

SRDF Foundations

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

- Identify the basic features of SRDF
- Describe the technical positioning of SRDF
- Describe the architectural components and theory of operations of SRDF
- List several SRDF management tools
- Describe several examples of how SRDF can be integrated into a Symmetrix environment

The objectives for this module are shown here. Please take a moment to read them.

The Business Tool

BENEFITS OF SRDF

© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 2

 Mercer Road
EMC Education Services

This section discusses SRDF as a business tool.

Symmetrix Remote Data Facility (SRDF)

- What is SRDF
 - It is a business tool that allows a client with Symmetrix based data centers to copy their data between the sites for a number of purposes
- SRDF offers various levels of Symmetrix based business continuance and disaster recovery solutions
- SRDF products offer the capability to maintain multiple, host-independent, mirrored copies of data
- SRDF allows for the Symmetrix systems to be in the same room, in different buildings within the same campus, or hundreds to thousands of kilometers apart

Symmetrix Remote Data Facility is a business tool that allows a client with Symmetrix based data centers to copy their data between the sites for a number of purposes.

SRDF – A Remote Data Replication Business Tool

- By maintaining copies of data in different physical locations, SRDF enables you to perform the following operations by integrating with your strategies for:
 - Disaster restart, Disaster restart testing
 - Recovery from planned outages, Remote backup
 - Data center migration, Data replication and mobility
- SRDF helps you meet your objectives
 - Recovery Point Objective
 - Point in time to which systems and data must be recovered after an outage
 - Recovery Time Objective
 - Period of time within which systems, applications, and functions must be recovered after an outage

By integration with business solutions for the tasks listed in the slide, SRDF becomes an integral part of the overall business solution strategy.


SRDF will help you meet your Recovery Point Objective and Recovery Time Objective.

The Recovery Point Objective, RPO, is the point in time to which systems and data must be recovered after an outage (for example, the end of the previous day's processing). RPOs are often used as the basis for the development of replication and backup strategies and as a determinant of the amount of data that may need to be recreated after the systems or functions have been recovered.

Recovery Time Objective, RTO, is the period of time within which systems, applications, and functions must be recovered after an outage (for example, one business day). RTOs are often used as the basis for the development of recovery strategies, and as a determinant as to whether or not to implement the recovery strategies during a disaster situation. Similar terms include Maximum Allowable Downtime.

SRDF Benefits

- Protect against local and regional site disruptions
 - Continuous data availability
 - Multiple remote recovery sites
 - Meet regulatory requirements
- Provide near-instant recovery
- Migrate, consolidate or distribute data across storage platforms
 - Data center consolidations
 - Technology refreshes
- Enable non-stop operations
 - Application restart across volumes



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 5

Mercer Road
EMC Education Services

The right remote replication solution can limit your exposure to planned and unplanned downtime by enabling operations at remote sites. Perhaps you need to provide your organization with efficient data replication to meet corporate or government standards, while still meeting your Total Cost of Ownership requirements. In addition, you need a flexible solution that changes as your needs change.


No matter what the challenge is, there is one underlying theme: data protection and faster business restart, in the event of a disaster or unplanned outage, is critical across the organization.

There are several methods of remote replication available, each with a different profile to help you manage and conquer your current business challenges, while enabling new processes and procedures that will help you gain a significant competitive advantage, something all businesses strive for in today's competitive marketplace.



The SRDF Product Family

- Offers solutions for many different disaster recovery and business continuance scenarios
- Consists of three base solutions
 - SRDF/S
 - SRDF/A
 - SRDF/DM
- Consists of four add-on solutions
 - SRDF/Star
 - SRDF/CG
 - SRDF/AR
 - SRDF/CE



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 6



The SRDF family of products offers solutions for many different disaster recovery and business continuance scenarios. The base solutions include SRDF/S, SRDF/A, and SRDF/DM. Add-On solutions include SRDF/Star, SRDF/CG, SRDF/AR, and SRDF/CE. For simplicity, this module refers to this product simply as "SRDF".

SRDF/Synchronous (SRDF/S) maintains a real-time (synchronous) copy of data at the logical volume level in Symmetrix or Symmetrix DMX™ systems.

SRDF/Asynchronous (SRDF/A) is another mode of remote replication that allows customers to asynchronously replicate data while maintaining a dependent write consistent copy of the data at the secondary site at all times. It uses Delta Set technology.

SRDF/Data Mobility (SRDF/DM) permits operation in SRDF adaptive copy mode only and is designed for data replication and/or migration between two or more Symmetrix systems.

Optional, Add-on Functionality

SRDF/Star provides advanced multi-site business continuity protection available. It enables concurrent SRDF/S and SRDF/A operations from the same source volumes with the ability to incrementally establish an SRDF/A session between the two remote sites in the event of a primary site outage.

SRDF/Consistency Groups (SRDF/CG) is a product offering designed to ensure the consistency of the data remotely mirrored by the SRDF operations, in the event of a rolling disaster.

SRDF/Automated Replication (SRDF/AR) is an automation solution that uses both SRDF and TimeFinder to provide a periodic asynchronous replication of a restartable data image.

SRDF/Cluster Enabler (SRDF/CE) for MSCS or VCS provides high availability and automated failover through storage-based replication and server clustering through SRDF/S and MSCS or VCS.

Internal Mechanisms

INTERNAL CHARACTERISTICS OF SRDF

Having briefly defined what SRDF is, we will now discuss the internal mechanisms of the tool.

Symmetrix Remote Data Facility

Data mirroring between physically separate Symmetrix systems

- Foundation for Symmetrix based disaster restart solutions
- Host operating systems, DBMS, file-system independent
- Bi-directional source-to-target(s) architecture
- Architecture uses delta technology (Track Tables)
 - only changed data copied




© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 8

Mercer Road
EMC Education Services

SRDF provides comprehensive business continuity and restart capabilities for both planned and unplanned outages. This online, host-independent, mirrored data solution, duplicates production site data on one or more physically separate target Symmetrix systems. Systems can be across the room, across the globe, or anywhere in between.

What makes this technology a leader in the industry is its use of track tables. While we will not get into track table details, the point is that we have a unique capability to send only changed information at the block level, thus allowing us to operate in several modes, such as those listed on the right, which we will address later in this presentation.




This shows the representation of the mirror positions for a DMX system when both the Source and the Target SRDF Logical Volumes have local protection (RAID-1).

In this diagram, the Target-R2 volume is also represented with 4 mirror positions and has local mirroring implemented, but only three of the mirror positions are used. The third mirror position (M3) is a pointer back to the Source Symmetrix Logical Volume, while the first (M1) and second (M2) mirror positions represent local mirrors. If a BCV is established with the R2 volume, then it will occupy the next available mirror position (M4).

Under normal circumstances, the R1 volume presents a Read-Write (RW) status to the host which accesses it, and the R2 presents Write-Disabled (WD) to its host.

SRDF Modes of Operation

- Modes
 - Synchronous Replication
 - Semi-Synchronous Replication
 - Adaptive Copy Replication
 - Asynchronous Replication
- Operational Modes can be changed dynamically using CLI or other EMC tools
- All methods of replication can co-exist within the Symmetrix array
- You can specify the operational method on a per device basis
- Performance, level of synchronization, and I/O serialization requirements determine appropriate operational mode



Performance Vs Availability

The modes of operation are Synchronous Replication, Semi-Synchronous Replication, Adaptive Copy Replication, and Asynchronous Replication.


