January 2021

Dell EMC SRDF storage agent configuration for IBM VM Recovery Manager

Setting up communication between EMC storage agent and KSYSNODE using SYMCLI software

Table of Contents

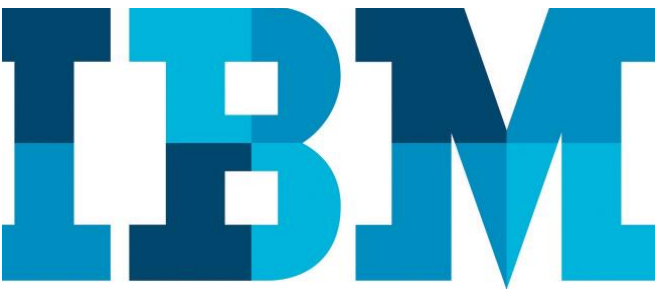

Overview

Challenge
Establishing communication between EMC SRDF storage and KSYS system requires a specific setting using SYMCLI software. Users are not aware of configuring the same for IBM VM Recovery Manager.

Solution
This paper provides detailed steps for configuring EMC SRDF storage agent using SYMCLI software and configuring KSYSNODE for communication to work with VM Recovery Manager

Introduction

Disaster recovery (DR) of applications and services is a key requirement to provide continuous business services. IBM® VM Recovery Manager DR for Power Systems is a disaster recovery solution that is easy to deploy and provides automated operations to recover the production site. This solution can support many storage replications. This paper explains the configuration process for a Dell EMC Symmetrix Remote Data Facility (SRDF) storage subsystem and storage controller to interact with a KSYS subsystem. This includes setting up a storage agent and monitoring disk replication steps by the storage agent created in the VM Recovery Manager solution.

Architecture overview: VM Recovery Manager solution and EMC storage agent

Disaster recovery and high availability (HA) solutions could be based on cluster-based technology or virtual machines (VMs) restart-based technology. Cluster HA and DR solutions typically deploy redundant hardware and software components to provide near real-time failover when one or more components in the configuration fails. The VM-based solution relies on the out-of-band monitoring and management option to restart the virtual machines in the event of hardware failure in an infrastructure. The VM Recovery Manager solution is based on the VM restart technology.

