

EMC Metro Smart DR Configuration & Integration with PowerHA

Table of Contents

Introduction:	1
Architecture Overview:	2
Use Case:	2
Prerequisites and Dependencies:	2
SRDF Storage Replication Configuration Steps for EMC Smart DR:	3
Integrating EMC SmartDR with PowerHA SystemMirror:	9
Conclusion:	12
Reference Links:	13
About the author:	13

Introduction:

PowerHA SystemMirror is a high availability clustering product which ensure applications are highly available and critical data is available in case of disaster scenarios.

PowerHA SystemMirror supports storages such as EMC, DS8K, XIV, SVC and Hitachi for data replication. Data is replicated to different locations using respective storage replication technology. PowerHA SystemMirror interacts to the storage replication component to allow workload IOs in the desired site. PowerHA SystemMirror retrieves the primary-secondary disk states from the replicated storages and initiates recovery as needed during failover and failback operations.

PowerHA supports EMC SRDF Synchronous (Sync) and Asynchronous (Async) replication. Sync replication is primarily used for short distance replication and Async replication is used for far distance replication.

EMC enabled new solution with PowerMAX storages called EMC Metro SmartDR, which is a Two-region High available (HA) Disaster recovery (DR) solution. It integrates SRDF/Metro (Metro) and SRDF/Async (SRDF/A) enabling HA DR for a metro session. This is achieved by closely coupling the SRDF/A sessions on each side of a Metro-pair to replicate to a single DR device. Also, it ensures that only a single SRDF/A session will be sending data to the DR site.

EMC SRDF/Metro SmartDR is supported only in PowerHA SystemMirror version 7.2.5 with Service Pack 1, or later.

Architecture Overview:

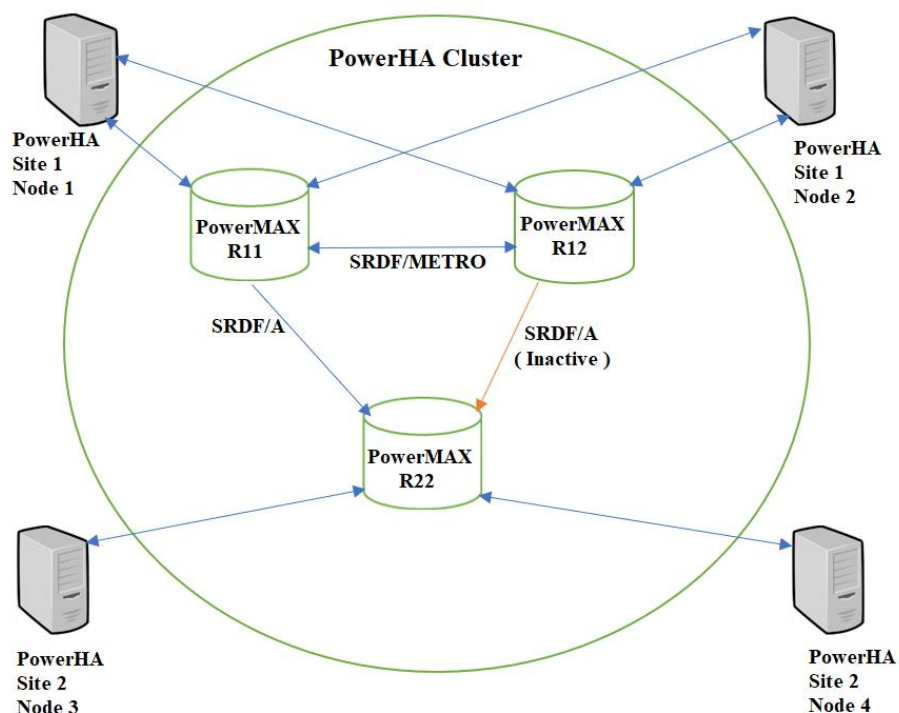


Figure 1. EMC Smart Metro DR integration with PowerHA

Use Case:

As shown in the Figure 1, Metro DR solution helps replicate IOs Synchronously between R11 and R12 and Asynchronously with R22. This model helps to continue IOs without disruption, when R11 or R12 fails and continues with other available storage. When both R11 and R12 are failed (In case of site1 storage failure), R22 become primary and continue with the workloads.