Synchronous Replication (SRDF/S) provides real-time mirroring of data between the source Symmetrix and the target Symmetrix systems.

Semi-Synchronous Replication (SRDF/S) writes data to the source system, completes the I/O, and then synchronizes the data with the target system.

Adaptive Copy Replication (SRDF/DM) transfers data from the source devices to the remote devices without waiting for an acknowledgments.

Asynchronous Replication (SRDF/A) places host writes into cycles or "chunks" and then transfers an entire chunk to the target system. These operational modes only affect those Symmetrix volumes that are remotely mirrored. I/O operations with local volumes occur as if they are in a non-SRDF environment as no updates are required to volumes at a different physical location.

There is a preferred mode of operation for each SRDF solution. These operational modes are selectable based on requirements of performance, distance, and speed of recovery. All methods of replication can co-exist in a Symmetrix array. SRDF/S and SRDF/DM allows you to specify the method on a per device basis. SRDF/A is set at the RA Group level, not the volume level. No special application coding is required and no CPU overhead is incurred when implementing SRDF.




SRDF Synchronous Mode is used primarily in SRDF campus environments. In this mode of operation, the Symmetrix maintains a real-time mirror image of the data of the remotely mirrored volumes.

Data on the source (R1) volumes and the target (R2) volumes are always fully synchronized at the completion of an I/O sequence through a first-in, first-out queue (FIFO) model. All data movement is at the block level with synchronized mirroring.

The sequence of operations is:

- The write I/O is received from host/server into cache of source and placed in FIFO queue
- The write I/O is transmitted to the cache of the target
- A receipt acknowledgment is provided by the target back to the cache of the source
- An ending status is presented to the host/server

The SRDF mode can be changed dynamically, depending on such factors as production activity level, time of day, bandwidth concerns, temporary re-structuring of environment, etc. It is very flexible.



SRDF Semi-Synchronous Mode is used primarily in extended distance environments. In this mode of operation, data on the remotely mirrored volumes are always synchronized between the source (R1) volume and the target (R2) volume, prior to initiating the next write operation to these volumes. As with Synchronous Mode, the data movement is at the block level and it works on a first in, first out model.


The sequence of operations is:

- The write I/O is received from the host/server into the cache of the source
- Ending status is presented to host/server and the write I/O is placed in FIFO queue
- The write I/O is transmitted to the cache of the target. Note: A second write will not be accepted on a Symmetrix source device until its target device has been synchronized.
- Once the device is synchronized, a receipt acknowledgment is provided by the target back to the cache of the source

The target is, at most, one I/O operation behind per source logical volume.

Semi-Synchronous Mode masks the impact of distance in the general case because it allows read operations while write operations are in transit. Again, SRDF uses a first-in, first-out queue.

It is important to note that Semi-Synchronous mode is not supported on DMX3.



SRDF Adaptive Copy Mode allows the source (R1) volumes and target (R2) volumes to be a few or many I/Os out of synchronization. Unlike Synchronous and Semi-Synchronous, it does not use the first in-first out model, target data is only usable after a full synchronization. This operational mode is not recommended for use when mirroring for disaster recovery/restart purposes, unless used with TimeFinder. This mode is used primarily for data migrations and data center moves

There are two types of adaptive copy: Write Pending Mode and Disk Mode. Disk Mode data movement is handled at the track level. Data from the cache is moved to the R1 volume, then keeps the track-level information as to what data is owed to the remote side so that information can be subsequently sent one track at a time to the source device. Pending Mode data accumulates in a local cache. A background process moves the data tasks to the source device and its corresponding target device. The advantage of this method is that it is typically faster to read data from cache than from disk. The disadvantage is that cache is temporarily consumed by the data until it moves to disk. Note that both methods allow write tasks to accumulate on the local side before being copied to the remote side.

The sequence of operations is:

- An I/O write is received from the host/server into the cache of the source
- An ending status is presented to the host/server
- I/O is placed in the SRDF queue
- I/O is destaged and processed using Write Pending Mode or Disk Mode
- The I/O is transmitted to the cache of the target
- A receipt acknowledgment is provided by the target back to the cache of the source

Summary – SRDF Level of Synchronization

- **Synchronous Mode**
 - Source = Target
- **Semi Synchronous Mode**
 - Source $\cong$ Target
 - At most Source is 1 I/O ahead of Target, per volume
- **Adaptive Copy**
 - Source $\neq$ Target
 - Source may be up to 99999 tracks per volume ahead of Target
- **SRDF/Asynchronous**

In review, when the Synchronous mode of operation has been chosen, the source data (R1) and target data (R2) will be guaranteed synchronized at any point in time. The R2 is a mirror of R1. None of the other modes of operation guarantee that the data is the same at any point in time.

Semi-Synchronous mode says that the source (R1) and target (R2) can be up to one I/O out of synch, and the data is serial since it arrives at the target Symmetrix in the same sequence that it arrived (from the host) in the source Symmetrix.


Adaptive Copy mode does not offer synchronization or serialization. As data is received in the source Symmetrix, it is queued for delivery to the remote Symmetrix. The order that it was received at the source Symmetrix is not maintained, and the source and target volumes can be many tracks out of synch.

SRDF/Asynchronous will be discussed later in this module.

Connectivity Options

SRDF CONNECTIONS AND CONFIGURATIONS

We will now discuss the connectivity options of SRDF.



A Remote Adapter is made up of a Remote Link Adapter (RLA) and a Remote Link Director (RLD). RLAs provide the link connection and fiber optic protocol support between local and remote Symmetrix units. An RLA is the same hardware component as a serial channel adapter board, but RLAs use SRDF microcode.


RLDs provide communications control between the local and remote Symmetrix units. An RLD is the same hardware component as a serial channel director (ESCON), but RLDs use SRDF microcode. This modification made to create an RLD adds functionality, but still allows the RLD to be used as a normal ESCON director, if desired. One of the things accomplished with the modifications is the ability to transmit Open Systems FBA data over the ESCON-based SRDF link.

Volumes are added to SRDF groups, this represents a logical connection between each Symmetrix and the volumes assigned to the group. One type of SRDF group is an RA group. The Remote Link Director may be assigned to the RA group. Groups are defined at configuration time.

The ESCON protocol (developed by IBM) was designed for mainframe to peripheral communications. It is, therefore, typically a one-way communication from mainframe to disk. Because of this design, some additional overhead is present when implementing SRDF with ESCON because SRDF can be implemented bi-directionally (sources and targets in both Symmetrix).



Remote Fibre Director (RFD)



- Uses the Fibre Channel protocol to facilitate data transfer between two Symmetrix systems
- Fibre adapter cards designated only as Remote Adapter Fibre (RAF) rather than "RA1" or "RA2" within ESCON
 - Fibre Channel communication is peer-to-peer rather than ESCON "master" and "slave"
 - Terms "Source Volume(s)" and "Target Volume(s)" remain however
- Ability to use a free RAF processor for Fibre Channel host connection
- Two types of RAF cards are available
 - Both ports support multi-mode fiber optic transmission (500 m max. with 50 micron cable)
 - One port supports multi-mode; one port supports single-mode fiber optic transmission

© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 17




A Remote Fibre Director (RFD) is an unmodified Fibre Channel Director running SRDF code.

