL state, the KSYS subsystem again moves the vm1_host1 virtual machine to the host1, which is the home host for the application. The KSYS subsystem continues this relocation process until the application status becomes NORMAL.
- For the VMs running on the Linux VM agent, the reboot operation might take longer time than expected, and the discovery operation might fail and display the following message: 'Restart has encountered error for VM VM_Name'.
  - Workaround: Re-run the discovery operation.

GUI limitations

- The VM Recovery Manager HA GUI does not support multiple sessions that are originating from the same computer.

- The VM Recovery Manager HA GUI does not support duplicate names for host group, HMC, host, VIOS, and VMs. If a duplicate name exists in the KSYS configuration, the GUI might have issues during host group creation or in displaying the dashboard data.
- The VM Recovery Manager HA GUI refreshes automatically after each topology change (for example, VM migration operation and host migration operation). After the refresh operation is complete, the default KSYS dashboard is displayed. You must expand the topology to view the log information in the Activity window for a specific entity.
- Any operation performed by a user from the command-line interface of VM Recovery Manager HA is not displayed in the activity window of the VM Recovery Manager HA GUI.

Miscellaneous

- The VM Recovery Manager HA solution does not support internet Small Computer Systems Interface (iSCSI) disk type. Only N_Port ID virtualization (NPIV) and virtual Small Computer System Interface (vSCSI) disk types are supported.
- In a user-defined cluster, if you want to add a host or VIOS to the environment, you must add it in the shared storage pool (SSP) cluster first. Then, you can add the host or VIOS to the KSYS cluster. Also, if you want to remove a host or VIOS from the environment, you must first remove it from the KSYS cluster and then remove it from the SSP cluster.
- VM Recovery Manager HA supports only cluster and detailed type snapshot
- After each manage VIOS operation and unmanage VIOS operation, you must perform the discovery operation.

Errors that the KSYS subsystem cannot handle

The KSYS subsystem automatically restarts the VMs only when the KSYS subsystem is certain of the failures. If the KSYS subsystem is unsure, it sends an alert message to the administrator to review the issue and to manually restart VMs, if required.

Sometimes, the KSYS subsystem cannot identify whether the host failure is real or the host failure is because of a partitioning network. The KSYS subsystem does not automatically restart VMs in the following example scenarios:

- When the KSYS subsystem cannot connect to the HMC to quiesce the failed VM (fencing operation) on the source host before restarting the VM on the target host. The fencing operation is required to ensure that the VM is not running on two hosts simultaneously.
- The host monitor module and the VIOS can monitor their own network and storage. Sometimes, network and storage errors are reported by the VIOS and these error events are notified to the administrator through email and text messages. In these cases, the KSYS subsystem does not move the VMs automatically to avoid false relocation.
- When a host group is spread across two buildings with storage subsystem technologies such as IBM SAN Volume Controller (SVC) HyperSwap, where HMCs, hosts and other required resources exist in each building and the KSYS LPAR is deployed on the backup building, the following scenarios cannot be automatically handled:
  - **Power failure in the main building**: The KSYS subsystem cannot connect to the HMCs and hosts in the main site. The KSYS subsystem detects the host failure and notifies the administrator.
  - **Issues in network and storage partitioning between the buildings**: The KSYS subsystem cannot connect to the HMCs, and therefore notifies the administrator about the host failure. The administrator must review the environment and decide whether to move the VMs. The VMs might be operating correctly on the main host. The administrator can rectify the network links between the hosts and the KSYS subsystem will start operating in normal mode.

Installing VM Recovery Manager HA

The VM Recovery Manager HA solution provides high availability (HA) management for IBM Power Systems servers with PowerVM virtualization. After you plan the implementation of VM Recovery Manager HA solution, you can install the VM Recovery Manager HA software. The VM Recovery Manager HA solution uses other subsystems such as Hardware Management Console (HMC) and Virtual I/O Server (VIOS) that must exist in your production environment.

The following figure shows the key components in the VM Recovery Manager HA solution:

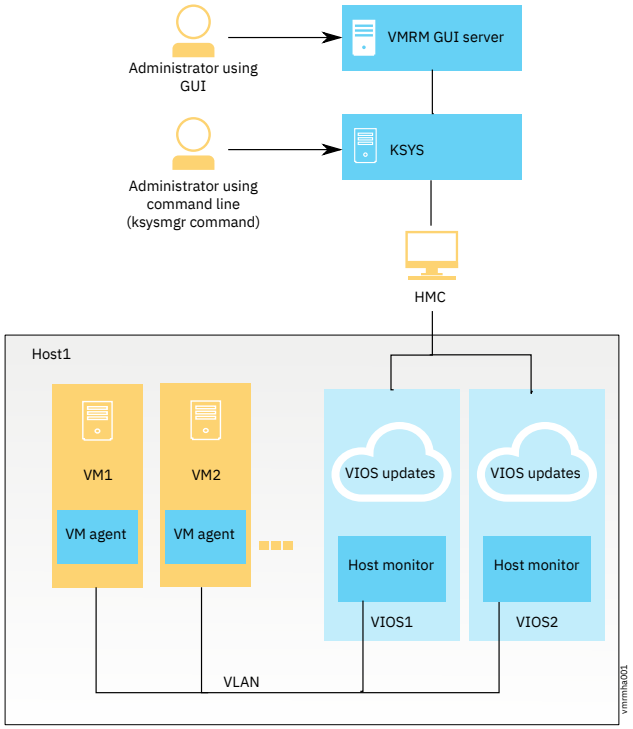

Figure 3. Components of the VM Recovery Manager HA solution

To install the VM Recovery Manager HA solution, you must first install the KSYS filesets. After the KSYS software is installed, the KSYS subsystem automatically monitors the health of hosts by enabling the host monitors in the VIOS partitions of each host that is part of the VM Recovery Manager HA management. You can optionally install the VM agents in the virtual machines that run AIX or Linux operating systems to monitor health of an individual virtual machine and applications that run in the virtual machines. You can also install the GUI server for the VM Recovery Manager HA solution to use the GUI by using a browser.

Complete the following procedures to install the VM Recovery Manager HA solution.

1. Install the VIOS interim fix.
2. Install the KSYS software.
3. Optional: Install the GUI server.
4. Optional: Install VM agents in the virtual machines.

Note: You must have root authority to perform any installation tasks.

Installing the VIOS interim fix

You must install the following VIOS interim fix on all VIOS instances that will be included in the KSYS subsystem: `https://aix.software.ibm.com/aix/efixes/IJ28933/IJ28933m1a.201106.epkg.Z`

Install the interim fix before you initialize the KSYS subsystem. Complete the following steps to install the VIOS interim fix:

1. Copy the interim fix to each of the VIOS instances.
2. Ensure that any cluster services are not active by running the **cluster -status** command. Stop any active cluster services by running the following command:

   ```
   clstartstop -stop -n clustername -m hostname
   ```

3. Run the following command in each of the managed VIOS instances:

   ```
   updateios -install -dev ifix_location -accept
   ```

   Follow the on-screen instructions. You might need to restart the system.
4. Verify whether the installation of interim fix is successful by running the following command:

   ```
   lssw
   ```

5. If the cluster services were stopped, start the cluster services by running the following command:

   ```
   clstartstop -start -n clustername -m hostname
   ```

Installing the KSYS software

You can use the **installp** command in the AIX LPAR to install KSYS filesets that are included in the package. Complete the following steps to install the KSYS software:

1. Ensure all the prerequisites that are specified in the Requirements topic are complete.
2. Navigate to the directory that contains the images that you want to install, and run the following command:

   ```
   installp -acFXYd fileset_location -V2 [-e filename.log] ksys.hautils.rte ksys.ha.license
   ksys.main.cmds
                     ksys.main.msg.en_US.cmds ksys.main.rte ksys.ui.agent ksys.ui.common
   ```

   The **-V2** flag enables the verbose mode of installation. Alternatively, you can use the `smit installp` command with the `all_latest` option to install all filesets in the directory.
3. Verify whether the installation of filesets is successful by running the following command:

   ```
   lslpp -l ksys.ha.license ksys.hautils.rte ksys.main.cmds
                     ksys.main.msg.en_US.cmds ksys.main.rte
   ```

   An output that is similar to the following example is displayed:

   ```
   Fileset                     Level  State      Description
     ----------------------------------------------------------------------------
   Path: /usr/lib/objrepos
   ksys.ha.license           1.5.0.0  COMMITTED  Base Server Runtime
   ksys.hautils.rte          1.5.0.0  COMMITTED  Base Server Runtime
   ksys.main.cmds            1.5.0.0  COMMITTED  Base Server Runtime
   ksys.main.msg.en_US.cmds  1.5.0.0  COMMITTED  Base Server Runtime
   ksys.main.rte             1.5.0.0  COMMITTED  Base Server Runtime
   ksys.ui.agent             1.5.0.0  COMMITTED  Base Server Runtime
   ksys.ui.common            1.5.0.0  COMMITTED  Base Server Runtime
   ```

4. Run the **/opt/IBM/ksys/ksysmgr** command to check the command line utility of the KSYS subsystem. The KSYS subsystem might take a few minutes to run the command for the first time. You can add the `/opt/IBM/ksys` directory to your *PATH* environment variable so that you can access the **ksysmgr** command easily.
5. After successful installation of KSYS filesets, enter the following command to check whether the class IDs are reserved:

```
cat /usr/sbin/rsct/cfg/ct_class_ids
IBM.VMR_HMC         510
IBM.VMR_CEC         511
IBM.VMR_LPAR        512
IBM.VMR_VIOS        513
IBM.VMR_SSP         514
IBM.VMR_SITE        515
IBM.VMR_SA          516
IBM.VMR_DP          517
IBM.VMR_DG          518
IBM.VMR_KNODE       519
IBM.VMR_KCLUSTER    520
IBM.VMR_HG          521
IBM.VMR_APP         522
IBM.VMR_CLOUD       523
```

6. If the IBM.VMR_APP or IBM.VMR_CLOUD class and its ID is not available in the output, contact IBM support to obtain a fix for APAR IJ29360.

Installing GUI server filesets

To use the VM Recovery Manager HA by using the graphical interface, you must install the GUI server fileset on a system to manage KSYS nodes by using the GUI. The logical partition, in which you want to install the GUI filesets, must be running IBM AIX 7.2 with Technology Level 2 Service Pack 1 (7200-02-01), or later. You can choose to install the GUI server fileset on one of the KSYS nodes. You must install the following GUI server filesets before you start the graphical interface. The GUI agent filesets are automatically installed along with the KSYS filesets. To install the GUI server filesets, complete the following steps:

1. Install the GUI server fileset based on the following scenarios:
   - If you are installing the GUI server filesets on one of the KSYS nodes, run the following command:

     ```
     installp -acFXYd fileset_location -V2 [-e filename.log] ksys.ui.server
     ```

   - If you are installing the GUI server filesets on a separate system that manages all the KSYS nodes, run the following command:

     ```
     installp -acFXYd fileset_location -V2 [-e filename.log] ksys.ha.license ksys.ui.server
     ksys.ui.common
     ```

2. Install the open source software packages, which are not included in the installed filesets, based on the following scenarios:
   - If the GUI server LPAR is connected to the internet, run the following command in the GUI server LPAR:

     ```
     /opt/IBM/ksys/ui/server/dist/server/bin/vmruiinst.ksh
     ```

     This command downloads and installs the remaining files that are not included in the filesets because these files are licensed under the General Public License (GPL).
   - If the GUI server LPAR is configured to use an HTTP proxy to access the internet, run the following command in the GUI server LPAR to specify the proxy information:

     ```
     /opt/IBM/ksys/ui/server/dist/server/bin/vmruiinst.ksh -p
     ```

     You can also specify the proxy information by using the *http_proxy* environment variable.
   - If the GUI server LPAR is not connected to the internet, complete the following steps:
     a. Copy the vmruiinst.ksh file from the GUI server LPAR to a system that is running the AIX operating system and that has internet access.
     b. Run the **vmruiinst.ksh -d /***directory* command where */directory* is the location where you want to download the remaining files. For example, /vmruiinst.ksh -d /tmp/vmrui_rpms.
     c. Download the following packages that are prerequisite packages for GUI server:

- info-4.13-3.aix5.3.ppc.rpm
- cpio-2.11-2.aix6.1.ppc.rpm
- readline-6.2-2.aix5.3.ppc.rpm
- libiconv-1.13.1-2.aix5.3.ppc.rpm
- bash-4.2-5.aix5.3.ppc.rpm
- gettext-0.17-6.aix5.3.ppc.rpm
- libgcc-4.9.2-1.aix6.1.ppc.rpm
- libgcc-4.9.2-1.aix7.1.ppc.rpm
- libstdc++-4.9.2-1.aix6.1.ppc.rpm
- libstdc++-4.9.2-1.aix7.1.ppc.rpm

d. Copy the downloaded files to a directory in the GUI server LPAR.

e. In the GUI server LPAR, run the **vmruiinst.ksh -i /***directory* command where */directory* is the location where you copied the downloaded files.

After the **vmruiinst.ksh** command is complete, a message displays a URL for the VM Recovery Manager HA GUI server. Enter the specified URL in a web browser in the GUI server LPAR and start using the VM Recovery Manager HA GUI.

Installing VM agents

VM agents are components that are installed in virtual machines (VMs) or logical partitions (LPARs). These optional agents offer robust monitoring of the VMs and applications that are running in VMs. You can manage HA applications in VMs through a lightweight application monitoring framework.

To install a VM agent in an AIX VM, go to Installing a VM agent in an AIX VM. For setting up a VM agent in Linux, see Installing a VM agent in a Linux VM.

Installing a VM agent in an AIX VM

1. Ensure all the prerequisites that are specified in the Requirements topic are complete.
2. Run the following command in the AIX virtual machine:

```
installp -acFXYd fileset_location -V2 [-e filename.log] ksys.vmmon.rte
```

3. Perform one of the following steps to verify whether the installation of VM agent is successful:
   a. Run the **lslpp** command.
   b. Ensure that the **ksysvmmgr** command and the binary file for the VM agent daemon exist in the following directories:
      - /usr/sbin/ksysvmmgr
      - /usr/sbin/ksys_vmmd
   c. Run the **lssrc -s ksys_vmm** command to verify whether the VM agent daemon is enabled. The status of the ksys_vmm subsystem must be Active in the output of this command.

Installing a VM agent in a Linux VM

To install the VM agent Red Hat Package Manager (RPM) packages in a Linux virtual machine, complete the following steps:

1. Ensure that the following Reliable Scalable Cluster Technology (RSCT) packages are installed in the Linux VM:
   - rsct.core
   - rsct.opt.storagerm
   - rsct.core.utils
   - rsct.basic

- DynamicRM

You can download the packages from the following link: http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/redhat/hmcmanaged/rhel7.html. For information about configuring the repository to easily install those packages, see Updating RSCT packages for PowerVM NovaLink.

2. Install the VM agent RPM packages based on the following Linux distributions in the virtual machine.

   In Red Hat Enterprise Linux (RHEL) (little endian) virtual machines, run the following command:

   ```
   rpm -ivh vmagent-1.3.0-1.0.el7.ppc64le.rpm
   ```

   In SUSE Linux Enterprise Server (SLES) (little endian) virtual machines, run the following command:

   ```
   rpm -ivh vmagent-1.3.0-1.0.suse123.ppc64le.rpm
   ```

3. Ensure RMC connection between the VMs and HMC. If the firewall is enabled on the RHEL VM, the RMC connection might be broken. Modify the firewall on the VMs to allow the RMC connection with the HMC. For details on modifying the firewall, see PowerLinux forum topic and Installing the PowerVM NovaLink software on a Red Hat Enterprise Linux partition topic.

Upgrading VM Recovery Manager HA

If you are upgrading the VM Recovery Manager HA solution from version 1.4 to version 1.5, you must install the packages that have updates to the existing software in your environment.

Prerequisites:

- You must have root authority to perform the installation tasks.
- When you install the new filesets, ensure that the existing version of the KSYS software is not running any active operations. The installation of the newer version of the KSYS software fails if the discovery, verification, or move operation is running in the existing version of the KSYS software.
- All KSYS-related operations such as discover, verification, move, and cleanup must be complete before you attempt to upgrade the existing version of the VM Recovery Manager HA solution to version 1.5.

Complete the following procedures to upgrade the VM Recovery Manager HA solution.

1. Upgrading the KSYS software.
2. Upgrading VM agents in the virtual machines.

Note: You must have root authority to perform any uninstallation tasks.

Upgrading the KSYS software

1. Copy the filesets to the location where the existing filesets are installed.
2. Decompress the filesets according to the guidelines that are provided with the package.
3. To install the filesets by using SMIT panel, complete the following steps:

   a. To open the SMIT panel, enter the following command:

   ```
   smit install
   ```

   b. In the **Install and Update Software** screen, select **Update Installed Software to Latest Level (Update All)**, and press Enter.

```
                          Install and Update Software

Move cursor to desired item and press Enter.

  Install Software
  Update Installed Software to Latest Level (Update All)
  Install Software Bundle
  Update Software by Fix (APAR)
  Install and Update from ALL Available Software
```

c. In the **Update Installed Software to Latest Level (Update All)** screen, change the values according to your situation. You must also accept new license agreements. Press `Enter` after you make all other changes.

4. Check the installation summary at the end of the installation output by scrolling to the end of the output. The output indicates whether the installation of your fileset was successful.

   If trace spooling is enabled and the trace file size is large, you must wait for a few minutes before you run the **`ksysmgr`** command. If the installation was not successful, check the reason of failure in the output.

Upgrading VM agents

Upgrade the VM agent RPM packages based on the following Linux distributions in the virtual machine:

- In Red Hat Enterprise Linux (RHEL) (little endian) virtual machines, run the following commands:

```
vmagent-1.5.0-2.0.el7.ppc64le.rpm
```

- In SUSE Linux Enterprise Server (SLES) (little endian) virtual machines, run the following command:

```
vmagent-1.5.0-2.0.suse123.ppc64le.rpm
```

These commands upgrade the VM agent software without modifying the current configuration of the virtual machines.

Uninstalling VM Recovery Manager HA

To uninstall the VM Recovery Manager HA solution, you can use the command-line interface to uninstall all the installed VM Recovery Manager HA filesets by running the **`installp -u`** command.

Complete the following procedures to uninstall the VM Recovery Manager HA solution.

1. Uninstall the KSYS software.
2. Uninstall the GUI server.
3. Uninstall VM agents in the virtual machines.

Note: You must have root authority to perform any uninstallation tasks.

Uninstalling the KSYS software

To uninstall the KSYS software, complete the following steps:

1. Choose one of the following options to remove the KSYS cluster:
   - Run the following command:

     ```
     ksysmgr delete cluster cluster_name
     ```

     This command removes the hosts from the host group and deletes the host group.
   - Complete the following steps:
     a. Remove all hosts from the existing host groups by running the following command:

        ```
        ksysmgr modify host_group host_group_name remove hosts=host_name
        ```

        All the corresponding virtual machines, Virtual I/O Servers, and Hardware Management Consoles will also be removed from the host group configuration.
     b. Delete the host group by running the following command:

        ```
        ksysmgr delete host_group host_group_name
        ```

c. Delete the host so that the VIOS remote copy programs (RCP), trunk adapters, and switches are removed by running the following command:

```
ksysmgr delete host host_name
```

2. Uninstall the KSYS filesets by running the following command:

```
installp -ug ksys.vmmon.rte
```

Uninstalling the GUI server filesets

If you had installed the GUI server filesets in the KSYS LPAR, you do not need to uninstall the GUI fileset separately. The GUI filesets will also be uninstalled with the KSYS filesets.

If you installed the GUI server filesets in a separate AIX LPAR, uninstall the GUI fileset by running the following commands:

```
installp -ug ksys.ui.server
installp -ug ksys.ui.common
```

Uninstalling VM agents

To uninstall the VM agent in an AIX VM, go to Uninstalling a VM agent in an AIX VM. To uninstall the VM agent in a Linux VM, see Uninstalling a VM agent in a Linux VM.

Uninstalling a VM agent in an AIX VM

1. Stop the VM agent module in the AIX virtual machine by running the following command:

```
ksysvmmgr stop
```

2. Uninstall the VM agent filesets from the AIX virtual machine by running the following command:

```
installp -u ksys*
```

Uninstalling a VM agent in a Linux VM

1. Stop the VM agent module in the Linux virtual machine by running the following command:

```
ksysvmmgr stop
```

2. Uninstall the VM agent package from the Linux virtual machine by running the following command:

```
rpm -e vmagent
```

Configuring VM Recovery Manager HA

After the VM Recovery Manager HA solution is installed, you must complete some mandatory configuration steps before you use the high availability feature of the VM Recovery Manager HA solution.

For a successful configuration, ensure that your environment meets all prerequisites that are specified in the “VM Recovery Manager HA requirements” on page 10 topic.

Setting up the KSYS subsystem

You can use the **ksysmgr** command or the VM Recovery Manager HA GUI to interact with the KSYS daemon to manage the entire environment for high availability.