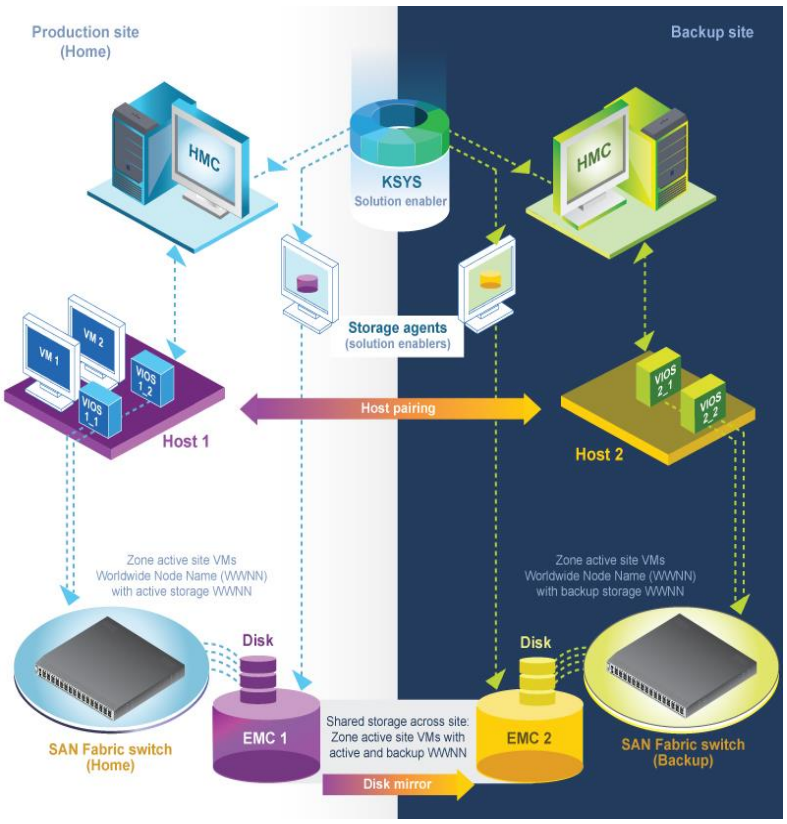

Figure 1. Generic setup for VM Recovery Manager DR

Figure 1 shows the basic setup for a solution to be implemented with many subsystems. The VM Recovery Manager solution provides a highly available environment by identifying a set of resources that are required for processing the VMs in a server during disaster situations.

The solution uses the following subsystems:

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

The solution provides support for the following storages subsystems (and more) and their respective peer-to-peer replication technologies.

- IBM System Storage SAN Volume Controller (SVC) / IBM Storwize®: Metro Mirror and Global Mirror
- Hitachi: True copy and universal copy
- EMC SRDF: Sync and async replication
- IBM System Storage DS8000®: Asynchronous replication
- IBM XIV storage system

Note: For this solution, the administrator must perform the setup manually at the storage level.

This paper describes how to configure the EMC storage agent, create SRDF group, disks, and replicate them in EMC. This paper also provides detailed steps to add a storage agent controller to a KSYS cluster.

EMC storage agent configuration

For configuring the EMC storage agent, the admin has to use two IBM AIX® installed partitions apart from the KSYS node. Two AIX installed partitions will be installed with solution enabler software provided by EMC. These storage agents configured will act as a medium for communication to the actual storage boxes. **An important requirement that each storage agent (AIX partition) should have one disk of size 1 GB from each storage (source and remote) to act as the gatekeeper device.**

Steps to create a storage agent

Perform the following steps on the storage agents (AIX installed LPARs) apart from the KSYS node:

1. Use two AIX installed partitions apart from the KSYS node. The partitions should be with a minimum of 0.1 CPU and 2 GB memory.
2. Install Dell EMC Solutions Enabler software on both the storage agents acting as the AIX partition. Command to install Solutions Enabler: se8407_install.sh -install -silent
   Note: In this white paper, se8407 is version of the Symmetric Command Line Interface (SYMCLI) software is mentioned. If the administrator has a later version, it will display that version. It will be a common file with install.sh along with solution enabler software.
3. Set SYMAPI_ALLOW_RDF_SYMFORCE = TRUE and SYMAPI_USE_RDFD = ENABLE in the /var/symapi/config/options file on both the storage agents.
4. Assign a minimum of 1 GB volume to each storage agent from storages boxes through Fibre Channel (FC) mapping which will be used as the gatekeeper device for communicating with the actual storage boxes
5. Run the /usr/symcli/bin/symcfg discover and /usr/symcli/bin/symcfg list command on both the storage agents. This provides details about the source storage and the remote storage with serial number for both the storage agents.
   If handshake error is seen while listing the configuration, update the **SYMAPI_SECURITY_LEVEL = SECURE** entry to **SYMAPI_SECURITY_LEVEL = ANY** in the /var/symapi/config/options file on both the storage agent nodes.

The following figure shows that ksys214 is the storage agent connected to the VMAX 508 storage.

```
(0) root @ ksys214: /usr/symcli/bin
# symcfg list

                                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      215040         1     14763
    000196800573 Remote      VMAX100K             5977      215040         0     13452
    000198701861 Remote      VMAX10K              5876       59392         0       286
```

The following figure shows that ksys215 is the storage agent connected to the VMAX 573 storage.

```
(0) root @ ksys215: /usr/symcli/bin
# symcfg list

                                S Y M M E T R I X

                                                  Mcode    Cache      Num Phys  Num Symm
    SymmID       Attachment  Model                Version  Size (MB)  Devices   Devices

    000196800573 Local       VMAX100K             5977      215040         1     13452
    000196800508 Remote      VMAX100K             5977      215040         0     14763
    000198701861 Remote      VMAX10K              5876       59392         0       286
```

Note: In this example, ksys214 is the source storage agent connected to VMAX 508 and ksys215 is target storage agent connected to VMAX 573. VMAX 508 and VMAX 573 are the actual storage boxes.