The additional overhead described with the ESCON protocol does not exist with the Fibre Channel standard. The Fibre Channel peer-to-peer relationship eliminates the potential 20% overhead induced by an ESCON "peripheral" requesting permission to talk to the ESCON "host" in ESCON-based SRDF. This increases the flexibility of SRDF in cases where it is desired to have source and target volumes located at each side of the SRDF link.

SRDF over Fibre Channel uses a Symmetrix RF director type. Remote Fibre Director comes in various port and processor combinations. The RF director is the same hardware as the FA host interface director, but is configured to run SRDF emulation code. It has the ability to use a free RAF processors for Fibre Channel host connections. RF directors use a SCSI 3 layer, running FC-4 specification Point-to-Point protocol. Connectivity in this environment may be via direct multimode cable up to 500 meters, or single-mode cable up to 10 km between RF directors in two Symmetrix systems. Fibre Channel switches may also be used for SRDF, with either multimode or single-mode cable segments between units.



GigE Remote Directors and Multiprotocol Channel Director (MPCD)



- Gigabit Ethernet (GigE) enables direct Symmetrix-to-IP network attachment
- GigE can be used in local or long-distance environments, no need for protocol-conversion devices
- Gigabit Ethernet adapter cards designated as Remote Adapter RE
- Communication is peer-to-peer
- Two types, MPCD & GigE card
 - MPCD for DMX, two-port and four-port models: combination FICON/GigE or iSCSI/GigE
 - MPCD has the ability to use ports for host interface emulation
 - GigE Remote Directors (RE), two-processor, two-port, for use in Symmetrix 8000

© 2006 EMC Corporation. All rights reserved. SRDF Foundations - 18 Mercer Road EMC Education Services

GigE provides Native IP support for any SRDF based product on Symmetrix systems. It is based on GigE technology that enables direct Symmetrix system-to-IP-network attachment. This increases the options for Symmetrix system to Symmetrix system connectivity, and allows a Symmetrix system to connect to existing Ethernet infrastructure and directly access high-speed data transmission conduits via IP.


Symmetrix DMX series systems provides this connectivity using the Multiprotocol Channel Director (MPCD). Unused SRDF ports can be configured with host interface emulation (iSCSI or FICON) while the other port is configured for SRDF emulation. Alternatively, both ports can be configured for SRDF GigE operation.

The Symmetrix 8000 series provides this connectivity through the GigE remote director.

The MPCD and GigE remote directors provide comparable functionality, with the exception of data compression, which is a feature of the MPCD only.

Note that GigE ports have higher latency than Fibre Channel ports, which will impact synchronous performance in campus configurations.

GigE, unlike Fibre Channel or ESCON, can be used in local or long-distance environments, without the need for protocol-conversion devices. Since the Symmetrix GigE director uses native IP protocol, traffic can be routed through an IP cloud of various connectivity devices over almost any distance.




SRDF offers three types of configuration: Uni-Directional, Bidirectional, and Dual Configuration.

- Uni-directional is a one way mirror relationship
- Bi-directional is a two-way mirror relationship
- Dual Configuration is two uni-directional configurations

Bi-Directional Configuration would be supported in a Campus Mode due to the overhead associated with RA1/RA2.

Additionally, these configurations can be implemented with Escon Directors, Remote Fibre Directors, or GigE Directors.


SRDF Flexible Configuration Options



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 20  Mercer Road
EMC Education Services

SRDF connections can be configured in many different ways. The current EMC Support Matrix should be referenced for information on the latest supported configurations.



Switched SRDF enables Symmetrix systems running SRDF to be connected through the Connectrix family of switches and simultaneously benefit from EMC's Enterprise Storage Network.

At the same time, switched SRDF enables the Symmetrix to be configured in either a "fan out" or "fan in" configuration. Although the SRDF volume pairing is fixed, that is, the source and target volumes are not changeable, the customer can share the bandwidth of Fibre Channel across fewer connections. This is possible due to the implicit multiplexing that occurs in the Fibre Channel protocol.

Users benefit in that fast, optics-based infrastructures can now be used in campus and metropolitan areas. Configurability is improved, while information access is made that much faster and easier.

More applications can now benefit from SRDF's disaster recovery, business continuance, improved risk management, and data replication characteristics.

Physical Link Connectivity

Connection Type	Cable	Distance	Product
ESCON	Direct Fiber, Multimode 62.5/125um cable.	3 Km.	Symmetrix 5000/8000/DMX
ESCON	With Extenders	66Km.	Symmetrix 5000/8000/DMX
ESCON	With DWDM	~200Km.	Symmetrix 8000/DMX
ESCON	T1/E1, T3/E3, ATM, IP	Unlimited	Symmetrix 5000/8000/DMX
Fibre	Direct Fiber, Multimode 50/125um cable	500 Meters	Symmetrix 5000/8000/DMX
Fibre	Direct Fiber, Single mode 9um cable	10 Km.	Symmetrix 8000/DMX
Fibre	FCSW to FCSW (single-mode thru-out path)	~60Km.	Symmetrix 8000/DMX
Fibre	DWDM	~200 Km.	Symmetrix 8000/DMX
GigE	Direct Fiber, Multimode 50/125um cable	500 Meters	Symmetrix 5000/8000/DMX
Fibre	With LAN/WAN	Unlimited	Symmetrix 5000/8000/DMX

Symmetrix systems, using SRDF, communicate over a variety of connectivity types.

The table above shows various cabling options and the distances associated with each. It is important to remember that connectivity types and distances are constantly changing. Refer to the latest connectivity material on the Support Matrix for updates. Different Symmetrix models support different connection types - refer to the SRDF Product/Connectivity Guide for reference.


Link Operations

SUSPEND AND RESUME OPERATIONS

This next section, Link Operations, discusses the Suspend and Resume Link operations.

Suspend Link

- Logically suspends mirror relationship between source and target volumes



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 24 Mercer Road
EMC Education Services

Suspending SRDF volumes on SRDF Link logically suspends the mirror relationship between source and target volumes.


The R1 is still Read/Write and the R2 is still Write Disabled to their locally attached hosts. New writes to the R1 device accumulate as invalid tracks owed to the R2 device.

Note the link is still active for other SRDF volumes, only the SRDF volumes "suspended on the Link" are affected.

When the suspend has completed successfully, the devices will be suspended on the SRDF links and their link status set to Not Ready (NR).

Resume Link

- Resumes mirror relationship between source and target volumes



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 25

Mercer Road
EMC Education Services

Resuming the link operation resumes I/O traffic across all link paths between the selected source (R1) and target (R2) volumes.

When resuming the volume on the link, accumulated invalid tracks from the R1 devices are sent to the R2 devices. Until these updates have completed, the R2 can not be considered in a recoverable state.


Disaster Recovery Operations

DISASTER RECOVERY WITH SRDF

Next, we will discuss Disaster Recovery Operations.

SRDF Failover

- Change status of source volumes to Write Disable
- Suspend link
- Change status of target volumes Read/Write



© 2006 EMC Corporation. All rights reserved.


SRDF Foundations - 27 Mercer Road
EMC Education Services

The Failover process makes the copy of data on the target Symmetrix volumes (R2) available to attached hosts. It could be utilized during a disaster scenario such as a host channel, Symmetrix, or site failure. It could also be used during maintenance activities.