The VM Recovery Manager HA solution monitors the hosts and the virtual machines when you add information about your environment to the KSYS configuration settings. Complete the following steps to set up the KSYS subsystem:

1. “Step 1: Initialize the KSYS cluster” on page 27.
2. “Step 2: Add HMCs” on page 28.
3. “Step 3: Add hosts” on page 28.
4. “Step 4: Create host groups” on page 28.
5. “Optional: Configure virtual machines” on page 29.
6. “Optional: Configure VIOS” on page 30.
7. “Step 5: Setting contacts for event notification” on page 31.
8. “Step 6: Enabling HA monitoring” on page 32.
9. “Step 7: Discovering and verifying the KSYS configuration” on page 32.
10. “Optional: Backing up the configuration data” on page 32.

Step 1: Initialize the KSYS cluster

The KSYS environment relies on Reliable Scalable Cluster Technology (RSCT) to create its cluster on the KSYS logical partition (LPAR). After you create the KSYS cluster, various daemons of RSCT and KSYS are activated. The KSYS node can then process the commands that you specify in the command line.

To create and initialize a KSYS cluster, complete the following steps in each of the KSYS LPARs:

1. Configure a cluster and add the KSYS node to the cluster by running the following command:

   ```
   ksysmgr add kyscluster cluster_name ksysnodes=ksysnode1 type=HA
   ```

2. Verify the KSYS cluster configuration by running the following command:

   ```
   ksysmgr verify kyscluster cluster_name
   ```

3. Deploy the KSYS cluster by running the following command:

   ```
   ksysmgr sync kyscluster cluster_name
   ```

   You can perform steps “1” on page 27 - “3” on page 27 by running the following command:

   ```
   ksysmgr add kyscluster cluster_name ksysnodes=ksysnodename1 sync=yes type=HA
   ```

4. Verify that the KSYS cluster is created successfully by running one of the following commands:

   ```
   # ksysmgr query kyscluster
   # lsrpdomain
   # lssrc -s IBM.VMR
   ```

The output message must display the state of the KSYS cluster as `Online`.

Step 2: Add HMCs

The KSYS interacts with the HMC for discovery, verification, monitor, recovery, and cleanup operations. HMCs provide details about the hosts and VIOS partitions that are managed by the HMCs. The VM Recovery Manager HA solution cannot be implemented without configuring the HMCs.

Note: The HMC user, whose user name and password details are provided to the KSYS, must have at least `hmcsuperadmin` privileges and remote access. The KSYS subsystem uses the Representational State Transfer (REST) API to communicate with the HMCs in the environment. Therefore, ensure that your environment allows HTTPS communication between the KSYS and HMC subsystems.

To add the HMCs to the KSYS configuration setting, complete the following steps in the KSYS LPAR:

1. Add the HMC with user name: `hscroot` and password: `xyz123` by running the following command:

   ```
   ksysmgr add hmc hmcname
         login=username
         password=password
         hostname|ip=hostname|ip
   ```

   For example, to add an HMC user with user name `hscroot` and an IP address, run the following command:

   ```
   ksysmgr add hmc hmc123 login=hscroot password=xyz123 ip=x.x.x.x
   ```

   To add an HMC user with user name `hscroot` and host name `hmc1.testlab.ibm.com`, run the following command:

   ```
   ksysmgr add hmc hmc123 login=hscroot password=xyz123 hostname=hmc1.testlab.ibm.com
   ```

2. Repeat step “1” on page 28 to add multiple HMCs.
3. Verify the HMCs that you have added by running the following command:

   ```
   ksysmgr query hmc
   ```

Step 3: Add hosts

After the HMCs are added to the KSYS subsystem, you can review the list of hosts that are managed by each HMC, and then identify the hosts that you want to add to the KSYS, for high-availability.

To add hosts to the KSYS configuration, complete the following steps in the KSYS LPAR:

1. Add the managed host to the KSYS by running the following command:

   ```
   ksysmgr add host hostname
         [uuid=uuid]
         [hostname|ip=hostname|ip]
   ```

   If the host is connected to more than one HMC, you must specify the Universally Unique Identifier (UUID) of the host. Hosts are identified by its UUID as tracked in the HMC. You can also use the **`ksysmgr query hmc`** command to identify the host name and the host UUID.
2. Repeat step “1” on page 28 for all hosts that you want to add to the KSYS subsystem.
3. Verify the hosts that you added by running the following command:

   ```
   ksysmgr query host
   ```

Step 4: Create host groups

You can group a set of hosts depending on your business requirements. Each host in the KSYS subsystem must be a part of a host group.

The KSYS subsystem creates a health monitoring Shared Storage Pool (SSP) cluster across the virtual I/O servers that are part of the host group. The health cluster monitors health of all virtual I/O servers across the cluster and retains the health data that is available to the KSYS subsystem by using a VIOS in the host group. The SSP cluster is used only by the KSYS. You must not use this SSP cluster for any other purpose. You can continue to use virtual Small Computer System Interface (vSCSI) or N_Port ID Virtualization (NPIV) modes of the cluster. However, if an SSP cluster exists in your environment, the KSYS subsystem does not deploy any new SSP clusters and instead, uses the existing SSP cluster for health management. However, if an existing SSP cluster is used, the KSYS subsystem might not support VIOS management.

The KSYS subsystem requires two disks to create the health monitoring SSP cluster across the Virtual I/O Servers in the host group. A disk of at least 10 GB is required to monitor the health of all hosts, called as a repository disk, and another disk of at least 10 GB is required to track the health data, called as a HA disk, for each host group. These disks must be accessible to all the managed Virtual I/O Servers on each of the hosts in the host group. You must specify the disk details when you create the host group or before you run the first discovery operation. You cannot modify the HA disk after the discovery operation is run successfully. If you want to modify the HA disk, you must delete the host group and re-create the host group with the HA disk details.

VM Recovery Manager HA supports automatic replacement of the repository disk. To automatically replace the repository disk, you must provide the details about backup repository disk. A maximum of six backup repository disks can be added for automatic replacement. When the storage framework detects failure of a repository disk, the KSYS subsystem sends an event notification. Then the KSYS sub system searches each disk on the backup repository list and locates a valid and active backup repository disk, and replaces the failed repository disk with the backup repository disk without any interruption. The failed repository disk will be placed as the last backup repository disk in the backup repository disk list. The failed backup repository disk can be reused after the disk failure is fixed and it becomes valid and active. The backup repository disk must meet all the VM Recovery Manager HA requirements for the automatic replacement of repository disk. For more information, see, “VM Recovery Manager HA requirements” on page 10.

If the backup repository disk is not specified, the automatic replacement feature is disabled. However, a failed repository disk can be replaced manually from the KSYS subsystem. For more information, see, Troubleshooting repository disk failure.

To create host group in the KSYS subsystem, complete the following steps in the KSYS LPAR:

1. Identify the available disks that you can designate as the repository disk and the HA disk for the SSP cluster and run one of the following commands:
   - `ksysmgr query viodisks vios=name1[,name2,..]`
   - `ksysmgr query viodisk hosts=host1[,host2,..]`
2. Create a host group and add the hosts and disks that you want in this host group by running the following command:

   ```
   ksysmgr add host_group hgname hosts= host1,host2
      repo_disk=diskuuid ha_disk=diskuuid
   backup_repo_disk=diskuuid1,diskuuid2…
   ```

   For repository disk failure issues, see Troubleshooting repository disk failure topic.
3. Repeat step “1” on page 29 for all host groups that you want to create in the KSYS subsystem.
4. Verify the host groups that you created by running the following command:

   ```
   ksysmgr query host_group
   ```

Optional: Configure virtual machines

When a host is added to the KSYS subsystem, all the virtual machines in the host are included by default in the HA management. If you do not want high availability for any of the virtual machines, you can exclude specific virtual machines from the HA management by running one of the following commands:

- ```
  ksysmgr unmanage vm name=vmname host=hostname | uuid=lparuuid |
                      ALL host=hostname | ALL host_group=hg_name
  ```
- ```
  ksysmgr unmanage vm vmname1|lparuuid1,...
  ```

You can include the VM back in the HA management at any time by using the **ksysmgr manage vm** command.

If you installed the VM agent for HA monitoring at the VM and application level, you can enable the HA monitoring by running the **ksysvmmgr start** command in the virtual machine. For more information about configuring the VM agent, see the "Setting up the VM agent" on page 38 topic.

Optional: Configure VIOS

When you add hosts to the KSYS subsystem, all the Virtual I/O Servers in the hosts are also added to the KSYS subsystem. The VM Recovery Manager HA solution monitors the hosts and virtual machines by using Virtual I/O Servers in the host.

The VM Recovery Manager HA solution requires at least two Virtual I/O Servers per host. You can have a maximum of 24 Virtual I/O Servers, spread across different hosts, in a single host group. If a host has more than 2 Virtual I/O Servers, you can exclude specific VIOS partitions from the HA management.

To exclude specific VIOS partitions from the HA management, complete the following steps:

1. Run the following command:

   ```
   ksysmgr unmanage vios viosname
   ```

   You can include the VIOS partition for the HA management at any time by using the **ksysmgr manage vios** *viosname* command.
2. Verify the existing Virtual I/O Servers by running the following command:

   ```
   ksysmgr query vios name
   ```

You can configure a specific LPAR and VIOS such that during each discovery operation, the KSYS subsystem fetches the size of the VIOS file system and the current file system usage in the VIOS. When the percentage of file system usage reaches the threshold value of 80%, the KSYS subsystem notifies you with a warning message so that you can make necessary updates to the VIOS file system.

The host monitor monitors the following file systems: `/, /tmp, /usr, /var, /home`. When the KSYS subsystem requests for the file system usage details, the host monitor responds with the details about the file system usage, which includes information about each file system and its usage. An event is generated when the file system usage surpasses the threshold value of the file system usage, also an event is generated when the file system usage comes under the threshold value.

VM auto-discovery

VM auto-discovery is a system-level property. You can disable or enabled this property. By default, this property is enabled.

By default, the KSYS subsystem manages all VMs automatically. The VM auto-discovery property allows the KSYS subsystem to manage or unmanage the newly created VMs and undiscovered VMs.

If the VM auto-discovery property is enabled, all VMs are managed automatically. If the auto-discovery property is disabled, all the newly created VMs on the KSYS managed hosts and the undiscovered VMs (existing VMs that are not yet discovered by the KSYS subsystem) will not be managed by the KSYS subsystem.

- To check whether the VM auto-discovery property is enabled or disabled to discover the resources across the site, run the following command:

  ```
  ksysmgr -a vm_auto_discovery query system
  ```

An output that is similar to the following example is displayed:

```
System-Wide Persistent Attributes
  vm_auto_discovery: enable
```

- To enable the VM auto-discovery property, run the following command:

```
ksysmgr modify system vm_auto_discovery=enable
```

- To disable the VM auto-discovery property, run the following command:

```
ksysmgr modify system vm_auto_discovery=disable
```

Step 5: Setting contacts for event notification

The KSYS subsystem tracks various events that occur in the environment, analyzes the situation, and notifies about any issues to the registered contacts. You must add the contact information to the KSYS subsystem so that you can receive notifications about any situation that might need your action.

You can add the following contact information for a specific user:

- Email address
- Phone number with phone carrier email address
- Pager email address

Note: The KSYS node must have a public IP address to send the event notifications successfully.

To register contact details to receive notification from the KSYS, run the following commands in the KSYS LPAR:

- To add an email address of a specific user to receive notification, enter the following command:

```
ksysmgr add notify user=username contact=email_address
```

For example,

```
ksysmgr add notify user=John contact=john.doe@testmail.com
```

You can add multiple email addresses for a specific user. However, you cannot add multiple email addresses simultaneously. You must run the command multiple times to add multiple email addresses.

- To add a specific user to receive an SMS notification, enter the following command:

```
ksysmgr add notify user=username
contact=10_digit_phone_number@phone_carrier_email_address
```

For example,

```
ksysmgr add notify user=John contact=1234567890@tmomail.net
```

You must specify the phone number along with the email address of the phone carrier to receive a short message service (SMS) notification. To determine the email address of your phone carrier, contact the phone service provider.

- To add a specific user to receive a pager notification, enter the following command:

```
ksysmgr add notify user=username contact=pager_email_address
```

For example,

```
ksysmgr add notify user=John contact=1234567890@SKYTEL.COM
```

Step 6: Enabling HA monitoring

You must enable HA monitoring for the KSYS subsystem to start monitoring the environment.

To enable HA monitoring, enter the following command:

1. Enable HA monitoring at system-level by running the following command:

```
ksysmgr modify system ha_monitor=enable
```

2. Enable HA monitoring at VM-level for each VM by running the following command:

```
ksysmgr modify vm vm1[,vm2,...] ha_monitor=enable
```

Step 7: Discovering and verifying the KSYS configuration

After adding various resources (HMCs, hosts, and host groups) to the KSYS subsystem, you must run the discovery operation. During the initial discovery operation, the KSYS subsystem creates the required high availability setup to monitor the VMs and hosts. The KSYS subsystem creates an SSP cluster based on the information that is specified in the configuration steps. During any subsequent discovery operations, the KSYS subsystem scans the environment for any changes to the environment and adapts to the modified environment. For example, when you add a host or when you run the Live Partition Mobility (LPM) operation from one host to another host that is outside of the current KSYS subsystem, the KSYS configuration settings are updated in the next discovery operation. By default, the KSYS subsystem automatically rediscovers sites once in every 24 hours during midnight. You can change this period by modifying the **auto_discover_time** system attribute.

After the KSYS subsystem discovers the resources, a verification is required to ensure that the virtual machines can be restarted on another host without any errors during a failover operation. The first discovery operation can take a few minutes because the SSP health cluster is deployed during the first discovery operation.

To discover and verify the configuration for a specific host group, complete the following steps:

1. Discover the resources by running the following command:

```
ksysmgr discover host_group hg_name
```

2. Verify the resources by running the following command:

```
ksysmgr verify host_group hg_name
```

You must run the discovery and verification commands each time you modify the resources in the KSYS subsystem. To perform both the discovery and verification operations, run the following command:

```
ksysmgr discover host_group hg_name verify=yes
```

Optional: Backing up the configuration data

As a best practice, after you configure the KSYS subsystem, back up the configuration data by using a snapshot. A snapshot preserves the configuration details of the KSYS environment at a specific point in time. Snapshots can be used when you upgrade a node or when an environment malfunctions because you can easily restore the configuration data without reconfiguring the hosts, HMCs, and other resources. When the KSYS LPAR fails, you can rebuild the KSYS LPAR by restoring the previously backed up configuration snapshot.

You can back up a cluster snapshot to save the KSYS cluster information or you can back up a basic snapshot to save details about other resources such as HMC, hosts, VIOS, excluded VMs. You can also choose to back up a detailed snapshot to save the application details. To save a snapshot of the current configuration data, run the following command:

```
ksysmgr add snapshot filepath=full_file_prefix_path|file_prefix
    type=CLUSTER|BASIC|DETAILED
```

You can restore the saved configuration snapshot by using the **ksysmgr restore snapshot filepath=filepath** command.

Setting up HA policies

After you set up the KSYS subsystem successfully, set up recovery policies to customize the default configuration settings to suit your high availability preferences.

Note: You must run the discovery and verification command after you set any policy.

The VM Recovery Manager HA solution provides the following options that you can customize:

HA monitoring

Turns on or turns off HA monitoring for the associated entity. The specified policy at the lowest resource level is considered first for HA monitoring. If you do not specify this policy for a resource, the policy of the parent resource is applied to the resource. For example, if you enable HA monitoring for the host group, HA monitoring is enabled for all virtual machines within the host group unless you disable HA monitoring for specific virtual machines.

You can enable HA monitoring for virtual machines only after you install the VM agent on each VM and start the VM agent successfully. For details, see Setting up the VM agent topic. If you do not set up the VM agent, the KSYS subsystem might return error messages for HA monitoring at VM-level.

To set the HA monitoring, run the following command:

```
ksysmgr modify system|host_group|vm name ha_monitor=enable|disable
```

ProactiveHA monitoring

ProactiveHA monitors every managed VM in the host group, the CPU utilization, and network packet loss during virtual machine or host monitor communication. When a VM's CPU utilization exceeds 90% or when network packet loss is detected on each of the VM's adapters during virtual machine or host monitor communication, an event is generated. The threshold for CPU utilization is predefined. By default, the ProactiveHA option is enabled.

Configuring network isolation events

The KSYS subsystem uses the network isolation feature to configure the VIOS `netmon` file, which is used by IBM Reliable Scalable Cluster Technology (RSCT) to monitor the network status. The KSYS subsystem generates the `NETWORK_ISOLATION_SUCESS` and the `NETWORK_ISOLATION_ERROR` events depending on whether the configuration of the VIOS `netmon` file succeeded. You can use the ksysmgr command to configure the IP addresses for the VIOS `netmon` file. After the discovery operation completes, the KSYS subsystem checks the configured IP addresses at site remote copy program (RCP) and generates a put message request for the host monitor to configure the VIOS `netmon` file. To add or delete the IP addresses for network isolation detection, run the following command:

```
ksysmgr modify system [network_isolation=<ip1,ip2,..|ALL> action=<add | delete>]
```

Restart policy

Indicates the KSYS subsystem to restart the virtual machines automatically during a failure. This attribute can have the following values:

- `auto`: If you set this attribute to `auto`, the KSYS subsystem automatically restarts the virtual machines on the destination hosts. The KSYS subsystem identifies the most suitable host based on free CPUs, memory, and other specified policies. In this case, the KSYS subsystem also notifies the registered contacts about the host or VM failure and the restart operations. This is the default value of the **restart_policy** attribute.
- `advisory_mode`: If you set this attribute to `advisory_mode`, the virtual machines are not restarted automatically after host or VM failures. In this case, the KSYS subsystem notifies the

registered contacts about the host or VM failures. The administrator must review the failure and manually restart the VMs on other hosts by using the **ksysmgr** commands.

To set the restart policy, run the following command syntax:

```
ksysmgr modify host_group name restart_policy=auto|advisory_mode
```

Host failure detection time

Indicates the time that the KSYS waits on a non-responsive host before the KSYS declares the host to be in an inactive state. This value is measured in seconds. The KSYS subsystem uses the specified time to ensure the health of the host and attempts to connect to the host before the KSYS declares the failure. After this duration, the virtual machines are restarted on another host that is located within the host group. The value of this attribute can be in the range 90 seconds – 600 seconds. The default value is 90 seconds.

To set the host failure detection time, run the following command:

```
ksysmgr modify system|host_group name
   host_failure_detection_time=time_in_seconds
```

VM failure detection speed

Represents the time that KSYS waits before KSYS declares the failure of a VM. You can select one of the following options: `fast`, `normal`, or `slow`. If you select `fast`, it means that the VM failure will be declared at the quickest time. The default value of this attribute is `normal`.

The time (in seconds) that KSYS waits is calculated based on the host failure detection time and the option that you specified:

- If you set the **vm_failure_detection_speed** attribute to `fast`, the VM failure detection time is calculated as follows: `Host failure detection time + VM threshold`
- If you set the **vm_failure_detection_speed** attribute to `normal`, the VM failure detection time is calculated as follows: `Host failure detection time + VM threshold*2`
- If you set the VM failure detection speed attribute to `slow`, the VM failure detection time is calculated as follows: `Host failure detection time + VM threshold*3`

where, `VM threshold` is a hardcoded value of 50 that is set in the KSYS subsystem. You cannot change this value.

Therefore, if you modify the **host_failure_detection_time** attribute, the value of VM failure detection time also changes. To set the VM failure detection speed, run the following command:

```
ksysmgr modify system|host_group|host|vm name
   vm_failure_detection_speed=fast|normal|slow
```

Failover priority

Specifies the order of processing of multiple VMs restart operations. For example, if a host fails and all the VMs must be relocated to other hosts in the host group, the priority of the VM determines which VM will be processed first. The supported values for this attribute are `High`, `Medium`, or `Low`. You can set this attribute at VM-level only. You must specify the UUID of the VM if you have two or more VMs with the same name. By default, all VMs in the host group have the priority of `Medium`.

To set the failover priority, run the following command:

```
ksysmgr modify vm name1[,name2,...] | filepath=filepath
priority=high|medium|low
```

Home host

Specifies the home-host of the virtual machine. By default, the KSYS subsystem sets this value initially to the host where the virtual machine was first discovered. You can change the home-host value of a virtual machine even when the virtual machine is running on another host. In such case, the specified home-host is used for all future operations. This attribute is useful when you get a host repaired after failure and you want to restart the virtual machines in its home-host.

To set the home-host value, run the following command:

```
ksysmgr modify vm name1[,name2...] homehost=hostname
```

Flexible capacity policy

Modifies the allocation of memory and CPU resources of a virtual machine when a virtual machine is moved from home-host to another host in the host group. You can set flexible capacity values based on the priority of a VM. You can set different flexible capacity values for various priorities of virtual machines: high, medium, and low. You must specify the flexible capacity values in percentage.