PowerHA is integrated with EMC Metro SmartDR solution to handle auto failover between Metro (R11/R12) to R22 and also does resynchronization back from Site2 (R22) to Site1 Metro (R11/R12), after failover.

Prerequisites and Dependencies:

Before going to EMC Metro SmartDR storage replication, you must ensure the following prerequisites and dependencies are met.

- Three PowerMAX storages need to be configured.
- All three storages need to have Fiber Channel (FC) connectivity. As shown in Figure 1, FC connectivity should exist among R11, R12 and R22.
- Node1 and Node2 should be connected to both Site1 storages, R11 and R12.

- Node3 and Node4 should be connected to Site2 storage, R22.
- Active link between R11-R12 and R11-R22. Inactive link between R12-R22.
- All three Powermax storages must be running PowerMax OS 5978.669.669 or higher.
- All three Powermax storages must be discoverable through Solutions Enabler 9.2.
- PowerMax R11, R12, and associated DR R22 disks must be the same capacity.
- User must have an installation of Symmetrix Command-line Interface (SYMCLI) software for AIX version 9.2.0.0, or later.

SRDF Storage Replication Configuration Steps for EMC Smart DR:

1. Create new RDF groups in all three storages.
 - i. Initially verify the configured RDF'S in all three storages with their respective RDF group name and number to avoid duplication.
 - ii. Use the below command to verify already configured groups in the storage.
#symcfg list -ra all -switched
2. Verify available ports and direction
 - i. Use the below command to verify ports and direction from storage.
symsan -sanrdf -sid <2237> -dir all list
3. Verify the free disk availability for configuration.
 - i. Use the below command to verify the free disks availability coming from storage.
symdev list pd
4. Adding RDF between two storages (From R11 to R12)
 - i. Command to add RDF from storage
symrdf addgrp -label <GrpLable> -rdfg <GrpNum> -sid <SID> -dir <DIR.Port,DIR.Port> -remote_rdfg <GrpNum> -remote_sid <SID> -remote_dir <DIR.Port,DIR.Port>

- Output of sample RDF creation:

```
(0) root @ Node1: /usr/symcli/bin
# symrdf addgrp -label rdf_grp18 -rdfg 18 -sid 000197902237 -dir 1E:11
-remote_rdfg 18 -remote_sid 000197902238 -remote_dir 1E:06
```

```
Execute a Dynamic RDF Addgrp operation for group
'rdf_grp18' on Symm: 000197902237 (y/[n]) ? y
```

```
Successfully Added Dynamic RDF Group 'rdf_grp18' for Symm:
000197902237
```

5. Verify whether the RDF Group is created or not.

- i. Run below command from storage.

#symcfg list -ra all -switched

Note: rdf_grp18 RDF Group should be available in PowerMAX R11 and R12 Array only.

6. Create RDF disk Replication (between R11 and R12 Storage).

- i. To create RDF disk replication, create a file with disk PVIDs of both primary and target storage.
- Output of sample disk pair from R11 to R12:

(0) root @ Node1: /usr/symcli/bin

cat /pair1.txt

00044 0003F

Note: * To get the available disk list use symdev list pd command

* Save the disk pair file only on home node.

7. Unmap the disk for R12 Array.

8. Adding disk pair to the RDF Group.

- i. Run below command from R11 storage.

symrdf createpair -sid <SID> -file <disk_id_file> -rdfig <GrpNum> -type

r1 -establish -metro

- Output of sample adding disk pair from R11:

(0) root @ Node1: /opt/emc/SYMCLI/bin

symrdf createpair -sid 237 -file /pair1.txt -rdfig 18 -type r1 - establish - metro

Execute an RDF 'Create Pair' operation for device file

'/pair1.txt' (y/[n]) ? y

An RDF 'Create Pair' operation execution is in progress for device file '/pair1.txt'. Please wait...

Create RDF Pair in (2237,004).....Started.

Create RDF Pair in (2237,004).....Done.

Mark target device(s) in (2237,004) for full copy from source....Started.

Devices: 0030-0031 in (2237,004).....Marked.

Mark target device(s) in (2237,004) for full copy from source....Done.

Merge track tables between source and target in (2237,004).....Started.

Devices: 0030-0031 in (2237,004).....Merged.

Merge track tables between source and target in (2237,004).....Done.