The Failover process causes the target (R2) volumes to take over read/write operations for source (R1) volumes. This operation halts all I/O activity to the Symmetrix unit containing the source (R1) volumes; this will write-disable the source (R1) volumes. This operation is typically performed when you need to transfer I/O operation from the source (R1) to the target (R2) volumes.

SRDF Failback

- Changes status of target volume to Write Disable, resumes link, synchronize R2 with R1, then write enables source volume




© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 28 

The Failback process, or returning control to the local host, resumes I/O operations with the source (R1) volumes after a period of performing I/O operations with the target (R2) volumes. This operation halts all I/O activity to the Symmetrix unit containing the target (R2) volumes; this will write-disable the target (R2) volumes. This operation is typically performed when returning to normal SRDF operation after a target takeover.

Updating the Source Volumes

- An optional step prior to resuming normal operations on the source volume is to synchronize R1 invalid tracks to a pre-specified level



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 29  Mercer Road
EMC Education Services

Update the source (R1) side with the changes from the target (R2) side while the target (R2) side is still operational to its local host. This brings the R1 side close to being synchronized with the R2 side before a failback.

Simultaneous Operations

OPERATING ON BOTH SIDES OF THE SRDF LINK


© 2006 EMC Corporation. All rights reserved.



Simultaneous Operations is when both sides of the SRDF mirrored pair are placed into the Read-Write Enabled state. When the mirrored pair is split, the link is logically suspended and the changes are tracked for both the source and target volumes. Reintegrating the pair can either update the source with target data, or update the target with source data.

Simultaneous Operations – Split

- Suspend link between source (R1) and Target (R2) Volumes
- Enables read and write operations on both R1 and R2 volumes



© 2006 EMC Corporation. All rights reserved.


SRDF Foundations - 31  Mercer Road
EMC Education Services

The split process stops remote mirroring between the source (R1) device and the target (R2) device. The target device is made available for local host operations.

On a split operation, the link is suspended between source (R1) and target (R2) volumes, then read and write operations are enabled on both R1 and R2 volumes. The R1 and R2 must be synchronized prior to the process, suspending the link and making R1 and R2 read/write.

Simultaneous Operations – Save Source Data

- Resume SRDF operation retaining data from source and overwriting any changed data on target



© 2006 EMC Corporation. All rights reserved.


SRDF Foundations - 32  Mercer Road
EMC Education Services

In this scenario, the user was upgrading their database software and has split an SRDF pair to upgrade/test the new code (R2). The database upgrade failed. They have decided to keep the source side (the non-upgraded side) as the production environment until they confer with the database company, try the upgrade, and test when the issues are sorted out.

They would save the source data to the target (R2) side. In open systems, this process is referred to as an establish. The link is brought back up, and changed tracks on the source side are flushed to the target. Tracks changed on the target are overwritten with the information from the source side, even if the source tracks were never changed. Tracks left untouched on both sides are not re-synched. The incremental operation is achieved by comparing the track table for the source to the table for the target.

Simultaneous Operations – Save Target Data

- Resumes SRDF operation retaining data on target and overwriting any changed data on source



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 33 Mercer Road
EMC Education Services

Alternatively, the upgrade could have been successful. Rather than re-upgrading the source side, it is quicker and easier to make the target side Write Disables, make the source side Read/Write, and flush all changed tracks on the target to the source and overwrite tracks changed on the source with tracks from the target. In open systems, this process is referred to as a restore.

In the example used, you could apply transaction logs to bring the newly upgraded environment up-to-date with transactions processed in production on the old database, while the upgrade was being tested.

Consistency Groups

DATABASE MANAGEMENT


This next section discusses the importance of Consistency Groups.



SRDF/Consistency Groups (SRDF/CG)

Maintaining Data Consistency

Mainframe, Open System, Windows, Enterprise



- Enabling technology for disaster restart solutions
- Data consistency preserved during rolling disaster

© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 35




SRDF Consistency Groups allow customers to define logical volume groups, which can be associated with a given workload. These groups of SRDF logical volumes will automatically be suspended in case of SRDF communications failures. The remote SRDF logical volumes will be consistent, even if these logical volumes span multiple Symmetrix systems, such as when a large database has its tables on one Symmetrix and its log files on another.

Features include the ability to do an “explicit trip” of the affected consistency group. An explicit trip means that SRDF transfers to the R2 side are programmatically stopped – leaving the R2 volumes in a “consistent” state as of the last transfer of information. As such, the R2 volumes can be made re-startable in a test bed or recovery procedure. Meanwhile, production data can continue to be written to the R1 side. The R2 side will be updated later.

Understanding Consistency

- **The Dependent Write I/O Concept**
 - An application concept where the start of one write is dependent on the completion of a previous write
 - A logical dependency, not a time dependency
 - Inherent in all Database Management Systems (DBMS)
 - ❖ Page (data) write is dependent write I/O based on a successful log write
 - Applications can also use this technology
 - Power failures create a dependent write consistent image
 - Restart transforms dependent write consistent to transactional consistent
- **Creating Dependent Write Consistency**
 - Database technologies
 - Hot Backup Mode, Suspend/Resume
 - Enterprise TimeFinder Consistent Split
 - Enterprise SRDF Consistency Groups


Dependent write is a write I/O operation which depends on a previous I/O write to be completed before its execution. All logging database management systems use this concept to maintain integrity. This is required for the protection against local power outages, loss of local channel connectivity, or storage devices. This is a logical dependency between I/Os built into database management systems, certain applications, middleware tools such as MQ Series, and operating systems.



It is important to understand rolling disasters in order to understand how EMC technology prevents them. On this diagram, there is a host, two source Symmetrix's, and one target Symmetrix. A Database Management System is running on the host. The logs of the DBMS are on the bottom source Symmetrix and the table space of the DBMS is on the top source Symmetrix.

In the example, the SRDF link from the bottom source Symmetrix and the target Symmetrix breaks. This is the beginning of the rolling disaster. However, the table space updates are propagated to the target Symmetrix, even though log updates are not.

Now the rolling disaster completes. When the restart of the DBMS occurs on the target side, it restarts fine. There are no errors, but data is corrupt. There is information in the table spaces that is not reflected in the log. This is known as 'data ahead of log'. If the table spaces and logs were reversed, it would be referred to as a 'log ahead of data'. The only way to detect that the data is corrupt is to be extremely familiar with the transactions and know what was processing at the exact time the rolling disaster occurred. This is almost impossible in most environments.



To prevent corruption of data, EMC created Consistency Groups to use with SRDF. ConGroups is a started task on the mainframe and is package in EMC Consistency Group for z/OS. On Open Systems, ConGroups work through the SYMCLI/SYMAPI and an RDF process daemon or SYMCLI / SYMAPI and PowerPath. A ConGroup contains a set of volumes, created by the customer, and needs consistency maintained. In this example, our host could be Mainframe or Open System and includes volumes "A", "C" and "E". PowerPath is used on Open Systems, EMC Consistency Group for z/OS is used on the mainframe.

With the ConGroup defined, our rolling disaster begins with the loss of the SRDF links from bottom source Symmetrix to the target Symmetrix. A sense code is sent back to the host stating that the data from volume "A" could not be propagated to its target side. The ConGroup started task on the mainframe, or the ConGroup definition in Open Systems, detects the sense code and works with IOS Mainframe or PowerPath for Open Systems to hold the I/O. The DBMS, or application, is not aware the I/O was held and created a Dependent Write Consistent copy, or DBMS restartable copy of the data on the target side. We have simulated a local power failure at the target side at the point of the beginning of the rolling disaster. After complete failure at the source side, the target side can be restarted and the DBMS can be restarted. This will provide transactional consistency.