For example, you can define the following flexible capacity values at the host group level: 100% CPU and 100% memory for high priority VMs, 70% CPU and 80% memory for medium priority VMs, and 60% CPU and 75% memory for low priority VMs. When a medium priority VM is migrated from its home-host to another host in the host group, its capacity is adjusted to 70% CPU and 80% memory. If the VM is restored back to its home-host, the VM is restored with 100% resources.

The flexible capacity policy does not consider I/O slots, adapters, and resources that are available in the hosts. You must ensure that all the I/O virtualization requirements of the VMs are met within the host group environment. Also, the flexible capacity policy is applicable only to VM relocation that is based on restart operations. LPM operations do not follow the flexible capacity policy.

To set the flexible capacity policy, run the following command:

```
ksysmgr modify system | host_group | site
     [memory_capacity=(1-100) | minimum | current_desired | none | default] [priority=low|
medium|high]
     [cpu_capacity=(1-100) | minimum | current_desired | none | default] [priority=low|
medium|high]
```

Affinity policies

Specifies affinity rules for a set of VMs that defines how the VMs must be placed within a host group during a relocation. The following affinity policies are supported:

- `Collocation`: Indicates that the set of VMs must always be placed on the same host after relocation.

  To set this option, run the following command:

  ```
  ksysmgr add collocation name vm=vmname1[,...]>
  ksysmgr modify collocation name policy=add|delete vm=vm1[,...]
  ```

- `Anticollocation`: Indicates that the set of VMs must never be placed on the same host after relocation.

  To set this option, run the following command:

  ```
  ksysmgr add anticollocation name vm=vmname1[,...]
  ksysmgr modify anticollocation name policy=add|delete vm=vm1[,...]
  ```

- `Workgroup`: Indicates that the set of VMs must be prioritized first based on the assigned priority.

  To set this option, run the following command:

  ```
  ksysmgr add workgroup name vm=vmname1[,...]
  ksysmgr modify workgroup name policy=add|delete vm=vm1[,...]
  ```

- `Host blacklist`: Specifies the list of hosts that must not be used for relocating a specific virtual machine during a failover operation. For a virtual machine, you can add hosts within the host group to the blacklist based on performance and licensing preferences.

  To set this option, run the following command:

  ```
  ksysmgr modify vm vmname
     blacklist_hosts=hostname[,...] policy=add|delete
  ```

  Note: All the VMs in the Workgroup must have the same priority.

When you set the affinity policies, ensure that the host group has sufficient capacity for the policies to be implemented during host failure, VM failure, or application failure. For example, if the host group contains only two hosts, you cannot set anti-collocation policy on VMs of a specific host because the host group does not contain multiple target hosts to restart the virtual machines.

Dependency between applications of virtual machines within a host group

VM Recovery Manager HA supports parent-child dependency and primary-secondary dependency.

- **Parent-child dependency**

  Defines dependency between applications, which have a parent-child structure across VMs. A parent application can have multiple children application. A child application also can have multiple parent applications. A child application can be a parent application for another child application. Similarly, you can define parent-child dependency between applications for up to four levels. Command operations take effect from the parent application to the child application that are in a parent-child application structure.

  For example, if app1 is the parent application for the app2 child application, and if you run a command to shut down the app1 parent application, the app2 child application will not be monitored.

  To establish dependency between the parent application and the child application, run the following command:

  ```
  ksysmgr add app_dependency <name>  type=<parent_child> app_list=<vm1:app1,vm2:app2[,...]>
  ```

- **Primary-secondary dependency**

  Defines dependency between applications, which have a hierarchical primary-secondary structure across VMs.

  To establish dependency between the primary application and the secondary application, run the following command:

  ```
  ksysmgr add app_dependency <name> type=<primary_secondary> app_list=<vm1:app1,vm2:app2>
  ```

  Note: The **app_list** attribute must have only two *vmname:appname* pairs for the primary-secondary structure of applications across VMs.

You can verify the dependency between applications across VMs. You can also delete the dependency between applications.

- To verify a dependency that you have created, run the following command:

  ```
  ksysmgr query app_dependency [name]
  ```

- To delete a dependency that you have created, run the following command:

  ```
  ksysmgr delete app_dependency [name]
  ```

Limitations

- When you shutdown or reboot a virtual machine manually, the dependent applications are not affected. The recovery of dependent applications are considered only when failure has occurred with the parent application, the virtual machine, or the host.

Fibre channel (FC) adapter failure detection

The KSYS subsystem monitors Fibre Channel (FC) adapter status. An event is generated if adapter failure is detected. To use this feature, you must enable the ProactiveHA feature.

The following four events display the status of a fibre channel (FC) adapter.

- SFW_ADAP_DOWN
- SFW_ADAP_UP
- SFW_PORT_DOWN

- SFW_PORT_UP

Modifying the KSYS configuration

Growing business requirements might need changes in your current configuration, for example, adding a specific host to your environment or modifying the attributes of an existing resource. After the VM Recovery Manager HA solution is implemented in your environment, the KSYS subsystem continues to monitor for any changes in the KSYS configuration. The KSYS subsystem scans the environment regularly and adapts to changes. However, if you want KSYS to immediately update the changes, you must run the discovery and verification operations to discover and validate the changes in the KSYS configuration.

You can change the resource attributes as shown in the following examples:

- To change the host group details, run the following command:

```
ksysmgr modify host_group <name> options
      [memory_capacity=<(1-100) | minimum | current_desired  | none>
        priority=<low | medium | high>]
      [cpu_capacity=<(1-100) | minimum | current_desired | none>
        priority=<low | medium | high>]
      [skip_power_on=<yes|no>]
      [network=<vlanmap | vswitchmap>  sites=<siteA,siteB>
        siteA=<#[,#,...] || all> siteB=<#[,#,...] || all> [dr_test=<yes|no>]
        [policy=delete]]
      [ha_disk=<ViodiskID>]
      [repo_disk=<ViodiskID>]
      [backup_repo_disk=<ViodiskID[,ViodiskID2...]> | backup_repo_disk=none]
      [ha_monitor=<enable | disable>]
      [proactiveha=<enable | disable>]
      [restart_policy=<auto | advisory_mode>]
      [vm_failure_detection_speed=<fast | normal | slow>]
      [host_failure_detection_time=<90-600>]
    modify => mod*, ch*, set
    host_group => hg, host_g*
```

- To update the HMC name, login credentials, or IP address, run the following command:

```
ksysmgr modify hmc hmcname
     [login=new_username]
     [password=new_password]
     [hostname|ip=new_hostname|new_ip]
```

- To update the attributes of virtual machines, run the following command:

```
ksysmgr modify vm <vmname[,vmname2,...]> | ALL host=<hostname> | ALL
host_group=<host_group_name> | file=<filepath>
      [uuid=<uuid>]
      [homehost=<hostname>]
      [DrTargetHost=<hostname | none>]
      [priority=<Low|Medium|High>]
      [log_level=<0|1|2|3>]
      [skip_power_on=<yes|no>]
      [ha_monitor=<enable | disable>]
      [proactiveha=<enable | disable>]
      [vm_failure_detection_speed=<fast | normal | slow>]
      [blacklist_hosts=<hostname[,hostname2...]> | [policy=<add | delete>]]
    modify => mod*, ch*, set
    vm => lp*, vm
```

Note:

– `ksysmgr modify vm ALL host_group=<host_group_name>` doesn't work for the **blacklist_hosts** attribute

- To update the contact information to receive the KSYS notification, run the following command:

```
ksysmgr modify notify
     oldcontact=old_username newcontact=new_username
 ksysmgr modify notify
     oldcontact=old_email_address newcontact=new_email_address
```

- To update the system tunable attributes, run the following command:

```
ksysmgr modify system attribute=new_value
```

Removing resources from the KSYS subsystem

You can remove the resources as shown in the following examples:

- To remove a host group, run the following command:

```
ksysmgr delete host_group hg_name
```

- If an HMC, which is included in the KSYS configuration, is not managing any hosts, you can remove the HMC from the KSYS configuration by running the following command:

```
ksysmgr delete hmc hmcname
```

- To remove a host, you must first remove the host group that contains the host, and then remove the host by running the following commands:

```
ksysmgr delete host_group hg_name
ksysmgr delete host hostname
```

Start and stop applications

- To start an application, run the following command:

```
ksysmgr start app [name=<appname1,[appname2,...]>] vm=<vmname>
```

- To stop an application, run the following command:

```
ksysmgr stop app [name=<appname1,[appname2,...]>] vm=<vmname>
```

Note: The name attribute can have multiple applications, which must belongs to the same VM. If you do not provide the name of the application for the name attribute, the command will start or stop all the applications of the VM.

Setting up the VM agent

If you configure the high-availability function at VM-level or application-level, you must set up the VM monitor (VMM) and the application monitoring framework (AppMon) in each VM for which VM failure detection is required.

Notes:

- You can install the VM agent to monitor the virtual machine and to monitor applications on the virtual machines that run only the following operating systems:
  - AIX Version 6.1, or later
  - PowerLinux
    - Red Hat Enterprise Linux (RHEL) (little endian) Version 7.4, or later
    - SUSE Linux Enterprise Server (SLES) (little endian) Version 12.3, or later
- Currently, the high-availability feature at VM-level or application-level is not supported for IBM i operating system or other Linux distributions. However, you can enable these virtual machines for host-level health management. In addition, you can perform manual restart and LPM operations on these VMs.
- You can configure the VM agent by using the **ksysvmmgr** command only. You cannot configure the VM agent by using the VM Recovery Manager HA GUI.

- Apart from any general application that you want the VM agent to monitor, the VM Recovery Manager HA solution supports the following type of applications that can be monitored by using the in-built scripts in the corresponding versions of operating systems:

Table 2. Version support matrix for application types and operating systems

Application type	AIX operating system	Linux operating system (RHEL and SLES)
ORACLE	Oracle Database 12.1, or later	Not supported
DB2®	IBM DB2 11.3, or later	IBM DB2 10.5, or later
SAP_HANA	Not supported	SAP HANA 2.0, or later
POSTGRES	Postgres 9.2.23 or later	Postgres 9.2.23 or later

After you install the VM agent filesets successfully, complete the following steps to set up the VM monitor in each guest virtual machine:

1. Start the VM monitor daemon in each virtual machine to start monitoring the virtual machines and applications by running the following command:

   ```
   ksysvmmgr start [vmm] [<ATTR#1>=<VALUE#1>][,<ATTR#n>=<VALUE#n>]
   ```

   For example,

   ```
   ksysvmmgr start vmm log=2
   ```

   To find more information about the attributes for the VM monitor daemon, see the **ksysvmmgr** command topic.

2. Register applications in the VM monitor daemon that must be monitored for high availability by running the following command:

   ```
   ksysvmmgr [-s] add app name <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
   ```

   For example,

   ```
   ksysvmmgr -s add app app1 monitor_script=/tmp/monitor1.sh
         start_script=/tmp/start1.sh stop_script=/tmp/stop1.sh
   ksysvmmgr -s add app oracleapp type=ORACLE instancename=orauser database=DBRESP
   ksysvmmgr -s add app sapapp type=SAPHANA instancename=S01 database=HDB01 configfile=/var/ksys/
   config/SAPHANAconfig.xml
   ksysvmmgr -s add app db2app type=DB2 instancename=db2inst1
   ksysvmmgr -s add app postgresapp type=POSTGRES instancename=postgres database=testdb
   configfile=/var/ksys/config/samples/POSTGRESconfig.xml
   ```

   After you add an application, if the application fails or stops working correctly, the VM agent attempts to restart the application in the same virtual machine for several times as specified in the **max_restart** attribute for the VM, which is set to 3 by default. If the application is still not working correctly, the KSYS subsystem notifies you about the issue. You can manually review the problem and restart the application.

3. Mark important applications as critical by running the following command:

   ```
   ksysvmmgr -s modify app app2 critical=yes
   ```

   When you mark an application as critical, if the application fails or stops working correctly, the VM agent attempts to restart the application for several times as specified in the **max_restart** attribute for the VM. If the application is still not working correctly, the KSYS subsystem notifies you about the issue and attempts to restart the virtual machine on the same host. If the application is still not working correctly, that is if the application status is displayed as RED when you run the **ksysmgr query app** command, the KSYS restarts the VM on another host within the host group based on the specified policies.

4. If some applications have dependencies, for example, if you need to specify a sequence of applications to start or stop, specify the dependencies by running the following command:

   ```
   ksysvmmgr [-s] add dependency <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
   ```

   For example,

   ```
   ksysvmmgr -s add dependency dependency_list=app1,app2,app3
             dependency_type=start_sequence

   ksysvmmgr -s add dependency dependency_list=app4,app5,app6
             dependency_type=stop_sequence

   ksysvmmgr -s add dependency dependency_list=app1,app2
            dependency_type=parent_child
   ```

   When you add a dependency, if the virtual machines are restarted on the same or another host, the applications in the virtual machine are started based on the specified dependency.

5. If you need to modify any attributes for the VM monitor daemon, applications, or application dependency, enter one of the following commands:

   ```
   ksysvmmgr [-s] modify [vmm] <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
   ksysvmmgr [-s] modify app <name> <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
   ksysvmmgr [-s] modify dependency <DEPUUID> <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
   ```

6. Save a snapshot of the VM monitor configuration by running the following command:

   ```
   ksysvmmgr backup vmm filepath=full_file_prefix_path
   ```

   You can restore the saved snapshot by using the **ksysvmmgr -s restore vmm filepath=***filepath* command.

Recovering hosts, VMs, and applications

After you configure the resources and policies in the KSYS subsystem, the KSYS continues to monitor the environment for any failures or issues. When any planned or unplanned outage occurs, the KSYS subsystem restarts the virtual machines on another host based on the specified policies.

The KSYS subsystem can be configured to monitor the hosts, virtual machines (VMs), and applications to perform recovery operations. By default, all VMs in the host group are managed. If you unmanage one or more VMs and run the discovery operation, the KSYS subsystem does not monitor and manage those VMs for high availability. Similarly, the KSYS subsystem does not monitor the specified resources for high availability if you disable HA monitoring for the entire system, a host group, or a host. However, you can monitor the health of VM and application only when you install the VM agent and enable HA monitoring in each VM that needs HA monitoring.

Recovering virtual machines in an unplanned outage

The KSYS subsystem performs the following types of recoveries for unplanned outages based on specified policies:

Automatic restart of virtual machines

When a host, VM, or critical application fails and the **restart_policy** attribute is set to auto, the KSYS subsystem restarts the virtual machines automatically on other hosts. The KSYS notifies you about the events; you do not have to take any actions.

However, if the KSYS subsystem could not successfully stop the VMs in the source host, the VMs are not restarted automatically. Also, if the KSYS subsystem identifies a problem, but cannot determine the issue, the VMs are not restarted on other hosts automatically to avoid unnecessary outage because of false failure detection. In both these cases, the KSYS subsystem notifies you about the problem. You must review the notified problem and then manually start the VMs, if necessary.

Manual recovery of virtual machines

When a host, VM, or critical application fails and the **restart_policy** attribute is set to `advisory_mode`, the KSYS notifies you about the issue. You can review the issue and manually restart the virtual machines on another hosts.

If you have configured the VM agent in each of your virtual machines, the KSYS notifies you when a virtual machine or a registered critical application fails or stops working correctly. In such cases also, you can restart the virtual machines on another hosts based on the specified policies.

To restart the virtual machines manually on another host, complete the following steps:

1. Restart specific virtual machines or all virtual machines in a host by running the following command:

   ```
   ksysmgr [-f] restart vm vmname1[,vmname2,….] [to=hostname|uuid]
   ```

   Or,

   ```
   ksysmgr [-f] restart host hostname|uuid [to=hostname|uuid]
   ```

   After the virtual machines are restarted successfully, the KSYS subsystem automatically cleans the VMs on the source host and on the HMC by removing the LPAR profile from the HMC.
2. If the output of the **ksysmgr restart** command indicates cleanup errors, clean up the VM details manually in the source host by running the following command:

   ```
   ksysmgr cleanup vm vmname host=hostname
   ```
3. If the restart operations fail, recover the virtual machine in the same host where it is located currently by running the following command:

   ```
   ksysmgr [-f] recover vm vmname
   ```

Planned migration of virtual machines

The KSYS subsystem uses the HMC-provided Live Partition Mobility (LPM) capability to support the planned HA management. You can also use HMC to perform the LPM operations and the KSYS adapts to the movements of the VMs within the host group as part of its regular discovery operation. If you plan for a host maintenance or an upgrade operation, you can move all the virtual machines to another host by using the LPM operation and also restore the virtual machines back to the same host after the maintenance or the upgrade operation is complete. You can also test whether the movement of virtual machines to another host will be successful by using LPM validation without moving the virtual machines. This validation is useful to avoid any errors that might occur during the relocation of virtual machines.

To migrate the virtual machines to another host by using the LPM operation, complete the following steps:

1. Validate the LPM operation without migrating the virtual machines by running one of the following commands:
   - To validate the LPM operation for specific virtual machines, run the following command:

     ```
     ksysmgr [-f] lpm vm vmname1[,vmname2,..] action=validate
     ```
   - To validate the LPM operation for all virtual machines in a specific host, run the following command:

     ```
     ksysmgr [-f] lpm host hostname|uuid action=validate
     ```

   If the output displays any errors, you must resolve those errors.
2. Migrate the virtual machines from the source host to another host by running one of the following commands:
   - To migrate specific virtual machines, run the following command:

     ```
     ksysmgr [-f] lpm vm vmname1[,vmname2,..] [to=hostname|uuid]
     ```

- To migrate all virtual machines in a specific host, that is to migrate all the VMs from the host, run the following command:

  ```
  ksysmgr [-f] lpm host hostname|uuid [to=hostname|uuid]
  ```

  When you run this command, the virtual machines are restarted on another host according to the specified policies in the KSYS configuration settings. If you have not specified the destination host where the virtual machines must be started, the KSYS subsystem identifies the most suitable host that can be used to start each virtual machine.

  If you have HMC Version 9 Release 9.3.0, or later, you can view the LPM progress as a percentage value.

3. Run the discovery and verify operations after each LPM operation to update the LPM validation state by running the following command:

   ```
   ksysmgr discover host_group hg_name verify=true
   ```

4. After the maintenance or upgrade activities are complete in the source host, restore all virtual machines by running the following command:

   ```
   ksysmgr restore host hostname|uuid
   ```

Configuring the sudo command

You can configure the sudo command in the VM Recovery Manager HA solution.

Prerequisites

The AIX operating system does not have sudo features by default. You must download the `sudorpm` package from the web and install it in the KSYS node.

Configuring the Sudo

Only a root user can configure sudo features in the AIX operating system. The username of the user must be added to the sudoers file for the user to run the sudo command.

Editing the sudoers file

To configure the sudo command, you can edit the sudoers file by using the `visudo` command. To enable the user to run the commands, in the sudoers file, under the user privilege specification, specify the username and commands. The user can run only the commands specified in the user privilege section for the user.

An example sudoers file follows:

```
##
## Host alias specification
##
##
## User alias specification
##
##
## Cmnd alias specification
##
#Cmnd_Alias SU = /usr/bin/su
##
## Uncomment to enable logging of a command's output, except for
## sudoreplay and reboot.  Use sudoreplay to play back logged sessions.
# Defaults log_output
# Defaults!/usr/bin/sudoreplay !log_output
# Defaults!/usr/local/bin/sudoreplay !log_output
# Defaults!REBOOT !log_output
##
## Runas alias specification
##
##
## User privilege specification
##
root ALL=(ALL) ALL
<username> ALL=(ALL) /opt/IBM/ksys/ksysmgr q vm, /opt/IBM/ksys/ksysmgr q host
## Allows people in group wheel to run all commands
# %wheel        ALL=(ALL)  ALL
```

Verifying the execute permissions for a user

To verify the execute permission that you provided for the user in the sudoers file, complete the following steps.

1. Log in to the user account for which you have provided the execute permission.
2. Run the following command:

   ```
   # sudo /opt/IBM/ksys/ksysmgr q host
   ```

   The command runs successfully and an output similar to the following example is displayed.

   ```
   ERROR: KSYS subsystem is currently offline, please sync ksyscluster to start KSYS "ksysmgr
   sync ksyscluster <name>"
   ```

3. Run the following command:

   ```
   # sudo /opt/IBM/ksys/ksysmgr q vm
   ```

   The command runs successfully and an output similar to the following example is displayed.

   ```
   ERROR: KSYS subsystem is currently offline, please sync ksyscluster to start KSYS "ksysmgr
   sync ksyscluster <name>"
   ```

4. Run the following command:

   ```
   # sudo /opt/IBM/ksys/ksysmgr q vios
   ```

   The command does not run successfully and an output similar to the following example is displayed.

   ```
   Sorry, <username> is not allowed to execute '/opt/IBM/ksys/ksysmgr q vios' as root on
   nuts007.ausprv.stglabs.ibm.com.
   ```

Since you provided the execute permission to the user for the first two commands in the example sudoers file (see the previous topic), the `ksysmgr q vm` and `ksysmgr q host`, commands ran successfully, and because you did not provide the execute permission to the user for the `ksysmgr q vios` command in the example sudoers file, this command did not run successfully, and a message stating that the user is not allowed to execute the command was displayed.

