

**SRDF/Star Software Uses and
Best Practices**

EMC Proven Professional Knowledge Sharing 2009



Bill Chilton
Chilton_Bill@emc.com

Table of Contents

SRDF/Star Software Uses and Best Practices	1
Disaster Recovery	4
What is an SRDF/Star Configuration?.....	4
What is a Star	6
Mirror Positions	7
What is Needed to Build an SRDF/Star Environment?	7
GNS Recommendations.....	8
Build the SRDF/Star Options File	8
Best Practices for SRDF/Star.....	10
Never Control C when in SRDF/Star	10
Do not Script an Unplanned Switch or Failover.....	10
When Faced with a Real Disaster: Isolate	11
Two Workstations at Each Site or Data Center Location (six servers in total)	11
Build a SRDF/Star Configuration for Testing Purposes.....	12
Have Enough Gatekeepers.....	12
Check Daemons	13
Star Naming Convention	13
Load Balancing Applications Across Data Centers.....	15
Eliminating Downtime while Working on Servers.....	16
Migrating Data with no Downtime while Staying Consistent.....	17
Whose Data should I Keep at Site A, B or C's.....	21
New Cascading SRDF/Star Feature.....	22
Cascaded SRDF	24
Location of Information in Powerlink	25
New Commands (for Cascading, Star and Cascading Star)	27
Troubleshooting Tips	28
Verbose -v.....	28
Turn your SYMAPI verbose logging on:	28
Running SRDF/Star in Debug Mode	28

Log Files	29
Check daemons	30
Allowable States for symstar Operations.....	30
Summary of symstar Control Operations* (taken from product guide)	32
Biography.....	35

Disclaimer: The views, processes or methodologies published in this compilation are those of the authors. They do not necessarily reflect EMC Corporation's views, processes, or methodologies

Disaster Recovery

How do you keep your information safe in a world filled with natural disasters and the continuous threat of manmade disasters? Companies and countries are developing disaster recovery plans to respond to these catastrophes. Initially, disaster plans were developed to enable data recovery at another site and have your systems back up and running in a matter of hours or days following data center destruction. Today, this time frame is unacceptable; our customers cannot afford to wait hours or days to come back online. They need their data centers up and running within minutes, even seconds.

Enterprises are ensuring that their information is available by building redundant data centers in different areas of the country, and in different countries in the world to ensure that if one area experiences problems, the data center in another part of the world will come back online. The largest financial institutions are building three redundant data centers to handle their information to ensure that their data is always available and online. They are positioning two data centers in close proximity to each other, with the third in a different region of the country or in a different country.

How can customers have their data centers up and running in a highly compressed time frame? One way is to adopt new recovery software such as EMC's SRDF[®]/Star, Cascaded SRDF and Cascaded SRDF/Star. In this article, I offer best practices and many tips on the best way to install and run these products at customer sites.

What is an SRDF/Star Configuration?

A Symmetrix Triangular Asynchronous Replication (STAR) Basic Configuration is composed of three arrays located in data centers that are united using Symmetrix SRDF/Star software. This software allows data to replicate concurrently from the primary or production worksite to the synchronous or secondary site located near the production site, and the asynchronous or remote site located anywhere in the world.

This technology provides consistent data protection and incremental data recovery in the event that the primary worksite experiences a disaster or data failure. SRDF/Star allows fast recovery and restartable worksites at either the secondary or remote site. The customer has the ability to determine which site they would like to make the new primary site in the case of a primary worksite failure. The customer can also determine which data they would like to keep and propagate at the new primary site (see whose data should I keep).

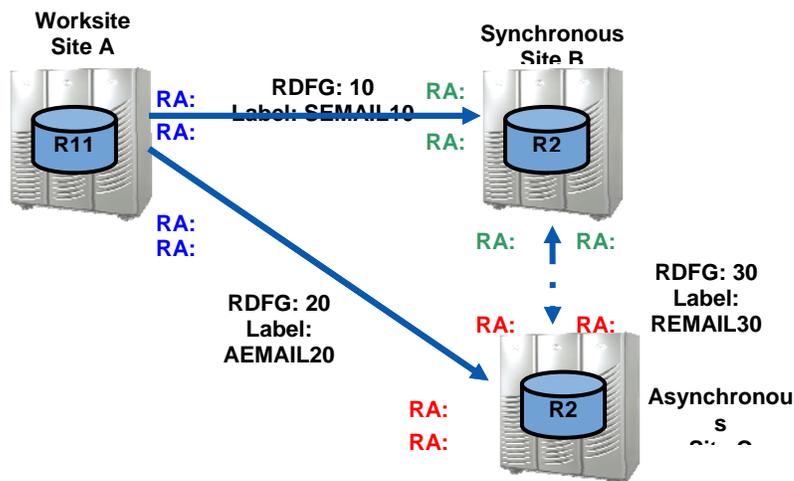
Data is always consistent since all data in SRDF/Star is transferred either synchronously or asynchronously. You can begin processing information at any of the sites should you have a disaster, and still maintain a disaster recovery site with consistent data at the other site. Data centers using SRDF/Synchronous must stay within two hundred miles of each other; the third data center using Asynchronous can be placed any where in the world depending on bandwidth and response time objectives.

In diagram 1, we see three arrays in a triangle. The worksite array (site A) is using concurrent SRDF to migrate data to the other two sites. Site A, the primary site, is sending data synchronously to site B (the synchronous site). At the same time, Site A is also sending data Asynchronously to Site C (the consistency site). This triangle is one of the ways that SRDF/Star got its name.

The worksite, primary array Site A, is where all the symcli commands are run to control the data being passed to the other two arrays. The synchronous site (Site B) must be located within 200 miles or 120 kilometers of Site A, all the data is being transmitted synchronously to this site meaning that both data centers Site A and Site B have exactly the same information, at the same time, at both sites. The Asynchronous site (Site C could be located anywhere in the world) and the data on this array is in a consistent mode. When you begin processing data at this data center site, the data and databases are restartable and will come online.

Continuing to look at SRDF/Star diagram 1, the dotted line between Site B and Site C is the recovery link or hollow leg. This is an empty RDF group that is defined in advance and used to track date changes between sites B and C. Should you decide to switch your worksite to Site B or Site C, this recovery group is used to determine which data is more current. This gives you the ability to decide which Site to switch to and what data you would like to keep at the new worksite. This can be the data from Site A, B or C, using the `symstar -c name -keep_data name` option. The `symstar -cg Name` query command will tell you which site has the latest and most current data.

SRDF/STAR Diagram 1



What is a Star

A star is a consistency group (CG) made up of concurrent Symmetrix® groups containing SRDF local devices and their remote mirror devices. These R1 devices replicate data between two sites with a recovery group on standby between sites B and C. A recovery group is an empty RDF group at the synchronous and asynchronous target sites that specifically matches each synchronous and asynchronous group at the workload site.