SRDF/A

ASYNCHRONOUS SRDF



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 39 

This next section discusses the SRDF/A mode of operation.



SRDF/Asynchronous (SRDF/A)

- High Performance Remote Replication for the DMX
- Asynchronous remote mirroring
 - Recoverable point-in-time copy
 - No impact to production applications
 - No distance limitations
- Operational savings through reduced bandwidth
 - Delta Sets vs. ordered writes
- Supports Mainframe and Open Systems
- Complements existing SRDF solutions
 - Meet a wide range of RPO and RTO service level requirements
- Mixed SRDF, SRDF-DM and SRDF/A
 - Share links and directors



Highest Performance Asynchronous Replication
In The Industry

© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 40



SRDF/Asynchronous (SRDF/A) is a remote mirroring solution for the Symmetrix DMX family. Its unique architecture delivers a remote mirroring solution that has no impact on production applications and no distance limitations. This unique architecture enables significant operational savings through reduced bandwidth requirements; you can size your bandwidth to your average peak workload, versus your absolute peak workload.

SRDF/A is a single solution supporting both Mainframe and Open Systems attach. It also complements SRDF solutions to meet mixed service level requirements. In fact, it can also share the same communication links as SRDF.

The unique architecture uses cache-based “Delta Sets”, which will be detailed shortly

SRDF/Asynchronous Environments

- All existing SRDF topologies (ESCON, Point-to-Point, Switched Fabric Fibre Channel, and Gig-E)
- Any host (MF, OS, Midrange) and any data type recognized by Symmetrix (FBA, CKD, etc.)
- Support for all currently EMC supported transports (DWDM, FCIP, ATM)
- Support for all environments with all Symmetrix-qualified channel extenders, switches etc. (SAN)

SRDF/A is supported by various operating systems, topologies, and hosts.

This encompasses all existing SRDF topologies:

- ESCON
- Fibre Channel
- Gig-E

Any host:

- Mainframe
- Open System

Any data type recognized by Symmetrix, including:

- FBA
- CKD

Key operating systems, including:

- UNIX (Sun, HP, IBM)
- Windows (NT, 2000, 2003)
- IBM Mainframe (z/OS)
- Support begins with Symm 5670 microcode and carries forward to future generations of Symmetrix

Terms & Concepts

- **Delta Set**


- A collection of incoming writes maintained in cache
 - Capture (Source, aka the N cycle)
 - Transmit (Source, aka the N-1cycle)
 - Receive (Target, aka the N-1cycle)
 - Apply (Target, aka the N-2 cycle)

- **Write Folding**

- Consolidates the writes in cache so that the Capture Delta Set can be sent across the SRDF/A link only once.
- This reduces overhead of transmission, increasing band width.

A Delta Set is the innovative architecture that SRDF/A is built upon. A Delta Set is simply a collection of writes that are part of an SRDF/A process. There are four types of Delta Sets which will be described in more detail later in this training. Their names are Capture, Transmit, Receive, and Apply. In some of our literature, you may see the term “cycles,” as in cycle N, cycle N-1, etc., as shown on the slide.

Applications tend to write data to the same location on a device. When this occurs, it is called locality of reference phenomenon. A process called Write Folding will identify the locality of reference phenomenon. Write Folding consolidates the writes in cache so data that is updated multiple times in the same Capture Delta Set is sent across the SRDF/A link only once. This reduces overhead of the transmission and increases the band width.



Let's take a very high level look at SRDF/A's innovative architecture. SRDF/A uses "Delta Sets" to maintain a group of writes over a short period of time. Delta Sets are the enabler of all the efficiencies that SRDF/A can deliver. There are four types of Delta Sets to manage the data flow process. SRDF/A's data flow can be summarized in simple steps.


Source side Delta Sets:

- Capture - Captures, in cache, all incoming writes to the source volumes involved in the SRDF/A group. Upon completion of the set, the Capture Delta Set is "write folded" and promoted to a Transmit Delta Set. A new, separate Capture Delta Set is then created to maintain the next Delta Set of writes.
- Transmit - Transfers its contents (only the last set of writes) from the source to the target system.

Target side Delta Sets:

- Receive - is on the target system and receives the data being transferred by the source-side Transmit Delta Set. Once received in its entirety, it is promoted to the Apply Delta Set.
- Apply - Applies the Delta Set's writes to the target volume to create the consistent, recoverable remote copy. This finishes the Delta Set cycle.

A consistent, recoverable copy is maintained at the remote location with each application of a Delta Set. Once the initial Delta Set cycle completes, it is simply repeated, Delta Set after Delta Set, providing a continuous checkpoint of Delta Sets. SRDF/A provides a solution for service level requirements that need data on the R2 side within the seconds to minutes range.



Achieving data consistency across multiple SRDF/A groups simply requires that the cycle process described above be coordinated among the participating Symmetrix. SRDF/A Multi Session Consistency coordinates the cycle of Capture, Transmit, Receive, and Apply across the multiple Symmetrix R1 source and R2 target relationships.

When the R1 Symmetrix finishes sending the inactive cycle to the R2 Symmetrix, the R1 Symmetrix will also send an internal message to the R2 side Symmetrix that the transfer is complete. In a similar fashion, when the R2 side Symmetrix finishes restore of the active cycle, it will send a message to the R1 Symmetrix that the restore is complete. The host will receive status messages indicating cycle processing readiness, the host will then coordinate the execution of the cycles by sending a cycle switch command to all R1 side Symmetrix.

The single host coordination point drives the cycle process in all participating Symmetrix systems. In a mainframe, this is provided by ResourcePak Base for z/OS executing Symmetrix Control Facility, SCF. In an open systems environment, cycle switching is provided by Solutions Enabler, executing MSC within a RDF process daemon.

 	
<i>SRDF/Asynchronous Features and Functionality</i>	
Technology	Delta Sets with Continuous Checkpoints
Systems supported	All DMX
Required bandwidth	Continuous support for “average” load, peaks absorbed in cache
Consistent, re-startable copy at all times	Yes
Multiple Symmetrix DMX system support	Yes
Primary host can read through to Target	Yes
Maximum distance between systems	trans-oceanic or trans-continental distances
Bi-directional mirroring	Yes
Support for dynamic and concurrent	Yes
RAID support options for Source/Targets	None, RAID-0, Parity RAID, RAID 1, RAID 5
Maximum SRDF/A groups per system	64
Management	Mainframe, Solutions Enabler & ControlCenter
Integration with TimeFinder	Full
Integration with EMC Clone	Full
Integration with EMC Snap	Limited

This chart details SRDF/A’s specific features and functionality.

SRDF/Asynchronous Considerations

- Will require additional cache to support “Delta Sets”
 - Two Delta Sets maintained at each location
 - ~.75 GB additional cache per 1 TB of data being remotely mirrored per DMX as a minimum, assuming bandwidth sized to peak write load
 - Much more cache could be required if bandwidth sized closer to average write I/O load on busy DMXs
- Bandwidth must be sized relative to cache for maximum performance/minimum Delta Set size

Some key considerations for SRDF/Asynchronous are:

- Each DMX system will also require additional cache to support the cache-resident Delta Sets. A general rule of thumb is ~.75 GB of additional cache per 1 TB of SRDF/A protected data.
- The balancing of cache to bandwidth for maximum performance is also critical.