Granting the execute permission for all commands

You can provide the execute permission to run all commands for a user on the KSYS node. Edit the *<username>* ALL attribute in the sudoers file as shown.

```
root ALL=(ALL) ALL
<useranme> ALL=(ALL) /opt/IBM/ksys/
```

Troubleshooting the sudo command configuration error

If the configuration of the sudo command is not successful, the `sudo -v` command displays the following error:

```
sudo: error in /etc/sudo.conf, line 0 while loading plugin "sudoers_policy"
sudo: unable to load /opt/freeware/libexec/sudo/sudoers.so:     0509-022 Cannot load module /opt/
freeware/libexec/sudo/sudoers.so.
        0509-150   Dependent module /usr/lib/libcrypto.a(libcrypto.so.1.0.2) could not be loaded.
        0509-152   Member libcrypto.so.1.0.2 is not found in archive
        0509-022 Cannot load module /opt/freeware/libexec/sudo/sudoers.so.
        0509-150   Dependent module /opt/freeware/libexec/sudo/sudoers.so could not be loaded.
sudo: fatal error, unable to load plugins
```

To resolve this error, export the library path `LIBPATH=/opt/freeware/lib:$LIBPATH` by running the `export` command.

VM agents

This section describes the VM agents that VM Recovery Manager HA supports.

DB2

Scripts:

The KSYS VM daemon uses the following scripts to start, stop, and monitor the DB2 application.

- `/usr/sbin/agents/db2/startdb2`
- `/usr/sbin/agents/db2/stopdb2`
- `/usr/sbin/agents/db2/monitordb2`

To add the DB2 VM agent, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] add app <NAME> type=DB2 instancename=<VALUE#1>
(database=<VALUE#2>) [<ATTR#n>=<VALUE#n>]
```

Example:

```
ksysvmmgr -s add app db2app type=DB2 instancename=db2inst1 database=dbone
```

In this example, *db2inst1* is the database instance owner and *dbone* is the database to monitor.

Attributes:

- **type**: While creating the DB2 application, the *type* attribute must have the value, DB2.
- **instancename**: The **instancename** attribute must be specified with the DB2 instance owner.

The **instancename** attribute is also passed as a parameter to the start and stop scripts.

Oracle

Scripts:

The KSYS VM daemon uses the following scripts to start, stop and monitor the oracle application.

- `/usr/sbin/agents/oracle/startoracle`
- `/usr/sbin/agents/oracle/stoporacle`

- `/usr/sbin/agents/oracle/monitororacle`

To add the Oracle VM agent, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] add app <NAME> type=ORACLE instancename=<VALUE#1>
database=<VALUE#2> [<ATTR#n>=<VALUE#n>]
```

Example:

```
ksysvmmgr -s add app oracleapp type=ORACLE instancename=orauser database=DBRESP
```

In this example, *orauser* is the Oracle username (SID) and *DBRESP* is the oracle system identifier (SID) or database name.

Attributes:

- **type**: While creating the ORACLE application, the **type** attribute must have the value, ORACLE.
- **instancename**: The **instancename** attribute must be specified with the ORACLE *username*.
- **database**: The database name must be specified with the ORACLE system identifier (SID).

POSTGRES

Scripts:

The KSYS VM daemon uses the following scripts to start, stop and monitor the POSTGRES application.

- `/usr/sbin/agents/postgres/startpostgres`
- `/usr/sbin/agents/postgres/stoppostgres`
- `/usr/sbin/agents/postgres/monitorpostgres`

To add the POSTGRES VM agent, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] add app <NAME> type=POSTGRES instancename=<VALUE#1>
database=<VALUE#2> configfile=<VALUE#3> [<ATTR#n>=<VALUE#n>]
```

Example:

```
ksysvmmgr -s add app postgresapp type=POSTGRES instancename=postgres database=testdb
configfile=/var/ksys/config/samples/POSTGRESconfig.xml
```

In this example, *postgres* is the instance owner and *testdb* is the database name.

Attributes:

- **type**: While creating the POSTGRES application, the **type** attribute must have the value, POSTGRES.
- **instancename**: The **instancename** attribute must be specified with the POSTGRES instance owner.
- **database**: The database name is optional.

  Note: If the database name is specified, the script monitors only the specified database. And if the database name is not specified, the script monitors all database of the postgres instance.
- **configfile**: The **configfile** attribute specifies the file path of the configuration file that stores settings of the application configuration. You must specify the path of the configuration file while adding the POSTGRES application. A sample configuration file, `POSTGRESconfig.XML` is provided in the `/var/ksys/config/samples` folder. You can use this sample file by updating the attribute values.

  If you do not specify appropriate values in the configuration file, you canto add the POSTGRES application agent.

The following table lists the description of the attributes present in the POSTGRES configuration file.

Table 3. POSTGRES attributes	
Attribute	**Description**
POSTGRESinstance ID	The ID of the POSTGRES instance owner. This attribute is mandatory in the configuration file.
data_directory	The file system location of the database configuration files. This attribute is specified while initializing the POSTGRES application (by using the `init` command with the `-D` option). This attribute is mandatory in the configuration file.

The flowing screen displays an example of a POSTGRES configuraiton file (`/var/ksys/config/samples/ POSTGREConfig.xml`) with example data:

```
<?xml version="1.0" encoding="UTF-8"?>
<POSTGRESConfig>
<!--NOTE: PLEASE UNCOMMENT THE REQUIRED ATTRIBUTES AND SET APPROPRIATE VALUES.-->
    <POSTGRESinstance id="pginst123">
        <data_directory>/var/lib/pgsql/data</data_directory>
    </POSTGRESinstance>

    <!--POSTGRESinstance id="instOwner"-->
        <!--data_directory>/var/lib/pgsql/data</data_directory-->
    <!--/POSTGRESinstance-->
</POSTGRESConfig>
```

SAP_HANA

IBM VM Recovery Manager HA for Power Systems supports the following SAP HANA configurations:

1. **SAP HANA scale up configuration with host-based replication**: you can create a replication between two SAP HANA nodes and add them to a VM agent. IBM VM Recovery Manager HA for Power Systems manages the SAP HANA nodes and replication between the two node, such as takeover of primary node failures.
2. **SAP HANA scale up configuration without replication**: You can install the SAP HANA DB and add it to a VM agent. The VM agent monitors the database status and manages any failures.

Scripts:

The KSYS VM daemon uses the following scripts to start, stop and monitor the SAP HANA application.

- `/usr/sbin/agents/saphana/startsaphana`
- `/usr/sbin/agents/saphana/stopsaphana`
- `/usr/sbin/agents/saphana/monitorsaphana`

To add the SAPHANA VM agent, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] add app <NAME> type=SAPHANA instancename=<VALUE#1>
database=<VALUE#2> configfile=<VALUE#3> [<ATTR#n>=<VALUE#n>]
```

Examples:

- To add the SAP HANA application without replication:

  ```
  ksysvmmgr -s add app sapapp1 type=SAPHANA instancename=S01 database=HDB01
  ```

- To add the SAP HANA application with replication:

  ```
  ksysvmmgr -s add app sapapp2 type=SAPHANA instancename=S01 database=HDB01
  configfile=/var/ ksys/config/samples/SAPHANAconfig.xml
  ```

In these examples, *S01* is the SAPHANA system ID, *HDB01* is the database name.

Attributes:

- **type**: While creating a SAP HANA application, the **type** attribute must have the value, SAPHANA.
- **instancename**: The **instancename** attribute must be specified with SAP HANA system ID.
- **database**: The **database** attribute must be specified with the SAP HANA database name.
- **configfile**: The **configfile** attribute specifies the file path of the configuration file, which stores settings of the application configuration. This attribute is not required for the SAP HANA application without replication configuration. But, you must specify the path of the configuration file while adding the SAP HANA application with replication configuration. A sample configuration file, POSTGRESconfig.XML is provided in the /var/ksys/config/samples folder. You can use this sample file by updating the attribute values. If you do not specify the configuration file path or appropriate values in the configuration file, the SAP HANA application will be added without replication configuration.

  The SAPHANAconfig.xml file contains the following attributes:
  instanceId, **replication**, **role**, **localsite**, **remotesite**, **secondarymode**, **virtualip**, **interfacename**, **executabledir**, **timeout** and **remotenode**.

  If the replication attribute of the instance is set to yes, you must specify values for all mandatory attributes (for example, **replication**, **role**, **localsite**, **remotesite**, **secondary mode,** and **remotenode**). If you do not specify values for all mandatory attributes, the SAP HANA application will not be added to the virtual machine.

  Few attributes are optional: **virtualip**, **interfacename**, **executabledir** and **timeout**.

The following table lists the description of the attributes present in the SAP HANA configuration file.

Table 4. SAP_HANA attributes

Attributes	Description
SAPinstance id	The SAP HANA system ID must be specified while installing the SAP HANA application.
replication	Defines whether the SAP HANA application runs in stand-alone or in replication mode. • The keyword yes indicates that the SAP HANA application is in replication mode • The keyword no indicates that the SAP HANA application is in stand-alone mode
role	The value of the attribute can be primary, sync, syncmem, or async. For a primary node, the value of this attribute is primary and for secondary node, the value of this attribute is sync, syncmem, or async.
localsite	The site name of the current node (the value that is specified while configuring the SAP HANA replication).
remotesite	The site name of the remote node (the value that is specified while configuring the SAP HANA replication).
secondarymode	The mode of secondary node. The value of the attribute can be syn, syncmem, or async.
remotenode	Hostname of the remote node. This name should be the same as the name shown in the output of the SAP command, hdbnsutil -sr_state.

Table 4. SAP_HANA attributes (continued)	
Attributes	**Description**
virtualip	The service or virtual IP address that can be aliased to **interfacename** and through which the primary or secondary node can be accessed by other applications. You can also configure the netmask using the **ipaddress/netmask** format for the interface. If a subnet mask is not provided, the **virtualip** takes the default value **ipaddress/32**.
subnet mask	The subnet mask can be used with the service or the **virtualip** address.
interfacename	The interface on which the specified virtual or service IP address will be aliased.
executabledir	The directory path where the shared libraries and executable files of SAP_HANA VM agent are present.
timeout	The duration of time within which the SAP_HANA command is expected to complete. This value is indicated in seconds. The default value is 120 seconds.

An example configuration file for the SAP HANA configuration without replication setup follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<HANAConfig>
<NOTE: PLEASE UNCOMMENT THE REQUIRED ATTRIBUTES AND SET APPROPRIATE VALUES.
If replication attribute is 'yes', then all below attributes are mandatory.>
<SAPinstance id="S01">
<replication>no</replication>
<role>primary</role>
<localsite>local.hostname.ibm.com</localsite>
<remotesite>remote.hostname.ibm.com</remotesite>
<secondarymode>sync/syncmem/async</secondarymode>
<virtualip>192.168.2.3/24</virtualip>
<interfacename>eth2</interfacename>
<remotenode>remotenode_name</remotenode>
<timeout>120</timeout>
<executabledir>/usr/sap/S01/HDB01/exe</executabledir>
</SAPinstance>
</HANAConfig>
```

An example configuration file for the secondary node SAP HANA configuration with replication follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<HANAConfig>
<NOTE: PLEASE UNCOMMENT THE REQUIRED ATTRIBUTES AND SET APPROPRIATE VALUES.
If replication attribute is 'yes', then all below attributes are mandatory.>
<SAPinstance id="S01">
<replication>yes</replication>
<role>syncmem</role>
<localsite>local.hostname.ibm.com</localsite>
<remotesite>remote.hostname.ibm.com</remotesite>
<secondarymode>syncmem</secondarymode>
<virtualip>192.168.2.3</virtualip>
<interfacename>eth2</interfacename>
<remotenode>remotenode_name</remotenode>
<timeout>120</timeout>
<executabledir>/usr/sap/S01/HDB01/exe</executabledir>
</SAPinstance>
</HANAConfig>
```

Setting up KSYS high availability through PowerHA SystemMirror

The KSYS subsystem is a major component in VM Recovery Manager HA, which monitors and manages the complete environment health. Hence, setting up high availability for the KSYS subsystem will be helpful to handle any scenario where the KSYS daemon hanged or the KSYS node went down. This high availability can be set up by managing the KSYS daemon by using PowerHA SystemMirror software. To manage the KSYS daemon through PowerHA SystemMirror software, PowerHA SystemMirror must be configured to monitor and manage the KSYS daemon by using custom scripts.

Prerequisite

- PowerHA SystemMirror 7.2.1, or later must be installed
- VM Recovery Manager HA 1.5, or later must be installed
- The `/etc/hosts` and `/etc/cluster/rhosts` files must be modified to include all PoweHA nodes
- The variable *CT_MANAGEMENT_SCOPE=2* must be defined in the `.profile` file for all nodes, and must be exported by running the `export` command
- When two-node PowerHA SystemMirror is integrated with the KSYS subsystem, each node needs to have minimum 30 GB space.

Concept

When the PowerHA SystemMirror cluster is created, Reliable Scalable Cluster Technology (RSCT) cluster is also created. The RSCT cluster creates the resource manager (RM) framework and allocates resources (including KSYS RM resources) to the KSYS subsystem. Hence, the KSYS subsystem can use the RSCT cluster and the resource manager framework instead of creating a new RSCT cluster. This ensures that the configuration settings and saved data or modifications to the data in one KSYS node reflects in the other KSYS node. The KSYS daemon can be monitored by custom scripts of the PowerHA SystemMirror resource group. The resource group remains online on one node at a time. If the KSYS node goes down, the resource group moves to a different node. The KSYS node, to which the resource group has moved, starts monitoring and managing the environment. This ensures high availability between KSYS nodes.

KSYS high availability through PowerHA SystemMirror

Complete the following procedures to set up the KSYS subsystem's high availability through PowerHA SystemMirror:

1. “Configuring PowerHA SystemMirror in VM Recovery Manager HA solution” on page 49
2. “Configuring the KSYS subsystem for high availability through PowerHA SystemMirror” on page 50

Configuring PowerHA SystemMirror in VM Recovery Manager HA solution

To create a PowerHA SystemMirror cluster in KSYS nodes, complete the following steps.

Note: Before starting the configuration procedure, ensure that the setup (hardware, VIOS, node, etc.) is ready and PowerHA SystemMirror and the KSYS subsystem is installed on your KSYS nodes.

1. Run the `phaksyssetup` setup script that is located at `/opt/IBM/ksys/samples/pha/`.
2. In the menu option of the script, select Standard Cluster (HA) for the KSYS node.
3. Specify the KSYS node names and the shared disk name for the repository.

After you have provided the required details, the sample script creates a Linked cluster.

The setup script creates the following PowerHA SystemMirror resources:

- Resource group: `ksysRG`
- Startup monitor: `ksysmonstartup`
- Startup monitor script: `/opt/IBM/ksys/samples/pha/startupmonksys`

- Long running monitor: `ksysmonlongrun`
- Long running monitor script: `/opt/IBM/ksys/samples/pha/longrunmonksys`
- Start script: `/opt/IBM/ksys/samples/pha/startksys`
- Stop script: `/opt/IBM/ksys/samples/pha/stopksys`
- File collection: `ksysfiles`

Configuring the KSYS subsystem for high availability through PowerHA SystemMirror

After the PowerHA SystemMirror cluster and PowerHA SystemMirror configuration is stable, you must add the KSYS resources to the KSYS configuration.

Before setting up the KSYS configuration, ensure that the following prerequisites are met:

- The `ksysmgr q cl` command should display the cluster name and type in HA KSYS nodes.
  If the `ksysmgr q cl` command does not display the correct cluster name and type, the `/var/ksys/config/ksysmgr.xml` file may have errors.
- On all KSYS nodes, IBM.VMR must be in an active state.
  To check the state of the IBM.VMR daemon, run `clcmd lssrc -s IBM.VMR` command from the PowerHA SystemMirror node.
- PowerHA SystemMirror resource group must be online on the group leader node.
  To view the group leader node, run the `lssrc -ls IBM.ConfigRM | grep -w Group Leader` command in the KSYS node.
  To view the group leader node's online status, run the `lssrc -ls IBM.ConfigRM | grep -w Group Leader` command in the KSYS node.

After the prerequisites are met, you can add KSYS resources such as sites, HMC, host, host group, and storage agents and perform regular operations.

Deleting the KSYS subsystem configuration and PowerHA SystemMirror cluster

To delete the KSYS subsystem configuration and PowerHA SystemMirror cluster, complete the following steps:

1. Set the PowerHA SystemMirror cluster to an offline state by running the `clmgr offline cl` command.
2. Start the peer domain by running the `startrpdomain` command.
   You can run the `lsrpdomain` command to view the status of the peer domain.
3. Delete all host groups from the KSYS configuration.
4. Delete the PowerHA SystemMirror cluster by running the `clmgr del cl` command.

Limitations

Consider the following limitations for the KSYS subsystem's high availability through PowerHA SystemMirror:

- To sync configuration settings between KSYS nodes, the IBM.VMR daemon must be active on both KSYS nodes.
- If a current group leader KSYS node fails, the next available KSYS node resumes control only after the IBM.VMR daemon is stopped and restarted.
- Only two KSYS nodes are supported in a PowerHA SystemMirror cluster for the KSYS subsystem's high availability.

Local database mode

You cannot update HA configuration (for example, disable/enable HA monitoring feature) when the VIOS is in local database mode.

Commands

This section provides information about all the commands used for VM Recovery Manager HA.

ksysmgr command

Purpose

The **ksysmgr** command provides a consistent interface to configure the controller system (KSYS) and to perform VM Recovery Manager HA operations. This command can be run from a terminal or a script.

Syntax

The basic format of the **ksysmgr** command follows:

```
ksysmgr [flags] ACTION CLASS [NAME] [ATTRIBUTES...]
```

For example: `ksysmgr -t -l max discover host_group Austin_hg verify=yes`

Alias usage

Alias is a shorthand definition for an operation and is defined by the most significant letters. The asterisk (*) in the aliases signify wildcard characters. For example, the alias value for the modify ACTION is `mod*`. If you type `modd`, the command still works. Aliases are provided for convenience from the command line and must not be used in scripts.

Log file

All **ksysmgr** command operations are logged in the `/var/ksys/log/ksysmgr.oplog` file, which includes the name of the command that was executed, start time, process ID for the **ksysmgr** operation, the command with arguments, and the overall return code. In addition, the `/var/ksys/log/ksysmgr.log` file tracks the internal activities of the **ksysmgr** command. The amount of information that is written to the `ksysmgr.log` file can be modified for each command by using the **-l** flag.

Notes:

- You must have root authority to run the **ksysmgr** command.
- Help information is available for the **ksysmgr** command from the command line. For example, when you run the **ksysmgr** command without any flags or parameters, a list of the available ACTIONs is displayed. If you enter **ksysmgr ACTION** in the command line without specifying any CLASS, the command results in a list of all the available CLASSes for the specified ACTION. Entering **ksysmgr ACTION CLASS** without specifying any NAME or ATTRIBUTES parameters might yield different results because some ACTION and CLASS combinations do not require any additional parameters. To display help information in this scenario, you can view the help information by appending the **-h** flag to the **ksysmgr ACTION CLASS** command.
- You cannot display help information from the command line for each ATTRIBUTE of the **ksysmgr** command.

Flags

You can use the following flags with the **ksysmgr** command:

ACTION
Describes the action to be performed. The ACTION flags are not case-sensitive. All ACTION flags provide a shorter alias. The following ACTION flags are available:

- add (alias: ad*, create, cr*, make, mk, bu*, bld)
- cleanup (alias: clean*)
- delete (alias: de*, remov*, rm, er*)
- discover (alias: di*)
- help (alias: hel*, ?)
- lpm (alias: none)
- manage (alias: man*, mg)
- modify (alias: mod*, ch*, set)
- query (alias: q*, ls, get, sh*)
- recover (alias: rec*)
- report (alias: rep*)
- restart (alias: resta*)
- restore (alias: rest*)
- sync (alias: syn*, pr*)
- unmanage (alias: unman*, umg)
- verify (alias: ver*)
- start (alias: none)
- stop (alias: none)
- refresh (alias: ref)

CLASS

Specifies the type of object on which the ACTION is performed. The CLASS flags are not case-sensitive. The following CLASS flags are supported:

- anti_collocation (alias: anti*)
- collocation (alias: collo*)
- event (alias: ev*)
- hmc (alias: hmcs, hmces)
- host (alias: serv*, mach*, cec*)
- host_group (alias: serv*, mach*, cec*, ho*)
- ksyscluster (alias: cl*, ksyscl*)
- notify (alias: rn, remote_not*, noti*)
- script (alias: scr*)
- snapshot (alias: snap*)
- system (alias: sys*)
- version (alias: vers*)
- viodisk (alias: viod*)
- vios (alias: vios*)
- vm (alias: lp*, vm*)
- workgroup (alias: workg*)
- app_dependency (alias: app_dep*)

NAME

Specifies a particular object, of type CLASS, on which the ACTION must be performed. The NAME flags are case-sensitive.