Resume RDF link(s) for device(s) in (2237,004).....Started.
Resume RDF link(s) for device(s) in (2237,004).....Done.

The RDF 'Create Pair' operation successfully executed for device file '/pair1.txt'.

9. Perform discover to sync the changes with other storage.

i. Use the below command to run discover:

symcfg discover

- Output of sample discover:

(0) root @Node1: /opt/emc/SYMCLI/bin

symcfg discover

This operation may take up to a few minutes. Please be patient...

Note: * The disk pair should appear on the particular disk on the local storage.

* The disk pair will not appear on the remote storage side since the disk is in unmap state.

* Check pair on disk using symdev list pd command.

* Pairs are listed as RDF+TDEV.

10. Map the disk back to the storage R12.

i. Connect to storage R12 and bring the disk back and run discover and sync from storage R12.

ii. Command to run discover:

symcfg discover

iii. Command to run sync:

symcfg sync

11. Create RDF group between R11 storage and DR storage R22.

i. Use below command to create RDF Group:

symrdf addgrp -label <GrpLable> -rdfg <GrpNum> -sid <SID> -dir <DIR.Port,DIR.Port> -remote_rdfg <GrpNum> -remote_sid <SID> -remote_dir <DIR.Port,DIR.Port>

- Output of sample RDF creation between R11 Storage and DR storage R22.

(0) root @ Node1: /opt/emc/SYMCLI/bin

symrdf addgrp -label rdf_grp19 -rdfg 19 -sid 000197902237 -dir 1E:11 -remote_rdfg 18 -remote_sid 000197902239 -remote_dir 1E:06

**Execute a Dynamic RDF Addgrp operation for group
'rdf_grp19' on Symm: 000197902237 (y/[n]) ? y
Successfully Added Dynamic RDF Group 'rdf_grp19' for Symm:
000197902237
Verify the group created using below command: rdf_grp19**

Note: Verify command symcfg list -ra all -switched to check RDF created or not.

12. Create a disk pair for RDF group (between R11 and DR storage R22).
 - i. To create RDF disk replication, create file with disk id of storage and target storage.
 - Output of sample disk pair between R1 and DR Storage R22.

(0) root @ Node1: /usr/symcli/bin

cat /pair2.txt

00044 0004C

Note: * To get the available disk list use symdev list pd command.

13. Adding disk pair to the RDF Group 2 (between R11 and DR storage R22)
 - i. Use below command from storage R11.


```
# symrdf createpair -sid <SID> -file <disk_id_file > -rdfg <GrpNum> -type  
r1 - establish -metro
```

 - Output of sample adding disk pair from R11.

(0) root @ Node1: /usr/symcli/bin

**# symrdf createpair -sid 237 -file /pair2.txt -rdfg 19 -type r1 -
establish**

**Execute an RDF 'Create Pair' operation for device file
'/pair2.txt' (y/[n]) ? y**

**An RDF 'Create Pair' operation execution is in progress for device
file '/pair2.txt'. Please wait...**

Create RDF Pair in (2237,019).....Started.

Create RDF Pair in (2237,019).....Done.

**Mark target device(s) in (2237,019) for full copy from
source....Started.**

Devices: 0044-0044 in (2237,019).....Marked.

**Mark target device(s) in (2237,019) for full copy from
source....Done.**

**Merge track tables between source and target in
(2237,019).....Started.**

Devices: 0044-0044 in (2237,019).....Merged.
Merge track tables between source and target in
(2237,019).....Done.
Resume RDF link(s) for device(s) in (2237,019).....Started.
Resume RDF link(s) for device(s) in (2237,019).....Done.

**The RDF 'Create Pair' operation successfully executed for device
file '/pair2.txt'.**

14. Create RDF group between R12 storage and DR storage R22.

- i. Command to create RDF Group:

```
# symrdf addgrp -label <GrpLable> -rdfg <GrpNum> -sid <SID> -dir  

<DIR.Port,DIR.Port> -remote_rdfg <GrpNum> -remote_sid <SID> -  

remote_dir <DIR.Port,DIR.Port>
```

- Output of sample RDF creation between R12 Storage and DR storage R22.