Star device groups use dynamic SRDF and can be configured online using the symconfigure command.

In our diagram, we built one Consistency group comprised of a synchronous group replicating data to Site B, an Asynchronous group replicating data to Site C, and a recovery group that is between Sites A and C that keeps track of the data that is on Sites B and C. Each group has a label, in our case the synchronous Site A to Site B the label name is SEMAIL10. The label between asynchronous Sites A and C is AEMAIL20. The recovery asynchronous group is labeled SEMAIL30.

Mirror Positions

You are using the four mirror positions available for your devices when building concurrent SRDF/Star. The device uses the first two mirror positions on your primary site, the third mirror position is used at the synchronous site (Site B), and the fourth mirror position is used at the asynchronous site (Site C). As a result, BCVs are not used in SRDF/Star configurations. Using EMC TimeFinder[®]/Clones is the preferred method of copying data internally; ensure that you have an EMC Clone license for this feature when building SRDF/Star. You cannot apply hot spares in your arrays to disks assigned to the SRDF/Star configuration, but this should not be an issue.

What is Needed to Build an SRDF/Star Environment?

You need three arrays to build an SRDF/Star environment. In most cases, this means three data centers. However, there are customers using two data centers with the synchronous arrays located in the same data center and the Asynchronous array located in a data center far away from the main data center. You will also need Solutions Enabler v6.3; use the latest version available. Make sure you have all your licenses installed. The licenses are usually located in the following file and location.

UNIX: /usr/emc/API/symapi/config/symapi_licenses.dat

Windows: C:\Program Files\EMC\symapi\config (CHECK)

Licences:

ConfigChange / Symmetrix
SRDF / Symmetrix
SRDFA / Symmetrix
TF/CG / Symmetrix
TF-Clone / Symmetrix
STAR / Symmetrix
TimeFinder / Symmetrix
DevMasking / Symmetrix
BASE / Symmetrix
SRDF/CG / Symmetrix

Set the following options in the /usr/emc/API/symapi/config/options file on UNIX and C:\Program Files\EMC\symapi\config\options on Windows

SYMAPI_ALLOW_RDF_SYMFORCE	= TRUE
SYMAPI_PARALLEL_RA_GROUPS	= ENABLE
SYMAPI_USE_GNS	= ENABLE
SYMAPI_USE_RDFD	= ENABLE

GNS Recommendations

GNS is recommended between the local control hosts, but at this time remote mirroring between sites is not supported in a concurrent SRDF environment. The symstar buildcg command was designed for this reason. This command builds the recovery links between the secondary sites. Run this command after running the symstar setup command and copying the definition file to the secondary sites B and C.

The definition file is created when you run the symstar setup command and is located in /var/EMC/symapi/def directory on UNIX in Windows this file is located in C:\Program Files\EMC\symapi\def. Once the definition file is created, copy it to workstations at Sites B and C and place it in the same directory as the Star options file on the local workstation. Otherwise, you will receive an error message when you switch to sites B or C and try to run commands at the new worksite.

Build the SRDF/Star Options File

You can create the SRDF/Star options file using notepad in Windows or VI in UNIX. This is where you define names for your workload, synchronous and asynchronous sites. We will discuss more on SRDF/Star naming conventions later in this article. The fourth line is the variable track for switching from adaptive code mode to Asynchronous

mode which is done automatically when there are 30000 tracks or approximately 30 seconds of data left to be transferred. Engineering suggests that you leave the timeout variable at 1800. The last line is a new feature developed in 5773 code and Solutions Enabler 6.5 called Cascaded and Cascaded/Star.

You can switch between SRDF/Star concurrent and SRDF/Cascading serial options. This feature is important if you need to work on your Asynchronous connection since it allows you switch to Cascading Star and back to Concurrent Star when the link is fixed.

Options File

```
SYMCLI_STAR_WORKLOAD_SITE_NAME=workload
SYMCLI_STAR_SYNC_TARGET_SITE_NAME=<sync_leg_name>
SYMCLI_STAR_ASYNC_TARGET_SITE_NAME=<async_leg_name>
SYMCLI_STAR_ADAPTIVE_COPY_TRACKS=30000
SYMCLI_STAR_ACTION_TIMEOUT=1800
SYMCLI_STAR_ALLOW_CASCADED_CONFIGURATION = YES
```

Best Practices for SRDF/Star

Never Control C when in SRDF/Star

Do not ctrl c out of a command when working in SRDF/Star; it is better to wait for the 30 minute timeout and then place a call to EMC support. Doing a control C in the process of setup, protect or connect can cause inconsistent data at your site.

Do not Script an Unplanned Switch or Failover

When first working with SRDF/Star, most customers want to script switching between sites if a disaster occurs. Using scripts for a planned switch is not an issue, but do you really want a script to automatically switch your worksite in an unplanned switch?

The first question that comes to mind in a real disaster or unplanned switch is who has the authority to tell the administrator to switch the primary site to one of the secondary sites? What data should you keep the synchronous or asynchronous data? Which data center should you switch to the synchronous or asynchronous data center? Should you isolate your sites instead of switching sites? There are a lot of questions you need to answer before making your switch in a true disaster. Hopefully, you have a disaster plan that addresses most of these questions and answers that you can use in a real disaster.

I was working on one SRDF/Star installation when the customer lost access to their data center but the equipment was still up and running. In this case, the customer continued running normally and never had to switch over to one of their other sites.

In another case, we had just finished building the customers first SRDF/Star (Exchange application) and switched to the new worksite. The next day they lost their data center and Exchange was the only application up and running.

It is best to have plans in place and in a real disaster make the decision at the right time (not scripted) to switch sites.

When Faced with a Real Disaster: Isolate

In a real disaster, if you lost your entire worksite and you needed to come up in either the synchronous or asynchronous site, I recommend isolating your two data centers. This is done in SRDF/Star using the `symstar -cg isolate -star` command. This protects your data so that you can come up on one of your sites to ensure that your data is not corrupt. You can then make a gold copy and bring SRDF/Star back online with your remaining data center.

This method is not necessary, but it does give you the ability to ensure that in a real disaster you have two independent data centers with two copies of data until you check your data out and make new gold copies.

Two Workstations at Each Site or Data Center Location (six servers in total)

I recommend that you have two control workstations at each data center to run the `symcli` commands and support the RDF daemon. You should have four servers (zoning in place for workstations to be switched if needed) at that location, and two at Site C if using two data centers with Sites A and B located in the same data center. These servers can be either Windows or UNIX; this is where you will run the `symcli` commands and RDF daemons. I recommend that no other applications run on the control servers that will interfere with the RDF daemon. The RDF daemons are responsible for checking the links and preventing Star from becoming disabled if a server loses connectivity or becomes unavailable which could cause SRDF control operations to stop and MSC to trip. Both servers are active and work like a cluster so that if one server loses connectivity, the other server and RDF daemon will continue to keep the links up and running.

