IBM Geographically Dispersed Resiliency for Power Systems

Version 1.2.0.0

# *Deployment Guide*

IBM

IBM Geographically Dispersed Resiliency for Power Systems

Version 1.2.0.0

# *Deployment Guide*

IBM

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 133.

**First edition (December 2017)**

This is the latest edition for Geographically Dispersed Resiliency for Power Systems Version 1.2.0.0 until otherwise indicated in a newer edition.

**© Copyright IBM Corporation 2017, 2019.**
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

© Copyright IBM Corp. 2017, 2019

# About this document

The Geographically Dispersed Resiliency for Power Systems™ solution is a set of software components that together provide a disaster recovery 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 GDR solution. This information provides system administrators with complete information about the following sections:

- Components that are used in the GDR solution.
- Planning the GDR implementation in your production environment and the minimum software requirements.
- Installing the GDR filesets.
- Configuring your environment to use the GDR solution.
- Recovering the virtual machines after a planned or unplanned outage by using the GDR solution.
- Administering other tasks associated with the GDR solution.
- Troubleshooting any issues associated with the GDR solution.
- Using the GDR 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 the 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 GDR

The command name and some flags in the GDR 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.

The case-sensitivity of the command name and the flags are listed as follows:

- The command name is case-sensitive.
- The **ACTION** flag is not case-sensitive.
- The **CLASS** flag is not case-sensitive.
- The **NAME** flag is case-sensitive.
- The **ATTR** flag and the corresponding *VALUE* parameter are case-sensitive.

**Correct example**

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

**Incorrect example**

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

© Copyright IBM Corp. 2017, 2019

## ISO 9000

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

# Overview for Geographically Dispersed Resiliency for Power Systems

Disaster recovery of applications and services is a key component to provide continuous business services. The Geographically Dispersed Resiliency for Power Systems (GDR) solution is a disaster recovery solution that is easy to deploy and provides automated operations to recover the production site. The GDR solution is based on the Geographically Dispersed Parallel Sysplex™ (GDPS®) offering concepts that optimizes the usage of resources. This solution does not require you to deploy the backup virtual machines (VMs) for disaster recovery. Thus, the GDR solution reduces the software license and administrative costs.

The following high availability (HA) and disaster recovery (DR) models are commonly used by customers:

- Cluster-based technology
- VM restart-based technology

Clustered HA and DR solutions typically deploy redundant hardware and software components to provide near real-time failover when one or more components fail. The VM restart-based HA and DR solution relies on an out-of-band monitoring and management component that restarts the virtual machines on other hardware when the host infrastructure fails. The GDR solution is based on the VM restart technology.

The following figure shows the conventional high-availability and disaster recovery solution that is based on the HA clustering model. The GDR solution does not use this model.

© Copyright IBM Corp. 2017, 2019

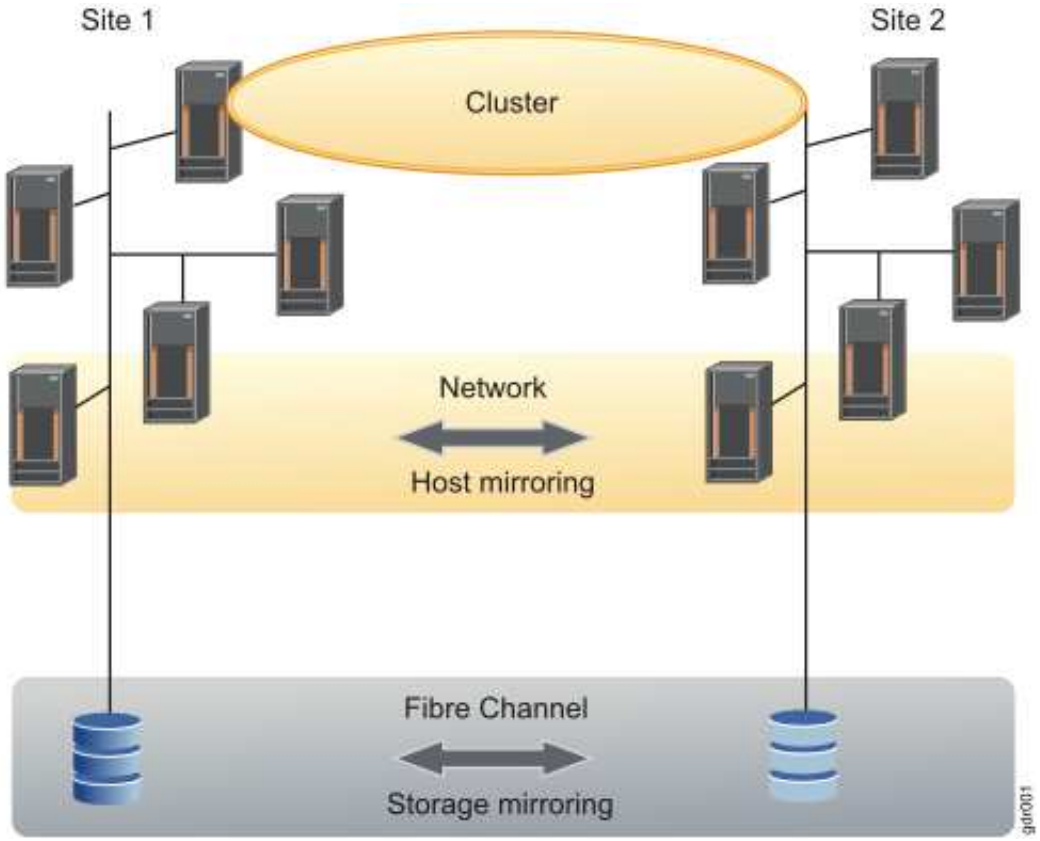

*Figure 1. Cluster-based disaster recovery model*

The following figure shows the disaster recovery solution that uses the VM restart-based technology. The GDR solution uses this model.

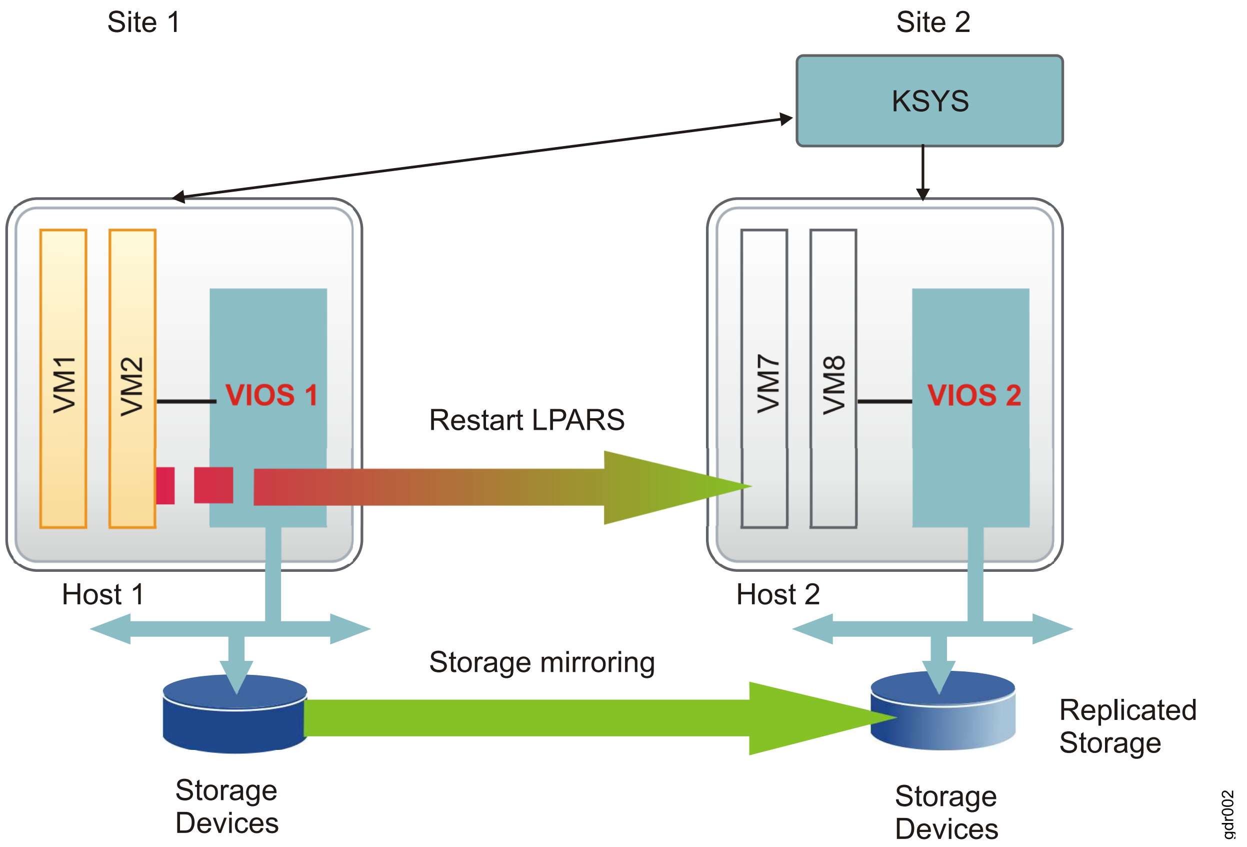

*Figure 2. VM restart-based disaster recovery model*

The following table identifies the differences between the conventional cluster-based disaster recovery model and the GDR solution:

*Table 1. Comparison between the cluster-based DR solution and the GDR solution*

| **Parameters** | **Cluster-based disaster recovery model** | **VM restart disaster recovery model that is used by the GDR solution** |
|---|---|---|
| Deployment method | Redundant hardware and software components are deployed in the beginning of implementation to provide near real-time failovers when some of the components fail. | With virtualization technology, many images of the operating system are deployed in a system. These virtual machines are deployed on physical hardware by the hypervisor that allocates and manages the CPU, memory, and I/O physical resources that are shared among the VMs. |
| Dependency | This solution relies on monitoring and heartbeat capabilities within the cluster to monitor the health of the cluster and take recovery action if a failure condition is detected. | This solution relies on an out-of-band monitoring software that works closely with the hypervisor to monitor the VM environment and to provide a disaster recovery mechanism for the VM environment. |
| Workload startup time | The workload startup time is faster because the virtual machines and the software stack are already available. | The VMs require additional time to restart in the backup environment. |
| Cluster administration required | Yes | No |
| Error coverage | Comprehensive. This solution monitors the entire cluster for any errors. | Limited. This solution monitors the servers and the virtual machines for errors. |
| Deployment simplicity | This solution must be set up in each VM. | Aggregated deployment at the site level. |

*Table 1. Comparison between the cluster-based DR solution and the GDR solution (continued)*

| Parameters | Cluster-based disaster recovery model | VM restart disaster recovery model that is used by the GDR solution |
|---|---|---|
| Protected workload type | Critical workloads can be protected by using this solution. | Less critical workloads can be selected for protection by using this solution. |
| Software license and administrative costs | This solution costs more because redundant software and hardware are required to deploy this solution. | This solution costs less because of optimized usage of resources. |

A disaster recovery implementation that uses a set of scripts and manual processes at a site level might take more time to recover and restore the services. The GDR solution automates the operations to recover your production site. This solution provides an easy deployment model that uses a controller system (KSYS) to monitor the entire virtual machine (VM) environment. This solution also provides flexible failover policies and storage replication management.

The following figure shows the architecture and the basic components that are used in the GDR solution:

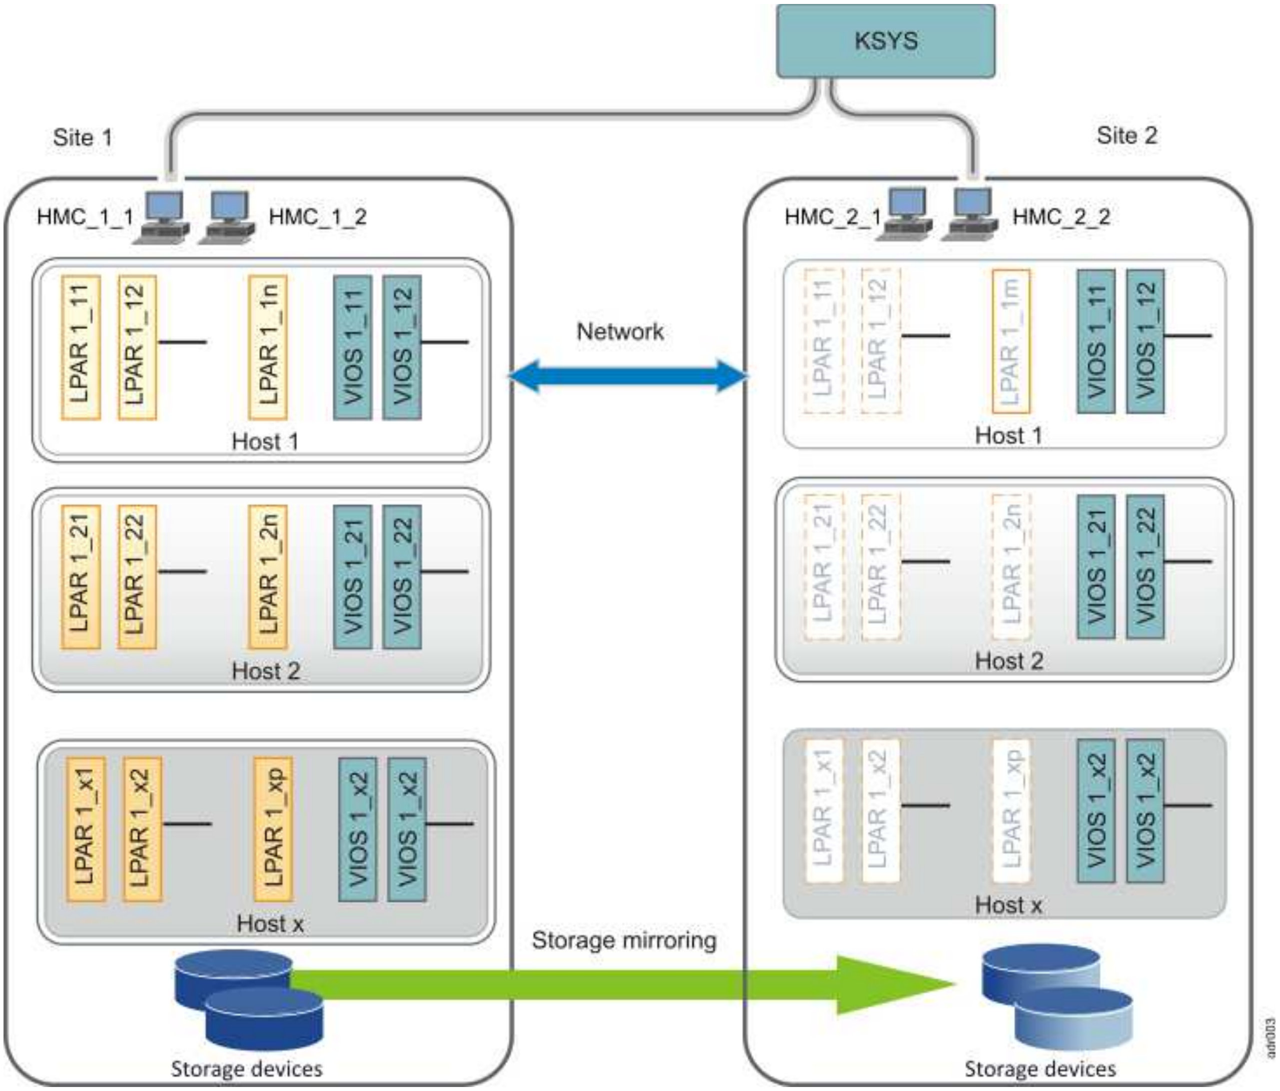

*Figure 3. An example of GDR solution architecture*

**Related information**:

Redbook: IBM Geographically Dispersed Resiliency for IBM Power Systems

# What's new in Geographically Dispersed Resiliency for Power Systems 1.2.0.0

Read about new or significantly changed information in GDR 1.2.0.0.

## How to see what's new or changed

In this PDF file, you might see revision bars (|) in the left margin that identifies new and changed information.

## April 2019

The following information is a summary of the updates that are made to GDR:

- Updated information about Hitachi fileset and log file location in the KSYS package topic.

## January 2018

The following information is a summary of the updates that are made to GDR:

- Updated information about Hitachi fileset and log file location in the KSYS package topic.

## December 2017

The following information is a summary of the updates that are made to GDR:

- GDR 1.2.0.0 supports host group based DR management. You can group a set of hosts with a logical name and perform DR operations (such as failover) at the host group level. For more information, see the following topics:
  - Hosts
  - Creating a host group
- You can test the disaster recovery operation without impacting the existing business operations or backup operations. For more information, see Failover rehearsal of the disaster recovery operation.
- GDR 1.2.0.0 supports Hitachi storage systems for storage management. For more information, see the following topics:
  - Hitachi storage systems
  - Registering Hitachi storage devices
- GDR 1.2.0.0 supports shared model of deployment, along with synchronous and asynchronous replicated environments. In the shared model, storage disks are shared across the sites and GDR does not perform any mirror management. For more information, see Managing the shared storage configuration.
- You can prioritize the virtual machines within a single host or among multiple hosts such that the virtual machines that run important workloads are considered first for the move operation. For more information, see Adding hosts to the KSYS subsystem.
- You can use the flexible capacity option to start the virtual machines on the backup site with a different capacity as compared to the active site. For more information, see Configuring the flexible capacity policies.
- You can create network mapping policies that contains a mapping of VLAN IDs or virtual switches for virtual machines when moved from the active site to the backup site. For more information, see Configuring the network mapping policy.

© Copyright IBM Corp. 2017, 2019

- The KSYS subsystem provides an option to use the current or previously saved LPAR or virtual machine (VM) profile to retry the recovery operation. For more information, see Recovering the failed virtual machines
- You can upgrade your GDR version to GDR 1.2.0.0 by using the following information: Upgrading the KSYS software.
- You can run automated checks as scripts at host-group level also. For more information, see Running scripts for additional checks.
- Updated the Managing the system attributes topic with new system-wide attributes.
- Updated the GDR restrictions topic with new limitations.
- Updated the Solving common problems topic with new troubleshooting scenarios.
- Updated the ksysmgr command topic with new usage syntax and examples.
- Updated the FAQ topic.

# GDR concepts

The GDR solution provides a highly available environment by identifying a set of resources that are required for processing the virtual machines in a server during disaster situations.

The GDR solution uses the following subsystems:

- Controller system (KSYS)
- Site
- Host
- Virtual machines (VMs) or logical partitions (LPARs)
- Storage
- Network
- Hardware Management Console (HMC)
- Virtual I/O Server (VIOS)

The following figure shows the GDR architecture:

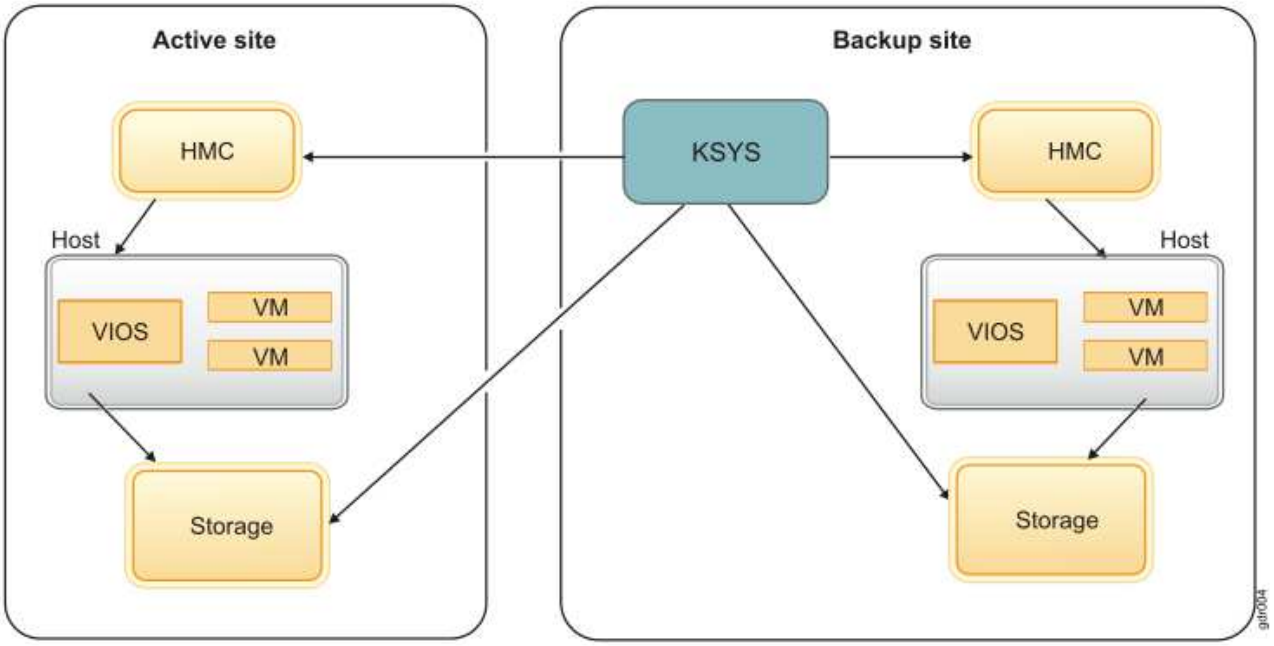

*Figure 4. Components of GDR solution*

## Controller system (KSYS)

The controlling system, also called *KSYS*, is a fundamental component in the Geographically Dispersed Resiliency for Power Systems (GDR) solution. It provides a single point of control for the entire environment managed by the GDR solution.

The KSYS cannot be affected by errors that can cause an outage in the production systems. Therefore, the KSYS must be self-contained and share a minimum number of resources with the production system. Ensure that you deploy the KSYS in the backup site so that the KSYS is isolated from any issues or

© Copyright IBM Corp. 2017, 2019

failure in the active site. In the GDR solution, the KSYS must have an out-of-band deployment in its own logical partition that must be running on AIX® 7 with 7200-01, or later.

The following figure shows a logical interaction between the KSYS and other physical components of the GDR solution:

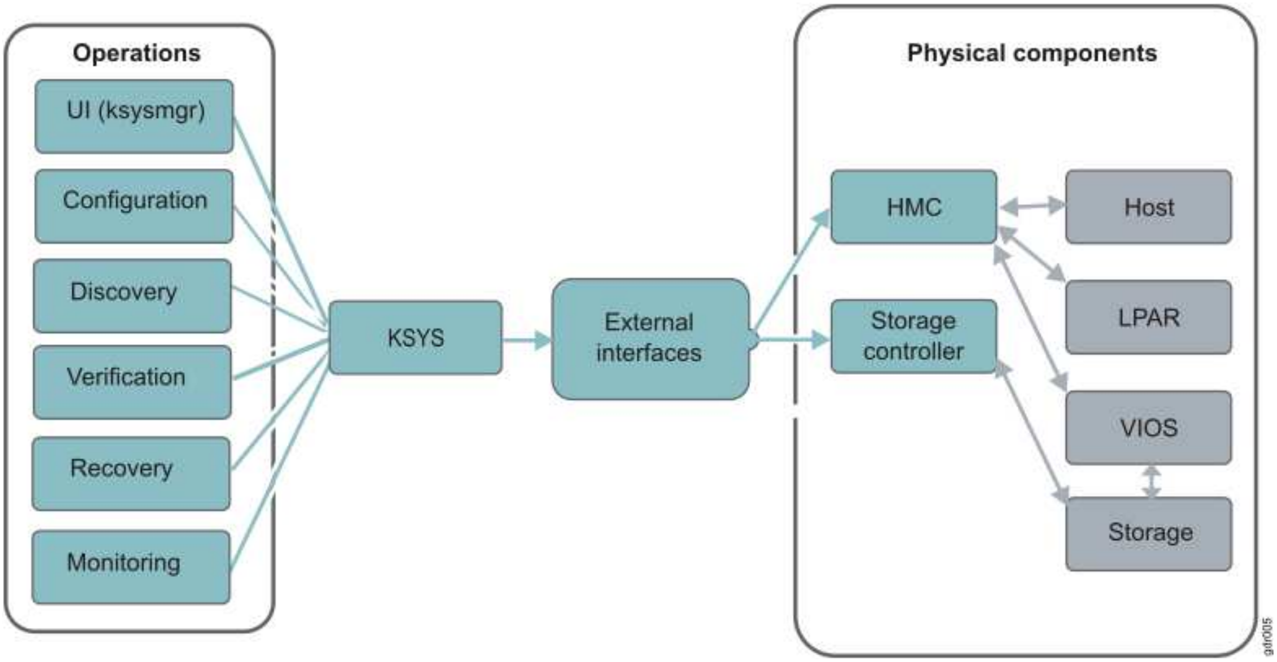

*Figure 5. Logical interaction between KSYS and other components*

The KSYS is responsible for recovery actions if a disaster or a potential disaster occurs. Therefore, the availability of KSYS is a fundamental requirement of the solution. The KSYS must be deployed in the backup site. The KSYS must remain operational even if the active site fails or if the disks located in the active site fail.

The KSYS constantly monitors the production environment for any unplanned outage that affects the production site or the disk subsystems. If an unplanned outage occurs, the KSYS analyzes the situation to determine the status of the production environment. When a site fails, the KSYS notifies the administrator about the failure. If the failure is severe, the administrator can initiate a site takeover. The KSYS pauses the processing of the data replication to ensure secondary data consistency and to process the site takeover.

The KSYS handles discovery, verification, monitoring, notification, and recovery operations to support disaster recovery for the GDR solution. The KSYS interacts with the Hardware Management Console (HMC) to collect configuration information of the managed systems. The KSYS interacts with the Virtual I/O Server (VIOS) through the HMC to obtain storage configuration information of the virtual machines. The KSYS also provides storage replication management and Capacity on Demand (CoD) management.

You can configure the KSYS by using the **ksysmgr** command. The **ksysmgr** command has the following format:

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

## | Security implementation for the KSYS subsystem

| The KSYS subsystem runs in an AIX logical partition (LPAR). You can customize security for this LPAR
| as per the AIX security requirements of your organization.

The KSYS management is enabled only for the root user in the KSYS LPAR. The KSYS subsystem does not communicate to any external systems except the following interfaces:

**HMC** The KSYS subsystem uses REST APIs to communicate to HMC. You must enable HTTPS in your environment for REST API communication.

**Storage subsystems**
The KSYS subsystem communicates to the storage subsystems by using APIs provided by the storage vendors. Refer to storage vendor documentation for any specific security requirements of these APIs.

**Related reference**:

"ksysmgr command" on page 113

## Hardware Management Console (HMC)

The controller system (KSYS) interacts with the HMC for processes that involve host, Virtual I/O Server (VIOS), and disk subsystem resources. These processes include discovery, monitoring, recovery, and cleanup.

The HMC collects information about the managed systems (hosts), VIOS, and logical partitions (LPARs) that can be managed by the KSYS. For example, the HMC collects information about the system processor, system memory, hardware, worldwide port name (WWPN) of the physical Fibre Channel adapter. The HMC also checks for VIOS capability for disaster recovery operations.

To ensure enhanced availability, you can configure dual Hardware Management Consoles in your environment. In this case, if one HMC is down or unreachable, the KSYS can use another configured HMC to collect the required information.

The HMC provides the Representational State Transfer (REST) application program interfaces (APIs) to KSYS to perform the following functions:

- Checks system capability for each operation.
- Collects information about the host state, LPAR state, VIOS state, and the IP addresses of the host, VIOS, and LPAR that KSYS can use for subsequent monitoring.
- Provides the disk mapping information to the VIOS in the backup site.
- Validates the backup site by checking whether the destination hosts are capable of the remote restart operation.
- Provides appropriate return codes to KSYS, so that KSYS can perform the required recovery actions.
- Cleans up disk mapping information in VIOS when the mapping information is no longer required.
- Cleans up the data and workload-related information from the primary site and saves the data in the KSYS data store.

**Related information**:

HMC REST APIs

## Sites

*Sites* are logical names that represent your sites. You must create sites at the KSYS level. All the Hardware Management Consoles, hosts, VIOS, and storage devices are mapped to one of the sites.

You can configure only two sites.

Sites can be of the following types:

**Active site (or primary site)**
Indicates the current site where the workloads are running at a specific time.

The active site type is dynamic. That is, the active site can be changed based on disaster recovery switches. The **ksysmgr** command can display and configure the current attributes of the active site.

**Backup site (or disaster recovery site)**
Indicates the site that acts as a backup for the workload at a specific time. During a disaster, or a potential disaster, workloads are moved to the backup site.

**Note:** This terminology reverses for sites after a planned disaster recovery operation.

You can create and configure sites by using the **ksysmgr** command.

The following figure shows the site attributes:

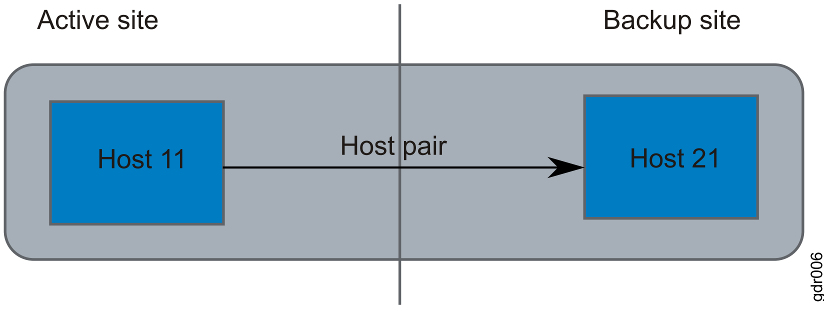

*Figure 6. Site attributes*

In this example, the `Site 1` and `Site 2` sites can switch back and forth as active and backup sites depending on where the logical partitions are currently located. You can set `Site 1` as the active site for the logical partitions in the `Host 1` host.

**Related reference**:

"ksysmgr command" on page 113

## Hosts

A *host* is a managed system in HMC that is primarily used to run the workloads. A host can also be called as a server or a Central Electronics Complex (CEC). Hosts are identified by its universal unique identifier (UUID) as tracked in the HMC.

The GDR solution uses the following host-specific terms:

**Host pair**
Indicates the set of hosts that are paired across the sites for high-availability and disaster recovery.

Each host in the host pair must meet all the resource requirements (for example, CPU, memory, VIOS-based virtual I/O aspects) to run the same workload in a disaster or a potential disaster.

**Host group**
Indicates a group of hosts that are logically chosen and named by the administrator.

You can group a set of hosts based on your business requirements. Each host must belong to a host group. For example, you can group the hosts that run similar type of workloads. Or, you can group the most important hosts together so that the monitoring and recovery operations can be performed for the set of hosts together and quickly. By grouping the most important hosts in their own group, in the event of a disaster, you can move the host group to the backup site first.

You can perform discovery, verification, move, and recovery operations at host group level. If you run these operations at site level, all the host groups are included in the operation. If a host group is already moved to the backup site, and then a site-level move operation is started, that host group is skipped from the operation.

The following figure shows the host-related configuration across sites:

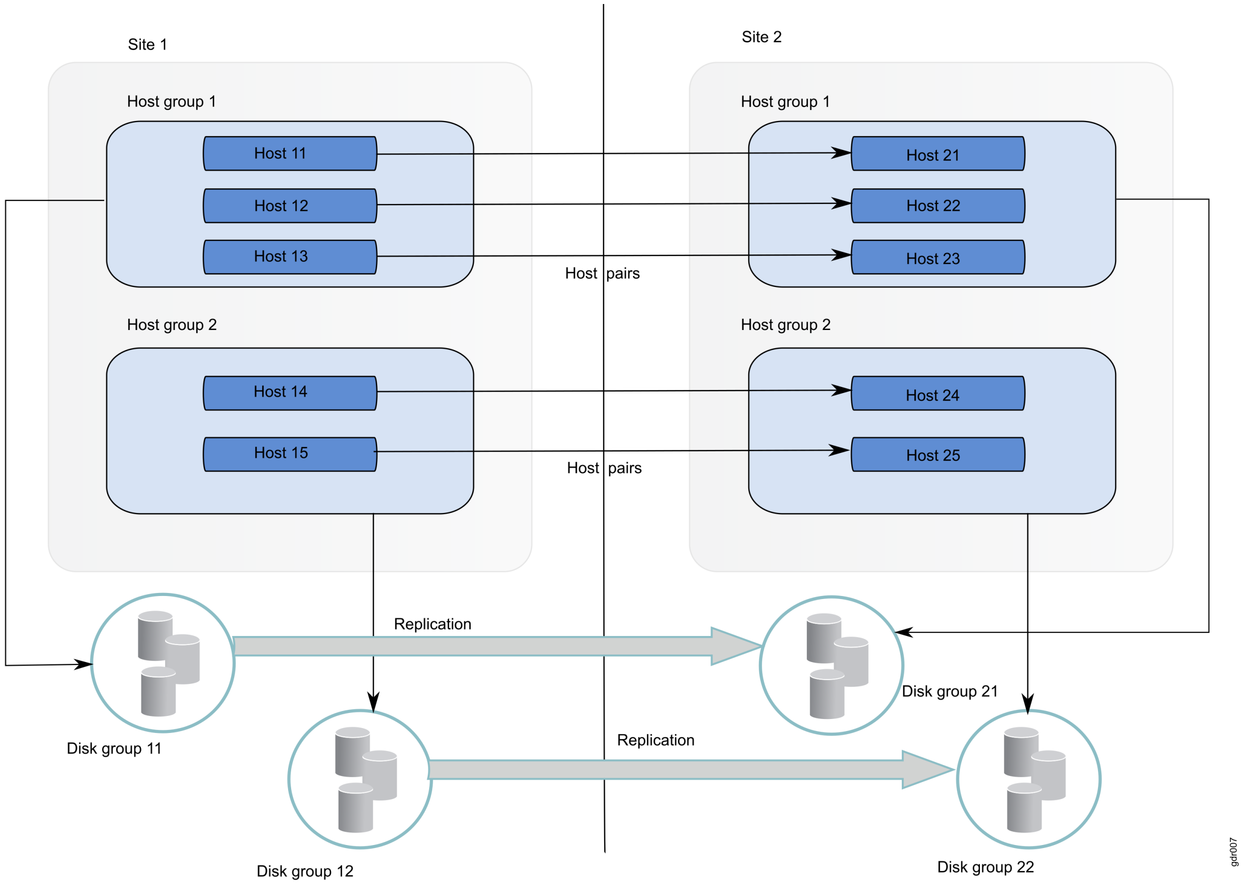

*Figure 7. Host configuration*

## Virtual machines

Virtual machines, also known as logical partitions (LPARs), are associated with specific Virtual I/O Server (VIOS) partitions for a virtualized storage to run a workload. A host can contain multiple virtual machines.

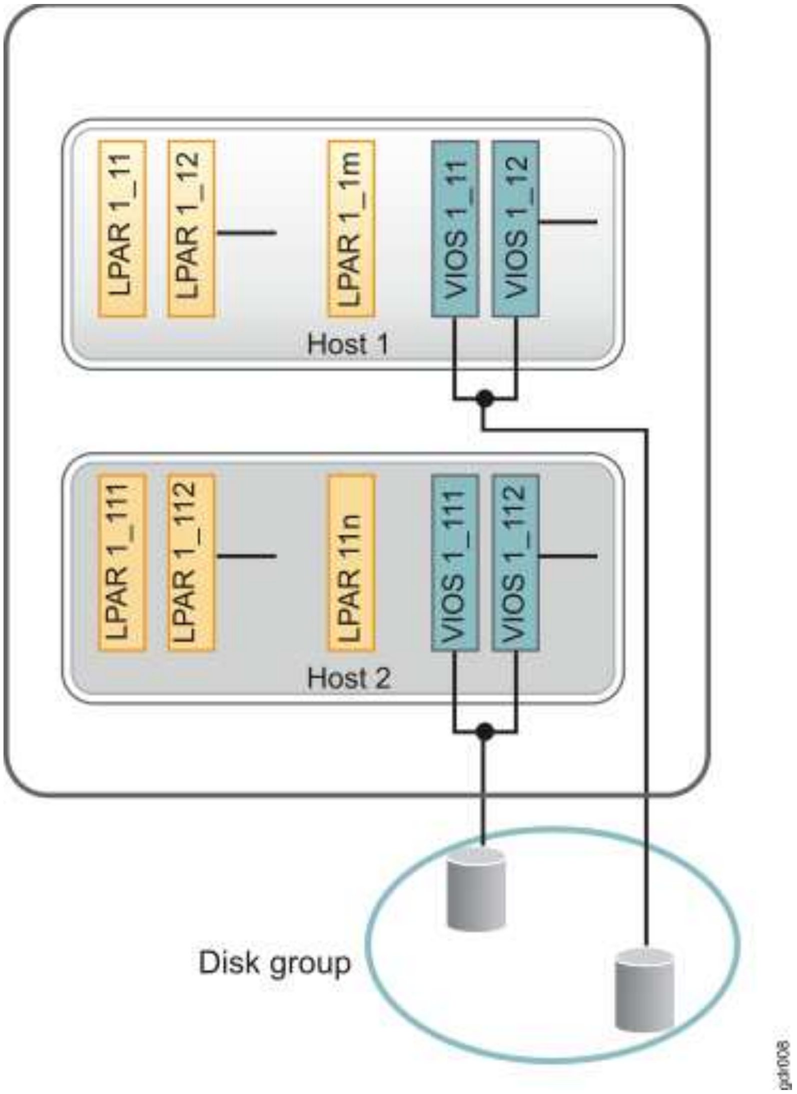

*Figure 8. Virtual machine configuration*

## Virtual I/O Server (VIOS)

The KSYS receives information about the storage configuration of the virtual machines from VIOS.

The storage area network (SAN) zoning and logical unit number (LUN) masking must be performed so that the VIOS can access the disks information. The KSYS interacts with the VIOS to obtain information about the disks that are provisioned to the client partitions. During the validation process, the data that is collected from the active site is used on the backup site to validate whether virtual machines can be moved to the backup site during disaster situations.

The KSYS interacts with the VIOS to get information about the LPAR storage. The KSYS also interacts with the VIOS when the KSYS starts the various virtual machines during the disaster recovery operations.

Therefore, the VIOS must have sufficient capacity to handle requests from the KSYS apart from handling the regular I/O activities that are in progress on the VIOS. If your environment contains multiple virtual machines that are clients of VIOS pairs, you might dedicate additional capacity during the disaster recovery operation. For example, you can have at least 0.1 CPU and 1 GB of memory in addition to the planned capacity for the VIOS.

You can back up and restore the virtual and logical configuration by using the **viosbr** command on the VIOS partitions. Also, you can collect, apply, and verify device settings in a VIOS run-time environment by using the VIOS rules management. The VIOS rules support consistent device settings on multiple VIOS partitions, and also improves usability of VIOS.

**Related information**:

viosbr command

VIOS rules management

## Storage agents

A disaster recovery solution requires an organized storage management because storage is a vital entity in any data center. The GDR solution relies on data replication from the active site to the backup site.

In the GDR solution, the data is replicated from the active site to the backup site by using storage replication. Depending on the type of storage devices in your environment, the initial storage configuration might involve installation of the storage controller software that is used to perform replication operations. The general storage management operations include starting, stopping, suspending, reversing, resyncing, pausing, and resuming the replication. For more information about installation of storage controller software, see the documentation from the storage vendor.

The following figure shows an example of storage device configuration in a site.

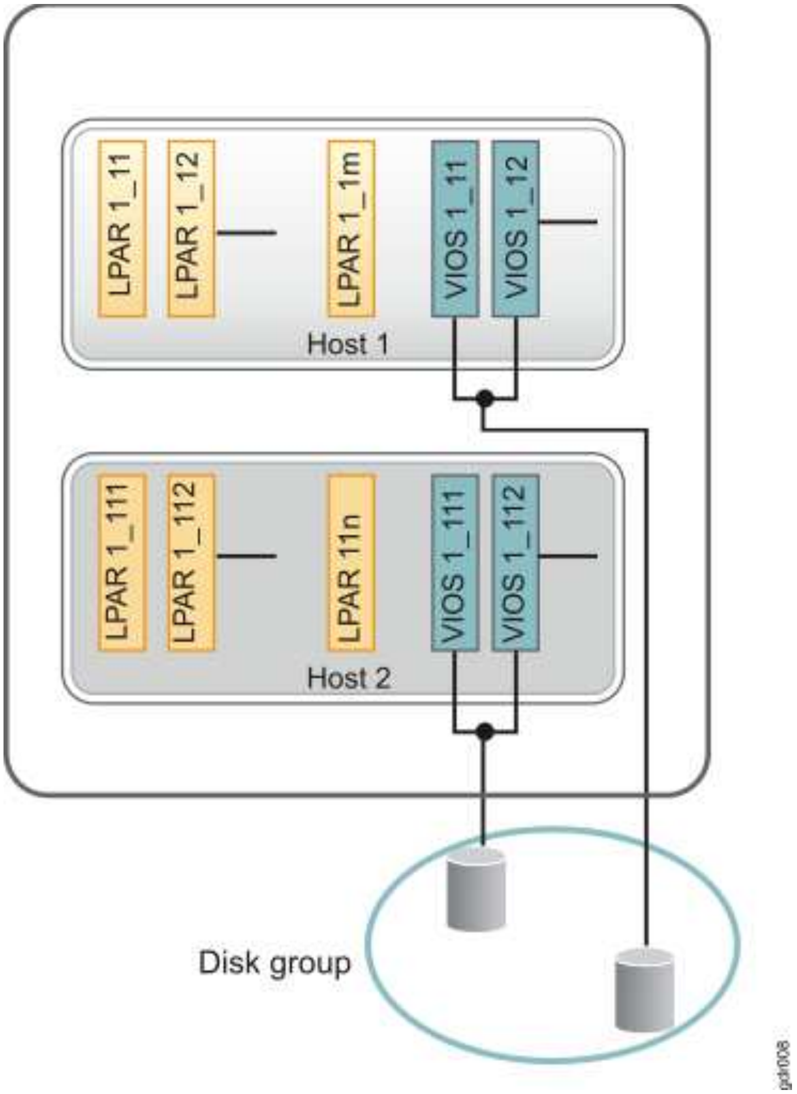

*Figure 9. Storage device configuration in a site*

The following figure shows an example of mapping storage devices across the active site and the backup site.