(0) root @ Node1: /opt/emc/SYMCLI/bin

```
# symrdf addgrp -label rdf_grp19 -rdfg 19 -sid 000197902238 -dir  

1E:06 -remote_rdfg 18 -remote_sid 000197902239 -remote_dir 1E:06
```

**Execute a Dynamic RDF Addgrp operation for group
'rdf_grp19' on Symm: 000197902238 (y/[n]) ? y
Successfully Added Dynamic RDF Group 'rdf_grp19' for Symm:
000197902238
Verify the group created using below command: rdf_grp19**

Note: Verify command symcfg list -ra all -switched to check RDF created or not.

15. Create metro DR group between R11 and R22

- i. Use the below command to create Metro DR group

```
#symmdr -sid <ArrayID> -name <EnvironmentName> -metro_rdfg  

<RDFGroup> -dr_rdfg <RDFGroup> environment -setup
```

- Output of sample Metro DR Group creation:

(0) root @ Node3: /opt/emc/SYMCLI/bin

```
# symmdr -sid 238 -name metrodr2 -metro_rdfg 4 -dr_rdfg 6 environment  

- setup
```

Execute a MetroDR 'Environment Setup' operation (y/[n]) ? y

A MetroDR 'Environment Setup' operation is in progress

for 'metrodr2'. Please wait...

```

environment attributes .....Started.
MetroR1_ArrayID: 000197902237, Metro_RDFG : 0004, DR_RDFG :
0005
MetroR2_ArrayID: 000197902238, Metro_RDFG : 0004, DR_RDFG :
0006
DR_ArrayID : 000197902239, MetroR1_RDFG: 0005, MetroR2_RDFG:
0006
DR Mode : Adaptive Copy, MetroDR Devs: 2
Set environment attributes .....Done.
Create RDF Pair(s) (MetroR2,DR).....Started.
Create RDF Pair(s) (MetroR2,DR).....Done.
Mark incremental copy (MetroR2,DR).....Started.
Mark incremental copy (MetroR2,DR).....Done.

The MetroDR 'Environment Setup' operation successfully executed
for 'metrodr2'.

```

Note: * Metro R1 to R2 RDF group number: 4
 * Metro R1 to DR RDF group number: 5
 * Metro R2 to DR RDF group number: 6
 * MetroDR name is user defined, here it is metrodr2.

16. Creating volume groups, logical volumes, and file systems on the disk where RDF defined on the primary site.

i. Use the below commands to create the following:

a. Volume Group

```
#/usr/sbin/mkvg -f -V 50 -y metro8_vg1 hdisk19 -> VG with
metro8_VG1
```

b. Logical Volume

```
# /usr/sbin/mklv -t jfs2 -y metro8_lv1 metro8_vg1 50
-> LV with name metro8-lv1
```

c. File System

```
# /usr/sbin/crfs -v jfs2 -d /dev/metro8_lv1 -m /metro8_fs8 -
A no -> FS with name metro8_fs8
```

d. To disable auto varyon

```
# /usr/sbin/chvg -a n -Q y metro8_vg1
```

e. To varyoffvg

```
# varyoffvg metro8_vg1
```


17. After creating the volume group, create the filesystem on the primary site of Node1.
 - i. Enable PVID on the Primary site nodes Node2 (**Connected to R12 Storage**).
 - Run following command to enable PVID:
#chdev -l hdiskXX -a pv=yes -g
 - ii. After enabling the PVID on the secondary site, PVID of a disk should be the same as that of primary site disk, between which relationship is formed.
18. Import the volume group on all the nodes of the primary site after enabling the PVID's.

importvg -V <Major Number> -y <VG Name> -f hdiskXX

 - Output of sample importvg:
#importvg -V 50 -y metro8_vg1 -f hdisk28 metro8_vg1

Note: * Varyoffvg is required because PowerHA will take control of the volume group after the PowerHA cluster services are started.
19. Import the volume group on all the nodes of the Secondary site after enabling the PVID's.
 - i. Before importing volume on the secondary site nodes metrodr should be in split state:
Use below command to keep metrodr in split state.