Build an SRDF/Star Configuration for Testing Purposes

I recommend building a test Star with at least two devices. This practice provides valuable knowledge transfer and experience that cannot be learned without building and testing in an SRDF/Star environment.

Running the commands in a test environment shows how various parameters can affect the outcome of a command that might differ from your expectation. You soon realize when working with SRDF/Star that every command has to have the `-star` parameter in the command or you receive an error message. Every command has to have the `symstar -cg` name in the command or it will fail, so do not use large complicated names when defining your worksite. You will learn how to use commands such as `halt`, `reset` and switching your workstation sites and even more importantly, how to switch the workstation back to the original data center. You will learn the difference between `isolate`, `disable` and `halt` parameters and what happens to your data in the background when running these commands. You will also learn that naming conventions are very important; you should take great care in deciding your naming conventions in advance.

In one case, I was called to a customer site where they were running SRDF/Star in production but the system administrator and manager were both concerned because they did not feel that they had the confidence to work on the product. We built a test star and once the customer started to work with the test setup, they not only felt confident but decided to build Stars for all their applications.

Have Enough Gatekeepers

You have no communication path to your array and no way of communicating your commands without a gatekeeper. With many gatekeepers, you have many paths to ensure that your commands are carried out. You should have at least one gatekeeper for every daemon that is running on the server and a gatekeeper for each RDF group. You will not be able to communicate if they are all busy so have more than needed at any time, and more available if necessary.

A friend who is a Symmetrix expert put it this way, “to ensure you have enough gatekeepers have one for every command you are issuing to the array at every server at the same time. A good rule of thumb is the more the merrier.”

Check Daemons

Make sure that necessary STAR daemons are running. Issue a command from all control servers to list the daemons that are running.

stordaeomon list

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

[*] storapid	EMC Solutions Enabler Base Daemon
[*] storgnsd	EMC Solutions Enabler GNS Daemon
[*] storrdfd	EMC Solutions Enabler RDF Daemon
storevntd	EMC Solutions Enabler Event Daemon
[*] storwatchd	EMC Solutions Enabler Watchdog Daemon

Star Naming Convention

When building one Star, you learn very quickly that naming conventions are very important and need to be thought out in advance. Working with one can be easy but once you realize how easy the product is to use and how safe and effective it is, you will start to put all your applications under the SRDF/Star software. As a result, you will end up with quite a few Stars in your SRDF/Star configuration.

When building SRDF/Star you need to decide on seven names (shown below) for each Star that you are using. I suggest that your Star or CG have a name that is useful and represents what application is under its control. I have listed an example below:

Star or CG name = EXFO SR: EX = Email FO= Front Office Application SR = Server name

Site A name = EMWS (EMAILWorksite)

Site B name = EMS10 (Email Sync group 10)

Site C name = EMA20 (Email Async group 20)

Sync label name = Sync10 (RDF sync group 10)

Async label name = Async20 (RGF async group 20)

Recovery group name = Rec30 (RDF recovery group 30)

Remember that you need to type the name every time you run an SRDF/Star command. I do not recommend making a variable name. If you intend to have more than 10 Stars use group 10 for the sync side, 30 for the async side and 50 for the recovery side otherwise a good rule is to use 10 for the sync side, 20 for the async side and 30 for the recovery side and to add one to the number for every SRDF/Star that you build. This way all SRDF/Star applications are in groups 10 -19 for synchronous, 20 – 29 for asynchronous, and 30 – 39 for the recovery leg.

There are customers in the field running Star configurations using the following naming conventions. They just add to the number incrementally when building new Stars. The naming convention is up to you and the customer. Below is an example:

CG = Star1

Site A = A1

Site B = B1

Site C = C1

Sync = sitesync1

Async = asyncsync1

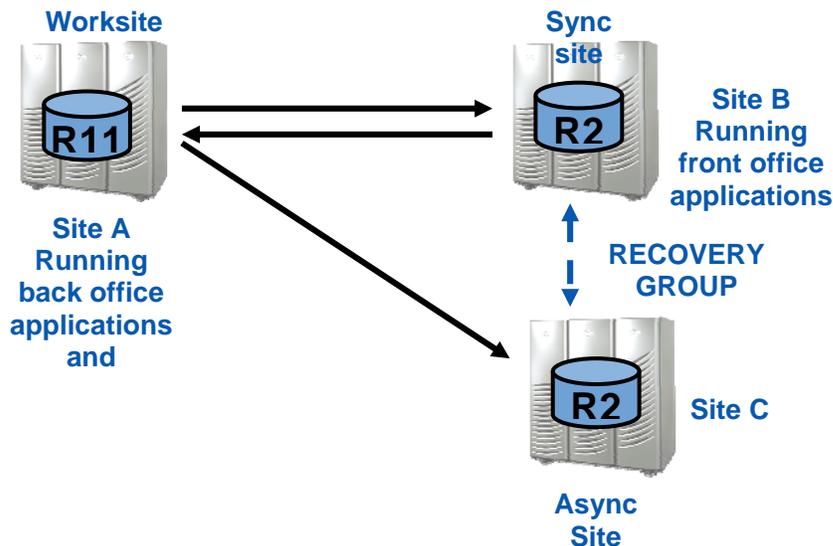
Recovery = recsite1

When building Stars in your SRDF/Star configuration, be aware that managing them can become confusing and cause serious issues if you do not know which Star you are working on when issuing commands. Try to keep the Stars to a manageable number, my suggestion is under 12 (between 6 and 8 is optimal). If you are creating more than 22, I recommend that you get an RPQ approved to ensure you will have no issues. You can run up to 64 Stars depending on the limits of groups allowed on the arrays you build your SRDF/Stars on.

Load Balancing Applications Across Data Centers

Although SRDF/Star was designed as disaster recovery software, once we started to install and run SRDF/Star at customer sites we found out that the product offers more than disaster recovery. One offering is load balancing applications between Sites A and B. SRDF/Star can help customers take advantage of their hardware at both sites. You can run half your applications at Site A and the other half of your applications at Site B giving you better performance and load balancing at both sites for all your applications while maintaining all of your data in synchronous mode.

SRDF/Star Load Balancing Diagram 2



The diagram illustrates that our customer is now running applications at sites A and B (sync sites). In this case, all front end applications are running in Site A and all back end office applications are running at the B site. This provides excellent use of hardware and improved application performance.