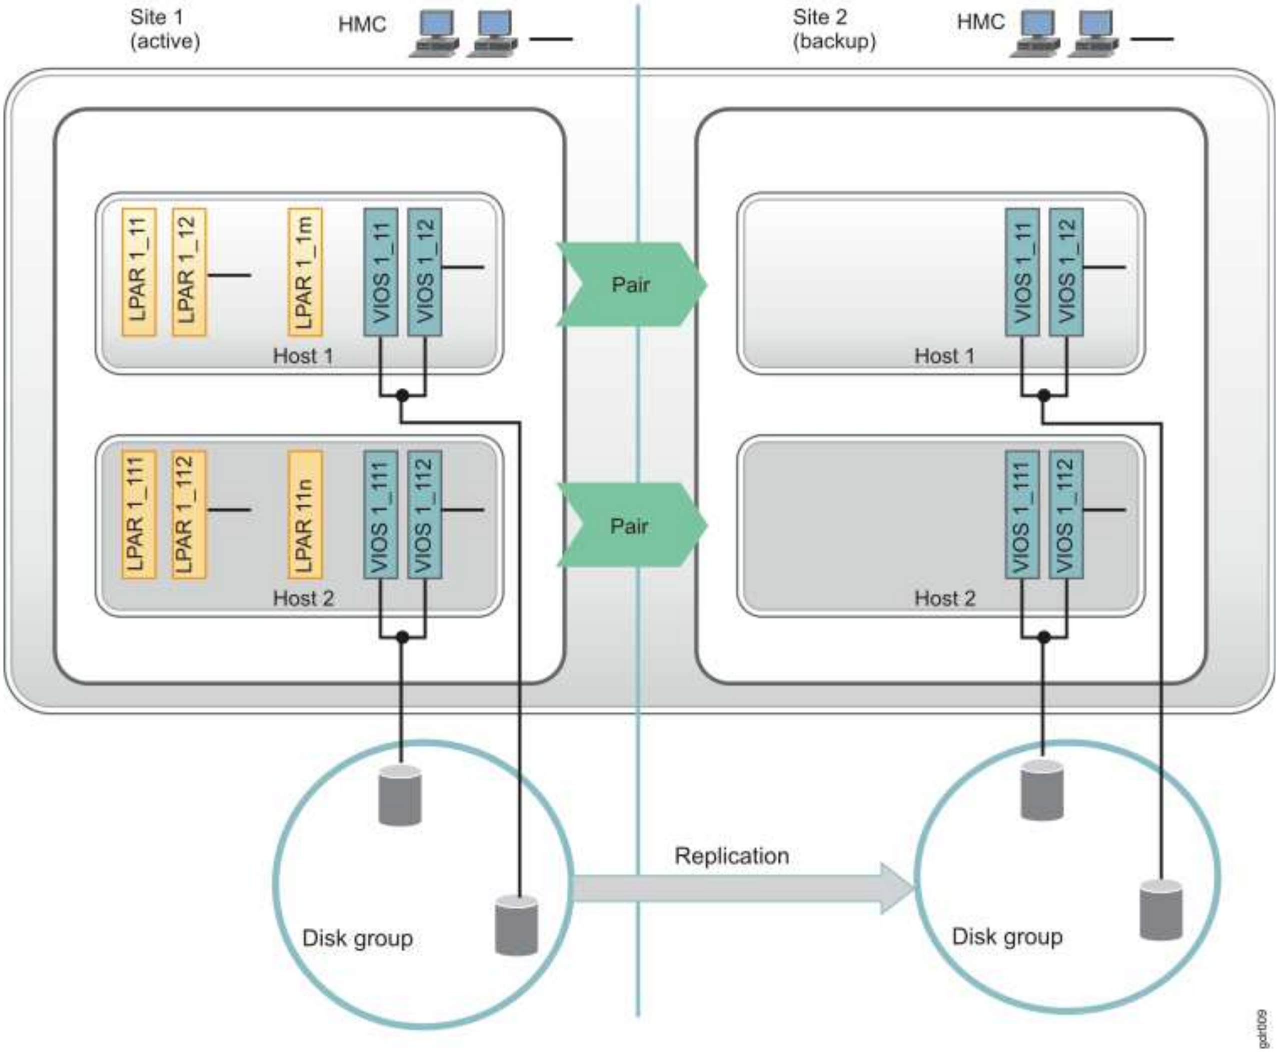

*Figure 10. Storage replication across sites*

After the initial storage setup, you must add a *storage agent* to the KSYS configuration that interacts with the storage subsystem. If you have a storage controller software, the storage agents interact with the storage controller to perform storage-specific operations in the disks.

The storage subsystem uses the following components for configuration and recovery operations:

**Storage controller**
A node that contains the software to manage the interaction between the disks and the hosts that are connected to the storage.

**Disk group**
Indicates a group of disks within a site.

**Consistency group**
Indicates a group of storage devices from multiple storage arrays to maintain the consistency of data. For more information about consistency groups, see Consistency groups.

**Disk pair**
Indicates the set of disks or disk groups that are paired across the sites for disaster recovery.

## EMC storage subsystem

The GDR solution implements disaster recovery with Symmetrix Remote Data Facility (SRDF) storage devices by using the EMC supplied Symmetrix command-line interface (SYMCLI). The GDR solution uses the Symmetrix application program interface (SYMAPI) server that runs on the EMC Solution Enabler server node for the SYMCLI operations.

To enable disaster recovery of EMC storage devices, complete the following steps before you implement the GDR solution:

1. Plan the storage deployment and mirroring functions that are necessary for your environment. This step is related to the applications and middleware that are deployed in the environment.
2. Use the EMC tools to configure and deploy the storage devices.
3. Use the SYMCLI interface to discover the storage devices that are deployed.

   **Note:** Complete step 1 and step 2 before you configure the KSYS for disaster recovery management of EMC storage devices. You can perform step 3 after the GDR implementation.

All EMC SRDF operations in the GDR solution are performed on a composite group, which is a group of disks that belong to multiple storage arrays. The consistency group that is enabled in the EMC storage devices for consistency is known as the *composite group*. The composite groups operate simultaneously to preserve the integrity and consistency of the dependent write operation of a data that is distributed across multiple arrays. Consistency for an SRDF replicated resource is maintained at the composite group level on the EMC storage device.

The GDR solution supports SRDF replicated resources in the following modes:

**SRDF/S (synchronous) replication**

In the synchronous mode, when the host issues a write operation to the source of the composite group, the EMC storage device responds to the host after the target EMC storage device acknowledges that it has received and checked the data.

**SRDF/A (asynchronous) replication**

In the asynchronous mode, the EMC storage device provides dependent write consistent, point-in-time data to the secondary storage devices that slightly lags in time from the primary storage devices. Asynchronous mode is managed in sessions. In the asynchronous mode, the data is transferred in predefined cycles (delta sets). You can change this default cycle time to change the time difference of dependent write operations on the secondary storage devices that suits your business requirement.

During the discovery phase, the storage subsystem uses the following process:

1. The KSYS interacts with the HMC to get the list of virtual machines and the corresponding VIOS information. The KSYS then interacts with the VIOS to fetch the storage disk information used by these virtual machines.
2. The KSYS verifies the disk pair to check whether the disk is set up for mirroring.
3. The KSYS checks whether the disks on the storage subsystem are part of any existing composite groups. If the disk is part of any existing composite group, the discovery operation fails.
4. The KSYS creates a composite group for each site that contains the corresponding disks and enables data consistency. The KSYS uses this composite group to change the replication direction of the corresponding disks during a planned recovery operation.

For the SRDF-capable storage subsystems of the EMC VMAX family, at least one SYMAPI client must be installed in each site on any virtual machine in the POWER® processor-based servers or any other systems. For more information about the installation of these components, see the documentation from the storage vendor.

In the GDR solution, the storage agents, which are added to the KSYS configuration, interact with the corresponding storage controllers in each site. The storage controller is called as the EMC Solution Enabler in the EMC configuration. The following figure shows the EMC Solution Enabler configuration for the GDR solution:

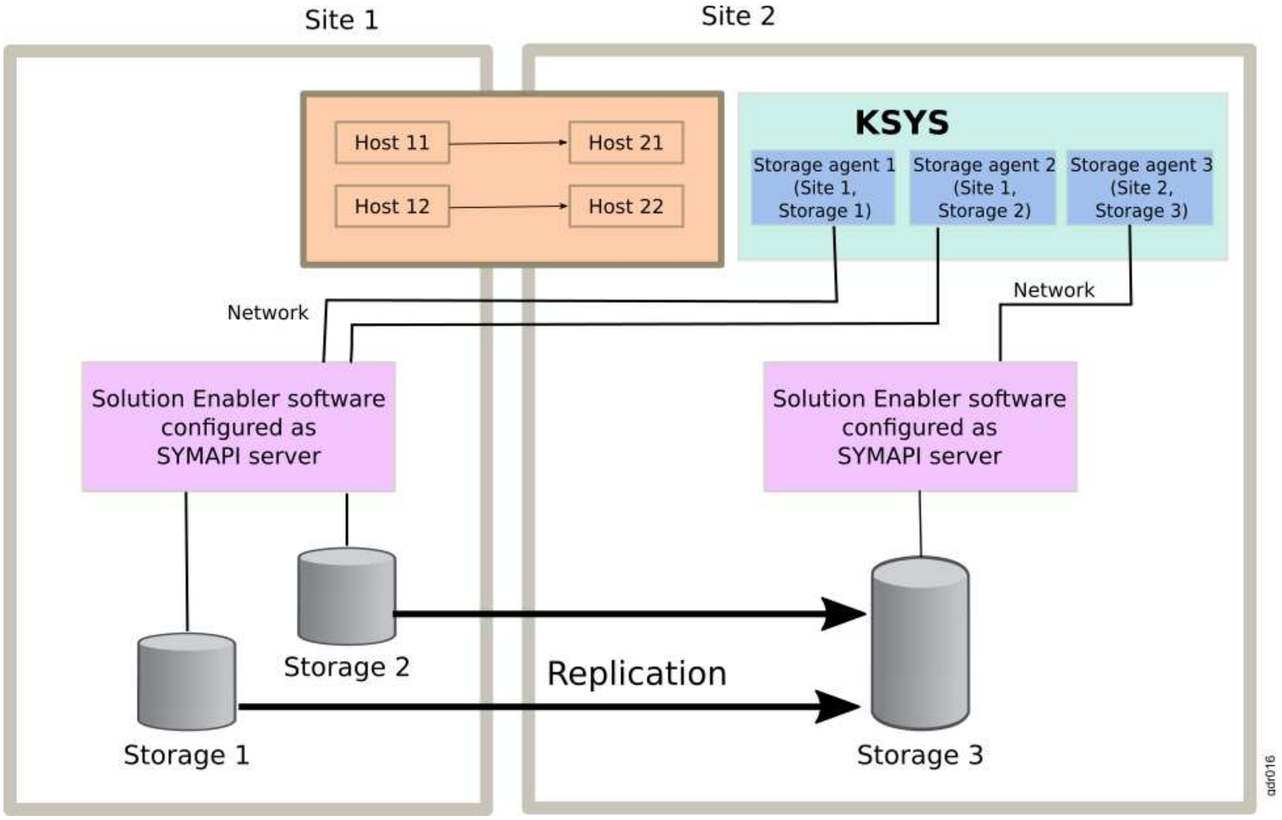
Site 1
Site 2
Host 11
Host 21
Host 12
Host 22
KSYS
Storage agent 1 (Site 1, Storage 1)
Storage agent 2 (Site 1, Storage 2)
Storage agent 3 (Site 2, Storage 3)
Network
Network
Solution Enabler software configured as SYMAPI server
Solution Enabler software configured as SYMAPI server
Storage 2
Replication
Storage 1
Storage 3
gdr016

*Figure 11. Relationship between storage agents and storage disks*

**Related information**:

EMC Solution Enabler CLI User Guide

## SAN Volume Controller system and Storwize system

Beginning with GDR version 1.1 Service Pack 1, the GDR solution supports the IBM® SAN Volume Controller (SVC) storage system and IBM Storwize® storage system. Both SVC system and Storwize storage systems follow the same host programming model for mirror management. Hence, the GDR solution supports both these storage subsystems by using the same configuration and management interfaces.

The SAN Volume Controller storage system integrates hardware and software components to control the mapping of storage data into volumes in a SAN environment. The SVC storage system includes the rack-mounted hardware components called *nodes*. Nodes are always installed in pairs. Each node pair is called an *I/O group*. A single node pair handles the I/O requests on a specific volume. A clustered system contains 1 - 4 I/O groups. A clustered system can have a maximum of eight nodes. When you set up the SVC system across multiple sites, the configuration is called *stretched system*. When you configure Metro Mirror or Global Mirror mode for data replication across sites, each site must have a separate SVC cluster.

A storage array consists of various physical drives that are logically grouped into Redundant Arrays of Independent Disks (RAID). Each storage system manages a set of logical unit numbers (LUNs) that correspond to storage arrays. The LUNs are mapped to the SVC system as groups of managed disks, which are assigned to a pool of virtual storage. The I/O groups convert the managed disks into storage pools, and then the storage pools are converted into one or more volumes. The volumes are assigned to hosts.

For more information about installation of the SVC system and Storwize system components, see the documentation from the storage vendor.

**Notes:**

- The GDR solution supports IBM SVC version 6.1.0, and later, and IBM Storwize V7000 7.1.0, and later.
- The GDR solution supports only a single SVC or Storwize cluster per host group because the SVC and Storwize storage systems do not support consistency groups across multiple storage systems within a host group. The SVC or Storwize cluster can contain a maximum of 8 nodes.
- Both the Metro Mirror (synchronous) and the Global Mirror (asynchronous) modes of data migration are supported.
- Each host that is associated with the SVC system and the Storwize system must have an installation of VIOS 2.2.5.20, or later. Additionally, the SVC or Storwize storage system requires HMC Version 8 Release 8.6.0 Service Pack 1.
- When you configure the SVC storage system by using the **mkrcrelationship** command, do not use the **-sync** option if the virtual machines are deployed in the active host after the synchronization relationship is created between disk volumes.
- When you add a virtual machine to be managed by the KSYS subsystem, you must ensure that the new virtual machine is associated with the storage volumes that have the same relationship type as the other existing virtual machines. That is, the relationship between the master disks and the auxiliary disks of the new virtual machine must be created to match the relationship between the master disks and auxiliary disks of other virtual machines. You can check the existing relationship information by using the **lsrcrelationship** command in the SVC storage system. After a relationship is created between the master disks and auxiliary disks, the storage administrator must start the data consistency of the relationship before you add the virtual machine to be managed by the KSYS subsystem.

**Note:** When you use the **-sync** option, the SVC storage system does not sync the backup site disks with the active site disks completely. Therefore, a disaster recovery operation from the active site to the backup site might fail. Also, the KSYS subsystem cannot check whether you used the **-sync** option. Hence, you must not use the **-sync** option to prevent unpredictable results.

The GDR solution manages disaster recovery across sites by grouping all the disks in a site and establishing a consistency group for the disks in the SVC or Storwize storage systems. During the discovery phase, the KSYS subsystem interacts with the HMC and the VIOS to fetch the storage disk information that is used by the virtual machines. The KSYS then creates a consistency group for each site and enables data consistency across the sites. The KSYS subsystem uses this consistency group to change the replication direction during a planned recovery operation. The KSYS subsystem also checks whether the disks on the storage subsystem are part of any existing consistency groups. If the disk is part of any existing consistency group, the discovery operation fails. In this case, you must ensure that the storage disks are not part of any consistency groups and then run the discovery operation again.

The KSYS subsystem uses a Secure Shell (SSH) application to communicate to the SVC or Storwize command line interface (CLI). You must configure this connection such that the authentication is established without entering the password so that the KSYS subsystem can communicate to the SVC or Storwize storage system without any intervention.

In the GDR solution, the storage agents, which are added to the KSYS configuration, interact with the corresponding storage devices in each site. The following figure shows the storage agent configuration for the SVC storage in the GDR solution.

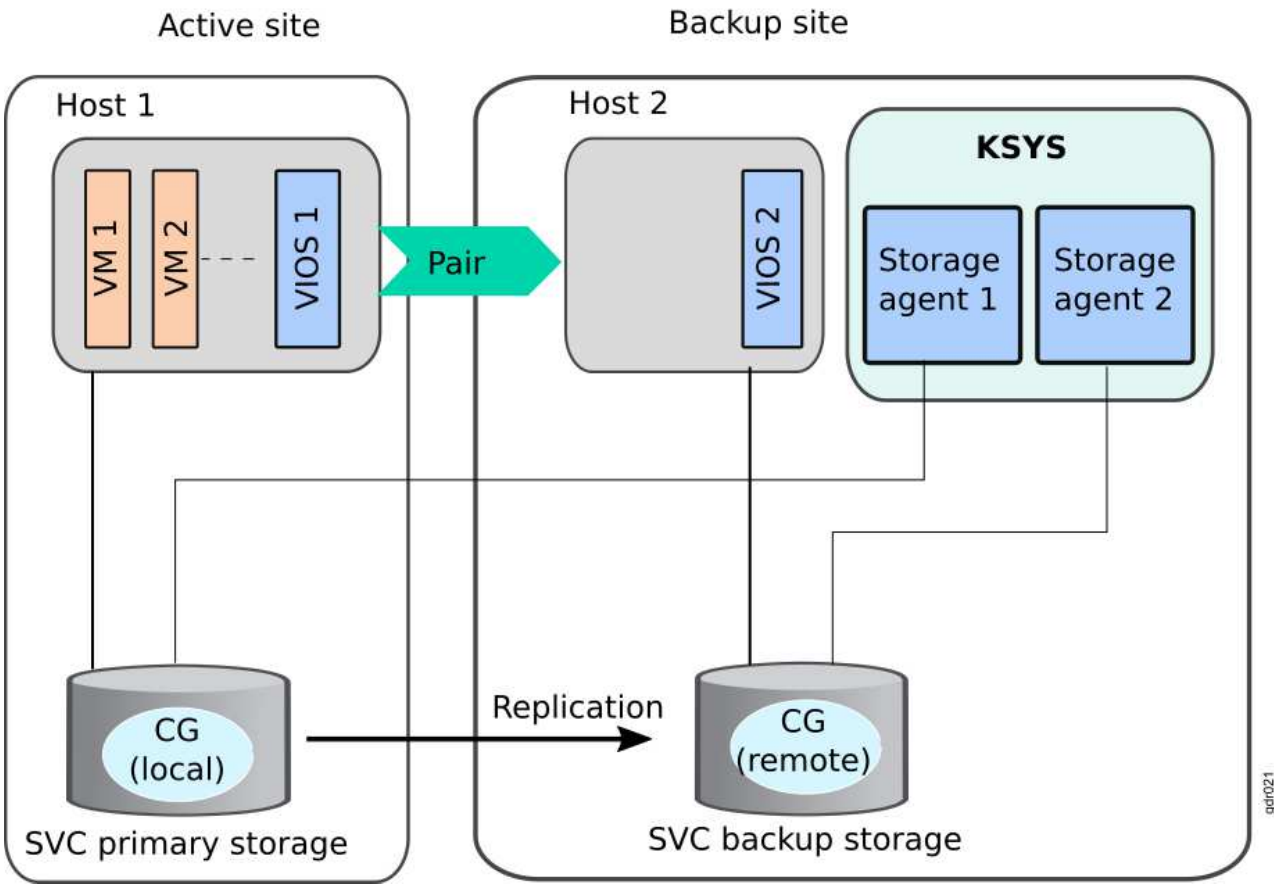

*Figure 12. Relationship between storage agents and SAN Volume Controller storage systems*

**Related information**:

IBM SAN Volume Controller

IBM Storwize V7000

mkrcrelationship command

lsrcrelationship command

Setting up an SSH client in SVC storage system

## DS8000 storage system

Beginning with GDR version 1.1 Service Pack 1, the GDR solution supports the IBM System Storage® DS8000® series storage system for storage management.

IBM System Storage DS8000 series is a high-performance and high-capacity series of disk storage that supports continuous operations. The DS8000 series that are supported by the GDR solution include the following models:

**DS8700**

The DS8700 storage system uses IBM POWER6® server technology, which provides high-performance and high-capacity for disk storage systems. The DS8700 storage system also provides improved performance by upgrading the processor and I/O interconnection technology.

**DS8800**

The DS8800 storage system is an advanced high-performance and high-capacity disk storage system. The DS8800 storage system supports IBM POWER6+ processor technology to provide higher performance.

All DS8000 series models consist of a storage unit and one or two management consoles depending on high availability configuration settings. You can use the graphical user interface (GUI) or the command line interface (CLI) to logically divide storage and use the built-in Copy Services functions.

**Notes:**

- The GDR solution supports only Global Mirror (asynchronous) mode of data replication across sites.
- The GDR solution supports only a single DS8000 storage system per site, which means that all storage devices must be configured as a single DS8000 cluster per site.
- Each host that is associated with the DS8000 series of storage devices must have VIOS 2.2.5.20, or later, installed. Additionally, the DS8000 storage system requires HMC Version 8 Release 8.6.0 Service Pack 1.

A *session* is a collection of disk volumes across multiple storage systems that are managed together to create consistent copies of data. All storage operations in the GDR solution are performed on a session ID that is a logical group of disks that are associated with the DS8000 series of storage array. During the discovery phase, the KSYS interacts with the HMC and the VIOS to fetch the storage disk information that is used by the virtual machines. The KSYS then creates a session in each site, adds all the disks to this session. The KSYS uses sessions to change the replication direction during a planned or unplanned recovery operation.

The GDR solution uses the DS8000 Series Command Line Interface (DSCLI) client to interact with the DS8000 storage array.

The following restrictions and limitations apply to the Global Mirror session in the DS8000 series of storage devices:

- If the disk volumes that are replicated span across multiple DS8700 Logical Sub Systems (LSS), the same Global Mirror session identifier must be used for all LSS.
- The Hardware Management Console (HMC) associated with the IBM DS8700 storage system that is used with the GDR solution, must be accessible by using TCP/IP on the KSYS node and both sites.
- Latest DSCLI client must be installed on the KSYS node. The GDR solution supports DSCLI version 7.7.51.48, and later.

For more information about the installation and configuration of the DS8000 series of storage devices, see the documentation from the storage vendor.

**Related information**:

IBM DS8000 Series

IBM DS8000 series migration information

## Hitachi storage systems

GDR Version 1.2, or later, supports disaster recovery (DR) for third-party vendor storage from Hitachi. Hitachi storage system supports asynchronous data replication for long-distance through Hitachi Universal Replicator (HUR) technology.

**Notes:**

- GDR 1.2 supports the following version of the Hitachi storage system: Hitachi Virtual Storage Platform (VSP) G1000 and Hitachi VSP G400.
- GDR 1.2 supports only asynchronous mode of storage replication.
- The DR failover rehearsal function is not supported on Hitachi storage systems.

To successfully use the HUR technology, you must correctly plan for its implementation. Before you continue, your environment must meet the following requirements:

- The GDR sites and the KSYS subsystem must be configured.
- HUR support must be configured.

The setup of GDR disaster recovery for Hitachi mirrored storage system involves the following steps:

1. Plan the storage deployment and replication necessary for your environment. This process is related to the applications and middleware that must be deployed in the environment that you want to include in the recovery management by the GDR solution.
2. Use the storage configuration tools that are provided by Hitachi to configure the storage devices you have defined in step 1, and deploy the devices.
3. Use the **ksysmgr** interface to discover the deployed storage devices and to define the disaster recovery policies for the applications or virtual machines that are using the mirrored storage.

Hitachi HUR storage management uses Command Control Interface (CCI) operations from the AIX operating system and GDR environment. The KSYS subsystem uses CCI to discover and integrate the Hitachi storage replication into the disaster recovery framework of GDR. This integration manages disaster recovery for applications that use the mirrored storage. In the GDR solution, the storage agents, which are added to the KSYS configuration settings, interact with the corresponding CCI software in each site.

**Hitachi Command Control Interface (CCI)**

The remote and in-system replication software from Hitachi requires the following CCI components to manage the disk pairs:

**Command device**

The command devices are located in storage systems. CCI uses the command device as the interface to the storage system from the host. The command device accepts commands from the host and executes them on the storage system. The command device is a dedicated logical volume.

**Hitachi Open Remote Copy Manager (HORCM)**

The HORCM is located in the CCI server. The HORCM operates as a daemon process. When activated, the HORCM refers to CCI configuration definition files that are also located on the CCI server. The HORCM instance communicates with the storage system and remote servers.

HORCM definition files describe the storage systems, pair volumes, and data paths. When you run a command, CCI uses the information in the HORCM files to identify target volumes of the command.

Two HORCM files are required for each pair. One file describes the primary volumes (P-VOLs), and the other file describes the secondary volumes (S-VOLs). The following figure shows an example setup of Hitachi storage configuration with two HORCM instances.

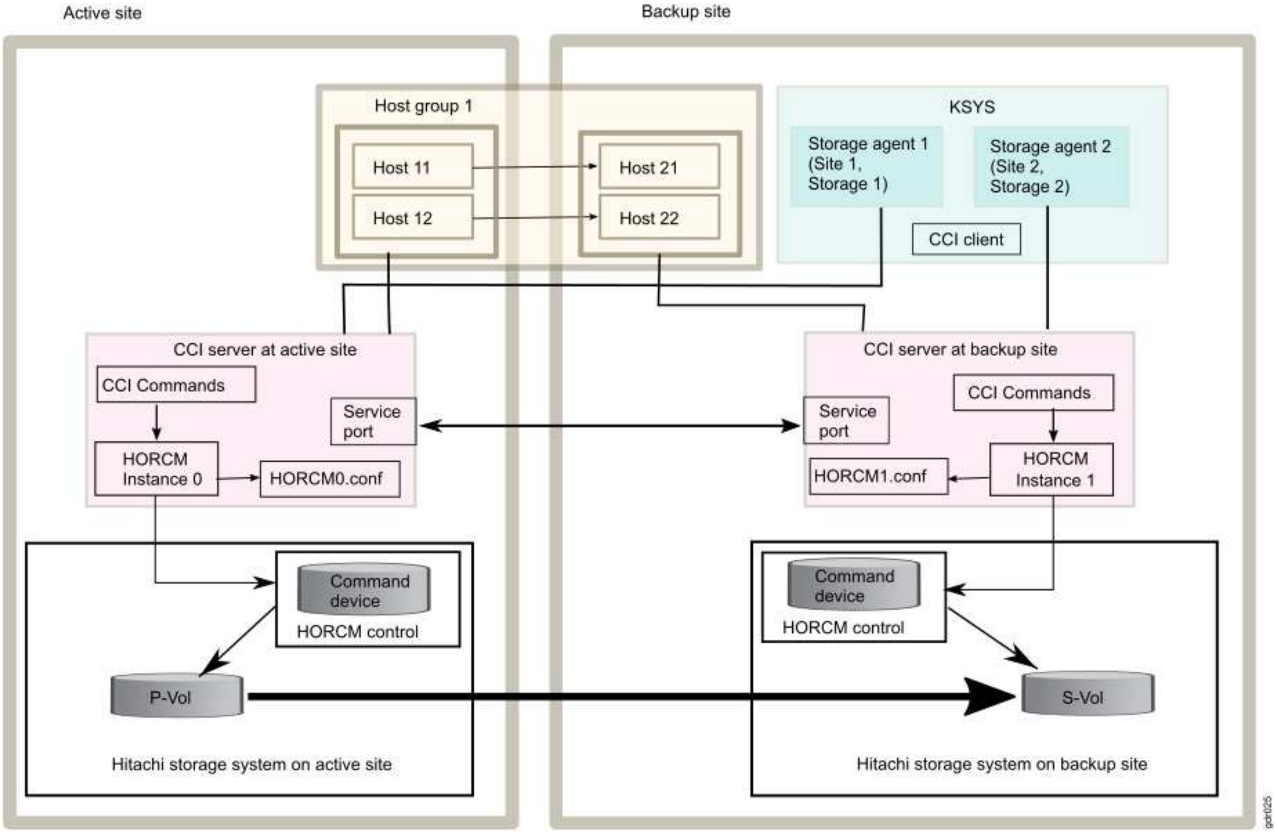

*Figure 13. Relationship between storage agents and storage disks*

### Requirements

The GDR solution requires the following components for Hitachi storage configuration:

- Command Control Interface (CCI) software for AIX. For more information about CCI, see Hitachi Command Control Interface (CCI) documentation that is maintained by Hitachi. The CCI software must also be installed on the KSYS LPAR to communicate with the CCI server.
- The GDR solution requires CCI version 01-39-03/04 with model RAID-Manager/AIX.

### Limitations

Consider the following limitations when you use the Hitachi storage systems in the GDR solution:

- Only Hitachi Universal Replicator (HUR) technology is supported for asynchronous mirroring. The synchronous mirroring is not supported.
- The device names (`dev_name` attributes) must map to logical devices and the device groups (`dev_group` attributes) must be defined under the `HORCM_LDEV` section in the `horcm.conf` file.
- The GDR solution uses the `dev_group` attribute for any basic operation that is related to the Hitachi storage systems. Some examples of basic operations are **pairresync**, **pairevtwait**, and **horctakeover**. If several device names exist in a device group, the device group must be consistency-enabled.
- For command devices, the **Command Device Security** option in the **Command Device Attributes** panel must be disabled. If the **Command Device Security** option is enabled, the required information, such as logical device (LDEV) ID, journal volume ID, consistency group ID, and volume type, are not displayed in any command output. The KSYS subsystem needs these information to monitor the state of storage subsystem disks.
- The command devices can be protected by using user authentication. However, in the GDR solution, best practice is to disable user authentication for all the Hitachi disks, including the disks designated as

command device, that are managed by the GDR solution. Otherwise, many of the disk management functions of the KSYS subsystem might report incorrect information and indeterminate behavior.

- The GDR solution does not trap Simple Network Management Protocol (SNMP) notification events for HUR storage. If an HUR link goes down when the hosts are functional and later the HUR link is repaired, you must manually resynchronize the disk pairs.
- The KSYS subsystem does not control the creation of disk pairs. You must create the disk pairs before you start the KSYS partition.
- Dynamic manipulation of disk groups is not supported for Hitachi storage systems. The KSYS subsystem might break the grouping of the disks in a disk group during the group manipulation operation. When a disk is removed from a disk group, the disk moves into simplex state.

  The following cases result in disk removal from a disk group. In these cases, the Hitachi storage subsystem removes the pair relationship.

  - Removing disks from a virtual machine
  - Performing Live Partition Mobility operation or Remote Restart operation across host groups
  - Restoring a snapshot
  - Unmanaging a virtual machine from the KSYS subsystem
  - Removing a host from host group and adding the host to another host group
  - Removing a cluster

  You must recreate the disk pair explicitly before you add the same disk or virtual machine to the GDR management.
- Hosts that are managed by the GDR solution cannot contain volume groups with both HUR-protected and non-HUR-protected disks. A host must contain an HUR-protected disk.
- All hosts within a site that are managed by the GDR solution must use the same HORCM instance.
- All disks within a site must belong to the same journal volume.

**Related information**:

Hitachi Universal Replicator User Guide

## Network

The network in your environment must already be configured for the existing resources that include hosts, HMCs, VIOS, and storage devices. The GDR solution requires that the KSYS node must be directly connected to the HMCs and the storage controllers in both the sites. The KSYS uses the HMC to interact with the hosts and VIOS, and the KSYS uses the storage controller to interact with the storage subsystem.

You can modify the KSYS system properties to enable or disable the network mapping function globally across the sites. By using the network mapping function, you can create VLAN ID or virtual switch mapping policies that contain a mapping of VLAN IDs or virtual switches that are assigned to the virtual machines when the virtual machines are moved from the active site to the backup site.

The recovery operation is not easy with the default settings that are used in the AIX operating system because of following issues:

*Table 2. AIX default setting issues for recovery of the virtual machines from the active site to the backup site*

| **Issue** | **Resolution** | **Related information** |
|---|---|---|
| Data, which includes the AIX root volume group, is replicated from the active site to the backup site. The AIX root volume group contains information about the disk, adapters, network, and other system-specific items. The data also includes the storage area network (SAN) World Wide Names (WWNs) that identify each SAN disk. Although the target disks are exact copies of the source disks, the disks are different. Therefore, the AIX operating system creates new devices when it starts. For example, if the source virtual machine has 10 devices, `hdisk0` - `hdisk9`, when the virtual machine is moved to the backup site, the `hdisk0` - `hdisk9` devices are marked as `defined` and the `hdisk10` - `hdisk19` storage devices are marked as `available`. In this case, the original disks, `hdisk0` - `hdisk9`, are called as `ghost` disks. These renamed disks might cause issue, especially when some scripts or tools depend on the original disk names. | On the source LPAR, set the AIX **`ghostdev`** parameter to `yes` to clean up the `ghost` devices. | For more information about setting this parameter, see "Configuring GDR" on page 45. |
| The naming convention of the disks that are used by the recovered virtual machine (`hdisk0`, `hdisk1`, and so on) might not match with the names of disks of the source virtual machine for various reasons. For example, the `hdisk5` disk in the source site might be displayed as `hdisk7` on the target site. This mismatch can cause difficulties in a disaster situation because the disks that were part of that volume group must be known to import any of the AIX volume groups. | On the recovered virtual machine in the backup site, run disaster recovery scripts that address the following issues:<br>• Set the local host name<br>• Set the network IP addresses<br>• Customize the host bus adapters (HBAs)<br>• Import the volume groups | For more information about the disaster recovery scripts, see "Running the disaster recovery scripts" on page 100. |

# Planning GDR

To implement the GDR solution, you must review your current disaster recovery plan and consider how the GDR solution can be integrated into your current environment. The GDR solution can coexist with some of the existing product offerings with a few exceptions. Therefore, planning the implementation prevents issues in the configuration and disaster recovery operations.

## Requirements for the GDR solution

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

Consider the following requirements for the GDR solution:

### Software requirements

- The KSYS logical partition must be running IBM AIX 7.2 with Technology Level 1 Service Pack 1 (7200-01-01), or later.

  If you are using AIX Version 7.2.0, or earlier in the KSYS logical partition, you must add the following class in the `/usr/sbin/rsct/cfg/ct_class_ids` file:

  ```
  IBM.VMR_HG   521
  ```
- You must install the OpenSSL software version 1.0.1.516, or later for the AIX operating system. You can download the OpenSSL software from https://www.ibm.com/marketing/iwm/iwm/web/reg/download.do?source=aixbp&lang=en_US&S_PKG=openssl&cp=UTF-8&dlmethod=http.

  **Note:** The latest version of the OpenSSL software is also included on the AIX base media.
- The GDR solution requires HMC Version 8 Release 8.7.1, or later.
- The GDR solution requires VIOS Version 2.2.6.00, or later, with all the subsequent patches.
- Each LPAR in the host must have one of the following operating systems:
  - AIX Version 6.1, and later
  - PowerLinux™
    - Red Hat Enterprise Linux (little endian, big endian) Version 7.2, or later
    - SUSE Linux Enterprise Server Version 12.1, or later
    - Ubuntu Linux distribution Version 16.04
  - IBM i Version 7.2, and later

    **Note:** The IBM i operating system is supported only with VIOS Version 2.2.5.20, or later.

### Configuration requirements

- You can configure only 2 sites: an active site and a backup site. These sites can be separated by unlimited distance.
- You can configure only 1 KSYS logical partition (LPAR). The KSYS LPAR must be placed in the backup site.
- 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.
- The virtual machines, which must be recovered during disaster situation, must be running on POWER7® processor-based server, or later, that are PowerVM-based systems managed by HMCs.
- The GDR solution supports the following storage devices:

© Copyright IBM Corp. 2017, 2019

**EMC storage system**
The GDR solution supports storage devices for the EMC VMAX family (VMAX1, VMAX2, and VMAX3). The EMC storage devices must be Symmetrix Remote Data Facility (SRDF)-capable. The EMC storage must have Solutions Enabler SRDF family Version 8.1.0.0 installed. Both SRDF/S (Synchronous) and SRDF/A (Asynchronous) replication modes are supported.

**IBM SVC and Storwize storage systems**
The GDR solution supports IBM SVC Version 6.1.0, and later, and IBM Storwize V7000 7.1.0, and later. Both Metro Mirror (synchronous) and Global Mirror (asynchronous) modes of data replication are supported across sites.

**IBM System Storage DS8000 series**
The GDR solution supports DS8700 or later DS8000 storage systems with DSCLI version 7.7.51.48, and later. Only Global Mirror (asynchronous) mode of data replication is supported across sites.

**Hitachi storage systems**
The GDR solution supports the Hitachi Virtual Storage Platform (VSP) G1000 and Hitachi VSP G400 with CCI version 01-39-03/04 and model RAID-Manager/AIX. Only asynchronous mode of storage replication is supported.

**Note:** The SVC, Storwize, and DS8000 storage systems are supported only with VIOS Version 2.2.5.20, or later and HMC Version 8 Release 8.6.0 Service Pack 1, or later. The Hitachi storage systems are supported with VIOS Version 2.2.6.00, or later and HMC Version 9 Release 9.1.0, or later.

## Network requirements

- All virtual machines (VMs) that are managed by the GDR solution must use virtual I/O resources through VIOS. The VMs must not be connected to a physical network adapter or any dedicated devices.
- The VIOS must have a Shared Ethernet Adapter (SEA) configuration to bridge to the same Ethernet network between the hosts located in the same site.
- The same virtual LAN (VLAN) must be configured across the site.
- Ensure redundant connection from the KSYS 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.

## Administrative requirements

If you are an administrator, ensure that the following prerequisites are met before you implement the GDR solution:

- The KSYS must be connected to all the HMCs across sites.
- All the VIOS partitions and disk pairs must be deployed correctly across sites.
- Storage replication must be set up correctly for various disks that are used in the GDR managed environment.
- Storage area network (SAN) connectivity and zoning must be configured so that VIOS can access the disks that are relevant to the hosts across the host pairs. For example, a disk `D1` that is connected to the VIOS of a host must have a mirror disk `D1_M` that is connected to the VIOS of the paired host in the backup site. Any connectivity issues can cause the GDR verification to fail.
- SAN connectivity must be performed such that the SAN fabrics of VIOS on the paired hosts do not connect with each other as shown in the following figure. The following figure shows such a configuration that meets the GDR requirement. The paired hosts, `Host 1` and `Host 2`, have `VIOS 1` and `VIOS 2` that have mutually exclusive SAN switches. You must have exclusive SAN switches because VIOS checks the SAN fabric that involves logging in to the connected switches. If VIOS from both sites are connected to the same switches, the login operation creates conflicts and verification of the configuration might fail.

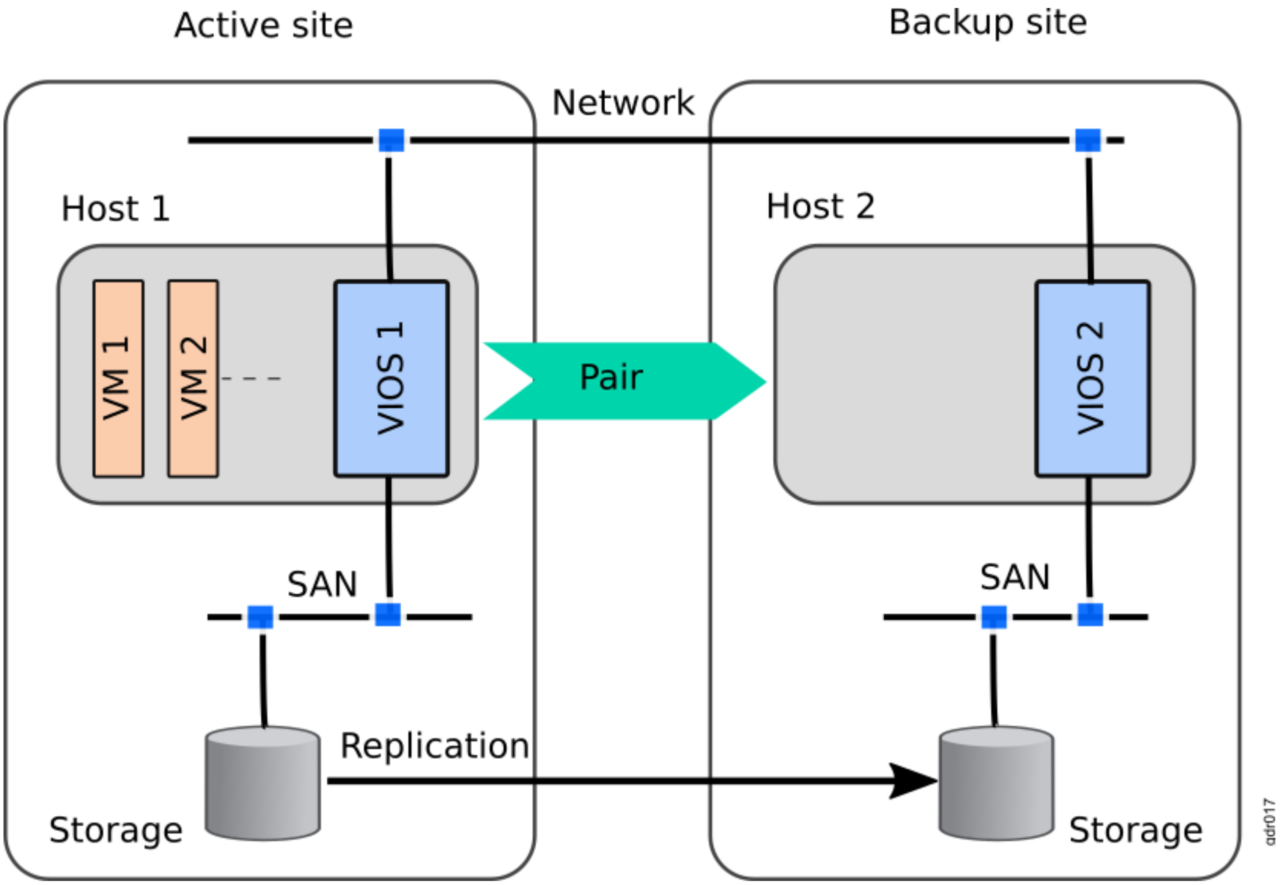

*Figure 14. SAN zoning requirement*

### Storage requirements

- Ensure that all the prerequisite software are installed on the same logical partition in which the KSYS software is installed. The prerequisite software includes the storage controller software that must be installed on the KSYS LPAR. The storage controller is the software component that you receive from the storage vendor that allows KSYS to contact storage devices and perform storage-specific operations.
- Irrespective of the type of storage, the disk size in the active and backup sites must be same.
- Ensure that the disks that belong to the virtual machines of an added host have mirror relationship with corresponding disks across the sites.
- Verify that the disks that are used for the virtual machines that are managed by the GDR solution must not be managed by any other disaster recovery solutions.
- Ensure that the storage disks do not belong to any existing composite group.
- For virtual Small Computer System Interface (vSCSI) disk mapping, ensure that you do not use Logical Volume Manager (LVM)-based disk partition management from VIOS.