Implementation Permutations


SRDF IMPLEMENTATIONS

We will now discuss some of the implementation configurations of the tool.

Campus Implementations

- The campus solution enables units to be up to 66 km apart

- **ESCON Direct Attach**
 - Distance of up to 3 km
 - Throughput up to 14 MB/s
- **Direct Fibre Channel Attach (Point-to-Point)**
 - Distance of up to 10 km
 - Throughput up to 55 MB/s
- **Switched Fibre Channel Attach**
 - Distance of up to 10 km
- **GigE**
 - Enables direct Symmetrix-to-IP network attachment
 - Allows Symmetrix to access existing Ethernet infrastructure




© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 48 Mercer Road
EMC Education Services

There are several SRDF Campus implementations available that enable units to be up to 66 km apart:

- **ESCON Direct Attach** — This connection type provides connectivity between two Symmetrix unit ports at a distance of up to 3 km for each cable segment. Throughput for each connection is typically rated at up to 14 megabytes per second (MB/s) maximum, and in most cases is relative to block size of the host write I/O operation.
- **Direct Fibre Channel Attach (Point-to-Point)** — This connection type provides connection via direct cable segment attached between two Symmetrix unit ports at a distance of up to 10 km for each cable segment (single mode hardware), or 500m for each cable segment (multimode hardware, default). Throughput for each connection is typically rated at up to 55 MB/s maximum, and in most cases is relative to the block size of the host write I/O operation.
- **Switched Fibre Channel Attach** — This connection type provides connections between two Symmetrix unit ports by way of Fibre Channel switch ports, at a distance of up to 10 km for each cable segment, or 500m for each cable segment (multimode hardware, default).
- **GigE Attach** — This connection type provides native IP support for any SRDF-based product on Symmetrix systems, which is based on Gigabit Ethernet (GigE) technology. This enables direct Symmetrix-to-IP network attachment. This increases the options for Symmetrix-to-Symmetrix connectivity, and allows a Symmetrix system to connect to an existing Ethernet infrastructure, and to directly access high-speed data transmission conduits via IP.




The use of a Metropolitan Area Network (MAN) is determined by the ability of the customer or carrier to install/provision “dark fiber” cable that allows the customer to use the entire bandwidth on that fiber.

MANs are useful for campuses with multiple data centers, ISPs that have several sites along a fiber optic network right-of-way, and XSPs that need to reach multiple customers in the area. This means that the maximum bandwidth is available and that the delay or latency is low and controllable. In the future, optical technology will expand to reach beyond the 200 km distance.

Extended Distance Implementations

- The Extended Distance, wide area network solution, enables units to be at trans-oceanic or trans-continental distances
- Communication options include:
 - T1/T3, E1/E3
 - Synchronous Optical Networks (SONET)
 - ATM-OC3
 - IP



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 50

Mercer Road
EMC Education Services

The Extended Distance Wide Area Network (WAN) solution enables units to be at trans-oceanic or trans-continental distances.


Typically, ATM, T3/E3, T1/E1, or Internet Protocol (IP) are offered by the carrier.

Communication options include:

- T1/T3, E1/E3
- Synchronous Optical Networks (SONET)
- ATM-OC3
- IP

Ideal SRDF Business Continuance Solution

- Determined by
 - Recovery Point Objective
 - Recovery Time Objective
 - How well your applications tolerates network latency



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 51

Mercer Road
EMC Education Services


When picking the ideal business continuance solution, the selection is determined by the distances separating your data centers, the Recovery Point Objective, Recovery Time Objective, and how well your applications tolerate network latency. Shorter distances reduce network latency, which allows using synchronous disk replication and data center mirroring. Business continuance solutions can be divided into three general categories, based on the distances between the local and remote data center:

- Campus solution is a limited subset of metro connectivity that uses fiber-optic cabling to transmit data over short distances using Symmetrix and Connectrix direct-attach capabilities. Typically the distance is smaller than 66 km
- Metropolitan Area Network (MAN) provides SRDF connectivity for distances typically less than 100 km, but up to 200 km or more. This environment is characterized by extremely low error rates and high bandwidth of fiber configurations, so error recovery and compression are less significant in choosing SRDF connectivity options.
- Extended Distance Wide Area Network (WAN) provides SRDF connectivity over long distances using telecommunications networks, such as IP, SONET, or ATM. WANs are differentiated from MANs as being non-lossless and limited bandwidth environments, where error-recovery, data buffering, and compression capabilities of connectivity.



SRDF Data Mobility (SRDF-DM)

- Economical 'Data Replication' solution not intended to be used for Disaster Recovery operations
- SRDF Data Mobility is just SRDF running in Adaptive Copy Disk Mode or Adaptive Copy Write Pending Mode
 - Adaptive Copy Disk Mode is recommended for optimal performance.
- The Primary Mode (Synchronous or Semi-Synchronous) is still set on the Symmetrix Logical Volume, but the user cannot disable the Secondary Mode
- Skew limit cannot be set to below 100 tracks per Symmetrix Logical Volume
- Minimum skew limit is set when the microcode is installed



© 2006 EMC Corporation. All rights reserved. SRDF Foundations - 52 

SRDF Data Mobility is an economical 'data replication' solution and is not intended for use as a disaster recovery solution. Unlike host level replication technologies, it does not require additional servers or network capability to perform unlimited distance data replication.


SRDF-DM uses the same user controls as SRDF (Solutions Enabler SYMCLI commands, and EMC ControlCenter). The only difference is:

- SRDF-DM uses the skew parameter; the skew parameter defines the number of tracks the source volume can lead the target volume. This limit cannot be set by the user to a value lower than 100 for SRDF-DM
- The user cannot disable the Secondary mode by setting a skew limit of zero. Setting the skew limit to zero would essentially mean running in Primary mode

These SRDF limitations are enforced by the microcode setting in the Source and Target Symmetrix.

Why not just use SRDF?

- SRDF-DM can be used for data movement between separate sites
- SRDF-DM uses storage infrastructure to move data between separate sites, while offloading the network and server workload
- SRDF-DM is an economical mobility solution, not requiring a full SRDF license
- It is suitable as a data replication solution where the cost of dedicated fast links for Synchronous Mode is not viable



Concurrent SRDF allows two remote SRDF mirrors of a single R1 device, e.g., use one remote copy for disaster recovery, and another for decision support or backup.

The above example shows a concurrent SRDF configuration in which the primary volume is communicating with one secondary volume in synchronous mode. Concurrently, the same primary volume is communicating with its other secondary volume in one of the adaptive copy modes (adaptive copy write-pending mode or adaptive copy disk mode).

Each Remote Link Director is assigned to an RA Group. With ESCON, only one RA group per RLD is allowed, but with Fibre Channel SRDF, the RA Groups can be defined to the same RLD.

Any combination of synchronous/semi-synchronous and adaptive copy is allowed with the following exceptions:


- One volume operating in synchronous mode and the other operating in semi-synchronous mode
- Both volumes operating in asynchronous mode

A write I/O from the host at the primary device side cannot be returned as completed until both remote Symmetrixes signal the local Symmetrix that the SRDF I/O is in cache at the remote side.