ATTRIBUTE=VALUE
Specifies an optional flag that has attribute pairs and value pairs that are specific to the ACTION and CLASS combination. Use these pairs to specify configuration settings or to run particular operations. Both the ATTRIBUTE and VALUE flags are case-sensitive.

-a {<ATTR#1>,<ATTR#2>,...}
Displays only the specified attributes. This flag must be used with the query ACTION flag. For example: `ksysmgr -a name,sitetype query site`.

-f
Overrides any interactive prompts and forces the current operation to be run.

-h
Displays help information.

-l low|med|high|max
Activates the following trace logging values:

low
Logs basic information for every **ksysmgr** operation. This is the default value of the **-l** flag.

med
Logs warning messages also.

high
Logs basic information messages also that can be used for demonstrations.

max
Performs high tracing operations such as adding the routine function and the utility function. It also adds a transaction ID to the entry messages of each function.

All trace data is written in the `ksysmgr.log` file. This flag is ideal for troubleshooting problems.

-v
Displays maximum verbosity in the output.

-i
Skips the interactive prompts from the ksysmgr command.

Comprehensive list of ksysmgr operations

Use the following information to find the syntax for all possible **ksysmgr** operations:

KSYS cluster configuration

- To create a KSYS cluster:

```
ksysmgr add kyscluster name ksysnodes=ksysnode1 type=HA
```

- To verify the KSYS cluster:

```
ksysmgr verify kyscluster name
```

- To synchronize the KSYS cluster:

```
ksysmgr sync kyscluster name
```

- To create, verify, and synchronize a KSYS cluster:

```
ksysmgr add kyscluster name ksysnodes=ksysnode1 type=HA sync=yes
```

- To query a KSYS cluster:

```
ksysmgr query kyscluster [name]
```

- To delete a KSYS cluster:

```
ksysmgr delete ksyscluster name
```

Note: When you delete the KSYS cluster, the virtual machine (VM) agent daemon becomes inoperative. In such cases, you must manually start the VM agent daemon.

HMC management

- To add an HMC:

```
ksysmgr add hmc hmcname    \
      login=username    \
      password=password    \
      hostname|ip=hostname|ip
```

- To modify the HMC details:

```
ksysmgr modify hmc hmcname    \
      [name=new_hmcname]    \
      [login=new_username]    \
      [password=new_password]    \
      [hostname|ip=hostname|ip]
```

- To query an HMC:

```
ksysmgr query hmc [hmcname]
```

- To delete an HMC:

```
ksysmgr delete hmc hmcname
```

- To sync updated HMC information:

```
ksysmgr refresh hmc [<hmcname>|<ALL>]
```

Host management

- To add a host:

```
ksysmgr add host hostname    \
      [uuid=uuid]    \
      [hostname|ip=hostname|ip]
```

- To modify the host details:

```
ksysmgr modify host hostname|uuid[,hostname2|uuid2,...]    \
      [name=new_hostname]    \
      [uuid=new_uuid]    \
      [hostname|ip=newhostname|ip]
```

- To query hosts:

```
ksysmgr query host [hostname]
```

- To delete a host:

```
ksysmgr delete host hostname
```

  You must first delete the host group to which the host belongs before deleting a host.

- To sync updated host information:

```
ksysmgr refresh host [<hostname>|<ALL>]
```

Host group configuration

- To add a host group:

```
ksysmgr add host_group hgname hosts= host1[,host2,…]
      [repo_disk=diskuuid] [ha_disk=diskuuid] [backup_repo_disk=diskuuid1, diskuuid2…]
```

  You can use the **ksysmgr query viodisk** command to identify the free disks that can be used in this step.

- To modify host group details:

```
ksysmgr modify host_group <name> <add | remove | options> [<hosts=<host1>,<host2>...>|
file=<filepath>] [memory_capacity=<(1-100) | minimum | current_desired | none> priority=<low
| medium | high>] [cpu_capacity=<(1-100) | minimum | current_desired | none> priority=<low |
medium | high>] [skip_power_on=<yes|no>] [ha_disk=<ViodiskID>] [repo_disk=<ViodiskID>]
[backup_repo_disk= none <ViodiskID[,ViodiskID2...]> | backup_repo_disk=none]
[ha_monitor=<enable | disable>] [proactiveha=<enable | disable>] [restart_policy=<auto |
advisory_mode>] [vm_failure_detection_speed=<fast | normal | slow>]
[host_failure_detection_time=<90-600>] modify => mod*, ch*, set host_group => hg, host_g*
```

- To modify the repository disk and the HA disk that are associated with a host group:

```
ksysmgr modify host_group hgname options
      [repo_disk=diskuuid] [ha_disk=diskuuid][backup_repo_disk=diskuuid1,diskuuid2…]
```

  You cannot modify the HA disk after the discovery operation, because the KSYS subsystem does not support the HA disk modification after the SSP cluster is created.

- To discover a host group:

```
ksysmgr discover host_group hg_name
```

- To verify a host group:

```
ksysmgr verify host_group hg_name
```

- To discovery and verify a host group:

```
ksysmgr discover host_group hg_name verify=true
```

- To query a specific host group or all host groups:

```
ksysmgr query host_group [hg_name]
```

- To clean up all virtual machines from a host group:

```
ksysmgr cleanup host_group hg_name
```

- To delete a host group:

```
ksysmgr delete host_group hg_name
```

Workgroup

- >| To query a workgroup:

```
ksysmgr query workgroup [name]
```

  |<

- >| To delete a workgroup:

```
ksysmgr delete workgroup name
```

|<

Application dependency between virtual machines

- To establish dependency between applications of a virtual machine within a host group, run the following command:

```
ksysmgr add app_dependency <name>
      app_list=<vmname1:appname1,vmname2:appname2[,...]>
      type=<parent_child|primary_secondary>
      [mode=<sync|async>]
    add => ad*, cr*, make, mk
    app_dependency => app_dep*
```

Note:

  - >|The **app_list** attribute must have only two *vmname:appname* pairs for the primary-secondary structure of applications across VMs.|<
  - >|By default, the mode is sync for the primary-secondary type of application dependency. |<

- >|To query application dependency:

```
ksysmgr query app_dependency
```

|<

- >|To start an application:

```
ksysmgr start app [name=<appname1,[appname2,...]>] vm=<vmname>
    vm => lp*, vm
```

Note:

  - You can enter multiple application names as value of the **name** attribute. But, all applications must belong to the single virtual machine.
  - If you do not provide any application name, all applications in the virtual machine start.

|<

- To stop an application:

```
ksysmgr stop app [name=<appname1,[appname2,...]>] vm=<vmname>
    vm => lp*, vm
```

Note:

  - You can enter multiple application names as value of the **name** attribute. But, all applications must belong to the single virtual machine.
  - If you do not provide any application name, all applications in the virtual machine stop.

HA monitoring policies

- To enable or disable HA management for the entire host group:

```
ksysmgr modify host_group|host|vm name
    ha_monitor=enable|disable
```

- To specify the time that KSYS waits on a non-responsive host before declaring the host to be in an inactive state:

```
ksysmgr modify system|host_group name
   host_failure_detection_time=time_in_seconds
```

The value of this attribute can be in the range 90 - 600 seconds. The default value is 90 seconds.

- To specify the time that KSYS will wait before declaring the failure of a VM:

```
ksysmgr modify system|host_group|host|vm name
vm_failure_detection_speed=fast|normal|slow
```

  You can select one of the following options: `fast` (140 seconds), `normal` (190 seconds), or `slow` (240 seconds). The default value of the **vm_failure_detection_speed** attribute is `normal`.

- To specify whether the virtual machines must be restarted automatically after a failure:

```
ksysmgr modify host_group name
restart_policy=auto|advisory_mode
```

  The default value of the **restart_policy** attribute is `auto`. If you set the attribute to `advisory_mode`, the virtual machines are not restarted automatically during failures. In this case, the KSYS subsystem notifies the registered contacts about the failures and the administrator must review the failure and manually restart the VMs on other hosts by using the **ksysmgr** commands.

- To modify the allocation of memory and CPU resources to a virtual machine when a virtual machine is moved from a host to another host:

```
ksysmgr modify host_group
     [memory_capacity=(1-100) | minimum | current_desired | none] [priority=low|medium|high]
     [cpu_capacity=(1-100) | minimum | current_desired | none] [priority=low|medium|high]
```

  Capacity adjustment occurs only when the VM is moving to a host that is not its home-host. Also, capacity adjustment is not possible when VM is moved to another host by using the Live Partition Mobility (LPM) operation. Capacity value is accepted in percentage. For example, a value of 89 or 890 means 89.0% of the original capacity must be deployed on the backup host of the host group when a VM is relocated.

Affinity policies

- To specify that the set of VMs must always be placed on the same host after relocation:

```
ksysmgr add collocation name vm=vm1[,..]>
ksysmgr modify collocation name policy=add|delete vm=vm1[,..]>
```

- To specify that the set of VMs must not be placed on the same host after relocation:

```
ksysmgr add anticollocation name vm=vm1[,...]>
ksysmgr modify anticollocation name policy=add|delete vm=vm1[,..]>
```

- To prioritize the set of VMs within the assigned priority:

```
ksysmgr add workgroup name vm=vm1[,...]>
ksysmgr modify workgroup name policy=add|delete vm=vm1[,...]>
```

- To specify a list of hosts that must not be used for relocating a specific virtual machine during a failover operation:

```
ksysmgr modify vm name blacklist_hosts=host1[,..] [policy=add|delete]
```

Virtual machines HA policies

- To include or exclude specific VMs from the KSYS configuration settings:

```
ksysmgr manage|unmanage vm vmname|vmuuid
```

  Note: When you stop managing a virtual machine, the VM agent daemon becomes inoperative. In such cases, you must manually start the VM agent daemon.

- To set the priority of a virtual machine or to specify the order of virtual machines for a specific operation:

```
ksysmgr modify vm name1[,name2,…] | filepath=filepath
priority=high|medium|low
```

where, the **filepath** parameter is an XML file that contains a list of virtual machines. An example of the XML file follows:

```
<?xml version="1.0"?>
<KSYSMGR><VM><NAME>VM1</NAME></VM>
<VM><NAME>VM2</NAME></VM>
<VM><NAME>VM3</NAME></VM>
</KSYSMGR>
```

For example, when you relocate all the VMs to another host, the priority of the VMs determines which VMs must be processed first.

- To set the home-host of a virtual machine:

```
ksysmgr modify vm name1[,name2...] homehost=hostname
```

- To query a specific or all virtual machines:

```
ksysmgr query VM [VM_name] [display_all=yes]
```

When you specify the **display_all** attribute, the output of this command includes the virtual machines of hosts that are not added to the KSYS subsystem but are registered in the added HMCs.

The `State` field in the output is a significant field that indicates the current status of the VM. If the `HA_monitor` option is enabled, the `Heartbeating_with` field indicates the VIOS name to which the VM is sending heartbeats. If the VM is not associated with any VIOS, this field will be not be available.

- To manually clean up a specific virtual machine after the move operations:

```
ksysmgr cleanup vm vmname host=hostname
```

- To restore a virtual machine:

```
ksysmgr restore vm <vmname>
```

- To query collocation:

```
ksysmgr query collocation [<name>]
```

- To query anticollocation:

```
ksysmgr query anticollocatoin [<name>]
```

VIOS management

- To query a specific VIOS or all virtual I/O servers:

```
ksysmgr query vios [vios_name] [display_all=yes]
```

When you specify the **display_all** attribute, the output of this command includes the VIOS of hosts that are not added to the KSYS subsystem but are registered in the added HMCs.

- To include or exclude a specific VIOS from the HA management:

```
ksysmgr manage|unmanage vios viosname
```

- To modify the VIOS logging level for specific components:

```
ksysmgr modify vios viosname[,viosname2,... | file=filepath
      [hsmon_daemon=0|1|2|3|4|5]
      [hm_api=0|1|2|3]
      [health_api=0|1|2|3]
```

Notification contacts

- To register contact details for notification from the KSYS:

```
ksysmgr add notify user=username contact=email_address
ksysmgr add notify user=username      \
       contact=10_digit_phone_number@phone_carrier_email_add
ksysmgr add notify user=username contact=pager_email_address
```

If you chose the **advisory_mode** option in the restart policy, you must register an email address to get notifications about the failures.

- To modify the contact details for notification from the KSYS:

```
ksysmgr modify notify oldcontact=old_username newcontact=new_username
ksysmgr modify notify oldcontact=old_email_address newcontact=new_email_address
```

- To query existing contact details:

```
ksysmgr query notify contact
```

- To delete a specific contact:

```
ksysmgr delete notify user=username
```

Event notification script management

- To register notification scripts for specific events and errors:

```
ksysmgr add notify script=full_path_script events=event_name
```

You can add a maximum number of 10 event notification scripts to the KSYS configuration settings.

- To modify a specific notification script:

```
ksysmgr modify notify oldscript=old_file_name newscript=new_file_name
```

- To query notification scripts:

```
ksysmgr query notify script
```

- To delete a specific notification script:

```
ksysmgr delete notify script=file_name
```

Script management

- To register scripts for specific KSYS operations such as discovery and verification:

```
ksysmgr add script entity=host_group
  pre_offline|post_online|pre_verify|post_verify=script_file_path
```

- To query the existing user scripts:

```
ksysmgr query script entity=host_group
```

- To delete an existing user script:

```
ksysmgr delete script entity=host_group
      [script_name=pre_offline|post_online|pre_verify|post_verify]
```

System-wide attributes

- To query system status:

```
ksysmgr query system
```

- To modify the system attributes:

```
ksysmgr modify system attribute=new_value
```

- To stop health monitoring for high-availability in the entire configuration:

```
ksysmgr modify system ha_monitor=disable
```

Configuration snapshots

- To create a backup of the KSYS environment:

```
ksysmgr add snapshot filepath=full_file_prefix_path
    type=CLUSTER|BASIC|DETAILED
```

- To restore the backed up configuration snapshot:

```
ksysmgr restore snapshot filepath=full_file_prefix_path
```

- To query snapshots:

```
ksysmgr query snapshot filepath=full_file_prefix_path
```

You cannot remove a snapshot by using the **ksysmgr** command. You must remove the snapshot manually by using standard operating system commands.

Cleanup operations

- To clean up a specific virtual machine in a specific host:

```
ksysmgr cleanup vm vmname host=hostname
```

Migration operations (LPM)

- To migrate a single VM or a set of VMs by using the Live Partition Mobility (LPM) operation:

```
ksysmgr [-f] lpm vm vm1[,vm2,..] [to=hostname|uuid]
```

- To validate an LPM operation without migrating virtual machines:

```
ksysmgr [-f] lpm vm vm1[,vm2,..] action=validate
```

- To migrate all VMs in a host:

```
ksysmgr [-f] lpm host hostname|uuid [to=hostname|uuid]
```

- To restore all VMs that belong to a specific host:

```
ksysmgr restore host hostname|uuid
```

Manual restart operations

- To relocate and restart a single VM or a set of VMs:

```
ksysmgr [-f] restart vm vm1[,vm2,..] [to=hostname|uuid]
```

- To relocate and restart all VMs in a host:

```
ksysmgr [-f] restart host hostname|uuid [to=hostname|uuid]
```

The manual restart commands must be used to relocate the VMs if you have set the **advisory_mode** option for the restart policy. You cannot restart a VM on the same host in which the VM is located by using the **ksysmgr restart** command. To restart a VM on the same host, use the restart option in the HMC. If the restart operations fail, you can attempt to recover the virtual machine in the same host where it is located.

- To recover a VM in the same host where the VM is located:

```
ksysmgr [-f] recover vm vmname
```

General queries

- To check the version of the KSYS software:

```
ksysmgr query version
```

- To query the system-wide configuration details:

```
ksysmgr query system
```

- To list the various types of events that are generated by the KSYS subsystem:

```
ksysmgr query event
```

- To list the free disks that are attached to or shared among various VIOS:

```
ksysmgr query viodisk vios=viosname
```

- To check the current activities and general health of the environment:

```
ksysmgr query system status
```

This command displays error messages as a verbose output.

- To monitor the activities that are running currently:

```
ksysmgr q system status monitor=yes
```

- To display the health status of the registered applications:

```
ksysmgr query app
```

You can use this command to check the status of the registered applications inside each virtual machine. The application status value GREEN indicates a stable application. The YELLOW value

indicates an intermediate state of application in which the application is attempted to be restarted. The RED value indicates permanent failure of an application.

The HA policies for an application can be configured in the VM only by using the **ksysvmmgr** command.

- To review the relocation plan for a VM or a host:

```
ksysmgr report system relocation_plan [vm=vm_list | host=host_list]
```

You can also use UUID of VMs or hosts instead of names.

- To print trace information from KSYS trace files to standard output:

```
ksysmgr trace logs=ksys|fde|fdelong|krest|krestlong|user|ALL
```

Examples

The following examples show a scenario where you deploy a host group for HA management, enable HA policies for the host group, and perform planned and unplanned relocation operations for the virtual machines, if required:

1. Identify the servers or hosts that will be part of the host group. Complete all the prerequisites that are specified in the Requirements topic.
2. Create, verify, and synchronize a KSYS cluster by running the following command:

```
ksysmgr add ksyscluster name
    ksysnodes=ksysnode1 type=HA sync=yes
```

3. Register all the HMCs in the environment by running the following command:

```
ksysmgr add hmc HMC1 ip=x.x.x.x login=username password=password
```

4. Create a host group by running the following command:

```
ksysmgr add host_group HG1 hosts=hostA[,hostB,...]
```

5. Identify the free disks that you can designate as repository disk and HA disk for the SSP cluster by running the following command:

```
ksysmgr query viodisk host_group=HG1
```

6. Modify the host group to add repository disk and HA disk by running the following command:

```
ksysmgr modify host_group HG1 options
     repo_disk=Disk1uuid ha_disk=Disk2uuid
```

7. Configure the failure detection time and other policies according to your requirements by running the following commands:

```
ksysmgr modify host_group HG1
  vm_failure_detection_speed=fast
  host_failure_detection_time=100
  restart_policy=advisory_mode
```

8. Configure an email address for event notifications by running the following command:

```
ksysmgr add notify user=John contact=john.doe@testmail.com
```

9. Enable HA monitoring for the host group by running the following command:

```
ksysmgr modify host_group HG1 options ha_monitor=enable
```

10. Discover and verify the added resources and policies by running the following command:

```
ksysmgr discover host_group HG1 verify=yes
```

11. If you want to upgrade a host or if you plan a host maintenance, migrate all the virtual machines to another host by running the following command:

```
ksysmgr lpm host hostA|uuid
```

12. After the upgrade operation or maintenance activity of the host is complete, restore the virtual machines back to the host by running the following command:

```
ksysmgr restore host hostA|uuid
```

13. If a host or a virtual machine fails, restart the virtual machines on another host by running the following command:

```
ksysmgr restart vm vm1[,vm2,...]
```

ksysvmmgr command

Purpose

The **ksysvmmgr** command provides a consistent interface for the VM agent to manage the virtual machine monitor (VMM) and the applications that are running in the virtual machine.

Syntax

The **ksysvmmgr** command uses the following basic format:

```
ksysvmmgr [flags] ACTION CLASS [NAME] [ATTRIBUTES...]
```

Description

The VM agent consists of the following subsystems:

VM monitor
This subsystem tracks the health of the VM by communicating periodically with the host monitor in the VIOS.

Application monitoring framework
This subsystem is an application monitoring framework that registers, starts, stops, and monitors the applications.

You can access the VM monitor and the application monitoring framework only through the command line by using the **ksysvmmgr** command. Graphical user interface is not available for such operations.

Log file

All **ksysvmmgr** command operations are logged in the `/var/ksys/log/ksysvmmgr.log` file, including the name of the command that was executed, the start and stop time of the command, and the user name of the user who initiated the command. You can use the **-l** flag to change the amount of information that is written to the log files. The **ksysvmmgr** command sends user messages that are transformed by using the catalog messages.

Flags

You can use the following flags with the **ksysvmmgr** command:

ACTION
Describes the action to be performed.

The ACTION flags are not case-sensitive. All ACTION flags provide synonyms, and each alias and synonym have shorter alias. For example, remove is an alias for delete, and remove can be abbreviated with rm. Aliases are provided for convenience from the command line and must not be used in scripts. The following ACTION flags are available:

- help (alias: h*)
- add (alias: a*, register, reg*)
- query (aliases: q*, list, get, li*, g*)
- modify (aliases: mod*, change, set, ch* se*)
- delete (aliases: unregister, remove, de*, unr* re*, rm)
- sync (alias: syn*)
- start (alias: on, enable, star*, en*)
- stop (alias: off, disable, sto*, di*)
- backup (alias: bac*)
- restore (alias: rest*)
- suspend (alias: unmonitor, sus*, unm*)
- resume (alias: monitor, resu*, mon*)
- snap (alias: snap, sna*)

CLASS
Specifies the type of object on which the ACTION is performed. The following CLASS objects are supported:

- vmm (alias: v*): The **vmm** CLASS is used by default. If you do not specify any CLASS, the action is performed on the VM monitor by default.
- app (alias: a*): The **app** CLASS is used to perform ACTION on applications. You must specify a NAME attribute to apply the ACTION on a specific application.
- dependency (alias:dep*): The dependency CLASS establishes a dependency relationship between the applications. You must specify the list of applications and dependency type to apply a dependency between applications.
- process (alias:none): The process class is used to perform ACTION on processes. You must specify a NAME attribute to apply the ACTION on a specific process.

