

VMAX3 Business Continuity Management Student Guide

Education Services
March 2015

VMAX3 Business Continuity Management



© Copyright 2015 EMC Corporation. All rights reserved.

Support Contact: [Education Services](#)

1

Welcome to VMAX3 Business Continuity Management.

Copyright ©2015 EMC Corporation. All Rights Reserved. Published in the USA. EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license. The trademarks, logos, and service marks (collectively "Trademarks") appearing in this publication are the property of EMC Corporation and other parties. Nothing contained in this publication should be construed as granting any license or right to use any Trademark without the prior written permission of the party that owns the Trademark.

EMC, EMC² AccessAnywhere Access Logix, AdvantEdge, AlphaStor, AppSync ApplicationXtender, ArchiveXtender, Atmos, Authentica, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, Avamar, Bus-Tech, Captiva, Catalog Solution, C-Clip, Celerra, Celerra Replicator, Centera, CenterStage, CentraStar, EMC CertTracker, CIO Connect, ClaimPack, ClaimsEditor, Claralert, cLARIiON, ClientPak, CloudArray, Codebook Correlation Technology, Common Information Model, Compuset, Compute Anywhere, Configuration Intelligence, Configuresoft, Connectrix, Constellation Computing, EMC ControlCenter, CopyCross, CopyPoint, CX, DataBridge, Data Protection Suite, Data Protection Advisor, DBClassify, DD Boost, Dantz, DatabaseXtender, Data Domain, Direct Matrix Architecture, DiskXtender, DiskXtender 2000, DLS ECO, Document Sciences, Documentum, DR Anywhere, ECS, elnput, E-Lab, Elastic Cloud Storage, EmailXaminer, EmailXtender, EMC Centera, EMC ControlCenter, EMC LifeLine, EMCTV, Enginuity, EPFM, eRoom, Event Explorer, FAST, FarPoint, FirstPass, FLARE, FormWare, Geosynchrony, Global File Virtualization, Graphic Visualization, Greenplum, HighRoad, HomeBase, Illuminator, InfoArchive, InfoMover, Infoscapes, Infra, InputAccel, InputAccel Express, Invista, Ionix, ISIS, Kazeon, EMC LifeLine, Mainframe Appliance for Storage, Mainframe Data Library, Max Retriever, MCx, MediaStor, Metro, MetroPoint, MirrorView, Multi-Band Deduplication, Navisphere, Netstorage, NetWorker, nLayers, EMC OnCourse, OnAlert, OpenScale, Petrocloud, PixTools, Powerlink, PowerPath, PowerSnap, ProSphere, ProtectEverywhere, ProtectPoint, EMC Proven, EMC Proven Professional, QuickScan, RAPIDPath, EMC RecoverPoint, Rainfinity, RepliCare, RepliStor, ResourcePak, Retrospect, RSA, the RSA logo, SafeLine, SAN Advisor, SAN Copy, SAN Manager, ScaleIO Smarts, EMC Snap, SnapImage, SnapSure, SnapView, SourceOne, SRDF, EMC Storage Administrator, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix, Symmetrix DMX, Symmetrix VMAX, TimeFinder, TwinStrata, UltraFlex, UltraPoint, UltraScale, Unisphere, Universal Data Consistency, Vblock, Velocity, Viewlets, ViPR, Virtual Matrix, Virtual Matrix Architecture, Virtual Provisioning, Virtualize Everything, Compromise Nothing, Virtuent, VMAX, VMAXe, VNX, VNXe, Voyence, VPLEX, VSAM-Assist, VSAM I/O PLUS, VSET, VSPEX, Watch4net, WebXtender, xPression, xPresso, Xtrem, XtremCache, XtremSF, XtremSW, XtremIO, YottaYotta, Zero-Friction Enterprise Storage.

Revision Date: March 2015

Revision Number: MR-1CP-VMAXBCM

Course Overview

Description	This specialist level course provides the knowledge required to deploy and manage VMAX3 array-based local and remote replication solutions for business continuity needs. Operational details and implementation considerations for EMC TimeFinder SnapVX and Symmetrix Remote Data Facility (SRDF) using Symmetrix Command Line Interface (SYMCLI) and Unisphere for VMAX are covered. Lab exercises will be performed in physical (Sun Solaris and Windows 2008) and virtualized (VMware ESXi) hosts attached to VMAX3 arrays.
Audience	This course is intended for EMC Personnel, EMC Partners, and EMC Customers who will be involved in implementing and managing local and remote replication solutions for Business Continuity using TimeFinder and SRDF in an Open Systems environment.
Objectives	Upon completion of this course, you should be able to: • Describe and perform TimeFinder SnapVX operations with SYMCLI and Unisphere for VMAX – in physical and virtualized environments • Describe and perform SRDF operations in Synchronous (SRDF/S) and Asynchronous (SRDF/A) modes using SYMCLI and Unisphere for VMAX – in physical and virtualized environments



© Copyright 2015 EMC Corporation. All rights reserved.

VMAX3 Business Continuity Management

2

This course provides the knowledge required to deploy and manage VMAX3 array-based local and remote replication solutions for business continuity needs. Operational details and implementation considerations for EMC TimeFinder SnapVX and Symmetrix Remote Data Facility (SRDF) using Symmetrix Command Line Interface (SYMCLI) and Unisphere for VMAX are covered. Lab exercises will be performed in physical (Sun Solaris and Windows 2008) and virtualized (VMware ESXi) hosts attached to VMAX3 arrays.

Agenda

Module/Lessons	Labs
Day 1 - Module: TimeFinder SnapVX Operations - Lesson – TimeFinder SnapVX Concepts and Operations - Lesson – Replicating a VMFS Datastore - Lesson – Replicating a VM Accessing RDM Hard Disks - Lesson – TimeFinder SnapVX Operations using Unisphere for VMAX - Module: SRDF/Synchronous Operations - Lesson – SRDF Initial Setup Operations - Lesson – SRDF Disaster Recovery Operations - Lesson – SRDF Decision Support Operations	(Sun Solaris Hosts / Windows 2008 Hosts) - Lab Exercise 1: TimeFinder SnapVX Operations (VMware ESXi Hosts) - Lab Exercise 1: TimeFinder SnapVX Replication of a VMFS Datastore (Sun Solaris Hosts / Windows 2008 Hosts) - Lab Exercise 2: SRDF/Synchronous Operations - Part 1 – Initial Setup and Basic Operations - Part 2 – Disaster Recovery Operations



Agenda

Module/Lessons	Labs
Day 2 - Module: SRDF/Synchronous Operations (Cont) - Lesson – SRDF/S Disaster Recovery for a VMFS Datastore - Lesson – Disaster Recovery for a VM Accessing RDM Hard Disks - Lesson – SRDF/S Operations – Unisphere for VMAX - Module: SRDF/Asynchronous Operations - Lesson – SRDF/A Concepts and Operations - Lesson – SRDF/A Resiliency Features - Lesson – SRDF/A Multi-session Consistency (MSC)	(Sun Solaris Hosts/Windows 2008 Hosts) - Lab Exercise 2: SRDF/Synchronous Operations Part 3 – Decision Support/Concurrent Access Operations (VMware ESXi Hosts) - Lab Exercise 2: SRDF/S Disaster Recovery for a VMFS Datastore (Sun Solaris Hosts / Windows 2008 Hosts) - Lab Exercise 3: SRDF/Asynchronous Operations

Module: TimeFinder SnapVX Operations

Upon completion of this module, you should be able to:

- Describe TimeFinder SnapVX concepts
- Perform TimeFinder SnapVX operations using SYMCLI and Unisphere for VMAX
- Replicate a VMFS datastore using TimeFinder SnapVX
- Replicate a VM accessing RDM hard disk using TimeFinder SnapVX



This module focuses on TimeFinder SnapVX local replication technology on VMAX3 arrays. Concepts, terminology, and operational details of creating snapshots and presenting them to target hosts is discussed. Performing TimeFinder SnapVX operations using SYMCLI and Unisphere for VMAX are presented. Use of TimeFinder SnapVX for replication in a virtualized environment is also presented.

Lesson: TimeFinder SnapVX Concepts and Operations

This lesson covers the following topics:


- TimeFinder SnapVX concepts
- Performing TimeFinder SnapVX operations



This lesson covers the concepts of TimeFinder SnapVX. Operational examples using SYMCLI are presented in detail.

TimeFinder SnapVX Overview

- Create local point-in-time copies (snapshots) of data without requiring pairing between Source and Target Volumes
 - Target Volumes are required only if host access to point-in-time data is desired
- Highly scalable
 - Up to 256 snapshots per Source Volume
- Highly space efficient
 - Sharing of point-in-time tracks among different snapshots



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

3

TimeFinder SnapVX provides a highly efficient mechanism for taking periodic point-in-time copies of source data without the need for target devices. Target devices are required only for presenting the point-in-time data to another host. Sharing allocations between multiple snapshots makes it highly space efficient. A write to the source volume will only require one snapshot delta to preserve the original data for multiple snapshots. If a source track is shared with a target or multiple targets, a write to this track will preserve the original data as snapshot delta and will be shared for all the targets. Write to the target will be applied only to the specific target.

TimeFinder SnapVX Terminology

- Source Volume
 - A VMAX3 device whose point-in-time copy is to be preserved
- Snapshot
 - Preserved point-in-time image of the Source Volume
- Snapshot Delta
 - Original Source Volume tracks at the point-in-time of the Snapshot that were preserved during host writes to the Source Volume
- Linked Target Volume
 - A VMAX3 device that is used to provide access to point-in-time data by linking it to a Snapshot
- Storage Resource Pool
 - A collection of data pools that are managed by FAST and provide physical storage for VMAX3 thin devices



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations


4

The terminology used in SnapVX is described in the slide. Note that all host accessible devices in a VMAX3 are thin devices.

Host writes to source volumes will create snapshot deltas in the SRP. Snapshot deltas are the original point-in-time data of tracks that have been modified after the snapshot was established.

SRP configuration must be specified when ordering the system, prior to installation. The source and target volumes can be associated with the same SRP or different SRPs. Snapshot deltas will always be stored in the source volume's SRP. Allocations owned by the source will be managed by its SLO. Allocations for the target will be managed by the target's SLO. Snapshot deltas will be managed by the Optimized SLO.

Re-direct on Write – Preserving Point-in-time Data



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

EMC²


5

When the snapshot is created, both the source device and the snapshot point to the location of data in the SRP. When a source track is written to, the new write is asynchronously written to a new location in the SRP. The source volume will point to the new data. The snapshot will continue to point to the location of the original data. The preserved point-in-time data becomes the snapshot delta. This is the re-direct on write mechanism.

Under some circumstances SnapVX will use Asynchronous Copy of First Write (ACOFW). This might be done to prevent degradation of performance for the source device. For example, if the original track was allocated on Flash drive, then it would be better to copy this down to a lower tier and accommodate the new write in the Flash drive.

Snapshot Generations and Time-to-Live

- **Generation Numbers:**
 - Any snapshot of a Source Volume with a unique name is assigned Generation Number 0
 - If a Source Volume has multiple snapshots with the same name, the latest snapshot is assigned Generation Number 0; the one prior to it will be Generation Number 1 and so on
- **Time-to-Live (TTL):**
 - At time of creation or later, an expiration time for the snapshot can be set



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations


6

Each snapshot is assigned a generation number. If the name assigned to the snapshot is reused, then the generation numbers are incremented. The most recent snapshot with the same name will be designated as generation 0, the one prior as generation 1, and so on. If each snapshot is given a unique name, they will all be generation 0. Terminating a snapshot will result in reassignment of generation numbers.

Time-to-live (TTL) can be used to automatically terminate a snapshot at a set time. This can be specified at the time of snapshot creation or can be modified later. HYPERMAX OS will terminate the snapshot at the set time. If a snapshot has linked targets, it will not be terminated. It will be terminated only when the last target is unlinked. TTL can be set as a specific date or as a number of days from creation time.

Accessing Point-in-time Data – Linked Targets

- Snapshot should be linked to a Target Volume if Host access to the point-in-time data is required
- Linked Targets can be in:
 - No Copy Mode (default)
 - Copy Mode
- Mode can be specified at the time of snapshot creation or modified later



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

7

A snapshot has to be linked to a target volume to provide access to point-in-time data to a host. The link can be in No Copy or Copy mode. Copy mode linked targets will provide full volume copies of the point-in-time data of the source volumes – similar to full copy clones. Copy mode linked targets will have useable copy of the data even after termination of the snapshot – provided the copy has completed.

A snapshot can have both no copy mode and copy mode linked targets. Default is to create No Copy mode linked targets. This can be changed later if desired.


Writing to a linked target will not affect the snapshot. The target can be re-linked to the snapshot to revert to the original point-in-time.

A snapshot can be linked to multiple targets. But a target volume can be linked to only one snapshot.

There is no benefit to have the no copy mode linked targets in an SRP different from the Source SRP. Writes to the Source volume will only create snapshot deltas which will be stored in the Source volume's SRP. The writes will not initiate any copy to the target.

A target volume that is larger than the source can be linked to a snapshot. This is enabled by default. The environment variable SYMCLI_SNAPVX_LARGER_TGT can be set to DISABLE to prevent this.

Linked Target – Undefined/Defined Tracks



- Undefined Tracks – Location of data for the Target has to be resolved through snapshot pointers
- Defined Tracks – Location of data for the Target points directly to the appropriate tracks in the SRP

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

8

When a snapshot is linked to a target, the process of defining the tracks for the target is initiated internally. In the undefined state, the location of data for the target has to be resolved through the pointers for the snapshot. In the defined state, data for the target points directly to the corresponding locations in the SRP.

Relinking and Unlinking Targets

- Relink operation will unlink the target from the current snapshot and link it to a different snapshot
 - Relink has to be between the same Source and Target devices
- Relink to the same snapshot will refresh the Target with the original point-in-time
 - Useful if target has been modified and we want to revert to the original snapshot
- Unlink operation disassociates the linked target from the snapshot
 - Copy mode linked targets will contain a full and useable copy of the source devices, after copy completes
 - No Copy mode linked targets will not contain valid data after unlinking



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

9

Relink provides a convenient way of checking different snapshots to select the appropriate one to access. A link between the snapshot of the source volume and the target must exist for the relink operation. Relink can also be performed with a different snapshot of the same source volume or a different generation of the same snapshot of the source volume. Unlink operation removes the relationship between a snapshot and the corresponding target. Copy mode linked targets can be unlinked after the copying completes. This will provide a full, independent useable point-in-time copy of the Source data on the Target device.

No copy mode linked targets can be unlinked at any time. After unlinking a no copy mode linked target, the target device cannot be considered as useable. This is because the target would have shared tracks with the source volume and/or the snapshot deltas. These tracks would not be available on the target after the unlink operation and hence render it un-useable.

Restore to Source

- Restore from snapshot
 - Snapshots can be directly restored to the Source volume
 - Source volume data will be set back to the point-in-time of the snapshot
 - Only changed data has to be restored from the snapshot delta – inherently differential operation
- Restore from linked target
 - Two step process
 - Create a snapshot of the linked target
 - Link this snapshot with the Source volume (which will now be a linked target)




© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

10

As the data on the Source volume from the host's perspective will be changing, the Source volume should be unmounted prior to the restore operation and then re-mounted. To restore from a linked target, a snapshot of it must be established; and this snapshot should be linked to the Source volume. The Source volume cannot be unlinked until copy completes. So the link should be created in copy mode.

Cascading Snapshots



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

EMC²

11

Snapshots of linked targets can be created. These can further be linked to other targets. This is referred to as cascading. There is no limit to the number of cascaded hops that can be created as long as the overall limit for SnapVX is maintained. However there may not be many practical uses for cascading, given the efficiency of SnapVX technology. Writes to linked target do not affect the snapshot. It always remains pristine – in effect the “gold” copy. If one must experiment with data on the linked target, there is no need to save a gold copy prior to this. When done with the experimentation, one can always refresh the target data with the original snapshot data by relinking. The linked target must be in a defined or copied state before a snapshot of it can be created. A cascaded snapshot can only be restored to the linked target that is in copy mode and has fully copied. If the linked target is in no copy mode, it cannot be unlinked without first terminating any snapshots that have been created from it. A linked target that has a cascaded snapshot must be fully copied before being unlinked. A snapshot with linked targets cannot be terminated.

Reserved Capacity

- Reserved Capacity is a percentage of Storage Resource Pool (SRP) that can be only allocated to new host writes
- When SRP gets to where only Reserved Capacity is left:
 - Snapshot will fail when new allocations for snapshot deltas are required
 - The snapshot has to be terminated
 - Copy to targets will halt
 - Copy will resume if free space is made available in the SRP or if the Reserved Capacity is lowered



Reserved capacity ensures that there will be sufficient capacity available in the SRP to accommodate new host writes. When the allocated capacity reaches the point where only reserved capacity remains, then SnapVX allocation for snapshot deltas and copy processes will be affected.

Expanding Storage Groups with Active Snapshots

- New volume(s) can be added to Storage Groups based on the growth of the application
- Existing snapshots will not include the new volume(s)
- New snapshots will include the newly added volume(s)
- If the SG is restored from an earlier snapshot, the new volume(s) will be set Not Ready to the host
 - The volume(s) will remain NR even after the restored session is terminated. User has to decide the best course of action to include these for the application and host
- Likewise if the SG of the linked targets has been expanded after the snapshot, these volume(s) will be set Not Ready to the host



Care needs to be exercised when expanding SGs with existing snapshots. If there are more volume(s) in the SG than are contained in the snapshot, then a restore from the snapshot will set these additional volume(s) to Not Ready. This is because these volume(s) were not present when the snapshot was taken. Of course subsequent snapshots after the SG expansion will contain all the volumes. Similarly if the linked target SG has been expanded and has more devices than the snapshot, then the additional volumes in the linked target SG will be set to Not Ready.

Creating Snapshots

```
C:\>symsnapvx -sid 483 -sg snapvx establish -name backup
```

```
Execute Establish operation for Storage Group snapvx (y/[n]) ? y
```

```
Establish operation execution is in progress for the storage group snapvx. Please wait...
```

```
  Polling for Establish.....Started.
```

```
  Polling for Establish.....Done.
```

```
  Polling for Activate.....Started.
```

```
  Polling for Activate.....Done.
```

```
Establish operation successfully executed for the storage group snapvx
```

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

EMC²

14

The most convenient and preferred way of performing TimeFinder SnapVX operations is using Storage Groups. In this example we are creating a snapshot named `backup` for the devices in the Storage Group `snapvx`.

Listing Snapshots

```
C:\>symsnapvx -sid 483 list -sg snapvx -detail
```

```
Storage Group (SG) Name      : snapvx
SG's Symmetrix ID           : 000196700483   (Microcode Version: 5977)
```

Sym Dev	Snapshot Name	Gen	Flgs		Snapshot Timestamp	Deltas (Tracks)	Total Non-Shared (Tracks)	Expiration
			FLRG					
00049	backup	0		Sun	Feb 22 10:06:37 2015	4557	1317	NA
	backup	1		Sun	Feb 22 10:05:40 2015	8033	1766	NA
	backup	2		Sun	Feb 22 10:00:22 2015	11146	6074	NA
						23736	9157	

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

15

We have created three successive snapshots using the same name. Note that each snapshot is given a generation number. As discussed earlier, the most recent snapshot is designated as generation 0. As there is workload on the source devices, the changes are accumulated in snapshot deltas. The non-shared tracks are unique to the specific snapshot. These are the tracks that will be returned to the SRP if the snapshot is terminated. As we did not specify a time-to-live during the establish operation, the Expiration Date is NA. Note that the output has been edited to fit the slide.

Setting Time-to-Live

```
C:\>symsnapvx -sid 483 -sg snapvx -snapshot_name backup -generation 0 set ttl -delta 5
```

```
C:\>symsnapvx -sid 483 list -sg snapvx -detail
```

```
Storage Group (SG) Name      : snapvx  
SG's Symmetrix ID           : 000196700483 (Microcode Version: 5977)
```

```
-----  
Sym          Flgs          Total  
Dev  Snapshot Name  Gen FLRG Snapshot Timestamp  Deltas  Non-Shared  
-----  
00049 backup      0 .... Sun Feb 22 10:06:37 2015  4557      1317  Fri Feb 27  
10:26:27 2015
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

16

We can set the time-to-live even after creating the snapshot. Output has been edited to fit the slide. The parameter `-delta` is used to specify the number of days for expiration from the time the snapshot was created.

Linking Snapshot to Target

```
C:\> symsnapvx -sid 483 link -sg snapvx -snapshot_name backup -gen 2 -lnsg tgtvx

Execute Link operation for Storage Group snapvx (y/[n]) ? y

Link operation execution is in progress for the storage group snapvx. Please wait...
Polling for Link.....Started.
Polling for Link.....Done.

Link operation successfully executed for the storage group snapvx

C:\>symsnapvx -sid 483 list -sg snapvx -detail -linked

Storage Group (SG) Name      : snapvx
SG's Symmetrix ID           : 000196700483   (Microcode Version: 5977)
```

Sym	Link	Flgs	Remaining	Done
Dev	Snapshot Name	Gen Dev FCMD Snapshot Timestamp	(Tracks)	(%)
00049	backup	2 0004F ...X Sun Feb 22 10:00:22 2015	16395	0



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

17

We can specify the generation of the snapshot that we want to link. The target device is contained in a Storage Group as well. It is specified with the `-lnsg` flag. The default for linking is the No Copy mode. So we see that no data has been copied yet. Furthermore, at this point in time, we are not writing to the target device either.

Flgs:

(F)ailed : F = Force Failed, X = Failed, . = No Failure

(C)opy : I = CopyInProg, C = Copied, D = Copied/Destaged, . = **NoCopy**
Link

(M)odified : X = Modified Target Data, . = **Not Modified**

(D)efined : **X = All Tracks Defined**, . = Define in progress

When Target Is Modified

```
C:\>symsnapvx -sid 483 list -sg snapvx -detail -linked
```

```
Storage Group (SG) Name      : snapvx  
SG's Symmetrix ID           : 000196700483   (Microcode Version: 5977)
```

Sym	Snapshot Name	Gen	Dev	Link	Flgs	FCMD	Snapshot	Timestamp	Remaining (Tracks)	Done (%)
00049	backup	2	0004F	..	XX		Sun Feb 22	10:00:22 2015	12173	25

12173

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

18

When target is written to, the original point-in-time snapshot is unaffected – it remains pristine. The % Done and Remaining (Tracks) indicates tracks that have been changed due to the writes. This helps with incremental operations later if desired.

Relinking Snapshot to Target

```
C:\>symsnapvx -sid 483 relink -sg snapvx -snapshot_name backup -gen 1 -lnsg tgtvx -nop
```

```
Relink operation execution is in progress for the storage group snapvx. Please wait...
Polling for Relink.....Started.
Polling for Relink.....Done.
```

```
Relink operation successfully executed for the storage group snapvx
```

```
C:\>symsnapvx -sid 483 list -sg snapvx -detail -linked
```

```
Storage Group (SG) Name      : snapvx
SG's Symmetrix ID           : 000196700483   (Microcode Version: 5977)
```

Sym	Link	Flgs	Remaining	Done				
Dev	Snapshot Name	Gen	Dev	FCMD	Snapshot	Timestamp	(Tracks)	(%)
00049	backup	1	0004F	...X	Sun Feb 22	10:05:40 2015	16395	0
							16395	

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

19

In this example we are relinking the target to snapshot generation 1 of the same source volume. The target volumes should be unmounted prior to relinking and then mounted back again to ensure that the host accesses correct data. We can link/relink different snapshots to target volumes to select the desired point-in-time.

Restoring from Snapshots

```
C:\>symsnapvx -sid 483 restore -sg snapvx -snapshot_name backup -gen 1 -nop
```

Restore operation execution is in progress for the storage group snapvx. Please wait...

 Polling for Restore.....Started.

 Polling for Restore.....Done.

Restore operation successfully executed for the storage group snapvx

```
C:\>symsnapvx -sid 483 list -sg snapvx -detail
```

Storage Group (SG) Name : snapvx
SG's Symmetrix ID : 000196700483 <Microcode Version: 5977>

Sym Dev	Snapshot Name	Gen	Flgs FLRG	Snapshot Timestamp	Total Deltas <Tracks>	Non-Shared <Tracks>	Expiration Date
00049	backup	0		Sun Feb 22 10:06:38 2015	8033	4793	Fri Feb 27 10:26:27 2015
	backup	1	.XX.	Sun Feb 22 10:05:41 2015	8033	1766	NA
					16066	6559	

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

20

Any available snapshot can be restored to the source volume. This will revert the data on the source volume to that specific point-in-time. As the data on the disk will be changing from the host's perspective, it is recommended to unmount the source volume prior to performing a restore operation. After the restore, the source volume can be mounted again to access the correct data. Terminating the restored session will leave the original snapshot intact. It will only terminate the restored session.

Terminating Snapshots

```
C:\>symsnapvx -sid 483 terminate -sg snapvx -snapshot_name backup -gen 2 -nop

Terminate operation execution is in progress for the storage group snapvx. Please wait...
  Polling for Terminate.....Started.
  Polling for Terminate.....Done.

Terminate operation successfully executed for the storage group snapvx
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

21

Snapshots that have linked targets in no copy mode cannot be terminated. One must first either unlink the target or change the mode to copy and terminate after copying completes. Terminating a snapshot that has a restored session would require terminating the restored session first, followed by terminating the snapshot.

SnapVX Emulation Support

- Supported
 - TimeFinder/Clone, TimeFinder VP Snap, and TimeFinder/Mirror
 - Source to Target pairing is required
 - Full volume copy
 - Snapshot deltas are not used
 - All existing VMAX restrictions at the latest level of Enginuity 5876 apply
- Cannot mix SnapVX and Emulation mode on same volumes
- Not supported
 - TimeFinder/Snap



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

22

TimeFinder SnapVX is the underlying technology that supports emulation mode for TimeFinder/Clone, TimeFinder/Mirror, and TimeFinder VP Snap commands. These emulations will be completely seamless and will automatically be invoked when performing TimeFinder/Mirror, Clone, and VP Snap operations. Emulation sessions will copy data directly from the source to target without using snapshot deltas. Emulation modes will use legacy Source-Target pairing. This will provide backwards compatibility with existing scripts that execute TimeFinder command and control operations. When legacy TimeFinder commands are used, SnapVX sessions are created in the background. All existing restrictions and session limits for these emulations are carried over from the latest version of Enginuity 5876. Emulation mode will not support the storage group (-sg) option.

VMAX3 volumes cannot be used as either SnapVX sources or link targets when participating in emulation sessions. Conversely, volumes cannot be used for emulation sessions when they are SnapVX sources or link targets. TimeFinder/Snap is no longer needed because of SnapVX point in time technology. SAVE devices do not exist in VMAX3 arrays.

Lesson Summary

During this lesson the following topics were covered:

- TimeFinder SnapVX concepts
- Performing TimeFinder SnapVX operations



This lesson covered concepts and terminology of TimeFinder SnapVX. Creating snapshots and other SnapVX related operations using SYMCLI were also covered.

Lesson: Replicating a VMFS Datastore

This lesson covers the following topics:

- Using TimeFinder SnapVX to create a snapshot of a VMFS datastore presented to the Primary ESXi server
- Linking the snapshot to a target and presenting it to the Secondary ESXi server



This lesson covers replicating a VMware VMFS datastore using TimeFinder SnapVX. A snapshot of the VMFS datastore presented to the Primary server will be created and linked to a target device. The target is then accessed on a Secondary ESXi server.