6. Make sure that the **storapid** and **storsrvd** daemons are running on both the storage agents.
   You can use the following command to list the daemons running on a storage agent.

```
# /usr/symapi/bin/stordaemon list

Available Daemons ('[*]': Currently Running):

[*] storapid        EMC Solutions Enabler Base Daemon
[*] storgnsd        EMC Solutions Enabler GNS Daemon
[*] storrdfd        EMC Solutions Enabler RDF Daemon
    storevntd       EMC Solutions Enabler Event Daemon
[*] storwatchd      EMC Solutions Enabler Watchdog Daemon
    storsrmd        EMC Solutions Enabler SRM Daemon
    storstpd        EMC Solutions Enabler STP Daemon
[*] storsrvd        EMC Solutions Enabler SYMAPI Server Daemon
```

Use the following command to start the **storsrvd** and **storapid** daemons in case they are not running.

```
/usr/symcli/bin/stordaemon start storsrvd
/usr/symcli/bin/stordaemon start storapid
```

Configuring KSYS node to communicate with EMC storage agent

Communication is required from the KSYS node to both the storage agents for VM Recovery Manager DR and HADR operations. Also, the Solutions Enabler software must be installed on the KSYS node.

Perform the following steps to configure the KSYS node to communicate with the EMC storage agents.

1. Install Solutions Enabler SYMCLI on the KSYS node in the same way as installed on the storage agent nodes.
2. Update the **netcnfg** file in the KSYS node to establish a secure communication between the KSYS node and the storage agents. This file will be generated after installing Solutions Enabler software. Make the following updates in the file with the storage agent host name and IP address:

   Entries to be updated in the /var/symapi/config/netcnfg file:

   SYMAPI_SITE_508 - TCPIP ksys214 10.40.24.156 2707 ANY
   SYMAPI_SITE_573 - TCPIP ksys215 10.40.24.157 2707 ANY

   Note: In this paper, the VMAX 508 and VMAX 573 models are used as the source and the target storages. VMAX 508 is the source storage (which has the storage agent created on the ksys214 AIX partition) and VMAX 573 is the target storage (which has storage agent created on ksys215 AIX partition).

3. Export the variable SYMCLI_CONNECT=SYMAPI_SECURE.
4. Make sure that the **storapid** and **storsrvd** daemons are running on KSYSNODE.
   Command to list the daemon running.
   # /usr/symapi/bin/stordaemon list
   If the daemons are not running, start the daemon using the following command:
   stordaemon start <daemon_name>
5. Run the /usr/symcli/bin/symcfg discover command to discover and connect with the storage agents
6. Set SYMAPI_ALLOW_RDF_SYMFORCE = TRUE and SYMAPI_USE_RDFD = ENABLE in the /var/symapi/config/options file on both the storage agents as well as KSYSNODE.
7. Export the variables to connect to the VMAX 508 and VMAX 573 storages remotely from the KSYS node through the storage agent.
   Export the following variables to connect remotely to the storage agents.export SERVICE_CONNECT_TYPE=REMOTE
   export SYMCLI_CONNECT=<label>

   For example:
   export SYMCLI_CONNECT=SYMAPI_SITE_573
   export SYMCLI_CONNECT=SYMAPI_SITE_508
8. Run the symcfg list command on the KSYS node to verify if the communication is established properly with both the storage agents. The following figure shows that ksys216 is the KSYS node and VMAX 508 is the local storage and VMAX 573 is the remote storage.

```
(0) root @ ksys216: /usr/symcli/bin
# symcfg list

                                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       215040         1     14763
    000196800573    Remote      VMAX100K          5977       215040         0     13452
    000198701861    Remote      VMAX10K           5876        59392         0       286
```

SRDF Disk replication for EMC storage

VM Recovery Manager solution supports EMC 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 disk of the composite group, the EMC storage device responds to the host and the target EMC storage device acknowledges that it has received and checked the data. Replication of data is synchronous and the data is copied in real time.