Eliminating Downtime while Working on Servers

With load balancing, you are provided enhanced application uptime and unlimited downtime for system maintenance at each of your synchronous sites with no negative performance impact on your applications. When you need to perform upgrades or maintenance on your servers or arrays, simply halt your Star to stop the application workload from writing to the R1 devices and completely synchronize both remote target sites. You then perform a symstar switch command to move your applications to the remote site while you work on the servers or arrays at the primary site. Switch back to your primary worksite when the work is completed.

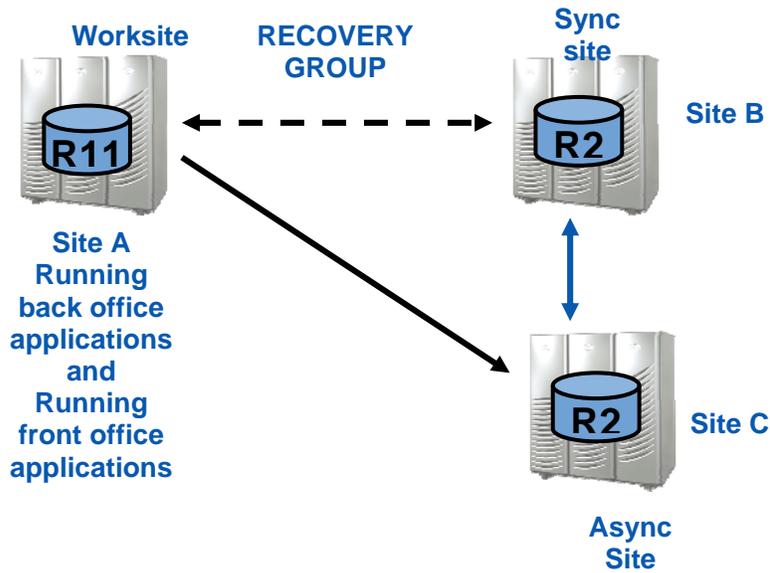
Applications can be switched to either Site A or Site B depending where you want them to live at that time. In the diagram below we have switched all applications running in Site B to Site A with the following commands: (see diagram 3)

```
symstar -cg <STARcg_name> halt -nop
```

```
symstar -cg <STARcg_name> switch -site <sync_site> -nop
```

During the maintenance window at Site B, all applications are running at Site A and are being Asynchronously transferred to Site C providing disaster recovery protection and maintaining consistency between Site A and Site C. Once the system maintenance is completed, the applications are switched back to Site B and normal SRDF/Star processing is continued within the three site disaster recovery.

SRDF/Star Eliminating Downtime Diagram 3

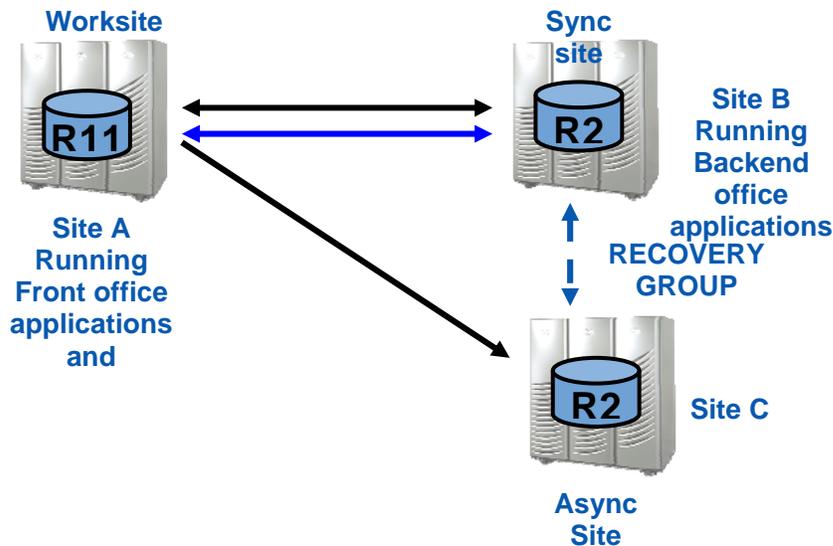


Migrating Data with no Downtime while Staying Consistent

It is easier to migrate data when running SRDF/Star. You can perform a migration with no downtime while staying consistent. In the example that follows, the customer has decided to purchase a new Symmetrix DMX4 for their data center at Site B. We will be moving one Star over to the new DMX4. After that migration is complete, we will move the rest of their Stars and application data from the DMX3 to the new DMX4.

Diagram 4, on the following page, shows the current customer site. The black line between sites A and B illustrates our Star applications, while the Blue line shows all the applications that are running bi-laterally between sites A and B.

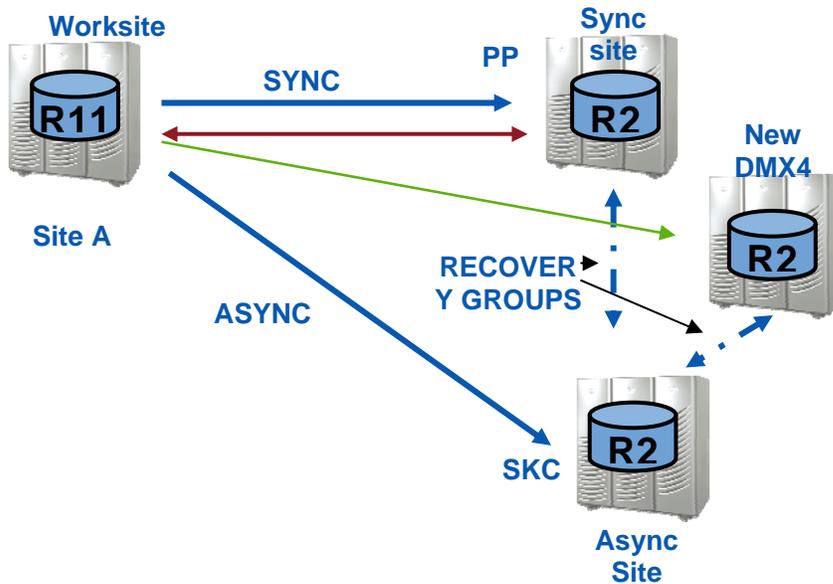
SRDF/Star Multiple Stars and Load Balancing Applications Diagram 4



The customer receives the new DMX4 at their data center. We setup the new DMX4 and begin to migrate the first Star application to it. The first step is to disable our Star application (in this case referred to as Star WEM011 running our email application). Data will continue to asynchronously transfer to Site C to keep our disaster recovery plan in effect with a restartable application at Site C. All of our other applications continue to transfer data to sites B and C with no impact on production.

The data migration full sync is started between Site A and Site B as shown by the green line to the new DMX4. We will use SRDF ACP_disk mode copy until most of the data is copied, at that time we will do a mode switch to SRDF/sync. Once the synchronization is completed and all the data is copied, we enable SRDF/Star for WEM011. Our first Star is migrated and is now enabled and running between sites A, B and C. See diagram 5 on the following page.