Primary ESXi Datastore

The screenshot shows the vSphere Client interface for a Primary ESXi64 VMWare ESXi, S.1.0, 799733. The left sidebar shows the 'Inventory' view with 'StudentVM01' selected. The main pane shows the 'Datastores' tab, which lists two datastores: 'datastore1' (Local VMWare Disk) and 'Production_Datastore' (EMC iSCSI Disk). The 'Production_Datastore' is highlighted, and its details are shown in the 'Datastore Details' section.

Path Selection	Properties	Extents
Round Robin (VM...)	Volume Label: Production_...	EMC iSCSI Disk (naa.60000...
	Datastore Name: Production_...	10.00
		Total Formatted Capacity 9.75

The 'Datastore Details' section also includes a 'Formatting' table:

Formatting
File System: VMFS 5.58
Block Size: 1 MB



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

25

Using the vSphere Client we find that the Primary ESXi server has access to Production_Datastore. Note the naa number of the device. We will use this number to correlate the device with the VMAX3 volume, using Unisphere for VMAX.

VM Resident on Production Datastore

The screenshot shows the VMware vSphere Datastore Browser interface. The main window displays the 'Production_Datastore' under the 'Datastores' tab. A 'Browse Datastore...' button is highlighted. The 'Datastore Details' pane on the left shows the location as /vmfs/volumes/54d22d1d and hardware acceleration as Supported. The 'Datastore Browser - [Production_Datastore]' window is open, showing a folder named 'StudentVM01'. The files within this folder are listed in the table below:

Name	Size	Provisioned Size
StudentVM01.vmx	1.84 KB	
StudentVM01.vmdk	1,887,232.00 KB	2,097,152.00 KB
vmware.log	91.56 KB	
StudentVM01.nvram	8.48 KB	
StudentVM01.vmx.f	0.26 KB	
StudentVM01.vmsd	0.00 KB	
vmx-StudentVM01-409862235...	93,184.00 KB	
StudentVM01.vmx.lck	0.00 KB	
StudentVM01.vmx~	1.82 KB	
StudentVM01-f44c0396.vswp	524,288.00 KB	

EMC


© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

26


Browsing the Production_Datastore shows that it has the folder StudentVM01. This folder contains the StudentVM01.vmx and other files pertaining to the VM StudentVM01.

StudentVM01 Uses only the Production Datastore



Summary view of Virtual Machine shows that it uses only the Production_Datastore. The VM is currently powered on.

StudentVM01 Contains Production Data



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

28

We can open a console to the StudentVM01 and examine the data. For the purposes of this example a folder named Production_data has been created on StudentVM01. The objective is to use TimeFinder SnapVX to take a snapshot of the VMAX3 device hosting the Production_Datastore. We have to identify a suitable target device accessible to a Secondary ESXi server. Then we can link the snapshot to the target device. Subsequently we should be able to power on a snapshot of the StudentVM01 on the Secondary ESXi Server.

Identifying Device Hosting Production Datastore

The screenshot shows the EMC Unisphere for VMAX interface. The top navigation bar includes 'Storage' and 'Hosts' tabs. The breadcrumb trail indicates the path: '000196800225 > Hosts > Masking View > primaryesxi'. A table on the left lists devices, with '00095' selected. The right pane displays the details for this device.

Name	Type	Emulation	Allocated
00095	TDEV	FBA	

Volume Identifier	N/A
Type	TDEV
Encapsulated Volume	No
Encapsulated WWN	N/A
Status	Ready
Reserved	No
Capacity (GB)	10
Capacity (MB)	10241
Capacity (Cylinder)	5462
Emulation	FBA
Symmetrix ID	000196800225
Symmetrix Volume ID	00095
HP Identifier Name	N/A
VMS Identifier Name	N/A
Nice Name	N/A
WWN	60000970000196800225533030303935

In Unisphere for VMAX, we navigate to the Masking View for the Primary ESXi Server and identify the device it has access to. In this example it is device 00095. Listing the details of this device shows the WWN for it. We can match this WWN with the naa number shown in slide 25 and conclude that the Primary ESXi Server has access to device 00095. This device is in SID:225 and its capacity is 10 GB. So in order to take a snapshot of this device and link it to a target, we have to identify a 10 GB device on SID:225 that has been masked to the Secondary ESXi Server.

vSphere Web Client

The screenshot displays the vSphere Web Client interface for the 'Production_Datastore'. The left sidebar shows 'Production_Datastore' selected under 'Virtual Machines'. The top navigation bar includes 'Summary', 'Monitor', 'Manage', and 'Related Objects' tabs. The main content area is divided into two sections: 'Storage System' and 'Storage Device'.

Storage System Details:

Property	Value
Name	000196800225
Product	VMAX
Model	VMAX100K
Revision	5977.497.471

powered by EMC

Storage Device Details:

Property	Value
SLO	Optimized
Workload Type	None
Parent Storage Group	primaryesi64_parent
Storage Group	primaryesi64_prod
Storage Resource Pool (SRP)	SRP_1
Raid Level	RAID-1
Total Capacity	10.00 GB
Available Capacity	10.00 GB
Used Capacity	0 bytes
Block Size	512 bytes
Is Thin Provisioned	true
WWN	60000970000196800225533030303935
Status	ONLINE
Type	TDEV
Name	00095
Volume Identifier	N/A
Recover Point Tagged	No

EMC

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

30

Alternatively if we have access to the vSphere Web Client with the EMC Storage Viewer plug-in, we can use it to correlate Production_Datastore with device 00095 as well.

Identifying Target LUN


The screenshot shows the EMC Unisphere for VMAX interface. On the left, a table lists devices, with '00094' highlighted. The main panel displays the details for this device. The 'Capacity (GB)' is 10, 'Symmetrix ID' is 000196800225, 'Symmetrix Volume ID' is 00094, and 'WWN' is 600009700001968002255330303934. The 'WWN' field is highlighted with a red box.

Name	Type	Emulation
00094	TDEV	FBA

Physical Name	N/A
Volume Identifier	N/A
Type	TDEV
Encapsulated Volume	No
Encapsulated WWN	N/A
Status	Ready
Reserved	No
Capacity (GB)	10
Capacity (MB)	10241
Capacity (Cylinder)	5462
Emulation	FBA
Symmetrix ID	000196800225
Symmetrix Volume ID	00094
HP Identifier Name	N/A
VMS Identifier Name	N/A
Nice Name	N/A
WWN	600009700001968002255330303934
DG Name	

Once again using Unisphere for VMAX, we navigate to the Masking View for the Secondary ESXi Server and identify the device it has access to. In this example it is device 0094. Listing the details of this device shows the WWN for it. We can match this WWN with the naa number reported in the vSphere Client for the Secondary ESXi Server shown in the next slide.

LUNs Accessible to Secondary ESXi Server



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

32

Using the vSphere Client we find that the Secondary ESXi server has access to a few devices. Note the naa number of the device highlighted. This correlates with the WWN of device 0094 as shown in the previous slide.

Create Snapshot of Production Device (1)

EMC Unisphere for VMAX V8.0.1.7

000196800225 > Data Protection > Protection Dashboard > Unprotected Storage Groups

Storage Group	Capacity (GB)
▶ app_server_parent	40
EMBEDDED_NAS_DM_SG	105.82
eNAS_SG	100.01
▶ esxi-88-34_parent	0.07
▶ esxi-88-36_parent	130.05
▶ esxi-88-46_parent	0.04
Finance_eNAS	100.01
▼ primaryesxi64_parent	10.01
primaryesxi64_gk	0.01
primaryesxi64_prod	1.0
▶ secondaryesxi65_parent	
sun-88-51_sg	
Training_eNAS	
▶ vcenter-88-20_parent	
▶ w2k8r2-88-62_parent	
▶ w2k8r2-88-63_parent	0.04

Context menu options for 'primaryesxi64_prod':

- Protect ...
- Manage Storage Groups ...
- Settings...
- Global Settings...
- About Adobe Flash Player 16.0.0.235...




Module: TimeFinder SnapVX Operations

33

In Unisphere for VMAX, SnapVX operations can only be performed on Storage Groups. As this is the first time we will be creating a snapshot for the Production Device, we navigate to Data Protection>Protection Dashboard>Unprotected Storage Groups. The Storage Group primaryesxi64_prod was created when the production device was masked to the Primary ESXi Server. Right click this Storage Group and select Protect.

Create Snapshot of Production Device (2)



In the Protect Storage Group wizard, we select Point In Time Protection Using SnapVX. The snapshot has been named `datastore_backup` and a 5 day expiration time has been set.

Link Snapshot to Target

```
C:\>symsnapvx -sid 25 link -snapshot_name datastore_backup -sg
primaryesxi64_prod -lsg secondaryesxi65_snap_tgt

Execute Link operation for Storage Group primaryesxi64_prod (y/[n]) ? y

Link operation execution is in progress for the storage group
primaryesxi64_prod. Please wait...
  Polling for Link.....Started.
  Polling for Link.....Done.

Link operation successfully executed for the storage group primaryesxi64_prod
```




© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

35


For the protection of users, Unisphere for VMAX will not permit linking a snapshot with a target Storage Group which is in a Masking View. In our example the target device 0094 has been placed in the Storage Group secondaryesxi65_snap_tgt. This Storage Group is a part of the Masking View for the Secondary ESXi server in order to enable access to the target device. So we use the CLI to perform the Link operation.

Rescan Secondary ESXi Server




Now we can rescan the Secondary ESXi Server. Choosing Rescan All will scan for new storage devices as well as VMFS volumes.

Mount Replica Datastore on Secondary ESXi (1)



After rescan completes, we use the Add Storage wizard to add the linked Target device. The Storage Type Disk/LUN is selected. This shows the VMFS Label to be Production_Datastore. This was the label given to the Production Datastore used by StudentVM01 on the Primary ESXi server. This indicates that this is the linked Target LUN. We choose this LUN and click Next.

Mount Replica Datastore on Secondary ESXi (2)



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

38

As the LUN is a replica, the wizard will offer Mount Options for VMFS. In this example we choose "Assign a new signature". Even though we are presenting the linked Target to a secondary ESXi server, it is a good practice to assign a new signature. We can then click Next and Finish.

Mount Replica Datastore on Secondary ESXi (3)

10.126.88.65 - vSphere Client

File Edit View Inventory Administration Plug-ins Help

Home Inventory Inventory

10.126.88.65 SecondaryESXi65 VMware ESXi, 5.1.0, 799733

Getting Started Summary Virtual Machines Resource Allocation Performance Configuration Local Users & Groups Events

Hardware

- Health Status
- Processors
- Memory
- Storage
- Networking
- Storage Adapters
- Network Adapters
- Advanced Settings
- Power Management

Software

- Licensed Features
- Time Configuration
- DNS and Routing
- Authentication Services

View: Datastores Devices

Datastores Refresh Delete Add Storage.

Identification	Device	Drive Type
datastore1 (1)	Local VMware Disk...	Non-SSD
snap-7e6fd57a-Production_Datastore	EMC iSCSI Disk (...)	Non-SSD

Datastore Details

snap-7e6fd57a-Production_Datastore 9.75 GB

Location: /vmfs/volumes/54d270b2-ae62d09f-16be-005056b20b17

Hardware Acceleration: Supported 3.25 GB

6.50 GB



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

39

As we chose "Assign new signature", the datastore on the Secondary ESXi server has snap-xxxxxxx as a prefix to the label.

Browse Datastore and Add VM to Inventory

The screenshot displays the vSphere Datastore Browser interface. The main window is titled "Datastore Browser - [snap-7e6fd57a-Production_Datastore]". The left pane shows the "Folders" section with a search bar and a folder named "StudentVM01". The right pane shows a list of files and folders within "StudentVM01". The file "StudentVM01.vmx" is selected, and a context menu is open with the "Add to Inventory" option highlighted. The "Datastore Details" pane on the left shows the datastore name "snap-7e6fd57a-Production_Datastore" and its location. The "Datastores" pane at the top shows the selected datastore and its properties: "Local VMware Disk..." and "Non-SSD".

Name	Size	Provisioned Size	Type
StudentVM01.vmx	1.04 KB	2,097,152.00 KB	Virtual Machine
StudentVM01.vmdk			Virtual Disk
vmware.log			Virtual Machine
StudentVM01.nvram			Non-volatile RAM
StudentVM01.vmx			File
StudentVM01.vmsd			File
vmx-StudentVM01-409			File
StudentVM01.vmx.lck			File
StudentVM01.vmx.vv			File

EMC²

© Copyright 2015 EMC Corporation. All rights reserved. Module: TimeFinder SnapVX Operations 40

We can now browse the replica Datastore. We see that it contains the folder StudentVM01. The Virtual Machine can now be added to the inventory of the Secondary ESXi server.

Add VM to Inventory

The image displays two sequential screenshots of the 'Add to Inventory' wizard interface.

First Screenshot: The 'Name' field is highlighted with a red box, showing the text 'StudentVM01_backup'. Below the field, a note states: 'Virtual machine (VM) names may contain up to 80 characters and they must be unique within each'.

Second Screenshot: The 'Resource Pool' selection screen is shown. A red box highlights the IP address '10.126.88.65' in the list of available resource pools. The text above the list reads: 'Select the resource pool within which you wish to run this virtual machine. Resource pools allow hierarchical management of computing resources within a host or cluster. Virtual machines and child pools share the resources of their parent pool.'




© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

41


We have given the name StudentVM01_backup. We are adding this to the inventory of the Secondary ESXi server. Then we can finish the Add to Inventory process.

Power-on VM on the Secondary ESXi Server



When trying to "Power On" the replica VM, we need to answer the Virtual Machine Message question. Here we choose "I copied it".

VM on the Secondary ESXi Server



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

43

We can open a console to the VM on the Secondary ESXi server and verify that this VM has the same data as the VM on the Primary ESXi server at the point in time of the snapshot.

Lesson Summary

During this lesson the following topics were covered:

- Creating a snapshot of a VMFS datastore presented to the Primary ESXi Server
- Presenting the linked snapshot Target to the Secondary ESXi Server



This lesson covered replicating a VMware VMFS datastore using TimeFinder SnapVX. A snapshot of the VMFS datastore presented to the Primary server was created and linked to a target device. The target was then accessed on a Secondary ESXi server.

Lesson: Replicating a VM Accessing RDM Hard Disks

This lesson covers the following topics:

- Using TimeFinder SnapVX to replicate a VM accessing an RDM hard disk
- Presenting the replicated VM to a Secondary ESXi Server



This lesson covers replicating a Virtual Machine accessing an RDM hard disk using TimeFinder SnapVX. A snapshot of the VM is created and linked to a target device. The target is then presented to a Secondary ESXi server on which the replica VM will be powered-on.

VM on Primary ESXi Server

The screenshot displays the vSphere Web Client interface for a VM named **RDM_VM**. The **Summary** tab is selected, showing the following details:


- General:**
 - Guest OS: CentOS 4/5/6 (64-bit)
 - VM Version: vmx-09
 - CPU: 1 vCPU
 - Memory: 2048 MB
 - Memory Overhead: 30.25 MB
 - VMware Tools: Not running (Not installed)
 - IP Addresses:
 - DNS Name:
 - State: Powered On
 - Host: PrimaryESX064
 - Active Tasks:
 - vSphere HA Protection: N/A
- Resources:**
 - Consumed Host CPU: 9 MHz
 - Consumed Host Memory: 2070.00 MB
 - Active Guest Memory: 0.00 MB
 - Provisioned Storage: 12.09 GB
 - Not-shared Storage: 12.09 GB
 - Used Storage: 12.09 GB
- Storage:**

Storage	Drive Type	Capacity	1:
datastore1	Non-SSD	15.00 GB	1:
- Commands:**
 - Power Off
 - Suspend
 - Reset
 - Edit Settings
 - Open Console




Summary information for RDM_VM shows that it is using datastore1 for its storage.

VM – Hard Disk Is an RDM



Examining the properties of the VM shows that the hard disk is an RDM in physical compatibility mode. It is only the mapping file that is stored on datastore1 along with other files that define the virtual machine.

Correlating Primary RDM with Symmetrix Device



We use the vSphere Web Client to correlate the RDM with the Symmetrix device. The EMC Storage Viewer view shows that it is device 091. This will be the source for our SnapVX snapshot.

Primary ESXi Server – Browse Datastore

The screenshot displays the vSphere Client interface for a Primary ESXi Server (ESXi, 5.1.0, 799733). The 'Configuration' tab is selected, showing the 'Datastores' section. A table lists the available datastores:

Identification	Device	Drive Type	Capacity
datastore1	Local VMware Disk...	Non-SSD	15
MEC Share /read	10.127.108.130/...	Unknown	246
		Non-SSD	


The 'Datastore Browser - [datastore1]' window is open, showing the contents of the 'RDM_VM' folder. The files listed are:

Name	Size	Type	Path	Modified
RDM_VM.vmx	2.12 KB	Virtual Machine	[datastore1] RDM_VM	2/4/20
RDM_VM.vmdk	10,487,040.00 K	Virtual Disk	[datastore1] RDM_VM	2/4/20
RDM_VM.nvram	8.48 KB	Non-volatile me...	[datastore1] RDM_VM	2/4/20
vmware-1.log	95.62 KB	Virtual Machine ...	[datastore1] RDM_VM	2/4/20
vmware.log	129.68 KB	Virtual Machine ...	[datastore1] RDM_VM	2/9/20
RDM_VM.vmxof	0.25 KB	File	[datastore1] RDM_VM	2/4/20
RDM_VM.vmsd	0.00 KB	File	[datastore1] RDM_VM	2/4/20
vmx-RDM_VM-3990187455-1....	93,184.00 KB	File	[datastore1] RDM_VM	2/4/20
RDM_VM.vmx.lck	0.00 KB	File	[datastore1] RDM_VM	2/4/20
RDM_VM.vmx<-	2.12 KB	File	[datastore1] RDM_VM	2/4/20
RDM_VM-edd56dbf.vswp	2,097,152.00 KB	File	[datastore1] RDM_VM	2/4/20

Datastore1 is on local storage. Browsing the datastore shows that it has the files that define the RDM_VM. The local storage will not be replicated by TimeFinder SnapVX. Only the RDM Symmetrix device presented to the VM will be replicated.



Primary ESXi Server – Download VMX File



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

50

We will first download the RDM_VM.vmx file. This file will be uploaded to a datastore on the Secondary ESXi Server.

Data on RDM_VM on the Primary ESXi Server



```
RDM_VM on PrimaryESXi64
File View VM

CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64

localhost login:
CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64

localhost login: root
Password:
Last login: Mon Feb  9 10:39:41 on tty1
[root@localhost ~]# ls
anaconda-ks.cfg  install.log  install.log.syslog
[root@localhost ~]# cd /
[root@localhost /]# ls
bin  dev  home  lib64  media  opt  PROD_DATA  sbin  srv  tmp  var
boot  etc  lib  lost+found  mnt  proc  root  selinux  sys  usr
[root@localhost /]# cd PROD_DATA
[root@localhost PROD_DATA]# ls
file1  file2  file3
[root@localhost PROD_DATA]# _
```

To release cursor, press CTRL + ALT



We can open a console to the RDM_VM. We have created a directory PROD_DATA and added some files to it.

Identifying LUN Presented to the Secondary ESXi Server

The screenshot shows the VMware vSphere Configuration tab for a Secondary ESXi server. The 'Configuration' tab is selected, and the 'Devices' view is active. A table lists the devices, with the first row highlighted. The 'Name' column contains 'EMC iSCSI Disk (naa.60000970000196800225533030303930)', which is circled. The 'Identifier' column shows 'naa.600009700001...'. Below the table, the 'Device Details' section shows the 'EMC iSCSI Disk (naa.60000970000196800225533030303930)' with its location, type, owner, and transport details.

Name	Identifier	Runtime Name
EMC iSCSI Disk (naa.60000970000196800225533030303930)	naa.600009700001...	vmhba33:0
EMC iSCSI Disk (naa.60000970000196800225533030303934)	naa.600009700001...	vmhba33:1
Local VMware Disk (mpx.vmhba1:C0:T0:L0)	mpx.vmhba1:C0:T0...	vmhba1:0
Local NECVMWar CD-ROM (mpx.vmhba32:C0:T0:L0)	mpx.vmhba32:C0:T0...	vmhba32:0
EMC iSCSI Disk (naa.60000970000196800225533030303631)	naa.600009700001...	vmhba33:2
EMC iSCSI Disk (naa.60000970000196800225533030303632)	naa.600009700001...	vmhba33:3

Device Details

EMC iSCSI Disk (naa.60000970000196800225533030303930)

Location: /vmfs/devices/disks/naa.60000970000196800225533030303930
 Type: disk
 Owner: NMP
 ID: naa.60000970000196800225533030303930
 Capacity: 10.00 GB
 Partition Format: Unknown
 Transport: iSCSI



From the Configuration tab for the Secondary ESXi server, we identify the naa number of the device. We can then correlate it with the WWN displayed in Unisphere for VMAX.

Device Detail in Unisphere for VMAX

The screenshot shows the EMC Unisphere for VMAX interface. The breadcrumb navigation path is: 000196800225 > Hosts > Masking View > secondaryesxi65_mv > Storage Group > secondaryesxi65_parent. The 'PROPERTIES' table lists various device attributes. Three rows are highlighted with red boxes: Symmetrix ID (000196800225), Symmetrix Volume ID (00090), and WWN (60000970000196800225533030303930).

EMC Unisphere for VMAX V8.0.1.7	
000196800225 ▾ Storage Hosts Data Protection Performance	
000196800225 > Hosts > Masking View > secondaryesxi65_mv > Storage Group > secondaryesxi65_parent	
PROPERTIES	
Name	00090
Physical Name	N/A
Volume Identifier	N/A
Type	TDEV
Encapsulated Volume	No
Encapsulated WWN	N/A
Status	Ready
Reserved	No
Capacity (GB)	10
Capacity (MB)	10241
Capacity (Cylinder)	5462
Emulation	FBA
Symmetrix ID	000196800225
Symmetrix Volume ID	00090
HP Identifier Name	N/A
VMS Identifier Name	N/A
Nice Name	N/A
WWN	60000970000196800225533030303930

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

EMC²

53

Unisphere for VMAX shows that the Symmetrix Volume is 090 and it is in SID:225. We will use this device as the target to link the snapshot of the Primary LUN.

Create Snapshot and Link to Target

```
C:\>symsnapvx -sid 225 establish -name rdm_backup -sg primaryesxi64_rdm
```

```
Execute Establish operation for Storage Group primaryesxi64_rdm (y/[n]) ? y
```

```
Establish operation execution is in progress for the storage group primaryesxi64_rdm. Please wait...
```

```
  Polling for Establish.....Started.
  Polling for Establish.....Done.
  Polling for Activate.....Started.
  Polling for Activate.....Done.
```

```
Establish operation successfully executed for the storage group primaryesxi64_rdm
```

```
C:\>symsnapvx -sid 225 link -snapshot_name rdm_backup -sg primaryesxi64_rdm -lnsg secondaryesxi65_rdm_snap_tgt
```

```
Execute Link operation for Storage Group primaryesxi64_rdm (y/[n]) ? y
```

```
Link operation execution is in progress for the storage group primaryesxi64_rdm. Please wait...
```

```
  Polling for Link.....Started.
  Polling for Link.....Done.
```

```
Link operation successfully executed for the storage group primaryesxi64_rdm
```

© Copyright 2015 EMC Corporation. All rights reserved.


Module: TimeFinder SnapVX Operations



54


We create a snapshot named `rdm_backup`. We then link this snapshot to the Storage Group that contains the Target LUN.

Rescan Secondary ESXi Server




Next we rescan the Secondary ESXi server for all storage. This will refresh the information about the linked Target.

Upload VMX File to Secondary ESXi Server




As mentioned earlier, the local datastore of the Primary ESXi server is not replicated using SnapVX. We will upload the RDM_VM.vmx file to a datastore on the Secondary ESXi Server. We have created a folder Backup_RDM_VM on the datastore and we upload the file to it. Note that these steps must be performed only the first time the VM is powered-on on the Secondary ESXi server.

Add VM to Inventory of Secondary ESXi Server




From the Datastore Browser, select the VMX file and click Add to Inventory.

Edit VM Settings – Remove Existing Hard Disk



Next we edit the VM settings and remove the existing hard disk. This is because the definition of the hard disk has been replicated from the Primary RDM. We have to point the VM on the Secondary ESXi server to the linked target RDM.

Add Linked Target RDM to the VM (1)



We choose Raw Device Mappings and select the linked Target we had identified earlier using the naa number.

Add Linked Target RDM to the VM (2)

The screenshot displays the 'Add Hardware' wizard in a VMware environment, specifically the 'Ready to Complete' step. The wizard is titled 'Add Hardware' and has a close button (X) in the top right corner. The main heading is 'Select a Datastore', with the instruction 'Onto which datastore do you want to map this LUN?'. Below this, there are two radio buttons: 'Store with Virtual Machine' (which is selected and highlighted with a black box) and 'Specify'. To the left of the radio buttons, there is a list of links: 'Device Type', 'Select a Disk', 'Select Target LUN', 'Select Datastore', 'Compatibility Mode', 'Advanced Options', and 'Ready to Complete'. Below the links, the 'Datastore' field is populated with 'datastore1' and 'snap-7e6f'. The 'Ready to Complete' section on the right contains the text 'Review the selected options and click Finish to add the hardware.' and a table of options:

Options:	
Hardware type:	Hard Disk
Create disk:	Use mapped system LUN
Virtual Device Node:	SCSI (0:0)
Disk mode:	Persistent
Target LUN:	EMC iSCSI Disk (naa.60000970000196800225533030303930)
Compatibility mode:	Physical
Mapped datastore:	Store with VM

At the bottom right of the wizard, there is an 'EMC' logo. The footer of the slide contains the text '© Copyright 2015 EMC Corporation. All rights reserved.', 'Module: TimeFinder SnapVX Operations', and the page number '60'.

Choose Store with Virtual Machine and finish the process.

Storage Viewer View of RDM_VM on Secondary ESXi Server

The screenshot shows the VMware vSphere Web Client interface. On the left, a tree view shows the hierarchy: vCenter > VCvApp79 > Franklin > Primary > 10.126.88.64 > RDM_VM. The main pane shows the 'RDM_VM' details, with the 'Monitor' tab selected. Under the 'Monitor' tab, the 'EMC Storage Viewer' sub-tab is active. It displays a table with the following data:

Name	Capacity	Location	Type
Hard disk 1	10.00 GB	datastore1 (1) (14.1 GB free)	RDM

Below the table, the 'Storage Details: Hard disk 1' section shows that details are hosted on EMC VNX datastores and are powered by EMC. A detailed table of properties is shown below:

Available Capacity	10.00 GB
Used Capacity	0 bytes
Block Size	512 bytes
Is Thin Provisioned	true
WWN	60000970000196800225533030303930
Status	ONLINE
Type	TDEV
Name	00090
Volume Identifier	N/A

We can use the vSphere Web Client to verify that the linked Target 090 is indeed presented as an RDM to RDM_VM on the Secondary ESXi server.




© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

61

Power on VM on the Secondary ESXi Server




© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

62

As in the case of VMFS, we choose "I copied it" for the Virtual Machine Message when we power-on the VM on the Secondary ESXi server.

Data Accessible to VM on Secondary ESXi Server



```
RDM_VM on SecondaryESXi65
File View VM

CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64

localhost login: root
Password:
Last login: Mon Feb  9 10:58:19 on tty1
[root@localhost ~]# cd /
[root@localhost /]# ls
bin  dev  home  lib64  media  opt  PROD_DATA  sbin  srv  tmp  var
boot  etc  lib  lost+found  mnt  proc  root  selinux  sys  usr
[root@localhost /]# cd PROD_DATA
[root@localhost PROD_DATA]# ls
file1  file2  file3
[root@localhost PROD_DATA]# _
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

63

We can open a console to the VM on the Secondary ESXi server and verify that it has access to the point in time data from the snapshot.

Lesson Summary

During this lesson the following topics were covered:

- Using TimeFinder SnapVX to replicate a VM accessing RDM hard disk
- Presenting the replicated VM to a Secondary ESXi Server



This lesson covered replicating a Virtual Machine accessing an RDM hard disk using TimeFinder SnapVX. A snapshot of the VM was created and linked to a target device. The target was then presented to a Secondary ESXi server on which the replica VM was powered-on.

Lesson: TimeFinder SnapVX Operations using Unisphere for VMAX

This lesson covers the following topics:

- Performing TimeFinder SnapVX operations using Unisphere for VMAX



This lesson covers performing TimeFinder SnapVX operations using Unisphere for VMAX.

Create Storage Group

Provision Storage

1 Create Storage

Storage Group Name: uni_snapvx

Storage Resource Pool: None

Service Level: None

Workload Type: None

Volumes:

Volume Capacity: GB

Avg. Response Time: System Optimized

Add Service Level

Total Capacity: 0 GB | Total Service Levels: 1

Set Host I/O Limits

< Back Next > Add to Job List Cancel Help



© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

66

TimeFinder Snapvpx operations are performed on Storage Groups in Unisphere for VMAX. Unisphere for VMAX does not support Device Group or Device Files for SnapVX operations. In our example, the number of devices are already in a Storage Group and Masking view for host access. We want to take a snapshot of just one of the devices. So we have to create a new Storage Group.

First we navigate to Array ID > Storage > Storage Group Dashboard > Storage Groups. From this page we select Create. This launches the Provision Storage wizard shown in the slide. We give the SG the name uni_snapvx. As the device is in another SG which is managed by FAST, we must select "None" for Storage Resource Pool as well as for Service Level. We can then select "Run Now". This will create an empty Storage Group named uni_snapvx.

Add Device to Storage Group (1)

000196700483 > Storage > Storage Groups Dashboard > uni_snapvx > Volumes

1 Find Volumes

Find volumes that match the following criteria

Capacity equal to GB

Volume ID (e.g. 001 or 001-0FF or 001,003-07F)

Volume Identifier Name

Additional Criteria

Select Category Attribute is equal to Value

Add Another

Clear All


☒ Include Volumes in Storage Groups

< Modify Criteria **Find Volumes >** Add to SG Cancel

© Copyright 2015 EMC Corporation. All rights reserved. Module: TimeFinder SnapVX Operations 67

Navigate to Array ID > Storage > Storage Groups Dashboard > uni_snapvx > Volumes and select "Add Volumes to SG". In the "Add Volumes to a Storage Group" wizard, we specify the device we want to add to the SG – 04A in our example. As noted in the previous slide, this device belongs to another SG. So we must select "Include Volumes in Storage Groups". Then select "Find Volumes".

Add Device to Storage Group (2)



Select the device and then select "Add to SG".

Create Snapshot (1)

000196700483 > Data Protection > Protection Dashboard > Unprotected Storage Groups

Storage Group	Capacity (GB)
▶ app_server_parent	40
▶ esxi-88-36_parent	0.04
▶ esxi-88-44_parent	12.08
▶ esxi-88-46_parent	0.04
▶ remotesxi74_parent	20.01
snapvx	2
sun-88-41-sg	0.04
tgt2vx	2
tgtvx	2
uni_snapvx	2
▶ vcenter-88-20_parent	0.04
▶ w2k8r2-88-72_parent	24.05
▶ w2k8r2-88-73_parent	12.04

Protect

Manage Storage Groups

EMC¹


© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

69

Navigate to Unprotected Storage Groups; select the Storage Group for which a snapshot should be created and then select "Protect". This will launch the Protect Storage Group wizard.

Create Snapshot (2)



Select "Point In Time Protection" and then select "Using SnapVX". Click "Next".

Create Snapshot (3)

1 Select Protection Type

2 Snapshot Details

3 Finish

Create New Snapshot

Reuse Existing Snapshot

Show Advanced >>

Name

backup

< Back Next > Add to Job List Cancel Help

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

71

We select "Create New Snapshot" and give it the name of backup. We can use the "Show Advanced" option to set TTL if required. Select "Next".

Create Snapshot (4)

1 Select Protection Type

2 Snapshot Details

3 Finish

Snapshot name: backup

Days to live: Manual

000196700483 > Data Protection > TimeFinder > TimeFinder/SnapVX > uni_snapvx

Snapshot	Creation Time	Linked	Restored
backup	Feb 23, 2015 8:48:58 AM (0)	--	--

< Back Next > Add to Job List Close Help

From the drop-down select "Run Now". This creates the snapshot as can be seen in the list shown.



Link Snapshot to Target (1)

000196700483 > Data Protection > TimeFinder > TimeFinder/SnapVX > uni_snapvx

Snapshot	Creation Time	Linked	Restored	Tracks	Non-Shared Tracks
backup	Feb 23, 2015 8:48:58 AM (0)			0	0

Create Snapshot Link Restore View Details >>

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

73

Select the snapshot and then select “Link”.

Link Snapshot to Target (2)

The screenshot shows the 'SnapVX - Link to Storage Group - backup' wizard. The 'Select existing target storage group' option is selected. A list of storage groups is displayed, with 'uni_tgtvx' highlighted. A success message dialog is also visible.

SnapVX - Link to Storage Group - backup

☐ New storage group target name

☒ Select existing target storage gr...

*** Storage Groups**

- uni_tgtvx
- snapvx
- tgtvx
- uni_tgtvx**

<< Hide Advanced

☐ Copy

☐ Force

SnapVX - Link to Storage Group - backup

Command successfully executed

Close

Add to Job List Cancel Help

EMC

© Copyright 2015 EMC Corporation. All rights reserved. Module: TimeFinder SnapVX Operations 74

The target device should also be in a Storage Group to perform the Link operation. For our example we have created a Storage Group named uni_tgtvx and added the target device to it. The procedure is the same as discussed earlier to create SG for the source device.

We select the "Select existing target storage group" in the wizard. This lists the candidate target SGs for the link operation. Note that if a Storage Group is a part of a Masking View, it will not be displayed in the list. This is to ensure that users do not accidentally select devices that are actively in use and corrupt the data. We can select "Run Now" from the drop-down to link the snapshot to the target.

Other SnapVX Operations

000196700483 > Data Protection > TimeFinder > TimeFinder/SnapVX > uni_snapvx

Snapshot	Creation Time	Linked	Restored
backup	Feb 23, 2015 8:48:58 AM (0)	✓	-

Create SnapshotLinkRestoreView Details>>

Relink

Unlink

Set Mode

Set Days to Live

Terminate

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

EMC²

75

The list now shows that the snapshot has linked target. Other SnapVX operations can be performed by selecting the snapshot name and then selecting ">>". From this list, we can choose the operation we want to perform.

Lesson Summary

During this lesson the following topics were covered:

- Performing TimeFinder SnapVX operations using Unisphere for VMAX



This lesson covered performing TimeFinder SnapVX operations using Unisphere for VMAX.

Lab:TimeFinder SnapVX Operations

This lab covers

- Creating Snapshots
- Accessing Snapshot data from a Secondary host
- Restoring to Source from Snapshots
- Restoring to Source from modified Target device

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

77

This lab covers creating and linking TimeFinder SnapVX snapshots to Target devices. It also covers restoring snapshot data to the source device as well as restoring modified target data back to the source device.

Lab Summary

This lab covered

- Creating TimeFinder SnapVX snapshots
- Linking the snapshot to a Target device and accessing Point-in-time snapshot data from a Secondary host
- Restoring a snapshot to the Source device
- Restoring modified Target data to the Source device

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

78

This lab covered creating and linking TimeFinder SnapVX snapshots to Target devices. It also covered restoring snapshot data to the source device as well as restoring modified target data back to the source device.

.

Lab:TimeFinder SnapVX Replication of a VMFS Datastore

This lab covers

- Identifying and correlating Source and Target devices with devices accessible by Primary and Secondary ESXi servers
- Creating a VMFS Datastore on the Source device and deploying a Virtual Machine
- Using Unisphere for VMAX to create a SnapVX snapshot of the Datastore and link the snapshot to Target
- Accessing the linked Target from the Secondary ESXi server and power-on Virtual Machine on the Secondary ESXi server

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

79

This lab covers performing TimeFinder SnapVX replication of a VMFS Datastore using Unisphere for VMAX and VMware vSphere client.

Lab Summary

This lab covered

- Identifying and correlating Source and Target devices with devices accessible by Primary and Secondary ESXi servers
- Creating a VMFS Datastore on the Source device and deploying a Virtual Machine
- Using Unisphere for VMAX to create a SnapVX snapshot of the Datastore and link the snapshot to Target
- Accessing the linked Target from the Secondary ESXi server and power-on Virtual Machine on the Secondary ESXi server

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: TimeFinder SnapVX Operations

80

This lab covered performing TimeFinder SnapVX replication of a VMFS Datastore using Unisphere for VMAX and VMware vSphere client.

Module Summary

Key points covered in this module:

- TimeFinder SnapVX concepts
- Performing TimeFinder SnapVX operations using SYMCLI and Unisphere for VMAX
- Replicating a VMFS datastore using TimeFinder SnapVX
- Replicating a VM accessing RDM hard disk using TimeFinder SnapVX



This module covered TimeFinder SnapVX local replication technology on VMAX3 arrays. Concepts, terminology, and operational details of creating snapshots and presenting them to target hosts were discussed. Performing TimeFinder SnapVX operations using SYMCLI and Unisphere for VMAX were presented. Use of TimeFinder SnapVX for replication in a virtualized environment was also presented.

This slide intentionally left blank.

Module: SRDF/Synchronous Operations

Upon completion of this module, you should be able to:

- Create Dynamic SRDF Groups and Dynamic SRDF Pairs
- Perform SRDF/S operations using SYMCLI and Unisphere for VMAX
- Perform Disaster Recovery of a VMFS datastore
- Perform Disaster Recovery of a VM accessing an RDM hard disk



This module focuses on SRDF operations in synchronous mode. Use of SYMCLI and Unisphere for VMAX to perform SRDF operations are presented in detail. Methods for performing DR operations in a virtualized environment for both VMFS datastore and RDM use cases are discussed.

Lesson: SRDF Initial Setup Operations

This lesson covers the following topics:

- Listing the SRDF environment
- Creating Dynamic RDF Groups
- Creating Dynamic RDF Pairs
- Grouping SRDF devices into SYMCLI device groups
- Changing SRDF mode, suspending and resuming RDF links

This lesson covers initial SRDF setup operations. Creating dynamic RDF groups and RDF pairs using SYMCLI is presented in detail. Basic SRDF operations are also discussed.

Dynamic SRDF Groups

```
C:\>symcfg list
```

S Y M M E T R I X

SymmID	Attachment	Model	Mcode Version	Cache Size (MB)	Num Phys Devices	Num Symm Devices
000196700483	Local	VMAX200K	5977	202752	18	882
000196800225	Remote	VMAX100K	5977	220160	0	1079

```
C:\>symcfg list -sid 483 -v
```

----Output Truncated-----

```
Switched RDF Configuration State      : Enabled
Concurrent RDF Configuration State    : Enabled
Dynamic RDF Configuration State      : Enabled
Concurrent Dynamic RDF Configuration : Enabled
```

----Output Truncated-----

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

3

SRDF groups define the relationships between the local SRDF director/ports and the corresponding remote SRDF director/ports. Any device that has been configured as an SRDF device must be assigned to an SRDF group. Storage Administrators can dynamically create SRDF groups and assign them to Fibre Channel directors or GbE directors. SRDF groups are also referenced as RA groups or RDF groups. Dynamic RDF Configuration State is Enabled by default.

In this example, we have a pair of VMAX3 arrays (SID:483 and SID:225) that are SRDF connected. The command has been executed from a host attached to SID:483. So from this perspective, SID:483 is the Local VMAX3 array and SID:225 is the Remote VMAX3 array. The Num Phys Devices column indicates that the host from which the command was executed has physical access to 18 devices on SID:483 and none on the other array. The Num Symm Devices column indicates the total number of devices that have been configured on the respective VMAX3 arrays.

A verbose listing of SID:483 shows that Dynamic RDF Configuration State is Enabled. This should be verified for SID:225 as well.

Dynamic RDF Configuration State is Enabled by default on VMAX3 arrays. The combination of the ability to dynamically create SRDF groups and dynamic device pairs enables you to create, delete, and swap SRDF R1-R2 pairs.

VMAX3 arrays will support up to 250 SRDF groups per array. VMAX3 arrays will support 250 SRDF groups per port.

List Remote Adapters

```
C:\>symcfg list -sid 483 -ra all
```

```
Symmetrix ID: 000196700483 (Local)
```

```
S Y M M E T R I X      R D F      D I R E C T O R S
```

Ident	Port	Remote SymmID	Local RA Grp	Remote RA Grp	Status Dir	Port
RF-1E	8	000196800225	1 (00)	1 (00)	Online	Online
RF-3E	8	000196800225	1 (00)	1 (00)	Online	Online



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

4

Before creating a new SRDF group, some information must be gathered. First you need to know the SRDF directors (Remote Adapters – RA) that are configured on the array. You also need to know the number of SRDF groups (RA groups) currently configured and their corresponding group numbers.

The listing shows that you have a pair of Remote Adapters (RF-1E and RF-3E) available. These are Fibre Remote Adapters (RF). Both the Remote Adapters are online. There is one SRDF Group that has been configured to use this pair of directors.

List Configured SRDF Groups

```
C:\>symcfg -sid 483 list -rdfg all
```

```
Symmetrix ID : 000196700483
```

S Y M M E T R I X R D F G R O U P S

Local			Remote		Group			RDFA Info			
-----			-----		-----			-----			
RA-Grp	LL (sec)	RA-Grp	SymmID	T	Name	Flags LPDS CHT	Dir Cfg	Flags CSRM	Cycle time	Pri	
-----			-----		-----			-----			
1 (0)	10	1 (0)	000196800225	D	Group_1	XX.. ..X	F-S	-IS-	15	33	

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

5

The command shown gives detailed information on the currently configured SRDF Groups. You can see that it is a Dynamic RDF Group named Group_1. The director configuration is Fibre Channel-Switched. We will look at the Group and RDFA Flag details in the SRDF/A module.

Legend:

Group (T)ype : S = Static, **D = Dynamic**

Director (C)onfig : **F-S = Fibre-Switched**, F-H = Fibre-Hub

G = GIGE, E = ESCON, T = T3, - = N/A

Display RDF Connectivity – symsan Command

```
C:\>symsan list -sid 483 -sanrdf -dir all
```

```
Symmetrix ID: 000196700483
```

Flags				Remote		
Dir:P	CS	S	Lnk	Symmetrix ID	Dir:P	WWN
01E:08	SO	O	C	000196800225	01E:07	5000097358038407
03E:08	SO	O	C	000196800225	02E:07	5000097358038447

Legend:

Director:

(C)onfig : **S = Fibre-Switched**, H = Fibre-Hub

G = GIGE, - = N/A

(S)tatus : **O = Online**, F = Offline, D = Dead, - = N/A

Port:

(S)tatus : O = Online, F = Offline, - = N/A

Link:

(S)tatus : C = Connected, P = ConnectInProg

D = Disconnected, I = Incomplete, - = N/A

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

6

The `symsan` command can be used to discover and display RDF connectivity between the VMAX3 arrays. This command helps in situations where RDF groups have not yet been created between the VMAX3 pairs.

In such cases where there are no RDF groups configured, the `symcfg` command in the previous slide will not return any connectivity information. The `symsan` command is particularly useful to determine the local and remote RDF directors, as well as the full serial number of the remote array. The full serial number of the remote array is required to create the first Dynamic RDF group. Subsequent RDF groups can be created by just specifying the last two digits of the remote array.

Creating a Dynamic RDF Group

```
C:\>symrdf addgrp -label SRDF_Sync -sid 483 -remote_sid 225 -dir 1E:8,3E:8 -remote_dir 1E:7,2E:7 -rdfg 10 -remote_rdfg 10
```

```
Execute a Dynamic RDF Addgrp operation for group 'SRDF_Sync' on Symm: 000196700483 (y/[n]) ? y
```

```
Successfully Added Dynamic RDF Group 'SRDF_Sync' for Symm: 000196700483
```

```
C:\>symcfg list -sid 483 -rdfg all
```

```
Symmetrix ID : 000196700483
```

S Y M M E T R I X R D F G R O U P S

Local		Remote		Group			RDFA Info		
LL									
RA-Grp	(sec)	RA-Grp	SymmID	T	Name	Flags LPDS CHT	Dir Cfg	Flags CSRM	Cycle time Pri
1 (0)	10	1 (0)	000196800225	D	Group_1	XX.. ..X	F-S	-IS-	15 33
10 (9)	10	10 (9)	000196800225	D	SRDF_Sync	XX.. ..X	F-S	-IS-	15 33



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

7

The `symrdf addgrp` command creates an empty Dynamic SRDF group on the source and the target Symmetrix and logically links them. The directors and the respective ports have to be specified in the command for VMAX3 arrays.

Note: The physical links and communication between the two arrays must exist for this command to succeed.

Note: The SRDF group number in the command (`-rdfg` and `-remote_rdfg`) is in decimal.

In the Symmetrix, it is converted to hexadecimal. The decimal group numbers start at 01 but the hexadecimal group numbers start at 00. Hence the hexadecimal group numbers will be off by one. We have created an RDF Group with the label `SRDF_Sync` and the RDF Group number 10 in decimal. Shown in parenthesis is the hexadecimal value 9. It would be convenient if the SRDF group numbers on the local and the remote arrays were identical, however, this is not a requirement.

Creating RDF Device Pairs

```
C:\>symrdf createpair -sid 483 -f pairs.txt -type R1 -rdfg 10 -establish
```

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

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

```
Create RDF Pair in (0483,010).....Started.  
Create RDF Pair in (0483,010).....Done.  
Mark target device(s) in (0483,010) for full copy from source....Started.  
Devices: 0059-005A in (0483,010).....Marked.  
Mark target device(s) in (0483,010) for full copy from source....Done.  
Merge track tables between source and target in (0483,010).....Started.  
Devices: 0059-005A in (0483,010).....Merged.  
Merge track tables between source and target in (0483,010).....Done.  
Resume RDF link(s) for device(s) in (0483,010).....Started.  
Resume RDF link(s) for device(s) in (0483,010).....Done.
```

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



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

8

The `symrdf createpair` command takes the dynamic capable device pairs listed in the text file (`pairs.txt`) and makes them R1-R2 pairs. On VMAX3 arrays, devices are created as Dynamic capable by default. By specifying `-establish` the newly created R2 devices are synchronized with the data from the newly created R1 devices. In this example the file contains the following entries:

`pairs.txt`

```
059 087
```

```
05A 088
```

The first column in the file lists the devices on the VMAX3 on which the command is executed (in our example `SID:483`) and the second column is the remote VMAX3 (in our example `SID:225`). Specifying `-type R1` makes the device in the first column to be R1s and the devices in the second column will become their corresponding R2s. The mode of operation for the SRDF pairs that are newly created is set to Adaptive Copy Disk Mode (discussed later in this module) by default.

Device Pairs Created

```
C:\>symrdf -sid 483 -f pairs.txt -rdfg 10 query
```

```
Symmetrix ID           : 000196700483 (Microcode Version: 5977)
Remote Symmetrix ID    : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number  : 10 (09)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
Standard		ST			LI		ST			
		A			N		A			
Logical	Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv			RDF Pair
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks			MACE STATE
-----					-----					-----
N/A	00059	RW	0	0	RW	00087	WD	0	0	CD.. Synchronized
N/A	0005A	RW	0	0	RW	00088	WD	0	0	CD.. Synchronized
-----					-----					-----
Total										
Track(s)			0	0				0	0	
MB(s)			0.0	0.0				0.0	0.0	

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC²

9

As noted earlier, the SRDF mode is by default set to Adaptive Copy Disk. The establish operation synchronizes data from the new R1 device to the new R2 device. The R1 devices are created on SID:483 and their corresponding R2 devices are created on SID:225. R1 device 059 is paired with R2 device 087. R1 device 05A is paired with R2 device 088.

Legend for MODE:

M(ode of Operation) : A = Async, S = Sync, E = Semi-sync, **C** = **Adaptive Copy**
: M = Mixed

A(daptive Copy) : **D** = **Disk Mode**, W = WP Mode, . = ACp off

C(onsistency State) : X = Enabled, . = Disabled, M = Mixed, - = N/A

(Consistency) **E**(xempt): X = Enabled, . = Disabled, M = Mixed, - = N/A

Deleting Device Pairing

- Removes the pairing information from the Symmetrix
- Must suspend RDF links before issuing `symrdf deletepair` command
 - Link state must be NR and pair state is Suspended, Split, or FailedOver
- Canceling dynamic SRDF pairings changes the type of the device group from R1 or R2 to Regular
- Devices in the device group are changed from RDF devices to RDF-capable standard devices and the SYMAPI database is updated



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

10

The `delete SRDF pairs` command cancels SRDF pairs in the device file specified. Before the `deletepair` can be invoked, the pair must be suspended first. The RDF Group is not deleted by this operation. If the RDF Group should be deleted, then the `symrdf removegrp` command should be used after the `deletepair` operation.

Example:

```
c:\symrdf suspend -sid 97 -file grp5.txt -rdfig 5
```

```
c:\symrdf deletepair -sid 97 -file grp5.txt -rdfig 5
```

Identify SRDF Devices Accessible to the Host (1)

```
C:\>symrdf list pd
```

```
Symmetrix ID: 000196700483
```

Local Device View											
Sym Dev	Sym RDev	RDF Typ:G	STATUS			MODES		R1 Inv Tracks	R2 Inv Tracks	RDF S T A T E S	
			SA	RA	LNK	MDATE				Dev	RDev
00059	00087	R1:10	RW	RW	RW	C.D1.		0	0	RW	WD
0005A	00088	R1:10	RW	RW	RW	C.D1.		0	0	RW	WD
Total											
Track(s)								0	0		
MB(s)								0.0	0.0		

Synchronized
Synchronized

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

11

The SYMCLI command `symrdf list pd` gives a list of all SRDF devices accessible to the host. In this example the host has access to 2 SRDF devices (059:05A). As can be seen under the RDF Typ:G column, the devices are type R1 and they have been created in SRDF Group 10. The mode of SRDF operation for these pairs is Adaptive Copy Disk Mode, and currently all the R1-R2 pairs are in a Synchronized state. The local R1 devices (the Sym Dev column of the output) 059:05A are paired with corresponding R2 devices (the RDev column of the output) 087:088.

Legend for MODES:

M(ode of Operation) : A = Async, S = Sync, E = Semi-sync, C = Adaptive Copy
: M = Mixed

D(omino) : X = Enabled, . = Disabled

A(daptive Copy) : D = Disk Mode, W = WP Mode, . = ACp off

(Mirror) T(ype) : 1 = R1, 2 = R2

(Consistency) E(xempt): X = Enabled, . = Disabled, M = Mixed, - = N/A

Identify SRDF Devices Accessible to the Host (2)

```
C:\>sympd list
```

```
Symmetrix ID: 000196700483
```

Device Name	Dir		Device				Cap
Physical	Sym	SA :P	Config	Attribute	Sts	(MB)	
\\.\PHYSICALDRIVE1	00031	03D:028	TDEV	N/Grp'd	RW	6	
\\.\PHYSICALDRIVE2	00032	03D:028	TDEV	N/Grp'd	RW	6	
\\.\PHYSICALDRIVE3	00033	03D:028	TDEV	N/Grp'd	RW	6	
\\.\PHYSICALDRIVE4	00034	03D:028	TDEV	N/Grp'd	RW	6	
\\.\PHYSICALDRIVE5	00035	03D:028	TDEV	N/Grp'd	RW	6	
\\.\PHYSICALDRIVE6	00036	03D:028	TDEV	N/Grp'd	RW	6	
\\.\PHYSICALDRIVE7	00049	03D:028	TDEV	N/Grp'd	RW	2049	
\\.\PHYSICALDRIVE8	0004A	03D:028	TDEV	N/Grp'd	RW	2049	
\\.\PHYSICALDRIVE9	0004B	03D:028	TDEV	N/Grp'd	RW	2049	
\\.\PHYSICALDRIVE10	0004C	03D:028	TDEV	N/Grp'd	RW	2049	
\\.\PHYSICALDRIVE11	0004D	03D:028	TDEV	N/Grp'd	RW	2049	
\\.\PHYSICALDRIVE12	0004E	03D:028	TDEV	N/Grp'd	RW	2049	
\\.\PHYSICALDRIVE13	00059	01D:028	RDF1+TDEV	N/Grp'd	RW	2049	
\\.\PHYSICALDRIVE14	0005A	01D:028	RDF1+TDEV	N/Grp'd	RW	2049	

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC²

12

The `sympd list` command gives the list of all the devices that the host can access on the array. This command is used to correlate the host physical device name with the Symmetrix device number. We see that the host addresses the R1 devices as `PHYSICALDRIVE13` and `PHYSICALDRIVE14`. To format the devices, create partitions, create and mount file systems (Volume management operations), the host's physical device names should be used.

Creating SRDF Device Groups – SYMCLI

- All devices in a device group must be in the same Symmetrix array
- All devices must be of the same type (RDF1, RDF2, RDF21, or Regular)

```
C:\>symdg create -type R1 syncdg
```

```
C:\>set SYMCLI_DG=syncdg
```

```
C:\>symdg addall dev -range 059:05A
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

13

A device group is a user created object for viewing and managing related Symmetrix devices. All devices in a device group should be on the same Symmetrix array. There are four types of device groups: RDF1, RDF2, RDF21, or REGULAR. When creating a device group, if a type is not specified explicitly, by default a device group of type REGULAR will be created. A device group, with the type REGULAR cannot contain SRDF devices; a device group of type RDF1 cannot contain R2 devices; and a device group of type RDF2 cannot contain R1 devices. When performing SRDF/S operations, SYMCLI commands can be executed for ALL devices in the device group or a subset of them. For SRDF/A operations, the commands should be executed for ALL devices in the SRDF Group.

Storage Administrators must create a device group with RDF1 or RDF2 for SRDF operations, as appropriate. In this example device group type R1 (RDF1) is created, so that the R1 devices 059 and 05A can be added to it. Note that the environment variable SYMCLI_DG has been set to the device group that was created. When this variable is set, subsequent commands to manage the device group will not need the -g <device_group_name> flag in the command syntax.

By default the device group definition is stored on the host where the `symdg create` command was executed. To manage the device group from other hosts connected to the same Symmetrix array, GNS daemon should be used.

Device Group Details (1)

```
C:\>syndg show syncdg|more
```

```
Group Name: syncdg
```

Group Type	: RDF1	(RDFA)
Device Group in GNS	: No	
Valid	: Yes	
Symmetrix ID	: 000196700483	
Group Creation Time	: Mon Feb 16 12:09:57 2015	
Vendor ID	: EMC Corp	
Application ID	: SYMCLI	
Number of STD Devices in Group	: 2	
Number of Associated GK's	: 0	
Number of Locally-associated BCV's	: 0	
Number of Locally-associated VDEV's	: 0	
Number of Locally-associated TGT's	: 0	

```
---Output Truncated-----
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

14

The `syndg show` command displays detailed group information for any specific device group. The device group (`srdfsdg`) contains 2 local standard devices. The device group type is RDF1. The Symmetrix serial number is also displayed. If there are any BCVs, VDEVs, etc. associated with or added to the group, this information is also displayed.