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). Users can change this default cycle time to change the time difference of dependent write operations on the secondary storage devices that suits business requirement.

Tasks such as disk allocation, host creation, and disk replication can be done using the symmetric user interface, graphical interface, or the command line interface.

Note: This paper shows devices creation and SRDF replication using command line.

1. **Create a host group, add initiator, and query for details.**

   Use the following commands to create a host group and to add a worldwide name (WWN) to the host group:

   ```
   symaccess -sid <storage_id> -name <host_group_name> -type initiator create
   symaccess -sid <storage_id> -name <host_group_name> -type initiator -wwn <WWN_of_parition> add
   ```

   The following example shows the creation of a host group for the gdrnova105 VM with WWN.

   ```
   symaccess -sid 000196800508 -name gdrnova105_HG -type initiator create
   symaccess -sid 000196800508 -name gdrnova105_HG -type initiator -wwn C0507606D6DF002A add
   ```

   Note: You need to perform similar steps on the target storage agent with a target storage ID.

   Use the following command to query the host group you created:

   ```
   symaccess -sid <storage_id> -type initiator -detail show <host_group_name>
   ```

   The following screen capture shows the query for the host group created for the VMAX 508 source storage.

```
(0) root @ ksys214: /usr/symcli/bin
# symaccess -sid 000196800508 -type initiator -detail show gdrnova105_HG

Symmetrix ID          : 000196800508

Initiator Group Name  : gdrnova105_HG
Last update time      : 01:07:03 AM on Fri Nov 09,2018

Group last update time: 01:07:03 AM on Fri Nov 09,2018

Port Flag Overrides   : No
Consistent Lun        : No

   Originator Port wwn  : c0507606d6df002a
   User-generated Name  : c0507606d6df002a/c0507606d6df002a
   FCID Lockdown        : No
   Heterogeneous Host   : No
   Port Flag Overrides  : No
   CHAP Enabled         : N/A
   Type                 : Fibre

   Originator Port wwn  : c0507606d6df002b
   User-generated Name  : c0507606d6df002b/c0507606d6df002b
   FCID Lockdown        : No
   Heterogeneous Host   : No
   Port Flag Overrides  : No
   CHAP Enabled         : N/A
   Type                 : Fibre
```

2. **Create a port group and query the details of the port group.**

Assigning devices (disks) to a host group and mapping to the virtual machines is done using ports. EMC has specific ports connected to SAN. User needs to get the details of the EMC ports with which Central Electronics Complex (CEC) is connected. While creating a port group, you need to give a specific EMC port number for allocating disks to a virtual machine. Use the following command to create a port group and to enable it.

symaccess -sid <storage_id> create -name <portgroup_name> -type port -dirport <dir_port>
set port <dir_port> ACLX=enable

Example (with the gdrnova105 VM):

symaccess -sid 000196800508 create -name gdrnova105_PG -type port -dirport 01D:009,02D:009
set port 01D:009 ACLX=enable

Note: You need to perform similar steps on the target storage agent with the target storage ID.

Use the following command to see the details of the port group you just created.

symaccess -sid <storage_id> -type port show <portgroup_name>

The following screen capture shows details about the port group query you created.

```
(0) root @ ksys214: /usr/symcli/bin
# symaccess -sid 000196800508 -type port show gdrnova105_PG

Symmetrix ID          : 000196800508

Port Group Name         : gdrnova105_PG
Last update time        : 01:07:03 AM on Fri Nov 09,2018

   Director Identification
     {
        Director
      Ident  Port   WWN Port Name / iSCSI Target Name
      ------ ---- ------------------------------------------------------
      FA-1D   009 500009735807f009
      FA-2D   009 500009735807f049
     }

   Masking View Names
     {
       gdrnova105_MV
     }
```

3. **Create a device (disk) and a storage group.**

You need to create a disk with the required capacity and then add that disk to a storage group. After creating the disk, you need to save the disk ID in a text file for SRDF replication.