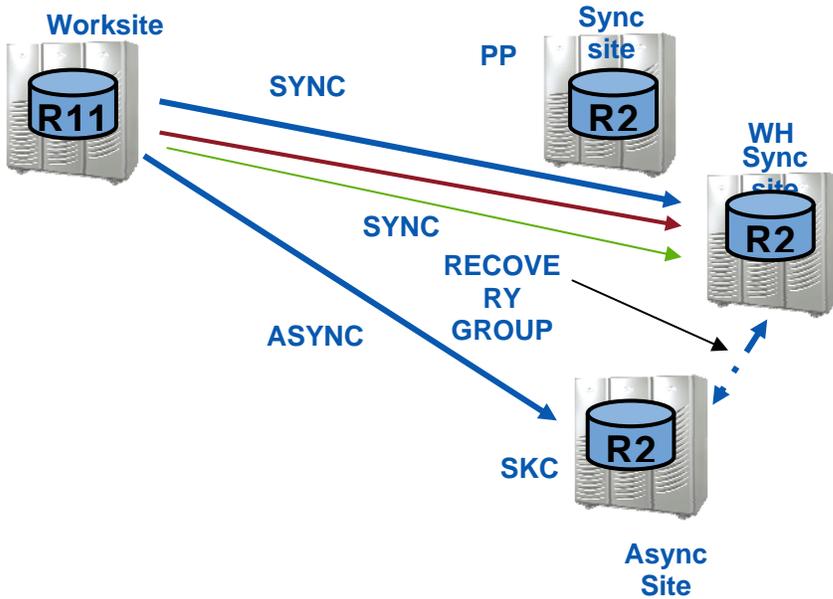
SRDF/Star Migrating Data Diagram 5



The customer will now migrate all of their Stars and application data from DMX3 at site A to the new DMX4. All of the Stars at the worksite are disabled while SRDF continues to Asynchronously transfer our data to site C for our consistent restartable copy. Once our Stars are disabled, we stop replication from our worksite to our sync site. We start a full sync to our new DMX4 using adaptive copy mode `acp_disk` while still using asynchronous replication to our Async site to ensure that we have a good, restartable copy of data.

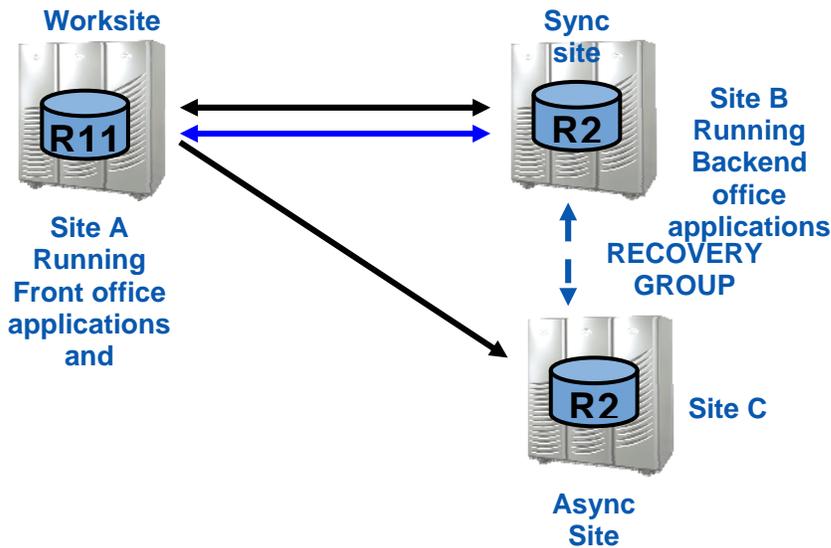
When most of the data is copied over to our synchronous site, we switch to synchronous copy mode SRDF/Sync and enable our SRDF/Stars that will provide us with three site disaster recovery while maintaining consistent data at all three sites.

SRDF/Star Migrating Stars and Applications Data Diagram 6



The customer enabled all of their Stars and we have now moved all of the customer's SRDF/Stars and application data from their DMX3 to their new DMX4 with no downtime. They continue to have their consistent copy at site C for their disaster recovery plan.

SRDF/Star Multiple Stars and Load Balancing Applications Diagram 7



Whose Data should I Keep at Site A, B or C?

You need to Switch from Site A to Site B but you want the latest data when you switch over to the new switch, how can you tell which site has the latest data? Run the symstar -cg <Star name> query -detail command to check which site has the latest and greatest data.

symstar -cg Starcg query -detail

```
Site Name                : Worksite

Workload Site            : WORKSITE
1st Target Site          : Site B
2nd Target Site          : Site C

Workload Data Image Consistent : Yes
System State:
{
  1st_Target_Site        : Halted
  2nd_Target_Site        : Halted
}

Last Action Performed    : N/A
Last Action Status       : N/A
Last Action Timestamp    : N/A

STAR Information:
{
  STAR Consistency Capable      : Yes
  STAR Consistency Mode         : NONE
  Synchronous Target Site      : Site B
  Asynchronous Target Site     : Site C
  Differential Resync Available : N/A
  R2 Recoverable                : N/A
  Asynchronous Target Site Data most Current : No **
}

1st Target Site Information:
{
  Site Name                : Site B
  RDF Consistency Capability : SYNC
  RDF Consistency Mode     : N/A
  Site Data Image Consistent : No
```

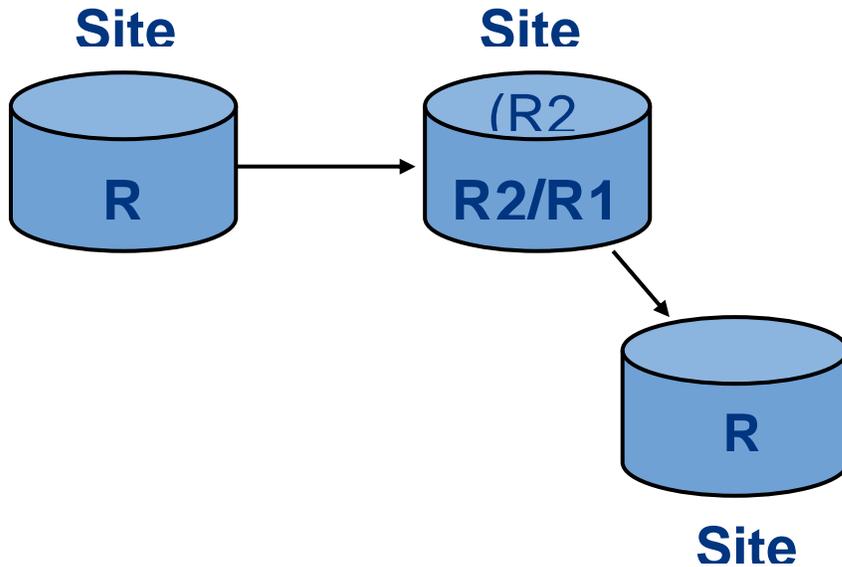
Once you stop processing or moving data between sites A, B and C for any reason, wait two minutes before running the symstar query command to check which Site has the latest information. You wait two minutes in case the link or processing stopped on the synchronous side (site B) while the asynchronous side (site C) continued processing data. It takes approximately two minutes for a full asynchronous cycle to get the latest data and to change the No to a Yes in the query command.