symmdr -sid <storage sid> -name <metrodr name> split -dr -force
 - Output of sample Metrodr split:
symmdr -sid 237 -name metrodr8 split -dr -force
 - ii. Import volume group on all secondary site nodes:
importvg -V <Major Number> -y <VG Name> -f hdiskXX
 - Output of sample importvg :
#importvg -V 50 -y metro8_vg1 -f hdisk28 metro8_vg1
 - iii. After importing volume group on all the secondary site nodes, the metrodr should be brought into normal state (consistent) using establish command:
symmdr -sid <storage sid> -name <metrodr name> establish -dr
 - Output of sample Metrodr establish:
symmdr -sid 237 -name metrodr8 establish -dr

Integrating EMC SmartDR with PowerHA SystemMirror:

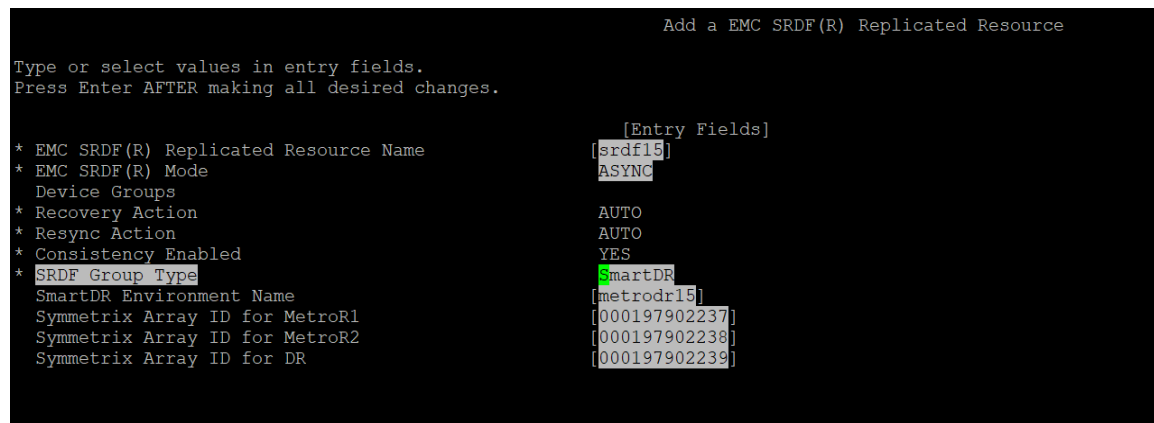
1. Adding SRDF SmartDR replicated resource to PowerHA

User must add an SRDF replicated resource before adding them to a resource group.

To add an SRDF SmartDR replicated resource, complete following steps:

- i. Open SMIT interface using the smitty hacmp command
- ii. Select the following option to navigate to the screen as shown in Figure 2.

**Cluster Applications and Resources -> Resources -> Configure EMC SRDF(R)
Replicated Resources -> Add EMC SRDF(R) Replicated Resource**



```

Add a EMC SRDF(R) Replicated Resource

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

* EMC SRDF(R) Replicated Resource Name      [srdf15]
* EMC SRDF(R) Mode                          ASYNC
Device Groups
* Recovery Action                          AUTO
* Resync Action                             AUTO
* Consistency Enabled                       YES
* SRDF Group Type                          SmartDR
SmartDR Environment Name                   [metrodr15]
Symmetrix Array ID for MetroR1              [000197902237]
Symmetrix Array ID for MetroR2              [000197902238]
Symmetrix Array ID for DR                   [000197902239]
```

Figure 2. Adding SRDF SmartDR replicated resource to PowerHA

Refer to the following description for the entries listed in figure 2.

* EMC SRDF(R) Replicated Resource Name: The name of the SRDF replicated resource that is defined to PowerHA SystemMirror as a composite group or a device group on the EMC storage device.

* EMC SRDF(R) Mode: Provide the EMC SRDF(R) Mode. The allowed values are SYNC and ASYNC. If you select SmartDR as the SRDF group Type, the default mode ASYNC is used for SRDF mirroring.

Device Groups: A list of device groups that is included in the preceding composite group or a device group in the replicated resource. If you select SmartDR as the SRDF Group Type, do not enter any value for this field.

* Recovery Action: Provide the Recovery Action. The allowed values are MANUAL and AUTO.

* Resync Action: If a site failover occurs, specify how you want to resynchronize the data between the sites after the failed site is brought back online. Enter **AUTO** if you want to automatically resynchronize the data between the sites as soon as the failed site is online. Enter **MANUAL** if you want to manually resynchronize the data between the sites. The default value for this field is **AUTO**.