## GDR restrictions

Consider the following restrictions for the GDR solution:

**SCSI reservations**

If the virtual machines in the active site are using Small Computer System Interface (SCSI)-2 or SCSI-3 persistent reservations to control access to a shared storage device (for example, if the value of the **reserve_policy** attribute is set to PR_shared in the AIX operating system), and a virtual machine is moved to the backup site, all the reservations are lost and policies are set to default reservation policies according to the storage and operating system associated with the virtual machine. It happens because the storage subsystems do not transfer reservations across mirror copies of data. In addition, the host or storage adapter identities, which are part of reservation management, also change across the active site and the backup site during the disaster recovery move operation.

**Example**: A PowerHA® SystemMirror® cluster is established across a set of virtual machines and disk fencing is enabled on the active site. In this case, PowerHA SystemMirror performs disk fencing by using SCSI-3 persistent reservations. Therefore, in the active site, each virtual machine must have stored its own key for disk fencing. When the GDR solution moves these virtual machines to the backup site, all these virtual machines start in the backup site, a cluster is established, reservation is enabled, and the keys are stored back into the storage subsystem.

**Date and time of virtual machines**

When virtual machines are moved from the active site to the backup site, the date and time in the virtual machine depends on the backup site environment in the following ways:

- For POWER7 and POWER8® processor-based systems: If time reference is set up for the system, for example, you set the reference of all virtual machines in the host based on a single source like a Virtual I/O Server (VIOS), and the move operation is performed, the virtual machines acquire the time and date information of the reference VIOS on the target host.
- For POWER8 processor-based systems: If simplified remote restart option is enabled for a virtual machine, and the virtual machine is moved to a target POWER8 processor-based system, the virtual machine retains its original time and date value from the active site.

**Change of disk names and disk queue depth values**

After the disaster recovery operation, the disk names are not sustained on the backup site. In case of replicated disks, the disk identifier changes when the virtual machines are started on the backup site and the disks are renamed based on the disk identifiers. Therefore, disk names are not sustained across sites. However, you can create customized scripts to be run after the disaster recovery operation and the disks can be renamed based on predefined information.

Similarly, as the disk identifier changes during the disaster recovery operation, the disks are added to the backup sites as new disks are plugged in. Therefore, custom queue depth values change on the backup site. The AIX operating system provides a tunable parameter in the **chdef** command to manage the queue depth as a global policy that works only for EMC storage systems and DS8000 series of storage systems. By using the **chdef** command, you can set the value of a predefined attribute for the entire AIX environment. Any new disks that are added after you run this command, can inherit the value from the specified global policy. For more information about the **chdef** command, see the **chdef** command in the IBM Knowledge Center.

**Dependency on VMR daemon**

If you stop the VMR daemon (IBM.VMR subsystem) forcefully for some reason and then start it again, the VMR daemon might take up to 2 minutes to become stable and perform any operations depending on the number of hosts and virtual machines in the KSYS configuration settings.

## Prerequisites for implementing the GDR solution

Before you start implementing the GDR solution, you must plan the resources and corresponding details for your production and backup sites. Identify the following information and have it available when you plan for the GDR implementation.

**KSYS node**

Identify the host and the logical partition in which you plan to create the KSYS node. The host must be located preferably in the backup site during the normal (non-disaster) conditions that is not managed under the KSYS subsystem. You must create the LPAR in the identified host that has IBM AIX 7.2 with Technology Level 1 Service Pack 1 (7200-01-01), or later, installed.

The KSYS node must be able to perform HTTPS-based communication to all HMCs on both sites. In addition, the KSYS node must be able to communicate to the storage subsystems on both sites by using the storage vendor methods.

**KSYS cluster**
Identify a name for your one-node KSYS cluster.

**Sites** Identify names for your active and backup sites.

**HMC** Identify the HMCs that you want to add in your active and backup sites. You can add two Hardware Management Consoles for a dual HMC configuration in your sites that ensures enhanced availability when one of the HMCs is down or unreachable.

Ensure that you have the following information available for each HMC that you plan to include in the GDR implementation:

- HMC name or IP address
- User name
- Password

**Hosts** Identify all the managed hosts in your production site that you want to add to the GDR implementation. Ensure that you plan to include the corresponding managing HMCs in the GDR solution. In addition, identify a corresponding managed host in the backup site that can be paired to each host in the active site. You must have the host name available for each host that you are planning to include in the GDR implementation.

**LPAR** Identify the LPARs that you want to include in the GDR implementation and install your applications as required. You can exclude the LPARs that you do not want to include in the GDR solution. You must have the LPAR name available for each LPAR that you are planning to include in the GDR implementation.

**Note:** The virtual machines must not be scheduled to restart automatically if the virtual machine is included in the GDR disaster recovery management.

**Virtual I/O Server (VIOS)**
The VIOS configuration in the active site hosts must ideally match across the backup site hosts that are paired together. If you have dual VIOS configuration in the active site hosts for VIOS redundancy, you should have a dual VIOS configuration in the backup site hosts that are paired. If one of the VIOS in the backup host is down and the host loses the dual VIOS configuration, you can move the VMs to the backup site when a disaster occurs. However, if you want to manually bring back the redundancy, change the LPAR profiles accordingly by using the HMC.

**Note:** For multiple VIOS configuration in virtual Small Computer System Interface (vSCSI) disk-mapping, ensure that the Virtual I/O Servers do not have any backend disk reservation.

During the verification phase, the GDR solution displays a warning message about any VIOS issues. For example, if any of the VIOS partitions is down, a warning message is displayed. The GDR solution skips this check if you use the **`lose_vios_redundancy`** attribute persistently. For more information about this option, see “Managing the system attributes” on page 85.

Even after you use this option, the source virtual machines can be moved to the backup host only when the VIOS in the target host can accommodate all the virtual adapters of the virtual machine. Problems might occur during the disaster recovery operation if one of the VIOS or Fiber Channel adapters are down. Review the following scenarios to determine the VIOS configuration issues:

**Scenario 1**
The following figure shows an N_Port ID virtualization (NPIV)-based VIOS configuration in which the source host contains 2 VIOS partitions that use 2-port Fibre Channel (FC) adapters. If the target host does not contain dual VIOS configuration or if one of the Virtual I/O Servers in the target host is not functioning, the virtual machine can be

moved from the source host to the target host only when the VIOS in the target host uses 4-port FC adapter.

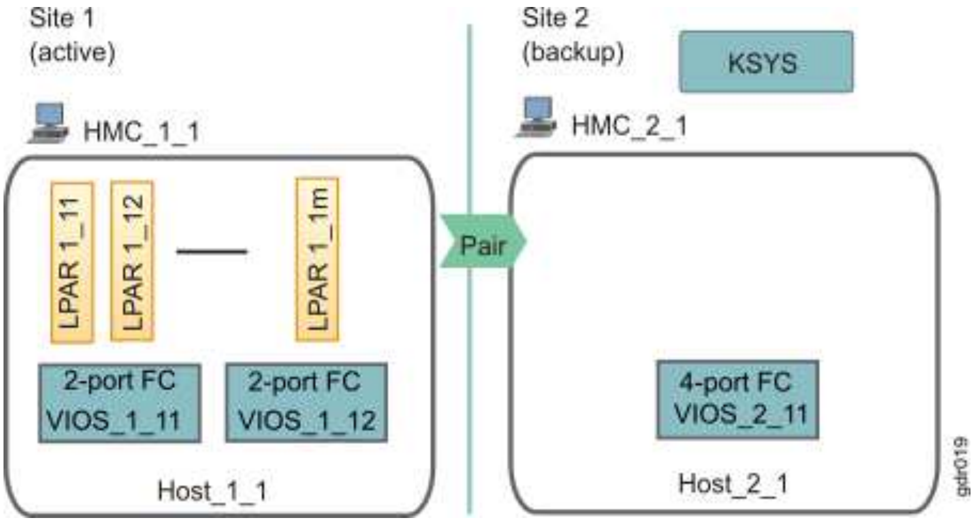

*Figure 15. VIOS configuration during disaster recovery: Scenario 1*

**Scenario 2**

The following figure shows an NPIV-based VIOS configuration in which the source host contains a VIOS partition (`VIOS_1_11`) that uses 2-port FC adapter. In this example, 70 virtual machines are running in the `Host_1_1`, where 64 VMs are mapped to the `fcs0` adapter and the remaining 6 VMs are mapped to the `fcs1` adapter.

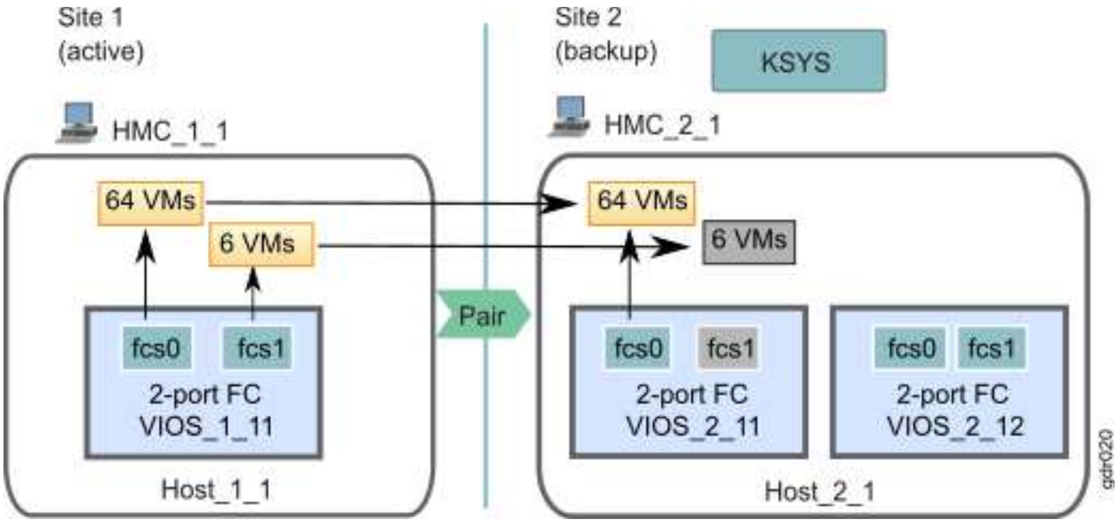

*Figure 16. VIOS configuration during disaster recovery: Scenario 2*

If the active site fails, the 70 virtual machines must be moved from the source host to the target host. The recovered virtual machines in the target host must ideally be using the VIOS partition (`VIOS_2_11`) that uses the 2-port FC adapter in the target host. However, if one of the adapters of the VIOS partition is not functional, the virtual machines that must be mapped to the nonoperational adapter are moved to the target host but remain inactive. Even when you have a dual VIOS configuration in the target host, the inactive virtual machines are not mapped to other available VIOS adapters.

In this case, the administrator must manually resolve the issues in the adapter and activate the virtual machines.

**Storage**

- Allocate the primary storage based on the current storage requirements.
- Map primary storage logical unit numbers (LUNs) to appropriate virtual machines and VIOS as required (vSCSI or NPIV).
- Allocate backup or mirror LUNs at the backup site.
- Identify the storage agent count and storage agent names based on the current storage configuration. The storage controller software that interacts with the storage subsystem must be installed on the KSYS node.
- Configure the physical connectivity and zoning of backup storage LUNs to appropriate adapters. It ensures the storage availability when virtual machines are started in the backup site.
- Set up replication relationships between primary storage LUNs and backup storage LUNs. Ensure to include virtual machine operating system LUNs in the replication. You do not need to set up consistency groups. The KSYS subsystem performs those operations.
- Have the storage administrative information ready to specify in the KSYS configuration settings (for example, storage agent count and name, user name, password, serial number of the storage disks, IP address of the storage controller server.)
- For EMC storage, review the existing SRDF composite groups. Verify that the storage disks, which are planned to be included in the GDR implementation, are not part of any existing composite groups.

**Email or contacts for notification**

Identify the contacts that must receive the notification if any failures or disaster occurs. You can have the following type of notifications:

- Email
- SMS
- Pager email

## Prerequisite for virtual machines that run on the Linux operating system

When a logical partition runs on the Linux operating system (SUSE or Ubuntu Linux distributions), the Linux operating system uses the Small Computer System Interface (SCSI) ID as the device ID for some file systems during installation. However, the SCSI ID might change when you recover the virtual machine from the active site to the backup site. In this case, the virtual machine cannot start in the backup site because of the change in SCSI ID. Therefore, you must replace the SCSI ID with the Universal Unique Identifier (UUID) in the `/etc/fstab` file and the **Boot loader** option after you install the Linux operating system in the virtual machine.

To replace the SCSI ID with UUID in the Linux virtual machine after installing the operating system, complete the following steps:

1. Identify the UUID of the required disk by running the following commands:
   a. Identify the disk that contains the Linux operating system by running the following command:

   ```
   linux:/usr/bin # fdisk -l
   ```

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

   ```
   Disk /dev/sda: 107.4 GB, 107375493120 bytes
   255 heads, 63 sectors/track, 13054 cylinders, total 209717760 sectors
   Units = sectors of 1 * 512 = 512 bytes
   Sector size (logical/physical): 512 bytes / 512 bytes
   I/O size (minimum/optimal): 512 bytes / 512 bytes
   Disk identifier: 0xbac70600
   ```

```
   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *        2048        2854         403+  41  PPC PReP Boot
/dev/sda2          417792     4626431     2104320   82  Linux swap / Solaris
/dev/sda3         4626432   209717247   102545408   83  Linux

Disk /dev/sdb: 107.4 GB, 107375493120 bytes
255 heads, 63 sectors/track, 13054 cylinders, total 209717760 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0xbac70600

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1   *        2048        2854         403+  41  PPC PReP Boot
/dev/sdb2          417792     4626431     2104320   82  Linux swap / Solaris
/dev/sdb3         4626432   209717247   102545408   83  Linux

Disk /dev/mapper/360000970000196800508533033323741: 107.4 GB, 107375493120 bytes
255 heads, 63 sectors/track, 13054 cylinders, total 209717760 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0xbac70600

                                        Device Boot      Start         End      Blocks   Id  System
/dev/mapper/360000970000196800508533033323741_part1   *        2048        2854         403+  41  PPC PReP Boot
/dev/mapper/360000970000196800508533033323741_part2          417792     4626431     2104320   82  Linux swap / Solaris
/dev/mapper/360000970000196800508533033323741_part3         4626432   209717247   102545408   83  Linux
```

In this example, the /dev/sda3 boot device is the disk that contains the Linux operating system.

b. List the corresponding UUID of the disks by running the following command:

```
linux:/dev/disk/by-id #
```

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

```
/dev/sda2: UUID="2d6e8edb-cc0e-4db1-9125-7d4ec8faf58d" TYPE="swap"
/dev/sda3: UUID="6187ca4a-1589-4f57-8c3e-33a4043450b8" TYPE="ext3"
/dev/sdb2: UUID="2d6e8edb-cc0e-4db1-9125-7d4ec8faf58d" TYPE="swap"
/dev/sdb3: UUID="6187ca4a-1589-4f57-8c3e-33a4043450b8" TYPE="ext3"
/dev/mapper/360000970000196800508533033323741_part2: UUID="2d6e8edb-cc0e-4db1-9125-7d4ec8faf58d" TYPE="swap"
```

In this example, you can identify the UUID of the /dev/sb3 disk.

2. Open and edit the /etc/fstab file to replace the SCSI ID of the disk with the UUID of the disk. For example:

```
linux:~ # cat /etc/fstab
/dev/mapper/360000970000196800508533033323741_part2 swap                 swap       defaults
        0 0
#/dev/disk/by-id/scsi-360000970000196800508533033323741-part3 /          ext3       acl,user_xattr
        1 1
/dev/disk/by-uuid/6187ca4a-1589-4f57-8c3e-33a4043450b8 /boot             ext3       acl,user_xattr
        1 1  ====> Replace SCSI ID with UUID
proc                 /proc                proc       defaults             0 0
sysfs                /sys                 sysfs      noauto               0 0
debugfs              /sys/kernel/debug    debugfs    noauto               0 0
devpts               /dev/pts             devpts     mode=0620,gid=5      0 0
```

3. Change the **Boot loader** option in the Linux virtual machine by using the Linux operating system setup and configuration tool, YaST. Go to **YaST** > **System** > **Boot loader**. Edit the root device and replace the SCSI ID with the UUID of the disk.

# Disaster recovery mechanism for the GDR solution

After you plan the details about how the GDR solution can integrate into your current environment, review the following flow chart that contains the high-level steps that are involved in the GDR implementation.

The following flow chart provides a summary of the entire GDR mechanism for disaster recovery:

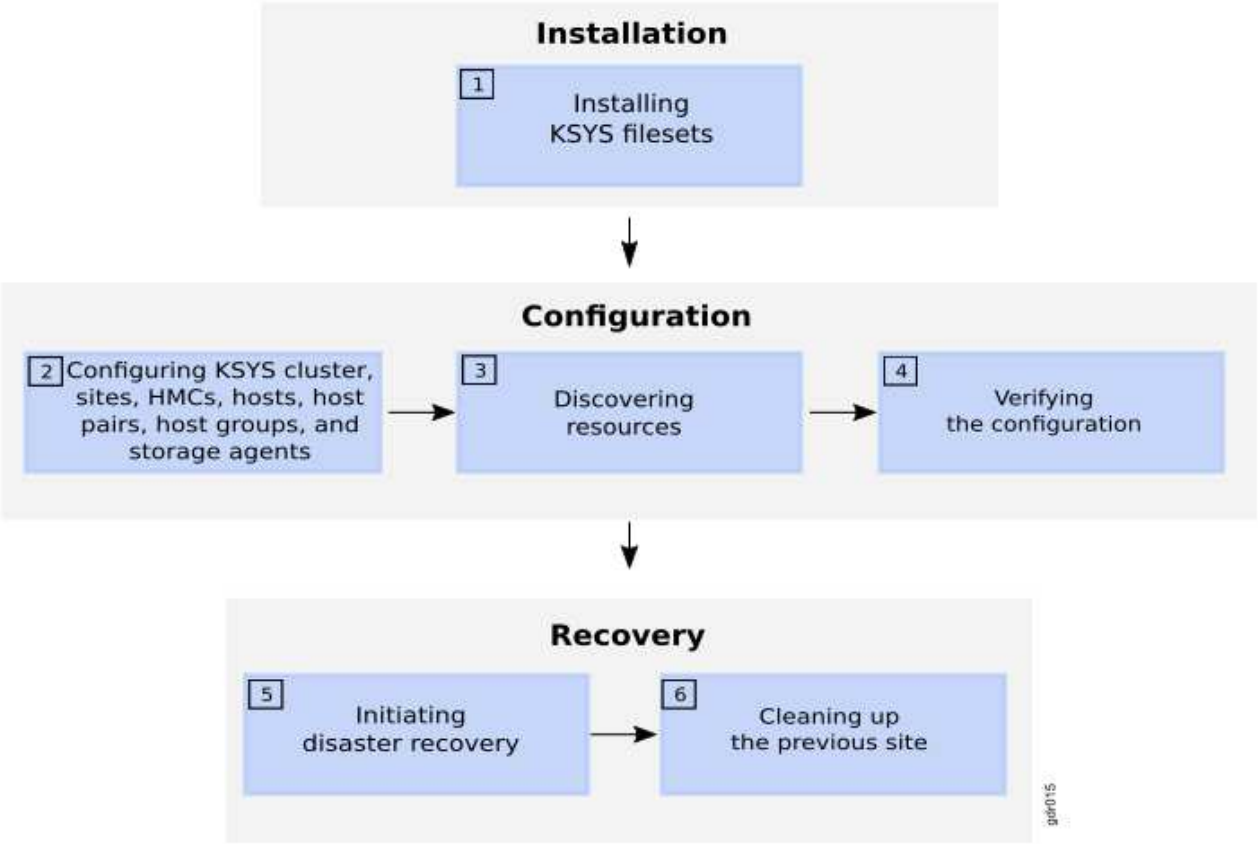
Installation
1
Installing KSYS filesets
Configuration
2
Configuring KSYS cluster, sites, HMCs, hosts, host pairs, host groups, and storage agents
3
Discovering resources
4
Verifying the configuration
Recovery
5
Initiating disaster recovery
6
Cleaning up the previous site
gdr015

*Figure 17. GDR solution: Disaster recovery mechanism*

## 1. Installation

The controlling system (KSYS) is the fundamental component in the GDR solution. Therefore, the KSYS filesets must be installed first.

The KSYS runs in an AIX 7.2.1 (or later) logical partition in the disaster recovery site. It controls the entire cloud environment for the GDR solution.

To manage the servers and data replication, the KSYS must be connected to all the managed servers through the HMC and out-of-band storage system connectivity to all associated primary and secondary disks.

© Copyright IBM Corp. 2017, 2019

## 2. Configuration

After the KSYS is installed and a one-node KSYS cluster is set up, you must configure all the other entities by using the KSYS interface.

You must complete the following procedures by using the **ksysmgr** command:

1. Create a one-node cluster for the KSYS node.
2. Create sites.
3. Add HMCs to the corresponding sites.
4. Add hosts to the corresponding sites.
5. Identify host pairs across the sites.
6. Create host groups.
7. Add storage agents to the corresponding sites.
8. Add contacts details for error notification.

## 3. Discovery

After the initial configuration is complete, the KSYS discovers all the hosts from all the host groups that are managed by the HMCs in both sites and displays the status.

During discovery, the KSYS subsystem monitors the discovery of all logical partitions (LPARs) or virtual machines (VMs) in all the managed hosts in the active site. The KSYS collects the configuration information for each LPAR, and displays the status, and also logs the status in the log files at the `/var/ksys/log/` directory.

The KSYS discovers the disks of each VM and checks whether the VMs are configured currently for the storage devices mirroring. If the disks are not configured for mirroring properly, KSYS notifies you about the volumes that are not mirrored. All volumes of a VM must be mirrored. Disks can be virtualized by using N_Port ID virtualization (NPIV), virtual SCSI (vSCSI), or combination of all these modes.

HMC collects information about the hosts, VIOS, and logical partitions that can be managed by the KSYS. For example, HMC collects information about the system processor, system memory, hardware, and worldwide port name (WWPN) of the physical Fibre Channel adapter. HMC also checks for VIOS capability for disaster recovery operations. HMC also collects the information about the host state, LPAR state, VIOS state, and IP addresses of the host, VIOS, and LPAR. HMC provides all this information to KSYS during the discovery phase.

## 4. Verification

In addition to the configuration validation that are initiated by you, the KSYS verifies and validates the environment periodically. The KSYS also verifies the configuration as part of the recovery process. In the verification phase, the KSYS fetches information from the HMC to check whether the backup site is capable to host the VMs during a disaster. The KSYS also verifies storage replication-related details and accessibility of the target disks. The verification is successful only if the storage area network (SAN) zones are configured properly on the target side.

If the verification fails as a part of the recovery process, the failure is considered as recovery failure.

## 5. Recovery

When any planned or unplanned outages occur, you must manually initiate the recovery by using the **ksysmgr** command that moves the virtual machines to the backup site. If you initiate a planned recovery, the storage replication direction is reversed from the current active site to the previously active site. If you initiate an unplanned recovery, the storage is failed over to the backup site and you must manually resynchronize the storage after the previously active site becomes operational.

## 6. Cleanup

After the disaster recovery phase, in case of a planned recovery, the KSYS automatically cleans up the source site of all the disk mapping and adapter information. In the case of an unplanned recovery, you must manually clean up the source site when the HMC and hosts in the previously active site become operational. If the VMs in the previously active site are still in active state, the VMs are first powered off, and then the cleanup operations are performed in the source site.

# Installing GDR

After you plan the details about the GDR implementation, you can install the GDR software. The GDR software contains the KSYS package that you must install on a logical partition in the backup site to manage the disaster recovery environment. The GDR solution uses other subsystems such as HMC, VIOS, and storage controllers that must exist in your current environment. You must understand concepts and planning information about the GDR solution before you install the GDR solution.

The GDR disaster recovery is enabled by using the following subsystems:

**KSYS or controller system**
The KSYS software is installed in an AIX LPAR. The KSYS monitors the entire environment and enables disaster recovery operations, if required.

**HMCs** HMCs in the active site and the backup site manage the IBM Power Systems servers.

**VIOS partitions**
Various VIOS partitions within hosts of the active site and the backup site virtualize and manage storage resources for the hosts' virtual machines.

**Storage controllers**
Storage systems in the active site and the backup site enable storage for the various virtual machines and the replication between the sites.

The following figure shows the deployment diagram of the GDR solution:

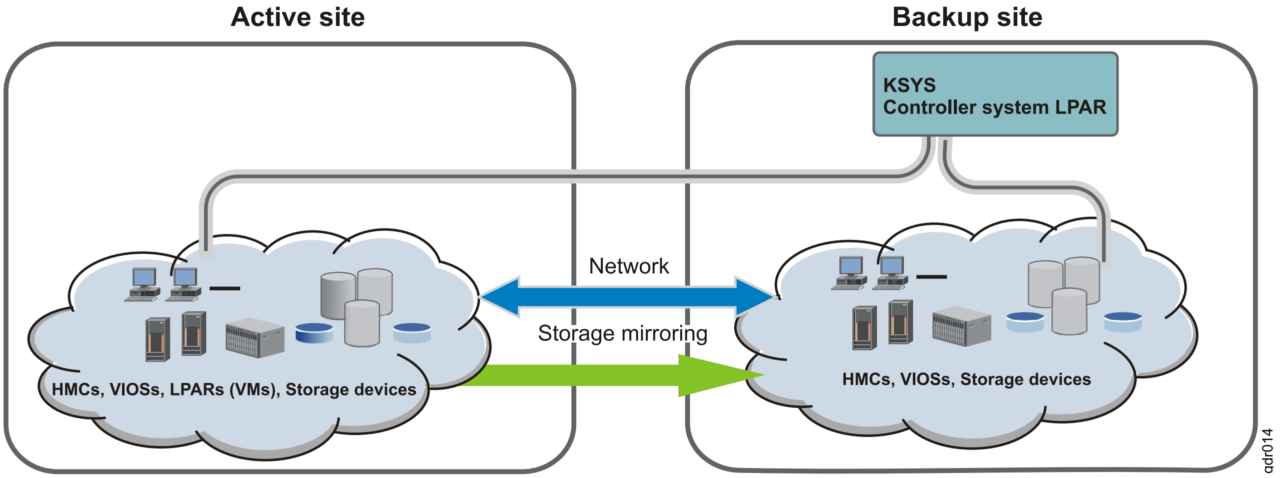

*Figure 18. GDR deployment diagram*

For more information about the installation and configuration of the HMC, VIOS, and storage subsystems, refer to the documentation of each of the subsystems.

The following section explains the packaging information and installation for the KSYS software. The KSYS software can be installed in any AIX LPAR.

**Note:** It is a best practice to dedicate this LPAR for KSYS purposes only. The LPAR, in which the KSYS software must be installed, must run IBM AIX 7.2 with Technology Level 1, or later.

© Copyright IBM Corp. 2017, 2019

When you install the KSYS software on a logical partition, the logical partition is referred to as the *KSYS node* that controls the entire environment for disaster recovery. To support disaster recovery for the GDR solution, the KSYS handles discovery, monitoring, notification, recovery, and verification aspects that are associated with the disaster recovery.

## KSYS package

The KSYS package consists of the filesets that can be installed on any AIX LPAR.

The KSYS package consists of the following filesets:

- `ksys.main.rte`
- `ksys.main.cmds`
- `ksys.mirror.emc.rte`
- `ksys.mirror.svc.rte`
- `ksys.mirror.ds8k.rte`
- `ksys.mirror.hitachi.rte`
- `ksys.main.msg.en_US.cmds`
- `ksys.license`

### Installation directories

When you install the KSYS filesets, all the necessary configuration and binary files are installed in the designated directories. Some of the important files and corresponding directories are listed in the following table.

*Table 3. Installation directories for configuration and binary files*

| Type of file | File name | Directory where the files are installed |
|---|---|---|
| KSYS administration command: **ksysmgr** binary file | `ksysmgr` | `/opt/IBM/ksys/` |
| CPU and memory capacity management command: **ksysrppmgr** binary file | `ksysrppmgr` | `/opt/IBM/ksys/` |
| Storage scripts | Multiple storage scripts | `/opt/IBM/ksys/storages/EMC/`<br>`/opt/IBM/ksys/storages/SVC/`<br>`/opt/IBM/ksys/storages/ds8k/`<br>`/opt/IBM/ksys/storages/Hitachi` |
| Sample files for configuration | • `data_collection`<br>• `setup_dr`<br>• `setup_dr_HBAs`<br>• `setup_dr_ethernet`<br>• `setup_dr_hadiskhbs`<br>• `setup_dr_hostname_ip`<br>• `setup_dr_vgs`<br>• `failover_config.cfg`<br>• `README`<br>• `setup_dr_hostname_ip_via_config_file` | `/opt/IBM/ksys/samples/site_specific_nw/AIX/` |
|  | • `postscript`<br>• `prescript` | `/opt/IBM/ksys/samples/custom_validation/` |
|  | `event_script_template` | `/opt/IBM/ksys/samples/event_handler/` |

*Table 3. Installation directories for configuration and binary files (continued)*

| Type of file | File name | Directory where the files are installed |
|---|---|---|
| Snap script directory | vmsnap | /usr/lib/ras/snapscripts/ |
| Log directory | events.log | /var/ksys/ |
| | • ksysmgr.log<br>• ksys.log<br>• ksys_srdf.log<br>• ds8k.log<br>• svc.log<br>• ksys_ccl.log | /var/ksys/log/ |

### Package size requirements

Most of the KSYS software is installed in the `/opt` file system. However, the KSYS creates multiple log files and trace files as part of first failure data capture (FFDC) information. The minimum disk space that is required for various file systems and directories are listed in the following table.

**Note:** The administrator must monitor and maintain the space required for the pre-installation and post-installation requirements.

#### Pre-installation or update

*Table 4. Disk space required before KSYS installation*

| File system or directory | Required disk space |
|---|---|
| /opt | 30 MB |

#### Post-installation or update

*Table 5. Disk space required after KSYS installation*

| File system or directory | Required disk space |
|---|---|
| /var | 200 MB |

## Installing the KSYS filesets

The KSYS filesets must be installed on the identified logical partition (LPAR), which is preferably in the backup site. You can run the **geninstall** command or use the System Management Interface Tool (SMIT) panel to install the filesets on the LPAR.

Before you install the KSYS filesets, ensure that the following prerequisites are met:

- You must have root authority to perform the installation tasks.
- Identify the host that you will use to create the KSYS node. The host must be preferably in the backup site during the normal (non-disaster) conditions.
- You must have already created an LPAR in the identified host that has IBM AIX 7.2 with Technology Level 1, or later, installed.
- Ensure that you have enough space in this LPAR so that KSYS filesets can be installed successfully. You might want to 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 in the identified LPAR, complete the following steps:

1. Download the GDR software components from Entitled Systems Support (ESS) website: http://www.ibm.com/servers/eserver/ess/index.wss
2. Copy the filesets to any location from where you want to install the filesets.

3. Uncompress the filesets as per the guidelines that are provided with the package.
4. Install the filesets by using one of the following steps:
   - To install the filesets by using command-line interface, enter the following command:

     **Note:** This method must be used when you have only KSYS-related software in the directory.

     ```
     geninstall -I "a -cgNQqwXY -J" -Z -d . -f File
     ```
   - 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 **Install Software**, and press `Enter`.

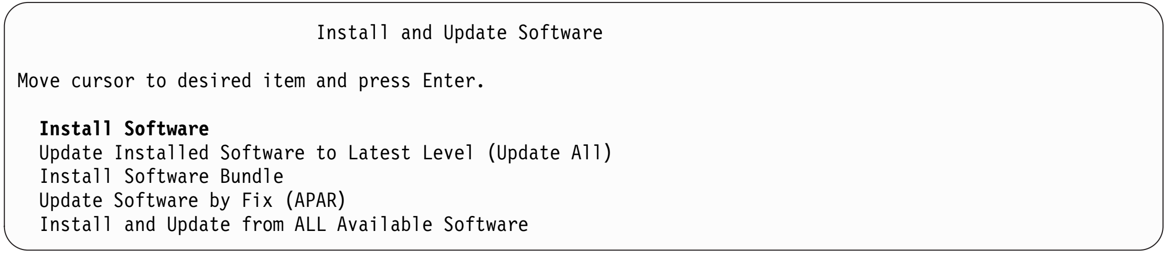

```
                            Install and Update Software

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 Install Software screen, change the values according to your situation. You must also accept new license agreements. Press `Enter` after you make all the changes.

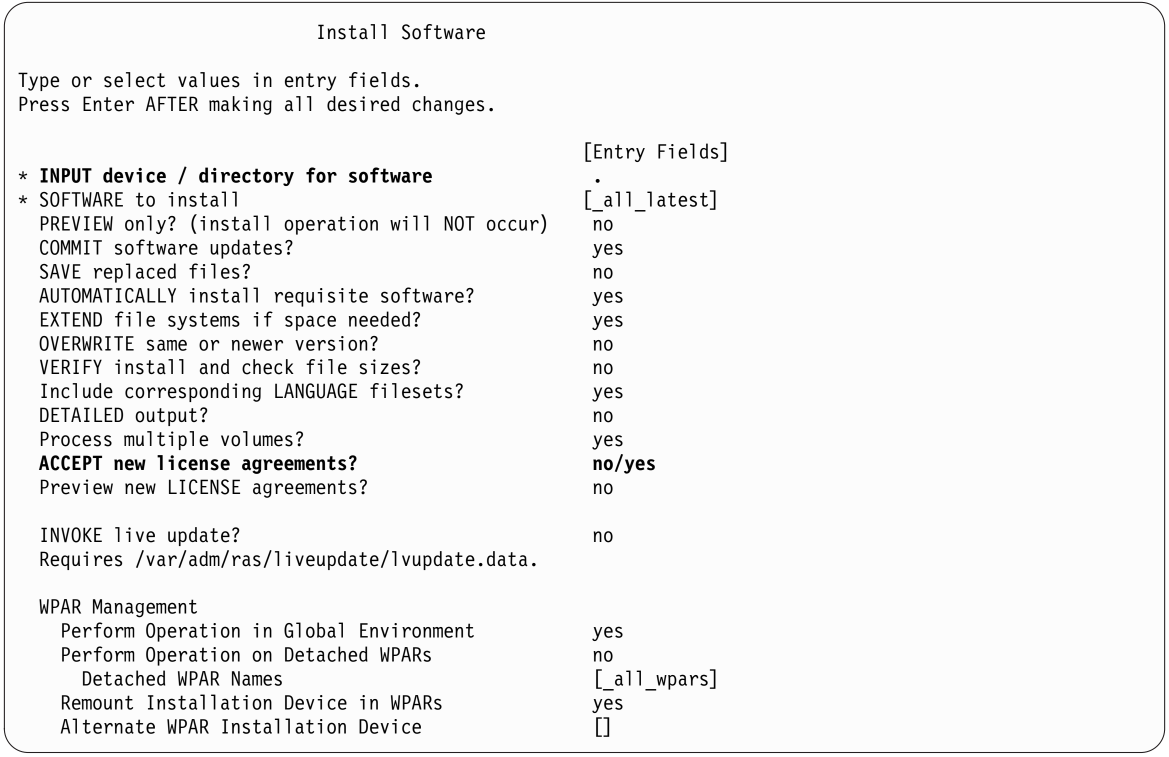

```
                            Install Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                     [Entry Fields]
* INPUT device / directory for software               .
* SOFTWARE to install                                [_all_latest]
  PREVIEW only? (install operation will NOT occur)    no
  COMMIT software updates?                            yes
  SAVE replaced files?                                no
  AUTOMATICALLY install requisite software?           yes
  EXTEND file systems if space needed?                yes
  OVERWRITE same or newer version?                    no
  VERIFY install and check file sizes?                no
  Include corresponding LANGUAGE filesets?            yes
  DETAILED output?                                    no
  Process multiple volumes?                           yes
  ACCEPT new license agreements?                      no/yes
  Preview new LICENSE agreements?                     no

  INVOKE live update?                                 no
  Requires /var/adm/ras/liveupdate/lvupdate.data.

  WPAR Management
    Perform Operation in Global Environment           yes
    Perform Operation on Detached WPARs               no
      Detached WPAR Names                             [_all_wpars]
    Remount Installation Device in WPARs              yes
    Alternate WPAR Installation Device                []
```

5. 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 the installation was not successful, check the reason of failure in the output. Contact IBM support, if necessary.

## Upgrading the KSYS software

If you have an earlier version of the KSYS software, you can upgrade the GDR solution to GDR version 1.2 by installing the latest KSYS filesets. The latest filesets will overwrite the existing software on the KSYS logical partition.

**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 operations must be complete before you attempt to upgrade the GDR version to version 1.2.
- If you attempted a move operation in the earlier GDR version and the move operation failed, the virtual machines might be in `Recovery` state. In this case, you must recover the virtual machines successfully or move the virtual machines to a proper state before upgrading the GDR version to version 1.2.

When you upgrade the existing GDR version to GDR Version 1.2, the KSYS configuration settings (for example, details about HMCs, hosts, host pairs, storage agents, and so on) are retained without any change. After the upgrade operation is complete, you must run a discovery operation before you add, delete, or modify any configuration settings in the KSYS subsystem.

Beginning with GDR Version 1.2, the KSYS subsystem supports host groups. Therefore, when you upgrade the existing GDR version to GDR Version 1.2, all the hosts in the KSYS configuration are added to a default host group called `Default_HG`. You can add and modify the host groups based on your requirements. If the virtual machines are moved to the backup site, and the GDR version is upgraded later, the `Default_HG` host group aligns with the new site automatically.

To upgrade the KSYS software in the existing KSYS LPAR, complete the following steps:

1. Download the GDR software components from Entitled Systems Support (ESS) website: http://www.ibm.com/servers/eserver/ess/index.wss
2. Copy the filesets to the location where the existing filesets are installed.
3. Decompress the filesets according to the guidelines that are provided with the package.
4. 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`.

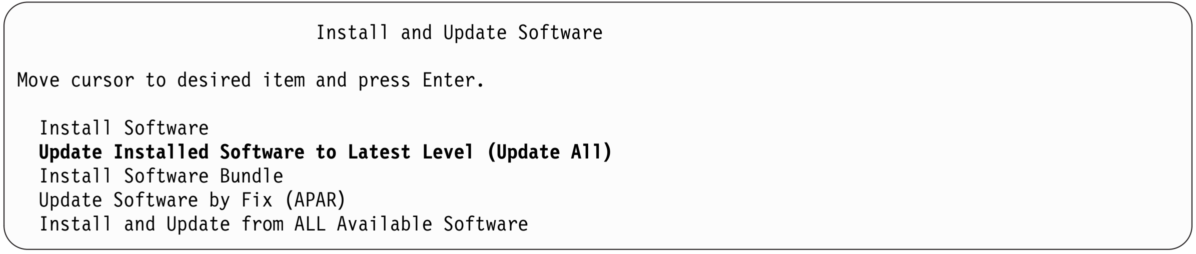

```
                           Install and Update Software

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.
5. 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 the installation was not successful, check the reason of failure in the output. Contact IBM support, if necessary.

# Configuring GDR

After you install the KSYS filesets, you can use the **ksysmgr** command to interact with the KSYS daemon to manage the entire environment for disaster recovery.

Review the GDR concepts to understand the associated entities for the GDR solution.

Before you configure the KSYS settings, you must complete the following tasks:

- Include the /opt/IBM/ksys path to the root directory so that you need not specify this path every time you run the **ksysmgr** command. To include this path, run the following command in the KSYS node:

  ```
  export PATH=$PATH:/opt/IBM/ksys
  ```

  Or, you can append this command to the .profiles file in the root directory.
- Verify that the KSYS node does not have an existing KSYS configuration by running the following command:

  ```
  lsrpdomain
  ```

  The output must indicate that no domain exists.
- Verify that the IBM.VMR resource class exists by running this command:

  ```
  lsrsrc | grep VMR
  ```

  Or,

  ```
  lssrc -a | grep VMR
  ```

  The output must display various IBM.VMR resource classes. If the **lsrsrc** command does not display any IBM.VMR resource classes, check the existing classes by running the following command:

  ```
  grep -i vmr /usr/sbin/rsct/cfg/ct_class_ids
  ```

  The output must display the IBM.VMR resource class.

  **Note:** You cannot configure the KSYS resources if the IBM.VMR resource class does not exist.
- If you are using AIX LPARs as virtual machines in the active site, set the **ghostdev** parameter in the AIX kernel. Thereby, after a disaster recovery movement to the backup site, all the ghost devices associated with the recovered LPARs are removed from the backup site during the AIX boot process. To set the **ghostdev** parameter, complete the following steps:
  1. Log in to the source virtual machine as a root user.
  2. Modify the AIX kernel by running the following command:

     ```
     # chdev -l sys0 -a ghostdev=yes
     ```
  3. Run the following commands:

     ```
     # savebase
     # bosboot -ad /dev/ipldevice
     ```

     **Note:** These commands must be run on the source virtual machines. Since all the data is replicated to the backup site, these changes are also propagated from the active site to the backup site.

## Flow chart for KSYS configuration

After the KSYS filesets are installed, you must complete the configuration steps to use the disaster recovery feature of the GDR solution. The configuration steps involve adding resources to the KSYS configuration.