NAME
Specifies the particular object, of type CLASS, on which the ACTION must be performed. The NAME flags are case-sensitive. You can use this flag only for the app CLASS.

ATTRIBUTES
Specifies an optional flag that has attribute pairs and value pairs that are specific to the ACTION and CLASS combination. Use these pairs to specify configuration settings or to run particular operations.

Both ATTRIBUTES and VALUE flags are case-sensitive. You cannot use the asterisk (*) character in the ACTION and CLASS names.

-a
Displays only the specified attributes. This flag is valid only with the query ACTION. Attribute names are not case-sensitive.

-f
Overrides any interactive prompts, forcing the current operation to be attempted, if allowed.

-h/-?
Displays help information.

-l 0|1|2|3
Activates the following trace log values for serviceability:

- 0: Updates the log file when an error is detected. This level of trace logging is the default value.
- 1: Logs warning messages also.
- 2: Logs basic information messages also that can be used for demonstrations.
- 3: Performs high tracing by logging the details about the routine function and the utility function. Traces the entry and exits for various functions.

All trace data is written into the `vmmgr.log` file. This flag is used for troubleshooting problems.

-s
Synchronizes the VM monitor daemon immediately by sending a notification to the daemon. The VM monitor daemon reloads the XML configuration when it receives this notification.

This flag is valid only with the `add`, `modify`, `delete`, `resume`, and `suspend` ACTIONS. By default, no notification is sent. To send a notification, use the **`ksysvmmgr sync`** command.

Attributes

The **`ksysvmmgr`** command configures the following classes and attributes:

vmm
When you start the VM monitor, the VM monitor daemon sends heartbeats to the host monitor, when requested by the host monitor, so that the KSYS subsystem can monitor the virtual machines. The VM monitor can have the following attributes:

version
Specifies the version of XML. This mandatory attribute is set to 1.0 for the current version of VM monitor and cannot be modified.

log
Specifies the log level of the VM monitor daemon. This attribute can have the following values:

- 0: Only errors are logged. It is the default value.
- 1: Warnings are also logged.
- 2: Informational messages are also logged.
- 3: Details of the operation are logged. This information is used for debugging.

period
Specifies the time duration in seconds between two consecutive occurrences of checks that are performed by the Application Management Engine (AME). By default, the value of this attribute is 1 second. The value of this attribute must be in the range 0 - 6. For best monitoring performance, do not modify the default value.

Application (app)
The application class contains the following mandatory attributes:

monitor_script
A mandatory script that is used by the VM agent to verify application health. This script is run regularly (based on the **`monitor_period`** attribute value) and the result is checked for the following values:

- 0: Application is working correctly.
- Any value other than 0: Application is not working correctly or has failed.

After several successive failures (based on the **`monitor_failure_threshold`** attribute value), the application is declared as failed. Based on the specified policies, the KSYS subsystem determines whether to restart the virtual machine.

stop_script
A mandatory script that is used by the VM agent to stop the application if the application must be restarted. The application can be restarted by successively calling the **stop_script** and **start_script** scripts.

start_script
A mandatory script that is used by the VM agent to start the application if the application must be restarted.

Exception: These scripts are not mandatory for the following application types: ORACLE, DB2, and SAPHANA.

The application class contains the following optional attributes:

monitored
Specifies whether the application is monitored by the KSYS subsystem. This attribute can have the following values:

- 1 (default): The application monitoring is active.
- 0: The application monitoring is suspended.

monitor_period
Specifies the time in seconds after which the application monitoring must occur. The default value of 30 seconds specifies that the **monitor_script** script is run by the VM agent every 30 seconds.

monitor_timeout
Specifies the waiting time in seconds to receive a response from the **monitor_script** script. The default value is 10 seconds, which means that the VM monitor waits for 10 seconds to receive a response from the **monitor_script** script after which the script is considered as failed.

monitor_failure_threshold
Specifies the number of successive failures of the **monitor_script** script that is necessary before the VM monitor restarts the application. A restart operation is performed by successively calling the **stop_script** and **start_script** scripts.

stop_stabilization_time
Specifies the waiting time in seconds to receive a response from the **stop_script** script. The default value is 25 seconds, which means that the VM monitor waits for 25 seconds to receive a response from the **stop_script** script after which the script is considered as failed.

stop_max_failures
Specifies the number of successive failures of the **stop_script** script that is necessary before the VM monitor considers that it cannot stop the application. The default value is set to 3.

start_stabilization_time
Specifies the waiting time in seconds to receive a response from the **start_script** script. The default value is 25 seconds, which means the VM monitor waits for 25 seconds to receive a response from the **start_script** script after which the script is considered as failed.

start_max_failures
Specifies the number of successive failures of the **start_script** script that is necessary before the VM monitor considers that it cannot start the application. The default value is set to 3.

max_restart
Specifies the number of cycles of successive VM restart operations that result in a monitoring failure before the daemon pronounces that restarting at VM level is insufficient. By default, this attribute is set to 3.

status
Specifies the dynamic status of application that is returned by Application Management Engine (AME). This attribute cannot be modified.

The color codes signifies the dynamic status of an application.

Table 5. Color code for dynamic application status	
Color code	**Application status**
Red	Permanent application failure
Orange	Initial application state
Yellow	Application is in normal state but failed more than two times within last 24 hours
Green	Application is in normal state
Gray	Application is not monitored by the VMM daemon
Blue	Application is in other intermittent states. For example, starting, stopping, or failing state.

version
Specifies the application version. This attribute does not have a default value.

critical
Marks the application as critical. The valid values are `Yes` and `No` (default). If you mark an application as critical, failure of the application may lead the VM to be rebooted or relocated by the KSYS subsystem.

type
Specifies the type of application. By default, the **type** attribute value is `CUSTOM` that indicates general applications. Other supported values are `ORACLE`, `DB2`, `POSTGRES` and `SAPHANA`. This attribute is case-sensitive and you must use uppercase characters. For these types of applications, if you do not specify start, stop, and monitor scripts, the internal scripts of the VM monitor are used.

instancename
Specifies the instance name for applications. This attribute is applicable only for agent applications, which are internally supported by the VM agent. The supported agent applications are: oracle, DB2, SAPHANA and POSTGRES. For example,

- If the application type is `ORACLE`, the **instancename** attribute must be specified with the Oracle user name.
- If the application type is `DB2`, the **instancename** attribute must be specified with the DB2 instance owner.
- If the application type is SAPHANA, the **instancename** attribute must be specified with the SAPHANA system id.
- If the application type is POSTGRES, the **instancename** attribute must be specified with the POSTGRES instance id.

database
Specifies the database that the applications must use. This attribute is applicable only for agent applications, which are internally supported by the VM agent. The supported agent applications are: oracle, DB2, SAPHANA and POSTGRES.For example,

- If the application type is `ORACLE`, the **database** attribute must be specified with the Oracle system identifier (SID).
- If the application type is `DB2`, the **database** attribute is not required.
- If the application type is SAPHANA, the **database** attribute must be specified with the SAP HANA database.
- If the application type is POSTGRES, the **database** attribute can be specified with the database name. If the database name is not specified, the script monitors all database of the POSTGRES instance.

appstarttype

Specifies the method in which the applications must be started. This attribute can have the following values:

- VMM: Specifies that the VM agent must start and monitor the applications.
- 0S: Specifies that the application must be started by the operating system or by a user.
- KSYS: Specifies that the application must be started or stopped by the KSYS subsystem. After the application is started by the KSYS subsystem, and eventually the application crashes, the VMM must restart the application.

From the ksysvmmgr command-line-interface (CLI), you can modify the **appstarttype** attribute of an application.

Further, if a VM daemon reboots, the VMM daemon starts all the VMM controlled applications, but the VMM daemon cannot start the KSYS controlled applications. Instead, the VMM daemon sends the status of the application to the KSYS subsystem. The KSYS subsystem determines whether to start or stop the KSYS controlled applications. The KSYS subsystem has privilege to modify the value of the **appstarttype** attribute of an application from KSYS to VMM, or vice versa.

The default value of the appstarttype attribute of an application is VMM.

Note: If the **appstarttype** attribute is modified from KSYS to VMM, then you must manually delete all the related application dependencies in the KSYS subsystem.

groupname

Specifies the **groupname** to which the application belongs. The default value of the **groupname** attribute is NULL. After a user creates a new group, the **groupname** of each application in the group is updated.

configfile

This file contains the application configurations settings for the supported agent applications . This attribute is used only by SAP HANA and POSTGRES agent applications. This attribute is blank for other agent applications.

Application dependency (dependency)

The **dependency** class contains the following mandatory attributes:

dependency_type

Specifies the type of dependency between applications. This attribute can have the following values:

- start_sequence: Specifies the order in which the applications must be started as mentioned in the **dependency_list** attribute. The **dependency_list** attribute must have more than one application for this dependency type.
- stop_sequence: Specifies the order in which the applications must be stopped as mentioned in the **dependency_list** attribute. The **dependency_list** attribute must have more than one application for this dependency type.
- parent_child: Specifies the parent-child relationship of the two specified applications in which one application is parent and the other is child. The parent application must start first and then the child application starts. You must stop the child application first and then stop the parent application. If the parent application fails, the child application also stops automatically. If the parent application recovers and starts, the child application started automatically.

dependency_list

Specifies the list of applications that have a dependency between them.

The **dependency** class also contains the following optional attributes:

strict

Specifies whether to continue the script or command if the dependency policy cannot be followed. If the **strict** attribute is set to Yes, the next application is not started until the previous application starts and is in the normal state. If the **strict** attribute is set to No, the next

application is started immediately after the first application is started irrespective of the state of the first application. This attribute is applicable only for the **start_sequence** dependency.

Process

The **process** class can have the following attributes:

start_script

This script is used by the process monitor to restart a process in a virtual machine. The restart operation is performed by successively calling the `stop_script` and the `start_script`.

stop_script

This script is used by the process monitor to stop an application.

monitor_period

Duration between two successive process monitor operations. This duration is displayed in seconds. The default values is 30 seconds.

stop_stabilization_time

The period of time before which no answer from the `stop_script` script is considered as timeout. The default values is 25 seconds.

stop_max_failures

The number of successive failures of the `stop_script` script after which it is considered that the processcannot be stopped. The default values is 3 successive failures.

start_stabilization_time

The duration of time before which no answer from the `start_script` script is considered as timeout. This duration is displayed in seconds. The default value is 25 seconds.

start_max_failures

The number of successive failures of the `start_script` script after which it is considered that the process cannot be started. The default values is 3 successive failures.

max_restart

The maximum number of restart cycles, after which the process monitor is considered as failed. The default value is 3.

Comprehensive list of ksysvmmgr commands

Use the following information to find the syntax for all possible **ksysvmmgr** operations:

VM monitor operations

- To display help information about the `vmm` class, run the following command:

```
ksysvmmgr -h vmm
```

- To start the VM monitor daemon, run the following command:

```
ksysvmmgr start [vmm] [<ATTR#1>=<VALUE#1>][,<ATTR#n>=<VALUE#n>]
```

- To stop the VM monitor daemon, run the following command:

```
ksysvmmgr stop [vmm]
```

- To query the details about the VM monitor daemon, run the following command:

```
ksysvmmgr [-l {0|1|2|3}] [-a <ATTR#1>[,<ATTR#2>,...]] query [vmm]
```

- To modify the VM monitor daemon attributes, run the following command:

```
ksysvmmgr [-s] modify [vmm] <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
```

- To back up the VM monitor configuration settings, run the following command:

```
ksysvmmgr backup vmm [<ATTR#1>=<VALUE#1>]
```

- To restore the VM monitor configuration settings, run the following command:

```
ksysvmmgr restore vmm <ATTR#1>=<VALUE#1>
```

- To notify the VM monitor daemon to synchronize with the contents of the XML configuration file, run the following command:

```
ksysvmmgr sync [vmm]
```

- To compress all log files, run the following command:

```
ksysvmmgr snap [vmm]
```

  This command creates a `*.pax.gz` file in the `/var/ksys/log/snap` directory. To read the contents of the file, you can unzip the file by using the following commands:

```
unzip *.pax.gz
pax -r -f *.pax
```

Application operations

- To display help information about the app class, run the following command:

```
ksysvmmgr -h app
```

- To add an application that must be monitored, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] add app <NAME> <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
```

- To add a critical application that must be monitored, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] add app <NAME> critical=yes
```

- To modify the attributes of a registered application, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] modify app <NAME> <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
```

- To query the details about a registered application, run the following command:

```
ksysvmmgr [-l {0|1|2|3}] [-a <ATTR#1>[,<ATTR#2>,...]] query app <NAME>
```

- To delete specific or all applications from the VM agent configuration settings, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] delete app <NAME>|<NAME1,NAME2,NAME3...>|<ALL>
```

- To suspend the monitoring of an application or all applications, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] suspend app [<NAME>]
```

- To resume the monitoring of an application or all applications, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] resume app [<NAME>]
```

- To start an application, run the following command:

```
ksysvmmgr start app <NAME>
```

- To stop an application, run the following command:

```
ksysvmmgr stop app <NAME>
```

- To group the applications, run the following command:

```
ksysvmmgr [-s] add group <NAME> <ATTR#1>=<VALUE#1>
```

- To modify the `applist` in a group, run the following command:

```
ksysvmmgr [-s] modify group <NAME> applist=<VALUE#1>
```

- To ungroup (delete a group), run the following command:

```
ksysvmmgr [-s] delete group <NAME>
```

Dependency operations

- To display help information about the dependency class, run the following command:

```
ksysvmmgr -h dep
```

- To add a dependency relationship between applications, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] add dependency <ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
```

- To query the details about an existing dependency between applications, run the following command:

```
ksysvmmgr [-l {0|1|2|3}] [-a <ATTR#1>[,<ATTR#2>,...]] query dep<DEPUUID>
```

- To modify a specific dependency relationship between applications, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] modify dependency <NAME>
<ATTR#1>=<VALUE#1>[,<ATTR#n>=<VALUE#n>]
```

- To delete a specific dependency relationship, run the following command:

```
ksysvmmgr [-s] [-l {0|1|2|3}] delete dep DEPUUID
```

Process operations

- To add a process for monitoring, run the following command:

```
ksysvmmgr [-s] add process process_name=<NAME> <ATTR#1>=<VALUE#1> [,<ATTR#n>=<VALUE#n>]
```

- To query a process, run the following command:

```
ksysvmmgr query process [process_name=<NAME>]
```

- To modify a process attribute, run the following command:

```
Ksysvmmgr [-s] modify process process_name=<NAME> <ATTR#1>=<VALUE#1> [,<ATTR#n>=<VALUE#n>]
```

- To suspend a process, run the following command:

```
ksysvmmgr suspend process process_name=<NAME>
```

- To resume monitoring a process, run the following command:

```
ksysvmmgr resume process process_name=<NAME>
```

- To delete a process (which is monitored), run the following command:

```
ksysvmmgr [-s] delete process process_name=<NAME>
```

An example scenario

The following examples show a scenario in which you start the VM monitor daemon, configure the VM agent by adding applications and dependencies:

1. To display the help information about the **vmm** class, run one of the following command:

   ```
   1. ksysvmmgr -h vmm
   2. ksysvmmgr help vmm
   ```

2. Start the VM monitor daemon by running the following command:

   ```
   ksysvmmgr start
   ```

3. To stop the VM monitor daemon, run the following command:

   ```
   ksysvmmgr stop
   ```

4. Modify the VM monitor daemon attributes by running the following command:

   ```
   ksysvmmgr modify log=2
   ```

5. Back up the VM monitor configuration settings by running the following command:

   ```
   ksysvmmgr backup vmm filepath=/var/ksys/config
   ```

   Note: Create the backup `HAMonitoring 2019-08-30 00:11:08.xml` in the `/var/ksys/config` path.

6. To restore the saved configuration settings, run the following command:

   ```
   ksysvmmgr restore vmm
      filepath=/var/ksys/config/HAMonitoring_2019-08-30_00:11:08.xml
   ```

   Note: stores the current configfile to /var/ksys/config/snapshot/HAMonitoring_backup.xml and the current confile is updated with the provided config filepath.

7. To collect the snapshot, run the following command:

   ```
   ksysvmmgr snap
   ```

8. To notify the VM monitor daemon to synchronize with the contents of the XML configuration file, run the following command:

   ```
   ksysvmmgr sync vmm
   ```

9. To query **vmm** parameters, run the following command:

   ```
   ksysvmmgr query vmm
   ```

10. To add an application (say app1), run the following command:

    ```
    ksysvmmgr -s add app app1 monitor_script=/tmp/monitor1.sh
            start_script=/tmp/start1.sh stop_script=/tmp/stop1.sh
    ```

11. To delete an application (say app1), run the following command:

    ```
    ksysvmmgr -s delete app app1
    ```

12. To suspend an application (say app1), run the following command:

    ```
    ksysvmmgr suspend app app1
    ```

13. To resume an application (say app1), run the following command:

```
ksysvmmgr resume app app1
```

14. To start an application (say app1), run the following command:

```
ksysvmmgr start app app1
```

15. To stop an application (say app1), run the following command:

```
ksysvmmgr stop app app1
```

16. To query all applications, run the following command:

```
ksysvmmgr query app
```

17. To query application (say app1), run the following command:

```
ksysvmmgr query app app1
```

18. To get the status of all applications, run the following command:

```
ksysvmmgr status app
```

19. To get the status of an application (say app1), run the following command:

```
ksysvmmgr status app app1
```

20. To group the applications, run the following command:

```
ksysvmmgr -s add group grp1 applist=app1,app2,app3
```

21. To modify the application list in the group, run the following command:

```
ksysvmmgr -s modify group grp1 applist=app2,app5
```

22. To ungroup (delete the group), run the following command:

```
ksysvmmgr -s delete group grp1
```

23. To display help information about the **dependency** class, run one of the following command:

```
ksysvmmgr -h dep
ksysvmmgr help dep
```

24. Notify the VM monitor daemon to synchronize with the contents of the XML configuration file, run the following command:

```
ksysvmmgr sync vmm
```

25. Add an application app1 by running the following command:

```
ksysvmmgr add app app1 monitor_script=/tmp/monitor1.sh
     start_script=/tmp/start1.sh stop_script=/tmp/stop1.sh
```

26. Modify the application attributes by running the following commands:

```
 ksysvmmgr modify app app1 new_name=newApp critical=yes
ksysvmmgr modify app app1 monitor_script=/apps/new_monitor_script
ksysvmmgr modify app app1 monitor_period=10
```

27. Create a **start_sequence** dependency with 3 applications app1, app2, app3 in the dependency list by running the following command:

```
ksysvmmgr -s add dependency dependency_list=app1,app2,app3
dependency_type=start_sequence
```

28. To create a **stop_sequence** dependency with 3 applications app4, app5, app6 in the dependency list, run the following command:

```
ksysvmmgr -s add dependency dependency_list=app4,app5,app6 dependency_type=stop_sequence
```

29. Create a **parent_child** dependency with applications app1 and app2 in the dependency list by running the following command:

```
ksysvmmgr -s add dependency dependency_list=app1,app2 dependency_type=parent_child
```

30. Display the dependency UUIDs by running the following command:

```
ksysvmmgr query dependency
```

31. To modify the details of the dependency that has UUID 1531835289870752764, run the following command:

```
ksysvmmgr -s modify dependency 1531835289870752764 dependency_list=app3,app1,app2
dependency_type=start_sequence
```

32. Display the details of the dependency that has UUID 1531835289870752764 by running the following command:

```
ksysvmmgr query dep 1531835289870752764
```

33. To delete the dependency that has UUID 1531835289870752764, run the following command:

```
ksysvmmgr -s delete dep 1531835289870752764
```

34. To display help information about the process class, run one of the following commands:

```
ksysvmmgr -h process
ksysvmmgr help process
```

35. To add a process to monitor, run the following command:

```
ksysvmmgr -s add process process_name=proc1 start_script=/tmp/startproc1.sh
stop_script=/tmp/stopproc1.sh
```

36. To query process proc1, run the following command:

```
ksysvmmgr query process ksysvmmgr query process process_name=proc1
```

37. To modify process attributes, run the following command:

```
ksysvmmgr -s modify process process_name=proc1 start_script=/tmp/new_startproc1.sh
```

38. To suspend a process, run the following command:

```
ksysvmmgr suspend process process_name=proc1
```

39. To resume a process, run the following command:

```
ksysvmmgr resume process process_name=proc1
```

40. To delete a process, run the following command:

```
ksysvmmgr -s delete process process_name=proc1
```

Troubleshooting VM Recovery Manager HA

To isolate and resolve problems in the VM Recovery Manager HA solution, you can use the following troubleshooting information.