SRDF/STAR – Symmetrix Triangular Aynchronous Replication

- One source, two targets
 - One target in Synchronous mode
 - Short distance, zero data lag
 - One target in Asynchronous mode
 - Longer distance, variable data lag, no performance impact
- SRDF/STAR: Remote Re-synchronization
 - Continued protection upon source failure




© 2006 EMC Corporation. All rights reserved. SRDF Foundations - 54 Mercer Road EMC Education Services

Building on the foundation of a near and a far site using SRDF/S and SRDF/A respectively, the remaining problem is that the far and near sites contain different data (the near site is a zero-data-loss situation, and the far site is seconds to minutes behind.)

To resolve these problems, a new product, SRDF/Star, has been created. SRDF/Star is a solution that provides a third copy of System Data in case one site was lost and not recoverable. SRDF/Star completes the connection to the two non-primary sites (the bunker and the long-distance site) with an SRDF/A link. If the primary site is lost to a catastrophic event, business may resume/continue using data from either the bunker or the long-distance site. SRDF/CG Consistency Group and SRDF/A Multi Session Consistency are used to integrate the SRDF/Star configuration.

SRDF Family replicated data can be used from whichever remaining Symmetrix (near or far site) is the most current. Realizing that determining which site has the better data for a restart is very difficult; therefore, EMC provides a utility to aid in the determination of the status of data in the two remaining sites. This allows you to make an incremental synchronization decision of which direction to synchronize in prior to a restart. This link between the near and far with SRDF/Star will maintain continuous protection for the remaining sites.



Prior to Dynamic SRDF, the R1 and R2 pairings were static and defined in the BIN during SDRF configuration. Any changes to SRDF device pairing required a new BIN file to be defined and loaded into the Source and Target Symmetrix.

Dynamic SRDF, available since 5x68 code, provides the capability to change device pairings as needed, without requiring a BIN file configuration change to be performed by EMC. R1 to R2 pairing will be dynamic and can be changed by the user, much like Standard to BCV relationships.

In this diagram, an R1 can have a remote mirror in Symmetrix B and then later be defined to mirror the data to Symmetrix C instead.

Another scenario, after breaking the R2 in Symmetrix B from its original R1, the R2 can be reused as the target for an R1 in Symmetrix C.

In conjunction with Symmetrix Manager, this provides the user with maximum flexibility to make ad-hoc changes to the remote mirroring protection for a volume.

Starting with Engenuity code 5669, a user can create Dynamic SRDF Groups and associate dynamic SRDF pairs to the group. Dynamic SRDF groups provide flexibility within the SRDF environment to change multiple remote mirroring connections for dynamic devices.

SRDF Integration with TimeFinder

Integrating SRDF & TimeFinder Operations with SRDF/Automated Replication


Let's explore how SRDF can be integrated with TimeFinder.



SRDF/Automated Replication (SRDF/AR)

SRDF/S & TimeFinder/Mirror

- Masks performance impact of distance
- Minimizes network costs
- Allows business restart site to be at any distance away from source



Extended Distance • **No performance Impact** • **Rapid Business Resumption**
Zero Data Loss • **Minimize communications cost**

© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 57 


SRDF/AR pauses transaction updates to a database, spanning multiple volumes or systems, while a point-in-time copy of the database is being made. This ensures that the database contents aren't altered during the copy operation so that the copy is recoverable and re-bootable. SRDF/AR provides a solution for service level requirements that need data on the R2 side within the minutes to hours range.

SRDF/AR and the automated Multi-Hop command eliminates the multiple, separate SRDF and TimeFinder commands previously required to implement these capabilities, when establishing BCVs for remote replication. The same command can be used for Consistent Split and Multi-Hop. In particular, the symreplicate command benefits Multi-Hop users greatly because it now makes a Multi-Hop configuration very easy to implement.

SRDF/Automated Replication

- The continuous movement of dependent write consistent data to a remote site in an asynchronous mode
 - Standard to BCV/R1 => R2 to BCV Single Hop
 - R1 => R2 to BCV/R1 => R2 to BCV Multi-Hop
- Utilizes Enterprise TimeFinder Consistent Split and SRDF Adaptive Copy Mode
- Parameters available to manipulate cycles
 - Mainframe
 - EMC Symmetrix Control Facility
 - UNIX/Windows
 - SYMCLI

SRDF/AR enables customers to deploy a long distance remote mirroring solution to optimize their business critical information processes and resources in a cost-effective way.




This environment is comprised of a source (production) site and a target (restart) site. The hosts that are attached can be mainframe, UNIX, and/or Windows.

The source site shows the standard volumes established with the BCV/R1 devices, which naturally suspends the BCV/R1 to R2 synchronization. The target site shows that the R2s are split from their BCVs.

The standard volumes at the source site represent the dependent write consistent copy.

The standard devices depict that logical consistency is maintained by utilizing TimeFinder Consistent Split.

The dependent write consistent copy could exist at any of these points throughout the cycle.



This slide represents a multi-hop configuration, which includes a source site, bunker site, and target site. In a multi-hop configuration, there is synchronous propagation from the source site to the bunker site, and adaptive copy propagation from the bunker site to the target site. There is usually an SRDF Consistency Group defined to maintain consistency, in the event of a rolling disaster between the source and bunker sites.

The function of the bunker site is to transform synchronous propagation to an adaptive copy propagation of dependent write consistent copies. The distance from the bunker site to the target site is usually a long distance and utilizes the adaptive copy mode of SRDF. The propagation between the bunker site and the target site was once a manual process and done without maintaining consistency.



SRDF Cluster Integration

- SRDF/Cluster Enabler for Microsoft Cluster Server
 - Available from EMC
- SRDF/Cluster Enabler for VERITAS Cluster Server
 - Available from EMC
- SRDF/Automated Availability Manager for EMC SRDF
 - EMC FullTime AutoStart
- EMC AutoSwap, redirects IO to the R2 devices in the event that the R1 devices are unreachable
- HP MetroCluster and Continental Cluster
 - Available from HP
- IBM HACMP/HAGEO
 - Available from IBM - requires user scripting



© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 61 

SRDF is increasingly being used in cluster environments. In a cluster, the server and application providers have usually worked out the issues of failing over from one server to another at the application level. This means they can restart the work on the surviving node in the cluster. However, if the failed node contained the disk storage and the working data being used by the application, the problem would be the surviving node could restart the applications, but only with new data.

SRDF's value in this scenario is in the mirrored data from one Symmetrix to another within the cluster. Using SRDF, the surviving node also has all the data because it was previously in an SRDF relationship. Procedures to make the R2 volumes read/write enabled have to be taken. In some cases, this task is automated. This is the case with SRDF/Cluster Enabler for MSCS, the Microsoft Cluster Server.

When the time comes to bring the failed node back online, SRDF provides the ability to fail back in the other direction and be ready for work very quickly, with up-to-the-second live data.

EMC AutoSwap works in a mainframe Parallel Sysplex environment with SRDF/S and SRDF/CG. AutoSwap will automatically make the R2 devices R/W and re-direct I/O to the R2 devices in the event that the R1 devices are unreachable by the local and remote nodes.

Management Options

MANAGING SRDF

This next section discusses the management options of SRDF.

EMC Education Services

EMC²
where information lives[®]

Accessing SRDF via EMC ControlCenter

- SRDF Manager within EMC ControlCenter
 - Easy, “point-and-click” access
 - Excellent for ad hoc SRDF operations

© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 63

Mercer Road
EMC Education Services

SRDF can be accessed via EMC ControlCenter. Through the graphical user interface (GUI) within EMC ControlCenter software, related devices are grouped together in device groups. SRDF operations may be performed on all devices in a device group using a single command, and group information is maintained in the SYMAPI Database.