Device Group Details (2)

Standard (STD) Devices (2):

{

LdevName	PdevName	Sym Dev	Device Config	Att. Sts	Cap (MB)
DEV001	\\.\PHYSICALDRIVE13	00059	RDF1+TDEV	RW	2049
DEV002	\\.\PHYSICALDRIVE14	0005A	RDF1+TDEV	RW	2049

}

---Output Truncated-----

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

15

Further information on the devices in the group, as well as relevant RDF information is also presented with the `symdg show` command. The two devices have been assigned (by default) the Logical Device Names of DEV001 and DEV002. The host (Windows OS in this example) addresses the two devices as PHYSICALDRIVE13 and PHYSICALDRIVE14.

symrdf Command Syntax

```
symrdf -g <device_group_name> <Action> [Options]
```

Actions

suspend	failback
resume	split
set mode	establish
failover	restore
update	

Users can perform a number of SRDF operations using host-based SYMCLI commands. Major SRDF operations or actions include: suspend, resume, failover, failback, update, split, establish, and restore.

Changing SRDF Mode of Operation

```
C:\>symrdf -g syncdg set mode sync -nop
```

```
C:\>symrdf -g syncdg query
```

```
Device Group (DG) Name      : syncdg
DG's Type                   : RDF1
DG's Symmetrix ID           : 000196700483 (Microcode Version: 5977)
Remote Symmetrix ID         : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number       : 10 (09)
```

Source (R1) View					Target (R2) View					MODE
ST					LI					
A					N					
Standard	Sym	T	R1 Inv	R2 Inv	K	Sym	T	R1 Inv	R2 Inv	RDF Pair
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MACE STATE
DEV001	00059	RW	0	0	RW	00087	WD	0	0	S...
DEV002	0005A	RW	0	0	RW	00088	WD	0	0	S...
Total										
Track(s)			0	0				0	0	
MB(s)			0.0	0.0				0.0	0.0	

Synchronized
Synchronized

EMC²

Module: SRDF/Synchronous Operations

17

The `symrdf set mode` command will change the SRDF operation mode. In this example the mode has been changed to Synchronous for these two R1-R2 pairs. This is indicated by the S in M column of the output. In the 'normal' operations of SRDF, the R1 device presents a Read/Write (RW) status to its host and the corresponding R2 device presents a Write-Disabled (WD) status to its host. Data written to the R1 device is sent over the links to the R2 Symmetrix. The meaning of the R1/R2 Inv(alid) Tracks will be discussed throughout this module.

Legend for MODE:

M(ode of Operation) : A = Async, **S** = **Sync**, E = Semi-sync, C = Adaptive Copy
: M = Mixed

A(daptive Copy) : D = Disk Mode, W = WP Mode, . = ACp off

C(onsistency State) : X = Enabled, . = Disabled, M = Mixed, - = N/A

(Consistency) **E**(xempt): X = Enabled, . = Disabled, M = Mixed, - = N/A

Suspending RDF Links

```
C:\>symrdf suspend -nop
```

```
C:\>symrdf query
```

```
Device Group (DG) Name      : syncdg
DG's Type                   : RDF1
DG's Symmetrix ID           : 000196700483 (Microcode Version: 5977)
Remote Symmetrix ID         : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number       : 10 (09)
```

Source (R1) View						Target (R2) View						MODE			
ST						LI		ST							
A						N		A							
Standard	Sym	T	R1	Inv	R2	Inv	K	Sym	T	R1	Inv	R2	Inv	RDF Pair	
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MACE	STATE				
DEV001	00059	RW	0	4266	NR	00087	WD	0	0	S...	Suspended				
DEV002	0005A	RW	0	4219	NR	00088	WD	0	0	S...	Suspended				
Total															
Track(s)			0	8485			0	0							
MB(s)			0.0	1060.6			0.0	0.0							

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC²

18

Suspend is a singular operation. Data transfer from the Source devices to the Target devices is stopped. The links for these devices are logically set to Not Ready (NR). This operation only affects the targeted devices in the device group. SRDF device pairs in other device groups and other SRDF Groups are not affected even if they share the same Remote Directors. Physical links and the RA communication paths are still available. New writes to the Source devices accumulate as invalid tracks to the R2 (R2 Inv Tracks). The R1s continue to be Read Write enabled and the R2s continue to be Write Disabled.

To invoke a suspend, the RDF pair(s) must already be in one of the following states:

Synchronized

R1 Updated

Resuming RDF Links

```
C:\>symrdf resume -nop
```

```
C:\>symrdf query
```

```
Device Group (DG) Name      : syncdg
DG's Type                   : RDF1
DG's Symmetrix ID           : 000196700483   (Microcode Version: 5977)
Remote Symmetrix ID         : 000196800225   (Microcode Version: 5977)
RDF (RA) Group Number       : 10 (09)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
Standard		ST			LI		ST			
		A			N		A			
Logical	Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv		RDF Pair	
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks		MACE STATE	
-----					-----					-----
DEV001	00059	RW	0	10718	RW	00087	WD	0	0 S...	SyncInProg
DEV002	0005A	RW	0	10422	RW	00088	WD	0	0 S...	SyncInProg
-----					-----					-----
Total										
Track(s)			0	21140				0	0	
MB(s)			0.0	2642.5				0.0	0.0	

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC²

19

Resume is a singular operation. To invoke this operation, the RDF pair(s) must already be in the Suspended state. Data transfer from R1 to R2 is resumed. The pair state will remain in SyncInProg until all accumulated invalid tracks for the pair have been transferred. Invalid tracks are transferred to the R2 in any order – so write serialization will not be maintained. The link is set to Read Write. The R1s continue to be Read Write enabled and the R2s continue to be Write Disabled.

Lesson Summary

During this lesson the following topics were covered:

- Listing the SRDF environment
- Creating Dynamic RDF Groups
- Creating Dynamic RDF Pairs
- Grouping SRDF devices into SYMCLI device groups
- Changing SRDF mode, suspending and resuming RDF links

This lesson covered initial SRDF setup operations. Creating dynamic RDF groups and RDF pairs using SYMCLI was presented in detail. Basic SRDF operations were also discussed.

Lesson: SRDF Disaster Recovery Operations

This lesson covers the following topics:

- SRDF Failover
- SRDF Update
- SRDF Failback



This lesson covers SRDF Disaster Recovery operations. Device and link states under different conditions are presented in detail. Host consideration when performing DR operations are also discussed.

SRDF Disaster Recovery Operations

- Failover:
 - Will make a copy of the data on target devices (R2s) available to the host accessing these devices on the target array
 - Invoked in the event of a disaster – Host, Symmetrix, or Site failure
 - Can be used for maintenance operations on the source site: Will provide data availability from the target devices, during host, Symmetrix, or Site maintenance
- Update:
 - Begins transfer of accumulated invalid tracks from the R2s to the R1s, while production work continues on the R2s
- Failback:
 - Resumes operation back on the primary host accessing data on the source devices (R1s), all changes made to the R2s when failed over will be transferred back to the source devices
 - Primary host can access the R1 devices as soon as the command completes without waiting for the data transfer to complete

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

22

The SRDF disaster recovery operations are:

- Failover from the source side to the target side, switching data processing to the target side
- Update the source side after a failover while the target side is still used for applications
- Failback from the target side to the source side by switching data processing to the source side

SRDF Failover

```
C:\>symrdf -g syncdg failover -nop
```

```
An RDF 'Failover' operation execution is  
in progress for device group 'syncdg'. Please wait...
```

```
Write Disable device(s) on SA at source (R1).....Done.  
Suspend RDF link(s).....Done.  
Read/Write Enable device(s) on RA at target (R2).....Done.
```

```
The RDF 'Failover' operation successfully executed for  
device group 'syncdg'.
```



Module: SRDF/Synchronous Operations

23

The failover operation can be executed from the source side host or the target side. This is true for all symrdf commands. In order to perform operations from the target side, a device group of type RDF2 should be created and the R2 devices should be added to it. In the event of an actual disaster, this is helpful as there would be no way of communicating with the source Symmetrix. The operation assumes disaster situation - makes all effort to enable data access on target Symmetrix:

- Will proceed if possible
- Will give message for any potential data integrity issue

As can be seen in the output, the R1 devices are write disabled, the SRDF links between the device pairs are logically suspended, and the R2 devices are read write enabled. Host accessing the R2 devices can now resume processing the application.

While in a true disaster situation when the source host/Symmetrix/site may be unreachable, the steps listed below would be recommended if performing a "graceful" failover to the target site.

If failing over for a Maintenance operation: For a clean, consistent, coherent point in time copy which can be used with minimal recovery on the target side some or all of the following steps may have to be taken on the source side:

- Stop All Applications
- Unmount file system (unmount or unassign drive letter to flush the filesystem buffers from the host memory down to the Symmetrix)
- Deactivate the Volume Group
- A failover leads to a write disabled state of the R1 devices. If a device suddenly becomes write disabled from a read/write state, the reaction of the host can be unpredictable if the device is in use. Hence the recommendation to stop applications, unmount filesystem/unassign drive letter, prior to performing a failover for maintenance operations.

Query after Failover

C:\>symrdf -g synctgtdg query

```
Device Group (DG) Name      : synctgtdg
DG's Type                  : RDF2
DG's Symmetrix ID          : 000196800225 (Microcode Version: 5977)
Remote Symmetrix ID        : 000196700483 (Microcode Version: 5977)
RDF (RA) Group Number      : 10 (09)
```

Target (R2) View					Source (R1) View					MODE	
-----					-----					-----	
		ST			LI	ST					
Standard		A			N	A					
Logical	Sym	T	R1 Inv	R2 Inv	K Sym	T	R1 Inv	R2 Inv		RDF Pair	
Device	Dev	E	Tracks	Tracks	S Dev	E	Tracks	Tracks	MACE	STATE	
-----					-----					-----	
DEV001	00087	RW	6042	0	NR 00059	WD	0	0	S...	Failed Over	
DEV002	00088	RW	6206	0	NR 0005A	WD	0	0	S...	Failed Over	
Total											
Track(s)			12248	0				0	0		
MB(s)			1531.0	0.0				0.0	0.0		

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC²

24

As seen in the output, the R1s are now Write Disabled, the links are set to Not Ready and the R2s are Read Write enabled. As application processing has been started using the R2 devices, we see that there are invalid tracks accumulating (R1 Inv Tracks) on the target Symmetrix. These are the changes that are being made to the R2 devices. When it is time to return to the source Symmetrix, these invalid tracks will be incrementally synchronized back to the source. The pair state is reflected as Failed Over.

Note that we have created a device group named synctgtdg on the remote host that is accessing the R2 devices. So we have created a device group of type RDF2 and have added the R2 devices to it. The query was executed on the remote host and shows the state from the perspective of the R2 devices.

SRDF Update

```
C:\>symrdf -g syntgtgdg update -nop
```

```
An RDF 'Update R1' operation execution is  
in progress for device group 'syntgtgdg'. Please wait...
```

```
Suspend RDF link(s).....Done.  
Merge device track tables between source and target.....Started.  
Devices: 0059-005A in (0225,010).....Merged.  
Merge device track tables between source and target.....Done.  
Resume RDF link(s).....Started.  
Resume RDF link(s).....Done.
```

```
The RDF 'Update R1' operation successfully initiated for  
device group 'syntgtgdg'.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

25

While the target (R2) device is still operational (Read Write Enabled to its local host), an incremental data copy from the target (R2) device to the source (R1) device can be initiated in order to *update the R1 mirror* with changed tracks from the target (R2) device. After an extended outage on the R1 side, a substantial amount of invalid tracks could have accumulated on the R2. If a failback is now performed, production starts from the R1. New writes to the R1 have to be transferred to the R2 synchronously. Any track requested on the R1 that has not yet been transferred from the R2 has to be read from across the links. This could lead to performance degradation on the R1 devices. The update operation helps to minimize this impact.

When performing an update, the R1 devices are still “Write Disabled”; the links become “Read Write” enabled because of the “Updated” state. The target devices (R2) remain “Read Write” during the update process.

The update operation can be used with the `-until` flag, which represents a skew value assigned to the update process. For example, you can choose to update until the accumulated invalid tracks is down to 30000. Then a failback operation can be executed.

Query after Update

C:\>**symrdf -g synctgtgdg query**

```
Device Group (DG) Name      : synctgtgdg
DG's Type                  : RDF2
DG's Symmetrix ID          : 000196800225   (Microcode Version: 5977)
Remote Symmetrix ID        : 000196700483   (Microcode Version: 5977)
RDF (RA) Group Number      : 10 (09)
```

Target (R2) View				Source (R1) View				MODE	
ST				LI				ST	
A				N				A	
Logical	Sym	T R1 Inv	R2 Inv	K Sym	T R1 Inv	R2 Inv	RDF Pair		
Device	Dev	E Tracks	Tracks	S Dev	E Tracks	Tracks	MACE	STATE	
DEV001	00087	RW	9554	0 RW	00059	WD	15469	0 S...	R1 UpdInProg
DEV002	00088	RW	9694	0 RW	0005A	WD	15153	0 S...	R1 UpdInProg
Total									
Track(s)		19248		0		30622		0	
MB(s)		2406.0		0.0		3827.8		0.0	

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

26

When the update operation is performed after a failover, the links become Read/Write enabled, but the Source devices are still Write disabled. Production work continues on the R2 devices.

SRDF Failback

```
C:\>symrdf -g syncdg failback -nop
```

```
An RDF 'Failback' operation execution is  
in progress for device group 'syncdg'. Please wait...
```

```
Write Disable device(s) on RA at target (R2).....Done.  
Suspend RDF link(s).....Done.  
Merge device track tables between source and target.....Started.  
Devices: 0059-005A in (0483,010).....Merged.  
Merge device track tables between source and target.....Done.  
Resume RDF link(s).....Started.  
Resume RDF link(s).....Done.  
Read/Write Enable device(s) on SA at source (R1).....Done.
```

```
The RDF 'Failback' operation successfully executed for  
device group 'syncdg'.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

27

When the source site has been restored, or if maintenance is completed, you can return production to the source site. The symrdf failback command will set the R2s Write Disabled, the link Read Write and the R1s Read Write enabled. Merging of the device track tables between the source and target is done. The SRDF links are resumed. The accumulated invalid tracks are transferred to the source devices from the target devices. So all changes made to the data when in a failed over state will be preserved. As noted earlier, the Primary host can access the R1 devices and start production work as soon as the command completes. If a track that has not yet been sent over from the R2 is required on the R1, SRDF can preferentially read that track from across the links.

As the R2s will be set to Write Disabled, it is important to shut down the applications using the R2 devices, and perform the appropriate host dependent steps to unmount filesystem/deactivate volume groups. If applications still actively access R2s when they are being set to Write Disabled, the reaction of the host accessing R2s will be unpredictable. In a true disaster, the failover process may not give an opportunity for a graceful shutdown. But a failback event should always be planned and done gracefully.

Query after Failback

C:\>symrdf -g syncdg query

```
Device Group (DG) Name      : syncdg
DG's Type                  : RDF1
DG's Symmetrix ID          : 000196700483 (Microcode Version: 5977)
Remote Symmetrix ID        : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number      : 10 (09)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
		ST			LI	ST				
Standard		A			N	A				
Logical	Sym	T	R1 Inv	R2 Inv	K Sym	T	R1 Inv	R2 Inv	RDF Pair	
Device	Dev	E	Tracks	Tracks	S Dev	E	Tracks	Tracks	MACE	STATE
-----					-----					-----
DEV001	00059	RW	15771	0	RW	00087	WD	9326	0 S...	SyncInProgress
DEV002	0005A	RW	15882	0	RW	00088	WD	9694	0 S...	SyncInProgress
Total										
Track(s)			31653	0				19020	0	
MB(s)			3956.6	0.0				2377.5	0.0	

EMC²

Module: SRDF/Synchronous Operations

28

As can be seen in the output, the R1s are set to Read Write, R2s are set to Write Disabled, and the links are set to Read Write. The pair states go into SyncInProgress. The accumulated invalid tracks have been transferred from the target array to the source array. Once all accumulated invalid tracks have been transferred, the pair state will go into Synchronized. Applications accessing the R2 devices must be stopped before a failback operation as the R2 devices are set to write-disabled. When a host suddenly loses RW access to a device while still actively accessing it, the results are unpredictable to say the least.

Lesson Summary

During this lesson the following topics were covered:

- SRDF Failover
- SRDF Update
- SRDF Failback



This lesson covered SRDF Disaster Recovery operations. Device and link states under different conditions were presented in detail. Host consideration when performing DR operations were also discussed.

Lesson: SRDF Decision Support Operations

This lesson covers the following topics:

- SRDF Split
- SRDF Establish
- SRDF Restore
- Concurrent SRDF

This lesson covers SRDF Decisions Support operations. Considerations for performing these operations are presented in detail. Concurrent SRDF where one R1 device is simultaneously paired with two R2 devices is also discussed.

SRDF Decision Support Operations

- Split
 - Enables accessing both the R1 and R2 device by their respective hosts
 - Suspends the links between the R1-R2 pairs
 - Read-Write enables the R2 device
- Establish
 - Resumes normal SRDF mirroring – source RW and target WD, link RW
 - Save source (R1) data – changes made to the R1 during split state are propagated to the R2; changes made to the R2 are discarded
- Restore
 - Resumes normal SRDF mirroring – source RW and target WD, link RW
 - Save target (R2) data – changes made to the R2 during split state are propagated to the R1; changes made to the R1 are discarded

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

31

The decision support operations for SRDF devices are:

Split an SRDF pair which stops mirroring for the SRDF pairs in a device group.

Establish an SRDF pair by initiating a data copy from the source side to target side. The operation can be a full or incremental.

Restore remote mirroring, which initiates a data copy from the target side to the source side. The operation can be full or incremental.

As noted in the slide title, these are decision support operations and are not disaster recovery/business continuance operations. In these situations, both the Source and Target sites are healthy and available.

SRDF Split

```
C:\>symrdf -g syncdg split -nop
```

```
C:\>symrdf -g syncdg query
```

```
Device Group (DG) Name      : syncdg
DG's Type                   : RDF1
DG's Symmetrix ID           : 000196700483 (Microcode Version: 5977)
Remote Symmetrix ID         : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number      : 10 (09)
```

Source (R1) View					Target (R2) View					MODE
ST					LI					
A					N					
Standard	Sym	T	R1 Inv	R2 Inv	K	Sym	T	R1 Inv	R2 Inv	RDF Pair
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MACE STATE
DEV001	00059	RW	0	16312	NR	00087	RW	10725	0	S...
DEV002	0005A	RW	0	16328	NR	00088	RW	10711	0	S...
Total										
Track(s)			0	32640				21436	0	
MB (s)			0.0	4080.0				2679.5	0.0	

EMC²

Module: SRDF/Synchronous Operations

32

The split command suspends the links between source (R1) and Target (R2) volumes. The source devices continue to be read write enabled. The target devices are set to read write enabled. It enables read and write operations on the target volumes. Writes to the R1 devices accumulate as R2 Inv(alid) Tracks – these are the tracks now owed to the R2 devices. Writes to the R2 devices accumulate as R1 Inv(alid) Tracks – these are the tracks owed to the R1 devices. The RDF Pair state reflects Split.

SRDF Establish

```
C:\>symrdf -g syncdg establish -nop
```

An RDF 'Incremental Establish' operation execution is in progress for device group 'syncdg'. Please wait...

```
Write Disable device(s) on RA at target (R2).....Done.
Suspend RDF link(s).....Done.
Mark target (R2) devices to refresh from source (R1).....Started.
Devices: 0087-0088 in (0483,010).....Marked.
Mark target (R2) devices to refresh from source (R1).....Done.
Merge device track tables between source and target.....Started.
Devices: 0059-005A in (0483,010).....Merged.
Merge device track tables between source and target.....Done.
Resume RDF link(s).....Started.
Resume RDF link(s).....Done.
```

The RDF 'Incremental Establish' operation successfully initiated for device group 'syncdg'.



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

33

Establish operation will resume SRDF remote mirroring. Changes made to the source while in a split state, will be transferred to the target. Changes made to the target are overwritten. The R2 devices are set to Write Disable. Hence applications should stop accessing the R2 devices, prior to performing an establish operation. The links are resumed.

Query after SRDF Establish

C:\>symrdf -g syncdg query

Device Group (DG) Name : syncdg
 DG's Type : RDF1
 DG's Symmetrix ID : 000196700483 (Microcode Version: 5977)
 Remote Symmetrix ID : 000196800225 (Microcode Version: 5977)
 RDF (RA) Group Number : 10 (09)

Source (R1) View					Target (R2) View					MODE
ST					LI					
A					N					
Logical	Sym	T	R1 Inv	R2 Inv	K	Sym	T	R1 Inv	R2 Inv	RDF Pair
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MACE STATE
DEV001	00059	RW	0	9045	RW	00087	WD	0	12966	S...
DEV002	0005A	RW	0	9309	RW	00088	WD	0	13291	S...
Total										
Track(s)			0	18354				0	26257	
MB(s)			0.0	2294.3				0.0	3282.1	

SyncInProg
SyncInProg

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

34

As can be seen in the query output, the states of the devices are reverted to their normal state (R1-RW; R2-WD) and the links are resumed (RW). Changes made to the R2 device during the split state are discarded. Changes made to the R1 device during the split state are propagated to the R2 device.

SRDF Restore

```
C:\>symrdf -g syncdg restore -nop
```

An RDF 'Incremental Restore' operation execution is in progress for device group 'syncdg'. Please wait...

```
Write Disable device(s) on SA at source (R1).....Done.
Write Disable device(s) on RA at target (R2).....Done.
Suspend RDF link(s).....Done.
Mark source (R1) devices to refresh from target (R2).....Started.
Devices: 0059-005A in (0483,010).....Marked.
Mark source (R1) devices to refresh from target (R2).....Done.
Merge device track tables between source and target.....Started.
Devices: 0059-005A in (0483,010).....Merged.
Merge device track tables between source and target.....Done.
Resume RDF link(s).....Started.
Resume RDF link(s).....Done.
Read/Write Enable device(s) on SA at source (R1).....Done.
```

The RDF 'Incremental Restore' operation successfully initiated for device group 'syncdg'.



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

35

Restore operation will resume SRDF remote mirroring. Changes made to the target while in a split state, will be transferred to the source. Changes made to the source are overwritten. The R2 devices are set to Write Disable. Hence applications should stop accessing the R2 devices, prior to performing an establish operation. The links are resumed. As data on the R1 devices will change without the knowledge of the host, access to R1 devices should be stopped prior to performing a restore operation. As soon as the command completes, R1 devices can be accessed again without waiting for synchronization to be completed. Any required track on the R1 that has not yet been received from the R2, will be read across the links.

Query after SRDF Restore

C:\>**symrdf -g syncdg query**

Device Group (DG) Name : syncdg
 DG's Type : RDF1
 DG's Symmetrix ID : 000196700483 (Microcode Version: 5977)
 Remote Symmetrix ID : 000196800225 (Microcode Version: 5977)
 RDF (RA) Group Number : 10 (09)

Source (R1) View					Target (R2) View					MODE
ST					LI					
A					N					
Logical	Sym	T	R1 Inv	R2 Inv	K	Sym	T	R1 Inv	R2 Inv	RDF Pair
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MACE STATE
DEV001	00059	RW	12868		0	RW	00087	WD	7450	0 S...
DEV002	0005A	RW	12117		0	RW	00088	WD	6732	0 S...
Total										
Track(s)			24985	0				14182	0	
MB(s)			3123.1	0.0				1772.8	0.0	

SyncInProg
SyncInProg

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

36

As can be seen in the query output, the states of the devices are reverted to their normal state (R1-RW; R2-WD) and the links are resumed (RW). Changes made to the R1 device during the split state are discarded. Changes made to the R2 device during the split state are propagated to the R1 device.

R1/R2 Personality Swap

- Changes the personality of the RDF devices
 - Current R1 becomes “New” R2; current R2 becomes “New” R1
- Data flow is from the “New” R1 (old R2) to the “New” R2 (old R1)
- Can be performed in one of two ways:
 - `symrdf swap`
 - `symrdf failover -establish`
- Useful for:
 - Disaster Recovery drills
 - Datacenter relocation
 - Maintenance operations on Local site hosts while continuing production work WITH Disaster Recovery protection
 - Selective “load balancing” for certain applications by swapping their device personalities and moving their workload to the other array



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

37

An R1/R2 personality swap (or R1/R2 swap) refers to when the RDF personality of the RDF device designations of a specified device group are swapped so that source R1 device(s) become target R2 device(s) and target R2 device(s) become source R1 device(s).

Sample scenarios for R1/R2 Swap

Symmetrix Load Balancing:

In today's rapidly changing computing environments, it is often necessary to redeploy applications and storage on a different Symmetrix without having to give up disaster protection. R1/R2 swap can enable this redeployment with minimal disruption, while offering the benefit of load balancing across two Symmetrix storage arrays.

Primary Data Center Relocation:

Sometimes a primary data center needs to be relocated to accommodate business practices. Businesses might want to test their Disaster Recovery readiness without sacrificing DR protection. R1/R2 swaps allow these customers to move their primary applications to their DR centers and continue to SRDF mirror back to their Primary data center.

Post-Failover Temporary Protection Measure:

If the hosts on the source side are down for maintenance, R1/R2 swap permits the relocation of production computing to the target site without giving up the security of remote data protection. When all problems have been solved on the local Symmetrix hosts, you have to failover again and swap the personality of the devices to go back to the original configuration.

Concurrent SRDF Devices


- An SRDF device with two SRDF mirrors is referred to as a Concurrent SRDF device
- There are three different types of Concurrent SRDF devices:
 - R11 – Each R1 mirror is paired with a different R2 mirror on two different remote Symmetrix arrays.
 - R21 – This device is the R2 mirror for an R1 device and also acts as a R1 mirror for another R2 device. This device is used in the secondary site of a Cascaded SRDF configuration.
 - R22 – Each R2 mirror is paired with a different R1 mirror on two different remote Symmetrix arrays. Only one of the R2 mirrors can be Read/Write on the links at a time.



The R21 devices are configured and used for Cascaded SRDF environments. The R22 devices are used in SRDF/Star environments. An R22 device has two R1 mirrors. However, it can receive data from just one of the R1 mirrors at a time.