Use the following command to create the device (disk):

symconfigure -sid <storage_id> -cmd "create dev count=1, size=30720 , emulation=FBA, config=TDEV ;" commit

VM Recovery Manager ignores the gatekeeper disk and the EMC Access Control Logix (ACLX) enabled disk though replicated. So, admin should not add the gatekeeper disk or the ACLX-enabled disk into remote data facility group (rdfg). EMC ACLX disks are special devices that act as the initial gatekeeper disks for managing actual storage environment through storage agents. These devices provide the control path for running the commands from the storage agent and to pass them on to the actual storage boxes. These disks are not data logical unit numbers (LUNs), and therefore, replication is not required for these disks. As ACLX disks are typically mapped to all Fibre Channel front-end adapter (FA) ports initially, all hosts attached can see these disks. So, it is recommended to unmap the ACLX disk.

Refer to the following sample command to check the ACLX disk present in the storage array.

```
# symdev list -aclx
Symmetrix ID: 000196800573

Device Name          Dir            Device
-------------------------  -------  ------------------------------------

                                       Cap
Sym   Physical         SA :P        Config      Attribute   Sts(MB)
-------------------------  -------  ------------------------------------
00001   Not Visible      01D:007 TDEV    N/Grp'd         ACLX WD      6
```

Run the following command to create the storage group.

symaccess -sid <storage_id> -name <storagegroup_name> -type storage create devs <diskid>

Note: Disk should be created on both storages. In this white paper, you can find generic commands that are run on the source storage agent. Similarly, users need to perform disk creation and storage group creation on the target storage agent using the target storage ID.

Refer to the following example of disk creation and storage group with the gdrnova105 VM.

```
symconfigure -sid 000196800508 -cmd "create dev count=1, size=30720, emulation=FBA, config=TDEV;" commit
symaccess -sid 000196800508 -name gdrnova105_SG -type storage create devs 0392A
```

output of disk creation:

```
Execute a symconfigure operation for symmetrix '000196800508' (y/[n]) ? y

    A Configuration Change operation is in progress. Please wait...

    Establishing a configuration change session.....Established.
    Processing symmetrix 000196800508
    Performing Access checks.....Allowed.
    Checking Device Reservations.....Allowed.
    Initiating COMMIT of configuration changes.....Queued.
    COMMIT requesting required resources.....Obtained.
    Step 006 of 011 steps.....Executing.
    Local: COMMIT.....Done.

    New symdev: 0392A [TDEV]
    Terminating the configuration change session.....Done.

    The configuration change session has successfully completed.
```

You can find the disk created using the symdev list command and storage group created using the symaccess command with attribute -type storage show as shown

```
symaccess -sid <storage_id> -type storage show <storagegroup_name>
```

The following screen capture shows the storage group created for gdrnova105 on the source storage agent. Masking view will be missing at first instance. This screenshot is captured after creating the masking view.

```
(0) root @ ksys214: /usr/symcli/bin
# symaccess -sid 000196800508 -type storage show gdrnova105_SG

Symmetrix ID                : 000196800508

Storage Group Name          : gdrnova105_SG
Last update time            : 01:16:56 PM on Wed Nov 13,2019
Group last update time      : 01:16:56 PM on Wed Nov 13,2019

   Number of Storage Groups : 0
   Storage Group Names      : None

   Devices                  : 0392A 0392D

   Masking View Names
     {
       gdrnova105_MV
     }
```

4. **Create a masking view on both the storage agents.**

EMC masking view also known as LUN masking, enables and assigns LUN to a host with one set of storage ports and initiators also known as WWN. To run the command to create a masking view, you need to know the storage group name, port group name, and the host group name created in previous steps.