Consider the following approach when you encounter an error or an error notification:

1. When you receive errors in the configuration or recovery operations, run the following command:

   ```
   ksysmgr query system status
   ```

   An output that is similar to the following example is displayed:

   ```
   Following error is returned for VM <vmname>:
   Partition migration failed with the following errors:
   <Error code> <Error message>
   <Error code> <Error message>
   Please review the error(s) and take any corrective actions
   ```

2. Review the suggested action and check whether you can resolve the issue.
3. If you cannot identify the cause of the error from the command output, review the log files and the trace files to diagnose the issue.
4. If you receive error notifications as an email or text message, review the `/var/ksys/events.log` file and check for the resolution.

Log files and trace files

If you receive any error message during the configuration and recovery operations by using the KSYS node, analyze the log files and trace files to find the component that is causing the error.

Analyzing the log files

Use the following log files that are available in the specified directory in your KSYS node to troubleshoot errors:

`/var/ksys/log/ksysmgr.log`
Contains the detailed processing information about each function when you run the **ksysmgr** command. The `ksysmgr.log` file contains the detailed processing information only when you specify the **`-l max`** flag along with the **ksysmgr** command.

`/var/ksys/log/ksysmgr.oplog`
Keeps a rolling record of all the **ksysmgr** operations that you ran for a specific period. All the commands that you entered are logged in this log file along with the date, time, and the transaction ID.

`/var/ksys/events.log`
Contains details about a specific event.

`/var/ksys/log/ksys_vmm.log`
Contains all the log information about each process thread such as the VM monitor daemon startup, shutdown, and restart processes, application reporting, communication between the VM monitor and host monitor daemons. The `/var/ksys/log/ksys_vmm.debug` file contains the same information in detail.

`/var/ksys/log/ksysvmmgr.log`
Contains the history of the **ksysvmmgr** commands that you ran in the command line interface. The `ksysvmmgr.log` file contains the detailed processing information only when you specify the **`-l 3`** flag along with the **ksysvmmgr** command. The `ksysvmmgr.critlog` file is a subset of the `ksysvmmgr.log` file and contains only the error information.

/var/ksys/log/ksys_vmm_ame.log
Contains log information about the application management engine. The `ksys_vmm_ame.debug` log file contains the same information in detail.

/var/ksys/log/ksys_vmm_vmH.log
Contains log information about the VM health monitor threads in the Linux VMs. The `ksys_vmm_vmH.debug` log file contains the same information in detail.

/var/ksys/log/lnx_vmmke.log
Contains log information about the kernel extensions in the Linux VMs.

/var/ksys/log/aix_vmmke.log
Contains log information about the kernel extensions in the AIX VMs.

/var/ksys/log/agents/oracle_agent/*logfiles*
Contains log information about the Oracle application processes that are registered in the VM agent configuration settings.

/var/ksys/log/agents/db2_agent/*logfiles*
Contains log information about the DB2 application processes that are registered in the VM agent configuration settings.

/var/ksys/log/agents/saphana/*logfiles*
Contains log information about the SAP HANA application processes that are registered in the VM agent configuration settings.

Use the following log files to troubleshoot the VM Recovery Manager HA GUI:

/opt/IBM/ksys/ui/server/logs/uiserver.log
Contains information about the VM Recovery Manager HA GUI server.

/opt/IBM/ksys/ui/agent/logs/uiagent.log
Contains information about the agent that is installed on each KSYS node.

/opt/IBM/ksys/ui/agent/logs/notify-event.log
Contains information about all VM Recovery Manager HA events that are sent from the GUI agent to the GUI server.

Analyzing trace files

Trace files contain the details about the processes that are running as part of the operations performed by the KSYS node. If you cannot identify the cause of the error by analyzing the log files, use the following resource manager (RM) trace files for problem determination:

/var/ct/*cluster_name*/log/mc/IBM.VMR/trace.user.*
Contains an abstract view of functions that are run for each operation. Review this log file to check the components that participated in an operation, the processing activity, the point where the operation failed, and the component that caused the failure.

/var/ct/*cluster_name*/log/mc/IBM.VMR/trace.ksys.*
Contains detailed information about the functions that are run for each operation.

/var/ct/*cluster_name*/log/mc/IBM.VMR/trace.fde.*
Contains information about the failure detection engine (FDE) processes.

/var/ct/*cluster_name*/log/mc/IBM.VMR/trace.krest.*
Contains abstract information about the HMC and VIOS processes.

/var/ct/*cluster_name*/log/mc/IBM.VMR/trace.krestlong.*
Contains detailed information about the HMC and VIOS processes.

You can use the **rpttr** command to format trace files. The syntax of the **rpttr** command follows:

```
# rpttr -f -o dict trace_file_names
```

For example,

```
#rpttr -o dict /var/ct/<cluster_name>/log/mc/IBM.VMR/trace.ksys.* >   <logfilename.log>
#rpttr -o dict /var/ct/<cluster_name>/log/mc/IBM.VMR/trace.krest.* >   <logfilename.log>
```

VM log files of a host group

You can get the VM log files of a host group or all host groups of a site by running the `export` command. You can use these log files to debug issues that might occur while starting a VM. To get VM log files, the KSYS node must be in the online state, the VM must be in the managed state, and HA monitoring must be enabled for the VM.

To collect the VM log files of a host group, run the following commands from the CLI of the KSYS node:

```
export VMRM_HGNAME=HA_HG1
```

```
snap -z "product_name=ksys_prod"
```

If the `export` command does not contain the host group name, the VM log files of all host groups will be generated.

Error notification for the KSYS events

The KSYS subsystem tracks various events that occur in the environment and saves the information in the `/var/ksys/events.log` file. The KSYS subsystem also sends emails and text notifications to the administrator if the contact information is registered while configuring the KSYS subsystem.

You can run the following command to list all events that can be notified:

```
ksysmgr query event
```

The events are categorized as critical errors, warnings, and informational events. To query all events of a specific event type, use the following command:

```
ksysmgr query event type=error|warning|info
```

The following table lists the common error events that are monitored by the KSYS subsystem:

Table 6. Common error events monitored by the KSYS subsystem

Events	Troubleshooting information
ADD_VIOS_TO_SSP_CLUSTER_FAIL	See "The SSP cluster is not created or updated successfully" on page 84
CREATE_SSP_CLUSTER_FAIL	
SSP_ATTRIBUTES_INIT_FAIL	
SSP_REGISTRY_FAIL	
SSP_RCP_DATA_FAIL	
REPOSITORY_DISK_FAILURE	
CREATE_CLIENT_ETH_ADAPTER_VLAN1_FAIL	See "Virtual switch, trunk adapter, or Ethernet adapters are not created successfully" on page 86
CREATE_CLIENT_ ETH_ADAPTER_VLAN2_FAIL	
SWITCH_CREATE_FAILED	
SWITCH_DELETE_FAILED	
TRUNK_ADAPTER_CREATE_FAILED	
TRUNK_ADAPTER_DELETE_FAILED	
FDE_FAIL_TO_REACH_VIOS	See "The host or VM failures are not detected correctly" on page 86
PM_VM_TO_CEC_MAP_TABLE_ERROR	See "The KSYS subsystem cannot find suitable target hosts" on page 83

Table 6. Common error events monitored by the KSYS subsystem (continued)	
Events	**Troubleshooting information**
LPM_FAILED	See "The LPM verification or LPM operation failed" on page 82
LPM_FAILED_DURING_VERIFICATION	
VM_RESTART_FAILED	
HOST_FAILURE_DETECTED	See "The restarted VM does not move to a stable state" on page 82
VM_FAILURE_DETECTED	
APP_FAILURE	See "The failed application does not move to a stable state after a restart operation" on page 89
APP_FAILURE_INITIATE_VM_RESTART	
APP_FAILURE_INITIATE_VM_RELOCATION	
HG_VERIFY_FAILED	See "The verification operation for a host group, host, or VM failed" on page 81
HOST_VERIFY_FAILED	
VM_VERIFY_FAILED	
ADD_HOST_MANAGED_SYSTEM_FAIL	See "The discovery operation for a host group, host, or VM failed" on page 81
NETWORK_INTERFACE_ADDED	
NETWORK_INTERFACE_ACTIVE	
NETWORK_INTERFACE_DELETED	
NETWORK_INTERFACE_FAILURE	

Solving common problems

This section describes the solutions to some problems that you might encounter when you use the VM Recovery Manager HA solution.

The discovery operation for a host group, host, or VM failed

Problem

The discovery operation failed or a VM is not discovered by the KSYS subsystem during the discovery operation.

Solution

1. Ensure that you have completed all the prerequisites that are specified in the Requirements section and the configuration steps that are specified in the Configuring section.
2. Check whether the **ha_monitor** attribute is enabled or disabled for site, host group, host, and VM by using the following commands:

```
lsrsrc IBM.VMR_SITE HAmonitor
lsrsrc IBM.VMR_HG Name HAmonitor
lsrsrc IBM.VMR_CEC Name HAmonitor
lsrsrc IBM.VMR_LPAR HaMonitor
lsrsrc IBM.VMR_LPAR Name HAmonitor
```

3. If the **HAMonitor** field shows `Disabled` or not set, enable the **ha_monitor** attribute by using one of the following commands:

```
chrsrc -c IBM.VMR_SITE HAmonitor="Enabled"
ksysmgr modify system|hg|host|vm [<name>] ha_monitor=enable
```

   If the **ha_monitor** attribute for a VM is not enabled at a VM-level, host-level, host group-level, or system-level, the VM is not considered for the discovery operation.
4. Ensure that you have started the VM monitor daemon by running the **ksysvmmgr start vmm** command. The VM agent can send heartbeats to the host monitor only when you start the VM monitor daemon.
5. Ensure that you have set all the HMC options that are specified in the Requirements section.

The verification operation for a host group, host, or VM failed

Problem

The `Phase` or the `PhaseDetail` attribute of a VM indicates a status that the verification operation failed.

Solution

During the verification operation, you can troubleshoot the possible issues as described in the following steps:

1. Analyze the verification operation flow in the `/var/ct/<cluster_name>/log/mc/IBM.VMR/trace.ksys.*` trace files and check the `Phase`, `PhaseDetail`, and `HAState` fields during each operation.
2. If any of the remote copy programs (RCP) for host group, host, or LPAR is deleted or does not exist, re-create the host group and add the hosts by using the **ksysmgr** command. Run the discovery operation and the verification operation to check whether the issue is resolved.
3. If the verification lock is acquired by some other process, the verification process might fail. Check the `trace.ksys.*` trace files to ensure that verification locks are not acquired by other threads before you start the verification operation.
4. The verification process for a VM cannot start if 32 threads are already running. Ensure that sufficient number of empty threads are available for VM task to complete in specified time.
5. Check the `/var/ct/<cluster_name>/log/mc/IBM.VMR/trace.krestlong.*` trace files to identify whether the LPM operation is successful. Rerun the verification operation, if required. Also, check whether the previous phase and `PhaseDetail` fields are cleared before you start the LPM operation.

The flexible capacity values cannot be calculated correctly

Problem

The verification operation failed because of an error in the calculated flexible capacity values.

Solution

1. If any of the remote copy programs (RCP) for host group, host, or LPAR is deleted or does not exist, the verification thread might fail. Recreate the host group and run the discovery operation.
2. If the calculated flexible capacity values are not correct, consider the following solutions:
   - Review the `/var/ct/<cluster_name>/log/mc/IBM.VMR/trace.ksys.*` trace files to check whether the policy map is created correctly by the policy manager of the target host and check the logical memory block (LMB) size that is calculated by the target host in the trace files.
   - Ensure that the specified capacity is between the minimum and maximum capacity for a VM, otherwise the VM moves to the target host with the same capacity.
   - Ensure that flexible capacity table is set properly in the host group by using the following commands:

```
#lsrsrc -s 'Name="hg_name"' IBM.VMR_HG FlexProcCapacityTable
#lsrsrc -s 'Name="hg_name"' IBM.VMR_HG FlexMemCapacityTable
#chrsrc -s 'Name="hg_name"' IBM.VMR_HG FlexMemCapacityTable={"100","70","50"}
#chrsrc -s 'Name="hg_name"' IBM.VMR_HG FlexProcCapacityTable={"100","70","50"}
#lsrsrc -s 'Name="hg_name"' IBM.VMR_HG Priority
#chrsrc -s 'Name="hg_name"' IBM.VMR_HG Priority="Medium"
```

The LPM verification or LPM operation failed

Problem

The VM cannot be moved to another host while performing the Live Partition Mobility (LPM) operation.

Solution

If the VM move operation failed while performing the LPM operation on the VMs from the KSYS subsystem or while restarting the VMs in another host, you must check the events, the VM state, and the log files to diagnose and resolve the issue. Use the following troubleshooting steps to diagnose the issue:

1. If you received an event notification, check the event details in the `/var/ksys/events.log` file and review the suggested action.
2. Identify the reason for the move operation (LPM or restart) failure in the `/var/ksys/log/ksysmgr.log` file.
3. Ensure that the RMC connection between the HMC and the VM is working. If the VM contains any firewall service, ensure that the RMC connection is allowed by the firewall. You can use the HMC **diagrmc** command to verify and correct the RMC connection.
4. If the move operation failed because of policy conflicts (such as collocation policy and anti-collocation policy), resolve the conflicts and run the move operation again. If the move operation failed because of insufficient capacity resources in the target host, increase the target host capacity and retry the move operation.
5. If the LPM or restore operation failed and the VM exists on both the source and target hosts, recover the VM by using the following command:

```
ksysmgr [-f] lpm vm vmname1 action=recover
```

6. Run the discovery and verification operations after each LPM operation to update the LPM validation state.

The restarted VM does not move to a stable state

Problem

When you restart a failed VM on another host, the VM does not move to a stable state.

Solution

The LPAR profile on the source host is deleted after the virtual machine is restarted successfully on the target host. However, if the virtual machine does not move to proper state, perform the following steps:

1. Restart the virtual machines in the target host by running the following command:

   ```
   ksysmgr [-f] restart vm|host [vmname1[,vmname2,...]]
   ```

2. If the restart operations fail, recover the virtual machine in the same host where it is located currently by running the following command:

   ```
   ksysmgr [-f] recover vm vmname
   ```

3. If the output of the restart command indicates cleanup errors, run the cleanup command manually to clean up the VM details in the source host by running the following command:

   ```
   ksysmgr cleanup vm vmname host=hostname
   ```

The KSYS subsystem cannot find suitable target hosts

Problem

The failed VMs cannot be restarted on other hosts successfully because the policy manager module cannot find suitable target hosts.

Solution

Whenever a VM or a host failure is identified by the KSYS subsystem, the policy manager module finds the best fit target host where the VM can be migrated based on various policies such as collocation, anti-collocation, affinity, priority, and blacklist. Sometimes, the policy manager module cannot find the target host for a VM because of policy conflict or resource check failures. Perform the following steps to troubleshoot this error:

1. Search for the string `Policymap VmtoCECMapTable` in the `/var/ct/<cluster_name>/log/mc/IBM.VMR/trace.user.*` trace file to find the VM-host mapping table.
2. Review the miscellaneous column of the VM-host mapping table to check the policies that have conflicts. Resolve the conflicts, run the discovery and verification operations, and then retry the recovery operations.

Restart operation failed because of a version mismatch

Problem

The failed VMs cannot be restarted on other hosts successfully because of version mismatch.

Solution

The major versions of host monitor filesets and VM monitor filesets must match for successful heartbeat operation. Run the following commands to identify the fileset versions of VM monitor and host monitor:

```
ksysmgr query vm
ksysmgr query vios
```

If any of the host monitor or VM monitor major version does not match, you must upgrade the minor version to any of the major versions.

The KSYS node restarted during the automatic restart operation

Problem

The VM restart operation is interrupted by the KSYS node restart operation.

Solution

If the KSYS node restarts during the automatic restart operation of virtual machines, the move operation is added to the policy manager module automatically and is resumed from the point where the move operation was interrupted before the KSYS node reboot operation. The move operation of virtual machines, whose restart operation is completed, are not added to the policy manager module. Check the `/var/ct/<cluster_name>/log/mc/IBM.VMR/trace.ksys.*` trace files to check the details of the operations.

The SSP cluster is not created or updated successfully

Problem

The KSYS subsystem cannot create an SSP cluster, cannot add VIOS to an existing SSP cluster, or cannot collect required SSP information.

Solution

- Re-create the host group to fix the initial SSP remote copy program (RCP).
- If the SSP cluster is not created successfully, search the following text in the `/var/ct/<cluster_name>/log/mc/IBM.VMR/trace.ksys.*` trace files: `Could not create single node cluster on VIOS`. Based on your analysis, perform the following steps:
  1. If any values of the input variables for the **`create_ssp()`** API are missing or incorrect, the problem might be in the KSYS configuration settings. Check and update the KSYS configuration settings and rerun the discovery operation.
  2. Check the `libkrest` logs for the **`kriSubmitCreateSSP()`** KREST API by examining the return code and error message to identify whether the problem is from the HMC or the VIOS.
  3. The HMC might get overloaded with multiple retry requests from the KSYS subsystem. Therefore, if you receive a message that HMC is busy, wait for some time and then retry the operation.
  4. Run the **`cluster -create`** command on the VIOS to identify whether the VIOS has any problems to create the SSP cluster. For more information about running the **`cluster`** command, see the cluster command documentation in VIOS.
- If the KSYS subsystem cannot add the VIOS to an existing SSP cluster, search the following text in the `/var/ct/<cluster_name>/log/mc/IBM.VMR/trace.ksys.*` trace files: `Could not add VIOS: xxxx to cluster xxx`. Based on your analysis, perform the following steps:
  1. If any values of the input variables for the **`add_SSP_node()`** API are missing or incorrect, the problem might be in the KSYS configuration settings. Check and update the KSYS configuration settings, and rerun the discovery operation.
  2. Check the **`kriSubmitaddSSPNode()`** API for the return code and error message to verify whether the problem is from the HMC or the VIOS. The KSYS subsystem uses the HMC REST APIs to handle the requests, therefore, HMC waits to get the acknowledgment of job completion from the VIOS. The retry operation failure does not necessarily mean that the request failed, if the error occurred in the first few retry operations.
  3. Run the **`cluster -addnode`** command on the VIOS to identify whether the VIOS has any problems to add node to the SSP cluster. For more information about running the **`cluster`** command, see the cluster command documentation in VIOS.
- If the KSYS subsystem fails to collect all the required SSP information, one or more attributes of the SSP cluster might be missing their values that are collected during the discovery operation. Check whether the storage pools are in operational state by running the **`cluster -status`** command in the VIOS. If any of the pools are not in operational state, the KSYS subsystem fails to collect the required SSP data. Refer to VIOS SSP documentation to fix the issue.
- Re-run the discovery operation so that the SSP data can be updated in the registry.

The repository disk failed

Problem

You receive a repository disk event, `REPOSITORY_DISK_FAILURE`.

Solution 1

When a repository disk fails, you can manually replace repository disk by running the `modify hg` command with a new repository disk ID. Run the `modify_Hg` command from the KSYS subsystem.

```
ksysmgr modify host_group <name> options [repo_disk=<ViodiskID>]
```

Solution 2

When a repository disk fails, you can manually replace repository disk by completing the following steps in the HMC GUI:

1. Log in to HMC GUI in a web browser as `hscroot` user.
2. Go to **Resources** > **All Shared Storage Pool Clusters**.
3. Select your cluster and click on the cluster name.
4. Click **Replace Disk**.
5. In the **Replace Repository Disk** panel, select one of the available free shared physical volumes as the new repository disk to replace the existing repository disk.
6. Click UUID value to validate the complete UUID and the local hdisk name on each VIOS.
7. Click **OK** to replace the repository disk. After the operation is complete, in the **Shared Storage Pool Cluster** window, click the repository disk UUID value to check whether it matches the selected new repository disk.
8. Run the discovery operation to update the KSYS configuration settings by running the following command:

   ```
   /opt/IBM/ksys/ksysmgr –t discover host_group <HGName>
   ```

9. After the discovery operation is complete, run the following command to verify whether the updated repository disk UUID in SSP remote copy matches the UUID in HMC:

   ```
   /opt/IBM/ksys/ksysmgr -v query host_group
   ```

   For example,

```
(0) root @ ksys305: /var/ksys
# /opt/IBM/ksys/ksysmgr -v q host_group
Name:                HG1
Hosts:               hk-8247-22L-2139F7A
                     ko-8284-22A-10FDC13
Memory_capacity:     Priority Based Settings
                     high:100
                     medium:100
                     low:100
CPU_capacity:        Priority Based Settings
                     high:100
                     medium:100
                     low:100
Skip_power_on:       No
HA_monitor:          enable
Restart_policy:      auto
VM_failure_detection_speed:    normal
Host_failure_detection_time:   90

SSP Cluster Attributes
Sspname:             KSYS_env30_ha_1
Sspstate:            UP
Ssp_version:         VIOS 3.1.0.00
VIOS:                kov2
                     hkv2
                     hkv1
                     kov1
Repo_disk:           01M0lCTTIxMDc5MDA2MDA1MDc2MzAzRkZEM0ZGMDAwMDAwMDAwMDAwMDQxNg==
HA_disk:             01M0lCTTIxMDc5MDA2MDA1MDc2MzAzRkZEM0ZGMDAwMDAwMDAwMDAwMDQxNw==
SspUuid:             d5e5d382-dc01-38e8-ab55-083ca3ffe826
XSD_version:         4.00
PoolName:            default_pool_1
PoolUuid:            000000000A281695000000005BD73628
```

Virtual switch, trunk adapter, or Ethernet adapters are not created successfully

Problem

The KSYS subsystem cannot create or delete virtual switch on a host, cannot create or delete trunk adapters on a VIOS, or cannot create or delete Ethernet adapters on VMs.