Concurrent RDF – R11 Devices

- One R1 can be paired with two R2 devices, concurrently
- Each of the two concurrent mirrors must belong to different SRDF groups (RA groups)



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

39

Concurrent SRDF allows two remote SRDF mirrors of a single R1 device. A concurrent R1 device has two R2 devices associated with it. Each of the R2 devices is usually in a different array. Any combination SRDF modes is allowed:

R11 ↔ R2 (Site B) in Synchronous mode and R11 ↔ R2 (Site C) in Asynchronous mode

R11 ↔ R2 (Site B) in Synchronous mode and R11 ↔ R2 (Site C) in Adaptive Copy Disk mode

R11 ↔ R2 (Site B) in Synchronous mode R11 ↔ R2 (Site C) in Synchronous mode

R11 ↔ R2 (Site B) in Asynchronous mode and R11 ↔ R2 (Site C) in Asynchronous mode

Each of the R1 ↔ R2 pairs are created in different SRDF Groups.

2 Synchronous remote mirrors: A write I/O from the host to the R11 device cannot be acknowledged to the host as completed until both remote arrays signal the local array that the SRDF I/O is in cache at the remote side.

SRDF swap is not allowed in this configuration.

Concurrent RDF – R11 Example

```
C:\>symrdf -g syncdg query -rdfg all
```

```
Device Group (DG) Name      : syncdg
DG's Type                  : RDF1
DG's Symmetrix ID          : 000196700483 (Microcode Version: 5977)
Remote Symmetrix ID        : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number    : 10 (09)
Remote Symmetrix ID        : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number    : 11 (0A)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
	ST				LI	ST				
Standard	A				N	A				
Logical	Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv			RDF Pair
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks		MACE	STATE
-----					-----					-----
DEV001	00059	RW	0	0	RW	00087	WD	0	0	S... Synchronized
		RW	0	0	RW	00089	WD	0	0	CD.. Synchronized
DEV002	0005A	RW	0	0	RW	00088	WD	0	0	S... Synchronized
		RW	0	0	RW	0008A	WD	0	0	CD.. Synchronized

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC²

40

For the purpose of illustration we show the R1 devices paired with two R2 devices on the same remote array. The real use for R11 devices will be to pair them with R2 devices on two different remote arrays, perhaps at two different locations.

In this example, R1 devices 059 and 05A on SID:483 are paired with R2 devices 087 and 088 on SID:225, as well as concurrently paired with R2 devices 089 and 08A on SID:225. This was accomplished by the following two commands:

```
C:\>symrdf addgrp -label SRDF_CONC -sid 483 -remote_sid 225 -dir 1E:8,3E:8 -
remote_dir 1E:7,2E:7 -rdfg 11 -remote_rdfg 11
```

A new RDF group (number 11) has been created.

```
C:\>symrdf createpair -sid 483 -rdfg 11 -f pairs2.txt -type r1 -est
```

Where the file pairs2.txt contains:

```
059 089
```

```
05A 08A
```

This specifies that R1 devices 059 and 05A should now be concurrently paired with R2 devices 089 and 08A as well.

Listing Concurrent RDF Devices

```
C:\>symrdf -sid 483 list -concurrent
```

```
Symmetrix ID: 000196700483
```

Local Device View										
Sym Dev	Sym RDev	RDF Typ:G	STATUS			MODES		RDF S T A T E S		
			SA	RA	LNK	MDATE	R1 Inv Tracks	R2 Inv Tracks	Dev	RDev Pair
00059	00087	R1:10	RW	RW	RW	S..l.	0	0	RW	WD Synchronized
	00089	R1:11	RW	RW	RW	C.Dl.	0	0	RW	WD Synchronized
0005A	00088	R1:10	RW	RW	RW	S..l.	0	0	RW	WD Synchronized
	0008A	R1:11	RW	RW	RW	C.Dl.	0	0	RW	WD Synchronized

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

41

The output shows that R1 device is now concurrently paired with R2 device 087 in RDF Group 10 as well as with R2 device 089 in RDF Group 11. Note that “one leg” {059←→087} is in Synchronous mode of SRDF and the “other leg” {059←→089} is in Adaptive Copy Disk mode. Likewise for the device pairs {05A←→088} and {05A←→08A}. If we want to change the “other leg” to Synchronous mode as well, then we can use the command `symrdf set mode sync -rdfg 11`.

So the way to deal with the two different legs is to call them out with the `-rdfg` flag and explicitly specify which leg we want to operate on.

SRDF Consistency Protection (1)

- SRDF consistency group is a composite group of SRDF devices that are enabled for consistency
 - If a source R1 device in the consistency group cannot propagate data to its corresponding R2 device, SRDF consistency suspends data propagation from all the R1 devices in the group.
- The SRDF daemon `storrdfd` provides consistency protection for:
 - SRDF/Synchronous RDF-Engineuity Consistency Assist (ECA) consistency groups
 - SRDF/Asynchronous Multi-session Consistency (MSC) consistency groups
 - Concurrent SRDF



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

42

SRDF consistency preserves the dependent-write consistency of devices within a consistency group by monitoring data propagation from source devices to their corresponding target devices. If a source R1 device in the consistency group cannot propagate data to its corresponding R2 device, SRDF consistency suspends data propagation from all the R1 devices in the group.

A Composite group must be created using the RDF consistency protection option (`-rdf_consistency`) and must be enabled using the `symcg enable` command for the RDF daemon to begin monitoring and managing the consistency group. Devices in a consistency group can be from multiple arrays or from multiple SRDF groups in the same array.

Consistency protection is managed by the SRDF Daemon which is a Solutions Enabler process that runs on a host with Solutions Enabler and connectivity to the array. Consistency protection is available for SRDF/S, SRDF/A and Concurrent SRDF modes. `storrdfd` ensures that there will be a consistent R2 copy of the database at the point in time in which a data flow interruption occurs.

SRDF Consistency Protection (2)

- For SRDF/S RDF-ECA, the SRDF daemon
 - Continuously polls SRDF/S sessions for data flow interruptions
 - If an R1 cannot propagate data to its R2, then it suspends RDF links for all devices in the consistency group
- For SRDF/A MSC, the SRDF daemon performs
 - Cycle switching and cache recovery for all SRDF/A sessions within a consistency group
 - Manages the R1 -> R2 commits for SRDF/A sessions in multi-cycle mode



RDF-ECA provides consistency protection for synchronous mode devices by performing suspend operations across all SRDF/S devices in a consistency group. SRDF/A MSC will be discussed in detail in the next module.

VMAX3 SRDF Coordination

- Enabled at the array level
 - Enabled by default
 - No software interface to DISABLE
 - Only VMAX3 to VMAX3
- R1 performance metrics transmitted across the link
 - R1 metrics are merged with the R2 metrics
 - R2 reflects R1 metrics
 - Allows FAST movement decisions for R2 data to account for R1 workload
- SLOs associated with R1 and R2 devices can be different

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

44

SRDF Coordination is enabled at the VMAX3 array level by default. There is no Solutions Enabler or Unisphere for VMAX interface to enable or disable SRDF Coordination.

Only VMAX3 to VMAX3 is supported. Performance metrics are periodically transmitted from R1 to R2, across the SRDF link. The R1 metrics are merged with the R2 metrics. This instructs FAST to factor the R1 device statistics into the move decisions that are made on the R2 device. Service Level Objectives (SLO) associated with R1 and R2 devices can be different.

VMAX3 Geometry Compatible Mode

- VMAX3 with HYPERMAX OS 5977 track size is 128K; VMAX with Enginuity 5876 track size is 64K
- If the cylinder count on VMAX is odd, a device of the exact same size cannot be created on a VMAX3
- To support the full suite of SRDF and TimeFinder features
 - R1 and R2 devices must be of the same size
 - Source and Target devices for TimeFinder must be of the same size
- VMAX3 supports a new device attribute – Geometry Compatible Mode (GCM)
 - Device GCM attribute is treated as half a cylinder size smaller than its actual size
 - Enables full functionality between VMAX3 and VMAX arrays for SRDF and TimeFinder replication



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

45

A GCM device is treated as half a cylinder smaller than its true configured size. The extra half cylinder is not addressable by a host and cannot be replicated. The GCM attribute can only be set for volumes on arrays running VMAX3. A volume with GCM attribute set is referred to as a GCM device and its size is referred to as the device's GCM size. The attribute can be set or unset manually using the set command in conjunction with symdev/symdg/symcg/symcg with a new -gcm option. For most operations, Solutions Enabler sets it automatically when required. For example, Solutions Enabler automatically sets the GCM attribute when restoring from a physically larger R2. It will be set automatically as part of the symrdf createpair operation.

Mixed RDF Modes on Remote Adapters

```
C:\>symqos -sid 483 -ra -dir 1E set io -sync 50 -async 40 -copy 10
```

```
C:\>symqos -sid 483 list -ra -io
```

```
Symmetrix ID:000196700483
```

```
RA IO State      : Enabled
```

```
System Defaults:
```

```
Synchronous IOs () : 70  
Asynchronous IOs () : 20  
Copy IOs ()       : 10
```

```
RDF Directors:
```

Ident	Flg	IO Percent		
	R	Sync	Async	Copy
RF-1E	X	50	40	10
RF-2E	.	70	20	10
RF-3E	.	70	20	10
RF-4E	.	70	20	10



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

46

RA CPU resource distribution for Synchronous, Asynchronous, and Copy modes can be set either system wide which will affect all RAs or can be set on a subset of RAs. The resource distribution can be enabled or disabled. The system defaults as seen in the slide are 70/20/10 for Sync/Async/Copy modes.

As shown here, for purpose of illustration, the distribution can be changed for one of the directors if necessary. In this case RA-1E has been changed to 50/40/10 for Sync/Async/Copy modes.

Legend for Flg:

(R)A IO Set: **X = Set**, . = Default, - = N/A

Lesson Summary

During this lesson the following topics were covered:

- SRDF Split
- SRDF Establish
- SRDF Restore
- Concurrent SRDF



This lesson covered SRDF Decisions Support operations. Considerations for performing these operations were presented in detail. Concurrent SRDF where one R1 device is simultaneously paired with two R2 devices was also discussed.

Lesson: SRDF/S Disaster Recovery for a VMFS Datastore


This lesson covers the following topics:

- Performing SRDF/S Disaster Recovery of a VMFS datastore
- Performing SRDF Failback and resuming operations on the Primary ESXi server



This lesson covers performing SRDF/S DR of a VMFS datastore. Datastore is created on the R1 device and synchronized with the R2 device in SRDF/S mode. The device pair is failed over and the copy of the datastore on the R2 device is accessed from an ESXi server connected to the remote VMAX3 array.

Datastore Presented to Primary ESXi Server



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

49

For this example we will use the Production_Datastore.


vSphere Web Client



vSphere Web Client shows that Production_Datastore has been created on Symmetrix device 095. The device is in SID:225 and is of the Type:RDF1+TDEV. This indicates that the Production_Datastore is on an R1 device.



Browse Datastore



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

51

Next we browse the Production_Datastore to determine the VM resident on it. This shows that StudentVM01 is the VM.

Datastore Used by VM

The screenshot shows the vSphere Client interface for a VM named 'StudentVM01'. The 'Summary' tab is selected, displaying various details:

- General:**
 - Guest OS: CentOS 4/5/6 (32-bit)
 - VM Version: vmx-09
 - CPU: 1 vCPU
 - Memory: 512 MB
 - Memory Overhead: 20.42 MB
 - VMware Tools: Not running (Not installed)
 - IP Addresses:
 - DNS Name:
 - State: Powered On
 - Host: PrimaryESX064
 - Active Tasks:
 - vSphere HA Protection: N/A
- Resources:**
 - Consumed Host CPU: 9 MHz
 - Consumed Host Memory: 532.00 MB
 - Active Guest Memory: 10.00 MB
 - Provisioned Storage: 2.59 GB
 - Not-shared Storage: 2.39 GB
 - Used Storage: 2.39 GB
- Storage:** A table showing the storage configuration:

Storage	Device	Drive Type
Production_Datastore	EMC ISC...	Non-SSD
- Commands:**
 - Power Off
 - Suspend
 - Reset
 - Edit Settings
 - Open Console




© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

52

Summary details of StudentVM01 shows that it is indeed using only Production_Datastore for storage. The VM is powered on.

Data on StudentVM01



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

53

We can open a console to StudentVM01. We have created a folder named Production_data. What we have determined so far:

- 1) Primary ESXi Server has access to Production_Datastore
- 2) Production_Datastore has been created on device 095 in SID:225
- 3) Device 095 is an SRDF R1 device
- 4) StudentVM01 uses only Production_Datastore for its storage
- 5) StudentVM01 contains Production_data

The objective is to perform an SRDF Failover of device 095. Access the corresponding R2 device from the Remote ESXi Server and power-on the VM on the Remote ESXi Server.

R2 Device Information and Device Group Details

C:\>**symrdf list**

Symmetrix ID: 000196800**225**

Local Device View

		STATUS		MODES		R1 Inv		R2 Inv		RDF S T A T E S		
Sym	Sym	RDF										
Dev	RDev	Typ:G	SA	RA	LNK	MDATE	Tracks	Tracks	Dev	RDev	Pair	
00091	00062	R1:1	RW	RW	RW	C.D1.	0	0	RW	WD	Synchronized	
00095	00061	R1:1	RW	RW	RW	S..1.	0	0	RW	WD	Synchronized	

000196800225 > Data Protection > SRDF > VMFS_dg

SRDF pair list: VMFS_dg

Group Type: RDF1 Show Group Details>>

Local		Hop2				
Source Volume	Source LDev	Group	Remote Symmetrix	Target Volume	State	Volume State
00095	DEV001	1	000196700483	00061	Synchronized	Ready

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

54

symrdf list command shows that R1 device 095 is paired with R2 device 061. A device group named VMFS_dg of type RDF1 has been created and 095 has been added to it. Details of this device group also shows that the R2 device is 061 and is in SID:483. The pair state is Synchronized and the SRDF mode is Synchronous. The process of creating a device group using Unisphere for VMAX will be presented in a later lesson in this module.

LUNs Presented to Remote ESXi Server

RemoteESXi74 VMware ESXi, 5.1.0, 799733

Getting Started Summary Virtual Machines Resource Allocation Performance Configuration Local Users & Groups Events Permissions

Hardware

- Health Status
- Processors
- Memory
- Storage
- Networking
- Storage Adapters
- Network Adapters
- Advanced Settings
- Power Management

View: Datastores Devices Refresh Rescan All...

Name	Identifier	Runtime Name
EMC iSCSI Disk (naa.60000970000196700483533030303630)	naa.600009700001...	vmhba33:CD
EMC iSCSI Disk (naa.60000970000196700483533030303631)	naa.600009700001...	vmhba33:CD
EMC iSCSI Disk (naa.60000970000196700483533030303632)	naa.600009700001...	vmhba33:CD
Local VMware Disk (mpx.vmhba1:C0:T0:L0)	mpx.vmhba1:C0:T0...	vmhba1:C0:
Local NECVMWar CD-ROM (mpx.vmhba32:C0:T0:L0)	mpx.vmhba32:C0:...	vmhba32:C0:
EMC iSCSI Disk (naa.60000970000196700483533030303546)	naa.600009700001...	vmhba33:CD

EMC


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

55

Configuration information for the Remote ESXi Server shows the LUNs accessible by it. As in the case of TimeFinder SnapVX, we have to correlate the naa name with the WWN using Unisphere for VMAX.

Unisphere for VMAX – R2 Device Detail



The screenshot shows the EMC Unisphere for VMAX interface. The breadcrumb navigation path is: 000196700483 > Hosts > Masking View > remoteesxi74_MV > Storage Group > remoteesxi74_p. The 'PROPERTIES' table lists various attributes of the device, with several key fields highlighted by red boxes: 'Type' (RDF2+TDEV), 'Symmetrix Volume ID' (00061), and 'WWN' (6000097000019670048353030303631).

PROPERTIES	
Name	00061
Physical Name	N/A
Volume Identifier	N/A
Type	RDF2+TDEV
Encapsulated Volume	No
Encapsulated WWN	N/A
Status	Write Disabled
Reserved	No
Capacity (GB)	10
Capacity (MB)	10241
Capacity (Cylinder)	5462
Emulation	FBA
Symmetrix ID	000196700483
Symmetrix Volume ID	00061
HP Identifier Name	N/A
VMS Identifier Name	N/A
Nice Name	N/A
WWN	6000097000019670048353030303631


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

56

We use Unisphere for VMAX to confirm that the correct R2 device has been presented to the Remote ESXi Server. As shown in the slide, we can match the WWN of device 061 with the naa number shown in the previous slide.

Primary ESXi Server – Shut Down VM and Remove from Inventory



To illustrate SRDF functionality with VMFS, we will perform a “graceful” failover. While in a true disaster this would not be possible. For this example we shut down RDFSudentVM and remove it from the inventory of the Primary ESXi server.

Primary ESXi Server – Unmount Datastore

The screenshot displays the vSphere Datastore browser interface. The 'Datastores' table lists three datastores: 'datastore1' (Local VMware Disk), 'NFS_Share (read only)' (NFS), and 'Production_Datastore' (EMC iSCSI Disk). The 'Production_Datastore' is selected, and a context menu is open with 'Unmount' highlighted. Below the table, the 'Datastore Details' for 'Production_Datastore' are shown, including its location and hardware acceleration status. A 'Confirm Datastore Unmount' dialog box is overlaid on the right, indicating that no virtual machines reside on the datastore and asking for confirmation to proceed with the unmount process.


Identification	Device	Drive Type	Capacity
datastore1	Local VMware Disk (mpx.vmhba1:C0...	Non-SSD	15 GB
NFS_Share (read only)	10.127.108.130:/export/software	Unknown	246 GB
Production_Datastore	EMC iSCSI Disk (naa.60000000000000000000000000000000)	Non-SSD	9 GB

Datastore Details
Production_Datastore
Location: /vmfs/volumes/54d22d1d-90045fef-2a52-000000000000
Hardware Acceleration: Supported

Confirm Datastore Unmount
For successful datastore unmount the following requirements must be fulfilled:
✓ No virtual machine resides on the datastore.
⚠ Please do not perform any configuration operations that may result in I/O to the datastore while the unmount is in progress.
Do you want to continue with the unmount process?
OK Cancel


Next we unmount the Production_Datastore. After this we will perform an SRDF Failover. CLI operations for this has been presented in an earlier lesson. Using Unisphere for VMAX to perform SRDF operations will be presented in a later lesson in the module.

Rescan Remote ESXi Server




We now rescan Remote ESXi Server for all storage. The process for accessing the R2 device from the Remote ESXi server is identical to the one we used for accessing the linked Target on the Secondary ESXi server.

Mount R2 Datastore on Remote ESXi Server (1)



After rescan completes, use the Add Storage wizard on the Remote ESXi server. The VMFS label `Production_Datastore` indicates that it is the datastore on the R2 device. Choose this LUN and click Next.


Mount R2 Datastore on Remote ESXi Server (2)



As the LUN is a replica, we get Mount Options. In this example we choose to keep the existing signature.



Mount R2 Datastore on Remote ESXi Server (3)



The screenshot shows the vSphere Client interface for a Remote ESXi server (RemoteESXi74 VMware ESXi, 5.1.0, 799733). The 'Configuration' tab is selected, and the 'Datastores' view is active. The 'Datastores' table lists the following information:


Identification	Device	Drive Type	Capacity	Reclaimable
datastore1 (2)	Local VMware Disk...	Non-SSD	15.00 GB	1
Production_Datastore	EMC iSCSI Disk (...)	Non-SSD	9.75 GB	

The 'Production_Datastore' is highlighted with a black box. The interface also shows a left-hand navigation pane with 'Hardware' selected, and a top menu bar with options like File, Edit, View, Inventory, Administration, Plug-ins, and Help.

© Copyright 2015 EMC Corporation. All rights reserved. Module: SRDF/Synchronous Operations 62


In the example of SnapVX we had assigned a new signature. So the datastore on the linked Target was named snap-xxxxx. In this example for SRDF/S, we chose to keep the existing signature. Hence the name has not been pre-fixed. The same name (Production_Datastore) as on the Primary ESXi server is retained.

Browse Datastore and Add VM to Inventory



Next we browse the datastore and add the VM to the inventory of the Remote ESXi server. We are retaining the same name for the VM. We choose the Remote ESXi server to host this VM on.

Power On VM on Remote ESXi Server




© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

64

As the VM is a replica, we select "I copied it" for the Virtual Machine Message, when we power it on.

VM on Remote ESXi Server



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

65

We can open a console to the VM on the Remote ESXi server and verify that it has the same data that was available on the Primary ESXi server at the time of the “graceful” failover.

VM on Remote ESXi Server – Add More Data



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

66


To simulate resuming production work from the R2 device, we add more data to the VM. Next we will do the steps necessary to failback the SRDF pair, resume production work on the R1 device, and verify that the data added in the failed over state is available back on the R1 device.

Shut Down VM on Remote ESXi Server and Remove from Inventory




Prior to performing an SRDF failback and resuming production work from the R1 device, we shut down the VM on the Remote ESXi server and remove it from inventory. Remember that the R2 device will be Write Disabled on a failback operation.

Unmount Datastore from Remote ESXi Server



Next we unmount the Production_Datastore from the Remote ESXi server. Now we can perform the SRDF Failback operation. Details of this using CLI has been presented in an earlier lesson. Using Unisphere for VMAX for SRDF operations will be presented in a later lesson in this module.

Rescan Primary ESXi Server



EMC


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

69

We rescan the Primary ESXi server for all storage.

Mount Datastore Back on Primary ESXi Server



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: SRDF/Synchronous Operations

70

As we chose to "Keep the existing signature" when we mounted the datastore on the Remote ESXi server after the rescan on the Primary ESXi server, we can still see the greyed out Production_Datastore that was unmounted prior to a failover. Right click the datastore and choose "Mount".

"Add Storage" wizard should be used to mount if the datastore if we had chosen "Assign new signature" on the Remote ESXi server.

Add VM to Inventory of Primary ESXi Server after SRDF Failback



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

71

Browse datastore and add the VM to the inventory of the Primary ESXi server. Remember that we had shut down the VM, removed it from inventory, and unmounted the datastore in order to perform a graceful failover. So we now reverse the steps by first mounting the datastore (previous slide), add VM to inventory (this slide), and power-on VM (details not shown – but we have done this a couple of times before).

Data on VM on Primary ESXi Server



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

72

We can see that data added to the VM when running on Remote ESXi server (after an SRDF Fail over) is available back to the VM on Primary ESXi server after performing an SRDF Failback operation.

Lesson Summary

During this lesson the following topics were covered:

- Performing SRDF/S Disaster Recovery of a VMFS datastore
- Performing SRDF Failback and resuming operations on the Primary ESXi server



This lesson covered performing SRDF/S DR of a VMFS datastore. Datastore was created on the R1 device and synchronized with the R2 device in SRDF/S mode. The device pair was failed over and the copy of the datastore on the R2 device was accessed from an ESXi server connected to the remote VMAX3 array.

Lesson: Disaster Recovery for a VM Accessing RDM Hard Disks

This lesson covers the following topics:

- Performing SRDF/S Disaster Recovery for a VM accessing RDM hard disks
- Performing SRDF Failback and resuming operations on Primary ESXi



This lesson covers performing SRDF/S DR for a VM accessing RDM hard disks. The RDF device pair is failed over and the VM is powered-on from the R2 device on an ESXi server connected to the remote VMAX3 array.

VM on Primary ESXi Server

10.126.88.64

RDM_VM
StudentVM01

RDM_VM

Getting Started Summary Resource Allocation Performance Events Console Permissions

General

Guest OS: CentOS 4/5/6 (64-bit)
VM Version: vmx-09
CPU: 1 vCPU
Memory: 2048 MB
Memory Overhead: 30.25 MB
VMware Tools: Not running (Not installed)
IP Addresses:
DNS Name:
State: Powered On
Host: PrimaryESXi01
Active Tasks:
vSphere HA Protection: N/A

Resources

Consumed Host CPU: 9 MHz
Consumed Host Memory: 2070.00 MB
Active Guest Memory: 0.00 MB
Refresh Storage Usage
Provisioned Storage: 12.09 GB
Not-shared Storage: 12.09 GB
Used Storage: 12.09 GB

Storage	Drive Type	Capacity
datastore1	Non-SSD	15.00 GB
NFS_Share	Unknown	246.20 GB

Network Type

EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

75

The virtual machine RDM_VM on the Primary ESXi server seems to be using datastore1 for its storage. It is powered on.

VM Hard Disk Is an RDM



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

76

Properties of the VM shows that its hard disk is an RDM in Physical Compatibility Mode. It is only the mapping file and other files that define this VM, that are stored on datastore1.

vSphere Web Client – EMC Storage Viewer

The screenshot displays the vSphere Web Client interface for an RDM_VM. The left sidebar shows 'VMs and Templates' and 'Datastores'. The main content area has tabs for 'Issues', 'Performance', 'Resource Allocation', 'Storage Reports', 'Tasks', 'Events', and 'EMC Storage Viewer'. The 'EMC Storage Viewer' tab is active, showing a table with one item: 'Hard disk 1' with a capacity of 10.00 GB, located in 'datastore1 (12.0 GB free)', and of type 'RDM'. Below this, the 'Storage Details: Hard disk 1' section shows various properties:

Property	Value
Block Size	512 bytes
Is Thin Provisioned	true
WWN	60000970000196800225533030303931
Status	ONLINE
Type	RDF1+TDEV
Name	00091
Volume Identifier	N/A
Recover Point Tagged	No
Symmetrix Volume ID	00091

The interface is powered by EMC.

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations



77

vSphere Web Client shows that the hard disk is an RDF1 device with the Volume ID 091.

Primary ESXi Server – Datastore

PrimaryESXi64 VMware ESXi, 5.1.0, 799733

Getting Started Summary Virtual Machines Resource Allocation Performance Configuration Local Users & Groups Events Permissions

Hardware

- Health Status
- Processors
- Memory
- Storage
- Networking
- Storage Adapters
- Network Adapters
- Advanced Settings
- Power Management

View: Datastores Devices

Datastores

Identification	Device	Drive Type	Capacity	Free
datastore1	Local VMware Disk (mpx...	Non-SSD	15.00 GB	12.03 G
MEC Share (read only)	10.177.108.130:/vmv...	Unknown	246.20 GB	92.37 G
				6.50 G

Datastore Browser - [datastore1]

Search

Folders

- RDM_VM

[datastore1] RDM_VM

Name	Size	Type	Path
RDM_VM.vmx	2.12 KB	Virtual Machine	[datastore1] RDM_VM
RDM_VM.nvram		Virtual Disk	[datastore1] RDM_VM
vmware-1.log		Non-volatile me...	[datastore1] RDM_VM
vmware.log		Virtual Machine ...	[datastore1] RDM_VM
RDM_VM.vmsd		File	[datastore1] RDM_VM
vmx-RDM_VM-39...		File	[datastore1] RDM_VM
RDM_VM.vmx.lck		File	[datastore1] RDM_VM
RDM_VM.vmx~		File	[datastore1] RDM_VM

1 object selected 2.12 KB


Suspend

EMC

© Copyright 2015 EMC Corporation. All rights reserved. Module: SRDF/Synchronous Operations 78

On the Primary ESXi server `datastore1` is the Local Storage. All files that define the `RDM_VM` reside on this datastore. The mapping file for the RDM is kept here as well. This datastore is local and is not replicated using SRDF. We browse the datastore and download the `RDM_VM.vmx` file. This file has to be uploaded to a datastore on the Remote ESXi server in order to build the VM on it.

Data on RDM_VM on Primary ESXi Server



```
RDM_VM on PrimaryESXi64
File View VM

CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64

localhost login: root
Password:
Last login: Thu Feb 12 12:40:01 on tty1
[root@localhost ~]# cd /
[root@localhost /]# ls
bin  dev  home  lib64  media  opt  PROD_DATA  sbin  srv  tmp  var
boot  etc  lib  lost+found  mnt  proc  root  selinux  sys  usr
[root@localhost /]# cd PROD_DATA
[root@localhost PROD_DATA]# ls
file1 file2 file3
[root@localhost PROD_DATA]# _
```

We have added some data to the VM on the Primary ESXi server. The objective is to failover the RDM LUN, access the R2 device as an RDM on the Remote ESXi server, bring up the VM on the Remote ESXi server, and access this data.

Shut Down VM on Primary ESXi Server

```
RDM_VM on PrimaryESXi64
File View VM

Shutting down...Stopping block device availability:
y: Deactivating block devices:
  [SKIP]: unmount of VolGroup-lv_swap (dm-1) mounted on [SWAP]
  [SKIP]: unmount of VolGroup-lv_root (dm-0) mounted on /
                                     [ OK ]
Stopping monitoring for VG VolGroup: 2 logical volume(s) in volume group "VolG
roup" unmonitored
                                     [ OK ]
Sending all processes the TERM signal...
                                     [ OK ]

To release cursor, press CTRL + ALT
```



We will perform a planned “graceful” failover. So we first shut down the VM on the Primary ESXi server.

SRDF Failover

```
C:\>symrdf -g RDM_dg failover -nop
```

```
---Output Truncated----
```

```
C:\>symrdf -g RDM_dg query
```

```
Device Group (DG) Name      : RDM_dg
DG's Type                   : RDF1
DG's Symmetrix ID           : 000196800225   (Microcode Version: 5977)
Remote Symmetrix ID         : 000196700483   (Microcode Version: 5977)
RDF (RA) Group Number      : 1 (00)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
		ST			LI		ST			
Standard		A			N		A			
Logical	Sym	T	R1 Inv	R2 Inv	K Sym	T	R1 Inv	R2 Inv	RDF Pair	
Device	Dev	E	Tracks	Tracks	S Dev	E	Tracks	Tracks	MACE STATE	
-----					-----					-----
DEV001	00091	WD	0	0	NR	00062	RW	0	0 S... Failed Over	

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC²

81

We use SYMCLI to perform an SRDF failover. Details of creating a device group, and adding the device have been presented earlier in this module. Note that the R2 device is 062 in Remote SID:483.