** In the red line above from the symstar query command you can tell which data is more up to date. If there is a No, the synchronous (site B) has the latest information; if there is a Yes, the asynchronous side has the latest information.

New Cascading SRDF/Star Feature

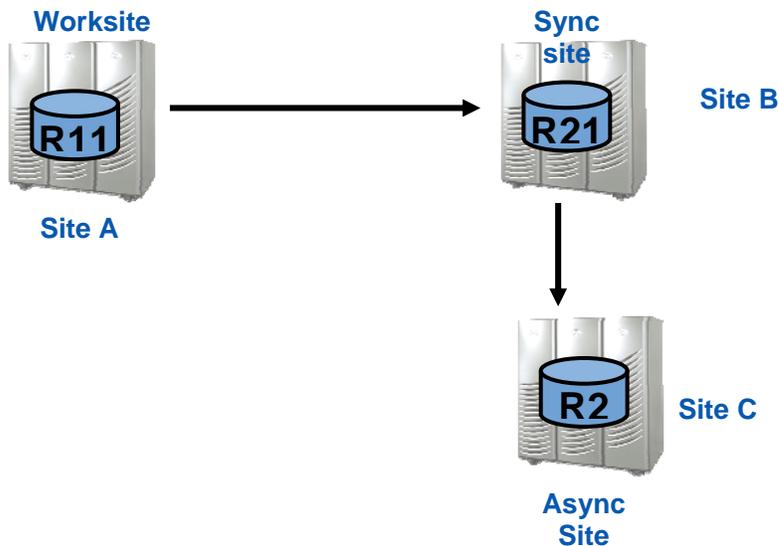
EMC's new R21 device (Solution Enabler 6.5 and Engenuity™ code 5773) offers new features for SRDF and SRDF/Star that improve our ability to transmit information between data centers for disaster recovery purposes, these are Cascaded and Cascading/Star. An R21 device as shown in diagram 8 on the following page sits at site B and acts as an R2 to our Site A's R1, and at the same time acts as an R1 for Site C's R2. This device acts as two devices simultaneously; in reality it is only one device that improves performance and reduces storage costs.

SRDF R21 Device - Diagram 8



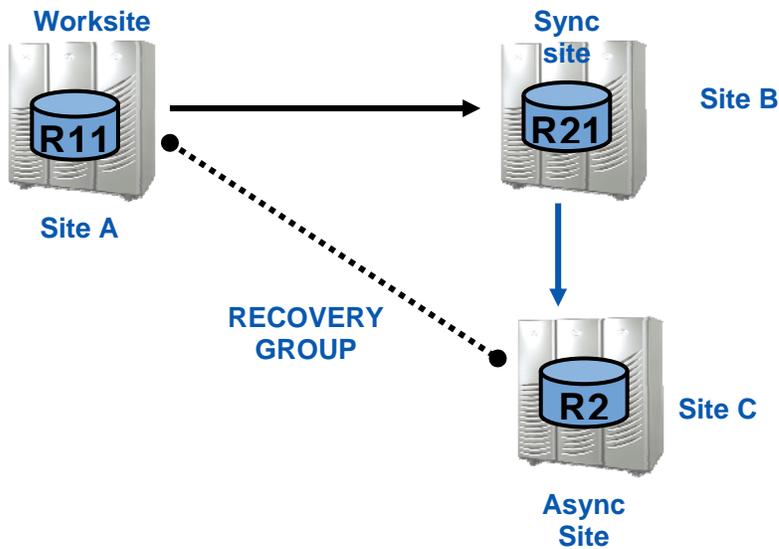
This R21 device allows for data pass thru and gives us the ability to transmit data across three data centers using synchronous, asynchronous and acp_disk mode while maintaining copies of data at all three sites. Cascading, unlike concurrent, sends the data to Site B which forwards the data to Site C using the new R21 device. In Concurrent SRDF, Site A sends the data to Site B using synchronous SRDF while at the same time Site A is transferring data to Site C using asynchronous SRDF. In the diagram below, R1 information at site A is transmitted to site B R21 which sends the information to Site C's R2. Cascading SRDF also gives the customer the ability to move dynamic SRDF devices between SRDF groups without requiring a full copy resynchronization.

Cascaded SRDF



You develop the new Cascading SRDF/Star feature by adding a recovery group between Site A and Site C (see diagram below). This feature allows you to change your SRDF/Star configuration to cascading SRDF/Star and back to SRDF/Star. Use the following command to change from cascaded RDF to Concurrent RDF mode: **Symstar reconfigure -site SiteC -path SiteA:SiteC -reset**

Some of the benefits of being able to switch between Cascading SRDF and Cascading SRDF/Star include the ability to reconfigure your configuration in response to any planned or unplanned events that may require relocation of your workload site (Site A) around extended network link outages between sites A and Site C.



Location of Information in Powerlink

How do you find documentation to get started with SRDF/Star? Powerlink is a good place to begin. Search the links below to find all the documentation that you need to start learning and building your own SRDF/Star. These links will take you to the TS Kit for SRDF/Star. This kit is the documentation that is built by the TSG organization to help the field learning about products, how they work and how to build them. There is a product guide, a level of efforts, a test and acceptance plan plus everything you need to work on SRDF/Star. Whenever you are looking for documentation I suggest you always check the TS Kits as each product that we sell has one.