© Copyright IBM Corp. 2017, 2019

The following flow chart provides a summary of the configuration steps:

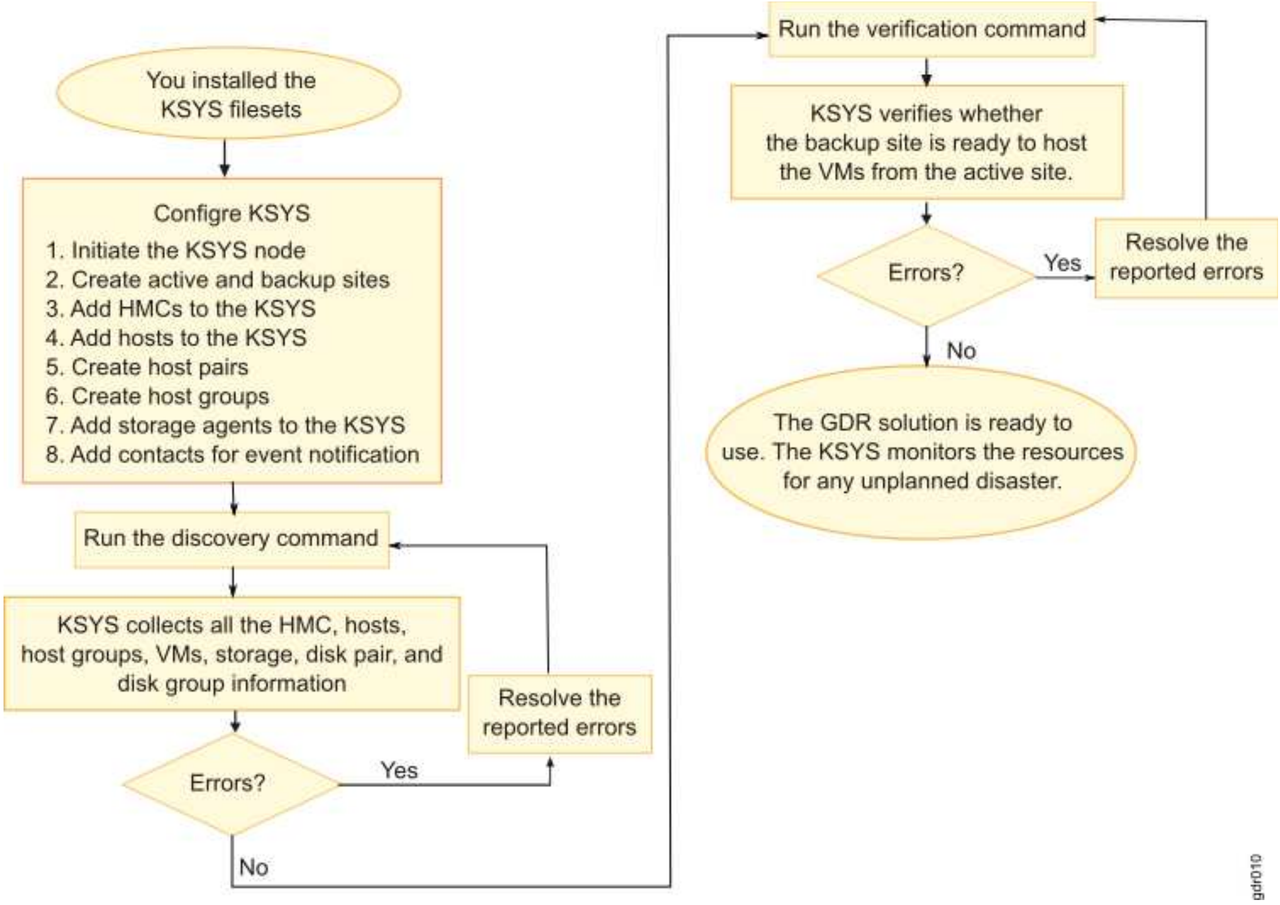
You installed the KSYS filesets
Configre KSYS
1. Initiate the KSYS node
2. Create active and backup sites
3. Add HMCs to the KSYS
4. Add hosts to the KSYS
5. Create host pairs
6. Create host groups
7. Add storage agents to the KSYS
8. Add contacts for event notification
Run the discovery command
KSYS collects all the HMC, hosts, host groups, VMs, storage, disk pair, and disk group information
Errors?
Yes
Resolve the reported errors
No
Run the verification command
KSYS verifies whether the backup site is ready to host the VMs from the active site.
Errors?
Yes
Resolve the reported errors
No
The GDR solution is ready to use. The KSYS monitors the resources for any unplanned disaster.
gdr010

*Figure 19. GDR solution: Installation and configuration*

## Setting up the KSYS subsystem

After the GDR software is installed, you must complete the some mandatory configuration steps to start using the disaster recovery feature of the GDR solution.

## Initiating the KSYS node

After the GDR software is installed on the KSYS LPAR, you must initiate and set up the KSYS environment before you configure the disaster recovery environment. The KSYS environment relies on Reliable Scalable Cluster Technology (RSCT) to create its cluster.

After you create the KSYS cluster, various daemons of RSCT and KSYS are activated. The KSYS LPAR can then process the commands that are required to configure the disaster recovery environment.

**Note:** In GDR Version 1.1, the KSYS operates as a one-node cluster.

To create and initiate a one-node KSYS cluster, complete the following steps:

1. Create a cluster and add the KSYS node to the cluster by running the following command.
```
ksysmgr add ksyscluster cluster_name
          ksysnodes=test_ksysnode1 type=DR
```
2. Verify the KSYS cluster configuration by running the following command.
```
ksysmgr verify ksyscluster cluster_name
```

3. Deploy the one-node KSYS cluster by running the following command.

   ```
   ksysmgr sync ksyscluster cluster_name
   ```

   **Note:** You can perform steps 1 - 3 by running the following command:

   ```
   ksysmgr add ksyscluster cluster_name ksysnodes=ksys_nodename
         sync=yes type=DR
   ```

   This command creates a cluster, adds the KSYS node to the cluster, verifies the cluster configuration, and deploys the one-node KSYS cluster.

4. Optional: Verify the KSYS cluster that you created by running the following commands:

   ```
   ksysmgr query ksyscluster
   Name:         test_ksys
   State:        Online

   # lsrpdomain
   Name          OpState     RSCTActiveVersion    MixedVersions    TSPort    GSPort
   test_ksys     Online      3.2.1.1              No               12347     12348

   # lssrc -a | grep VMR
    IBM.VMR          rsct_rm          9961878      active
   ```

   **Note:** These commands do not display any output until you run the **ksysmgr sync** command.

## Creating sites

You must create sites that are used to map all the HMCs, hosts, and storage devices. You must create an active site where the workloads are currently running and a backup site that acts as a backup for the workloads during a disaster or a potential disaster situation.

*Sites* are logical names that represent your sites. A site name can be any American Standard Code for Information Interchange (ASCII) string that is limited to 64 characters. A site name cannot contain any special characters or spaces.

By default, the active site is the home site. You can configure only two sites. When you create sites, the replication type of the site is asynchronous by default. After you create sites, you can change the type of storage replication to synchronous.

To create sites for the GDR solution, complete the following steps in the KSYS logical partition:

1. Create an active site (`Site1`) by running the following command:

   ```
   ksysmgr add site Site1 sitetype=home
   ```

2. Create a backup site (`Site2`) by running the following command:

   ```
   ksysmgr add site Site2 sitetype=backup
   ```

3. Optional: Verify the sites that you created by running the following command:

   ```
   ksysmgr query site
   ```

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

   ```
   Replication type for site(s): async
   Name:            Site2
   Sitetype:        BACKUP

   Name:            Site1
   Sitetype:        HOME
   ```

4. If you want to change the replication type of a site from the default `async` value to the `sync` value, enter the following command:

   ```
   ksysmgr modify system replication_type=sync sites=Site1,Site2
   ```

## Adding HMCs to the KSYS

The KSYS interacts with the HMC for discovery, verification, monitoring, recovery, and cleanup operations. HMCs that are configured in both the active and backup sites provide details about the hosts and VIOS partitions that are managed by the HMCs in each site. The GDR solution cannot be implemented without configuring the HMCs. Therefore, you must provide the HMC details to the KSYS.

**Note:** The HMC user, whose user name and password details are provided to the KSYS, must have `hmcsuperadmin` privileges with remote access enabled.

To add the HMCs to a specific site, complete the following steps in the KSYS LPAR:

1. Add the HMC (for example, HMC name: `Site1_HMC1`, user name: `hscroot`, password: `xyz123`) that manages the host or hosts in the active site (`Site1`) by running the following command:

   ```
   ksysmgr add hmc Site1_HMC1 login=hscroot password=xyz123
           hostname=Site1_HMC1.testlab.ibm.com site=Site1
   ```

   **Note:** If you do not want to enter the password in the command line, you can omit the password field. In this case, the **ksysmgr** command prompts for the password later.
2. Add the HMC (for example, HMC name: `Site2_HMC1`, user name: `hscroot`, password: `xyz123`) that manages the host or hosts in the backup site (`Site2`) by running the following command:

   ```
   ksysmgr add hmc Site2_HMC1 login=hscroot password=xyz123
           hostname=Site2_HMC1.testlab.ibm.com site=Site2
   ```
3. Repeat Step 1 and Step 2 for all the HMCs that you want to add to the KSYS.
4. Optional: Verify the HMCs that you added by running the following command:

   ```
   ksysmgr query hmc
   ```

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

   ```
   Name:            Site2_HMC1
   Site:            Site2
   Ip:              9.xx.yy.zz
   Login:           hscroot
   HmcTimeout:      0
   Maxjobs:         0
   ViosTimeout:     0
   SwXSDVersion:    V8R8.6.0

                           Managed Host List:

   Host Name                               Uuid
   =========                               ====
   d25m1-9179-MHB-100141P                  82e8fe16-5a9f-3e32-8eac-1ab6cdcd5bcf
   d67m3_8231-E2D-068ED7H                  74931f30-e852-3d47-b564-bd263b68f1b1
   kumquat_9179-MHD-105E67P                c15e9b0c-c822-398a-b0a1-6180872c8518
   r5r3m1                                  f6cbcbda-8fec-3b6e-a487-160ca75b2b84
   rar1m3-9117-MMD-1016AAP                 4ce17206-4fde-3d5a-a955-dbf222865a77
   gsk1_8233-E8B-1000ADP                   346f184d-bace-36f5-97b5-3955c62a6929
   Hawk-8233-E8B-SN1000A7P                 a977d52e-fd3a-325d-bd02-124663066cac
   rar1m6                                  ae115482-3a50-32f3-935a-7ff9be433e33
   rar1m5-9109-RMD-106A79R                 b3880199-3b8b-3ade-b360-f76146c2d7f3
   r5r3m2                                  26c7c48e-3b81-363b-85d0-e110ebc43b15
   ============================================================================


   Name:            Site1_HMC1
   Site:            Site1
   Ip:              9.xx.yy.zz
   Login:           hscroot
   HmcTimeout:      0
   Maxjobs:         0
   ViosTimeout:     0
   SwXSDVersion:    V8R8.6.0
   ```

```
                         Managed Host List:

Host Name                               Uuid
=========                               ====
e10m4_8286-42A-21E0A7V                  caffee0a-4206-3ee7-bfc2-f9d2bd3e866f
pbrazos_9119-MME-21BBC47                6ce366c5-f05d-3a12-94f8-94a3fdfc1319
orange_9179-MHD-107895P                 67ff62ec-ecb5-3ad4-9b35-0a2c75bb7fe4
============================================================================
```

**Related concepts**:

"Requirements for the GDR solution" on page 27
Before you plan the implementation of the GDR solution, you must understand the other entities and resources that the GDR solution requires for disaster recovery operations.

## Adding hosts to the KSYS

The KSYS monitors and manages the disaster recovery operations across sites. The KSYS requires that each host must be paired to another host across sites. This type of pairing allows the virtual machines to move from one host to another host across sites. You must plan the host pairing across sites in advance, and then implement the pairing.

**Note:** The GDR solution supports POWER7 processor-based server, or later.

**Host pairing guidelines**

- Each host in a pair must belong to different sites.
- Paired host in the backup site must always have enough resources to host the managed virtual machines from the active site. For example, if the current active site is `Site_1` and the `Host_1` host in that site is running 55 virtual machines that requires 100 cores and 512 GB memory, you can pair the `Host_1` host from the active site with the `Host_2` host on the backup site only when the `Host_2` host contains at least 100 cores CPU and 512 GB memory. If the capacity is not sufficient in the backup host, the validation issues warnings.
- If you use the Enterprise Pool capacity option, the `Host_2` host can borrow required resources from the Enterprise Pool before the disaster recovery movement.
- You can pair the POWER8 host and the POWER7 hosts. For example, you can pair a POWER8 host on the active site with a POWER7 host on the backup site if the POWER7 system has enough resources to host the virtual machines from the POWER8 host.

After the HMCs are added to the KSYS, you can review the list of managed hosts by each HMC, and then identify the hosts that you want to add to the KSYS for disaster recovery. You must connect the source hosts and target hosts to different HMCs across sites. If you connect the source hosts and target hosts to the same HMC, it leads to an invalid configuration in the KSYS subsystem and can cause failures in disaster recovery operations.

When a host is added to a host group, all the virtual machines in the host are included by default in the disaster recovery management scope. However, the disaster recovery management starts only after you configure the subsystems and run the discovery and verification operations. Therefore, if you plan to exclude a set of virtual machines after adding the hosts, you can unmanage those virtual machines, and then run the discovery and verification operations.

You can prioritize the virtual machines within a single host or among multiple hosts such that the virtual machines that run important workloads are considered first for the move operation. You can set a priority of `high`, `medium`, or `low` for the virtual machines. When you run a move operation, the **ksysmgr** command initiates the operation for the virtual machines that has the highest priority.

**Tip:**

- Obtain the host name by querying the HMC. You can copy the host name and use it when you run commands.
- If you use PowerHA SystemMirror for high availability and disaster recovery in your environment, see "PowerHA SystemMirror and GDR coexistence" on page 97.

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

1. Add the managed host (for example, host name: `Site1_host1`), which is running the workload, to the KSYS by running the following command:

   ```
   ksysmgr add host Site1_host1 site=Site1
   ```

   The `Site1_host1` host, which is a managed host for the `Site1_HMC1` HMC, is added in the active site.
2. Add the backup host (for example, host name: `Site2_host1`), which acts as a backup host, to the KSYS by running the following command:

   ```
   ksysmgr add host Site2_host1 site=Site2
   ```

   The `Site2_host1` host, which is a managed host for the `Site2_HMC1` HMC, is added in the backup site.
3. Repeat Step 1 and Step 2 for all hosts that you want to add to the KSYS.
4. Optional: Verify the hosts that you added by running the following command:

   ```
   ksysmgr query host
   ```

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

   ```
   Name:          Site2_host1
   UUID:          c15e9b0c-c822-398a-b0a1-6180872c8518
   FspIp:
   Pair:          None
   Site:          Site2
   VIOS:          Site2_VIOS1
                  Site2_VIOS2
   HMC:           Site2_HMC1

   Name:          Site1_host1
   UUID:          67ff62ec-ecb5-3ad4-9b35-0a2c75bb7fe4
   FspIp:
   Pair:          None
   Site:          Site1
   VIOS:          Site1_VIOS1
                  Site1_VIOS2
   HMC:           Site1_HMC1
   ```
5. If you want to exclude some virtual machines during a recovery operation, run the following command for each virtual machine that you want to exclude:

   ```
   ksysmgr unmanage VM_name
   ```
6. If you want to specify a priority for a specific virtual machine for discovery, verification, and move operations, run the following command:

   ```
   ksysmgr modify VM name1[,name2,name3,...] | file=filepath
      [uuid=uuid_value]
      [host=hostname]
      [priority=low|medium|high]
   ```

   where, the **`file`** parameter is an XML file that contains a list of virtual machine references.

## Creating host pairs

After the hosts are added to the KSYS, identify the hosts that must be paired across the active site and the backup site. Each backup host in the host pair must meet all the resource requirements so that the backup host can run the same workload in a disaster or a potential disaster situation.

To pair the hosts across the sites in the GDR solution, complete the following steps in the KSYS LPAR:

1. Pair the identified host (for example, host name: `Site1_host1`) in the active site to the identified host (for example, host name: `Site2_host1`) in the backup site by running the following command:

   ```
   ksysmgr pair host Site1_host1 pair=Site2_host1
   ```

2. Repeat Step 1 for all the host pairs that you want to create.
3. Optional: Verify the host pair that you created by running the following command:

   ```
   ksysmgr query host
   ```

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

   ```
   Name:           Site2_host1
   UUID:           c15e9b0c-c822-398a-b0a1-6180872c8518
   FspIp:
   Pair:           Site1_host1
   Site:           Site2
   VIOS:           Site2_VIOS1
                   Site2_VIOS2
   HMC:            Site2_HMC1

   Name:           Site1_host1
   UUID:           67ff62ec-ecb5-3ad4-9b35-0a2c75bb7fe4
   FspIp:
   Pair:           Site2_host1
   Site:           Site1
   VIOS:           Site1_VIOS1
                   Site1_VIOS2
   HMC:            Site1_HMC1
   ```

## Creating host groups

You can group a set of hosts based on your business requirements. For example, you can group the hosts that run similar type of workloads. You can also group important hosts together so that the monitoring and recovery operations can be performed for the set of hosts together and quickly. In disaster situations, you can move a host group separately to the backup site.

**Guidelines to manage host groups:**

- A host must already be added to the KSYS configuration settings and the host must have been paired with a backup host in the backup site.
- Each host in a site must be a part of a host group. If a host is not added to any host groups, the host is automatically added to the `Default_HG` host group during the discovery operation.
- If you add or remove hosts from a host group, you must run a discovery operation to manage or unmanage all the virtual machines from the recovery management. The modified host group displays the correct list of managed virtual machines only after a discovery operation.
- If we remove all hosts from `Default_HG`, the disk group corresponding to `Default_HG` is not removed. The disk groups are retained with the removed hosts.
- The corresponding hosts in the backup site that are paired with the active site hosts are grouped logically within the same host group. For example, if `host1` in the active site is paired with `host2` in the backup site and you create a host group `hg1` with `host1`, then `host2` is automatically added to the host group `hg1`.
- Each host group must be associated with a separate disk group. The disks in the disk group must not be shared among different host groups. The disk groups are named in the following format:

  ```
  VMRDG_{peer_domain_name}_{site_name}_{host_group_ID}
  ```

  However, the disk group name must not exceed the maximum number of characters that is supported for a consistency group at storage level.
- For SAN Volume Controller and DS8000 series of storage systems, host groups can span across a single type of storage. Multiple host groups can use same type of storage disks but the disks must not be shared among different host groups.

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

1. To create a host group and add the existing hosts that you want to include in this host group, run the following command:

   ```
   ksysmgr add host_group hg1 site=Site_1 hosts=Host_11,Host_12,Host_13
   ```

   Note that all the hosts in the backup site that are paired with the specified hosts are also added to the host group.
2. Repeat Step 1 for all host groups that you want to create in the KSYS subsystem.
3. Optional: To verify the host groups that you added, run the following command:

   ```
   ksysmgr query host_group hgname
   ```

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

   ```
   Name:               hg1
   Active Site Hosts:  Host_11
                       Host_12
                       Host_13
   Backup Site Hosts:  Host_21
                       Host_22
                       Host_23
   cpu_capacity:       none
   memory_capacity:    none
   skip_power_on:      No
   Site:               Site_1
   Vswitchmap:         Not currently set
   Vlanmap:            Not currently set
   ```
4. To add or remove hosts from the existing host groups, run the following command:

   ```
   ksysmgr modify host_group hg_name add | remove
         hosts=host1,host2... | file=filepath
   ```

   Where, the **`file`** parameter is an XML file that contains a list of hosts. An example of the XML file is as follows:

   ```
   <KSYSMGR><HOST><NAME>host1</NAME></HOST></KSYSMGR>
   <KSYSMGR><HOST><NAME>host2</NAME></HOST></KSYSMGR>
   <KSYSMGR><HOST><NAME>host3</NAME></HOST></KSYSMGR>
   ```
5. To modify the capacity-related attributes for all the hosts in a host group, run the following command:

   ```
   ksysmgr modify host_group hg_name options
         [memory_capacity=(Whole Number > 1) | minimum | current_desired | none]
         [cpu_capacity=(Whole Number > 1) | minimum | current_desired | none]
         [skip_resource_check=yes|no]
         [skip_power_on=yes|no]
   ```

   For more information about flexible capacity policies, see "Configuring the flexible capacity policies" on page 74.

## Adding storage agents to the KSYS

In the GDR solution, data is replicated from the active site to the backup site by using storage replication. The KSYS manages the storage subsystems for data mirroring as a part of the disaster recovery operations. To manage the storage subsystems, the KSYS uses the APIs that are provided by the storage subsystem. You must register the various storage devices in all sites with the KSYS as storage agents so that KSYS can monitor and manage the data replication across sites. The storage agents interact with the storage devices or the storage controller software in each site depending on the type of storage in your environment.

**Note:** All the prerequisite software must be installed on the same logical partition in which the KSYS software is installed. If the storage subsystem uses a storage controller software, the storage controller software must also be installed on the KSYS LPAR. The storage controller is the software component that you receive from the storage vendor that allows KSYS to contact storage devices and perform replication operations. For example, identify mirroring pairs, create disk groups, and reverse mirroring.

**Related concepts**:

"Requirements for the GDR solution" on page 27
Before you plan the implementation of the GDR solution, you must understand the other entities and resources that the GDR solution requires for disaster recovery operations.

## Registering EMC storage devices

For EMC storage, the storage agent use the SYMAPI commands to interact with the EMC Solution Enabler software to manage the EMC storage devices. You can use the SYMAPI commands (for example, **symcfg list**) to determine the 12-digit serial number and the IP address of the storage device.

To add the storage agents in the GDR solution, complete the following steps in the KSYS LPAR:

1. Add the storage agent (for example, name: `Site1_storage1`, user name: `abc`, password: `123`, serial number: `000196800508`, IP address: `10.xx.yy.zz`) to the active site by running the following command:

   ```
   ksysmgr add storage_agent Site1_storage1 login=abc password=123 site=Site1
   serialnumber=000196800508 storagetype=emc ip=10.xx.yy.zz
   ```

2. Add the storage agent (for example, name: `Site2_storage1`, user name: `abc`, password: `123`, serial number: `000196800573`, IP address: `10.xx.yy.zz`) to the backup site by running the following command:

   ```
   ksysmgr add storage_agent Site1_storage1 login=abc password=123 site=Site2
   serialnumber=000196800573 storagetype=emc ip=10.xx.yy.zz
   ```

3. Repeat Step 1 and Step 2 for all storage arrays that you want to add to the KSYS. For example, if a site contains two storage arrays, you must add two storage agents to the KSYS.
4. Optional: Verify the storage agents that you added by running the following command:

   ```
   ksysmgr query storage_agent
   ```

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

   ```
   Name:            Site2_storage1
   Serialnumber:    196800573
   Storagetype:     EMC
   Site:            Site2
   Ip:              10.xx.yy.zz
   Login:           abc

   Name:            Site1_storage1
   Serialnumber:    196800508
   Storagetype:     EMC
   Site:            Site1
   Ip:              10.xx.yy.zz
   Login:           abc
   ```

## Registering SVC and Storwize storage devices

For IBM SAN Volume Controller (SVC) storage system and Storwize storage system, the storage agent uses specific storage scripts and their corresponding interfaces to interact with the storage devices. When you add a storage agent to the KSYS subsystem, you must specify the cluster ID, storage login user name, and the IP address of the storage subsystem.

You can use the SVC or Storwize GUI or the following command to get the cluster ID of storage disks:

```
ssh usrid@ipaddr svcinfo lscluster
```

To add the storage agents in the GDR solution, complete the following steps in the KSYS LPAR:

1. Add the storage agent (for example, name: `Site1_storage1`, user name: `abc`, cluster ID: `007`, IP address: `10.xx.yy.zz`) to the active site by running the following command:

```
ksysmgr add storage_agent Site1_storage1 login=abc site=Site1
clusterid=007 storagetype=svc ip=10.xx.yy.zz
```

**Note:** The value of the **storagetype** attribute must be svc for both the SVC system and the Storwize system.

2. Add the storage agent (for example, name: Site2_storage1, user name: abc, cluster ID: 008, IP address: 10.xx.yy.zz) to the backup site by running the following command:

   ```
   ksysmgr add storage_agent Site1_storage1 login=abc site=Site2
   clusterid=008 storagetype=svc ip=10.xx.yy.zz
   ```

3. Optional: Verify the storage agents that you added by running the following command:

   ```
   ksysmgr query storage_agent
   ```

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

   ```
   Name:            Site2_storage1
   Clusterid:       008
   Storagetype:     SVC
   Site:            Site2
   Ip:              10.xx.yy.zz
   Login:           abc

   Name:            Site1_storage1
   Clusterid:    007
   Storagetype:     SVC
   Site:            Site1
   Ip:              10.xx.yy.zz
   Login:           abc
   ```

## Registering DS8000 series storage devices

For DS8000 series storage system, the storage agent uses specific storage scripts and their corresponding interfaces to interact with the storage devices. When you add a storage agent for the DS8000 storage system, you must specify the serial number or the storage ID, storage login user name, password, and the IP address of the storage subsystem.

To plan for the deployment of the IBM DS8000 series across two sites, complete the following prerequisites:

- Add the path name of the dscli client to the PATH environment variable for the root user on the KSYS node.
- Identify the disk volumes on the DS8000 storage system that contain the application data that you want to be included for disaster recovery.
- Ensure that sufficient number of disk volumes and Fibre Channel ports are available on the storage systems for the active site and backup site to allow a mirror path or PPRC path between the storage units.
- Verify that a FlashCopy® relationship is established for each disk volume on the backup site.
- Verify that all the data volumes that need to be mirrored are accessible to all relevant hosts. The DS8000 disk volumes must be zoned so that the FlashCopy volumes cannot be accessed by the KSYS node.
- Ensure that the KSYS node can access all HMCs by using the Internet Protocol network.

To add the storage agents in the GDR solution, complete the following steps in the KSYS LPAR:

1. Add the storage agent (for example, name: Site1_storage1, user name: abc, serial number: IBM.2107-75LY981, IP address: 10.xx.yy.zz) to the active site by running the following command:

   ```
   ksysmgr add storage_agent Site1_storage1 login=abc site=Site1
   serialnumber=IBM.2107-75LY981 storagetype=ds8k ip=10.xx.yy.zz
   ```

2. Add the storage agent (for example, name: Site2_storage1, user name: abc, serial number: IBM.2107-75LY982, IP address: 10.xx.yy.zz) to the backup site by running the following command:

```
ksysmgr add storage_agent Site1_storage1 login=abc site=Site2
serialnumber=IBM.2107-75LY982 storagetype=ds8k ip=10.xx.yy.zz
```

3. Optional: Verify the storage agents that you added by running the following command:

   ```
   ksysmgr query storage_agent
   ```

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

   ```
   Name:            Site2_storage1
   Serialnumber:    IBM.2107-75LY981
   Storagetype:     DS8K
   Site:            Site2
   Ip:              10.xx.yy.zz
   Login:           abc

   Name:            Site1_storage1
   Serialnumber:    IBM.2107-75LY982
   Storagetype:     DS8K
   Site:            Site1
   Ip:              10.xx.yy.zz
   Login:           abc
   ```

## Registering Hitachi storage devices

For Hitachi storage systems, the storage agents use the Command Control Interface (CCI) to interact with the CCI server to manage the storage devices. GDR 1.2.0 supports only asynchronous mode of storage replication.

For Hitachi storage systems, the KSYS subsystem requires the Hitachi Open Remote Copy Manager (HORCM) instance for each storage agent. You must specify the instance parameter to indicate the storage agent for the HORCM instance.

To add the storage agents in the GDR solution, complete the following steps in the KSYS LPAR:

1. Add the storage agent (for example, name: `Site1_storage1`, serial number: `441108`, IP address: `10.xx.yy.zz`) to the active site by entering the following command:

   ```
   ksysmgr add storage_agent Site1_storage1 instance=14 site=Site1
   serialnumber=441108 storagetype=hitachi ip=10.xx.yy.zz
   ```

2. Add the storage agent (for example, name: `Site2_storage1`, serial number: `357558`, IP address: `10.xx.yy.zz`) to the backup site by entering the following command:

   ```
   ksysmgr add storage_agent Site1_storage1 instance=15 site=Site2
   serialnumber=357558 storagetype=hitachi ip=10.xx.yy.zz
   ```

3. Optional: Verify the storage agents that you added by entering the following command:

   ```
   ksysmgr query storage_agent
   ```

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

   ```
   Name:            Site1_storage1
   Serialnumber:    441108
   Storagetype:     Hitachi
   Site:            Site1
   Ip:              10.xx.yy.zz
   Instance:        14

   Name:            Site2_storage1
   Serialnumber:    357558
   Storagetype:     Hitachi
   Site:            Site2
   Ip:              10.xx.yy.zz
   Instance:        15
   ```

| 
## Setting up contacts for event notification

The KSYS tracks various events that occur in the environment, analyzes the situation, and notifies you about any issues or potential disaster through the registered contacts. You must provide the contact details to the KSYS so that you can receive notifications about any situation that might need your action.

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

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

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

You must specify the phone number along with the phone carrier email address to receive a short message service (SMS) notification. To find your phone carrier email address, contact your phone service provider or visit the following website: http://www.emailtextmessages.com.

**Note:** The logical partition, in which the KSYS software is installed, must have a public IP address to send the event notifications successfully.

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

- To add an email notification for a specific user, enter the following command:

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

  For example,

  ```
  ksysmgr add notify user=John contact=john.doe@testmail.com
  ```
- To add an SMS notification for a specific user, 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
  ```
- To add a pager notification for a specific user, enter the following command:

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

  For example,

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

## Discovering resources

After you add the various HMCs, hosts, host groups, and storage subsystems to the KSYS subsystem for disaster recovery management, you must run the **ksysmgr** command to discover all the hosts that are managed by the HMCs in both the active and the backup sites. During the discovery process, the KSYS subsystem captures the configuration information of the active site and its relationship with the backup site and prepares the backup site to perform disaster recovery operations later.

During the initial discovery operation, the KSYS subsystem uses this configuration information to gather the list of VMs from all the host groups across sites and the corresponding disks for disaster recovery management. During any subsequent discovery operations, the KSYS subsystem scans the environment for any changes to the environment (for example, addition of a new VM, addition of a disk to VM, Live Partition Mobility (LPM) movement of a VM from one host to another host, and so on) and adapts to the modified environment.

The KSYS subsystem interacts with the HMC to retrieve the details about the disks of each VM and to check whether the VMs are currently set up for the storage devices mirroring. If the disks are not set up for mirroring properly, the KSYS subsystem notifies you about the volume groups that are not mirrored. All volume groups of a VM must be mirrored. Disks can be available over N-Port ID Virtualization (NPIV), virtual SCSI (vSCSI), and combination of all these modes.

The KSYS subsystem identifies and stores the Universally Unique Identifier (UUID) of the boot disk for each virtual machine during the discovery operation. The KSYS subsystem also stores the information about the corresponding replicated boot disks in the backup site. When you initiate a disaster recovery operation, the KSYS subsystem uses this information to boot the virtual machines with the corresponding boot disks on the paired host in the backup site. For example, if a virtual machine in the active site has multiple bootable disks, the KSYS subsystem restarts the virtual machine by using the corresponding boot disk in the backup site.

**Note:** The GDR 1.1 Service Pack 1, and later, identifies and stores the boot disk information only for POWER8 processor-based servers. The GDR solution requires HMC Version 8 Release 8.6.0 Service Pack 1 to support this feature. If your production environment contains an older version of host or HMC, the KSYS subsystem cannot store boot disk information and the virtual machines will restart in the backup site by using the first disk in the System Management Services (SMS) menu.

If the configuration is modified, for example, if a logical partition or a storage device is added, the KSYS subsystem rediscovers the active site, identifies the changes in the configuration, and marks the changes in its registries. The KSYS subsystem monitors this new environment for any disaster situations.

**Note:** For EMC storage subsystem, the Gatekeeper and Access Control Logix (ACLX) devices are ignored by the KSYS node during the discovery operation.

By default, the KSYS subsystem automatically rediscovers sites once in every 24 hours. You can change this period by modifying the **`auto_discover_time`** attribute. However, if you modified the configuration by adding or removing any resource, and you want the KSYS subsystem to rediscover the resources immediately, you can manually run the **ksysmgr discover** command. If you run the discovery operation for a site, the KSYS subsystem might take a few minutes to discover all virtual machines from all the host groups across both the sites and to display the output. To save time, you can run the discovery operation for a specific host group that contains the hosts that you modified.

**Tip:** To avoid configuration information loss on the KSYS node because of any node failure events, back up your current configuration settings as a snapshot after you complete the initial configuration of sites, hosts, host pairs, host groups, HMCs, and storage devices.

To discover resources in the KSYS configuration settings, complete one of the following steps in the KSYS LPAR:

- To discover all the resources across both sites, run the following command:

  ```
  ksysmgr discover site site_name
  ```

  The KSYS subsystem discovers all the hosts and virtual machines from all the host groups across both the sites. Therefore, it might take a few minutes to discover all the hosts and to display the output.
- To discover all the hosts in a specific host group, run the following command:

  ```
  ksysmgr discover host_group hg_name
  ```

**Related concepts**:

"Saving and restoring snapshots" on page 91
Use the **ksysmgr** command to save the configuration snapshots. The snapshots are saved in an XML format. When you create a snapshot, the **ksysmgr** command appends the date and time to the specified file name to follow the *`filename.DateTime`* name convention. By default, the snapshot files are saved in the `/var/ksys/snapshots` directory. However, you can specify the path where the snapshot files must be saved.

**Related reference**:

"Managing the system attributes" on page 85
After you synchronize the KSYS cluster by using the **ksysmgr sync ksyscluster** command, the KSYS subsystem sets up the default system-wide persistent attributes. The KSYS subsystem uses these system-wide persistent attributes for activities such as automatic rediscovery of the resources, notification of critical events, removal of duplicate notification.

## Verifying the configuration

After the KSYS subsystem discovers the resources, it monitors the entire environment for any disaster situations and verifies whether the configuration setting is valid for both sites. This verification is required to ensure that the backup site is ready to host the workloads from the active site during a site switch for disaster recovery operation. If you have set a priority to specific virtual machines, the verification operation is initiated for the virtual machines that have the highest priority.

To validate the configuration setting in both sites, complete one of the following steps in the KSYS node:

- To validate the configuration in the active site, run the following command:

  ```
  ksysmgr verify site Site1
  ```

  The KSYS subsystem validates the configuration settings on all the hosts and virtual machines from all the host groups across both sites. You can perform both the discovery and verification operations by running the following command:

  ```
  ksysmgr discover site site_name verify=yes
  ```

- To validate the configuration for a specific host group, run the following command:

  ```
  ksysmgr verify host_group hg_name
  ```

  or,

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

After the validation is complete, you can run the following command to query the IBM.VMR_LPAR resource class to ensure that the virtual machines are ready to be moved if a disaster occurs:

```
lsrsrc IBM.VMR_LPAR
```

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