LUNs Presented to Remote ESXi Server

The screenshot shows the 'RemoteESXi74 VMware ESXi, 5.1.0, 799733' configuration page. The 'Configuration' tab is selected, and the 'Devices' sub-tab is active. A table lists the devices presented to the ESXi server. The third row, 'EMC iSCSI Disk (naa.60000970000196700483533030303632)', is highlighted. The table has columns for Name, Identifier, and Runtime Name.

Name	Identifier	Runtime Name
EMC iSCSI Disk (naa.60000970000196700483533030303630)	naa.600009700001...	vmhba33:CC
EMC iSCSI Disk (naa.60000970000196700483533030303631)	naa.600009700001...	vmhba33:CC
EMC iSCSI Disk (naa.60000970000196700483533030303632)	naa.600009700001...	vmhba33:CC
Local VMware Disk (mpx.vmhba1:C0:T0:L0)	mpx.vmhba1:C0:T0...	vmhba1:C0:
Local NECVMWar CD-ROM (mpx.vmhba32:C0:T0:L0)	mpx.vmhba32:C0:...	vmhba32:CC
EMC iSCSI Disk (naa.60000970000196700483533030303546)	naa.600009700001...	vmhba33:CC

EMC

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

82

From the configuration of the Remote ESXi server we can list the devices presented to it. Note the highlighted device and its naa number. We will use Unisphere for VMAX to verify that this indeed corresponds to device 062 on SID:483.

Unisphere for VMAX

000196700483 > Storage > Volumes > 00062

PROPERTIES	
Name	00062
Physical Name	N/A
Volume Identifier	N/A
Type	RDF2+TDEV
Encapsulated Volume	No
Encapsulated WWN	N/A
Encapsulated Device Flag	None
Status	Ready
Reserved	No
Capacity (GB)	10
Capacity (MB)	10241
Capacity (Cylinder)	5462
Emulation	FBA
AS400 Gatekeeper	No
Symmetrix ID	000196700483
Symmetrix Volume ID	00062
HP Identifier Name	N/A
VMS Identifier Name	N/A
Nice Name	N/A
WWN	60000970000196700483533030303632




© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

83

Examining the Volume Information in Unisphere for VMAX shows the WWN for device 062. This WWN matches the naa number noted in the previous slide. So, the Remote ESXi server indeed has access to the R2 device.

Rescan Remote ESXi Server



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

84

As before we will rescan the Remote ESXi server for all storage.

Create VM on Remote ESXi Server



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

85

Next we browse the datastore on the Remote ESXi server and add a folder named RDM_VM to it. After this we upload the RDM_VM.vmx file to this folder. The steps on this slide and the next few slides have to be done the first time the RDM device failed over. As we will upload the RDM_VM.vmx file, edit settings to point to the appropriate R2 RDM, and store the mapping information with the VM, subsequent failover operations will not require this process to be repeated.

Add VM to Inventory of Remote ESXi Server



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

86

Next we add the virtual machine to the inventory of the Remote ESXi server. Right click the RDM_VM.vmx file in the Datastore Browser and select Add to Inventory.

Edit VM Settings – Remove Existing Hard Disk



Next we remove the hard disk from the virtual machine. This is because this existing hard disk is actually pointing to the R1 device. We will remove it and replace it with the RDM that points to the R2 device.

Add R2 Device as RDM to VM on Remote ESXi

The image displays two screenshots of the VMware Add Hardware wizard, illustrating the steps to add an R2 device as an RDM to a VM on a Remote ESXi host.

Step 1: Select a Disk

The first screenshot shows the "Add Hardware" wizard with the "Select a Disk" step. The "Device Type" is set to "Disk". The "Select a Disk" section is active, and the "Raw Device Mappings" option is selected under "Select the type of disk to use.".

Step 2: Select and Configure a Raw LUN

The second screenshot shows the "Add Hardware" wizard with the "Select and Configure a Raw LUN" step. The "Which LUN would you like to use for this raw disk?" question is displayed. A table lists available LUNs:

Name	Identifier	Path ID
EMC iSCSI Disk (naa.60000970000196700483533030303630)	naa.6000097...	vmhba33
EMC iSCSI Disk (naa.60000970000196700483533030303631)	naa.6000097...	vmhba33
EMC iSCSI Disk (naa.60000970000196700483533030303632)	naa.6000097...	vmhba33
EMC iSCSI Disk (naa.60000970000196700483533030303634)	naa.6000097...	vmhba33

The LUN with identifier "naa.60000970000196700483533030303632" is selected.

Step 3: Select a Datastore


The third screenshot shows the "Add Hardware" wizard with the "Select a Datastore" step. The "Onto which datastore do you want to map this LUN?" question is displayed. The "Store with Virtual Machine" option is selected under "Select the datastore on which to store the LUN mapping.".

EMC

© Copyright 2015 EMC Corporation. All rights reserved. Module: SRDF/Synchronous Operations 88

Using the Add Hardware wizard, we add the device as an RDM. We select the device we had correlated earlier as being the R2 device.


Add R2 Device as RDM to VM on Remote ESXi Server



Select Physical compatibility mode and finish the process of adding the hard disk.




Power on VM on Remote ESXi Server



As usual we answer the Virtual Machine Message with "I copied it", when we power it on.

Verify Data on VM on the Remote ESXi Server



```
RDM_VM (2) on RemoteESXi74
File View VM

CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64

localhost login: root
Password:
Last login: Thu Feb 12 12:45:43 on tty1
[root@localhost ~]# cd /
[root@localhost /]# ls
bin  dev  home  lib64  media  opt  PROD_DATA  sbin  srv  tmp  var
boot  etc  lib  lost+found  net  proc  root  sellinux  sys  usr

[root@localhost /]# cd PROD_DATA
[root@localhost PROD_DATA]# ls
file1  file2  file3
[root@localhost PROD_DATA]# _
```




© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

91

We can launch a console to the VM and verify that we have access to the data we created on the R1 device prior to the failover.

Add New Data



```
RDM_VM (2) on RemoteESXi74
File View VM

CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64

localhost login: root
Password:
Last login: Thu Feb 12 13:46:12 on tty1
[root@localhost ~]# cd /
[root@localhost /]# ls
bin  dev  home  lib64  media  opt  PROD_DATA  root  selinux  sys  usr
boot  etc  lib  lost+found  net  proc  r2_new_data  sbin  srv  tmp  var

[root@localhost /]# cd r2_new_data
[root@localhost r2_new_data]# ls
r2_new_file1  r2_new_file2  r2_new_file3
[root@localhost r2_new_data]# shutdown +1 -P

Broadcast message from root@localhost
(/dev/tty1) at 13:53 ...

The system is going down for power off in 1 minute!
```

EMC

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

92

To simulate production work on the Remote site while in a failover scenario, we now add some more files to the R2 device. After this we shut down the VM and perform an SRDF failback operation. The objective is to verify that the data added on the R2 side will be available on the R1 when we power on the VM back again on the Primary ESXi server.

SRDF Failback (1)

```
C:\>symrdf -g RDM_dg query
```

```
Device Group (DG) Name      : RDM_dg
DG's Type                   : RDF1
DG's Symmetrix ID           : 000196800225   (Microcode Version: 5977)
Remote Symmetrix ID         : 000196700483   (Microcode Version: 5977)
RDF (RA) Group Number       : 1 (00)
```

Source (R1) View					Target (R2) View					MODE
ST					LI	ST				
A					N	A				
Logical	Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv		RDF Pair	
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks		MACE STATE	
DEV001	00091	WD	0		0 NR	00062	RW	32308	0 S... Failed Over	

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC

93

Query shows that there are accumulated invalid tracks owed back to the R1 as we performed production work on the R2 device.

SRDF Failback (2)

```
C:\>symrdf -g RDM_dg failback -nop
```


```
An RDF 'Failback' operation execution is  
in progress for device group 'RDM_dg'. Please wait...
```

```
Write Disable device(s) on RA at target (R2).....Done.  
Suspend RDF link(s).....Done.  
Merge device track tables between source and target.....Started.  
Devices: 0091-0091 in (0225,001).....Merged.  
Merge device track tables between source and target.....Done.  
Resume RDF link(s).....Started.  
Resume RDF link(s).....Done.  
Read/Write Enable device(s) on SA at source (R1).....Done.
```

```
The RDF 'Failback' operation successfully executed for  
device group 'RDM_dg'.
```

Using SYMCLI we perform an SRDF failback operation.

Rescan Primary ESXi Server and Power on VM



```
RDM_VM on PrimaryESXi64
File View VM
CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64
localhost login: root
Password:
Last login: Thu Feb 12 13:48:44 on tty1
[root@localhost ~]# cd /
[root@localhost /]# ls
bin  dev  home  lib64  media  opt  PROD_DATA  root  selinux  src  usr
boot  etc  lib  initrd.img  mod  proc  r2_new_data  sbin  srv  tmp  var
[root@localhost /]# cd r2_new_data
[root@localhost r2_new_data]# ls
r2_new_file1  r2_new_file2  r2_new_file3
[root@localhost r2_new_data]# _
```

To release cursor, press CTRL + ALT



We can now rescan the Primary ESXi server, power on the VM, and verify that all the data added to the R2 in the failed over state is available to the VM which is now using the R1 device in RDM.

Lesson Summary

During this lesson the following topics were covered:

- Performing SRDF/S Disaster Recovery for a VM accessing RDM hard disks
- Performing SRDF Failback and resuming operations on Primary ESXi



This lesson covered performing SRDF/S DR for a VM accessing RDM hard disks. The RDF device pair was failed over and the VM was powered-on from the R2 device on an ESXi server connected to the remote VMAX3 array.

Lesson: SRDF/S Operations – Unisphere for VMAX

This lesson covers the following topics:

- Creating Dynamic SRDF Groups and SRDF Pairs
- Creating Device Groups
- Performing SRDF/S Operations



This lesson covers performing SRDF operations using Unisphere for VMAX.

Creating Dynamic SRDF Groups (1)

The screenshot displays the EMC Unisphere for VMAX V8.0.1.7 web interface. The top navigation bar includes icons for Storage, Hosts, Data Protection, Performance, Databases, System, and Support. The breadcrumb trail indicates the current location: 000196700483 > Data Protection > Replication Groups and Pools > SRDF Groups. A table lists the configured SRDF Groups:

Group	SRDF Group Label	Remote SRDF Group	Remote Symmetrix	Volume Count
1 (0)	Group_1	1 (0)	000196800225	2
10 (9)	SRDF_Sync	10 (9)	000196800225	2
11 (A)	SRDF_CONC	11 (A)	000196800225	2

Below the table, there are buttons for 'Create Group', 'Create Pairs', 'Edit', 'View Details', and a double arrow '>>'. The 'Create Group' button is highlighted with a red box. The bottom of the interface shows the EMC logo, copyright information (© Copyright 2015 EMC Corporation. All rights reserved.), the module name (Module: SRDF/Synchronous Operations), and the page number (98).

To list the currently configured SRDF Groups, navigate to SID>Data Protection>Replication Groups and Pools>SRDF Groups. We see that there are currently 4 SRDF Groups created on SID:483. Click Create Group to launch the wizard.

Creating Dynamic SRDF Groups (2)

Create SRDF Group (000196700483)

* Communication Protocol: FC

SRDF Group Label: uni_rdfg

LOCAL

Symmetrix ID: 000196700483

* SRDF Group Number: 2

Director: RF-1E:8, RF-3E:8

REMOTE

* Remote Symmetrix ID: 000196800225

* Remote SRDF Group Number: 2

Remote Director: RF-1E:7, RF-2E:7

Show Advanced >>

OK Cancel Help

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC

99

The Unisphere for VMAX, Create SRDF Group dialog is shown on the slide.

Choose the desired Communication protocol **FC** or **GigE**, and enter an RDF group label. Choose a Remote Symmetrix ID, enter the desired RDF group number for both the source and remote Symmetrix arrays. Choose the RDF director:Ports that will be part of this group and then click OK to create the RDF Group. The new RDF group will appear in the SRDF Groups listing. This is equivalent to the command line syntax:

```
symrdf addgrp -label uni_rdfg -sid 483 -remote_sid 225 -dir 1E:8,3E:8 -  
remote_dir 1E:7.2E:7 -rdfg 2 -remote_rdfg 2
```

Creating Dynamic SRDF Pairs

The screenshot displays the EMC Unisphere for VMAX v8.0.1.7 interface. The main window shows the 'SRDF Groups' page with a table of existing groups. A 'Create Pairs' button is highlighted with a red box and a circled '1'. Below the table, a 'Create SRDF Pairs' dialog box is open, showing configuration options: Mirror Type (R1), SRDF Mode (Adaptive Copy), Adaptive Copy Mode (Enabled: Disk Mode), Number of Volumes in range (1), Local Start Volume (05B), and Remote Start Volume (08B). A 'Show Advanced' link is visible. The 'Establish' radio button is selected. A 'Confirmation' dialog box is also shown, asking 'Create selected pairs now?' with 'Yes' and 'No' buttons. The 'Yes' button is highlighted with a red box and a circled '3'. The 'Create SRDF Pairs' dialog box has a circled '2' next to the 'Include Larger R2' checkbox.

Group	SRDF Group Label	Remote SRDF Group	Rem
1 (0)	Group_1	1 (0)	0001
2 (1)	uni_rdfg	2 (1)	0001
10 (9)	SRDF_Sync	10 (9)	0001
11 (A)	SRDF_CONC	11 (A)	0001

To create Dynamic RDF pairs in Unisphere for VMAX, navigate to the SRDF Groups page. Click the RDF group that you want create RDF Pairs in and then click the Create Pairs button to launch the Create Pairs dialog.

In the dialog, choose the RDF Mirror Type (R1 or R2), RDF Mode. Then choose the number of devices that will form the RDF pairs and the starting volume in the local and remote array (if necessary, use the Select button to help pick the correct volume). In this example the local mirror type will be R1, the RDF mode will be adaptive copy disk (which is the default).

Click the Show Advanced link to see additional options. The slide shows the advanced options. In this example Establish box has been checked. Click OK and answer in the affirmative for the Confirmation. This is equivalent to the command syntax:

```
symrdf createpair -sid 483 -f pairs.txt -rdfig 2 -type R1 -establish
```

Where:

pairs.txt

05B 08B

Setting SRDF Group Attributes (1)

EMC Unisphere for VMAX V8.0.1.7

000196700483 > Data Protection > Replication Groups and Pools > SRDF Groups

Group	SRDF Group Label	Remote SRDF Group	Remote Symmetrix	Volume Count
1 (0)	Group_1	1 (0)	000196800225	2
2 (1)	uni_rdfg	2 (1)	000196800225	1
10 (9)	SRDF_Sync		000196800225	2
11 (A)	SRDF_CONC		000196800225	2

SRDF/A Setting
SRDF/A Pacing Setting
Swap Groups
Delete

Create Group Create Pairs Edit View Details >>

EMC

© Copyright 2015 EMC Corporation. All rights reserved. Module: SRDF/Synchronous Operations 101

From the SRDF Groups page, select the SRDF Group and click ">>". View attributes that can be set and other actions on this SRDF Group are displayed.

Setting SRDF Group Attributes (2)

The image displays two side-by-side dialog boxes from the EMC SRDF/Synchronous Operations module.

Set SRDF/A Settings

- Group: 2 (uni_rdfg)
- Minimum Cycle Time (seconds): 15
- Session Priority ('1' is Highest): 33
- Transmit Idle: ☒ Enabled, ☐ Use Both Sides

Buttons: OK, Cancel, Help

Set SRDF/A Pacing Settings

- Group: 2 (uni_rdfg)
- Pacing Delay (usec): 50000
- Threshold: 60
- Autostart Group Pacing: ☐ Enabled
- Group Pacing Activate/Deactivate: No change (dropdown menu)

Buttons: OK, Cancel, Help

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

102

Choose "SRDF/A Setting", "SRDF/A Pacing Setting". Each of the choices will launch a specific dialog. Make the desired changes in the specific dialog and click OK.

Creating Device Group (1)

000196700483 > Data Protection > Replication Groups and Pools > Device Groups

Name	Group Type

Create Edit Delete View Details >>

Create Device Group

1 Device Group Name and Type

2 Select Source

3 Select Target and Replication Technology

4 Select Target Manually

5 Summary

Symmetrix 000196700483

* Device Group Name uni_rdfdg

* Device Group Type

Regular
R1
R2
R21
Any

< Back Next > Finish Cancel

EMC

© Copyright 2015 EMC Corporation. All rights reserved. Module: SRDF/Synchronous Operations 103

From SID>Data Protection>Replication Groups and Pools>Device Groups page, click Create. This launches the Create Device Group wizard. Give a name for the Device Group. For SRDF it is important to select the appropriate Device Group Type. In this example, we choose type R1.

Creating Device Group (2)

Create Device Group

1 Device Group Name and Type

2 Select Source

3 Select Target and Replication Technology

4 Select Target Manually

5 Summary

* Select Source: **Select volumes manually**
Select storage group

Select Source Vol Type: STD

Name	Capacity (GB)	Status	Volume Confi
00059	2.00	Ready	RDF1+TDEV
0005A	2.00	Ready	RDF1+TDEV
0005B	2.00	Ready	RDF1+TDEV

Add to Group

< Back Next > Finish Cancel

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

104

For Select Source, choose Select volumes manually. If we want to add all devices in a Storage Group to this device group, then we could Select storage group. Select Source Vol Type as STD.

Click the device to add to the Device Group and click Add to Group. In this example we are adding the R1 device 05B which we had created as an SRDF pair (with the R2 being 08B).

Creating Device Group (3)

The screenshot shows the 'Create Device Group' wizard with five steps. Step 2, 'Select Source', is active. It displays a table of available devices. The table has columns for 'Volume T...', 'Name', 'Capacity (...)', and 'Status'. Two devices are listed: 00059 and 0005A. Device 0005B is highlighted in blue. Below the table is an 'Add to Group' button. At the bottom of the wizard, there are four buttons: '< Back', 'Next >', 'Finish' (which is circled), and 'Cancel'.

Volume T...	Name	Capacity (...)	Status
	00059	2.00	Ready
	0005A	2.00	Ready
	0005B	2.00	Ready

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

105

Note that the device has been moved down from the list. Click Finish. This will create a device group and add device 05B to it. The equivalent command syntax would be:

```
symdbg create -type R1 unirdfdg
```

```
symdbg -g unirdfdg add dev 05B
```

SRDF Operations (1)

000196700483 > Data Protection > SRDF

Source Group	Standard	BCV State	Group
uni_rdfdg	1	0 Synchronized	RDF

EstablishSplitSuspendRestoreView Storage GroupsView Details>>

ResumeFailoverFailbackSwapEnable ConsistencyDisable ConsistencyDelete PairMoveActivate SRDF/ADeactivate SRDF/AINvalidateReadyNot ReadyR1 UpdateRW EnableWrite DisableRW Disable R2sRefreshSet Mode

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved. Module: SRDF/Synchronous Operations 106

All SRDF operations are performed from SID>Data Protection>SRDF page. Select the device group to be managed. Clicking the (>>) button shows the exhaustive list of operations that can be performed from Unisphere for VMAX.

SRDF Operations (2)

Set Mode - Device Group

Source Group

uni_rdfdg

Source Type

All

☐ Use 2nd hop

Group

2

SRDF Mode

Synchronous

Adaptive Copy Mode

Enabled: Disk Mode

AC Skew

65535

Use Consistent

☐

Show Advanced >>

Add to Job List

Cancel

Help

Add to Job List

Run Now

Failover - Device Group

Device Group

uni_rdfdg

Source Type

All

☐ Use 2nd hop

Group

2

Copy

☒ None ☐ Use Establish ☐ Use Restore

☐ Use Remote Side ☐ Use Immediate

Show Advanced >>

Add to Job List

Cancel

Help

Add to Job List

Run Now

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

107

Selecting Set Mode from the operations list gives the dialog. We can change the mode (in this example to Synchronous) and click on Run Now. Similarly if Failover is selected, we get the corresponding dialog and we can execute an SRDF Failover operation for the devices in the device group.

Lesson Summary

During this lesson the following topics were covered:

- Creating Dynamic SRDF Groups and SRDF Pairs
- Creating Device Groups
- Performing SRDF/S Operations

This lesson covered performing SRDF operations using Unisphere for VMAX.

Lab: SRDF/Synchronous Operations

This lab covers

- SRDF Initial Setup and Basic Operations – Creating Dynamic RDF Groups, RDF pairs
- SRDF/S Disaster Recovery Operations – Failover, Failback
- SRDF/S Decision Support/Concurrent Access Operations – Split, Establish, Restore

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

109

This lab covers creating dynamic RDF groups, RDF pairs, basic SRDF operations as well as SRDF Disaster Recovery and Decision Support operations.

Lab Summary

This lab covered

- SRDF Initial Setup and Basic Operations – Creating Dynamic RDF Groups, RDF pairs
- SRDF/S Disaster Recovery Operations – Failover, Failback
- SRDF/S Decision Support/Concurrent Access Operations – Split, Establish, Restore

This lab covered creating dynamic RDF groups, RDF pairs, basic SRDF operations as well as SRDF Disaster Recovery and Decision Support operations.

Lab: SRDF/S Disaster Recovery for a VMFS Datastore

This lab covers

- Identifying and correlating SRDF R1 and R2 devices with devices accessible by Primary and Remote ESXi servers
- Creating a VMFS Datastore on the R1 device and deploying a Virtual Machine
- Using Unisphere for VMAX to create a Device Group, add the R1 device and perform SRDF Failover operation
- Accessing the SRDF R2 device from the Remote ESXi server and power-on Virtual Machine on the Remote ESXi server

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Synchronous Operations

EMC

111

This lab covers performing SRDF/S Disaster Recovery for a VMFS Datastore.

Lab Summary

This lab covered

- Identifying and correlating SRDF R1 and R2 devices with devices accessible by Primary and Remote ESXi servers
- Creating a VMFS Datastore on the R1 device and deploying a Virtual Machine
- Using Unisphere for VMAX to create a Device Group, add the R1 device and perform SRDF Failover operation
- Accessing the SRDF R2 device from the Remote ESXi server and power-on Virtual Machine on the Remote ESXi server

This lab covered performing SRDF/S Disaster Recovery for a VMFS Datastore.

Module Summary

Key points covered in this module:

- Creating Dynamic SRDF Groups and Dynamic SRDF Pairs
- Performing SRDF/S Operations using SYMCLI and Unisphere for VMAX
- Performing Disaster Recovery of a VMFS datastore
- Performing Disaster Recovery of a VM accessing an RDM hard disk



This module covered SRDF operations in Synchronous mode. Use of SYMCLI and Unisphere for VMAX to perform SRDF operations were presented in detail. Methods for performing DR operations in a virtualized environment for both VMFS Datastore and RDM use cases were discussed.

This slide intentionally left blank.

Module: SRDF/Asynchronous Operations

Upon completion of this module, you should be able to:

- Describe SRDF/Asynchronous remote replication on a VMAX3 array
- Perform SRDF/A operations
- Describe SRDF/A resiliency features
- Manage SRDF/A multi-session consistency



This module focuses on SRDF/Asynchronous mode of remote replication. Concepts and operations for SRDF/A in single and multi-session modes are presented. SRDF/A resiliency features are also discussed.

Lesson: SRDF/A Concepts and Operations


This lesson covers the following topics:

- Multi-cycle mode for SRDF/A on VMAX3 arrays
- SRDF/A – System level and group level attributes
- Adding/removing device pairs to/from active SRDF/A session
- Monitoring SRDF/A



This lesson covers SRDF/A multi-cycle mode on VMAX3 arrays. The attributes that can be set for SRDF/A at a system and group level are discussed in detail. Methods for adding/removing RDF device pairs to/from active SRDF/A sessions and monitoring SRDF/A are also presented.

SRDF/A Multi-cycle Mode – VMAX3



SRDF/A Multi-Cycle Mode (MCM) allows more than two capture cycles on the R1 side.


When the `minimum_cycle_time` has elapsed the data from the capture cycle will be added to a transmit queue and a new capture cycle will occur. The transmit queue is a feature of SRDF/A. It provides a location for R1 captured cycle data to be placed so a new capture cycle can occur.

The capture cycle will occur even if no data is transmitted across the link. If no data is transmitted across the link the capture cycle data will again be added to the transmit queue. The transmit queue holds the data until it is transmitted across the link. The transmit cycle will transfer the data in the oldest capture cycle to the R2 first and then repeat the process.

The benefit of this is to capture controlled amounts of data on the R1 side. Each capture cycle will occur at regular intervals and will not contain large amounts of data waiting for a cycle to occur.