Run the following command to create a masking view:

symaccess -sid <storage_id> create view -name <maskingview_name> -sg <storagegroup_name> -pg <portgroup_name> -ig <hostgroup_name>

Example of creating a masking view for the gdrnova105 VM.

symaccess -sid 000196800508 create view -name gdrnova105_MV -sg gdrnova105_SG -pg gdrnova105_PG -ig gdrnova105_HG

You can see detailed output of the masking view using the following command:

symaccess -sid <storage_id> show view <maskingview_name>

The output shows information of the host group with WWN, port group with director port, storage group with LUNs for a VM. The following screen capture shows the masking view query output for vm gdrnova105

Note: This needs to be done on both the storage agents.

```
(0) root @ ksys214: /usr/symcli/bin
# symaccess -sid  000196800508 show view gdrnova105_MV

Symmetrix ID                  : 000196800508

Masking View Name             : gdrnova105_MV
Last update time              : 01:07:03 AM on Fri Nov 09,2018
View last update time         : 01:16:56 PM on Wed Nov 13,2019

Initiator Group Name          : gdrnova105_HG

   Host Initiators
     {
       WWN   : c0507606d6df002a
               [alias: c0507606d6df002a/c0507606d6df002a]
       WWN   : c0507606d6df002b
               [alias: c0507606d6df002b/c0507606d6df002b]
     }

Port Group Name               : gdrnova105_PG

   Director Identification
     {
        Director
      Ident  Port    WWN Port Name / iSCSI Target Name
      ------ ----  ------------------------------------------------------
      FA-1D    009 5000097358O7f009
      FA-2D    009 5000097358O7f049
     }

Storage Group Name            : gdrnova105_SG

   Number of Storage Groups : 0
   Storage Group Names      : None

Sym                                          Host
Dev      Dir:Port  Physical Device Name      Lun   Attr  Cap(MB)
------   --------  ----------------------   ----  ----  -------
0392A    01D:009   Not Visible                0          57600
         02D:009   Not Visible                0
0392D    01D:009   Not Visible                1          57600
         02D:009   Not Visible                1
                                                       -------
Total Capacity                                          115200
```

5. **Create an SRDF group.**

   You need to create an SRDF group for establishing replication between the source storage and the target storage disks. To create an SRDF group, a unique group ID and a name are required. EMC VMAX has a maximum limit of 255 IDs. Before creating an SRDF group, you need to make sure that the group ID and the name are not used by any other existing groups on both the storages.

   Use the following command to list all the IDs used in storages. From this output you can find out which ID is unused. You need to run this command on both source and target storages.

   ```
   /usr/symcli/bin/symcfg list -ra all -switched
   ```

   After identifying the unused SRDF group ID, create the SRDF group with the source and target director port, as well as the source and target storage ID.

   Command to create an SRDF group.
   Note: You must run the SRDF group creation command only on the source storage.

   ```
   /usr/symcli/bin/symrdf addgrp -label <rdfg_group_name> -rdfg <rdfg_id> -sid <source_storage_id> -dir <source_dir_port> -remote_rdfg <remote_rdfg_port> -remote_sid <remote_storage_id> -remote_dir <remote_dir_port>
   ```

   Example of creating the SRDF group with the ID as **222** and the name, **test_group**.

   ```
   0) root @ ksys214: /usr/symcli/bin
   # /usr/symcli/bin/symrdf addgrp -label test_group -rdfg 222 -sid 000196800508 -dir 01E:27 -remote_rdfg 222 -remote_sid 000196800573 -remote_dir 01E:27

   Execute a Dynamic RDF Addgrp operation for group
   'AUT_A' on Symm: 000196800508 (y/[n]) ? y

   Successfully Added Dynamic RDF Group 'test_group' for Symm: 000196800508
   ```

   After creating the group, verify it using the `/usr/symcli/bin/symcfg list -ra all -switched` command.

6. **Create SRDF disk replication.**

   You need to create SRDF disk replication for data mirroring between the source disk and the target disk. IDs of the source and the target storage disks is required to create replication. Create text file as shown in the following example:

   Note: Command to create disk replication needs to be run on the source storage agent only.