```
        Name               = "xxx"
        LparUuid           = "59C8CFxx-4Bxx-43E2-A0CE-F028AEB5Fxxx"
        LparIPList         = {}
        SiteCleanupTastList = {}
        ActiveSiteID       = 80708xxxx
        LCB                = {  }
        BootDiskList       = {}
        CecUuid            = "6ce366c5-f05d-3a12-94f8-94a3fdfcxxxx"
        ErrMsg             = ""
        Phase              = "READY_TO_MOVE"
        PhaseDetail        = 4194305
        Memory             = "4352"
        Processors         = "0.1"
        ActivePeerDomain   = "vmdr"
```

If an error occurs during configuration validation, review the error details in the ErrMsg field. The Phase field is set as READY_TO_MOVE after a successful verification operation.

## Daily checks by KSYS

The KSYS subsystem checks the active site in the GDR environment daily to ensure that any change in the configuration setting or resources is discovered and verified by the KSYS subsystem. The daily verification checks ensure that the workloads are always ready to be moved to the backup site if any disaster occurs.

The following table lists the checks that are performed by the KSYS subsystem:

*Table 6. Checks performed by KSYS*

| Check type | Description |
|---|---|
| Check the capacity of the backup host to host a disaster recovery failover | In the GDR solution, the backup hosts must be in the standby mode. The backup hosts must have enough capacity (CPU cores and memory) to host all the virtual machines from the partner host in the active site.<br>The KSYS subsystem adds all the CPU and memory that are consumed by the virtual machines in the active site and compares the amount to the capacity of the host in the backup site. If the backup site does not have enough capacity, the KSYS subsystem generates the `HOST_CAPACITY_CHECK` event notification. Therefore, you must ensure that the backup host has enough capacity available. |
| Pre-verification and post-verification scripts | If you configure scripts by using one of the following commands, the KSYS subsystem executes these scripts during daily validation:<br>• `ksysmgr add script entity=site pre_verify=`*`script_path`*<br>• `ksysmgr add script entity=site post_verify=`*`script_path`*<br>These customized scripts are run before or after the discovery and verification processes depending on how you specified the scripts to be run. For more information about scripts, see "Running scripts for additional checks" on page 99 |
| Disk space check for the `/tmp` and `/var` directories | The KSYS operations require enough space in the `/var` directory to log the background daemon and storage messages. The KSYS operations also uses the `/tmp` directory to add or delete temporary files.<br>The KSYS checks whether the space usage in the `/var` and `/tmp` directories has reached or exceeded 90%. If the space usage in the `/var` or `/tmp` directories exceeds 90%, the KSYS generates the `TMP_USAGE_CROSSED_90PERCENT` or `VAR_USAGE_CROSSED_90PERCENT` event notification. |
| EMC storage validation | During storage agent verification, the KSYS subsystem queries EMC storage devices on both the active and backup sites, and checks for the composite groups and the disks that are associated with the composite groups.<br>The KSYS subsystem checks whether the composite group exists. It also checks the composite group state, the composite group disk, the RDF group disk, and the replicated disk size. In addition, the KSYS also pings the local storage devices. If any discrepancy in the storage validation is detected, the KSYS generates the `STORAGE_UNREACHABLE` event notification. |

*Table 6. Checks performed by KSYS (continued)*

| Check type | Description |
|---|---|
| Disaster recovery verification | The KSYS subsystem performs this verification daily to check the overall health of the active and backup sites if a disaster recovery operation is required at any time.<br>This verification consists of the following checks:<br>• Logical partition validation to check whether the backup hosts are capable of running the workloads.<br>• CPU and memory resources validation to check whether the backup hosts have required amount of CPU cores and memory resources.<br>• Storage validations<br>• Network validations |

**Related concepts**:

"Notification for the KSYS events" on page 105
The KSYS subsystem tracks various events that occur in the environment and saves the information in log files. The KSYS subsystem also sends emails and text notifications to the administrator if the contact information is registered on the KSYS configuration by using the **ksysmgr add nofity** command.

# Recovering the virtual machines during a disaster

After the verification phase, the KSYS continues to monitor the active site for any failures or issues in any of the resources in the site. When any planned or unplanned outages occur, if the situation requires disaster recovery, you must manually initiate the recovery by using the **ksysmgr** command that moves the virtual machines to the backup site.

## Failover rehearsal of the disaster recovery operation

The KSYS subsystem can perform a failover rehearsal at the backup site in the disaster recovery environment, without disrupting the production workloads or the storage replication from the active site to the backup site.

The failover rehearsal feature is supported for the following storage subsystems:

- EMC SRDF family of storage system (for example, VMAX)
- SAN Volume Controller (SVC) and Storwize family of storage systems
- DS8000 storage system

**Notes:**

- The DR failover rehearsal feature is supported only for disaster recovery and not high availability.
- The DR failover rehearsal feature is not supported for heterogeneous storage systems and shared storage model of deployment.
- The DR failover rehearsal feature is not supported for Hitachi storage systems.

The failover rehearsal feature is useful to rehearse the disaster recovery operation without performing a real DR failover and to test the readiness of the entire environment. It provides you the flexibility to perform DR testing more frequently. These failover rehearsal operations allow you to perform various workload-related tests that include write operations in the virtual machines (VMs) for a longer period of time.

You can test the disaster recovery operation at host group level or at the entire site level. In a test disaster recovery operation, the virtual machines in the active site continue to run the existing workloads and are not shut down. The storage replication between the storage devices in the active site and the backup site is also not impacted by the test disaster recovery operation.

Since the VMs continue to run in the active site when duplicate test VMs are started on the backup site as part of the failover rehearsal operation, you must ensure network isolation between the active site VMs and the test VMs that are started on the backup site. You can change the VLANs of the test VMs by using the network attribute of the **ksysmgr modify** command that helps to isolate the network on the backup site or you can use some other method to achieve network isolation between the sites.

The following figure shows a high-level flow diagram for a DR failover rehearsal as compared to a regular DR failover operation.

© Copyright IBM Corp. 2017, 2019

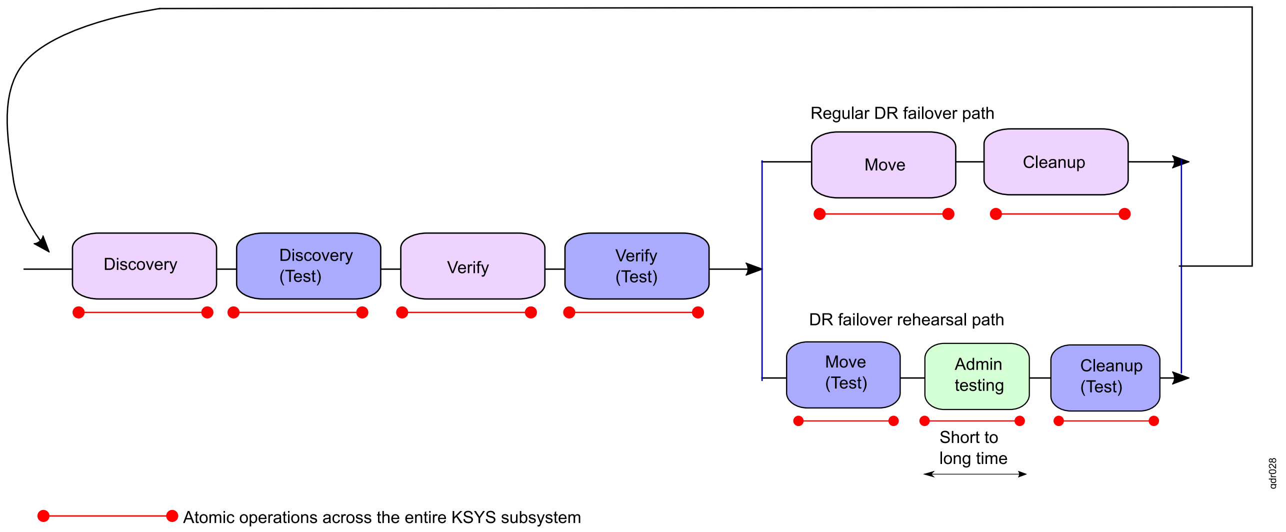

*Figure 20. DR failover rehearsal flow*

**Notes:**

- You can perform DR failover rehearsal for a single host group, multiple host groups, or the entire site. If you are running the failover rehearsal operation at host group level, you must start the operation sequentially for each host group. If you are running the failover rehearsal operation at site level, all the host groups are handled in parallel.
- After the DR failover rehearsal starts for a host group, you cannot perform regular move operations on the same host group or on the entire site for the entire duration of DR failover rehearsal time.
- You must not perform Live Partition Mobility operation of virtual machines into or out of the host group that is under the test operation for the entire duration of DR failover rehearsal. When the LPM operation is started directly through the HMC, the KSYS subsystem cannot stop those operations. Therefore, if a host group is under DR failover rehearsal and a VM is moved into or out of the host group as part of the LPM activity, the results are unpredictable.
- You cannot change the configuration settings for host groups that are under DR failover rehearsal operation.
- For periods of time that are marked as atomic operation in the figure, you cannot start DR failover rehearsal operation or a regular move operation for any other host group or site.
- You can perform regular move operation for the other host groups other than the host groups that are in DR failover rehearsal mode. Thereby you can recover host group HG1 while simultaneously testing host group HG2.
- If a real disaster occurs during a planned DR test time and you want to perform a real recovery related to a disaster, you can quit the DR failover rehearsal mode by executing the cleanup step.
- The cleanup step in regular DR move operation is different than the DR test move operation. If failover rehearsal operation is performed, you must manually perform the cleanup steps.
- For a regular unplanned DR move operation, the cleanup step is mandatory. The cleanup step must be performed after the move operation is complete at the earlier-active site.
- When the test-cleanup operation is in progress, you cannot perform any other operations on the KSYS subsystem. You cannot work on the test VMs also because the test VMs are deleted as part of the cleanup process.
- When you perform the rehearsal cleanup operation, the VM console that is opened in the HMC must be closed.

- If you want to save a snapshot of the current configuration settings of your KSYS environment, you must save a detailed snapshot so that the tertiary disk values are also retained in the snapshot. If you want to get the tertiary disk information after restoring a basic snapshot, you must run the discovery operation with the **`dr_test`** flag.
- The error messages for any failover rehearsal operations are displayed only in the output message of the operation. The error messages are not displayed in the output of the **ksysmgr query system status** command.

The following figure shows an example of failover rehearsal of all the virtual machines in a host:

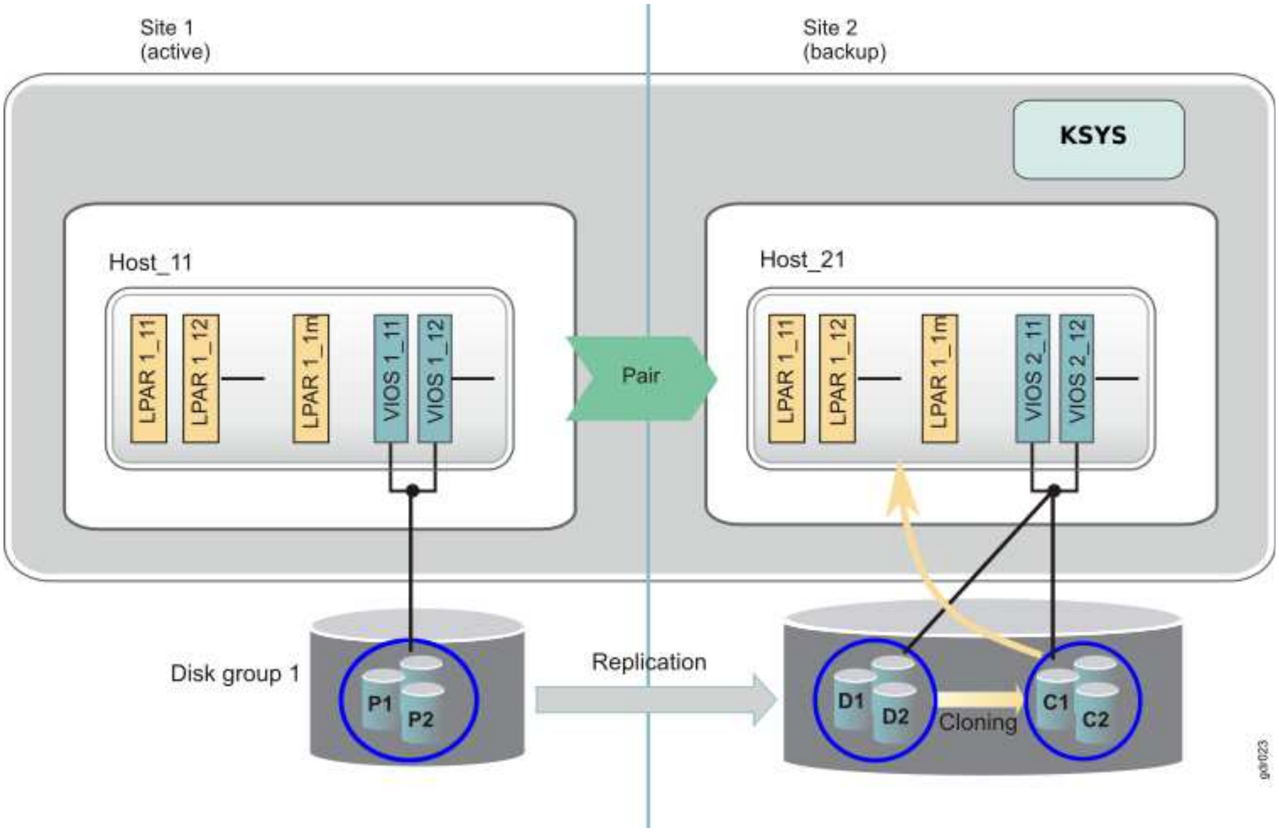

*Figure 21. Example for failover rehearsal of the disaster recovery operation*

### Storage prerequisite

The storage administrator must have mapped all the hosts in the backup site to the backup storage disks (D1, D2, and so on). The storage administrator must also have created a set of clone disks (C1, C2, and so on) that are of the same number and size as the active site storage disks (P1, P2, and so on) and backup storage disks (D1, D2, and so on). The cloning (D1-C1, D2-C2, and so on) must be started from the backup storage disks to the clone disks. The storage administrator can set up the cloning relationship by using interfaces (command-line or graphical user interface) that are provided by specific storage vendors. Refer to documentation from the storage vendor for more details about allocation of storage disks and establishing relationship with secondary copy of data on the backup site. The following table lists the tools that are necessary for establishing cloning relationship for various storage systems.

*Table 7. Storage vendors and the corresponding cloning feature*

| Storage vendor | Clone feature | Reference |
|---|---|---|
| EMC SRDF family of storage system (for example, VMAX) | `symclone` | EMC Solutions Enabler CLI User Guide |
| SAN Volume Controller (SVC) and Storwize family of storage systems | `Flashcopy` | Managing Copy Services in SVC |
| DS8000 storage system | `Flashcopy` | Redbook: IBM System Storage DS8000 Copy Services Scope Management and Resource Groups |

## Performing the DR failover rehearsal operation

The disaster recovery (DR) failover rehearsal operation contains the following steps:

*Table 8. DR failover rehearsal flow*

| S. No. | Step | Administrator action | KSYS operations |
|---|---|---|---|
| 1. | Pre-test environment setup (must be performed only once) | 1. Configure the disks that are required for cloning and will be used as third copy of disks.<br>2. Ensure correct disk connectivity and mappings on the backup site.<br>3. Establish clone relationships between second and third copies of disks on the backup site.<br>4. Ensure network isolation for VMs that are started on the backup site and for VMs on active site. | None. |
| 2. | Discover the DR test environment | Run the **ksysmgr discover** command in test mode by entering the following command:<br>`ksysmgr discover host_group\|site` *name* `dr_test=yes` | The KSYS subsystem discovers the test environment and prepares the environment for the test move operation.<br>During this step, data from the second copy of disks is copied to third copy of disks. After you run this step, you must run commands that are provided by the specific storage vendor to check the completion of the copy operation. You can start the failover test operation only after the copy operation is complete. |
| 3. | Verify the test environment | Run the **ksysmgr verify** command in test mode by entering the following command:<br>`ksysmgr verify host_group\|site` *name* `dr_test=yes` | This step checks the test environment and confirms that the environment is ready for the test-move operation. You can perform this step any number of times, if required. |
| 4. | Start the VMs in the DR test environment | Run the **ksysymgr move** command in test mode by entering the following command:<br>`ksysmgr move host_group\|site` *name* `dr_test=yes` | • This step starts the VM on the backup site systems by using the third (test) copy of disks. The KSYS subsystem displays any issues and error messages.<br>• The move operation changes storage level mapping. The move operation maps third copy of disks to target VIOS and unmaps the secondary disks. The cleanup operation unmaps the third copy of disks and maps the secondary disks to target VIOS. |
| 5. | Perform testing activities | • Access the newly started VMs and perform all the read or write tests in the VMs.<br>• Ensure that network traffic is isolated during this period of test. | None. |

*Table 8. DR failover rehearsal flow (continued)*

| S. No. | Step | Administrator action | KSYS operations |
|---|---|---|---|
| 6. | Cleanup of DR test environment | Run the `ksysymgr cleanup` command:<br>`ksysmgr cleanup host_group\|site` *`name`* `dr_test=yes` | • Deletes all the test VMs and configures the storage settings in VIOS to map back to secondary disks.<br>• Retains the clone relationships so that DR test can be performed again, if required, after performing a new discovery in test mode.<br>• If the test move operation fails, you must run the test cleanup operation to correct the configuration settings for the next DR operation. Some failures might need manual intervention and troubleshooting. |

The detailed description of DR failover rehearsal flow follows:

**1. Pre-test environment setup (must be performed only once)**

Before you enable the DR test mode, you must set up the following prerequisites once:

1. Ensure that disks that are related to third clone copy exist on the backup storage subsystem. Establish clone relationship between the second copy disks and the third copy (test) disks. You can perform this step once and retain disks for any future testing.
2. Ensure that connectivity to all the relevant disks (second and third copy disks) exist to the VIOS on the backup site systems because the backup site hosts and Virtual I/O Servers cannot access the active site disks. Poor connectivity can cause problems during the boot operation of test VMs.
3. Ensure proper network isolation of the DR test environment from the active site VMs. If network is not isolated, the DR rehearsal operation can impact the production environment. The KSYS subsystem supports specific VLANs for setting up the DR test environment. You can plan and predefine the VLANs that must be deployed with test VMs when the DR testing is performed.

**Note:**

- The clone disks configuration is required on both sites if you intend to perform DR testing on any of the two sites. The KSYS subsystem enables the DR failover rehearsal mode only on the backup site. But backup site can be any one of the two sites, depending on the previous failover operation.
- Virtual switches and VLAN IDs that are specific to DR failover rehearsal operation are created on the target site based on mapping details as a part of the DR test operation. Network issues that are related to these fields are not validated in the verification operation. Therefore, you must specify VLAN IDs that have proper physical connectivity.
- Only a single clone relationship must exist for each secondary disk during the test-discovery operation. The secondary disk cannot have a multiple-clone relationship with any other disks.

  For DS8000 storage system, the FlashCopy clone copies are already created for Global Mirror type of data replication that are used internally within the storage system. The KSYS subsystem differentiates the clone copy disks with Global Mirror FlashCopy disks by using the **`TargetWriteEnabled`** field.

  Use the following command to create the FlashCopy clone relation between secondary disks and clone disks:

  ```
  mkflash -persist -dev sid LUN1:LUN2
  ```

**2. Discover the objects and relationships that are specific to the test environment**

In this step, you request the KSYS subsystem to discover and enable the DR failover rehearsal or DR test mode. The KSYS subsystem displays information about the DR test mode when you run various reporting commands.

To discover and initialize the backup environment in the DR failover rehearsal mode, enter the following command:

```
ksysmgr discover host_group|site name dr_test=yes
```

**Note:**

- During this step, some of the setup activities are also performed. The KSYS subsystem requests the storage subsystem to copy secondary disk data (D2) to third disk (D3). Because the copy operations can take few minutes, depending on the number of disks and sizes of the disks, you must use commands that are provided by the specific storage vendors to check the completion of the copy operation. The following table lists some command examples for various storage systems.

*Table 9. Storage vendors and corresponding command examples*

| **Storage vendor** | **Command example** |
|---|---|
| EMC SRDF family of storage system (for example: VMAX) | `symclone verify -f` *`file_with_pairs`* |
| SAN Volume Controller (SVC) and Storwize family of storage systems | `svcinfo lsfcmap \| grep` *`disk_UUID`* |
| DS8000 storage system | `lsflash -l 1023:1024 ...` |

- The KSYS subsystem also identifies all the necessary objects for the test environment and indicates whether the environment is ready for the test-verify or test-move operation. If any failure occurs, you must correct the issue and retry the discovery operation.
- The test-discovery step initiates the disk copy operation. Therefore, the test-discovery step must be repeated if there is long elapsed time between previous test-discovery operation and the test-move operation.
- Any configuration changes after a test-discovery operation invalidates the previous discovery operation. If configuration settings are changed after a test-discovery operation, you must perform another test-discovery operation.
- You can run the test-discovery operation as many times as needed.

For information about how to troubleshoot common problems in the test-discovery step, refer to the following sections:

- "The test-discovery step in the DR failover rehearsal operation is failing with the error message: Tertiary storage copy is missing." on page 110
- "The test-discovery step in the DR failover rehearsal operation is failing with the error message: Storage agent is not accessible." on page 110

**3. Verify the test environment readiness for test-move operation**

Validate the backup environment for DR test mode by entering the following command:

```
ksysmgr verify host_group|site name dr_test=yes
```

This step checks the test environment and confirms that the environment is ready for the test-move operation. If any failure occurs, you must correct the issue and retry the verify operation. You can perform this step any number of times, if required.

**Notes:**

- The test-verify operation validates the configuration settings only at storage level. Issues, such as target VIOS down and other resource gaps are not detected as a part of this operation.

Hence, you must invoke a regular verify operation for verification of the entire environment. You can also refer to the periodic automatic validations that occur once in every 24 hours.

- You must validate the VLAN and virtual switch network mapping policies and attributes across both the sites. If the specified configuration is not valid in the backup site, the disaster recovery rehearsal operation fails.

**4. Start the VMs in the DR test environment**

In this step, you request the KSYS subsystem to restart all the VMs for the entire site or a specific host group in the test environment. The KSYS subsystem attempts to restart the VMs on the backup hosts in backup site. To start the VMs in the backup hosts in a DR test mode, enter the following command:

```
ksysmgr move host_group|site name dr_test=yes
```

The KSYS subsystem initiates the test-move operation by breaking the cloning relationship (D1-C1, D2-C2, and so on) in the storage subsystem. The cloned storage disks are mapped to the newly started virtual machines. The storage replication between the active site storage disks and the backup site storage disks (P1-D1, P2-D2, and so on) is retained.

The KSYS subsystem displays any errors that occur while starting the virtual machines. The KSYS subsystem restarts the VMs according to the various policies that are configured for the regular DR environment. For example, if you configured for flexible capacity, the VMs are started by using the flexible capacity policies. Similarly, if you have configured the network mapping policy, the VMs will be started with the configured test VLANs.

After the VMs are started in the backup site, you must manage the test environment. Because the backup systems are used for failover test, regular move operation cannot be performed for the host group or the site that is undergoing the test operation. You can perform regular planned or unplanned DR failover only after the test VMs are cleaned up and the KSYS subsystem operates in normal mode.

**5. Perform testing activities**

In this step, you can use the test VMs that are started as part of the DR rehearsal process. You can perform both read and write test operations on the third cloned copy disks that are mapped to the VMs. If you have ensured network isolation, you can also start the workload in VMs and run test operations as required to validate the middleware on the backup site.

There is no time limit for performing the failover rehearsal operation. However, during the DR test mode period, you cannot perform a regular DR move operation from the active site to the backup site because the backup hosts are being used up by the DR rehearsal process. You can quit the DR failover rehearsal mode, clean up the DR test environment, and then perform the regular DR move operations.

**6. Cleanup of DR test environment**

After your testing activities on the VMs are complete, you can request the KSYS subsystem to clean up the test environment and resume the regular DR mode for the host group or site that is undergoing the test operation.

To clean up the DR test environment and to retain the normal DR backup state, enter the following command:

```
ksysmgr cleanup host_group|site name dr_test=yes
```

When you perform the cleanup operation in the DR failover rehearsal mode, the KSYS subsystem deletes the test VMs and resets the environment back to point to the second copy of disks and moves the environment to regular DR mode.

When you clean up the virtual machines in the backup site, the following steps are performed by the KSYS subsystem:

- Third copies of test disks (C1, C2, and so on) are unmapped from the virtual machines.

- The previously mapped second copies of storage disks (D1, D2, and so on) are mapped again to the virtual machines.
- The cloning is restarted between the storage disks (D1-C1, D2-C2, and so on) for future test.

After the cleanup operation is complete, the KSYS subsystem starts operating in normal mode and any other operation (for example, discovery, test-discovery, verify, test-verify) can be performed for the specific host group.

## Flow chart for disaster recovery

After the KSYS configuration is verified, the KSYS monitors the active site for any resource failures in the site and notifies you about the unplanned outage.

The following flow chart provides a summary of the disaster recovery steps:

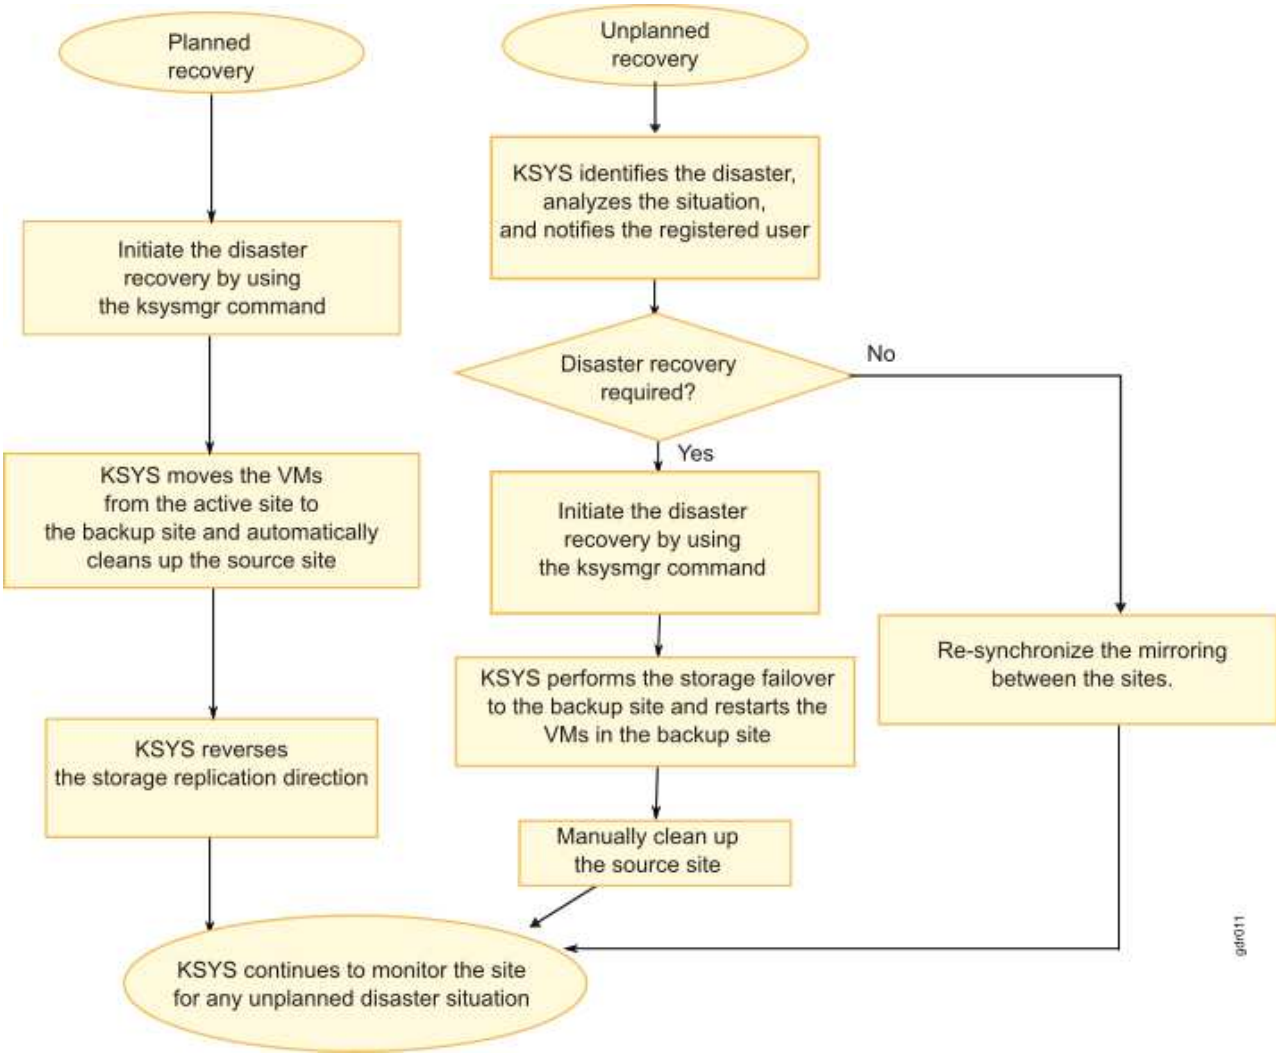
Planned recovery
Unplanned recovery
KSYS identifies the disaster, analyzes the situation, and notifies the registered user
Initiate the disaster recovery by using the ksysmgr command
Disaster recovery required?
No
Yes
KSYS moves the VMs from the active site to the backup site and automatically cleans up the source site
Initiate the disaster recovery by using the ksysmgr command
Re-synchronize the mirroring between the sites.
KSYS performs the storage failover to the backup site and restarts the VMs in the backup site
KSYS reverses the storage replication direction
Manually clean up the source site
KSYS continues to monitor the site for any unplanned disaster situation
gdr011

*Figure 22. GDR solution: Disaster recovery process*

## Initiating the disaster recovery

In a planned recovery, you can initiate the site switch by using the **ksysmgr** command. In an unplanned outage, the KSYS subsystem analyzes the situation and notifies you about the disaster or potential disaster. Based on the information about the disaster, you can determine whether a site switch is required.

When issues occur in the replication of storage subsystem, the KSYS preserves the storage consistency information and notifies you about the issue by using the specified notification method. For example, email, text message, or pager email. You can evaluate the health of the system by using the KSYS information, HMCs, and applicable storage vendor tools to determine whether the situation requires a true disaster recovery.

If the situation requires disaster recovery, you must manually initiate the recovery by switching the site from the active site to the backup site. When you switch sites, the virtual machines or the LPARs that were running on the active site are restarted automatically on their corresponding hosts in the backup site that were paired earlier.

For a planned recovery, the storage replication direction is reversed from the current active site to the earlier active site. After the site switch is complete, the KSYS automatically cleans up the source site from where the switch was initiated. For an unplanned recovery, the storage is not replicated back to the previously active storage. Therefore, after the problems in the previously active site are resolved, you must manually resynchronize the storage. A resynchronization is necessary for any subsequent planned move operation. Also, you must manually clean up the source site from where the switch was initiated after the HMC and hosts become operational.

If the virtual machines, which are moved from the active site to the backup site within the paired hosts, have multiple boot disks, the KSYS subsystem uses the stored information about the boot disk list during the discovery operation to boot the virtual machines by using the corresponding boot disks in the backup site. The boot disk information is available in the `IBM.VMR_LPAR` resource class for each virtual machine. Before you initiate a disaster recovery operation, ensure that the boot disks in the active site are replicated to corresponding disks in the backup site and the KSYS subsystem contains HMC Version 8 Release 8.6.0 Service Pack 1. This feature is supported only for POWER8 processor-based servers. If your production environment contains an older version of host or HMC, the KSYS subsystem cannot store boot disk information and the virtual machines will restart in the backup site by using the first disk in the System Management Services (SMS) menu.

To recover the virtual machines, complete the following steps in the KSYS node:

1. Optional: Verify whether a disk pair and disk group exist before you initiate the recovery.
   a. To verify that the disk pair exists between the active site and the backup site, run the following command:
      ```
      ksysmgr query disk_pair
      ```
   b. To verify that the composite disk group exists in both the active site and the backup site, run the following command:
      ```
      ksysmgr query disk_group
      ```
2. Switch the site from the active site (`Site1`) to the backup site (`Site2`) for planned or unplanned recovery by running the following command:
   ```
   ksysmgr move site from=Site1 to=Site2 dr_type=planned|unplanned cleanup=yes|no
   ```

   If you do not specify the **`dr_type`** attribute, the **ksysmgr** command starts a planned recovery by default. The logical partitions automatically restart in the backup site.

   In case of a planned recovery, the KSYS automatically cleans up the source site from where the switch was initiated. In case of an unplanned recovery, you must manually clean up the source site after the HMC and hosts become operational. However, if you specify the **`cleanup=no`** attribute during a planned recovery, the virtual machines are not cleaned up from the source site.

   **Note:** If you start a move operation for a specific host group, and the move operation fails, the state of the host group becomes `RECOVERY_VM_ONLINE`. In this case, you cannot start the move operation for the entire site unless the failed host group is recovered completely. Therefore, you must recover the failed host groups before you attempt a move operation for the entire site. Otherwise, you can continue to move the other host groups to the backup site individually.

3. For an unplanned recovery, clean up the source site (Site1) manually by running the following command:

   ```
   ksysmgr cleanup site Site1
   ```

   For an unplanned recovery, you can also clean up each virtual machine separately by running the following command:

   ```
   ksysmgr cleanup vm vmname
   ```

4. For an unplanned recovery, resynchronize the storage data from the active site to the backup site by running one of the following commands:

   ```
   ksysmgr resync site active_site_name
   ksysmgr resync host_group active_hg_name
   ```

   If the unplanned move operation was at site level, you must run the **ksysmgr resync** command at site level. Similarly, if a virtual machine was moved to the backup site in an unplanned move operation at host group level, you must run the **ksysmgr resync** command at host group level.

## Recovering the failed virtual machines

After the site switch operation is complete, if some virtual machines were not moved successfully, you can use the **ksysmgr** recover command to move the failed virtual machines to the new active site.

When you run the **ksysmgr recover** command for a specific host group, the KSYS subsystem attempts to move all the failed virtual machines from the current site to the target site.

**Note:** The **ksysmgr recover** command can be used only when the reverse replication of storage is successful. You can use this command only at host group level.

When you perform the recovery operation after a move operation of a VM has failed, the KSYS subsystem provides an option to use the current or previously saved LPAR or virtual machine (VM) profile to retry the recovery operation. The LPAR profiles of VMs are backed up regularly based on the configuration settings. Additionally, each time the administrator changes the configuration settings for an LPAR and runs the discovery operation, a backup of the LPAR profile file is created with the relevant timestamp.

When you run the **ksysmgr recover** command, the command interface prompts you whether you want to recover the virtual machine to the backup site by using its current or default LPAR profile or an LPAR profile from the backup profiles list. If you want to restart the virtual machine with a previously backed up LPAR profile, select yes, and then the command interface prompts you to select the LPAR profile file based on the time stamp. The KSYS subsystem fetches the backup profile and uses the configuration settings that are specified in the selected LPAR profile to restart the virtual machine. If you select no as the response for the command interface prompt, the virtual machine is restarted on the backup host with the existing configuration settings of the LPAR.

After the site switch operation is complete, if the storage is successfully switched to the backup site, the failed virtual machines in the previously active site are not linked to the most recent disks. Therefore, the **ksysmgr recover** command moves the virtual machine configuration without affecting the storage because the storage is already moved.

To recover the failed virtual machines after the move operation is complete, enter the following command:

```
ksysmgr recover host_group host_group_name
```

## Moving the virtual machines by using the force option

In some scenarios, when you modify the KSYS configuration, the KSYS subsystem might have discovered the configuration change as part of its daily check. However, the resources might not have been verified to check whether the virtual machines can be moved to the backup site. In these scenarios, if an unplanned recovery is required, you can use the force option to move the virtual machines from the active site to the backup site.

To move the virtual machines to the backup site by using the force option, run the following command:

```
ksysmgr move -f from=site1 to=site2  force=true
```

**Note:** If a disaster occurs in the source site and the virtual machines are in the `init` state, you must initiate the disaster recovery operation with the force option by running the following command:

```
ksysmgr move -f from=site1 to=site2  force=true dr_type=unplanned
```

# Administering GDR

These topics provide a list of the tasks you can perform to maintain and monitor the resources in the Geographically Dispersed Resiliency for Power Systems solution.

## Generating the KSYS system report

After you add all the required resources to the KSYS configuration, you can generate the KSYS system report that summarizes the entire KSYS environment. Instead of running different commands to query each resource, you can generate a consolidated report to view all the added resources in the KSYS configuration.

To generate the KSYS system report, run the following command:

```
ksysmgr report system
```

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

```
Current enviornment:
Number of sites:2
        Site_1
        Site_2
Number of storage agents:2
        salocal
        saremote
Number of HMCs:2
        Site1_HMC1
        Site2_HMC1
Number of hosts:8
        xxx_8233-E8B-1000ADP
        xxx-8233-E8B-06DA57R
        xxx-9179-MHD-SN107895P
        xxx-8408-E8D-SN21ED67T
        rxxx1m6
        xxx1m5
        xxx_9179-MHD-105E67P
        xxx-8233-E8B-06DA56R
Number of VIOS:13
        doit2v2
        doit1v2
        gsk1v2
        doit1v1
        rar1m5v1
        kumquatv1
        orangev1
        doit2v1
        gsk1v1
        rar1m6v1
        rootbeerv1
        orangev2
        kumquatv2
Number of VMs:201
Total managed cores: 30
Total managed memory: 76
```

You can also generate an organized report by using the verbose option:

```
ksysmgr -v report system
```

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

```
active site=site_1
backup site=site_2
site=site_1
     hmc=hmc1
          host=hosta
               number vm: 50
               total memory: 22
               total cores: 34
          host=hostb
               number vm: 40
               total memory: 22
               total cores: 31
     hmc=hmc2
          host=hostc
               number vm: 40
               total memory: 22
               total cores: 11
     storage=mainstoragea
     storage=secstorageb

site=site_2
     hmc=hmc3
          host=host2a
     hmc=hmc4
          host=host2c
          host=host2b
     storage=mainstorage2a
```

## Configuring the flexible capacity policies

By using the flexible capacity policies, you can configure the memory and processor resources such that you can move the virtual machines (VMs) from the source host to the target host even when the target host does not have the same quantity of memory and processor resources as the source host. You can set the flexible capacity policies for the backup site.

You can use the flexible capacity option to start the virtual machines on the backup site with a different capacity as compared to the active site. The active site resources are used as reference to calculate the percentage of resources that must be assigned during recovery of virtual machines on the backup site.

The flexible capacity policies can be set at virtual machine, host, host group, or site level.

The flexible capacity policies can be used in the following cases:

- To perform the disaster recovery operations with different amount of resources on the backup site.
- To test the disaster recovery operation by using the failover rehearsal function with fewer or more resources on the backup site.
- To pair hosts across the active site and the backup site even when the available resources for both hosts differ.

When a logical partition (LPAR) is created in HMC, a profile is created for the LPAR that includes the resource limits such as **minimum**, **desired**, and **maximum** memory and processors that you want to allocate for that LPAR. If flexible capacity policies are not specified, by default, the **desired** value of resources are used by KSYS.

The following figure shows the various flexible capacity policies:

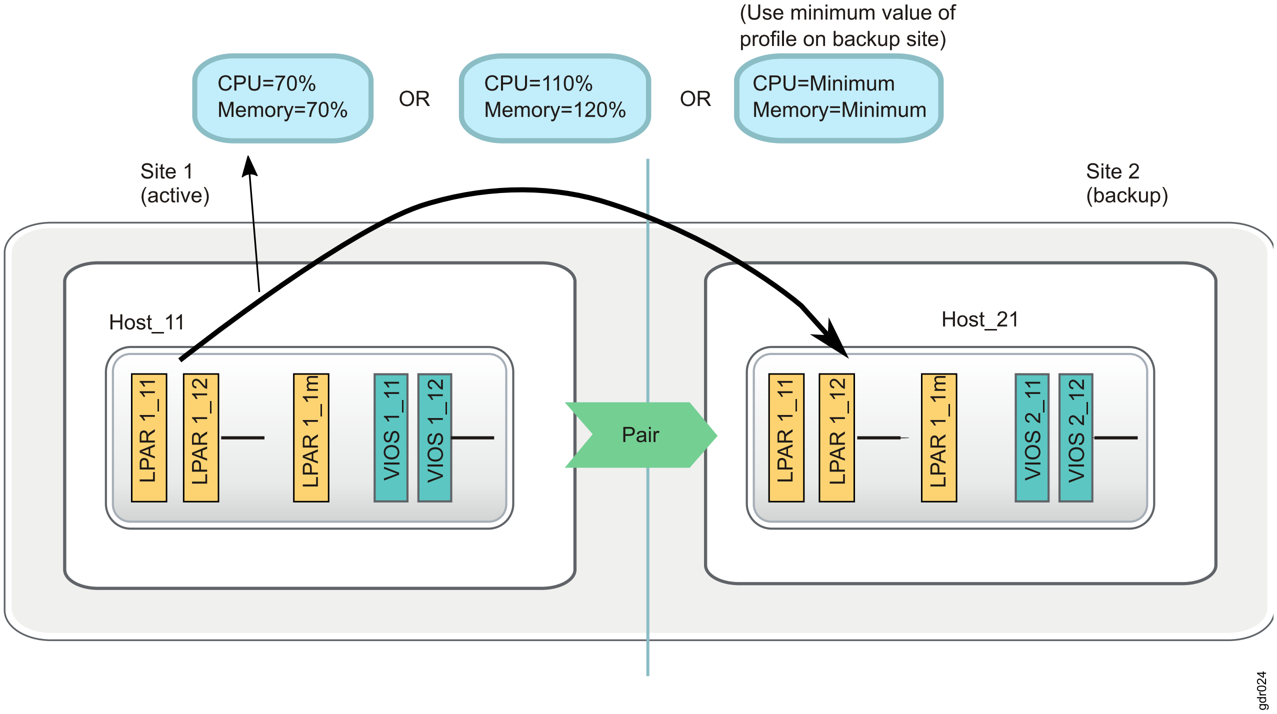

*Figure 23. Flexible capacity management for backup site*

The active site always uses the required resources that are defined in the LPAR profiles. You can set up flexible capacity policies only for the backup site. Various policies can be deployed by using the following tunable parameters:

**Percentage of the existing capacity value**

This policy specifies the percentage of resources that must be allocated to the VM after recovery on the backup site as compared to the active site.

The following examples describe the percentage based resource allocation:

- If a VM is running on the active site with 10 GB memory and 2 dedicated processors, and if you specify the flexible capacity policy with 50% CPU resource and 70% memory resource, then after the recovery of that VM on the backup site, the VM will be running with 1 processor and 7 GB memory. Memory is calculated and rounded off to the nearest multiples of memory region size of the target host. If dedicated processor is allocated to a VM, the calculated values are rounded off to the closest decimal.
- If a host is running 10 VMs, and each VMs has 4 GB memory and 1.0 shared processor units each on the active site, and if you specify the flexible capacity policy with 50% CPU resource and 50% memory resource at the host level, then after the recovery of that host on backup site, all the VMs will be running with 2 GB memory and 0.5 processor units each.

Similarly, we can configure the flexible capacity values at host group level and site level also. The percentage value that you specify in this policy must always be greater than 1.

**Minimum value**

This value sets the resource value of backup site LPAR (when the LPAR is started on the backup site during the recovery operation) to the minimum value that is defined in the LPAR profile on the active site.

**Current desired**

This value sets the resource value of backup site LPAR (when the LPAR is started on the backup site during the recovery operation) to the desired value that is defined in the LPAR profile on the active site.

**None** Resets the existing value. When you set the value to `none`, VM capacity management does not follow any flexible capacity management. CPU and memory resources across the active site and backup site match the values that are defined in the LPAR profile.

**Notes:**

- When you set the reduced or increased capacity value to a percentage of the existing capacity value, the resulting value must be within the range of the minimum and maximum values that are specified in the LPAR profile in the HMC. Otherwise, the **ksysmgr** command will return an error.
- The priority of virtual machine is also considered when the virtual machine is moved from the active host to the backup host with reduced capacity.
- The precedence order of the flexible capacity policy is: virtual machines, hosts, host groups, and then site.

For example, if you set the following values for the **`memory_capacity`** attribute at various levels:

```
Site: memory_capacity = "none"
     Host_group1: memory_capacity="50"
        Host1: memory_capacity="200"
        Host2: memory_capacity="none"
        Host3: memory_capacity="current_desired"
```

then, the virtual machines will be restarted in the target host by using the following quantities of resources:

- All the VMs in `Host1` will be restarted in the target host with 200% capacity.
- All the VMs in `Host2` will be restarted in the target host with reduced 50% capacity because the memory resource value is read from the host group (`Host_group1`) level.
- All the VMs in `Host3` will be restarted in the target host with the required quantity of resources (as specified in the LPAR profile). In this case, even though the flexible capacity value is configured at the host group level, the value is overwritten with the `current_desired` values.

The performance of the virtual machines might be decreased but the virtual machines can be restarted in the target host that has the reduced capacity. You can use the reduced capacity function during planned outages, where the virtual machines are temporarily moved to the target host and then moved back to the source host after the system maintenance or upgrade operation is complete. When the virtual machines are moved back to the home site, the original capacity settings of the virtual machine are retrieved and the virtual machines are run with the initial performance.

If you do not want the HMC to check and compare the resource capacity between the source host and the target host during the verify operation, you can set the **`skip_resource_check`** parameter to `yes`.

If you do not want the virtual machines to start automatically after they are moved to the target host, you can set the **`skip_power_on`** parameter to `no`.

### Virtual machine configuration

To configure reduced capacity settings for a group of the virtual machines, enter the following command in the KSYS LPAR:

```
ksysmgr modify vm vmname[,vmname2,...] | file=filepath
      [uuid=uuid]
      [host=hostname]
      [priority=low|medium|high]
      [memory_capacity=<(Whole Number > 1)> | minimum | current_desired | none]
      [cpu_capacity=<(Whole Number > 1)> | minimum | current_desired | none]
      [skip_resource_check=yes|no]
      [skip_power_on=yes|no]
```

For example:

```
ksysmgr modify vm file=ListOf3VMs.xml memory_capacity=60 cpu_capacity=50 skip_resource_check=yes
```

where, the `ListOf3VMs.xml` file contains a list of three virtual machines as shown in the following example:

```
<?xml version="1.0"?>

<KSYSMGR>
<VM><UUID>335406DA-CE5D-4A7A-B015-02B1F8D36A8C</UUID></VM>
<VM><NAME>VM039</NAME></VM>
<VM><NAME>VM038</NAME><HOST>HYDERABAD_HOST1</HOST></VM>
</KSYSMGR>
```

When you query the details about the `VM038` virtual machine, an output that is similar to the following example is displayed:

```
# ksysmgr query vm VM038

VM:
Name:                VM038
UUID:                335406DA-CE5D-4A7A-B015-02B1F8D36A8C
Host:                Site1_Host1
State:               READY
Priority:            Medium
skip_resource_check: Yes
skip_power_on:       No
memory_capacity:     60
cpu_capacity:        50
```

**Host-level configuration**

To configure reduced capacity settings for all the virtual machines in a group of hosts, enter the following command in the KSYS LPAR:

```
ksysmgr modify host hostname[,hostname2,...] | file=filepath
      [uuid=uuid]
      [memory_capacity=<(Whole Number > 1)> | minimum | current_desired | none]
      [cpu_capacity=<(Whole Number > 1)> | minimum | current_desired | none]
      [skip_resource_check=yes|no]
      [skip_power_on=yes|no]
```

For example:

```
# ksysmgr modify host Site1_Host1,Site2_Host1 memory_capacity=50 cpu_capacity=50
      skip_power_on=yes skip_resource_check=yes
```

When you query the details about the hosts, an output that is similar to the following example is displayed:

```
# ksysmgr query host

Name:                Site1_Host1
UUID:                21b4b05f-9b84-349c-9ce9-xxxxxxx
FspIp:   10.40.1.xxx
Pair:                Site2_Hostx
cpu_capacity:        50
memory_capacity:     50
skip_resource_check: Yes
skip_power_on:       Yes
Site:                Site1
VIOS:                pepsiv2
                     pepsiv
HMCs:                vmhmc5

Name:                Site2_Host1
UUID:                7d35be3a-a9b3-3cdf-a31e-xxxxxxx
FspIp:   10.40.1.xxx
```

```
Pair:                Site1_Hostx
cpu_capacity:        50
memory_capacity:     50
skip_resource_check: Yes
skip_power_on:       Yes
Site:                Site2
VIOS:                colav2
                     colav1
HMCs:                vmhmc3
```

### Host group-level configuration

To configure reduced capacity settings for all the hosts in a host group, enter the following command in the KSYS LPAR:

```
ksysmgr modify host_group hgname[,hgname2,...]
      [memory_capacity=<(Whole Number > 1)> | minimum | current_desired | none]
      [cpu_capacity=<(Whole Number > 1)> | minimum | current_desired | none]
      [skip_resource_check=yes|no]
      [skip_power_on=yes|no]
```

For example:

```
# ksysmgr modify host_group Site1_HG1,Site1_HG2 memory_capacity=50 cpu_capacity=50
      skip_power_on=yes skip_resource_check=yes
```

### Site-level configuration

To configure reduced capacity settings for all the virtual machines in the active site, enter the following command in the KSYS LPAR:

```
ksysmgr modify site sitename1[,sitename2,...] | file=filepath
      [name=newsitename]
      [memory_capacity=<(Whole Number > 1)> | minimum | current_desired | none]
      [cpu_capacity=<(Whole Number > 1)> | minimum | current_desired | none]
      [skip_resource_check=yes|no]
      [skip_power_on=yes|no]
```

For example:

```
# ksysmgr modify site Site_1 memory_capacity=current_desired cpu_capacity=none skip_power_on=yes
```

When you query the details about the site, an output that is similar to the following example is displayed:

```
# ksysmgr query site Site_1

Name:                Site_1
Sitetype:            ACTIVE
cpu_capacity:        current_desired
memory_capacity:     111
skip_resource_check: No
skip_power_on:       Yes
```

## Configuring the network mapping policy

A virtual LAN (VLAN) is created by assigning artificial LAN identifiers (VLAN IDs) to the datagrams that are exchanged through the physical network. Hosts that are located on the same VLAN represent a subset of the hosts that are located on the physical network. Hosts that belong to the same subnet allows communication without any physical device. The subnets are separated when the hosts in a subnet have different VLAN IDs.

When a virtual Ethernet adapter is created in an HMC, a virtual Ethernet switch port is configured simultaneously. The virtual machines within a host, which need to communicate with other virtual machines for workload operations, are configured to have the same VLAN IDs. Similarly, some virtual machines in your host environment might be isolated from other virtual machines through a private network and might have different VLAN IDs.

For example, consider a host in the active site that contains two virtual machines that use the following VLAN IDs: VLAN1, VLAN12, VLAN13, and VLAN5. If you want these virtual machines to start in the backup site with VLAN IDs: VLAN1, VLAN22, VLAN23, and VLAN5, you can set a VLAN policy that modifies the VLAN ID from VLAN12 to VLAN22, and from VLAN13 to VLAN23 when virtual machines are moved from the active site to the backup site. Therefore, when you move the virtual machines across sites, the virtual machines are restarted in the target site with the assigned VLAN IDs as shown in the following figure:

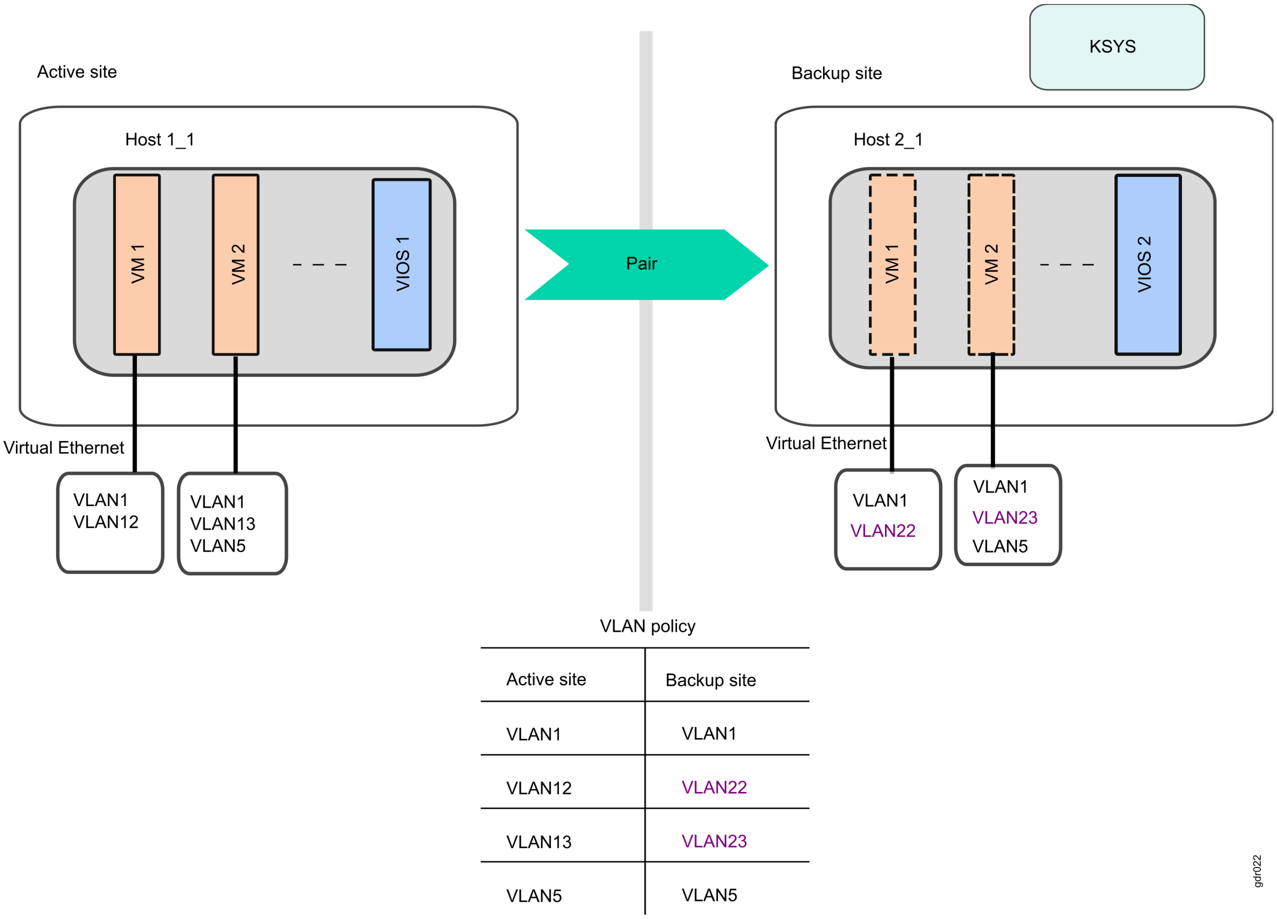

| Active site | Backup site |
|---|---|
| VLAN1 | VLAN1 |
| VLAN12 | VLAN22 |
| VLAN13 | VLAN23 |
| VLAN5 | VLAN5 |

*Figure 24. Example of network mapping policy configuration*

**Notes:**

- You can modify the KSYS system properties to enable or disable the network mapping policy for all virtual machines across the sites.
- When you map VLAN IDs at host-level, the VLAN IDs are applied to all the virtual machines of that host.

You can create VLAN ID or virtual switch mapping policies that contains a mapping of VLAN IDs or virtual switches that are assigned to virtual machines when the virtual machines are moved from the active site to the backup site. These policies are useful in the following cases:

- In a disaster situation, when you move source hosts or host groups to the backup site, the hosts must have the same VLAN ID, otherwise the recovery operation fails. If the target site is configured with a different VLAN ID, you must set a VLAN policy for source hosts to acquire the same VLAN ID when virtual machines are restarted in the backup site for a successful recovery operation.

- During the test operation for the disaster recovery, when you move hosts or host groups to the backup site in the test mode, if you do not specify a VLAN ID or virtual switch, the virtual machines are started with the same VLAN ID or virtual switch in the backup site as the existing virtual machine in the source site. If both source and target hosts have same VLAN ID, it can result in an IP conflict.

### System-level network mapping policy

To enable the network mapping policy and to create network mapping policy for all hosts and host groups across the active site and the backup site, enter the following command in the KSYS LPAR:

```
ksysmgr modify system network_mapping=<enable | disable>
network=<vlanmap | vswitchmap> sites=<siteA,siteB>
       siteA=<#,[#,...]> siteB=<#,[#,...]>]
```

For example:

```
ksysmgr modify system  network_mapping=enable
   network=vlanmap sites=siteA,siteB
   siteA= VLAN1,VLAN12,VLAN13,VLAN5
   siteB= VLAN1,VLAN22,VLAN23,VLAN5
```

### Site-level network mapping policy

To enable the network mapping policy and to create network mapping policy for all hosts and host groups in a specific site, enter the following command in the KSYS LPAR:

```
ksysmgr modify site <sitename[,sitename2,...]> | file=<filepath>
[network=<vlanmap | vswitchmap>  backupsite=siteB
  sitename=<#[,#,...] || all> siteB=<#[,#,...] || all> [dr_test=<yes|no>]
```

For example:

```
ksysmgr modify site site1 network=vlanmap backupsite=site2
  site1=1,2,3 site2=4,5,6 dr_test=yes
```

### Host-group level network mapping policy

To create a mapping policy of VLAN ID or virtual switches for all the hosts in a host group across sites, enter the following command in the KSYS LPAR:

```
ksysmgr modify host_group <name> options
      network=<vlanmap | vswitchmap>  sites=<siteA,siteB>
       siteA=<#,[#,...]> siteB=<#,[#,...]>
```

For example:

```
ksysmgr modify host_group HG1 options
     network=vswitchmap sites=site1,site2
     site1=vswitch1,vswitch2
     site2=vswitch2,vswitch1
```

### Host-level network mapping policy

To create a VLAN ID mapping policy for all virtual machines in a host across sites, enter the following command in the KSYS LPAR:

```
ksysmgr modify host <hostname[,hostname2,...]> | file=<filepath>
     network=<vlanmap | vswitchmap>  sites=<siteA,siteB>
       siteA=<#,[#,...]> siteB=<#,[#,...]>
```

For example:

```
ksysmgr modify host host_1_2, host 2_2
    network=vlanmap sites=Site1, Site2
    site1= VLAN1,VLAN12,VLAN13,VLAN5
    site2= VLAN1,VLAN22,VLAN23,VLAN5
```

## Modifying the KSYS configuration

Growing business requirements might need changes in your current configuration, for example, adding a specific virtual machine or an entire host to your environment. After the GDR solution is implemented in your environment, the KSYS subsystem continues to monitor any changes in the KSYS configuration. If you want to modify the current configuration in your environment, you can run the **ksysmgr discover** command to discover and validate the changes in the KSYS configuration immediately.

**Notes:**

- If you want to add a new virtual machine to the KSYS configuration, ensure that the corresponding host and the managing HMC are also added to the KSYS configuration. Also, you must add a corresponding host and managing HMC to the backup site so that you can create a host pair across sites.
- When you add a host to the KSYS configuration, by default all the virtual machines are included in the KSYS configuration.
- If you want to add a disk array to the KSYS configuration, you must add corresponding storage agents to the corresponding site. Also, you must add equal number of storage disks with the same disk size to the backup site so that disk pair can be created by the KSYS subsystem during the discovery operation.

To add a specific VM to the KSYS configuration after the initial KSYS subsystem setup, complete the following steps:

1. Add the managing HMC and the managed host, which contains VM, to the active site by running the following command:

   ```
   ksysmgr add hmc Site1_HMC3 login=hscroot hostname=Site1_HMC3.testlab.ibm.com site=Site1

   ksysmgr add host Site1_host3 site=Site1
   ```

2. Add the managing HMC and the managed host to the backup site by running the following command. These managing HMC and managed hosts can be used to create host pairs.

   ```
   ksysmgr add hmc Site2_HMC3 login=hscroot hostname=Site2_HMC3.testlab.ibm.com site=Site2

   ksysmgr add host Site2_host3 site=Site2
   ```

3. Create a host pair between these added hosts by running the following command:

   ```
   ksysmgr pair host Site1_host3 pair=Site2_host3
   ```

4. If you want to exclude some virtual machines during a recovery operation, run the following command for each VM that you want to exclude:

   ```
   ksysmgr unmanage VM_name
   ```

5. If the virtual machines use storage disks that are not yet added to the KSYS configuration, add the corresponding disks to the existing storage arrays in the corresponding sites.

   **Note:** You must add a corresponding disk in the backup site so that the KSYS can create a disk pair across sites for replication during the discovery operation.

6. If you are adding a storage array to a site, you must create storage agents in the KSYS configuration by running the following command:

   ```
   ksysmgr add storage_agent sa_site1_array3 site=Site1
                   serialnumber=000196xxx storagetype=emc ip=1.2.xx.xx

   ksysmgr add storage_agent sa_site2_array3 site=Site2
                   serialnumber=000197xxx storagetype=emc ip=1.2.xx.xx
   ```

7. Discover and verify these added resources by running the following command:

```
ksysmgr discover site Site1 verify=yes
```

The KSYS starts to monitor the modified configuration in the active site.

## Modifying the attributes of the added resources

In addition to adding new resources, you can also modify the attributes of an existing resource in the KSYS configuration.

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

- To change the site name, use the following syntax:

  ```
  ksysmgr modify site old_site_name
        [name=new_site_name]
  ```

- To update the HMC name, login credentials, or IP address, use the following syntax:

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

- To change the storage agent details, use the following syntax:

  ```
  ksysmgr modify storage_agent old_sa name=new_sa
  ```

  **Note:** You can change only the name of an existing storage agent.

- To update the contact details for the KSYS notification, use the following syntax:

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

  For more information about the KSYS notifications, see "Managing the KSYS notifications" on page 83.

- To update the system tunable attributes, use the following syntax:

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

  For more information about the system tunable attributes, see "Managing the system attributes" on page 85.

- To update the priority of virtual machines, use the following syntax:

  ```
  ksysmgr modify VM name1[,name2,name3,...] | file=filepath
     [uuid=uuid_value]
     [host=hostname]
     [priority=low|medium|high]
  ```

**Note:** You cannot update any attributes of an existing host in the KSYS configuration. Any modifications in the host details are registered in the HMC, and the KSYS subsystem discovers these changes during the discovery operation.

# Removing resources from the KSYS configuration

If a specific virtual machine is not required to be covered when you are using the GDR disaster recovery solution, you can remove the resource and the associated resources from the KSYS configuration.

## Removing a virtual machine

When you remove a host from the KSYS configuration, all the virtual machines are removed. Therefore, if you want to remove a specific virtual machine, instead of removing the host from the KSYS configuration, you must exclude the virtual machine from the KSYS configuration so that the virtual machine is not moved to the backup site during a disaster recovery operation. You can exclude a specific virtual machine by using the `ksysmgr unmanage vm_name` command.

### Removing a host

To remove a host, you must first remove the host from the host group, and then break the associated host pair by using the following command:

```
ksysmgr modify host_group hg_name remove hosts=host1
ksysmgr pair host=host1 pair=none
```

After the host pair is broken, use the following command to remove the hosts in both the active site and the backup site:

```
ksysmgr delete host hostname
```

All the virtual machines that are running in the host are removed from the KSYS configuration. You must also remove the host in the opposite site that was paired to the removed host.

If you add or remove hosts from a host group, you must run a discovery operation to manage or unmanage all the virtual machines from the recovery management. The modified host group displays the correct list of managed virtual machines only after a discovery operation. If we remove all hosts from `Default_HG`, the disk group of `Default_HG` is not removed. The disk groups are retained with the removed hosts.

### Removing a host group

To remove a host group, you must first remove all the hosts from the host group, and then delete the host group by using the following command:

```
ksysmgr delete host_group hg_name
```

### Removing an HMC

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

```
ksysmgr delete hmc hmcname
```

### Removing a storage agent

If you want to remove storage disks from the active site, ensure that the storage disks in the backup site that were paired with the storage disks in the active site are also removed. Otherwise, the disk pair might cause discovery errors. If you remove an entire storage array from a site, you must remove the storage agent that is associated with that storage array by using the following syntax:

```
ksysmgr delete storage_agent storage_agent_name
```

## Managing the KSYS notifications

After you add the contact details to the KSYS configuration for the first time, you can modify it later depending on the changing requirements.

To modify the contact information, use the following commands:

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

For example, to change the user name of John to Dave, and to change the email address, enter the following command:

```
ksysmgr modify notify oldcontact=John newcontact=Dave
ksysmgr modify notify oldcontact=john@gmail.com newcontact=dave@gmail.com
```

To delete all the contact information for a specific user, use the following command:

```
ksysmgt delete notify user=username
```

For example,

```
ksysmgt delete notify user=John
```

To query all the registered contact details, use the following command:

```
ksysmgr query notify contact
```

The output might be similar to the following example:

```
User:           Mark Smith
Contact:       mike@mike.com

User:           joe
Contact:       joe@gmail.com
```

## Running scripts for specific events

You can create scripts for specific events. When an event occurs, you can enable the script to be run as part of the event. By using scripts, you not only get notified for a specific event, but you can also collect details about the event, take corrective actions, and handle the processes after the event completion. For more information about scripts, you can refer to the event script samples that are located in the /opt/IBM/ksys/samples/event_handler/event_script_template directory.

You must specify the full path of the script to add a script for notification configuration. When the event occurs, the KSYS subsystem validates the existence and the authorization of the scripts.

To add a script, use the following command:

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

For example,

```
ksysmgr add notify script=/tmp/script.sh events=HMC_DOWN
```

To modify a script, use the following command:

```
ksysmgr modify notify oldscript=old_script_file_path_name newscript=new_script_file_path_name
```

For example,

```
ksysmgr modify notify oldscript=/tmp/script.sh newscript=/tmp/newscript.sh
```

To remove a script, use the following command:

```
ksysmgr delete notify script=script_file_path_name
```

For example,

```
ksysmgr delete notify script=/tmp/script.sh
```

To query a script, use the following command:

```
ksysmgr query notify script
```

## Notification message

Even if you set the KSYS configuration to not receive any event notifications, the messages are logged in the /var/ksys/events.log notification log file. An example of the notification message follows:

```
HMC_DOWN event has occurred. Details are as follows:
     Event:             HMC_DOWN
     Type:              Critical Error Event
     Time:              Tue Jul 19 00:35:32 CDT 2016
```

```
Entity Affected:    HMC
Resource Affected:  vmhmc1
Description:        0000-132 Error - HMC x.x.x.x is down.
Suggestion:         Please try to restart.
```

## Managing the system attributes

After you synchronize the KSYS cluster by using the **ksysmgr sync ksyscluster** command, the KSYS subsystem sets up the default system-wide persistent attributes. The KSYS subsystem uses these system-wide persistent attributes for activities such as automatic rediscovery of the resources, notification of critical events, removal of duplicate notification.

By default, the KSYS subsystem sets up the following system attributes:

**auto_discovery_time**

Specifies the time interval in hours after which the KSYS subsystem rediscovers the environment automatically for any new or modified resources. By default, the value of this attribute is 24 hours, which means that the KSYS subsystem discovers the resources and updates its database about the virtual machines in a period of 24 hours.

The minimum acceptable period is 24 hours. The HMC and VIOS are also involved in the rediscovery process to update information about the virtual machines. Therefore, if your environment is large (for example, hundreds of logical partitions in the data center), you might want to increase this period to reduce load on the HMC, VIOS, and the underlying I/O subsystems. Any configuration changes to the hosts, virtual machines, disks, and any other entities (for example, addition of new disks to a virtual machine) are captured when the rediscovery occurs after the specified period. This attribute also specifies the span of time during which any changes in the configuration setting can be lost if a disaster occurs before the rediscovery.

**lose_vios_redundancy**

Specifies to start the virtual machines in another site without the dual VIOS setup in the target host. By default, this attribute is set to `no`, which means that the dual VIOS setup is maintained during the disaster recovery operation.

If your currently active site, which has dual VIOS configuration, fails, and one of the VIOS in the target host of the backup site is not functioning, you might want to recover the virtual machines with only one VIOS on the backup site and continue the workloads that are running in the active site. In this case, you can set this attribute to `yes`. However, if the virtual machines are started with single VIOS configuration on the backup site, and later if you want to move the virtual machines back to the previous site that has a dual VIOS configuration, you must manually add the second VIOS into the configuration. For more information, see Configuring VIOS partitions for a dual setup.

**auto_reverse_mirror**

Resets the storage data mirroring after a site recovery operation.

Currently, this attribute is always set to **no**. You must resynchronize the storage data mirroring manually by using the **ksysmgr resync** command.

**notification_level**

Enables or disables the notification for different types of events. This parameter supports the following values:

**low (default)**

Only critical error events are notified.

**medium**

Critical and warning error events are notified.

**high** All events, which include informational events, are notified.

**disable**
None of the events are notified.

**dup_event_processing**
Reduces duplicate event notifications. The email and script notifications that are related to the duplicate events are also disabled. This parameter can have the following values:

**yes (default)**
Notifies about only those events that are not repeated in the last 24 hours.

**no** Notifies all the messages.

**replication_type**
Specifies the storage replication mode across the sites. This parameter can have the following values:

**async** Specifies the asynchronous replication mode in which data is transferred across sites in predefined timed cycles or in delta sets.

**sync** Specifies the synchronous replication mode in which the storage subsystem acknowledges to the host that it has received and checked the data.

You must also specify the source and the target sites that are associated with the storage replication operation.

**network_mapping**
Enables or disables the network mapping policies for all hosts and host groups across the active site and the backup site.

**network**
Creates network mapping policy for all hosts and host groups. This attribute can be specified only when the **`network_mapping`** attribute is set to `enable`.

You can create either a VLAN mapping policy or a virtual switch mapping policy. This attribute must be specified as follows:

```
[network=<vlanmap|vswitchmap> sites=<siteA,siteB>
      siteA=<#,[#,...]> siteB=<#,[#,...]>
```

To query the information or status related to system-wide attributes, use the following command:

```
ksysmgr query system [properties | status [monitor=<no|yes>]]
```

### Examples for querying system properties and status

- To query the existing values of the system attributes, run the following command:

  ```
  ksysmgr query system properties
  ```

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

  ```
  System-Wide Persistent Attributes
  auto_discovery_time  ="00:00" hours
  lose_vios_redundancy ="no"
  auto_reverse_mirror  ="yes"
  notification_level   ="low"
  dup_event_processing ="yes"
  replication_type     ="Asynchronous"
  network_mapping      ="enable"
  network              =vlanmap sites=siteA,siteB
        siteA= VLAN1,VLAN12,VLAN13,VLAN5
        siteB= VLAN1,VLAN22,VLAN23,VLAN5
  ```

- To query the system status, run the following command:

  ```
  ksysmgr query system status
  ```

  The output might look similar to the following sample:

```
Discovery is in progress for Site Site1.
 Discovery is in progress for Site Site2.
 Please use "ksysmgr query system status monitor=yes " to track the progress of the operation.
```

If you specify the **monitor=yes** parameter, the KSYS monitors and displays all the current actions on the KSYS node. The actions include discovery, verification, movement, and recovery operations, which continue to progress even when the command is killed or when you close the command window. The **monitor** parameter also displays the encountered errors that are logged by the virtual machines, hosts, disk pairs, and disk groups. For example:

```
ksysmgr query system status monitor=yes
        Discovery is in progress for Site Site1.
        Discovery is in progress for Site Site2.
        Monitoring status...
Running discovery on entire site, this may take few minutes...
...
```

To modify the system-wide attributes, use the following command:

```
ksysmgr modify system <attribute>=<new_value>
```

**Examples for modifying the system-wide attributes**

- To enable the KSYS subsystem to automatically rediscover the resources once in two days, run the following command:
  ```
  ksysmgr modify system auto_discover_time=48
  ```
- To change the notification level of your system to receive notification for all critical errors and warnings for all events, run the following command:
  ```
  ksysmgr modify system notification_level=medium
  ```
- To change the duplicate event processing option to receive notification for all events, even if the events are duplicated, run the following command:
  ```
  ksysmgr modify system dup_event_processing=yes
  ```
- To change the storage replication mode between two sites from asynchronous mode to synchronous mode, run the following command:
  ```
  ksysmgr modify system replication_type=sync sites=SiteA,SiteB
  ```

## Managing the shared storage configuration

The GDR solution manages disaster recovery across two sites based on storage replication across the sites. However, the GDR solution also supports a mode of deployment in which disks are shared across sites. In this case, the KSYS subsystem does not manage any storage subsystems. The disks can be shared across sites that are separated by short distances (0 - 100 km). The storage technologies (for example, Storwize HyperSwap®) perform synchronous mirroring across sites and abstract the mirroring from the hosts. These storage technologies provide shared disk type of deployment for hosts across sites.

**Restriction:**

- Because the storage is shared, the N_Port ID virtualization (NPIV) and other similar ports are visible to VIOS on both sites. It might cause problems that are related to SAN login and disk validations. Hence, HMC and VIOS-related checks are not performed in the shared deployment mode. Therefore, the administrator must set up the sites appropriately, considering the storage, network, and so on, and must maintain the configuration settings. Any misconfiguration might result in errors during a disaster recovery.
- The DR failover rehearsal operation is not supported for shared storage deployments.

### Shared storage without replication management

When the storage device in your environment is a single storage system that is split or separated by distance in two different sites as stretched systems, the storage replication management is hidden from the hosts and VIOS partitions. The storage subsystem is displayed as a single shared storage across the two sites. In this case, the storage recovery and replication are performed entirely by the storage platform. You do not need storage agents in the KSYS subsystem to interact with the storage devices. Therefore, you do not need to consider the disk pair and disk group mappings for the shared storage configuration.

When you move the virtual machines from the active site to the backup site, the KSYS subsystem considers the storage subsystem as unmirrored shared storage and starts the virtual machines on the backup site. If the disaster recovery operation is unplanned, the storage subsystem performs the entire storage recovery.

**Notes:**

- The GDR solution supports this type of shared storage only for sites that are spanned across short distances. Also, the storage subsystem must provide shared storage characteristics.
- The GDR solution does not support heterogeneous storage systems for this type of shared mode configuration. You must deploy same type of storage systems across the sites in your environment.

The following figure shows an example of shared storage configuration with Storwize HyperSwap based technology:

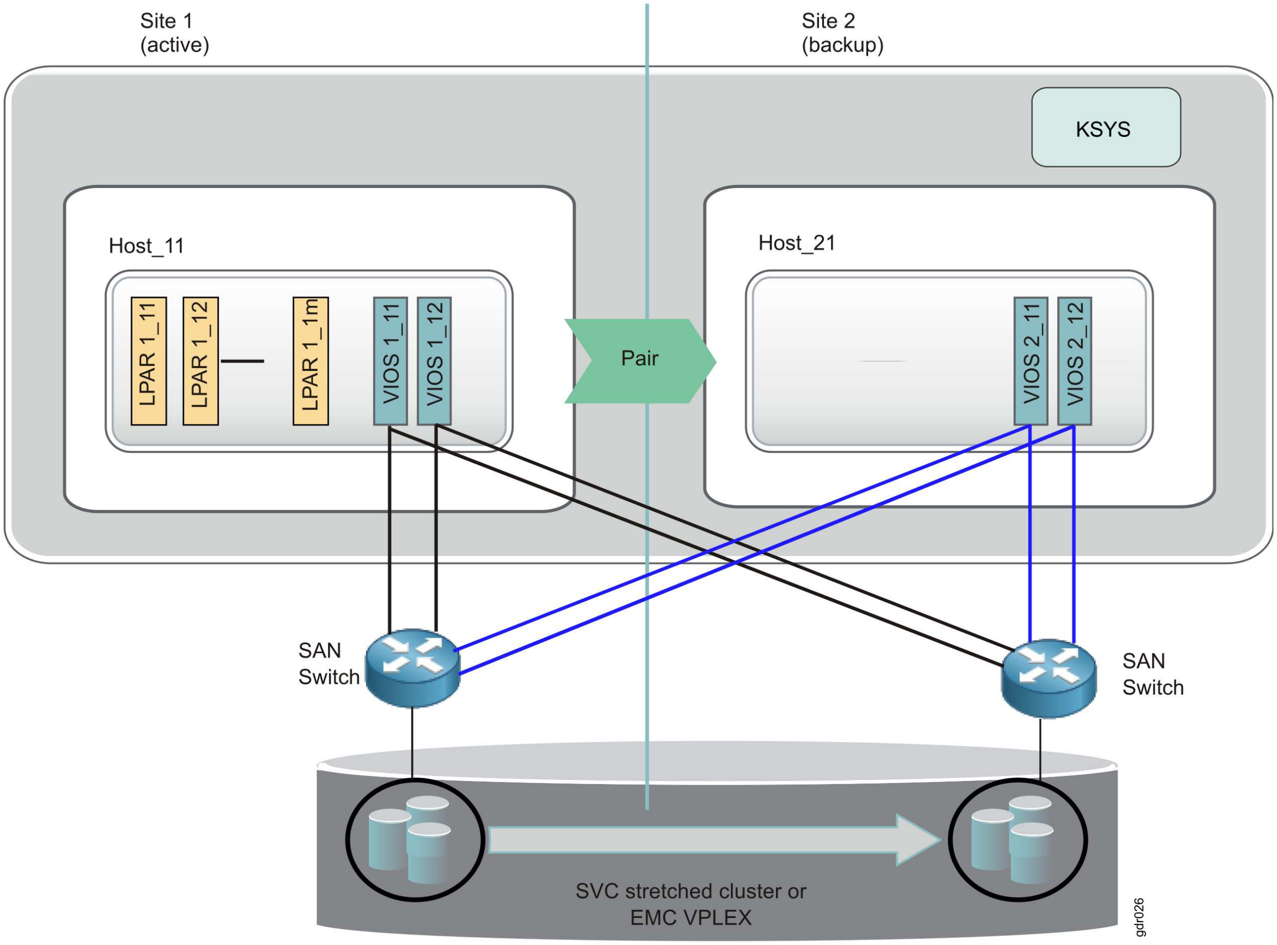

*Figure 25. Shared storage configuration without KSYS-based replication management*

### Shared storage with SSP mirror management

VIOS can manage storage efficiently across multiple hosts in an environment by using the Shared Storage Pool (SSP) technology. Shared storage provides various features that support mirroring. Mirroring enables continuous storage availability across sites against storage failure in the site. If your environment is managed entirely by SSP for both system (for example, rootvg) and data disks, then you can spread the SSP cluster across two buildings and achieve protection against storage failures. This type of deployment is suited for short distances (less than 1 km).

SSP-based deployments belong to the shared storage category. That is, the KSYS subsystem does not manage the storage replication. Because the SSP-based storage exhibits the shared storage model, you do not need to consider the disk pair and disk group mappings for storage configuration.

When you move the virtual machines from the active site to the backup site, the KSYS subsystem does not perform any storage-specific preparation (unlike the conventional replicated storage systems) and starts the virtual machines on the backup site.

**Note:** The GDR solution supports shared storage with SSP mirror management only for sites that are spanned across short distances.

The following figure shows an example of shared storage configuration with mirror management by SSP.

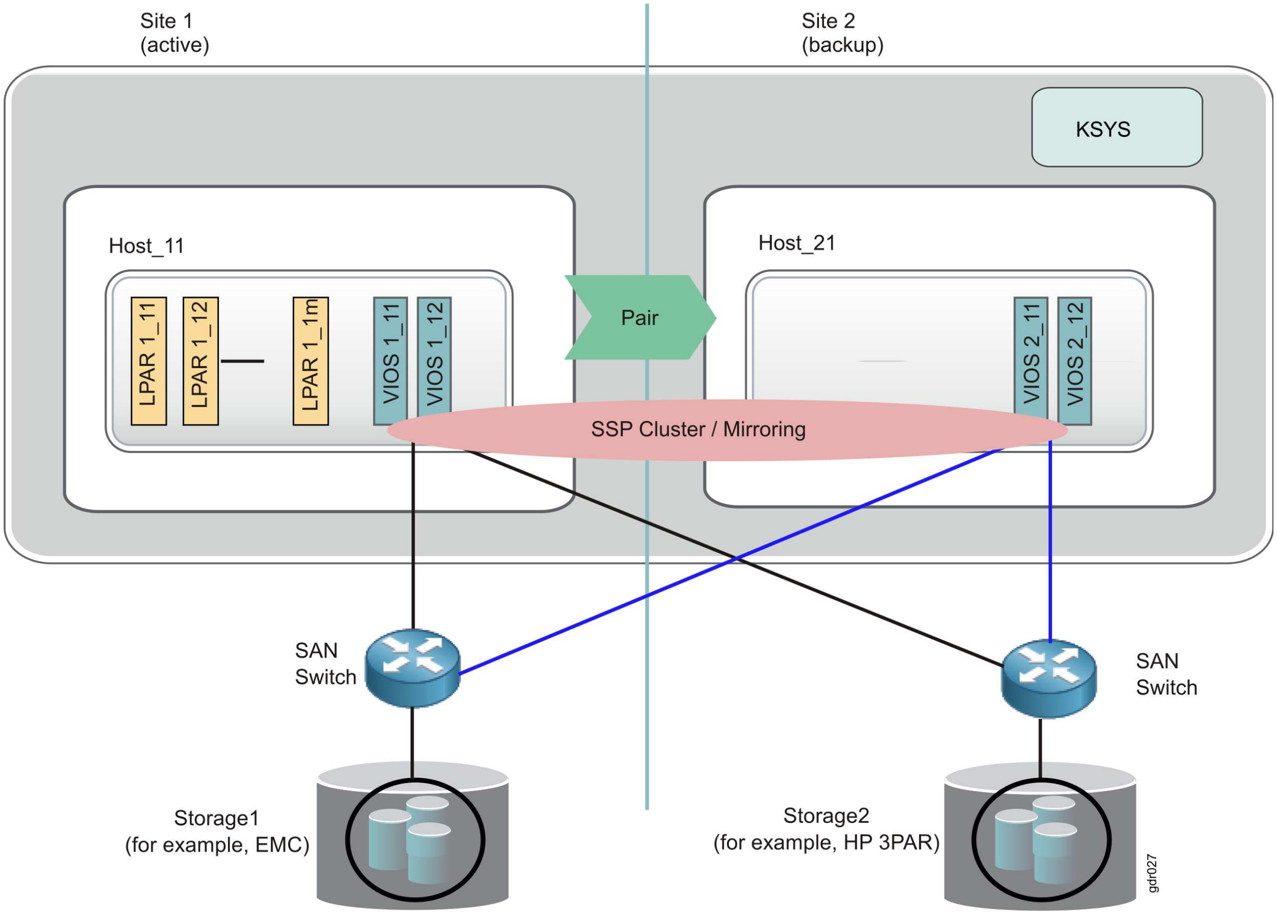

*Figure 26. Shared storage configuration with SSP mirror management*

## Backing up and restoring the configuration data

You can back up all the current configuration settings of your KSYS environment as a snapshot. A *snapshot* preserves the configuration details of the KSYS environment at a specific point in time. For example, a snapshot file contains the information about the existing sites, details about the managing HMCs and the managed hosts in a specific site, and the storage device details in the site. You should back up your current configuration settings after you configure the sites, hosts, HMCs, and storage devices initially.

If you save a snapshot of the current KSYS configuration settings, you can restore the configuration settings later by applying the snapshot on the KSYS configuration. The snapshots are useful during node upgrades or environment malfunctions because snapshots eliminate the need to reconfigure the sites, hosts, HMCs, and storage devices. For example, if the KSYS node must be reinstalled, you can use a snapshot and do not have to re-create sites, hosts, and other resources.

You can save the following types of snapshots:

**Cluster configuration data**

Backs up the core KSYS deployment data that is used to create the KSYS cluster. The cluster configuration snapshot contains the following information:

- Cluster name
- KSYS node name

**Basic configuration data**

Backs up the cluster configuration data along with the data to configure sites, hosts, HMCs, and storage agents. It also includes information about the LPARs in a host that are excluded from the configuration. However, it does not include detailed information about disk pairs and disk groups for each LPAR.

The basic configuration snapshot contains the following information:

- Sites
  - Site name
  - Site type
- Storage agent
  - Name
  - IP address
  - Storage host name
  - User name
  - Password
  - Site name
  - Site ID
  - Storage type
  - Serial number
- HMCs
  - User name
  - Password
  - IP address
  - Name (logical name)
  - Site ID
  - Site name
- Hosts

  - Name
  - UUID
  - FSP ID
  - FSP host name
  - Host's partner's UUID

**Detailed configuration data**
Backs up all the basic configuration data along with the detailed LPAR data that is determined by discovering resources in a site such as disk pairs and disk groups. The detailed configuration snapshot contains information about sites, storage agents, HMCs, hosts, disk pairs, and disk groups. If you had configured tertiary disk in the backup site for failover rehearsal of disaster recovery operation, the detailed configuration snapshot contains the tertiary disk values.

If you capture a snapshot on GDR Version 1.1.0.1, and you want to restore the configuration settings on GDR version 1.2.0, you can restore only **cluster** and **basic** type of snapshots. To restore **detailed** type of snapshot, the GDR version must be the same as the version it was captured.

## Saving and restoring snapshots

Use the **ksysmgr** command to save the configuration snapshots. The snapshots are saved in an XML format. When you create a snapshot, the **ksysmgr** command appends the date and time to the specified file name to follow the *filename.DateTime* name convention. By default, the snapshot files are saved in the /var/ksys/snapshots directory. However, you can specify the path where the snapshot files must be saved.

**Notes:**

- You must ensure that the /var file system has enough space for the snapshot files before you back up the configuration data.
- If the KSYS node must be reinstalled, you must save the snapshot files in a different location so that the snapshot files can be used later for restoring the configuration settings by using the **ksysmgr** command. You must use the **ftp** command to copy the snapshot files to another system and to copy the snapshot files from another system to the KSYS node after the installation is complete.
- You cannot restore a snapshot file that is captured in a different site. To restore a snapshot file, the snapshot file must be captured on the same site.
- If you want to get the tertiary disk information after restoring a basic snapshot, you must run the discovery operation with the **dr_test** flag.
- If you connect the source hosts and target hosts to the same HMC, it leads to an invalid configuration in the KSYS subsystem. If you had saved a snapshot of such an invalid configuration, the restore operation fails.

Use the following steps to manage the snapshots:

- To save a snapshot, use the following command syntax:

  ```
  ksysmgr add snapshot filepath=full_file_prefix_path|file_prefix type=cluster|basic|detailed ]
  ```

  For example, to back up basic configuration data of your KSYS environment once in a week such that no more than 5 backup files exist, run the following command:

  ```
  ksysmgr add snapshot filepath=/home/ksysdir/myksysbackup type=basic
  ```

  The **ksysmgr** command saves the snapshot after archiving and compressing the file.
- To view an existing snapshot, use the following command syntax:

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

  For example:

  ```
  ksysmgr query snapshot filepath=/home/ksysdir/myksysbackup_2016_06_23_04:54:30.xml.tar.gz
  ```

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

```
# ksysmgr query snapshot filepath=myksysbackup_2016_06_23_04:54:30.xml.tar.gz
---- Snapshot Contents ----
File:       /home/ksysdir/myksysbackup_2016_06_23_04:54:30.xml
Type:       BASIC
Version:    1.00
Date:       2016-06-23
Time:       04:54:30
---------------------------

Cluster:
--------
Name:      test_cluster
Node:       test_KSYS_node
---------------------------
```

Before you restore a saved snapshot on a KSYS node, in which the operating system is reinstalled, or on another logical partition that must be used as KSYS LPAR, you must ensure that the *HOST* variable is set. You can set the *HOST* variable as shown in the following example:

```
#  export HOST=host_name
```

- To restore the configuration data on a KSYS node, use the following command syntax:

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

  For example:

  ```
  ksysmgr restore snapshot filepath=/home/ksysdir/myksysbackup_2016_06_23_04:54:30.xml.tar.gz
  ```

  This command decompresses and unarchives the snapshot file, and then applies the configuration settings to the KSYS node.

**Related reference**:
"ksysmgr command" on page 113

## Managing the CoD resources

If you use capacity management solutions in your environment to reduce the capacity on the backup hosts during non-disaster times, you can use the GDR solution to manage the resource allocations during a disaster recovery failover operation.

To manage the resource allocations in a disaster recovery environment, the GDR solution provides a resource pool provisioning (RPP) command that is called the **ksysrppmgr** command. The **ksysrppmgr** command optimizes available resources on the managed hosts. This command also minimizes your resource costs by optimizing the local consumption of pool resources.

Before you run the **ksysrppmgr** command to allocate resources to a managed host, you can run the **ksysrppmgr** command in check mode to simulate the execution and analyze the results.

The GDR solution can manage resource allocation for the following capacity management solutions:

- Power Enterprise Pool
- Elastic (On/Off) Capacity-on-Demand

**Related reference**:
"ksysrppmgr command" on page 126

### GDR and Power Enterprise Pool

Power Enterprise Pool provides flexibility to hosts that operate together as a pool of resources. Mobile activations can be assigned to any host in a predefined pool and the resources can be reassigned within a pool.

If you use Power Enterprise Pool in your environment for capacity management, review the following scenarios to determine how you can manage the resource allocations by using the GDR solution:

## Scenario 1: Using Enterprise Pools across sites

In this scenario, the Enterprise Pool is shared across sites as shown in the following figure:

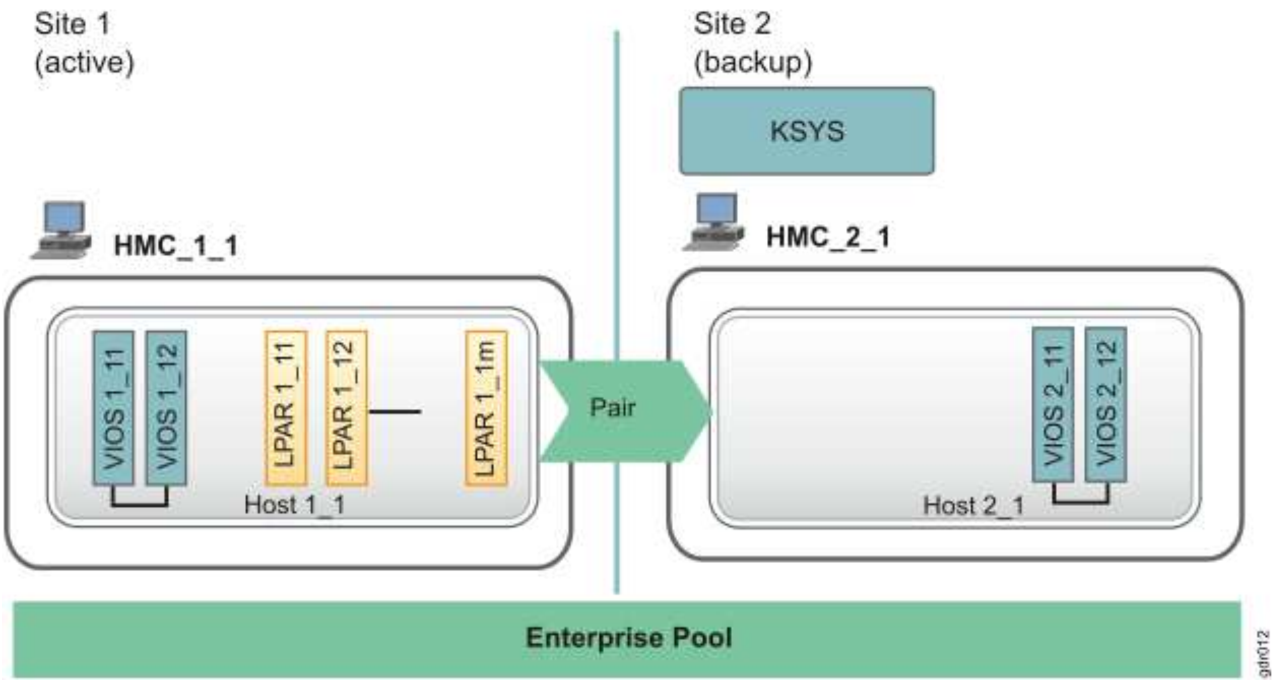

*Figure 27. Power Enterprise Pool usage across the sites*

When the active site fails, complete the following steps before you initiate the site-switch operation:

1. In the KSYS node, authenticate all the HMCs by running the **hmcauth** command:

   ```
   hmcauth -u hmcuser -p password -a HMC_1_1
   hmcauth -u hmcuser -p password -a HMC_2_1
   ```

2. Check whether the required number of processors and memory can be available in the backup host that does not use Elastic (On/Off) CoD by running the following command:

   ```
   ksysrppmgr -o check -h :HMC_1_1:hmcuser -h :HMC_2_1:hmcuser \
   -m Host_2_1:set:n:<memory_amount>:<no_of_processors> -r
   ```

   If the return code of the command is 0, all the requests can be fulfilled. If the return code is 1, at least one request has failed.

3. If the resource requests are not fulfilled, release the resources that are used by the virtual machines of `Site_1`, and return the resources to the Enterprise Pool either by using the HMC or by running the following command in the KSYS node:

   ```
   ksysrppmgr -o execute -h :hmc1_1:hmcuser -m host_1_1:set:n:0:0
   ```

   For more information about the steps to release or allocate resources by using HMC, see Using Power Enterprise Pools.

4. Allocate the required amount of resources to the hosts in the `Site_2` site by using the HMC or by running the following command in the KSYS node:

   ```
   ksysrppmgr -o execute -h :HMC_2_1:hmcuser \
   -m Host_2_1:set:n:<memory_amount>:<no_of_processors> -r -v
   ```

The target host on the backup site now contains all the required resources to host the recovered virtual machines.

5. Verify whether the virtual machines can be moved to the backup site by running the following command:

```
ksysmgr verify site site_2
```

**Note:** If the backup site does not have sufficient processors and memory, the verify operation fails with a warning message. You can use the force (**-f**) option to move the virtual machines to the backup site with the existing configuration.

6. Initiate the disaster recovery by switching the site from Site_1 to Site_2 by running the following command in the KSYS node:

```
ksysmgr move from=site_1 to=site_2 dr_type=planned
```

## Scenario 2: Using Enterprise Pool within the backup site

In this scenario, the Enterprise Pools are shared across hosts within the backup site. In the example, as shown in the figure, Host_1_1 in the active site is paired to Host_2_1 in the backup site. Host_2_2 is another host in the backup site that is running low priority virtual machines. When the active site fails, you can allocate some resources from the Host_2_2 to the Host_2_1 to run the recovered virtual machines.

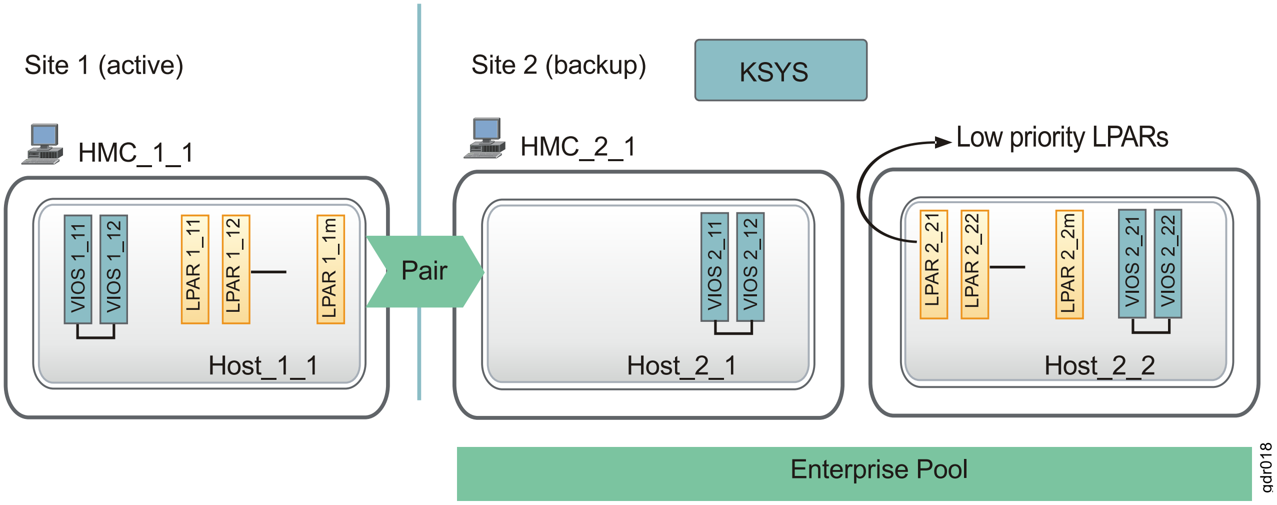

*Figure 28. POWER Enterprise Pool usage within the backup site*

Before you initiate the site-switch operation, complete the following steps:

1. In the KSYS node, authenticate all the HMCs by running the **hmcauth** command:

```
hmcauth -u hmcuser -p password -a hmc_2_1
```

2. Check whether the required number of processors and memory can be available in the backup host that does not use Elastic (On/Off) CoD by using the HMC or by running the following command in the KSYS node:

```
ksysrppmgr -o check -h :HMC_1_1:hmcuser -h :HMC_2_1:hmcuser \
-m Host_2_1:set:n:<memory_amount>:<no_of_processors> -r -v
```

If the return code of the command is 0, all the requests can be fulfilled. If the return code is 1, at least one request has failed. For more information about the steps to release or allocate resources by using HMC, see Using Power Enterprise Pools.

3. If the output indicates that the request cannot be fulfilled, reduce the resources that are allocated to Host_2_2 that runs low priority virtual machines and return the resources to the Enterprise Pool either by using HMC or by running the following command in the KSYS node:

```
ksysrppmgr -o execute -h :hmc_2_1:hmcuser \
-m host_2_2:set:n:<memory_amount>:<no_of_processors>
```

4. Allocate the resources to the Host_2_1 host by running the following command:

   ```
   ksysrppmgr -o execute -h :hmc_2_1:hmcuser \
   -m host_2_1:set:n:<memory_amount>:<no_of_processors>
   ```

5. Verify whether the virtual machines can be moved to the backup site by running the following command:

   ```
   ksysmgr verify site site_2
   ```

   **Note:** If the backup site does not have sufficient processors and memory, the verify operation fails with a warning message. You can use the force (**-f**) option to move the virtual machines to the backup site with the existing configuration.

6. Initiate the disaster recovery by running the following command:

   ```
   ksysmgr move from=site_1 to=site_2 dr_type=planned
   ```

## GDR and Elastic (On/Off) CoD

Elastic Capacity-on-Demand (formally known as On/Off CoD) provides short-term CPU and memory activation capability for fluctuating peak processing requirements.

If the resource requirements are not met even after you allocate maximum number of resources from the Enterprise Pool, you can allow the Elastic (On/Off) CoD usage that activates temporary resources for a specified number of days. The following figure shows an example of Elastic (On/Off) CoD usage within a site:

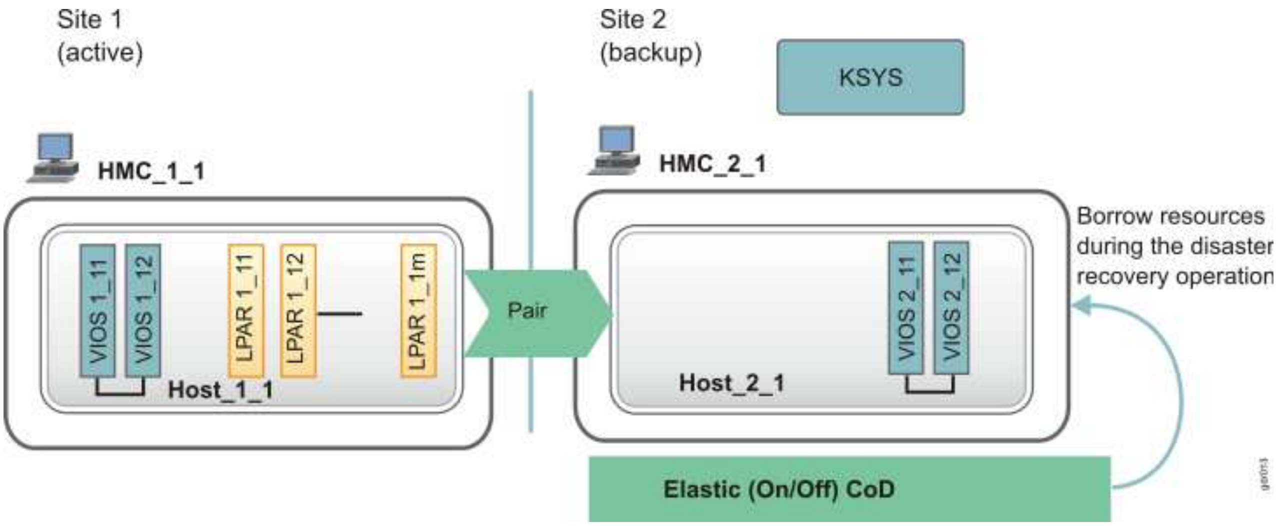

*Figure 29. Elastic (On/Off) CoD usage for a host*

When the active site fails, complete the following steps to allow Elastic (On/Off) CoD to manage the resources before you initiate a site-switch operation:

1. Authenticate all the HMCs by running the **hmcauth** command:

   ```
   hmcauth -u hmcuser -p password -a hmc_2_1
   ```

2. Identify the amount of available processor and memory units that are required by the backup hosts and determine whether you want to authorize the use of Elastic (On/Off) CoD to achieve the request.
3. If the resource requests are not met by using Enterprise Pool, allow the use of Elastic (On/Off) CoD. For example, to request 2.5 CPUs and 10500 MB memory for the host_2_1 host and to use Elastic (On/Off) CoD for 5 days, enter the following command:

   ```
   ksysrppmgr -o execute -h :hmc_1_1:hmcuser -h :hmc_2_1:hmcuser -m host_2_1:set:y5:10500:2.5 -r
   ```

4. Verify whether the virtual machines can be moved to the backup site by running the following command:

   ```
   ksysmgr verify site site_2
   ```

   **Note:** If the backup site does not have sufficient processors and memory, the verify operation fails with a warning message. You can use the force (**-f**) option to move the virtual machines to the backup site with the existing configuration.
5. Initiate the disaster recovery operation by running the following command in the KSYS node:

   ```
   ksysmgr move from=site_1 to=site_2 dr_type=planned
   ```

# GDR coexistence with other products and offerings

Review the following scenarios that explain how the GDR solution coexists with other solutions that are deployed in the same environment.

## Live Partition Mobility and GDR coexistence

The Live Partition Mobility (LPM) allows a running LPAR to be relocated from one host to another host within a site. If the GDR solution is implemented in your environment, when you relocate a virtual machine from a host to another host within a site by using the LPM feature, the KSYS subsystem detects this change as part of its daily check. A discovery operation is required for the modified virtual machine to ensure its successful recovery during a disaster situation.

**Note:** The KSYS can identify this movement as an LPM operation only when the source host and the target host for the LPM operation are already added to the KSYS configuration. If one of the source or target hosts are not part of the KSYS configuration, the KSYS detects this movement as an addition or deletion of a virtual machine.

An LPM operation within a single host group does not result in any storage-level changes. An LPM operation across multiple host groups removes the LPAR disks from the source storage group and adds the disks into the target storage group.

By default, any changes in the KSYS configuration are discovered in the span of 24 hours. However, if a disaster recovery operation is required before the discovery operation runs, the movement of this newly-moved virtual machine from the active site to the backup site might fail. Therefore, you must run a discovery and verification operation in the site after the LPM operation is completed.

**Related concepts**:

"Discovering resources" on page 56
After you add the various HMCs, hosts, host groups, and storage subsystems to the KSYS subsystem for disaster recovery management, you must run the **ksysmgr** command to discover all the hosts that are managed by the HMCs in both the active and the backup sites. During the discovery process, the KSYS subsystem captures the configuration information of the active site and its relationship with the backup site and prepares the backup site to perform disaster recovery operations later.

## IBM PowerVC and GDR coexistence

IBM Power® Virtualization Center (PowerVC) is a virtualization management offering that is used to create and manage virtual machines on IBM Power Systems servers by using PowerVM® or PowerKVM hypervisor. The PowerVC offering can coexist with the GDR solution. This information lists the areas where resource conflicts might occur and how the GDR solution handles the conflicts.

The PowerVC solution can trigger the following features as part of its regular or disaster recovery operations:

**Live Partition Mobility (LPM)**
: A running virtual machine is moved from a source host to a target host.

**Remote restart**
A running virtual machine is moved from a faulty source host and restarted in a target host.

**Maintenance mode switchover**
All virtual machines are moved from a host to another host for maintenance purpose on the source host.

**Capacity management**
Memory or CPU core resources can be reassigned to a host within a pool.

**Note:** All these resource changes do not cause a conflict for the GDR solution because any such movements are discovered by the KSYS in the next scheduled discovery. You can also manually run a discovery operation after any such modification is complete.

**Storage area network (SAN) zoning**
You can use the PowerVC solution to perform the SAN zoning for your environment. However, you must ensure that the SAN is connected according to the GDR requirement for the configuration of virtual machines and VIOS partitions.

**AIX Live Update**
The AIX Live Update function eliminates downtime that is associated with installing the interim fixes for the AIX operating system. This function allows workloads to remain active during a Live Update operation and the operating system can use the interim fix immediately without restarting the entire system.

The Live Update function changes the MAC address for a virtual machine after the Live Update operation is complete. The Live Update function results in a surrogate LPAR to exist in the environment for a few minutes. The Live Update function renames the original LPAR with a naming convention, `originalLparName_lku0`.

**Note:** You must not perform a Live Update operation when you are running a disaster recovery operation by using the GDR solution.

For more information about the Live Update operation, see the Live Update topic.

## PowerHA SystemMirror and GDR coexistence

The IBM PowerHA SystemMirror software provides cluster-based high availability (Standard Edition) and disaster recovery (Enterprise Edition) solutions. The GDR solution can operate with PowerHA SystemMirror Version 7.1.0 if you follow the guidelines that are required to deploy both the solutions together.

### Disaster recovery by using PowerHA SystemMirror Enterprise Edition

If you are using the PowerHA SystemMirror Enterprise Edition to perform disaster recovery for some of the virtual machines in your environment, you do not need to deploy the GDR solution for those virtual machines. In this case, you must exclude those virtual machines from the GDR disaster recovery management. Use the **ksysmgr unamange** command to exclude the PowerHA virtual machines from the GDR configuration before you discover resources in the GDR solution.

### High availability by using PowerHA SystemMirror Standard Edition

PowerHA SystemMirror Standard Edition is deployed within a site. PowerHA creates a cluster of a set of virtual machines within the active site for high availability management. If you are configuring such a cluster within the active site of the GDR environment, consider the following guidelines:

- Include all the virtual machines in the cluster to the GDR disaster recovery management.
- Perform a test failover of the cluster to the backup site to validate whether that cluster starts correctly on the backup site.

If you deployed PowerHA SystemMirror Version 7.1.0 (or later) clusters in the environment, some additional steps are required to start the cluster on the backup site by using the GDR solution after the operating system in the virtual machines are started during the recovery operation. The following disk details are required for the backup disks on each site:

- Name
- Physical volume identifier (PVID)
- Universal Unique ID (UUID)

To obtain the disk information of the cluster nodes, complete the following steps on any one of the cluster nodes in the source site immediately after the GDR solution is implemented in your environment:

1. Run the following command to obtain the name and UUID of each disk in the active site cluster:
   ```
   /usr/lib/cluster/clras dumprepos
   ```
   The same information is available on each node of the cluster.
2. Run the following command to obtain the PVIDs:
   ```
   lspv -u | egrep -w "<disk1>|<disk2>|..."
   ```
   where *disk1* and *disk2* are the disk names as displayed in the previous output.
3. Identify and differentiate the disks in the primary repository from the disks in the backup repository.

To restart the PowerHA SystemMirror Version 7.1.0 or Cluster Aware AIX (CAA) clusters on the backup site after the disaster recovery operation is complete, complete the following steps on any one of recovered virtual machine that is used to start the cluster. You can use a single node to perform all the following steps:

1. If the recovered virtual machines are not active, run the following command with the identified PVIDs to get the corresponding disks:
   ```
   lspv -u | egrep -w "<pvid1>|<pvid2>|..."
   ```
   where *pvid1* and *pvid2* are the PVIDs that are listed in the active site.

   **Note:** The disks in the active site and the backup site share common PVIDs.
2. Save the type (primary or backup), name, PVID, and UUID for each disk in the active and backup sites and identify the corresponding disks between sites. For example, `hdiskA` with PVID `B` and UUID `C` mirrors to `hdiskX` with PVID `B` and UUID `Y`.
3. Identify the disks that must be used as the primary repository and the backup repository in the backup site.
4. Remove any CAA information from the backup site repository by using the following command:
   ```
   CAA_FORCE_ENABLED=true rmcluster -fr <backup_repos_name>
   ```
   This command removes only the CAA data.
5. Run the following command to write the repository disk information by using the information from the CAA repository backup file:
   ```
   chrepos -c <backup_repos_name>
   ```
6. If you were using non-native AIX Multipath I/O (MPIO) disks (for example, EMC PowerPath), run the following command to register the disks with the CAA and AIX disk drivers:
   ```
   clusterconf -d
   ```
7. Run the following command for each backup repository disk in the previous site:
   ```
   chrepos -x <old_backup_UUID>,<new_backup_name | new_backup_UUID>
   ```
   For example, if the backup repository disk `hdisk1` with PVID `X` and UUID `Y` mirrors to the backup repository disk `hdisk5` with PVID `X` and UUID `Z`, run the following command:

```
chrepos -x Y,hdisk5
```

or,

```
 chrepos -x Y,Z
```

8. Run the following command to start the cluster in the backup site:

   ```
   clusterconf
   ```

**Related information**:

IBM PowerHA SystemMirror V7.1 for AIX documentation

## Running scripts for additional checks

If you want KSYS to perform some additional checks that are specific to your environment, you can add those checks as scripts. You can specify whether those scripts must be run before or after the discovery, verification, or disaster recovery operations.

For example, if you use the Enterprise Pool resource management, you can add a customized script to update the backup host capacity and revert to older capacity values after the verification is complete. To monitor the workload that is running on the virtual machines on a specific criteria, you can add script to check the workload before and after the verification.

You can run the scripts at site and host group levels. The sample scripts are available in the `/opt/IBM/ksys/samples/custom_validation/` directory.

### Running scripts before or after discovery and verification operations

You can use the following attributes to add a script for additional checks during verification operations:

**pre_verify**

When you specify a script with this attribute, the script is run before any discovery and verification operations. You can add a `pre_verify` script by entering the following command syntax:

```
ksysmgr add script entity=site|host_group pre_verify=script_path
```

**post_verify**

When you specify a script with this attribute, the script is run after any discovery and verification operations. You can add a `post_verify` script by entering the following command syntax:

```
ksysmgr add script entity=site|host_group post_verify=script_path
```

The following environment variables are set when you run the `pre_verify` and `post_verify` scripts:

**GDR_OPERATION_TYPE**

Specifies the type of the operation that is run. This environment variable can have the following values:

**GDR_DISCOVER**

This value is set if the discovery operation is run. For example, when you run the following command:

```
ksysmgr discover site=Site1
```

**GDR_DISCOVER_AND_VERIFY**

This value is set if the discovery operation is run along with the verification operation. For example, when you run the following command:

```
ksysmgr discover site=Site1 verify=true
```

**GDR_VERIFY**

This value is set if the verification operation is run. For example, when you run the following command:

```
ksysmgr verify site=Site1
```

**GDR_OPERATION_MODE**

Specifies the mode of discovery operation. This environment variable is not valid for verification operation. This environment variable can have the following values:

**GDR_QUICK_DISCOVERY**

This value is set for the automatic hourly discovery operation that is run by the KSYS subsystem.

**GDR_DEEP_DISCOVERY**

This value is set for the automatic daily discovery operation that is run by the KSYS subsystem or a manual discovery operation.

### Running scripts before or after the disaster recovery operations

You can use the following attributes to add a script for additional checks during disaster recovery operations:

**pre_offline**

When you specify a script with this attribute, the script is run before the virtual machines are shut down in the primary site. You can add a `pre_offline` script by entering the following command syntax:

```
ksysmgr add script entity=site|host_group pre_offline=script_path
```

**post_offline**

When you specify a script with this attribute, the script is run after all the virtual machines are shut down in the primary site. You can add a `post_offline` script by entering the following command syntax:

```
ksysmgr add script entity=host_group post_offline=script_path
```

**pre_online**

When you specify a script with this attribute, the script is run before the storage replication direction is reversed and before the virtual machines are restarted on the target site. You can add a `pre_online` script by entering the following command syntax:

```
ksysmgr add script entity=host_group pre_online=script_path
```

**post_online**

When you specify a script with this attribute, the script is run after the virtual machines are restarted on the target site. You can add a `post_online` script by entering the following command syntax:

```
ksysmgr add script entity=site|host_group post_online=script_path
```

For the scripts that are run at host group level, the name of the host group is also passed to the scripts by the KSYS subsystem. You can look at the sample scripts to understand how to execute operations for different host groups from the same script.

## Running the disaster recovery scripts

During a disaster recovery event, when the logical partitions (LPARs) move from the active site to the backup site, the IP addresses, subnet, and other network-related attributes change. If you want the backup environment to be the same as the source environment for the LPARs, you can use the disaster recovery scripts that collect the information from the source LPARs and reconfigures the backup LPARs to match the system name, adapter configuration, network parameters, volume group information, and clustering configuration.

The disaster recovery scripts are custom scripts that are available in the KSYS package. Run these scripts in the virtual machines to collect required information about the source LPAR and to use the collected information to re-create or import the environment in the recovered LPAR.

Consider the following prerequisites before you use the disaster recovery scripts:

- The **data_collection.ksh** script must be run successfully on every LPAR in the source site.
- The **setup_dr.ksh** script must be run successfully on every LPAR in the backup site.
- All of the cluster LPARs must be running and must have network capability.

The disaster recovery scripts and sample files are described in the following table:

*Table 10. Disaster recovery scripts and configuration files*

| Disaster recovery scripts or files | Description |
|---|---|
| **data_collection.ksh** | Collects the following information about the source environment:<br>• System host name<br>• Network adapter information<br>• Host bus adapter (HBA) configuration<br>• Domain Name System (DNS) server and domain<br>• LPAR attributes<br>• Volume group attributes and hard disk attributes<br>• AIX kernel (`sys0`) configuration<br>**Note:** This script must be installed on the source LPAR in the `/usr/local/bin` directory. The **data_collection.ksh** script must be run on the source LPARs regularly.<br>The **data_collection.ksh** script collects system information and places it in the following locations:<br>**`/usr/local/dr/data` directory**<br>Contains system-customized information.<br>**`/usr/local/dr/data_default` directory**<br>Contains information about the default parameters for each device. |
| **setup_dr.ksh** | Reconfigures the environment of the backup LPAR to be the same as the source LPAR.<br>**Note:** All scripts must be installed in the `/usr/local/bin/` directory of the source LPARs so that these scripts are also available in the backup LPARs during a disaster. You must run this script on the backup LPAR during a disaster recovery event.<br>Consider the following prerequisites before you run this script:<br>• The LPAR must be started and running in the target site with the rootvg disk.<br>• Root passwords must be acquired and used to start and log in to the backup LPARs.<br>The **setup_dr.ksh** script calls other scripts automatically to perform the following tasks:<br>• Reconfigure the HBA adapters of the backup LPAR to be the same as the source LPAR.<br>• Reconfigure the Ethernet adapter of the backup LPAR by reading the contents of the `failover_config.cfg` configuration file and set the host name, IP address, and the base network of the backup LPAR.<br>• Reconfigure any additional Ethernet adapters on the backup LPAR by using the appropriate IP addresses.<br>• Import any volume groups from the source LPAR to the backup LPAR. |

*Table 10. Disaster recovery scripts and configuration files (continued)*

| Disaster recovery scripts or files | Description |
|---|---|
| `failover_config.cfg` | Contains sample configuration file for information about the backup LPAR.<br>You must manually edit this file and fill appropriate information about the AIX operating system configuration in the backup LPAR that are listed as follows:<br>• IP address of LPAR at the source site<br>• IP address of LPAR at the backup site<br>• Network netmask that must be used at the backup site<br>• DNS server that must be used at the backup site<br>• Network domain name that must be used at the backup site<br>• Default gateway IP address that must be used at the backup site<br>**Note:** The updated `failover_config.cfg` file must be placed in the `/usr/local/dr/data` directory of the source LPAR. |

After you complete the initial configuration on the source LPAR, complete the following steps:

1. Copy the script and sample files from the following location in the KSYS node:
   ```
   /opt/IBM/ksys/samples/
   ```
2. Install these script files in the following location of the source LPARs:
   ```
   /usr/local/bin/
   ```
3. Edit the `failover_config.cfg` configuration file with the appropriate LPAR information and place it in the `/usr/local/dr/data` directory.
4. Run the **data_collection.ksh** script regularly to collect the environment information. You can set your LPARs to run this script daily by using the AIX `cron` utility. For example, you can add the following line as a `crontab` entry to run this script daily at 1 AM:
   ```
   00 01 * * * /usr/local/dr_collection
   ```

When you move the source LPARs to the backup site during a disaster event, you can use the disaster recovery scripts to retain the environment of the source LPARs. The procedure to run the disaster recovery steps in the backup site assumes that the following prerequisites are met:

- All the disaster recovery scripts are installed on the source LPARs.
- The `failover_config.cfg` configuration file is manually edited with the appropriate backup LPAR information.
- The **data_collection.ksh** script is successfully run on the source LPARs.
- The `failover_config.cfg` configuration file is updated with the appropriate host name, IP address, netmask, name server that corresponds to the backup LPAR.
- The disk replication and split operations copied all of the source data to the backup site.
- The disk storage is available to the appropriate LPAR on the backup site.
- A disaster occurred, and the source LPARs are moved to the backup LPARs.
- The appropriate root passwords are acquired for the backup LPARs.
- The appropriate DNS server is available at the backup location.
- The backup LPARs are identified and are accessible through the HMC.

To run the disaster recovery scripts in the backup LPARs, complete the following steps:

1. To check the `/usr/local/dr/data` directory for date and time stamps, run the following command:
   ```
   ls -la /usr/local/dr/data
   ```

   Verify whether the replication was complete.
2. If the console messages are large, route the output to a file by running the following command:
   ```
   swcons /tmp/console.log
   ```

3. Run the **setup_dr.ksh** script to reconfigure the backup LPAR host name, network IP addresses, and HBAs, and to import the volume groups.

   **Note:** The LPARs must be completely recovered by now.
4. Verify the LPAR configuration to confirm the changes by completing the following checks:
   a. Run the **hostname** command to verify the LPAR host name.
   b. Run the **ifconfig -a** command to verify the LPAR IP address.
   c. Run the **df** or **mount** command to verify that all local file systems are mounted. Mount any unmounted file systems by using the **mount** command, if required.
   d. Run the **lsps -a** command to display all paging spaces and their status. For unavailable paging spaces, run the **swapon** command with the appropriate paging space to set the paging space to the active state.

# Troubleshooting GDR

Use this information to troubleshoot the GDR solution.

## Notification for the KSYS events

The KSYS subsystem tracks various events that occur in the environment and saves the information in log files. The KSYS subsystem also sends emails and text notifications to the administrator if the contact information is registered on the KSYS configuration by using the **ksysmgr add nofity** command.

You can run the following command to list all the events that can be notified:

```
ksysmgr query event
```

The events are categorized as critical errors, warning, 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 events that are monitored by the KSYS:

*Table 11. Notification for KSYS events*

| Events | Event type | Description | Suggested actions |
|---|---|---|---|
| HMC_UNREACHABLE | Critical | HMC is down or not reachable | |
| STG_UNREACHABLE | Critical | Storage subsystem is down or not reachable | Check the network between the KSYS node, EMC Solution Enabler, and the storage disks. Also, check whether the EMC Solution Enabler server is down. |
| HMC_REACHABLE | Informational | HMC recovered and is now reachable | |
| VIOS_RMC_STATE_DOWN | Critical | The RMC connectivity between HMC and VIOS has problems | |
| INSUFFICIENT_HOST_CAPACITY | Critical | Backup host does not have sufficient capacity to perform a successful disaster recovery failover | |
| VIOS_FAILURE | Critical | VIOS has failed | |
| VM_CONFIG_COLLECTION_FAILURE | Critical | Configuration data collection failed for the VM | |
| DAILY_VERIFY_FAILED | Critical | Daily verification has failed | |
| REPLICATION_FAILURE | Critical | Storage reports replication problem | Verify the replication between the local and the remote disks. |
| MIRROR_RELATIONSHIP_MISSING | Critical | Disk has no mirror pair | Create a mirror pair for the disk. |
| HOST_FAILURE | Critical | Host failure has occurred | |
| FILESYSTEM_SPACE_WARNING | Warning | File system is reaching full condition | Free up some space in the `/var` location by removing old log files or any unwanted files. |
| HMC_LOGIN_FAILURE | Critical | HMC login failed | |

*Table 11. Notification for KSYS events (continued)*

| Events | Event type | Description | Suggested actions |
|---|---|---|---|
| DISK_VALIDATION_FAILURE | Critical | Disk group validation failed | Investigate the disk groups and resolve the issue. |
| VIOS_DELETED | Warning | VIOS deletion has been detected | |
| VM_NOT_ACTIVE | Informational | VM is not active | |
| DUPLICATE_VMS | Informational | VM exists on multiple hosts | |
| VM_MOVE | Informational | VM has moved from one host to another | |
| DAILY_VERIFY_COMPLETE | Informational | Daily verification has completed successfully | |
| HOST_IN_INVALID_STATE | Informational | Host is in an invalid state | |
| VM_STORAGE_COLLECTION_FAILURE | Critical | Storage information collection has failed for the VM | Check VIOS for more details. |
| VM_DISCOVERED_ON_HOST | Informational | A new VM is discovered in a host | |
| VM_DELETED_FROM_HOST | Warning | A VM has been deleted from a host | |
| VM_NOT_FOUND | Critical | A VM does not exist | |
| VM_NOT_ACTIVE | Informational | VM is not active | |

## Analyzing the log files

If you receive errors while you run the **ksysmgr** command, you can analyze the log files to diagnose the issue. Determine the software component that is causing the problem by analyzing the log files. You can find the **ksysgmr** command log files in the `/var/ksys/log/` directory.

When you run the **ksysmgr** command, the following types of log files are created:

**`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.

**`ksysmgr.log`**
Contains the detailed processing information about each function when you run a **ksysmgr** command.

**Note:** The `ksysmgr.log` file contains the detailed processing information only when you specify the **`-l max`** flag when you run the **ksysmgr** command.

**`ksys_srdf.log`**
Contains the detailed processing information about all the EMC storage-specific functions along with the date and time.

If large number of virtual machines are included in the **ksysmgr**-related operations, you can increase the size of the trace files to accommodate more log information. You can increase the size of trace file by using the `/var/ct/cfg` file as follows:

```
/var/ct/cfg
# cat trace.conf
IBM.VMR:
    pat        = /var/ct/.*/log/mc/IBM.VMR/.*
    spooling   = OFF
    pages      = 4
    dest       = /tmp/ksys/rm/
    #size      = 8192000
    size       = 11534336
```

This example changes the size of trace files from 8 MB to 11 MB.

## Example for the `ksysmgr.oplog` file and the `ksysmgr.log` file

This example shows samples for the `ksysmgr.oplog` file and the `ksysmgr.log` files.

For instance, you can run the following commands to add a KSYS cluster called `test_cluster` and then, you can run the commands to verify and sync the cluster:

```
ksysmgr add cluster test_cluster ksysnodes=xxx.xx.xx.ibm.com // Creates a cluster called test_cluster.
ksysmgr verify ksyscluster test_cluster                      // Verifies the cluster.
ksysmgr -l max sync ksyscluster test cluster                 // To Sync the cluster.
                                                             // However, the cluster name is misspelled.
                                                             //       The output displays error.
ksysmgr -l max sync ksyscluster test_cluster                 // Syncs the cluster.
                                                             //       The output indicates success.
```

In this scenario, the contents of the `ksysmgr.oplog` file might be similar to the following sample:

```
8/3/2106 23:50:12   15401408    ksysmgr add cluster test_cluster ksysnodes=xxx.xx.xx.ibm.com
8/3/2106 23:54:35   15401222    ksysmgr verify ksyscluster test_cluster
8/3/2106 23:54:59   15401238    ksysmgr -l max sync ksyscluster test cluster
8/3/2106 23:55:5    10551612    ksysmgr -l max sync ksyscluster test_cluster
```

The contents of the `ksysmgr.log` file might be similar to the following sample. Because the **`-l max`** flag was not specified in the first two commands, the details for those commands are not logged in the `ksysmgr.log` file.

```
ksysmgr_setenv()[134]:  ENTERING with (ksysmgr_TRANSACTION_ID, 15401238, 809)
main()[811]:    Setting ksysmgr_TRANSACTION=15401238 in the environment.
main()[855]:    Processing the action input, "sync".
expand_action()[430]:   ENTERING with (sync, ksysmgr -l max sync ksyscluster test cluster)
expand_action()[472]:   Converted user action "sync" to "sync" using the "exact match" alias "sync".
main()[884]:    The fully expanded action is "sync".
main()[892]:    Processing the class input, "ksyscluster".
expand_class()[161]:    ENTERING with (ksyscluster, ksysmgr -l max sync ksyscluster test cluster)
expand_class()[224]:    Converted user class "ksyscluster" to "ksyscluster" using the "exact match"
                         alias "ksyscluster".
expand_class()[249]:    Expanded class: "ksyscluster")
main()[925]:    The fully expanded class is "ksyscluster".
main()[932]:    Processing the "object" input, "test".
main()[944]:    Processing ATTR=VALUE input, "cluster".
main()[954]:    Remaining space for attributes is 3001 bytes.
sync_ksyscluster()[541]:    ENTERING
sync_ksyscluster()[557]:    START PROCESSING INPUTS FROM USER
--------------------------------------------------
ksysmgr_setenv()[134]:  ENTERING with (ksysmgr_TRANSACTION_ID, 10551612, 809)
main()[811]:    Setting ksysmgr_TRANSACTION=10551612 in the environment.
main()[855]:    Processing the action input, "sync".
expand_action()[430]:   ENTERING with (sync, ksysmgr -l max sync ksyscluster test_cluster)
expand_action()[472]:   Converted user action "sync" to "sync" using the "exact match" alias "sync".
main()[884]:    The fully expanded action is "sync".
main()[892]:    Processing the class input, "ksyscluster".
expand_class()[161]:    ENTERING with (ksyscluster, ksysmgr -l max sync ksyscluster test_cluster)
expand_class()[224]:    Converted user class "ksyscluster" to "ksyscluster" using the "exact match"
                        alias "ksyscluster".
expand_class()[249]:    Expanded class: "ksyscluster")
main()[925]:    The fully expanded class is "ksyscluster".
main()[932]:    Processing the "object" input, "test_cluster".
sync_ksyscluster()[541]:    ENTERING
sync_ksyscluster()[557]:    START PROCESSING INPUTS FROM USER
sync_ksyscluster()[590]:    DONE PROCESSING INPUTS FROM USER
runSystemCommand()[198]:    RUNNING COMMAND:/usr/bin/grep -q -F $(/usr/bin/hostname)
                            /var/ct/cfg/ctrmc.acls
runSystemCommand()[198]:    RUNNING COMMAND:/usr/sbin/rsct/bin/mkrpdomain test_cluster
                            $(/usr/bin/hostname) > /dev/null 2>&1;
```

```
runSystemCommand()[198]:    RUNNING COMMAND:/usr/sbin/rsct/bin/startrpdomain test_cluster>
                            /dev/null 2>&1;
runSystemCommand()[170]:    ERROR: 1
runSystemCommand()[170]:    ERROR: 1
runSystemCommand()[170]:    ERROR: 1
runSystemCommand()[198]:    RUNNING COMMAND:/usr/bin/stopsrc -s IBM.VMR > /dev/null 2>&1;
                            /usr/bin/sleep 5;
                            /usr/sbin/rsct/bin/rmcctrl -z > /dev/null 2>&1;
                            /usr/sbin/rsct/bin/rmcctrl -A > /dev/null 2>&1;
                            /usr/bin/startsrc -s IBM.VMR > /dev/null 2>&1;
                            /usr/bin/sleep 5; /usr/sbin/rsct/bin/rmcctrl -A > /dev/null 2>&1;
runSystemCommand()[198]:    RUNNING COMMAND:/usr/bin/lssrc -g rsct_rm | /usr/bin/grep IBM.VMR |
                            /usr/bin/awk '{print $4}' | /usr/bin/grep -q -F -- 'active';
sync_ksyscluster()[663]:    EXITING SUCCESSFULLY
```

### Example for the `ksys_srdf.log` file

This example shows a sample of the `ksys_srdf.log` file:

For instance, you initiate a planned recovery from the active site to the backup site, and problems occur during the storage reverse replication. The output of the command displays the error. However, if you want to know the exact step where this problem occurred, you can review the `ksys_srdf.log` file as follows:

```
...
Wed Aug 17 04:45:03 CDT 2016 reverse_emc_srdf_cg: srdf_cg_query : Successfully Processed
                             /var/ksys/log/data-0-1Amy3-839-1471427099.xml
Wed Aug 17 04:45:03 CDT 2016 reverse_emc_srdf_cg: 119 : CG: VMRDG_Cheese50_vmdr_Site1 in Asynchronous
Wed Aug 17 04:45:03 CDT 2016 reverse_emc_srdf_cg: 136 : Info: CG VMRDG_Cheese50_vmdr_Site1 in PLAN
Wed Aug 17 04:45:03 CDT 2016 reverse_emc_srdf_cg: 150 : SID:196800573 : SOURCE for
                             VMRDG_Cheese50_vmdr_Site1
Wed Aug 17 04:45:03 CDT 2016 reverse_emc_srdf_cg: 161 : Info: CG VMRDG_Cheese50_vmdr_Site1 in
                             State:Consistent Type:RDF2
Wed Aug 17 04:45:03 CDT 2016 reverse_emc_srdf_cg: CG_DisableConsistency : CG:VMRDG_Cheese50_vmdr_Site1
                             Start
Wed Aug 17 04:45:03 CDT 2016 reverse_emc_srdf_cg: srdf_cg_cmd_exec : Command Executing : 11
Wed Aug 17 04:45:04 CDT 2016 reverse_emc_srdf_cg: log_srdf_msg : Error:
                             All the devices in the SRDF/A session must be managed together
                             when the devices are in async mode
Wed Aug 17 04:45:04 CDT 2016 reverse_emc_srdf_cg: CG_DisableConsistency : CG: VMRDG_Cheese50_vmdr_Site1
                             disable consistency failed:1
Wed Aug 17 04:45:04 CDT 2016 reverse_emc_srdf_cg: CG_EnableAcpMode : CG:VMRDG_Cheese50_vmdr_Site1 Start
Wed Aug 17 04:45:04 CDT 2016 reverse_emc_srdf_cg: srdf_cg_cmd_exec : Command Executing : 13
Wed Aug 17 04:45:05 CDT 2016 reverse_emc_srdf_cg: CG_EnableAcpMode : Failed with status 1
Wed Aug 17 04:45:05 CDT 2016 reverse_emc_srdf_cg: log_srdf_msg : Error:
An RDF Set 'ACp Disk Mode ON' operation execution is in
progress for composite group 'VMRDG_Cheese50_vmdr_Site1'. Please wait...
...
```

In this example, it is clear that the issue occurred during the `CG_DisableConsistency` phase.

## Solving common problems

When you receive errors for any **ksysmgr** commands, the command output shows the error and the suggested resolution. However, if you are not able to determine the issue, review the following approaches to diagnose the issue.

**The discovery operation failed.**

If a problem occurs during the discovery steps, you must analyze the `ksysmgr.log` file for any failures. After the discovery is complete, you can query the `IBM.VMR_LPAR` resource class to confirm the successful completion of discovery operation:

```
lsrsrc IBM.VMR_LPAR
```

The output might be similar to the following sample:

```
Name                = "xxx"
LparUuid            = "59C8CFxx-4Bxx-43E2-A0CE-F028AEB5Fxxx"
LparIPList          = {}
SiteCleanupTastList = {}
ActiveSiteID        = 80708xxxx
LCB                 = {  }
BootDiskList        = {}
CecUuid             = "6ce366c5-f05d-3a12-94f8-94a3fdfcxxxx"
ErrMsg              = ""
Phase               = "READY"
PhaseDetail         = 4194305
Memory              = "4352"
Processors          = "0.1"
ActivePeerDomain    = "vmdr"
```

In case of any errors in the discovery operation, the `Phase` field is set as `VERIFY` and the `ErrMsg` field indicates the error details. The `Phase` field is set as `READY` after a successful discovery operation.

**The discovery operation failed with the `getlcb` error.**

The cause for this error might be that the virtual machine's Fibre Channel port in the storage area network (SAN) fabric is zoned with a storage port that does not provide any logical unit numbers (LUNs) to the virtual machine. You can resolve this error by completing one of the following steps:

- Ensure that the virtual machine is zoned only with those storage ports that provide LUNs to the virtual machines.
- Run the **cfgmgr** command in the virtual machine where the `getlcb` failure occurred and then run the discovery operation again.

**The discovery operation failed indicating that the storage disk was already a part of an existing composite group.**

If any of the storage disks in the GDR solution are already a part of an existing composite group, the discovery operation cannot complete successfully. A storage disk in the GDR solution must be associated with a single composite group, which is asynchronously consistent. Remove the older composite groups, and run the discovery operation again.

**The discovery operation failed indicating that the disk group is created in only one site.**

Review the `/var/ksys/log/ksys_srdf.log` file for any consistency-enabled issue. Ensure all the disks that belong to a Remote Data Facility (RDF) group are also a part of the composite group.

**The verification phase failed.**

After the validation is complete, you can query the `IBM.VMR_LPAR` resource class to ensure that the virtual machines are ready to be moved during a disaster situation:

```
lsrsrc IBM.VMR_LPAR
```

The output might be similar to the following sample:

```
Name                = "xxx"
LparUuid            = "59C8CFxx-4Bxx-43E2-A0CE-F028AEB5Fxxx"
LparIPList          = {}
SiteCleanupTastList = {}
ActiveSiteID        = 80708xxxx
LCB                 = {  }
BootDiskList        = {}
CecUuid             = "6ce366c5-f05d-3a12-94f8-94a3fdfcxxxx"
ErrMsg              = ""
Phase               = "READY_TO_MOVE"
PhaseDetail         = 4194305
Memory              = "4352"
Processors          = "0.1"
ActivePeerDomain    = "vmdr"
```

In case of any errors in configuration validation, review the error details in the `ErrMsg` field. The `Phase` field is set as `READY_TO_MOVE` after a successful verification operation.

**The test-discovery step in the DR failover rehearsal operation is failing with the error message: `Tertiary storage copy is missing.`**

This error occurs when one or more of the third copy of disks required for cloning the backup storage data is missing. For each backup (2nd copy) disk, a corresponding tertiary (third copy) disk must exist in the backup site. Check the availability and accessibility of the tertiary disks in the storage subsystem. You can also check the status of cloning relationship by using commands that are provided by the specific storage vendor.

**The test-discovery step in the DR failover rehearsal operation is failing with the error message: `Storage agent is not accessible.`**

This error occurs because of a problem in communication between the KSYS subsystem and storage subsystem. Check for any hardware issues. For example, ensure proper connectivity between all the subsystems. Also, identify the issue by analyzing resource manager trace log files.

**The HMC interface indicates that the LPAR has no Resource Monitoring Control (RMC) connection or the RMC is inactive.**

Check whether the LPAR properties also indicate an RMC issue between the HMC and VIOS. The RMC connectivity issue can occur because of the security mode that is set in the LPAR. The security mode for both the HMC and LPAR must be set to the same value.

For example, list the security mode for LPAR by running the following command:

```
/usr/sbin/rsct/bin/lssecmode
```

The output might look like the following sample:

```
Current Security Mode Configuration
Compliance Mode : nist_sp800_131a
Asymmetric Key Type : rsa2048_sha256
Symmetric Key Type : default
```

Similarly, list the security mode for the HMC by running the following command:

```
/usr/sbin/rsct/bin/lssecmode
```

The output might look like the following sample:

```
Current Security Mode Configuration
Compliance Mode : none
Asymmetric Key Type : rsa512
Symmetric Key Type : default
```

In this case, the LPAR has the `nist_xxx` security mode enabled, but the HMC has no security mode. This mismatch can occur if another HMC was connected or a security mode was set before any reset operation was started.

**Errors occur when you register an EMC storage agent.**

Check the `/var/symapi/config/netcnfg` file to determine whether the configuration contains at least two EMC subsystems.

**You want to view all the storage disks in the active site and the backup site.**

Run the following command to list all the storage disks:

```
# symcfg list
```

The output might be similar to the following sample:

```
                                S Y M M E T R I X

                                       Mcode    Cache      Num Phys  Num Symm
    SymmID       Attachment  Model     Version  Size (MB)  Devices   Devices
```

```
000196800508 Local       VMAX100K  5977      217088        65      9964
000194901326 Remote      VMAX-1SE  5876       28672         0       904
000196800573 Remote      VMAX100K  5977      217088         0      7275
000198701861 Remote      VMAX10K   5876       59392         0       255
```

**The output of ksysmgr query command is not updated for hosts, HMCs, and VIOS even when the entities are updated.**

Sometimes, the **ksysmgr query** command displays static information even when the hosts or HMCs are modified. To update the **ksysmgr query** command output dynamically, complete the following steps:

1. Unpair the hosts across the sites by using the following command:

   ```
   ksysmgr pair host host_name pair=none
   ```

2. Remove the hosts from the site by using the following command:

   ```
   ksysmgr remove host host_name
   ```

3. Add the new HMC to the site by using the following command:

   ```
   ksysmgr add hmc hmc_name login=login_name password=login_password
          ip=ip_address  site=site_name
   ```

4. Add the corresponding managed hosts to the site by using the following command:

   ```
   ksysmgr add host host_name
   ```

**The disaster recovery move operation failed.**

When a disaster occurs and you initiate a move operation, the KSYS subsystem coordinates the start of the virtual machines on the backup site. During this process, LPAR profiles are created through HMCs on the backup site hosts. During the LPAR profile creation, if any errors occur such that the KSYS cannot communicate with HMCs, the move operation might fail. At that time, any virtual machines that are partially created on HMC require manual restart. For the rest of the virtual machines, you can use the **ksysmgr recover** command to recover and start the virtual machines.

**An unplanned move operation from the active site to the backup site is successfully completed and the specified flexible capacity policy is followed. Later, another unplanned move operation from the backup site to the source site fails. When the virtual machines are recovered to the source site, the virtual machines are started but without any change in the processor and memory values, (that is, without following the flexible capacity policy).**

This situation can happen when the active hosts and backup hosts are connected to the same HMC. You must connect the source hosts and target hosts to different HMCs to continue unplanned move operations in case of source HMC failures.

# Commands

The following commands are commonly used to obtain information about the KSYS environment or to run a specific function.

For complete information about command's capabilities and restrictions, see the man page.

To view the man page information for a command, enter the following command:

```
man command_name
```

## ksysmgr command

### Purpose

The **ksysmgr** command provides a consistent interface to configure the controller system (KSYS) and to perform Geographically Dispersed Resiliency for Power Systems operations by using a terminal or script.

### Syntax

```
ksysmgr [-v] [-f] [-l {low|max}]
[-a {<ATTR#1>,<ATTR#2>,...}] <ACTION> <CLASS> [<NAME>]
[-h | <ATTR#1>=<VALUE#1> <ATTR#2>=<VALUE#2> ...]

ksysmgr [-v] [-f] [-l {low|max}]
[-a {<ATTR#1>,<ATTR#2>,...}] <ACTION> <CLASS> [<NAME>]
<ATTR#1>=<VALUE#1> <ATTR#2>=<VALUE#2> ...]