* Consistency Enabled: Specify whether consistency is enabled or not. The allowed values are YES, NO. If you select SmartDR as the SRDF Group type, the default value YES is used for this field.

* SRDF Group Type:

- If devices are managed as a composite group, select **CompGroup**.
- If devices are managed as a device group, select **DevGroup**.
- If devices are managed as a Metro/SmartDR environment, select **SmartDR**.

SmartDR Environment Name: Enter the environment name for storage that is used during Metro/SmartDR configuration.

Symmetrix Array ID for MetroR1: Enter Symmetrix Array ID for MetroR1 Storage. This is numeric value.

Symmetrix Array ID for MetroR2: Enter Symmetrix Array ID for MetroR2 Storage. This is numeric value.

Symmetrix Array ID for DR: Enter Symmetrix Array ID for DR Storage. This is numeric value.

2. Adding SRDF SmartDR replicated resource to Resource Group

The SRDF SmartDR replicated resources must be added to a resource group before they can be managed by PowerHA SystemMirror.

To add a SRDF SmartDR replicated resource to a PowerHA SystemMirror resource group, complete following steps:

- i. Open SMIT interface using the `smitty hacmp` command.
- ii. Select the following option to navigate to the screen as shown in Figure 3.

Cluster Applications and Resources -> Resource Groups -> Change/Show Resources and Attributes for a Resource Group

```

Change/Show All Resources and Attributes for a Resource Group

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

[MORE...36]                                     [Entry Fields]

Primary Workload Manager Class                  []          +
Secondary Workload Manager Class                []          +

Miscellaneous Data                             []          +
WPAR Name                                     []          +
User Defined Resources                         [ ]          +
SVC PPRC Replicated Resources                 []          +
EMC SRDF(R) Replicated Resources               [srdf15]     +
DS8K Series Global Mirror Replicated Resources █          +
XIV Replicated Resources                      []          +
TRUOCOPY Replicated Resources                  []          +
DS8K Series Metro Mirror (In-band) Resources  []          +
[BOTTOM]

```

Figure 3. Adding SRDF SmartDR replicated resource to Resource Group

Note: * To configure SRDF/Metro storage devices as replicated resources, user must use Never Fallback as the fallback policy for PowerHA® SystemMirror® resource groups.

3. Verifying SRDF SmartDR Configuration

After adding the SRDF SmartDR replicated resource to the PowerHA resource group, perform verification and synchronization from the node where SRDF SmartDR replicated resources are added to PowerHA. Cluster will be synchronized, and changes will be propagated on all nodes of cluster. Resource Group will come online based on policies, once cluster services are started and **STABLE**.

- i. Verify PowerHA configured smartDR values using “odmget HACMPsr”.

```

# odmget HACMPsr
HACMPsr:
  SRDFCgName = ""
  phaSRDFname = "srdf15"
  SRDFMode = "ASYNCR"
  DeviceGroups = "metrodr15_000197902237_000197902238_000197902239"
  RecoveryAction = "AUTO"
  ConsistencyEnabled = "YES"
  SRDFGrpType = "SDR"
  ResyncAction = "AUTO"

```

Figure 5. odmget HACMPsr will list all srdf configured to PowerHA

Conclusion:

This article explains end to end procedure for EMC Metro DR storage configuration and integration with PowerHA systemMirror.

Reference Links:

<https://www.ibm.com/docs/en/powerha-aix/7.2?topic=resources-installing-srdf-filesets>

<https://www.ibm.com/docs/en/powerha-aix/7.2?topic=resources-adding-srdf-replicated-resource>

<https://www.ibm.com/docs/en/powerha-aix/7.2?topic=resources-planning>

About the author:

Aruna Sree Chittigala has been working in the IT industry for the last 16 years and has expertise in PowerHA High Availability, Disaster Recovery with storage replicate and host-based replication. She also worked in PowerVS Cloud testing with GLVM. Aruna is Currently working as Test Architect for the products PowerHA and Cloud. You can reach Aruna Sree Chittigala @ Aruna.Chittigala1@in.ibm.com.

Md Azhar is an software engineer in the PowerHA® SystemMirror® product team. He has more than 3 years of experience in the IBM Power® platform and has knowledge on disaster recovery and high availability.