   In the following example, a file named **disk_id_file** is created.

   ```
   (0) root @ ksys214: /usr/symcli/bin
   # cat disk_id_file
   0392A 002BE
   ```

   0392A is the disk ID of the source storage and 002BE is the disk ID of the target storage.

   Command to create synchronous type replication.

   ```
   symrdf -file <disk_id_file> createpair -type R1 -sid <source_storage_id> -rdfg <rdfg_id> -establish -rdf_mode sync -nop -force
   ```

Command to create asynchronous type replication

```
symrdf -file <disk_id_file>  createpair -type R1 -sid  <source_storage_id> -rdfg <rdfg_id> -establish -rdf_mode async -nop -force
```

Example of creating sync replication.

```
(0) root @ ksys214: /usr/symcli/bin
#  symrdf -file disk_id_file createpair -type R1 -sid 000196800508 -rdfg 222 -establish -rdf_mode sync -nop -force

An RDF 'Create Pair' operation execution is in progress for device
file 'disk_id_file'. Please wait...

   Create RDF Pair in (0508,222).....Started.
   Create RDF Pair in (0508,222).....Done.
   Mark target device(s) in (0508,222) for full copy from source....Started.
   Devices: 0392A-002BE in (0508,222).....Marked.
   Mark target device(s) in (0508,222) for full copy from source....Done.
   Merge track tables between source and target in (0508,222).....Started.
   Devices: 0392A-002BE in (0508,222).....Merged.
   Merge track tables between source and target in (0508,222).....Done.
   Resume RDF link(s) for device(s) in (0508,222).....Started.
   Resume RDF link(s) for device(s) in (0508,222).....Done.

The RDF 'Create Pair' operation successfully executed for device
file 'disk_id_file'.
```

Disk replication created can be verified using the symrdf list command.

Command to see disk replication: "symrdf list | grep "R1:<rdfg_id"

Example of disk replication for SRDF group with ID 222:

```
(0) root @ ksys214: /usr/symcli/bin

# symrdf list | grep "R1:222"

0392A 002BE   R1:222 RW RW  RW  S1.E       0       0 RW  WD   Synchronized
```

Ensure that in the symrdf list command output, the state of replication is *Synchronized*. If the state is other than *Synchronized*, the KSYS subsystem will display an error message during discovery operation.

Configuration for DR rehearsal

EMC SRDF has a symclone feature which creates the mirror copy of an existing target disk on the target storage. This extra copy will be used by VM Recovery Manager to activate the VM on the target site. EMC Symmetrix symclone generally supports three operations: Create, activate, and terminate. Create operations will create a clone of the target disk using the symclone create command. Activate operation will activate the clone using the symclone activate command. Size of the target disk and the extra disk used for cloning should be the same. Configuration of clone copy is required only on the target storage for DR rehearsal.

Create a text file with the disk ID of the target disk and the clone disk for creating and activating the clone. File name can be any user-defined name. In this paper, clone_disk_id is the file name used for reference.

Command for symclone create operation

```
/usr/symcli/bin/symclone -sid <target_storage_id> -f <clone_disk_id> create -diff -nop -force
```

Command to check the status of mirror between target disk and extra clone disk

```
symclone -sid 573 list | grep -E "disk_ids"
```

The following example shows the status of mirroring between the target disk and the clone disk using the symclone list command:

```
# symclone -sid 573 list | grep -E "0307E|02F8D|03075|03078|0307C|0307D"  ====> disk names
02EC8          409605 02F8D          X.X. Created
02EC9          409605 03075          X.X. Created
02ECB          409605 03078          X.X. Created
02ECF          409605 0307C          X.X. Created
02EC0          409605 0307D          X.X. Created
02F6C          409605 0307E          X.X. Created
```

After the clone is created, the status of mirroring will be shown in the symclone list output. You need to activate mirroring between the clone and the target disk using the symclone activate command and wait until the status shows *Copied*. If the mirroring state is not *Copied,* then the verify operation for DR test in VM Recovery Manager will fail.