Another benefit is data that is sent across the SRDF link will be smaller in size and should not overwhelm the R2 side. The R2 side will still have two delta sets, the receive and the apply.

SRDF/A Cycle Switch – Legacy Mode



In previous versions of Enginuity an active SRDF/A session had capture and transmit cycles on the R1 side and receive and apply cycles on the R2 side. The factors that governed a cycle switch were: the minimum cycle time has expired, the transmit delta set has been completely transferred and the apply delta set was completely applied. The creation of a new capture cycle was dependent on the transmit cycle completing its commit of data from the R1 side to the R2 side.

SRDF System Attributes

```
C:\>symcfg -sid 483 list -v|more
```

```
Symmetrix ID: 000196700483 (Local)  
Time Zone   : Eastern Standard Time
```

```
Product Model           : VMAX200K  
Symmetrix ID            : 000196700483  
  
Microcode Version (Number) : 5977 (17590000)  
Microcode Registered Build : 1  
Microcode Date           : 12.02.2014  
  
Microcode Patch Date     : 12.02.2014  
Microcode Patch Level    : 497
```

```
--Output Truncated--
```

```
Switched RDF Configuration State : Enabled  
Concurrent RDF Configuration State : Enabled  
Dynamic RDF Configuration State : Enabled  
Concurrent Dynamic RDF Configuration : Enabled
```

```
--Output Truncated--
```

```
SRDF/A Maximum Host Throttle (Secs) : 0  
SRDF/A Maximum Cache Usage (Percent) : 75  
SRDF/A DSE Maximum Capacity (GB)    : NoLimit
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

5

As noted in the SRDF/S module, the RDF Configuration States are all Enabled by default. The use of Host Throttle, Maximum Cache Usage and DSE maximum Capacity attributes will be explained later in this lesson.

SRDF Group – SRDF/A Attributes

```
C:\>symcfg -sid 483 list -rdfg 20 -rdfa
```

```
Symmetrix ID : 000196700483
```

S Y M M E T R I X R D F A G R O U P S

RA-Grp	Group	Flags	Cycle	Pri	Thr	Transmit	Delay	Thr	GRP	DEV	FLG
	Name	CSRM TDA	time			Idle Time	(usecs)	(%)	SAU	SAU	P
20 (13)	SRDF_Asyn1	-IS- XIX	15	33	50	000:00:00	50000	60	I.-	---	X



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

6

The SRDF/A attributes of an RDF Group can be listed with the `-rdfa` option as shown in the slide. Note that all attributes displayed here are default values. The RDF Group has just been created and no modification to the attributes has been done.

List of an Individual RDF Group

```
C:\>symrdf -sid 483 list -rdfg 20
```

```
Symmetrix ID: 000196700483
```

Local Device View											
Sym	Sym	RDF	STATUS			MODES		R1 Inv	R2 Inv	RDF S T A T E S	
Dev	RDev	Typ:G	SA	RA	LNK	MDATE	Tracks	Tracks	Dev	RDev	Pair
0005B	00089	R1:20	RW	RW	RW	C.D1.	0	0	RW	WD	Synchronized
0005C	0008A	R1:20	RW	RW	RW	C.D1.	0	0	RW	WD	Synchronized
Total							-----	-----			
Track(s)							0	0			
MB(s)							0.0	0.0			

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

7

The list of devices in SRDF Group 20 is displayed here. The two devices 05B and 05C are R1 devices on the local Symmetrix 483. They are in Adaptive Copy Disk mode of SRDF operations, and they are currently Synchronized with their remote mirrors. As noted in the SRDF/S module, the default mode for newly created SRDF pairs is Adaptive Copy Disk. The displays in this and the previous slide are the results of the following operations (seen earlier in SRDF/S module):

Create RDF Group:

```
symrdf addgrp -label SRDF_Asyn1 -sid 483 -remote_sid 225 -rdfg 20 -
remote_rdfg 20 -dir 1E:08,3E:08 -remote_dir 1E:07,2E:07
```

Create RDF device pairs:

```
symrdf createpair -sid 483 -rdfg 20 -f pairs.txt -type R1 -establish -g
asyncdg1 (the -g asyncdg1 option adds the newly created RDF device pairs to a SYMCLI
device group -asyncdg1)
```

pairs.txt

05B 089

05C 08A

Transitioning to SRDF/A Mode

- From Synchronous mode:
 - If the devices are in Synchronized state, then the R2 data is already consistent. Enabling SRDF/A immediately provides consistent data on the R2.
- From Adaptive Copy Disk mode:
 - Any invalid tracks owed to the R2 are synchronized. Two cycle switches after Synchronization, SRDF/A provides consistent data on the R2.



SRDF/A can be enabled when the device pairs are operating in any of the listed modes. In the case of Adaptive Copy to SRDF/A transition, it takes two additional cycle switches after resynchronization of data for the R2 devices to be consistent.

Example: Synchronous Mode to SRDF/A

```
C:\>symrdf -g asyncdgl query
```

```
Device Group (DG) Name      : asyncdgl
DG's Type                   : RDF1
DG's Symmetrix ID          : 000196700483   (Microcode Version: 5977)
Remote Symmetrix ID        : 000196800225   (Microcode Version: 5977)
RDF (RA) Group Number      : 20 (13)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
Standard		ST			LI	ST				
		A			N	A				
Logical	Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv			RDF Pair
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks			MACE STATE
-----					-----					-----
DEV001	0005B	RW	0		0 RW 00089	WD	0		0 S...	Synchronized
DEV002	0005C	RW	0		0 RW 0008A	WD	0		0 S...	Synchronized

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

9

Any SRDF/A operation (with the exception of consistency exempt, discussed later in the module) must be performed on ALL devices in an SRDF group (RA group). This means that all devices in an SRDF group must be in the same SRDF Device group as well. This is in contrast with SRDF/S, where operations can be performed on a subset of devices in an SRDF group.

Example: Synchronous Mode to SRDF/A

```
C:\>symrdf -g asyncdg1 set mode async -nop
```

```
C:\>symrdf -g asyncdg1 enable
```

```
C:\>symrdf -g asyncdg1 query -rdfa
```

```
---Output Truncated---
```

```
RDFA Session Status : Active - DSE
```

```
---Output Truncated---
```

```
R2 Data is Consistent : True
```

```
---Output Truncated---
```

```
Remote Symmetrix ID : 000196800225 (Microcode Version: 5977)
```

```
RDF (RA) Group Number : 20 (13)
```

Source (R1) View				Target (R2) View				MODE
-----				-----				-----
	ST			LI	ST			
Standard	A			N	A			
Logical Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv		RDF Pair
Device Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks	MACE	STATE
-----				-----				-----
DEV001	0005B RW	0		0 RW 00089 WD	0	0	A.X.	Consistent
DEV002	0005C RW	0		0 RW 0008A WD	0	0	A.X.	Consistent

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

EMC²

10

The mode of SRDF operation is set to Asynchronous for the device pairs in the device group asyncdg1. SRDF/A consistency is enabled. symrdf query -rdfa gives detailed information about the SRDF/A state of the device group. As described earlier transition from Synchronous to Asynchronous mode is immediate. The consistency state of the R2 devices is displayed in the query: True. The RDF pair state reflects Consistent.

Example: Adaptive Copy Disk Mode to SRDF/A

C:\>symrdf -g asyncdgl query

Device Group (DG) Name : asyncdgl
 DG's Type : RDF1
 DG's Symmetrix ID : 000196700483 (Microcode Version: 5977)
 Remote Symmetrix ID : 000196800225 (Microcode Version: 5977)
 RDF (RA) Group Number : 20 (13)

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
Standard		ST			LI	ST				
		A			N	A				
Logical	Sym	T R1 Inv		R2 Inv	K Sym	T R1 Inv	R2 Inv		RDF Pair	
Device	Dev	E Tracks		Tracks	S Dev	E Tracks	Tracks	MACE	STATE	
-----					-----					-----
DEV001	0005B	RW	0	100	RW	00089	WD	0	0	CD.. SyncInProg
DEV002	0005C	RW	0	76	RW	0008A	WD	0	0	CD.. SyncInProg
Total					-----					
Track(s)			0	176				0	0	
MB(s)			0.0	22.0				0.0	0.0	

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

11

In this example, the device pairs are in SRDF Adaptive Copy Disk Mode (**C.D.**). There are R2 invalid tracks.

Example: Adaptive Copy Disk Mode to SRDF/A

```
C:\>symrdf -g asyncdg1 set mode async -nop
```

```
C:\>symrdf -g asyncdg1 enable
```

```
C:\>symrdf -g asyncdg1 query -rdfa
```

```
---Output Truncated-----
```

```
RDFA Session Number      : 19
RDFA Cycle Number        : 3
RDFA Session Status      : Active - DSE
```

```
---Output Truncated-----
```

```
R2 Data is Consistent    : False
```

```
---Output Truncated-----
```

```
Remote Symmetrix ID      : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number    : 20 (13)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
		ST			LI		ST			
Standard		A			N		A			
Logical	Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv		RDF Pair	
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks		MA	CE STATE
-----					-----					-----
DEV001	0005B RW		0		0 RW 00089 WD		0		0 A..X.	SyncInProg
DEV002	0005C RW		0		0 RW 0008A WD		0		0 A..X.	SyncInProg

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

EMC²

12

The transition into SRDF/A is immediate (**A..X.**), and the group has been enabled for consistency. However, the pair state is **SyncInProg**. R2 device does not have consistent data until the pair state is synchronized and then at least two cycle switches have completed. Consistency of the R2 data is also displayed by the highlighted field in the output: "R2 Data is Consistent". As can be seen it is "False" currently.

Legend for MODE:

M(ode of Operation) : **A** = **Async**, S = Sync, E = Semi-sync, C = Adaptive Copy

: M = Mixed

A(daptive Copy) : D = Disk Mode, W = WP Mode, . = ACp off

C(onsistency State) : **X** = **Enabled**, . = Disabled, M = Mixed, - = N/A

(Consistency) **E**(xempt): X = Enabled, . = Disabled, M = Mixed, - = N/A

Adding Devices to an Active SRDF/A Session

- Create a new SRDF device pair in a temporary SRDF Group and synchronize them
 - Synchronization can be done using the `-establish` option with the `createpair` operation
- Verify synchronization and suspend the device pair
- Move the device pair from the temporary SRDF Group into the active SRDF/A Group
 - Use the `-cons_exempt` flag with the `movepair` operation
- Resume the device pair that has been moved
- Wait for the pair state to change from Consistency Exempt to Consistent



With consistency exempt, the existing devices in an active SRDF/A session need not be suspended when adding new devices to the session. Consistency is maintained for the existing devices. The new devices are excluded from the consistency calculation until they are synchronized, move into a consistent state, and the consistency exempt attribute has been removed. Enginuity automatically clears the consistency exempt status. There is no CLI to do this. It is critical to wait for the new devices to go into a Consistent RDF Pair state, before using the R1 device for application data. As long as the Consistency Exempt attribute is set, data on the R2 is not guaranteed to be consistent with the primary data on the R1. Devices that have the Consistency Exempt attribute set can be controlled independently of the other devices in the active SRDF/A session. The operations are limited suspend, resume, and establish.

Query the Active SRDF/A Session

```
C:\>symrdf -g asyncdgl query -rdfa
```

```
--Output Truncated----
```

```
Remote Symmetrix ID      : 000196800225      (Microcode Version: 5977)
```

```
RDF (RA) Group Number    : 20 (13)
```

Source (R1) View				Target (R2) View				MODE
-----				-----				-----
Standard		ST		LI		ST		
		A		N		A		
Logical	Sym	T R1 Inv	R2 Inv	K Sym	T R1 Inv	R2 Inv		RDF Pair
Device	Dev	E Tracks	Tracks	S Dev	E Tracks	Tracks	MACE	STATE
-----				-----				-----
DEV001	0005B	RW	0	0 RW	00089	WD	0	0 A.X. Consistent
DEV002	0005C	RW	0	0 RW	0008A	WD	0	0 A.X. Consistent

```
---Output Truncated---
```

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

14

RDF Group 20 has a pair of devices that are currently in an active SRDF/A session. The objective is to add another SRDF pair to this group without affecting the consistency of the current SRDF/A session.

Create a new SRDF Pair

```
C:\>symrdf addgrp -label temp -sid 483 -remote_sid 225 -dir 1E:08,3E:08 -remote_dir 1E:07,2E:07 -rdfig 21 -remote_rdfg 21
```

```
C:\>symrdf createpair -sid 483 -f new_pair.txt -rdfig 21 -type r1 -rdf_mode sync -establish
```

```
C:\>symrdf -sid 483 -f new_pair.txt -rdfig 21 query
```

```
Symmetrix ID           : 000196700483      (Microcode Version: 5977)
Remote Symmetrix ID    : 000196800225      (Microcode Version: 5977)
RDF (RA) Group Number  : 21 (14)
```

Source (R1) View				Target (R2) View				MODE
-----				-----				-----
	ST			LI	ST			
Standard	A			N	A			
Logical Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv		RDF Pair
Device Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks	MACE	STATE
-----				-----				-----
N/A	0005D RW	0		0 RW	0008B WD	0	0 S...	Synchronized

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

EMC²

15

A new SRDF device pair is created in a different SRDF Group (Group 21). The pair state has been synchronized.

Move the New Device Pair into the Active SRDF/A Session

```
C:\>symrdf -sid 483 suspend -rdfg 21 -f new_pair.txt
```

```
C:\>symrdf -sid 483 movepair -cons_exempt -f new_pair.txt -rdfg 21 -new_rdfg 20
```

```
Execute an RDF 'Move Pair' operation for device file  
'new_pair.txt' (y/[n]) ? y
```

```
An RDF 'Move Pair' operation execution is in progress for device  
file 'new_pair.txt'. Please wait...
```

```
Move RDF Pair from (0483,021) to (0483,020).....Started.
```

```
Move RDF Pair from (0483,021) to (0483,020).....Done.
```

```
The RDF 'Move Pair' operation successfully executed for device  
file 'new_pair.txt'.
```

The EMC logo is a blue square with the letters "EMC" in white, followed by a small superscript "1".

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

16

We next suspend the link for the new device pair and move it from RDFG 21 (where it was created) to RDFG 20 which has the active SRDF/A session. We use the `-cons_exempt` flag for the `movepair` operation.

Query the SRDF/A Session

```
C:\>symrdf -g asyncdg1 query -rdfa
```

```
Device Group (DG) Name      : asyncdg1
DG's Type                   : RDF1
DG's Symmetrix ID           : 000196700483 (Microcode Version: 5977)
RDFA Session Status         : Active - DSE
RDFA Consistency Exempt Devices : Yes
--Output Truncated-----
Remote Symmetrix ID         : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number       : 20 (13)
```

Source (R1) View				Target (R2) View				MODE
-----				-----				-----
Standard	Sym	ST		LI	ST			
		A		N	A			
Logical		T R1 Inv	R2 Inv	K Sym	T R1 Inv	R2 Inv		RDF Pair
Device	Dev	E Tracks	Tracks	S Dev	E Tracks	Tracks	MACE	STATE
-----				-----				-----
DEV001	0005B	RW	0	0 RW 00089	WD	0	0 A.X.	Consistent
DEV002	0005C	RW	0	0 RW 0008A	WD	0	0 A.X.	Consistent
DEV003	0005D	RW	0	0 NR 0008B	WD	0	0 A.XX	Suspended

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

17



Query shows that the SRDF/A session is still active and that the group contains Consistency Exempt Devices. As can be seen, the mode indicates the Consistency Exempt attribute for the new device that has been added to the SRDF/A session. The existing devices continue to be in a Consistent state.

Legend for MODE:

M(ode of Operation) : **A** = **Async**, S = Sync, E = Semi-sync, C = Adaptive Copy
: M = Mixed

A(daptive Copy) : D = Disk Mode, W = WP Mode, . = ACp off

C(onsistency State) : **X** = **Enabled**, . = Disabled, M = Mixed, - = N/A

(Consistency) **E**(xempt): **X** = **Enabled**, . = Disabled, M = Mixed, - = N/A

Resume Link for the New Device Pair

```
C:\>symrdf -sid 483 resume -rdfg 20 -f new_pair.txt
```

```
C:\>symrdf -g asyncdg1 query -rdfa
```

```
Remote Symmetrix ID      : 000196800225    (Microcode Version: 5977)
RDF (RA) Group Number    : 20 (13)
```

Source (R1) View					Target (R2) View					MODE					
-----					-----					-----	-----				
		ST			LI	ST									
Standard		A			N	A									
Logical	Sym	T	R1	Inv	R2	Inv	K	Sym	T	R1	Inv	R2	Inv	RDF	Pair
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MACE	STATE				
-----					-----					-----					
DEV001	0005B	RW		0	0	RW	00089	WD		0	0	A.X.	Consistent		
DEV002	0005C	RW		0	0	RW	0008A	WD		0	0	A.X.	Consistent		
DEV003	0005D	RW		0	0	RW	0008B	WD		0	0	A.XX	SyncInProg		

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

18

We can now resume the link for the newly added device pair. Note that the query shows the state to be SyncInProg and the Consistency Exempt attribute is still set. We have to wait until the pair state goes to Consistent and the Consistency Exempt attribute is cleared, before using this device for workload.

Verify Consistency

```
C:\>symrdf -g asyncdgl query -rdfa
```

```
Remote Symmetrix ID       : 000196800225   (Microcode Version: 5977)
RDF (RA) Group Number    : 20 (13)
```

Source (R1) View						Target (R2) View						MODE			
-----						-----						-----	-----		
		ST				LI	ST								
		A				N	A								
Standard		T	R1	Inv	R2	Inv	K	Sym	T	R1	Inv	R2	Inv		RDF Pair
Logical	Sym														
Device	Dev	E	Tracks	Tracks			S	Dev	E	Tracks	Tracks			MAC	STATE
-----						-----						-----			
DEV001	0005B	RW		0		0	RW	00089	WD		0		0	A.X.	Consistent
DEV002	0005C	RW		0		0	RW	0008A	WD		0		0	A.X.	Consistent
DEV003	0005D	RW		0		0	RW	0008B	WD		0		0	A.X.	Consistent

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

EMC²

19

The Consistency Exempt flag has been removed. The pair state has moved to Consistent. Now it is safe to use the device for workload.

Removing Devices from an Active SRDF/A Session

- Suspend the device pair in the active SRDF/A session
 - Use the `-cons_exempt` flag with the suspend operation
 - Would require `-force` option if Consistency has been enabled for the SRDF/A session
- Move the device pair to a different SRDF Group
 - Use the `-cons_exempt` flag again if moving to another group with active SRDF/A session



To remove a device pair, use the `-cons_exempt` flag to first suspend the link for the devices. Then use the `movepair` operation to move the devices out of the active SRDF/A session to a different SRDF Group.

SRDF/A Configuration Parameters

- Array-wide parameters:
 - Maximum Cache Usage
 - `symconfigure -sid 483 -cmd "set symmetrix rdfa_cache_percent=50;" commit`
 - Maximum Host Throttle
 - `symconfigure -sid 483 -cmd "set symmetrix rdfa_host_throttle_time=2;" commit`
 - DSE Maximum Capacity
 - `symconfigure -sid 483 -cmd "set symmetrix dse_max_cap=1000;" commit`
- SRDF Group level parameters:
 - Minimum Cycle Time
 - `symrdf -sid 483 -rdfg 20 set rdfa -cycle_time 3`
 - Session Priority
 - `symrdf -sid 483 -rdfg 20 set rdfa -priority 20`



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

21

The array-wide parameters are set using `symconfigure` command as shown in the slide. The Group parameters for SRDF/A can be set using the `symrdf` command as shown.

Session priority = The priority used to determine which SRDF/A sessions to drop if cache becomes full. Values range from 1 to 64, with 1 being the highest priority (last to be dropped).

Minimum Cycle Time = The minimum time to wait before attempting an SRDF/A cycle switch. Values range from 1 to 59 seconds, minimum is 3 for MSC. The default minimum cycle time is 15 seconds.

SRDF/A Array-wide Parameters

- `rdfa_cache_percent`:
 - Defaults to 75, with a range of valid values from 0 to 100 percent
 - It is the percentage of the Max# of System Write Pending Slots available to SRDF/A. The purpose is to ensure that other applications can utilize some of the WP limit
 - When SRDF/A hits its WP cache limit it will be forced to drop SRDF/A sessions to free up cache
 - Setting it lower reserves WP limit for non-SRDF/A cache usage. Setting it higher, allows SRDF/A to potentially use more of the cache WP limit, potentially creating performance problems for other applications
- `rdfa_host_throttle_time`:
 - Defaults to 0, with a range of valid values from 0 to 65535
 - If >0, this value overrides the `rdfa_cache_percent` and `session_priority` settings
 - When the System WP Limit is reached, throttling will delay a write from the host until a cache slot becomes free
 - The value is the number of seconds to throttle host writes before dropping SRDF/A sessions. A value of 65535 means wait forever
- `dse_max_cap`:
 - Specifies the maximum number of GB in the specified SRP that can be used by DSE
 - Defaults to "No Limit"



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

22

Each Symmetrix has an array-wide Max # of System Write Pending Slots limit (generally calculated as 75% of available cache slots).

The purpose of this limit is to ensure that cache is not filled with Write Pending (WP) tracks, potentially preventing fast writes from hosts, because there is no place to put the I/O in cache.

SRDF/A creates WP tracks as part of each cycle.

Monitoring SRDF/A

- Using `symstat`:

- `symstat -type cache -reptype rdfa -rdfig 20 -i 60`
- `symstat -type cycle -reptype rdfa -rdfig all -i 60`
- `symstat -type request -reptype rdfa -rdfig 20 -i 60`

- Using `symevent`:

- `symevent list -error`

Examples are displayed in the following slides.

Monitoring SRDF/A

```
C:\>symstat -type cycle -reptype rdfa -rdfg 20 -i 60
```

SRDF/A Session Cycle Summary Information

Symmetrix ID : 000196700483
Timestamp : 14:14:16

		Session		Cycle Time (sec)				Cycle Size			
		RA	FLGS	Active	Last						
Device	Group Name	Grp	TAS	Cycle#	Min	Avg	Last	Switch	Active	Inact	Shared

RaGrpNum_20		20	1YA	177	15	15	15	10	18996	0	0

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations



24

This output was captured from the R1 side. The Active Cycle is the Capture and the Inactive is the Transmit, as this output is from the R1 (source) perspective.

From the R2 perspective the Active Cycle is Apply and the Inactive is Receive. The Cycle Size attribute Shared is applicable to Concurrent SRDF with both legs in SRDF/A mode. It represents the amount of shared cache slots between the two legs.

Legend for Session Flags:

T(ype) : 1 = RDF1, 2 = RDF2
A(SYNC) : Y = Yes, N = No
S(tatus) : A = Active, I = Inactive, - = N/A

Legend for the Attribute of Cycle Size:

RDF1: Active = Capture Inactive = Transmit

RDF2: Active = Apply Inactive = Receive

Monitoring SRDF/A

```
C:\>symstat -type cache -reptype rdfa -rdfg 20 -i 60
```

SRDF/A Session Cache Summary Information

```
Symmetrix ID           : 000196700483
Timestamp              : 14:27:30

System Write Pending Limit : 722439 (88.19 GB)
Cache Slots available for all SRDF/A sessions : 541829 (66.14 GB)
Total Local Write Pending Count : 23234
Total System Write Pending Count : 48356
Total System SRDF/A DSE Used Tracks : 0 (0.00 GB)
Total System Cache Slots In Use : 17604 (2.15 GB)
Total System Cache Slots Shared : 0 (0.00 GB)
Total System Cache Full Percentage : 2.0
```

		Session		Rdfa DSE			Cache Slots		Cache	
		RA	Flgs						Full	
Device	Group	Name	Grp	TAS	Pri	Thr	UsedTrks	In Use	Shared	(%)

RaGrpNum_20	20	1YA	33	50			0	17604	0	6.0

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations



25

The Cache Slots available for all SRDF/A sessions is 75% of the System Write Pending Limit (541829/722439).

Lesson Summary

During this lesson the following topics were covered:

- Multi-cycle mode for SRDF/A on VMAX3 arrays
- SRDF/A – System level and group level attributes
- Adding/removing device pairs to/from active SRDF/A session
- Monitoring SRDF/A

This lesson covered SRDF/A multi-cycle mode on VMAX3 arrays. The attributes that can be set for SRDF/A at a system and group level were discussed in detail. Method for adding/removing RDF device pairs to/from active SRDF/A sessions and monitoring SRDF/A were also presented.

Lesson: SRDF/A Resiliency Features

This lesson covers the following topics:

- Transmit Idle
- Delta Set Extension
- Group-level Write Pacing
- Recovery after link loss

This lesson covers SRDF/A resiliency features such as Transmit Idle, Delta Set Extension, and Group-level Write Pacing. Method for recovering after a link loss is also discussed.

Transmit Idle

- Transmit Idle is enabled by default when Dynamic SRDF Groups are created
- Keeps SRDF/A sessions active in the event of a “temporary link loss”
- When links are restored, data transmission resumes
- When all links are lost:
 - Data transmission from source to target is halted
 - Cycle switching continues if Source and Target arrays are VMAX3 (Multi-cycle Mode)
 - Transmit queue depth increases
 - Data accumulates in cache until SRDF/A cache usage reaches the DSE threshold
 - When maximum DSE capacity is reached, the session is dropped



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

28

SRDF/A Transmit Idle is a feature of SRDF/A that provides it with the capability of dynamically and transparently extending the Capture, Transmit, and Receive phases of the SRDF/A cycle while masking the effects of an “all SRDF links lost” event.

Without the SRDF/A Transmit Idle feature, an “all SRDF links lost” event would normally result in the abnormal termination of SRDF/A. SRDF/A would become inactive. The SRDF/A Transmit Idle feature has been specifically designed to prevent this event from occurring. Transmit Idle is enabled by default when dynamic SRDF groups are created. When all SRDF links are lost, SRDF/A still stays active.

If the Source AND the Target arrays are VMAX3 – then cycle switching continues. Multiple transmit delta sets accumulate on the source side. With VMAX3 arrays, Delta Set Extension is enabled by default. DSE will use the designated Storage Resource Pool.

Delta Set Extension

- Extends cache available for SRDF/A by off-loading cycle data from cache to disk
- VMAX3 arrays are preconfigured with one or more Storage Resource Pools (SRP) prior to installation
- One SRP is designated for DSE allocations and will support DSE for all SRDF/A sessions in the array
 - The default SRP for DSE is the default SRP for FBA devices
- Data is paged to disk when array Write Pending count crosses the DSE threshold
 - Default threshold is 50% of array Write Pending limit
- When conditions become normal, data is read back from disk to cache and transmitted to the Target array



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

29

With VMAX3, DSE pools no longer need to be configured by a user. Instead when SRDF/A spills tracks it will use a Storage Resource Pool (SRP) designated for use by DSE. Autostart for DSE is enabled by default on both the R1 or R2 sides. When running SRDF/A MCM, smaller cycles on the R2 side eliminate the need for DSE on the R2 side. Autostart is enabled on the R2 side in case there is a personality swap. Managing a DSE pool or associating a DSE pool with an SRDF group is no longer needed with VMAX3 arrays.

DSE Maximum Capacity

```
C:\>symcfg -sid 483 list -v|more
-----Output Truncated-----
SRDF/A Maximum Host Throttle (Secs) : 0
SRDF/A Maximum Cache Usage (Percent) : 75
SRDF/A DSE Maximum Capacity (GB) : NoLimit
-----Output Truncated-----

C:\>symconfigure -sid 483 -cmd "set symmetrix dse_max_cap=5;" commit

Execute a symconfigure operation for symmetrix '000196700483' (y/[n]) ? y

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

Establishing a configuration change session.....Established.
Processing symmetrix 000196700483
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Committing configuration changes.....Started.
Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations



30

The default maximum capacity for DSE is “No Limit”. In this example we are setting the maximum capacity to be 5 GB. This is now reflected in the `symcfg` output as shown below:

```
SRDF/A DSE Maximum Capacity (GB) : 5
```

Verify Designated SRP for DSE

```
C:\>symcfg -sid 483 list -srp
```

STORAGE RESOURCE POOLS

Symmetrix ID : 000196700483

C A P A C I T Y					
Name	Flg DR	Usable (GB)	Allocated (GB)	Free (GB)	Subs (%)
SRP_1	FX	57481.3	455.4	57026.0	0
Total		57481.3	455.4	57026.0	0

Legend:

Flags:

(D)efault SRP : F = FBA Default, . = N/A

(R)DFA DSE : **X = Usable**, . = Not Used

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

31

Listing of the SRP shows that SRP_1 is designated for DSE use.