        ACTION={add|modify|delete|query|manage|unmanage|...}\n\
        CLASS={ksyscluster|site|hmc|host|...}\n\

ksysmgr {-h|-?} [-v] [<ACTION> [<CLASS>]]

ksysmgr [-v] help
```

The basic format for using the **ksysmgr** command is as follows:

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

**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 is different 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 of the **ksysmgr** command's individual ATTRIBUTES.

© Copyright IBM Corp. 2017, 2019

## Description

All **ksysmgr** command operations are logged in the /var/ksys/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 overall return code.

## Flags

**ACTION**

Describes the action to be performed.

**Note:** The **ACTION** flags are not case-sensitive. All **ACTION** flags provide a shorter alias. For example, rm is an alias for delete. Aliases are provided for convenience from the command line and must not be used in scripts.
The following **ACTION** flags are available:

**Note:** The asterisk (*) in the aliases signify wildcard characters. For example, for the modify ACTION, the alias value is mod*. If you type modd, the command still works.

- query (alias: q*, ls, get, sh*)
- add (alias: ad*, create, cr*, make, mk, bu*, bld)
- delete (alias: de*, remov*, rm, er*)
- modify (alias: mod*, ch*, set, sets)
- verify (alias: ver*)
- sync (alias: syn*, pr*)
- restore (alias: rest*)
- manage (alias: man*, mg)
- unmanage (alias: unman*, umg)
- discover (alias: di*)
- help (alias: hel*, ?)
- move (alias: mov*, mv, swi*)
- recover (alias: rec*)
- pair (alias: map)
- cleanup (alias: clean*)

**CLASS**

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

The following **CLASS** objects are supported:

- ksyscluster (alias: cl*, ksyscl*)
- site (alias: sit*)
- hmc (alias: hmcs, hmces)
- host (alias: serv*, mach*, cec*, ho*)
- host_group (alias: serv*, mach*, cec*, ho*)
- vios (alias: vi*)
- vm (alias: lp*, vm*)
- disk_pair (alias: dp, disk_p*)
- disk_group (alias: dg, disk_g*)
- storage_agent (alias: storage_a*, sta)
- version (alias: vers*)
- notify (alias: rn, remote_not*, noti*)

- snapshot (alias: snap*)
- script

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

**ATTR=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 **ATTR** 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|max**
Activates the following trace logging values for troubleshooting purposes:

**low (default)**
Logs basic information for every **ksysmgr** operation.

**max** Performs high tracing operations such as adding the routine function and the utility function. Adds transaction ID to the function's entry messages.

**Note:** All trace data is written into the ksysmgr.log file. This flag is ideal for troubleshooting problems.

**-v** Displays maximum verbosity in the output.

## Exit status

**RC_UNKNOWN (Exit value = -1)**
Result is not known. This value is used as an initializer.

**RC_SUCCESS (Exit value = 0)**
No errors are detected. The operation is successful.

**RC_ERROR (Exit value = 1)**
An error occurred.

**RC_NOT_FOUND (Exit value = 2)**
The specified resource does not exist or cannot be found.

**RC_MISSING_INPUT (Exit value = 3)**
Required input is missing.

**RC_INCORRECT_INPUT (Exit value = 4)**
Detected input is incorrect.

## Examples

- To get help information about the ksyscluster class, enter the following command:

  ```
  ksysmgr help ksyscluster
  ```

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

```
Available actions for ksyscluster:
         add
         delete
         query
         sync
         verify
```

## Cluster configuration examples

- To add a KSYS cluster, use the following command syntax:

  ```
  ksysmgr add ksyscluster clustername ksysndoes=nodename
       type=DR [sync=true]
  ```

  For example,

  ```
  ksysmgr add ksyscluster SampleClusterName ksysnodes=ksysnode1
       type=DR sync=true
  ```
- To query the KSYS cluster, enter the following command:

  ```
  ksysmgr query cluster
  ```

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

  ```
  Name:            vmr
  State:           Online
  ```
- To remove a KSYS cluster, use the following command syntax:

  ```
  ksysmgr delete ksyscluster clustername
  ```

  When you delete a KSYS cluster, the **ksysmgr** command prompts for your confirmation. The **ksysmgr** command also recommends to create a snapshot at this point. You can override these prompts by using the **-f** flag. For example,

  ```
  ksysmgr -f delete ksyscluster SampleClusterName
  ```

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

  ```
  WARNING: This action will remove all configuration and destroy the KSYS setup,
  it is recommended to create a backup "ksysmgr add snapshot -h"
  Consistency group cleanup successful
  Peer domain stopped successfully
  Peer domain was removed successfully
  ```

## Site configuration examples

- To create a site in the KSYS subsystem, use the following command syntax:

  ```
  ksysmgr add site sitename sitetype=home|backup
  ```

  For example,

  ```
  ksysmgr add site SiteA sitetype=home
  ```

  No output is displayed. The command returns with the exit value of 0.
- To query the details about all existing sites, enter the following command:

  ```
  ksysmgr query site
  ```

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

  ```
  Replication type for site(s): async
  Name:             SiteA
  Sitetype:        HOME

  Name:             SiteB
  Sitetype:        BACKUP
  ```

The site that is currently running the active virtual machines is labeled as home. By default, the replication type of the site is async.

- To query the details about a specific site, use the following command syntax:

  ```
  ksysmgr query site sitename
  ```

  For example,

  ```
  ksysmgr query site
  ```

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

  ```
  Name:                site2
  Sitetype:            BACKUP
  Host_groups:         None

  Name:                site1
  Sitetype:            HOME
  Host_groups:         None
  cpu_capacity:        none
  memory_capacity:     none
  skip_resource_check: No
  skip_power_on:       No
  vswitchmap:          none
  drvswitchmap:        none
  vlanmap:             none
  drvlanmap:           1/4,2/5,3/6
  ```

- To discover a site, use the following command syntax:

  ```
  ksysmgr discover site sitename
  ```

  For example,

  ```
  ksysmgr discover site SiteA
  ```

  The KSYS subsystem discovers all the hosts and virtual machines from all the host groups across both the sites. Discovery operation might take a few minutes to complete.

- To modify the configuration information of a site, use the following command syntax:

  ```
  ksysmgr modify site <sitename[,sitename2,...]> | file=<filepath>
  [name=<newsitename>]
  [memory_capacity=<(Whole Number >= 1)> | minimum | current_desired | none]
  [cpu_capacity=<(Whole Number >= 1)> | minimum | current_desired | none]
  [skip_resource_check=<yes|no>]
  [skip_power_on=<yes|no>]
  [network=<vlanmap | vswitchmap>  backupsite=siteB
    sitename=<#[,#,...] || all> siteB=<#[,#,...] || all> [dr_test=<yes|no>]
  [policy=<delete>]]
  ```

  For example,

  ```
  ksysmgr modify site site1 network=vlanmap backupsite=site2
    site1=1,2,3 site2=4,5,6 dr_test=yes
  ```

- To change the replication type of a site from the default async value to the sync value, enter the following command:

  ```
  ksysmgr modify system replication_type=sync sites=Site1,Site2
  ```

## HMC configuration examples

- To add an HMC to the KSYS configuration, use the following command syntax:

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

```
      site=site_name
      [hmctimeout=value]
      [maxjobs=value]
      [SwXSDVersion=value]
```

For example,

```
ksysmgr add hmc PrimaryHmcName login=hscroot password=xxx ip=86.xx.xx.xx site=SiteA
```

No output is displayed. The command returns with the exit value of 0.

You can also run the command without specifying the password in the command line. In this case, you can enter the password as hidden characters when the command prompts for the password. For example,

```
ksysmgr add hmc PrimaryHmcName login=hscroot ip=86.xx.xx.xx site=SiteA
Enter Password for hmc: ***********
Re-Enter Password: ************
```

- To query the details about all HMCs, enter the following command:

```
ksysmgr query hmc
```

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

```
Name:            PrimaryHmcName
Site:            SiteA
Ip:               86.xx.xx.x
Login:           hscroot
                        Managed Host List:
Name                                    Uuid
================                        ================
cola_8286-42A-2120DEW                   7d35be3a-a9b3-3cdf-a31e-80958bd2b9c8
pepsi_8286-42A-2120DFW                  21b4b05f-9b84-349c-9ce9-d03f0e78f9f7
========================================================================

Name:            BackupHmcName
Site:            SiteB
Ip:              9.3.18.34
Login:           hscroot
                        Managed Host List:
Name                                    Uuid
================                        ================
pepsi_8286-42A-2120DFW                  21b4b05f-9b84-349c-9ce9-d03f0e78f9f7
cola_8286-42A-2120DEW                   7d35be3a-a9b3-3cdf-a31e-80958bd2b9c8
========================================================================
```

- To query the details about a specific HMC, use the following command syntax:

```
ksysmgr query hmc <hmcname>
```

For example,

```
ksysmgr query hmc PrimaryHmcName
```

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

```
Name:            PrimaryHmcName
Site:            SiteA
Ip:               86.75.30.9
Login:           username
                        Managed Host List:
Name                                    Uuid
================                        ================
cola_8286-42A-2120DEW                   7d35be3a-a9b3-3cdf-a31e-80958bd2b9c8
pepsi_8286-42A-2120DFW                  21b4b05f-9b84-349c-9ce9-d03f0e78f9f7
========================================================================
```

- To modify the details of a specific HMC, use the following command syntax:

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

For example, to modify the login and password details of the PrimaryHmcName HMC, enter the following command:

```
ksysmgr modify hmc PrimaryHmcName login=scott password=tiger
```

No output is displayed. The command returns with the exit value of 0.

- To remove an HMC from the KSYS configuration, use the following command syntax:

```
ksysmgr delete hmc <hmcname>
```

For example,

```
ksysmgr delete hmc PrimaryHmcName
```

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

```
HMC resource removed
```

## Host configuration examples

- To add a host to the KSYS resource manager, use the following command syntax:

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

For example,

```
ksysmgr add host Site1_host1 site=Site1
```

No output is displayed. The command returns with the exit value of 0.

- To list details about all hosts, use the following command syntax:

```
ksysmgr query host
```

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

```
Name:            Site1_host1
UUID:            21b4b05f-9b84-349c-9ce9-d03f0e78f9f7
FspIp:           10.xx.1.xxx
Pair:            Site1_host1
Site:            Site1
VIOS:            Site1_VIOS1
                 Site1_VIOS2
HMCs:            Site1_HMC1

Name:            Site2_host1
UUID:            7d35be3a-a9b3-3cdf-a31e-80958bd2b9c8
FspIp:           10.40.1.161
Pair:            Site1_host1
Site:            Site2
VIOS:            Site2_VIOS1
                 Site2_VIOS2
HMCs:            Site2_HMC1
```

- To list details about a specific host, use the following command syntax:

```
ksysmgr query host hostname
```

For example,

```
ksysmgr query host Site1_host1
```

- To pair two hosts across sites, use the following command syntax:

  ```
  ksysmgr pair host hostname pair=hostname
  ```

  For example,

  ```
  ksysmgr pair host Site1_host1 pair=Site2_host1
  ```

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

  ```
  Site1_host1 on site1 has been paired
  with Site2_host1 on Site2 successfully.
  ```
- To unpair two hosts across sites, use the following command syntax:

  ```
  ksysmgr pair host hostname pair=none
  ```

  For example,

  ```
  ksysmgr pair host Site1_host1 pair=none
  ```

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

  ```
  Site1_host1 on site1 has been unpaired
  with Site2_host1 on Site2 successfully.
  ```
- To remove a host from the KSYS resource manager, use the following command syntax:

  ```
  ksysmgr delete host hostname
  ```

  For example,

  ```
  ksysmgr delete host Site1_host1
  ```

## Host group configuration examples

- To create a host group in the active site and to add the existing hosts to this new host group, use the following command syntax:

  ```
  ksysmgr add host_group hg_name
        site=sitename
        hosts=host1,[host2,...] | file=filepath
  ```

  where, the **file** parameter is an XML file that contains a list of hosts.
- To add or remove hosts from the existing host groups, use the following command syntax:

  ```
  ksysmgr modify host_group hg_name add | remove
  hosts=host1,[host2,...] | file=filepath
  ```
- To modify the capacity-related attributes for all the hosts in a host group, use the following command syntax:

  ```
  ksysmgr modify host_group hg_name options
        [memory_capacity=(Whole Number > 1) | minimum | current_desired | none]
        [cpu_capacity=(Whole Number > 1) | minimum | current_desired | none]
        [skip_resource_check=yes|no]
        [skip_power_on=yes|no]
  ```
- To create a network mapping policy of VLAN ID or virtual switches for all the hosts in a host group, across sites, use the following command syntax:

  ```
  ksysmgr modify host_group hg_name options
        network=vlanmap | vswitchmap  sites=siteA,siteB
         siteA=<#,[#,...]> siteB=<#,[#,...]>
  ```
- To discover and verify all the hosts in a specific host group, use the following command syntax:

  ```
  ksysmgr discover host_group hg_name verify=true
  ```
- To delete a host group, use the following command syntax:

  ```
  ksysmgr delete host_group hg_name
  ```
- To query the details about all host groups or a specific host group, use the following command syntax:

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

## LPAR configuration examples

- To include or exclude a specific virtual machine from the KSYS configuration, use the following command syntax:

  ```
  ksysmgr -h manage|unmanage VM_name
  ```

  The excluded virtual machine is not moved to the backup site when a site-switch operation is initiated.

  **Notes:**

  – After including or excluding a virtual machine, you must run the discovery and verification commands to rediscover the resources and validate the modified configuration setting.

  – If you use SAN Volume Controller storage system, before you include a specific virtual machine to the KSYS subsystem, you must ensure that the new virtual machine is associated with the storage volumes that have the same relationship type as the other existing virtual machines. For more information about this restriction, see "SAN Volume Controller system and Storwize system" on page 19.

- To update the priority of virtual machines, use the following syntax:

  ```
  ksysmgr modify VM name1[,name2,name3,...] | file=filepath
     [uuid=uuid_value]
     [host=hostname]
     [priority=low|medium|high]
  ```

  where, the **file** parameter is an XML file that contains a list of virtual machine references.

## Storage configuration examples

- To add a storage agent to a site, use the following command syntax:

  ```
  ksysmgr add storage_agent storage_agent_name
        hostname=hostname>|ip=ip
        site=sitename
        storagetype=type
        serialnumber=number | clusterid=number
  ```

  For example,

  ```
  ksysmgr add storage_agent StorageNamePrimary
  site=Site1
  serialnumber=000196xxx
  storagetype=emc
  ip=1.2.xx.xx
  ```

- To list storage agent details, use the following command syntax:

  ```
  ksysmgr query storage_agent storage_agent_name
  ```

  For example,

  ```
  ksysmgr query storage_agent StorageNamePrimary
  ```

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

  ```
  Name:             StorageNamePrimary
  Serialnumber:     00019xxxx
  Storagetype:     EMC
  Site:            Site1
  Ip:               1.2.xx.xx
  ```

- To remove a storage agent from a site, use the following command syntax:

  ```
  ksysmgr delete storage_agent storage_agent_name
  ```

- To resynchronize the storage data after an unplanned recovery from the active site to the backup site, use one of the following commands:

  ```
  ksysmgr resync site active_site_name
  ksysmgr resync host_group active_hg_name
  ```

If the unplanned move operation was at site level, you must run the **ksysmgr resync** command at site level. Similarly, if a virtual machine was moved to the backup site in an unplanned move operation at host group level, you must run the **ksysmgr resync** command at host group level.

## Discovery and verification examples

- To discover the resources in a site, use the following command syntax:

  ```
  ksysmgr discover site sitename
  ```

  For example,

  ```
  ksysmgr discover site site1
  ```

- To discover resources and to verify the KSYS configuration, use the following command syntax:

  ```
  ksysmgr discover site sitename verify=true
  ```

  For example,

  ```
  ksysmgr discover site site1 verify=true
  ```

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

  ```
  Running discovery on entire site, this may take few minutes...
          Discovery has started for VM VM_1
          Configuration information retrieval started for VM VM_1
          Discovery has completed for VM VM_1
          Disk Group creation on storage subsystem started for Site site1
          Disk Group creation on storage subsystem started for Site site2
          Disk Group creation on storage subsystem completed for Site site1
          Disk Group creation on storage subsystem completed for Site site2
  Discovery has finished for site1
  1 out of 1 VMs have been successfully discovered
  ```

## Script configuration examples

- To add a script for automatic execution before or after the discovery and verification operations, use the following command syntax:

  ```
  ksysmgr add script entity=site|host_group pre_verify|post_verify=script_path
  ```

  **Note:** The `pre_verify` and `post_verify` scripts can be run only at site or host group level.

- To add a script for automatic execution before or after the disaster recovery move operation, use the following command syntax:

  ```
  ksysmgr add script entity=site|host|host_group
      pre_offline|post_offline|pre_online|post_online=script_path
  ```

## Events query examples

- To query the events of a specific type, use the following command syntax:

  ```
  ksysmgr query event [type=error|warning|info]
  ```

  For example,

  ```
  ksysmgr query event type=error
  ```

- To query the events of all types, run the following command:

  ```
  ksysmgr query event
  ```

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

  ```
  Event Name                    Description
  ---------------------------------------------------------------
  HMC_UNREACHABLE               HMC is down or not reachable
  STG_UNREACHABLE               Storage subsystem is down or not reachable
  HMC_REACHABLE                 HMC has recovered and is now reachable
  VIOS_RMC_STATE_DOWN           HMC to VIOS RMC connectivity seems to be having problems
  ```

```
INSUFFICIENT_HOST_CAPACITY     Backup host does not have sufficient capacity
                               for a successful DR failover
VIOS_FAILURE                   VIOS seems to have failed
VM_CONFIG_COLLECTION_FAILURE   Configuration data collection failed for the VM
DAILY_VERIFY_FAILED            Daily verification checks have failed
REPLICATION_FAILURE            Storage reports replication problem
MIRROR_RELATIONSHIP_MISSING    Disk has no mirror pair
HOST_FAILURE                   Host failure has occurred
FILESYSTEM_SPACE_WARNING       Filesystem is reaching full condition
VM_MOVE                        VM has moved from one host to another
DAILY_VERIFY_COMPLETE          Daily verification checks have completed successfully
HOST_IN_INVALID_STATE          Host is in invalid state
VM_STORAGE_COLLECTION_FAILURE  Storage information collection has failed for the VM
HMC_LOGIN_FAILURE              HMC login failed
DISK_VALIDATION_FAILURE        Disk Group validation failure
VIOS_DELETED                   VIOS deletion has been detected
VM_NOT_ACTIVE                  VM does not seem to be active
DUPLICATE_VMs                  VM exists on multiple hosts
```

## Notification configuration examples

- To add an email, pager, or SMS notification for a specific user, use the following command syntax:

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

  For example,

  ```
  ksysmgr add notify user=John contact=john.doe@testmail.com
  ksysmgr add notify user=John contact=1234567890@tmomail.net
  ksysmgr add notify user=John contact=1234567890@SKYTEL.COM
  ```

- To modify an email address, pager email address, or SMS number for a specific user, use the following command syntax:

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

  For example,

  ```
  ksysmgr modify notify oldcontact=John newcontact=Dave
  ksysmgr modify notify oldcontact=john@gmail.com newcontact=dave@gmail.com
  ```

- To query all the registered contact details, use the following command:

  ```
  ksysmgr query notify contact
  ```

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

  ```
  User:            Mark Smith
  Contact:       mike@mike.com

  User:            joe
  Contact:       joe@gmail.com
  ```

- To delete all the contact information for a specific user, use the following command syntax:

  ```
  ksysmgt delete notify user=username
  ```

  For example,

  ```
  ksysmgt delete notify user=John
  ```

- To add a script for a predefined set of notifications and subsequent actions for a specific event, use the following command syntax:

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

  For example,

  ```
  ksysmgr add notify script=/tmp/script.sh events=HMC_DOWN
  ```

- To modify a script, use the following command syntax:

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

  For example,

  ```
  ksysmgr modify notify oldscript=/tmp/script.sh newscript=/tmp/newscript.sh
  ```

- To remove a script, use the following command syntax:

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

  For example,

  ```
  ksysmgr delete notify script=/tmp/script.sh
  ```

- To query a script, use the following command:

  ```
  ksysmgr query notify script
  ```

## System-wide attributes configuration

- To query details about system-wide attributes, enter the following command:

  ```
  ksysmgr query system
  ```

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

  ```
  System-Wide Persistent Attributes
  auto_discovery_time  =24 hours
  lose_vios_redundancy ="no"
  auto_reverse_mirror  ="yes"
  notification_level   ="high"
  dup_event_processing ="no"
  replication_type     ="Asynchronous"
  ```

- To enable the KSYS subsystem to automatically rediscover the resources twice a day, enter the following command:

  ```
  ksysmgr modify system auto_discover_time=12
  ```

- To change the notification level of your system to receive notification for all critical errors and warnings of all events, enter the following command:

  ```
  ksysmgr modify system notification_level=medium
  ```

- To change the duplicate event processing option to receive notification for all events, even if the events are duplicated, enter the following command:

  ```
  ksysmgr modify system dup_event_processing=yes
  ```

- To change the storage replication mode between two sites from asynchronous mode to synchronous mode, enter the following command:

  ```
  ksysmgr modify system replication_type=sync sites=SiteA,SiteB
  ```

- To enable the network mapping function and to create network mapping policy for all hosts and host groups across the active and the backup site, enter the following command:

  ```
  ksysmgr modify system network_mapping=enable network=vlanmap|vswitchmap sites=siteA,siteB
         siteA=<#,[#,...]> siteB=<#,[#,...]>]
  ```

## Disaster recovery failover rehearsal examples

- To discover, set up, and enter into DR failover rehearsal mode, use the following command syntax:

  ```
  ksysmgr discover host_group|site name dr_test=yes|no
  ```

  For example:

  ```
  ksysmgr discover host_group hg1 dr_test=yes
  ```

- To check whether the test environment is ready for the test-move operation, use the following command syntax:

  ```
  ksysmgr verify host_group|site name dr_test=yes|no
  ```

For example:

```
ksysmgr verify host_group hg1 dr_test=yes
```

- To start the VMs in the DR test environment, use the following command syntax:

  ```
  ksysmgr move host_group|site name from=Site1 to=Site2 dr_test=yes|no
  ```

  For example:

  ```
  ksysmgr move host_group hg1 from=Site1 to=Site2 dr_test=yes
  ```

- To clean up the DR test environment and to retain the normal DR backup state, use the following command syntax:

  ```
  ksysmgr cleanup host_group|site name dr_test=yes|no
  ```

  For example:

  ```
  ksysmgr cleanup host_group hg1 dr_test=yes
  ```

## Disaster recovery operation examples

- To initiate a site-switch operation, use the following command syntax:

  ```
  ksysmgr move site|host_group hg_name
        from=sitename
        to=sitename
        [force=true|false]
        [lose_vios_redundancy=yes|no]
        [dr_type=planned|unplanned]
        [cleanup=yes|no]
  ```

  Where, the default value of the **force** attribute is `false`, the default value of the **lose_vios_redundancy** attribute is `no`, and the default value of the **dr_type** attribute is `planned`. For example:

  ```
  ksysmgr move site from=Site1 to=Site2 dr_type=planned cleanup=no
  ```

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

  ```
  Site move started for Site1 to Site2
  Shutdown on Site1 site has started for VM Site1_VM
  Shutdown on Site1 site has completed for VM Site1_VM
  Storage mirror reversal has started
  Mirroring will be setup from Site2 to Site1
  Storage mirror reversal has completed
  Restart on Site2 site has started for VM Site1_VM
  Restart on Site2 site has completed for VM Site1_VM
  Configuration cleanup successful on Site1 site for VM Site1_VM
  Move has completed for VM Site1_VM
  Site move completed from Site1 to Site2
  1 out of 1 Vms have been successfully moved from Site1 to Site2
  Site2 is now the active site
  ```

- To recover failed virtual machines after the move operation is complete, use the following command syntax:

  ```
  ksysmgr recover host_group name
  ```

- To clean up a site, use the following command syntax:

  ```
  ksysmgr cleanup site|host_group|VM name
  ```

  If you do not specify the **cleanup** attribute, for a planned recovery, the KSYS subsystem automatically cleans up the source site from the location where the switch was initiated. For an unplanned recovery, you must manually clean up the source site after the HMC and hosts become operational. For example:

  ```
  ksysmgr cleanup site SiteA
  ```

## Snapshot examples

- To save a snapshot of the KSYS cluster configuration and resources, use the following command syntax:

```
ksysmgr add snapshot filepath=full_file_prefix_path|file_prefix type=cluster|basic|detailed ]
```

For example,

```
ksysmgr add snapshot filepath=/home/ksysdir/myksysbackup type=basic
```

- To view a snapshot, use the following command syntax:

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

  For example,

  ```
  ksysmgr query snapshot filepath=/home/ksysdir/myksysbackup_2016_06_23_04:54:30.xml.tar.gz
  ```

- To restore the configuration data on a KSYS node, use the following command syntax:

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

  For example:

  ```
  ksysmgr restore snapshot filepath=/home/ksysdir/myksysbackup_2016_06_23_04:54:30.xml.tar.gz
  ```

  This command decompresses and unarchives the snapshot file, and then applies the configuration settings to the KSYS node.

- To delete a snapshot, use the following command syntax:

  ```
  ksysmgr delete snapshot filepath=full_file_prefix_path
  ```

  For example,

  ```
  ksysmgr delete snapshot filepath=/home/ksysdir/myksysbackup_2016_06_23_04:54:30.xml.tar.gz
  ```

# ksysrppmgr command

## Purpose

Manage the Power Enterprise Pool resources and the Elastic Capacity on Demand (CoD) resources, which was formerly known as On/Off CoD. The **ksysrppmgr** command adjusts available resources on the managed hosts; you do not need to check the current available resources. This command also minimizes your resource costs by optimizing the local consumption of pool resources. For example, this command can convert Elastic (On/Off) CoD resources to Enterprise Pool CoD resources automatically. The **ksysrppmgr** command can also be used if the current level of resources is sufficient for your requirements. In this case, the command might release unnecessary resources.

## Syntax

```
ksysrppmgr -o action
  -h [hmc_name]:hmc_uri:hmc_user
  [-m ms_name:ms_action:[onoff_use(n|y[nb_days])]:[mem_amount]:[proc_amount]]
  [-M ms_uuid:ms_action:[onoff_use(n|y[nb_days])]:[mem_amount]:[proc_amount]]
  [-e enterprisepool_name]
  [-E enterprisepool_uuid]
  [-v] [-r] [-l none|logfile_path] [-p logfile_prefix]
```

## Description

**Note:** To use the **ksysrppmgr** command, you must authenticate to the HMCs by using the **hmcauth** command. The **ksysrppmgr** command communicates to the HMC through the REST APIs. Therefore, the APIs must be activated on the existing HMCs. At least one HMC is necessary to execute resources requests.

When you run the **ksysrppmgr** command, the HMC retrieves information about the topology of the managed hosts and resources. Then, the resource requests are processed. You do not need to check the current level of available resources to run the **ksysrppmgr** command. The **ksysrppmgr** command automatically determines whether a resource acquisition or a resource release is required to fulfill your resource request.

You can run the **ksysrppmgr** command in the following modes:

- Check mode
- Execution mode

In both the modes, you must first explore the topology and resource levels by logically connecting all the entities. When you explore the topology, you can identify the Enterprise Pools that are associated with a managed host. During the topology exploration, the REST universally unique identifier (UUID) and information about the entities are also read. For example, you can identify whether the Elastic (On/Off) CoD is enabled for a specific managed host. After you explore the topology, you can use the results to optimize the requirements of the subsequent requests.

When you run the **ksysrppmgr** command in check mode, all the execution steps are performed except that all the HMC operations are simulated after the initial processing. The initial process contains the topology exploration. During the topology exploration, the HMC reads the status of resources for all entities (pools and managed hosts) that are involved in the operation, because the current values must be used as base values for all the simulation. The command execution in the check mode might show some differences as compared to the results in the execution mode when you use Elastic (On/Off) CoD facility because the HMC might contain some information about the current day activations that are available only in execution mode.

When you run the **ksysrppmgr** command that contains multiple resource requests, the requests are executed sequentially. If you want to execute the resource requests simultaneously, you can run the **ksysrppmgr** command separately for each managed host. In this case, the topology cost must be paid for each command execution.

If you are using the Power Enterprise Pools facility, the **ksysrppmgr** command handles the conflicts in resource requests. If the same managed host has contradictory requests, the available amount of resources cannot be guaranteed but will be at least at the lowest request value, and at most at the highest request value. For example, if one CPU is available in the `host_1` host, and if you run three **ksysrppmgr** request simultaneously for the `host_1` host each requesting two CPUs, three CPUs, and five CPUs. When all the three **ksysrppmgr** commands are run, the number of available CPUs cannot be predicted. However, the number of available CPUs will be between the lowest and the highest request, that is, in the range 2 - 5. Therefore, you might avail 2, 3, 4, or 5 CPUs depending on the priority of the threads indicated by the system.

If you are using the Elastic (On/Off) CoD facility, the conflicts are handled based on resource priority and therefore, the method is thread-safe. For example, a conflict of request during a resource acquisition is retried until success and a conflict of request during a resource release is dropped immediately. In this case, the amount of resource that is released might not be as expected, but this change is not treated as an error.

You must specify the correct Enterprise Pool to decrease the topology exploration cost of Enterprise Pools. If the Enterprise Pool is not specified in the **ksysrppmgr** command, the **ksysrppmgr** command checks all the Enterprise Pools that are identified by the HMCs, which are specified in the command, to find the correct pool for the specified managed hosts. If one or more Enterprise Pools are specified in the command, only those Enterprise Pools are checked to get the required information. Therefore, the time spent in topology exploration is reduced. However, if you are not sure of which Enterprise pools can be used, you can run the command to request the resource without specifying any Enterprise Pools.

When the **ksysrppmgr** command ends successfully, and if no other command was run for the same managed host with different resource levels, the requested resources are made available to the specified host. If you specify the **-r** flag in the command, you can open the log file to check whether the requested amount of resources match the existing amount of available resources at the end of command execution.

By default, the log files for the **ksysrppmgr** command are located in the `/var/ksys/log/capmgr.log` file. You can configure the log file settings to overwrite the existing logs or delete the old log files after the log files reach a specific limit. You can also set a log file prefix to analyze the log files if more than one execution information is available in the log files.

When you use the **-r** flag to check the topology and resource status, you can filter the log files by using the following command:

```
cat logfile | grep -e "##" -e "@@" | cut -d " " -f 2-
```

## Flags

**-o *action***
Specifies the action to be performed. The following action flags are available:

**execute (alias: e)**
Runs the resource requests.

**check (alias: c)**
Simulates whether the resource requests can be completed.

**-h [*hmc_name*]:*hmc_uri*:*hmc_user***
Specifies the HMC details.

The Unique Resource Identifier (URI) and the user name of the HMC are mandatory parameters. The *hmc_name* parameter is optional and is used only if the HMC REST API page about HMC information returns information about other HMCs too, which might not occur.

All colons (:) in this flag are required, even if the optional parameters are not filled. This flag can be specified multiple times to add multiple HMCs.

**-m *ms_name*:*ms_action*:[onoff_use(n|y[*nb_days*])]:[*mem_amount*]:[*proc_amount*]**
Specifies the name of the managed host that must be configured for CoD resources.

The host name and action parameters are mandatory. Currently, the only allowed value for the *ms_action* parameter is **s** or **set**. This parameter matches the amount of requested resource to the amount of available resource. The unit to specify the memory resource is MB.

All colons (:) in this flag are required, even if the optional parameters are not filled. This flag can be specified multiple times to add multiple hosts.

**-M *ms_uuid*:*ms_action*:[onoff_use(n|y[*nb_days*])]:[*mem_amount*]:[*proc_amount*]**
Specifies the universally unique identifier (UUID) of the managed host that must be configured for CoD resources.

All other parameters are similar to the parameters of the **-m** flag.

**-e *enterprisepool_name***
Specifies the name of the Enterprise Pool that must be monitored.

If the **-e** or **-E** flags are not specified, all the enterprise pools are monitored by default. This flag can be specified multiple times.

**-E *enterprisepool_uuid***
Specifies the UUID of the Enterprise Pool that must be monitored.

If the **-e** or **-E** flags are not specified, all the enterprise pools are monitored by default. This flag can be specified multiple times.

**-v** Specifies that you want to display detailed processing information on your screen.

This flag displays the information about the managed host and return codes, one per line, at the end of the execution.

**-r** Specifies that you want detailed report.

This flag enables logging operation that includes the overall status of topology and resources before and after execution. It also includes a timing analysis on the overall execution.

If you specify this option, more REST operations might be triggered, which require more time to complete.

**-l** Specifies the log file that must be used.

By default, the `/var/ksys/log/capmgr.log` file is used for the logging operation. You can use the `none` value to remove all the logging operations. The `libhmc` log files are stored in the *`LOGFILE`*`.librpp_last_rest.log` file or in the *`LOGFILE.LOGPREFIX`*`.librpp_last_rest.log` file if the log prefix is specified. The `libhmc` log files overwrite previous log files, if the same prefix or no prefix was used.

**-p** Specifies the log prefix that must be used in the log files.

This flag specifies the prefix string that is determined at the beginning of each log string, along with other prefix information. This flag can also be used to avoid overwriting the `libhmc` log files.

## Return values

Because you can run the **ksysrppmgr** command on multiple hosts simultaneously, you cannot get a complete error report in the return value. For a more detailed error report for each managed host, you must run the **ksysrppmgr** command with the **-v** flag. The following error codes can be returned:

**0** The operation was successful.

**1** An error occurred.

**2** Indicates command line parsing error

### Error codes for verbose output

**0** The operation was successful.

**1** Conflict with other operations. The operation will be attempted again.

**2** Parsable error. This error might be associated to HSCL, or REST. Check the log files for more details.

**3** HSCL error. Check the report or log file for HSCL code and more information.

**4** REST error. Check the report or log file for REST code and more information.

**5** `libhmc` error. Check the report or log file for `libhmc` code and more information.

**6** Parameter error. Check log files for more information.

**7** Invalid output. Check log files for more information.

**8** Check log files for more information.

**9** Operation cannot be performed. Computation for resource pool provisions (RPP) determined that the operation will fail.

**10** Internal RPP error. Retry the operation. If problem persists, check log files.

## Examples

1. To check whether 3 CPUs can be available on the `host_1` host and whether 2 GB memory can be available on the `host_2` host while restricting the use of Elastic (On/Off) CoD, enter the following command:

```
# ksysrppmgr -o c -h :hmc1:hmcuser -h :hmc2:hmcuser -m host_1:s:n::3 -m host_2:s:n:2048: -r
```

In execution mode, if the return code of the command is 0, all the requests are successful. If the return code is 1, at least one request has failed. To view detailed information about the failure, you can either use the **-v** option for detailed printed results, or check the log file by using the filter command.

2. To request 3 CPUs from the `host_1` host and 2 GB memory from the `host_2` host, to restrict the Elastic (On/Off) CoD usage, and to print detailed information, enter the following command:

   ```
   # ksysrppmgr -o e -h :hmc1:hmcuser -h :hmc2:hmcuser -m host_1:s:n::3 -m host_2:s:n:2048: -r -v
   ```

   The return value of 0 indicates that all operations are successful, and the resource requests are satisfied. The **-v** flag prints a more detailed result that contains the final status for each managed host:

   ```
   host_1:0
   host
   ```

3. To request 2.5 CPUs and 10500 MB memory for the `host_1` host and to allow Elastic (On/Off) CoD usage for 5 days, enter the following command:

   ```
   # ksysrppmgr -o e -h :hmc1:hmcuser -h :hmc2:hmcuser -m host_1:s:y5:10500:2.5 -r
   ```

4. To release all possible resources from the `host_1` host to the pool and to optimize costs by converting Elastic (On/Off) CoD into Enterprise Pool, enter the following command:

   ```
   # ksysrppmgr -o e -h :hmc1:hmcuser -m host_1:s:y5:0:0
   ```

   If you want to convert Elastic (On/Off) CoD facility to Enterprise Pool CoD without modifying the current level of available resources, set the currently available resources as the target levels.

5. To specify multiple Power Enterprise Pools to be monitored by the **ksysrppmgr** command, enter the following command:

   ```
   # ksysrppmgr -o e -h :hmc1:hmcuser -m host_1:s:n:0:0 -m host_2:s:n:0:0 -e EP_1 -e EP_2
   ```

# Frequently asked questions (FAQs)

If you have questions about the GDR solution, review the following list of answers to some frequently asked questions.

**How many sites can be configured by using the GDR solution?**
Currently, only two sites can be configured. An active site and a backup site.

**Which POWER servers can be included for disaster recovery by using the GDR solution?**
POWER7 processor-based server, or later.

**Which type of storage devices are supported by the GDR solution?**

- EMC Symmetrix Remote Data Facility (SRDF) capable systems: Both SRDF/A (Asynchronous) and SRDF/S (Synchronous) replication modes are supported.
- SAN Volume Controller (SVC) system and Storwize system: Both the Metro Mirror (synchronous) and the Global Mirror (asynchronous) modes of data replication are supported.
- DS8000 storage system: Only Global Mirror (asynchronous) mode of data replication is supported.
- Hitachi storage systems: Only asynchronous mode of storage replication is supported.

**Note:** The SVC, Storwize, and DS8000 storage systems are supported only with VIOS Version 2.2.5.20, or later and HMC Version 8 Release 8.6.0 Service Pack 1, or later. The Hitachi storage systems are supported with VIOS Version 2.2.6.00, or later and HMC Version 9 Release 9.1.0, or later.

**Does this solution support Shared Storage Pool (SSP)?**
Yes. The GDR solution also supports shared storage configuration. For more information, see “Managing the shared storage configuration” on page 87.

**Does the GDR solution support any flexible capacity management solutions?**
Yes. The GDR solution supports Enterprise Pool Capacity on Demand (CoD) and Elastic CoD (formerly known as On/Off CoD) features. The GDR solution manages the memory and processor requests before starting the disaster recovery operation. For more information, see “Managing the CoD resources” on page 92.

The GDR solution also supports flexible capacity policies by using the KSYS command interface. For more information, see “Configuring the flexible capacity policies” on page 74.

**Which operation systems are supported for the virtual machines in the hosts?**

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

See “Requirements for the GDR solution” on page 27 for more information.

**Can the GDR solution work in the presence of other high availability and disaster recovery solutions in the current environment?**
Yes, if the cluster management of the virtual machines that are covered under other high availability and disaster recovery solutions are handled separately and outside of the KSYS subsystem. A specific virtual machine can be included in only one high availability and disaster recovery solution.

**What is the minimum version of HMC that must be used for this solution?**
See “Requirements for the GDR solution” on page 27.

**What is the minimum version of VIOS that must be used for this solution?**
See “Requirements for the GDR solution” on page 27.

© Copyright IBM Corp. 2017, 2019

**What is KSYS?**

The KSYS (also known as controlling system) is a controlling software for the disaster recovery operation that operates in an AIX Version 7.2.1 (or later) logical partition in the backup site and controls the entire cloud environment for the GDR solution. The KSYS is responsible for carrying out recovery actions if a disaster or a potential disaster occurs.

**Can I test the disaster recovery operation without an actual disaster?**

Yes. The GDR solution performs failover rehearsal at the backup site in the disaster recovery environment, without disrupting the production workloads or the storage replication from the active site to the backup site. For more information, see “Failover rehearsal of the disaster recovery operation” on page 61.

# 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:

© Copyright IBM Corp. 2017, 2019

*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.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

# Index

© Copyright IBM Corp. 2017, 2019

Printed in USA