Accessing SRDF via Solutions Enabler

- **Solutions Enabler SRDF Feature for Open Systems**
 - Based on SYMAPI
 - Consistent command syntax
 - Functions implemented using binary commands
 - Used for scripting on Open Systems
- **Symmetrix Control Facility for z/OS and SRDF Host Component for z/OS started tasks.**
 - Operator Commands for single commands
 - ISPF Panel menus
 - Batch JCL for automation with multiple commands.
- **Device Groups / Composite Groups**
 - A collection of devices assigned to a named group to provide a more manageable object to query status and impart control operations
 - Devices can be associated as either a device group or a composite group

Open Systems Solutions Enabler allows access through the command line interface (CLI) and through Open Systems scripting.

Mainframe Symmetrix Control Facility for z/OS and SRDF Host Component for z/OS started tasks allows operator commands for single commands, ISPF Panel menus, and Batch JCL for automation with multiple commands.

Device Groups and Composite Groups hold collections of devices. The groups are used to manage devices pairs, query status and impart control operations. Device Groups define devices only on a single Symmetrix array. Composite Group defines devices that can span multiple Symmetrix arrays. A composite group provides greater flexibility than a device group because the devices span multiple Symmetrix array.

SRDF Foundations

BUSINESS JUSTIFICATION

We will now discuss the business justification of SRDF.

EMC Education Services

EMC²
where information lives[®]

The CIO's Information Storage and Management Requirements

Lower TCO
Utilization, consolidation,
automated management

**Access and
availability**
Information growth >50%
annually

Compliance
>16,000 regulations
worldwide

Simplicity
Hundreds of information
management tools

Protection/recovery
Matching availability
to requirements

The right local replication solution will help you overcome these challenges


© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 66 Mercer Road
EMC Education Services

As EMC's customers work to support their overall business, they have to be particularly focused on information storage and management.

In addition, there are new compliance requirements, new levels of protection and recovery. It's not just about backup; it's actually how customers can restore and get back into business at the different service levels based on the application need.

All of this needs to be done at lower cost. Every customer continues to be asked to do either more with less or more with the same. This is not going to change. You will continue to need to extract the most value out of every investment you make. You're going to look for fast payback periods. You're going to look for a measurable and demonstrable return on investment. The right remote replication solution will help you overcome these challenges.



The key to determining which remote replication solution is best for a particular application is a thorough understanding of service levels. Specifically, you will need to balance your requirements for performance, functionality, availability, and economics against capacity needs, bandwidth requirements, and overall Total Cost of Ownership. Every solution will have benefits and risks, and all of these factors must be balanced in order to decide the best solution for a given problem.

EMC Education Services

EMC²
where information lives[®]

Remote Replication Considerations

The diagram illustrates the components of remote replication. It shows a 'Source' server on the left and a 'Target' server on the right. A solid blue arrow points from the Source to the Target, representing the replication path. Along this path, four starburst shapes highlight key considerations: 'Response time' (with a clock icon), 'Infrastructure' (with a server rack icon), 'Comms' (with a network cable icon), and 'Application consistency' (with a document icon). A dashed orange line also connects the Source and Target servers.

- **Response time**—Application impact of replication
- **Infrastructure**—Costs to enable solution
- **Communications**—Monthly line and bandwidth costs
- **Application consistency**—Application inter-dependency

© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 68

Mercer Road
EMC Education Services

Today's demands require customers to invest in more complex data infrastructures. These data processing environments require a focused attention to such components as application response times, application consistency, and the costs associated with the infrastructure and communications.

Service Level Requirements			Solutions		Limitations & Requirements	
RPO	RTO	RGO	Product / Solution	Mode	Max. Distance	Link Bandwidth
No data loss	Medium	Multi-location	SRDF/STAR	Synchronous and Asynchronous	Unlimited	Highest
No data loss	Fastest	Short distance	SRDF/S, Congroup & AutoSwap	Synchronous mode with Continuous Availability	~200Km	High
No data loss*	Fast	Long distance	<u>SRDF/AR</u> <u>Multi Hop</u>	Synchronous mode + Consistent Point in Time Copy	Up to 200Km for synch. Unlimited for adcopy.	High Low
Sec / Min	Fast	Long distance	SRDF/A with MSC	Continuous Asynchronous Copy	unlimited	Medium or High
Min / Hours	Fast	Long distance	<u>SRDF/AR</u> <u>Single Hop</u>	Consistent Point in time copy	unlimited	Low

© 2006 EMC Corporation. All rights reserved.

SRDF Foundations - 69


 Mercer Road
EMC Education Services

If an outage occurs on Source Symmetrix, when it comes to business continuity and remote mirroring, the time it takes to get the business running again - Recovery Time Objective (RTO) - is always “fast”, no matter which EMC replication solution is deployed.

What is the customer’s recovery point objective (RPO)? How much data exposure can their business tolerate in return for such characteristics as distance, performance, and economics?

EMC has the right solution for each particular requirement. If the customer can’t tolerate any data exposure, we have the industry leader for synchronous mirroring - SRDF. However, as with any synchronous solution, there are characteristics that must be understood – distance is limited by application time-outs and speed of light issues and bandwidth must be sized for peak, workload at all times.

We can also deliver solutions that combine SRDF with TimeFinder to create single-hop and multi-hop environments for specialized needs. These solutions offer different RPOs and have different requirements for bandwidth, supported distances, etc.

SRDF/A provides a solution for service level requirements that need RPOs in the seconds to minutes range. SRDF/AR provides a solution for service level requirements that need RPOs in the minutes to hours range.

- RPO = Recovery point Objective, how much data can be lost
- RTO = Recovery Time Objective, How quickly can data be recovered
- RGO = Recovery Geographic Objective

SRDF – Benefits Review

- **Protect against local and regional site disruptions**
 - Continuous data availability
 - Multiple remote recovery sites
 - Meet regulatory requirements
- **Migrate, consolidate or distribute data across storage platforms**
 - Data center consolidations
 - Technology refreshes
- **Enable non-stop operations**
 - Application restart across volumes

The right remote replication solution can limit the exposure to planned and unplanned downtime, enabling non-stop operations. Or perhaps you need to provide your organization with efficient data replication to meet corporate or government standards, while still meeting your Total Cost of Ownership requirements. In addition, you need a flexible solution that changes as your needs change.

No matter what the challenge is, there is one underlying theme: data protection and faster business restart, in the event of a disaster or unplanned outage, is critical across the organization.

There are several methods of remote replication available. Each has a different profile to help you manage and conquer your current business challenges, while enabling new processes and procedures that help gain a significant competitive advantage. This is something all businesses strive for in today's competitive marketplace.

Module Summary

Key points covered in this module:

- SRDF consists of three base solutions
 - SRDF/S
 - SRDF/A
 - SRDF/DM
- SRDF consists of four add-on solutions
 - SRDF/Star
 - SRDF/CG
 - SRDF/AR
 - SRDF/CE
- SRDF implementations are configured as Campus, MAN, and Extended Distance implementations
- SRDF primary modes of operation are Synchronous, Semi-synchronous and SRDF/Asynchronous

These are the key points covered in this training. Please take a moment to review them.