The new TS Kit Remote Replication Guides

http://powerlink.emc.com/km/appmanager/km/secureDesktop?_nfpb=true&_pageLabel=query2&internalId=0b014066800bf6ea&_irrt=true

The new TS Kit Practitioner's Guide for the Open Systems Replication is at

http://powerlink.emc.com/km/live1/en_US/Offering_Technical/Technical_Documentation/PG_SRIS_PSCUSPRDSVC.pdf?mtcs=ZXZlbnRUeXBIPUttQ2xpY2tDb250ZW50RXZlbnQsZG9jdW1lbnRJZD0wOTAxNDA2NjgwMzI2MWRILGRvY3VtZW50VHlwZT1wZGYsbmF2ZU5vZGU9MGIwMTQwNjY4MDBiZjZiYQ

You can also look at the TS Kit's Test & Acceptance Plan for Cascaded SRDF:

http://powerlink.emc.com/km/live1/en_US/Offering_Technical/Technical_Documentation/TA_SRIS_PSCUSPRDSVC_cascaded.doc?mtcs=ZXZlbnRUeXBIPUttQ2xpY2tDb250ZW50RXZlbnQsZG9jdW1lbnRJZD0wOTAxNDA2NjgwMzI2MTZkLGRvY3VtZW50VHlwZT1tc3c4LG5hdmVOb2RIPTBiMDE0MDY2ODAwYmY2ZWE

The new Solutions Enabler SRDF CLI Product Guide:

http://powerlink.emc.com/km/live1/en_US/Offering_Technical/Technical_Documentation/300-000-877_a10_elcncnt_o.pdf?mtcs=ZXZlbnRUeXBIPUttQ2xpY2tDb250ZW50RXZlbnQsZG9jdW1lbnRJZD0wOTAxNDA2NjgwMzEyZDg0LG5hdmVOb2RIPVNVZndhcmVEb3dubG9hZHMtMg,

Also, look up the Education Services IMPACT module titled

["Engenuity 5773 Hardware Focus SRDF IMPACT"](#)

New Commands (for Cascading, Star and Cascading Star)

-hop2

Indicates the device is two hops away

With an R1 -> R21 -> R2 configuration, issuing a query -hop2 from an RDF1 CG indicates that the query should show the relationship of the R21->R2 device pairs. Thus the query will display the R21 device from the R1 mirror point of view (and vice versa for RDF2 CG). To see both hops of the RDF1 or RDF2 CG that contains devices in a cascading RDF relationship, use the symrdf -cg query command with the -hop2 and the -detail options.

-Star

Once Star is enabled you need the -star in your commands

-set_mode

rdf_mode acp_disk -star -nop

-Halt reset

If you do a symstar halt and do not have workstations or access to other sites

Symstar reconfigure -site dallas -path newyork:dallas -reset
Changes from cascaded RDF to concurrent RDF mode

Symstar -cg name protect -site name -keep_data name
Protect turns on RDF consistency, in cascading you must protect sync leg first

Symstar -cg name connect -site name -keep_data name
Connect brings the devices to (RW) on the RDF link, in cascading you must connect sync site first

-detail Gives more information
Symstar -cg name query -detail -star

-wkload
Allows the user to specify the current workload site

-V verbose
Provides more detailed listing

Troubleshooting Tips

Verbose -v

When an error is encountered in a symstar command, repeat the command with the -verbose option

The following two features, logging and debug mode, were written by EMC Engineering*

Turn your SYMAPI verbose logging on:

```
root@ /: export SYMCLI_VERBOSE=1
```

Check to see if your verbose log directives have taken place:

```
root@:/: symcli -def
```

Current settings of the SYMCLI environmental variables:

```
SYMCLI_VERBOSE      : 1
```

Running SRDF/Star in Debug Mode

1. Run `symstar -cg cgr_s1 query -detail` and capture the output. Ensure STAR is operating normally.
2. Set the environment variable `SYMAPI_DEBUG=-1`, and the variable `SYMAPI_DEBUG_FILENAME=/var/symapi/log/symapi-test1109.log`.
`symcfg list -def` (verify it got done)
3. Edit the file `/var/symapi/config/daemon-options`, and add the following lines:

```
storgnsd:debug=debug  
storrdfd:debug=debug  
storapid:debug=debug
```

4. In `/var/symapi/log`, rename the current day's log files:

```
storapid-20061109.log  
storrdfd-20061109.log  
storgnsd-20061109.log  
symapi-20061109.log  
symstar-20061109.log
```

In the daemon_options set

```
stor?????:logfile_type = dated
```

for each of the daemons. When the daemons are restarted they should begin writing to the dated file format.

5. Restart the following daemons:

```
storapid  
storrdfd  
storgnsd
```

6. Run a symconfigure –sid 370 query and make sure that you’re seeing output in the /var/symapi/log/symapi-test1109.log file.

7. Run /apps/drm/bin/final_shutdown_STAR.pl --verbose. If we don’t get the error, then we don’t have the problem. If we do get the error, collect all the log files that we just generated (listed above) and send to Engineering for further analysis.

8. Undefine the variables set in step 2, remove the lines set in the daemon_options file in step 3, and restart the same daemons from step 4.

Log Files

Read and interpret the error messages that you find in your log files to help you determine what could have caused an issue. Use commands to gather more information such as the symstar query command and analyze the output. You should be able to draw a diagram of your SRDF/Star configuration from this output. Check and read your logs!

symstar-yymmdd.log

Key log with detail about every symstar operation

symapi-yymmdd.log

General SYMAPI log including information on all active API commands (including all RDF actions)

storrdfd.log0, 1

RDF Daemon log

Swaps between 0 and 1

Check daemons

stordaemon show storrdfd

Daemon State : Running
Daemon Start Time : Tue Apr 25 20:31:11 2006
Version : V6.2-724 (0.0)
Auto-Restart by Watchdog : Enabled

Total Number of Connections : 3
Number of Active Connections : 3
Total Number of Requests : 2008

Config options :
Syscall Path : Base Daemon
GK Pool Size : 20
Dedicated GKs per Symm : 2
Base Daemon Connections : 20
Main Thread Interval : 30 sec
Group Monitor Interval : 15 sec
DB to Disk Interval : 120 sec
MSC Threads per Symm : 10
MSC READY Retries : 90
MSC Host Cleanup : Enabled
ECA Monitor Interval : 1 sec
ECA Trip Threads per Symm : 1

Allowable States for symstar Operations

The SRDF/Star system components need to be in an allowable state to perform a symstar command. Otherwise, a message is returned stating that SRDF/Star is not in a state that permits the particular operation that you are attempting to perform.

For example, the three component states are Protected when SRDF/Star is running normally:

System State:

```
{  
1st_Target_Site : Protected  
2nd_Target_Site : Protected  
STAR : Protected  
}
```

When SRDF/Star is enabled, the STAR line displays either Protected or Tripped.

The following sections list the allowable states for each SRDF/Star control operation. The System State is displayed with the symstar query command.

Protect The symstar protect command can be executed only on a target site that is in the Connected state. For example, if the current state of the sync target site is PathFail, you cannot protect the sync site without first performing a connect action.

Unprotect The symstar unprotect command can be executed when the specified target site is in the Protected state. If you are running the workload at the async target site (Site C), only one path can be protected (in async mode) at any one time. You can, however, toggle the protection between Site A and Site B. That is, you can unprotect one and then protect the other.

Enable Once both target sites are in the protected states you can execute the symstar enable command. When the workload site has been switched to the asynchronous target site (site C) you cannot enable as the action is blocked when running at the asynchronous site.