Command for symclone activate operation.

/usr/symcli/bin/symclone -sid <target_storage_id> -f <clone_disk_id> activate -nop -force

Adding storage agents in KSYS cluster

Storage agents needs to be added to the KSYS cluster for collecting disk and replication information from the actual storages. Adding EMC storage agents will be performed using entries added in the netcnfg file in each storage agent. This will require storage agent hostname or IP, storage serial number, and site name created in the KSYS cluster.

Command to add a storage agent in the KSYS node:

Ksysmgr add storage_agent <sa_name> login=<storage_agent_user> site=<site_name> serialnumber=<storage_agent_id> storagetype=EMC hostname=<storage_Agent_hostnameorip>

While adding a storage agent, if there is any error or failure, check all the setup steps properly as mentioned in the earlier sections. Make sure that both the storage agents daemons are active and in healthy state to communicate with the actual storage boxes.

After adding the storage agent to the KSYS node, perform the KSYS discovery operation to get the disk replication details. In VM Recovery Manager, the storage agents that are added to the KSYS configuration interact with the storage controllers in each site. If replication is of the *synchronous* type at the storage level, update the KSYS system replication type to *sync* using the ksysmgr modify system command.

Adding and removing disks from VM

During KSYS discovery operations, a composite group is created in EMC storage agents. The composite group often goes into error state while removing or adding disks in VM. If admin wants to remove or add a new disk to the VM managed by VM Recovery Manager, admin needs to perform specific steps in the storage agents to ensure that the composite group is in error-free state in the subsequent KSYS operations.

Steps to follow if a new disk is added to VM:

1. Map the disk to the VM.
2. Add the disk to the SRDF group in the storage agent.
3. Run the discovery operation at the KSYS node.

Steps to follow if disk is removed from VM:

1. Unmap the disk from the VM.
2. Remove the LUN /disk from the composite group.
3. Remove the LUN / disk from the SRDF group.
4. Run the discovery operation at the KSYS node.

To unmanage a VM which is part of work group after removing a LUN, perform the following steps:

1. Remove the VM from the workgroup using the ksysmgr modify workgroup command. This will directly remove the LUNs related to that VM from the composite group without any error.
2. Unmanage the VM using the ksysmgr unmanage vm command.
3. Remove the LUN / disk from the SRDF group in the storage agent.
4. Run the discovery operation at the KSYS node.

Conclusion

This white paper describes the procedure to configure EMC storage agents and establish communication between the storages and the KSYS node using EMC solution enabler software. It also describes the steps for creating a device in the storage agents and creating SRDF replication between the source storage and the target storage disks.

About the authors

Dishant Doriwala is a DR component test lead in the VM Recovery Manager product team. He has more than 7 years of experience working with the IBM Power platform including IBM PowerHA® SystemMirror® and VM Recovery Manager. You can reach Dishant at dishantdoriwala@in.ibm.com.

Srikanth Thanneeru is an advisory software engineer and is currently working as the test lead in the IBM VM Recovery Manager product team. Srikanth has around 10 years of experience with IBM AIX operating system functional testing. His areas of expertise include file systems, kernel, shared storage pools, high availability, and disaster recovery solution. You can reach Srikanth at sreekkanth@in.ibm.com.

Adhish Kapoor is an advisory software engineer and is currently working as the development lead in the IBM VM Recovery Manager product team. Adhish has expertise in storage and clustering technologies. His experience included development and deployment of various complex HA/DR solutions, SAN/NAS products. You can reach him at Adhish.Kapoor@in.ibm.com.

Neha Jain does functional verification testing in the VM Recovery Manager product team. She has more than 3 years of experience in the IBM Power platform. She has knowledge on disaster recovery and high availability, and has expertise with IBM i and IBM System Storage™ DS8000® storage. You can reach Neha at nehajain29@in.ibm.com.

IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America