Query after a Temporary Link Loss (1)

```
C:\>symrdf -g asyncdgl query -rdfa
```

```
Device Group (DG) Name      : asyncdgl
DG's Type                   : RDF1
DG's Symmetrix ID           : 000196700483   (Microcode Version: 5977)

RDFA Session Number         : 19
RDFA Cycle Number           : 907
RDFA Session Status        : Active - DSE
RDFA Consistency Exempt Devices : No
RDFA Minimum Cycle Time     : 00:00:03
RDFA Avg Cycle Time         : 00:00:03
RDFA Avg Transmit Cycle Time : 00:00:03
Transmit Queue Depth on R1 Side : 68
Tracks not Committed to the R2 Side: 426702
Time that R2 is behind R1   : 00:03:24
R2 Image Capture Time       : Thu Feb 19 12:49:47 2015
R2 Data is Consistent       : True
RDFA R1 Side Percent Cache In Use : 84
RDFA R2 Side Percent Cache In Use : 0
R1 Side DSE Used Tracks      : 115248
```

-----Output Continued on Next Slide-----

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

EMC²

32

SRDF/A session is still active. The transmit queue depth on the R1 side increases as cycle switches continue in MCM. DSE spillover has started as can be seen from the **R1 Side DSE Used Tracks**.

Query after a Temporary Link Loss (2)

---Output Continued from Previous Slide-----

```
R2 Side DSE Used Tracks      : 0
R1 Side Shared Tracks       : 0
Transmit Idle Time        : 00:03:14

Remote Symmetrix ID         : N/A           (Microcode Version: N/A)
RDF (RA) Group Number      : 20 (13)
```

Source (R1) View					Target (R2) View					MODE	
-----					-----					-----	
ST					LI					ST	
Standard	Sym	T R1 Inv	R2 Inv		N	A					
Logical	Dev	E Tracks	Tracks		K Sym	T R1 Inv	R2 Inv			RDF Pair	
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks		MACE	STATE	
-----					-----					-----	
DEV001	0005B RW	0	0 RW	00089 NA	NA	NA	NA	A.X.	TransIdle		
DEV002	0005C RW	0	0 RW	0008A NA	NA	NA	NA	A.X.	TransIdle		

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

33

The session has been in transmit idle for a little over 3 minutes and of course the pair state is reflected as TransIdle.

DSE Utilization

```
C:\>symcfg -sid 483 list -srp -demand -type slo
```

STORAGE RESOURCE POOLS

Symmetrix ID : 000196700483

Name : SRP_1
Usable Capacity (GB) : 57481.3
SRDF DSE Allocated (GB) : 5.3
Snapshots Allocated (GB) : 1.3

SLO Name	Subscribed		Allocated	
	(GB)	(%)	(GB)	(%)

Optimized	485.9	0	454.0	93
Platinum	20.0	0	0.0	0
Gold	20.0	0	0.0	0

Total	525.9	0	454.0	86

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

EMC²

34

We can also see the a little over 5 GB has been allocated so far for DSE from the designated SRP (SRP_1 in this case).

Group-level Write Pacing

- VMAX3 introduces enhanced Group-level write pacing
- Paces host I/Os to the DSE transfer rate for an SRDF/A session
- Requires HYPERMAX OS on the R1 side
 - R2 can be HYPERMAX OS or Enginuity 5876
- Responds to only to spillover rate on the R1 side
 - Not affected by spillover on the R2 side
- Device-level write pacing is not supported on VMAX3 arrays



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

35

VMAX3 introduces enhanced group-level pacing. Enhanced group-level pacing paces host I/Os to the DSE spill-over rate for an SRDF/A session.

When DSE is activated for an SRDF/A session, host-issued write I/Os are throttled so their rate does not exceed the rate at which DSE can offload the SRDF/A session's cycle data. The system will pace at the spillover rate until the usable configured capacity for DSE on the SRP reaches its limit.

At this point, the system will then either drop SRDF/A, or pace to the link rate option. To drop or pace is user definable.

All existing pacing features are supported and can be utilized to keep SRDF/A sessions active. Enhanced group-level pacing is supported between VMAX3 arrays and VMAX arrays running Enginuity 5876 with fix 67492.

Activate Group-level Write Pacing

```
C:\>symrdf -rdfg 20 activate -rdfa_wpace -nop
```

An RDF 'Write Pacing Activate' operation execution is in progress for device group 'asyncdg1'. Please wait...

Activate SRDF/A Write Pacing on (0483,020)Done.

The RDF 'Write Pacing Activate' operation successfully executed for device group 'asyncdg1'.

```
C:\>symrdf -sid 483 -rdfg 20 set rdfa_pace -wp_autostart on -nop
```

An RDF Set 'Attributes' operation execution is in progress for RDF group 20. Please wait...

SRDF/A Set Group Pacing Autostart (0483,020).....Started.

SRDF/A Set Group Pacing Autostart (0483,020).....Done.

The RDF Set 'Attributes' operation successfully executed for RDF group 20.

The EMC logo, consisting of the letters "EMC" in white on a blue square background.

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

36

The command examples show activating Group-level Write Pacing and setting Autostart "on".

Verify Group-level Write Pacing

```
C:\>symcfg -sid 483 list -rdfg 20 -rdfa
```

```
Symmetrix ID : 000196700483
```

```
S Y M M E T R I X   R D F A   G R O U P S
```

										Write Pacing		
RA-Grp	Group	Flags	Cycle	Pri	Thr	Transmit	Delay	Thr	GRP	DEV	FLG	
	Name	CSRM TDA	time			Idle Time	(usecs)	(%)	SAU	SAU	P	
20 (13)	SRDF_Asyn1	.AS- XAX		3	33	50 000:00:00	50000	60	AXX	---	X	

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

37

Details of the SRDF Group shows the current status of Group-level write pacing.

Write Pacing Flags :

(GRP) Group-Level Pacing:

(S)tatus : **A = Active**, I = Inactive, - = N/A

(A)utostart : **X = Enabled**, . = Disabled, - = N/A

S(U)pported : **X = Supported**, . = Not Supported, - = N/A

(DEV) Device-Level Pacing:

(S)tatus : A = Active, I = Inactive, - = N/A

(A)utostart : X = Enabled, . = Disabled, - = N/A

S(U)pported : X = Supported, . = Not Supported, - = N/A

(FLG) Flags for Group-Level and Device-Level Pacing:

Devs (P)aceable : X = All Devices, . = Not all devices, - = N/A

Recovering after Extended Loss of Links

- It is recommended to make a “Gold Copy” of the R2 prior to starting any resynchronization.
- In the event of an extended loss of link, a large number of R2 invalid tracks can build up on the R1 side.
- It is advisable to enable SRDF/A after the two sides are synchronized.
- Resynchronization prior to enabling SRDF/A can be performed by:
 - Setting SRDF mode to Adaptive Copy Disk Mode



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

38

As noted earlier, during resynchronization, R2 does not have consistent data. A copy of the consistent R2 data prior to resynchronization can safeguard against unexpected failures during the resynchronization process. When the link is resumed, if there are a large number of invalid tracks owed by the R1 to its R2, it is recommended that SRDF/A not be enabled right away. Enabling SRDF/A right after link resumption causes a surge of traffic on the link due to (a) shipping of accumulated invalid tracks, and (b) the new data added to the SRDF/A cycles. This would lead to SRDF/A consuming more cache and reaching the System Write Pending limit. If this happens, SRDF/A would drop again. Like with SRDF/S, resynchronization should be performed during periods of relatively low production activity.

Resynchronization in Adaptive Copy Disk mode minimizes the impact on the production host. New writes are buffered and these, along with the R2 invalids, are sent across the link. The time it takes to resynchronize is elongated.

Resynchronization in Synchronous mode impacts the production host. New writes have to be sent preferentially across the link while the R2 invalids are also shipped. Switching to Synchronous is possible only if the distances and other factors permit. For instance, if the norm is to run in SRDF/S and toggle into SRDF/A for batch processing (due to higher bandwidth requirement). In this case, if a loss of links occurs during the batch processing, it might be possible to resynchronize in SRDF/S.

In either case, R2 data is inconsistent until all the invalid tracks are sent over. Therefore, it is advisable to enable SRDF/A after the two sides are completely synchronized.

Recovery Example (1)

```
C:\>symrdf -g asyncdg1 query -rdfa
```

```
RDFA Cycle Number      : 0
RDFA Session Status    : Inactive
```

```
Remote Symmetrix ID    : N/A (Microcode Version: N/A)
RDF (RA) Group Number  : 20 (13)
```

Source (R1) View						Target (R2) View						MODE		
-----						-----						-----		
		ST				LI	ST							
Standard		A				N	A							
Logical	Sym	T	R1	Inv	R2	Inv	K	Sym	T	R1	Inv	R2	Inv	RDF Pair
Device	Dev	E	Tracks		Tracks		S	Dev	E	Tracks		Tracks		MACE STATE
-----						-----						-----		
DEV001	0005B	RW		0	16395		NR	00089	NA			NA	A.X.	Partitioned
DEV002	0005C	RW		0	16395		NR	0008A	NA			NA	A.X.	Partitioned

Total		-----				-----								
Track(s)		0				32790				NA NA				
MB(s)		0.0				4098.8				NA NA				

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

EMC²

39

In this example there is a workload on the devices in SRDF/A enabled state. A permanent loss of link places the devices in a **Partitioned** state. Production work continues on the R1 devices and the new writes arriving for the R1 devices are marked as invalid or owed to the R2. SRDF/A is dropped and session is Inactive. To get to this state, even the maximum DSE capacity should have been exceeded. So there is no choice but to drop SRDF/A.

Recovery Example (2)

```
C:\>symrdf -g asyncdgl query -rdfa
```

```
RDFA Cycle Number      : 0
RDFA Session Status    : Inactive
Time that R2 is behind R1 : 00:07:30
R2 Image Capture Time  : Fri Feb 20 12:37:16 2015
R2 Data is Consistent  : True
```

```
Remote Symmetrix ID      : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number   : 20 (13)
```

Source (R1) View					Target (R2) View					MODE				
-----					-----					-----				
		ST					LI							
		A					N							
Standard														
Logical	Sym	T	R1	Inv	R2	Inv	K	Sym	T	R1	Inv	R2	Inv	RDF Pair
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MACE	STATE			

DEV001	0005B	RW		0	16395	NR	00089	WD		0	0	A.X.	Suspended	
DEV002	0005C	RW		0	16395	NR	0008A	WD		0	0	A.X.	Suspended	

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

EMC

40

When the links are restored, the pair state goes to Suspended. Even though the flags indicate SRDF/A mode, the session status is Inactive. Also note that R2 Data is Consistent. This is because the data would be consistent up to the last (N-M-2) Apply cycle. However there are accumulated R2 Invalid tracks that are owed to the R2 side.

Recovery Example (3)

```
C:\>symrdf -g asyncdg1 disable -nop
```

```
C:\>symrdf -g asyncdg1 set mode acp_disk -nop
```

```
C:\>symrdf -g asyncdg1 resume -nop
```

```
C:\>symrdf -g asyncdg1 query -rdfa
```

```
Remote Symmetrix ID      : 000196800225    (Microcode Version: 5977)
RDF (RA) Group Number    : 20 (13)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
		ST			LI		ST			
Standard		A			N		A			
Logical	Sym	T	R1 Inv	R2 Inv	K Sym	T	R1 Inv	R2 Inv		RDF Pair
Device	Dev	E	Tracks	Tracks	S Dev	E	Tracks	Tracks	MACE	STATE

DEV001	0005B	RW	0	8091	RW	00089	WD	0	0	CD.. SyncInProg
DEV002	0005C	RW	0	8086	RW	0008A	WD	0	0	CD.. SyncInProg



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

41

As mentioned, we will next place the device group in Adaptive Copy Disk mode. As consistency was enabled when the links were lost, we have to first disable consistency before changing the mode to Adaptive Copy Disk. The RDF pair state is still Suspended. Next we resume the links. Once the RDF pair state goes to Synchronized, the mode can be changed to Asynchronous and consistency enabled.

```
symrdf set mode async
```

```
symrdf enable
```

Failover/Failback in SRDF/A Mode

- If the primary site fails, data on R2 is consistent up to the last Apply cycle (N-M-2):
 - Partial data in the Receive cycle is discarded
- SRDF failover procedure can then be executed and the workload can be started on the R2 devices:
 - Consistency protection should be disabled prior to issuing `symrdf failover` without the `-force` option
- Failback procedure after the primary site has been restored is identical to Synchronous SRDF:
 - After `symrdf failback` command completion, workload can be restarted on the R1 devices. SRDF/A can be enabled



Again, it is advisable to make a copy of the R2 prior to executing a failback operation. When workload is resumed on the R1 devices immediately after a failback, accumulated invalid tracks have to be synchronized from the R2 to the R1, and new writes must be shipped from the R1 to R2. If there is an interruption now, data on the R2 is not consistent. Even though SRDF/A can be enabled right after a failback, for reasons stated earlier, it should be enabled after the SRDF pairs entered the Synchronized state.

Lesson Summary

During this lesson the following topics were covered:

- Transmit Idle
- Delta Set Extension
- Group-level Write Pacing
- Recovery after link loss

This lesson covered SRDF/A resiliency features such as Transmit Idle, Delta Set Extension, and Group-level Write Pacing. Method for recovering after a link loss was also discussed.

Lesson: SRDF/A Multi-session Consistency

This lesson covers the following topics:

- Multiple independent SRDF/A groups
- Multi-session Consistency (MSC)
- Managing MSC



This lesson covers managing SRDF/A Multi-session Consistency.

Independent SRDF/A Groups (1)

C:\>symrdf -g asyncdgl query -rdfa

RDFA Cycle Number : 124
 RDFA Session Status : Active - DSE
 RDFA Consistency Exempt Devices : No
 RDFA Minimum Cycle Time : 00:00:03

Remote Symmetrix ID : 000196800225 (Microcode Version: 5977)

RDFA (RA) Group Number : 20 (13)

Source (R1) View						Target (R2) View						MODE		
-----						-----						-----	-----	
		ST				LI	ST							
Standard		A				N	A							
Logical	Sym	T	R1	Inv	R2	Inv	K	Sym	T	R1	Inv	R2	Inv	RDF Pair
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MACE	STATE			

DEV001	0005B	RW		0		0	RW	00089	WD		0		0	A.X. Consistent
DEV002	0005C	RW		0		0	RW	0008A	WD		0		0	A.X. Consistent

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

45

Devices in RDF Group 20 are in an active SRDF/A session. The pair state is Consistent. The current cycle number for this group is 124.

Independent SRDF/A Groups (2)

```
C:\>symrdf -g asyncdg2 query -rdfa
```

```
RDFA Cycle Number           : 94
RDFA Session Status         : Active - DSE
RDFA Consistency Exempt Devices : No
RDFA Minimum Cycle Time     : 00:00:10
```

```
Remote Symmetrix ID         : 000196800225 (Microcode Version: 5977)
RDF (RA) Group Number       : 30 (1D)
```

Source (R1) View					Target (R2) View					MODE
-----					-----					-----
		ST			LI	ST				
Standard		A			N	A				
Logical	Sym	T R1 Inv	R2 Inv		K Sym	T R1 Inv	R2 Inv		RDF Pair	
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks		MACE STATE	
-----					-----					-----
DEV001	0005D	RW	0		0 RW	0008B	WD	0	0 A.X.	Consistent
DEV002	0005E	RW	0		0 RW	0008C	WD	0	0 A.X.	Consistent

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

46

Devices in RDF Group 30 are in an active SRDF/A session. The pair state is Consistent. The current cycle number for this group is 94. The two groups switch cycles independently of each other.

Independent SRDF/A Groups (3)

```
C:\>symrdf -g asyncdgl query -rdfa
```

```
RDFA Cycle Number      : 291
RDFA Session Status    : Active - DSE
```

```
RDF (RA) Group Number : 20 (13)
```

Source (R1) View					Target (R2) View					MODE	
-----					-----					-----	
Standard		ST			LI	ST					
		A			N	A					
Logical	Sym	T Rl Inv	R2 Inv		K Sym	T Rl Inv	R2 Inv			RDF Pair	
Device	Dev	E Tracks	Tracks		S Dev	E Tracks	Tracks		MACE	STATE	
-----					-----					-----	
DEV001	0005B	RW	0		0 RW	00089	NA	NA	NA A.X.	TransIdle	
DEV002	0005C	RW	0		0 RW	0008A	NA	NA	NA A.X.	TransIdle	



© Copyright 2015 EMC Corporation. All rights reserved.


Module: SRDF/Asynchronous Operations

47

Loss of links for RDF Group 20 causes the pair states to go into Transmit Idle.

SRDF/A Multi-Session Consistency (MSC)

- Manages multiple SRDF/A sessions logically as if they were a single session:
 - RDF Daemon for Open Systems
 - Sessions can be within or across Symmetrix arrays
 - Ensures a complete, restartable point-in-time copy on the remote side



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

49

If one or more source (R1) devices in an SRDF/A MSC enabled SRDF consistency group cannot propagate data to their corresponding target (R2) devices, then the MSC process suspends data propagation from all R1 devices in the consistency group, halting all data flow to the R2 targets. RDF daemon (storrdpd) performs cycle-switching and cache recovery operations across all SRDF/A sessions in the group. This ensures that a consistent R2 data copy of the database exists at the point-in-time any interruption occurs. If a session has devices from multiple Symmetrix arrays, then the host running storrdpd must have access to all the arrays to coordinate cycle switches. It would be recommended to have more than one host with access to all the arrays running the storrdpd daemon. In the event one host fails, the surviving host can continue with MSC cycle switches.

A composite group must be created using the RDF consistency protection option (`-rdf_consistency`) and must be enabled using the `symcg enable` command before the RDF daemon begins monitoring and managing the MSC consistency group.

SRDF/A MSC

- RDF Daemon coordinates cycle switching of the SRDF/A MSC group sessions as a single entity:
 - Responsible for detecting 'failure' conditions that would cause data on the R2 side to become inconsistent
 - When a 'failure' condition is detected, the cycle switching for all SRDF/A sessions in the group are stopped in a manner that leaves the R2 side with consistent data



The RDF process daemon maintains consistency for enabled composite groups across multiple arrays for SRDF/A with MSC. For the MSC option (-rdf_consistency) to work in an RDF consistency-enabled environment, each locally-attached host performing management operations must run an instance of the RDF daemon (storrdfd). Each host running storrdfd must also run an instance of the base daemon (storapid). Optionally, if the Group Naming Services (GNS) daemon is also running, it communicates the composite group definitions back to the RDF daemon. If the GNS daemon is not running, the composite group must be defined on each host individually.

How does Multi Session Consistency work?

In MSC, the Transmit cycles on the R1 side of all participating sessions must be empty, and also all the corresponding Apply cycles on the R2 side. The switch is coordinated and controlled by the RDF Daemon.

All host writes are held for the duration of the cycle switch. This ensures dependent write consistency. If one or more sessions in MSC complete their Transmit and Apply cycles ahead of other sessions, they have to wait for all sessions to complete, prior to a cycle switch.

SRDF/A MSC Operations

- Set `SYMAPI_USE_RDFD = ENABLE` in options configuration file
- Create a Composite Group (CG) with the `-rdf_consistency` option:
 - Group definition is passed to the RDF Daemon as a 'candidate group'.
 - If the Daemon is not already running, it is started automatically.
- Add all of the devices in the multiple SRDF/A sessions to the CG.
- Put all CG devices into Async mode:
`symrdf -cg <CGname> set mode async`
- Enable CG devices for consistency protection:
`symcg -cg <CGname> enable`
 - The RDF Daemon is notified that the group should now be monitored.
 - Enable command must be done after the devices are put into Async mode.
- When the devices become RW on the link, the RDF Daemon:
 - Starts performing cycle switching.
 - Actively monitors the health of the group to maintain R2 data consistency



The option to use RDF daemon has to be enabled in the SYMAPI options file. RDF daemon should of course be started and running. Managing MSC requires the creation of Composite Groups. When the Composite Group is enabled, the cycle switching is now controlled by the RDF daemon.

SRDF/A MSC

```
C:\>symrdf -g asyncdg1 disable -nop
```

```
C:\>symrdf -g asyncdg2 disable -nop
```

```
C:\>symdg dg2cg asyncdg1 msc_cg -rdf_consistency
```

```
C:\>symdg dg2cg asyncdg2 msc_cg -rdf_consistency -rename
```

```
C:\>symcg list
```

COMPOSITE GROUPS

Name	Type	Valid	Number of			Number of			
			Symms	RAGs	DGs	Devs	BCVs	VDEVs	TGTs
msc_cg	RDF1	Yes	1	2	0	4	0	0	0

```
C:\>symcg -cg msc_cg enable -nop
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

52

The objective is to manage our two SRDF/A groups as a single entity using MSC. We first disable consistency for the two groups and then add them to a consistency group as shown in the slide. Next we enable MSC for the CG.

Query of the Composite Group

```
C:\>symrdf -cg msc_cg query -rdfa
```

```
Composite Group Name      : msc_cg
Composite Group Type      : RDF1
Number of Symmetrix Units : 1
Number of RDF (RA) Groups : 2
RDF Consistency Mode     : MSC
```

```
RDFA MSC Info
```

```
{
  MSC Session Status      : Active
  Consistency State       : CONSISTENT
}
```

```
RDF (RA) Group Number    : 30 (1D)
```

```
RDFA Info:
```

```
{
  Cycle Number            : 15
  Session Status          : Active - MSC - DSE
}
```

```
RDF (RA) Group Number    : 20 (13)
```

```
RDFA Info:
```

```
{
  Cycle Number            : 15
  Session Status          : Active - MSC - DSE
}
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

53

The cycle numbers for the two groups have been reset to be the same. MSC has been enabled.

SRDF/A MSC at Work

```
C:\>symrdf -cg msc_cg query -rdfa
```

```
RDF (RA) Group Number      : 30 (1D)
Source (R1) View           Target (R2) View      MODE
-----
Standard      ST          LI          ST
              A            N            A
Logical  Sym   T R1 Inv  R2 Inv  K Sym   T R1 Inv  R2 Inv  RDF Pair
Device  Dev   E Tracks  Tracks  S Dev   E Tracks  Tracks  MACE STATE
-----
DEV003  0005D RW          0          0 NR 0008B WD          0          0 A.X. Suspended
DEV004  0005E RW          0          0 NR 0008C WD          0          0 A.X. Suspended
```

```
RDF (RA) Group Number      : 20 (13)
Source (R1) View           Target (R2) View      MODE
-----
Standard      ST          LI          ST
              A            N            A
Logical  Sym   T R1 Inv  R2 Inv  K Sym   T R1 Inv  R2 Inv  RDF Pair
Device  Dev   E Tracks  Tracks  S Dev   E Tracks  Tracks  MACE STATE
-----
DEV001  0005B RW          0  16394 NR NA      NA      NA      NA A.X. Partitioned
DEV002  0005C RW          0  16393 NR NA      NA      NA      NA A.X. Partitioned
```

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations



54

The loss of links for RDF Group 20 results in suspending the links for RDF Group 30, even though its links are still available. Note that the output is very verbose and has been edited to show the relevant details for this example. When the links are restored, recovering from this state can be accomplished with:

```
symrdf -cg rdfa_msc_cg establish
```

Once the invalid tracks are marked, merged, and synchronized, MSC protection is automatically re-instated; i.e., user does not have to issue `symcg -cg rdfa_msc_cg enable` again.

MSC Cleanup

- Cleanup is automatically performed by the RDF Daemon on the R1 side, if the link to the R2 side is available.
- If the link is unavailable (total site failure on the R1 side), then invocation of any SRDF command, such as symrdf failover or split, from the R2 side performs the automatic cleanup.

Cleanup is automatically performed by the RDF Daemon if the link to the R2 side is available.

Lesson Summary

During this lesson the following topics were covered:

- Multiple independent SRDF/A groups
- Multi-session Consistency (MSC)
- Managing MSC

This lesson covered managing SRDF/A Multi-session Consistency.

Lab: SRDF/Asynchronous Operations

This lab covers

- Single Session SRDF/Asynchronous groups
- Configuring Concurrent SRDF
- Configuring and managing SRDF/A Multi-session Consistency

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

57

This lab covers setting SRDF/Asynchronous mode of operation for SRDF device pairs and enabling consistency protection. It also covers configuring Concurrent SRDF with one leg in SRDF/Synchronous mode and the other in SRDF/Asynchronous mode. Configuring and managing SRDF/A Multi-session consistency is covered as well.

Lab Summary

This lab covered

- Single Session SRDF/Asynchronous groups
- Configuring Concurrent SRDF
- Configuring and managing SRDF/A Multi-session Consistency

This lab covered setting SRDF/Asynchronous mode of operation for SRDF device pairs and enabling consistency protection. It also covered configuring Concurrent SRDF with one leg in SRDF/Synchronous mode and the other in SRDF/Asynchronous mode. Configuring and managing SRDF/A Multi-session consistency was covered as well.

Module Summary

Key points covered in this module:

- SRDF/Asynchronous remote replication on a VMAX3 array
- SRDF/A operations
- SRDF/A resiliency features
- SRDF/A Multi-session Consistency



This module covered SRDF/Asynchronous mode of remote replication. Concepts and operations for SRDF/A in single and multi-session modes were presented. SRDF/A resiliency features were also discussed.

This slide intentionally left blank.

Course Summary

Key points covered in this course:

- TimeFinder SnapVX operations with SYMCLI and Unisphere for VMAX – in physical and virtualized environments
- SRDF operations in Synchronous (SRDF/S) and Asynchronous (SRDF/A) modes using SYMCLI and Unisphere for VMAX – in physical and virtualized environments

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: SRDF/Asynchronous Operations

61

This course covered performing TimeFinder SnapVX and SRDF operations in Synchronous and Asynchronous modes for business continuity management on VMAX3 arrays.



This concludes the Training. Thank you for your participation.