Disable The symstar disable command can be executed when the System State indicates that SRDF/Star is protected.

Isolate This command can be used to isolate one of the sites from the other two sites. Isolate is useful when you have a real disaster and you want to come up on one site while isolating the other site to ensure that there is no data corruption at any of the sites. The symstar isolate command can be executed only when the remote target site is in the Protected state.

Reset You can only execute the symstar reset command when the remote target site is in the PathFail state. Use the reset action after correcting the physical cause of a transient fault. Unless the -force option is specified, the reset action is rejected if the RDF path to the site is in a Partitioned state.

Summary of symstar Control Operations* (taken from product guide)

You perform SRDF/Star setup and recovery operations using EMC's host-based symstar commands. Table 1 summarizes the SRDF/Star operations.

Table 1 SRDF/Star Operations Summary

Control Operation	symstar Action Argument	Results	Host Site: Workload(W) or Target (T)
Begin SRDF synchronization	connect	Starts the SRDF data flow in Adaptive Copy Disk mode.	W
Suspend SRDF synchronization	disconnect	Suspends the SRDF data flow and transitions the path to Adaptive Copy Disk mode.	W
Set up RDF consistency protection to target site	protect	Synchronizes devices linked to the remote site and sets up RDF consistency protection to the specified target site.	W
Discontinue RDF consistency protection to a target site	unprotect	Suspends RDF consistency protection to the specified target site.	W
Enable full SRDF/Star protection	enable	Enables complete SRDF/Star protection across the three sites.	W
Disable full SRDF/Star protection	disable	Disables SRDF/Star protection across the three sites.	W
Halt the SRDF/Star system	halt	Prepares the system for a planned switch of the workload to a target site. This action write-disables the R1 devices, drains all invalid tracks and MSC cycles so that SiteA=SiteB=SiteC, suspends RDF links, disables all consistency protection, and sets Adaptive Copy Disk mode.	W
Isolate a target site from the SRDF/Star configuration	isolate	Isolates one of the target sites from the SRDF/Star configuration and makes its R2 devices RW-enabled to the target site hosts.	W
Reset after a transient fault	reset	Cleans up internal metadata and Symmetrix cache at the remote site after a transient fault (for example, loss of connectivity to the sync or async target site).	W
Clean up after a disaster fault at workload site	cleanup	Cleans up internal metadata and Symmetrix cache at the remote site after failure at the workload site.	T
Display SRDF/Star status	query	Displays the status of a given SRDF/Star configuration.	W/T
Display the SRDF/Star internal definition	show	Displays the contents of the internal definition file for a given SRDF/Star configuration.	W/T
Switch workload operations to a target site	switch	Transitions workload operations to a target site after a workload site failure or as part of a planned event.	T

SRDF/Star and Cascaded SRDF/Star Aggregate Operations* (from product guide)

Operation	Functions Involved
Setup	Build Star internal definition – Reads and validates the composite group definition and Options file, and generates the Star internal definition file.
Buildcgr	Build the host Composite Group from internal definition – Reads the Star internal definition file and creates the RDF1 or RDF2 composite groups.
Query	Display the current State of the Star configuration – Displays the status of the Star configuration and the last action performed.
Connect	Make an SRDF connection and start data flow – Performs the necessary commands to reconfigure the RDF devices and then starts the data flow in Adaptive Copy Disk mode. Transitions the targeted remote site into a “connected” state.
Disconnect	Turn off Star protection – Performs the necessary commands to suspend the SRDF data flow and transitions the targeted remote site into a ‘Disconnected’ state.
Protect	Turn on SRDF Consistency protection – Waits for the devices to get close to being synchronized and transitions the devices into “final” RDF mode (that is, synchronous or asynchronous). Then it activates RDF consistency protection. When consistency is reached, the remote site is transitioned to a “protected” state.
Protect	Turn on SRDF Consistency protection – Waits for the devices to get close to being synchronized and transitions the devices into “final” RDF mode (that is, synchronous or asynchronous). Then it activates RDF consistency protection. When consistency is reached, the remote site is transitioned to a “protected” state.
Unprotect	Turn off SRDF consistency protection – Deactivates RDF consistency protection to the specified target site and changes the RDF mode to adaptive copy disk. Transitions the targeted remote site back into a “connected” state.
Enable	Turn on Star protection – Performs necessary commands to enable complete SRDF/Star protection across the three sites. Creates/activates SDDF bitmaps at remote sites, enables Star consistency protection, and transitions the Star configuration to “enabled” state.
Disable	Turn off Star protection – Performs necessary commands to disable Star protection. Deactivates/deletes the SDDF bitmaps, disables the Star consistency protection, and transitions the system out of the “Star protected” state.

Operation	Functions Involved
Show	Display Star internal definition – Displays the contents of the internal definition file for a given SRDF/Star configuration. Depending on the options selected by the user, the display includes information on Symmetrix resources and, optionally, all of the devices in the configuration.
Isolate	Stop SRDF protection, make devices RW at remote site – Isolates a remote target site from the Star environment by disabling the RDF consistency to the remote target site and perform a RDF “split” operation to make the remote target site’s R2 devices RW to the host. Transitions the remote target site into an “Isolated” state.
Reset	Perform reset after a transient failure condition – Provides the necessary cleanup to recover from a transient fault to either Synchronous or Asynchronous target sites. Transitions the remote target site from “PathFail” to a “disconnected” state.
Cleanup	Perform cleanup after a failure condition – Performs the MSC cleanup at the asynchronous remote target and allows for “gold copy” to be captured before starting any resynchronization as part of the switch command processing. Transitions the remote target from a “PathFail;CleanReq” to a “PathFail” state.
Switch	Prepare SRDF devices to run the workload at remote site - Switches the workload to a remote site (either synchronous or asynchronous remote target site). The switch is allowed from “PathFail; PathFail; Tripped” and “Halted;Halted” states. Moves the workload site and leaves remote sites in a “connected” or “disconnected” state. Provides both failover and failback.
Switch	Prepare SRDF devices to run the workload at remote site - Switches the workload to a remote site (either synchronous or asynchronous remote target site). The switch is allowed from “PathFail; PathFail; Tripped” and “Halted;Halted” states. Moves the workload site and leaves remote sites in a “connected” or “disconnected” state. Provides both failover and failback.
Halt	Halt the SRDF/Star – Stops the application workload from writing to the R1 devices and completely synchronizes both remote target sites. This is allowed when both remote sites are in a “connected” or “disconnected” state. Transitions the system into a state where both remote sites are “halted.” Allows for a planned switch or failover to a remote site with no data resynchronization needed between the sites. The “reset” option allows workload to be restarted at the “original” workload site. You cannot perform Halt if a link or site is not in Star configuration.

*Chart has been copied from the SRDF/Star product guide

Biography