Solution

1. Ensure that the KSYS network configuration is set properly by performing the following steps:
   a. Each managed host must contain a switch for the KSYS configuration that is created during the discovery operation. Check the switch UUID that is stored in the KSYS registry by running the **lsrsrc IBM.VMR_CEC SwitchUUID** command.
   b. Each VIOS must contain two trunk adapters with VLAN IDs 101 and 102 in the associated host. Check the `ActiveHM` fields for the VIOS by running the **lsrsrc IBM.VMR_VIOS ActiveHM AdapterMACAddress AdapterSlot** command. Other virtual I/O servers show the `ActiveHM` field as blank.
   c. Each managed virtual machine must contain two virtual adapters. Check the MAC addresses and UUIDs that are saved in the registry by running the **lsrsrc IBM.VMR_LPAR AdapterMACAddress1 AdapterUUID1 AdapterMACAddress2 AdapterUUID2** command.
2. Rerun the discovery operation. If the required KSYS network is still not configured, manual intervention might be required.
3. If you find the following errors in the `traces.ksys.*` trace file:

   ```
   “Dynamic add of virtual I/O resources failed”
   “The specified slot contains a device or devices that are currently configured.”
   “This virtual switch cannot be deleted since the following virtual networks are using this
   virtual switch…”
   ```

   Investigate these errors at the HMC level and resolve them as specified in the error message.
4. Access the HMC and run the **chhwres** command to create or delete switches and adapters. After you verify the configuration on the HMC, run the discovery operation to update the KSYS configuration settings.

The host or VM failures are not detected correctly

Problem

The failure detection engine (FDE) module of the KSYS subsystem cannot connect to the VIOS or is not working correctly.

Solution

The Failure Detection Engine (FDE) module is a KSYS module that detects the health of a virtual machine and initiates the relocation request if the VM must be moved. The most common reasons for monitoring failures by the FDE module are as follows:

- The HA monitoring at the site or system level is disabled.
- The HA monitoring at the VM level is disabled.
- The KSYS daemon might be repeatedly restarting.

Diagnose the issue by performing the following steps:

1. Check whether the FDE module is active by performing the following steps:
   a. Check whether the HA monitoring is enabled by searching the `Trace Started – Pid` string in the `/var/ct/< cluster_name>/log/mc/IBM.VMR/trace.fde.*` trace files. For example:

   ```
    [00] 06/12/18 ____ 14:14:29.274292  ****************** Trace Started - Pid = 6553828
   **********************
   [00] 06/12/18 _VMR 14:15:20.910220 DEBUG FDEthread.C[125]: Monitoring Enabled.
   ```

b. Review the **HAmonitor** attribute settings by using the **lsrsrc** command in the KSYS node as follows:

i) Check the persistent values that are specific to the HA monitoring and are saved in the VMR_SITE class by running the following command:

```
# lsrsrc -c IBM.VMR_SITE HAmonitor hostFDT VMFDS VMthreshold FDEpollInt
Resource Class Persistent Attributes for IBM.VMR_SITE
resource 1:
        HAmonitor   = "Enabled"
        hostFDT     = 90
        VMFDS       = "Normal"
        VMthreshold = 50
        FDEpollInt  = 20
```

The high-availability monitoring is enabled at a global level based on the **HAmonitor** value.

ii) Check the persistent values that are specific to the HA monitoring and are saved in the VMR_HG (host group) class by running the following command:

```
# lsrsrc IBM.VMR_HG Name HAmonitor hostFDT
Resource Persistent Attributes for IBM.VMR_HG
resource 1:
        Name      = "HG1"
        HAmonitor = "Enabled"
        hostFDT   = 0
```

If the failure detection time for host (hostFDT value) is 0, the value specified at the site level is used.

iii) Check the persistent values that are specific to the HA monitoring and are saved in the VMR_CEC (host) class by running the following command:

```
# lsrsrc IBM.VMR_CEC HAmonitor VMFDS
Resource Persistent Attributes for IBM.VMR_CEC
resource 1:
        HAmonitor = "Enabled"
        VMFDS     = ""
```

iv) Check the persistent values that are specific to the HA monitoring and are saved in the VMR_VIOS class by running the following command:

```
# lsrsrc IBM.VMR_VIOS Name ViosUuid CecUuid ActiveHM CAAstate HMrespSec HMresponsive
CAAstateReason MonitorMode SlowIO
Resource Persistent Attributes for IBM.VMR_VIOS
resource 1:
        Name           = "lasagnav1"
        ViosUuid       = "7B5099B5-019E-443F-98EB-E04A680D6DA6"
        CecUuid        = "ac420f59-cdbf-3ab8-b523-37f66d461741"
        ActiveHM       = 102
        CAAstate       = "UP"
        HMrespSec      = 1
        HMresponsive   = "Yes"
        CAAstateReason = ""
        MonitorMode    = "GLOBAL"
        SlowIO         = 0
```

If a VIOS is running in a LOCAL mode, the MonitorMode field is set to LOCAL. If the host monitor is not operating correctly, the MonitorMode field is set to DOWN.

v) Check the persistent values that are specific to the HA monitoring and are saved in the VMR_LPAR class by running the following command:

```
lsrsrc IBM.VMR_LPAR Name LparUuid CecUuid HMstate HAmonitor
     HBmissed HMstateHM1 HBmissedHM1 HMstateHM2 HBmissedHM2 notAvailHM1 notAvailHM2
...
resource 8:
        Name        = "romano001"
        LparUuid    = "2C55D2BB-1C50-49F1-B1A3-5C952E7070C7"
        CecUuid     = "caffee0a-4206-3ee7-bfc2-f9d2bd3e866f"
        HMstate     = "STARTED"
        HAmonitor   = "Enabled"
        HBmissed    = 0
        HMstateHM1  = ""
        HBmissedHM1 = 0
        HMstateHM2  = ""
        HBmissedHM2 = 0
        notAvailHM1 = 0
        notAvailHM2 = 0
```

The variables HMstateHM1, HBmissedHM1, HMstateHM2, HBmissedHM2, notAvailHM1, and notAvailHM2 are applicable only for the LOCAL mode. The HMstateHM1, HMstateHM2, HBmissedHM1, and HBmissedHM2 variables store the state of the VM as observed from VIOS1 and VIOS2. If the notAvailHM1 and notAvailHM2 variables are set to 1, it implies that no data was available for this VM from the VIOS.

2. Check whether the FDE module is requesting health information from the VIOS and whether it obtained data from the VIOS by performing the following steps:
   a. Identify the VIOS that is associated with the request. For example:

```
[00] 06/12/18 _VMR 14:15:20.910261 DEBUG FDEthread.C[190]: Use VIOS 1F5D7FFC-34BD-45B6-
BD4F-101512D9BD2A for polling
```

   b. Check whether a REST request was initiated. For example:

```
[00] 06/12/18 _VMR 14:15:21.096723 DEBUG VMR_HMC.C[6728]:
getQuickQuery: Calling kriSubmitQuickQuery!. HMC:9.3.18.186,
viosUuid: 1F5D7FFC-34BD-45B6-BD4F-101512D9BD2A
```

```
[00] 06/12/18 _VMR 14:16:04.250468 DEBUG VMR_HMC.C[6617]:
getNeedAttn: Calling kriSubmitNeedAttn!. HMC:9.3.18.186,
viosUuid: 1F5D7FFC-34BD-45B6-BD4F-101512D9BD2A
```

   c. Check whether the REST request was successful. For example:

```
[00] 06/12/18 _VMR 14:16:05.537662 DEBUG VMR_HG.C[10768]:
FDE doNeedAttn success GLOBAL_DATA
```

   d. Determine the VIOS health packet content. For example:

```
[00] 06/12/18 _VMR 14:16:05.537635 DEBUG VMR_HMC.C[6666]: JobOutput
[00] 06/12/18 _VMR 14:16:05.537635 DEBUG <VIO><Response>
… XML nodes here with data inside the Response node …
```

3. Identify the actions taken by the FDE module, if any, by performing the following steps:
   a. Search for the string Task added to check whether the FDE module has passed the tasks to other components. For example:

```
 [00] 06/13/18 _VMR 13:18:02.631206 DEBUG needAttn.C[918]:
RESYNC HM TASK ADDED: vios 1f5d7ffc-34bd-45b6-bd4f-101512d9bd2a
```

   If the FDE module passed the task, the task is added to the KSYS queue. The trace.ksys.* trace files might contain further details.
   b. Check whether a move operation is initiated by searching the RECOVERY TASK ADDED for LPAR string. If you cannot find this string, the VM has not met the criteria for a move operation. For example, threshold on missed heartbeats has not reached:

```
[15] 06/11/18 _VMR 12:34:08.266355 DEBUG VMR_LPAR.C[14541]:
ssetHBmissed 46 for romano001: 2C55D2BB-1C50-49F1-B1A3-5C952E7070C7
```

c. Check whether the FDE module enabled the local mode. For example:

```
[06] 06/11/18 _VMR 09:18:48.906817 DEBUG FDEthread.C[209]:
    Did not find a VIOS - Going into local database mode
[06] 06/11/18 _VMR 09:18:48.906874 DEBUG FDEthread.C[679]:
    Use VIOS 6F97A18C-3738-4DE6-901A-96A338A3BA80 for local DB VLANID 101
[06] 06/11/18 _VMR 09:18:48.907018 DEBUG FDEthread.C[679]:
    Use VIOS 50C3E089-2254-4322-9B98-57038A701813 for local DB VLANID 102
[06] 06/11/18 _VMR 09:18:48.907065 DEBUG VMR_HG.C[10841]:
    FDE performing doNeedAttn LOCAL_DATA
```

In the global mode, the request is sent to the VIOS and the FDE module waits for a response. The response is parsed and the FDE module either takes action or moves the task to the KSYS subsystem. The local mode provides information about when heartbeat was missed.

You cannot restore a previously backed up KSYS configuration snapshot

Problem

When you attempt to restore a previously backed up KSYS configuration snapshot, you receive error messages indicating that the restore operation is not successful.

Solution

1. Search the `/var/ksys/log/ksysmgr.log` file for any of the following text when you receive error messages during the snapshot operations to find the cause of the error: `add snapshot`, `removing old configuration`, `creating new cluster`, `creating HMC, host, host group`, and so on.
2. Ensure that the existing KSYS node is not in the corrupted state. If the KSYS node is corrupted, reinstall all the KSYS filesets.
3. Query the snapshots to check whether all resource attribute values are set correctly by using the following command:

   ```
   ksysmgr query snapshot filepath=filename
   ```

   The default location for a saved snapshot is `/var/ksys/snapshots/`.
4. If you receive a host group creation error, one of the HA disks (`ha_disk`) and repository disks (`repo_disk`) might not be available. In this case, check and resolve the disk availability.
5. If you receive error messages about cluster type, check whether you have set the type of the cluster. After a cluster is created and the `IBM.VMR` daemon is started, set the **`ClusterType`** persistent attribute for the IBM.VMR_SITE class by running the following command:

   ```
   chrsrc -c IBM.VMR_SITE 'ClusterType="HA|DR"'
   ```
6. Ensure that the `IBM.VMR` daemon is in the active state. If not, reinstall the daemon.

The failed application does not move to a stable state after a restart operation

Problem

The VM agent subsystem cannot restart the failed application successfully.

Solution

1. Run the **`ksysvmmgr query app <NAME>`** command to check the state and UUID of an application. The application is in one of the following states:

   UNSET

   State of an application when the application monitoring starts, but its status is not set.

TO_START
State of an application when the application monitoring has failed. The application is successfully stopped and must be started.

NORMAL
State of an application when the application is monitored properly.

NOT_MONITORED
State of an application when the application is not monitored because the daemon is not started or because the application monitoring is suspended.

FAILING
State of an application when the application is receiving monitor script error. The application has not yet failed because the number of successive failures to trigger a restart operation is not reached.

TO_STOP
State of an application when the application monitoring has failed, the threshold frequency of application monitoring is passed, the application has failed and must be restarted (first stopped, then restarted).

NOT_STOPPABLE
State of an application when the application cannot be stopped. Although the stop script is run, the stop operation fails continuously or times out.

NOT_STARTABLE
State of an application when the application cannot be started. Although the start script is run, the start operation fails continuously.

ABNORMAL
State of an application when an abnormal condition occurs during monitoring, stopping or starting operations. For example, monitor, stop, or start scripts are not found or cannot be run.

FAILURE
State of an application when the application can be restarted but remains in the failure state even after successful restart operations.

2. Search the UUID of the associated application in the `/var/ksys/log/ksys_vmm.log` file to get more information about the application failure such as heartbeat requests, VM removal requests, and application reporting.
3. Ensure that you have provided sufficient inputs to the application agents. The VM agent supports the following application agents:

   ORACLE

   a. Ensure that you provide the correct instance name and oracle database name to the Oracle agent scripts. For example: `oracle` (instance name) and DBRESP (database name).
   b. Ensure that you specify the correct `listener.ora` file in the `ORACLE_HOME/TNS_ADMIN` location for the listener processes to work.
   c. Ensure that the specified start, stop, and monitor scripts are working correctly with the database.
   d. Analyze the `/var/ksys/log/agents/oracle_agent/<logfilename>` file to diagnose the agent script failures. These log files contain information about any missing attribute or parameter.

   DB2

   a. Ensure that you provide the correct DB2 instance owner name to the DB2 agent scripts. For example: `db2inst1` (instance owner).
   b. Ensure that you create the DB2 database before running any script. The scripts monitor the created database for the instance owner.
   c. Analyze the `/var/ksys/log/agents/db2_agent/<logfilename>` file to diagnose the agent script failures. These log files contain information about any missing attribute or parameter.

SAPHANA

a. Ensure that you provide the correct instance name and the database number to SAP HANA agent scripts. For example: S01 (instance name) and HDB01 (database number).

b. Ensure that you specify the application version, instance name, and database number while adding the application. Otherwise, the application version field remains empty.

c. Analyze the log files in the `/var/ksys/log/agents/saphana/` directory to diagnose the agent script failures. These log files contain information about any missing attribute or parameter.

d. Ensure that you have marked the application as critical by using the **`ksysvmmgr modify app <NAME> critical=yes`** command. The KSYS subsystem restarts a failed application only when you mark the application as critical. When a critical application in a VM reports a permanent failure state, diagnose the issue in the VM by checking the `ksys_vmm.log` file. When a non-critical application fails, the KSYS subsystem flags this application as failed and notifies you to take further action.

Core dump error in the POSTGRES database

Verify whether the POSTGRES database is running on all the VIOS nodes. If the database not running, run the following command to restart the POSTGRES database.

```
start - vdba -cm -start
```

You cannot log in to the VM Recovery Manager HA GUI

Problem

You cannot log in to the VM Recovery Manager HA GUI.

Solution

1. Check for issues in the `/opt/IBM/ksys/ui/server/logs/uiserver.log` file.
2. If you received an error message, `Permission missing on Smuiauth: login will not be done`, verify that the **`smuiauth`** command is installed correctly. Also, verify that the **`smuiauth`** command has the correct permissions by running the **`ls -l`** command from the `/opt/IBM/ksys/ui/server/lib/auth/smuiauth` directory. An example output follows:

   ```
   -r-x------    1 root     system        21183 Jun 11      21:48
   ```

3. Verify that you can run the **`smuiauth`** command successfully by running the command along with the **`-h`** flag.
4. Verify that the pluggable authentication module (PAM) framework is configured correctly by locating the following lines in the `/etc/pam.conf` file:

   ```
   smuiauth        auth       required     pam_aix
   smuiauth        account    required     pam_aix
   ```

   The PAM is configured when you install the `ksys.ui.server` fileset.

You cannot register a KSYS node in the VM Recovery Manager HA GUI server

Problem

You cannot register a KSYS node in the VM Recovery Manager HA GUI server.

Solution

1. Check for issues in the `/opt/IBM/ksys/ui/server/logs/uiserver.log` file by performing the following steps:
   a. If SSH File Transfer Protocol (SFTP)-related signatures exist in the log file, such as `Received exit code 127` while establishing SFTP session, a problem exists with the SSH communication between the VM Recovery Manager HA GUI server and the KSYS node that you are trying to add.
   b. From the command line, verify that you can connect to the target system by using SFTP. If you cannot connect, verify that the daemon is running on the GUI server and the target node by running the **ps -ef | grep -w sshd | grep -v grep** command.
   c. Check the SFTP subsystem configuration in the `/etc/ssh/sshd_config` file and verify that following path is correct.

   ```
   Subsystem       sftp    /usr/sbin/sftp-server
   ```

   If the path is not correct, you must enter the correct path in the `/etc/ssh/sshd_config` file, and then restart the `sshd` subsystem.
2. Check for issues in the `/opt/IBM/ksys/ui/agent/logs/agent_deploy.log` file on the target cluster.

Node server down: GUI fails to start

Problem

The VM Recovery Manager HA GUI server is not working correctly.

Solution

If the applications are not running correctly, the node server status might be causing the issue. Run the **ps -ef | grep node** command to check the status and run the **startsrc -s vmruiserver** command to start the node server.

Unplanned system reboot causes fallover attempt to start GUI

Problem

You cannot access the VM Recovery Manager HA GUI because of an unplanned GUI server or GUI agent node reboot operation.

Solution

During the system node reboot operation, you cannot access the GUI. Run the **lssrc -s vmruiserver** command to check the status of the `vmruiserver` subsystem.

```
#lssrc -s vmruiserver
Subsystem       Group     PID     Status
vmruiserver     vmrui             inoperative
```

If the status of the `vmruiserver` subsystem is displayed as `inoperative`, run the **startsrc -s vmruiserver** command to restart the UI server node from the command line. You can then access the GUI and register the agent nodes again.

Unsuccessful Deployment: Dependency file missing during installation

Problem

A dependency file is missing during the installation of the GUI server and the GUI agent filesets.

Solution

Determine the missing file from the log files that you received by using the **installp -e** flag and install that dependency file from a certified host.

You cannot stop or start the GUI server and GUI agent processes

Problem

You cannot stop or start the GUI server and agent processes.

Solution

- **GUI server**: Stop the GUI server by running the following command: `stopsrc -s vmruiserver`.

  Restart the GUI server by running the following command: `startsrc -s vmruiserver`. If you are starting the GUI server for the first time after installing GUI server, run the **`vmruiinst.ksh`** command. For information about running this command, see "Installing GUI server filesets" on page 21.
- **GUI agent**: Stop the GUI agent process by running the following command in the guest VM: `stopsrc -s vmruiagent`. This command unregisters the KSYS node from the GUI server and the KSYS node will no longer be accessible from the GUI server.

  Restart the GUI agent by running the following command: `startsrc -s vmruiagent`. This command registers the KSYS node again.

Collecting diagnostic data to contact IBM Support

If the current error symptom does not match any of the existing troubleshooting scenarios, you must collect diagnostic data and contact IBM Support.

You must provide the following information when you contact IBM Support:

- Details about the failing operation.
- When the error or problem is occurring in an operation, along with the timestamp.
- The managed host names and virtual machine names that are associated with the problem.

To collect the data and log files for IBM Support, complete the following steps:

1. Collect the diagnosis data as soon as the problem is detected by running the following command in the KSYS LPAR:

   ```
   # snap vmsnap
   ```

   This command stores all the important log files and trace files in a compressed file at the following location: `/tmp/ibmsupt/ksys.pax.Z`.
2. Collect the data from the guest virtual machines by running the following command:

   ```
   # ksysvmmgr snap
   ```

   This command creates a `*.pax.gz` file in the `/var/ksys/log/snap` directory.
3. Depending on the issues, collect data from the Virtual I/O Servers and collect portable executable debug (PEDBG) data from the associated Hardware Management Consoles. For more information, see the following IBM Support web pages:
   - Collecting snap data on VIOS partition
   - Collecting PEDBG from the HMC
4. Submit the problem to IBM Support by using one of the following methods:
   - IBM Support Portal (https://www.ibm.com/mysupport/s/?language=en_US).
   - Contact IBM Support in your region. For contact information for your country, see the Directory of worldwide contacts website (http://www.ibm.com/planetwide/).

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.