



# **Veritas Storage Foundation For Oracle RAC with EMC SRDF**

**A Disaster Recovery Solution**

*- by -*

**Naveen Williams  
Symantec Principal SQA Engineer**

**September 11, 2007**



## TABLE OF CONTENTS

Overview.....	5
Disaster Recovery .....	7
Storage Foundation for Oracle RAC (SFRAC) .....	7
Key SFRAC Benefits.....	7
A Typical SFRAC Environment with GCO / SRDF.....	9
Key Components .....	9
Environment Setup Steps.....	9
For Each Cluster’s Non-Replicated Shared Storage.....	10
For Local Storage for Oracle and CRS Binaries on each Node.....	10
Software Installation and Configuration Steps.....	11
Failure and Recovery Scenarios.....	16
Intra-Cluster Failover and Failback .....	16
Inter-Cluster Failover and Failback.....	17
Symmetrix Remote Data Facility (SRDF).....	21
SRDF Links.....	21
Unidirectional / Bidirectional Data Transfer.....	21
<i>SRDF Topologies</i> .....	23
<i>SRDF Device Types and Groups</i> .....	24
SRDF Setup and Management .....	25
Conclusion .....	30



## Overview

This whitepaper describes Symantec's Veritas *Storage Foundation for Oracle RAC* clustering and replication technologies. It examines EMC's SRDF data replication to disaster recovery sites and provides technical details on setting up a Disaster Recovery solution for Oracle Real Application Clusters (RAC). This technical discussion only considers synchronous mode SRDF with unidirectional configurations.



## Disaster Recovery

Effective disaster recovery reestablishes mission-critical applications and associated data to normal operational states, usually at remote locations. Enterprise databases such as Oracle RAC are often essential components of mission-critical applications. To achieve disaster recovery typically requires replicating primary processing site data to remote sites using hardware-based or host-based replication methods.

- In *host-based replication*, hosts write to local storage and transmit writes to the remote site using a network. *Veritas Volume Replicator (VVR)* is an example of host-based replication.
- In *hardware-based replication*, storage arrays typically write to a local disk and send data to a remote array. Some products providing hardware-based replication are EMC's *SRDF*, Hitachi's *True Copy*, IBM's *Metro Mirror*, Network Appliance's *Snap Mirror*, and EMC's *Mirror View*.

## Storage Foundation for Oracle RAC (SFRAC)

Symantec's *Storage Foundation for Oracle RAC (SFRAC)* is a storage management infrastructure that provides high availability and disaster recovery technologies for implementing robust, manageable, and scalable Oracle Real Application Clusters (RAC).

SFRAC resides at the top of the I/O management stack and consists of:

- *Veritas Cluster Server (VCS)*
- *Cluster File System (CFS)*
- *Cluster Volume Manager (CVM)*
- Other key features as listed below.

### Key SFRAC Benefits

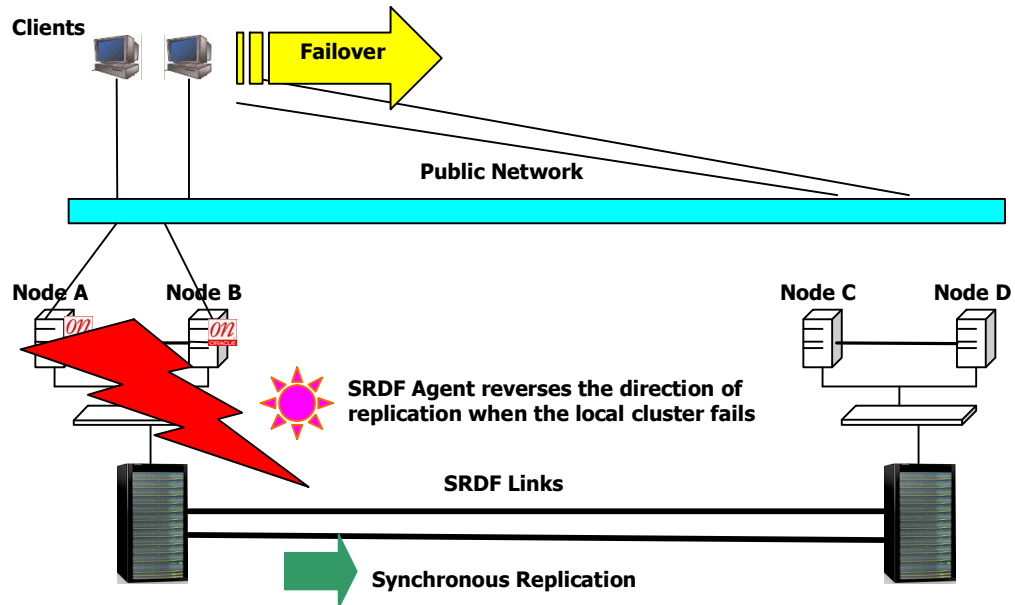
SFRAC provides:

- Shared storage management with a highly simplified Cluster File System (CFS)
- Higher performance and reduced complexity through Cluster Volume Manager (CVM)
- Dynamic Multi-pathing (DMP) that helps increase disk availability and I/O load balancing across all available paths
- Significantly improved Oracle disk I/O performance with Veritas Oracle Disk Manager (ODM)
- VCS Inter-process communication (IPC) for RAC across the private network links in a cluster. Load Balancing and link failovers are handled efficiently using low latency transport (LLT)
- I/O fencing to prevent data corruption on shared storage
- PrivNIC Agent provides high availability of private interconnects for Oracle Cluster Ready Services (CRS)

- Oracle Cluster Synchronization Services Daemon (CSSD) Application Agent manages startup and shutdown of CRS
- Global Cluster Option (GCO) for seamless failover to a remote cluster site during disaster recovery
- VCS Agents to monitor and manage database instances, virtual IPs, listeners, diskgroups, volumes and file system mounts
- Replication Agents like SRDF to monitor, manage, failover to the remote site and change the direction of the replication



## A Typical SFRAC Environment with GCO / SRDF



### Key Components

#### Replication Agent (EMC SRDF)

- Monitors and manages state of the Symmetrix SRDF device groups.
- The replication direction can change to make the remote cluster s the source and the local cluster the target.

#### Wide Area Connector (WAC)

- Registers with the local VCS engine and facilitates inter-cluster communication
- Manages WAN heartbeat (below) and notifies VCS in case of a heartbeat failure

#### Heartbeats

- WAC started processes that monitor remote and local clusters and storage

#### Global Group

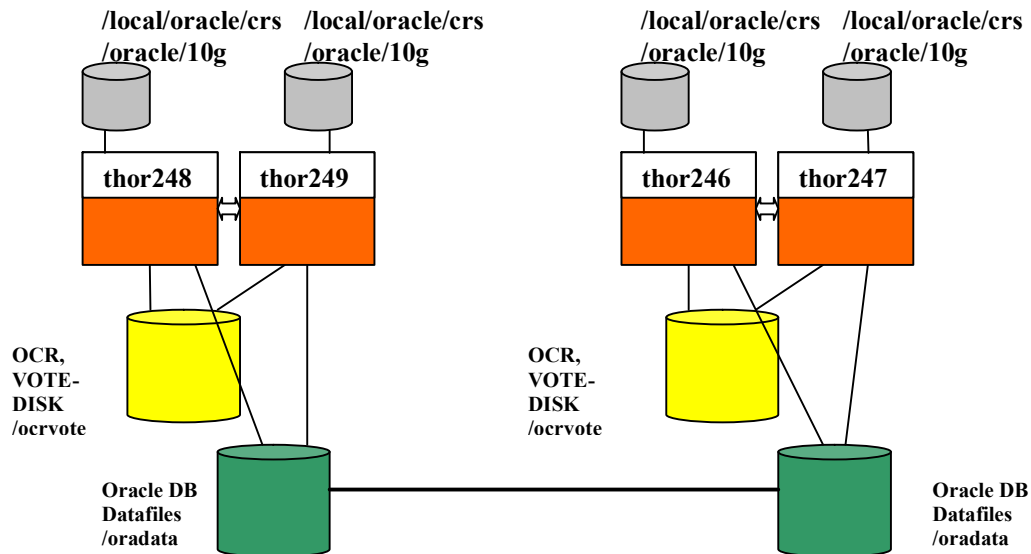
- A service group that enables wide-area failover. A group can be set to failover automatically when the local cluster or storage breaks.

### Environment Setup Steps

#### For Hardware and Replicated Shared Storage

1. Identify hardware and storage requirements.
2. Install the Symmetrix array CLI package.
3. Based on the RAC database size, identify the number of EMC disks/devices to replicate to the remote EMC storage array. The Oracle database files reside on these disks.
4. Create Device groups and RDF pairs.

For further information refer to this paper's SRDF section.



### **Sample Setup**

**Cluster1 nodes = thor248, thor249**

**Cluster2 nodes = thor246, thor247**

**CRS\_HOME = /local/oracle/crs (Installed locally on each node)**

**ORACLE\_HOME = /oracle/10g (Installed locally on each node)**

**OCR and VOTE-DISK = /ocrvote (Cluster File System using non-replicated disks)**

**Oracle DB Data files = /oradata (Cluster File System using replicated disks)**

### **For Each Cluster's Non-Replicated Shared Storage**

Identify CRS OCR (Oracle Cluster Registry) and VOTE-DISK shared storage files at the local and the remote clusters. OCR and VOTE-DISK files are not replicated to the remote cluster. These files are shared within the cluster.

### **For Local Storage for Oracle and CRS Binaries on each Node**

Identify local storage for the Oracle RDBMS and CRS installation.

## Software Installation and Configuration Steps

1. Create disk groups, volumes and CFS using Shared Replicated (/oradata) and Shared Non-Replicated disks (/ocrvote).
2. Install SFRAC with GCO option on primary and remote cluster. Install CRS and Oracle RDBMS local storage on both the clusters. Create the database on replicated storage on the primary cluster.
3. Configure Global Cluster Option by setting heartbeat, wac, gcoip and remotecluster resources in VCS. The following example displays the main.cf contents:

```
include "CFSTypes.cf"
include "CUMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "SRDFTypes.cf"

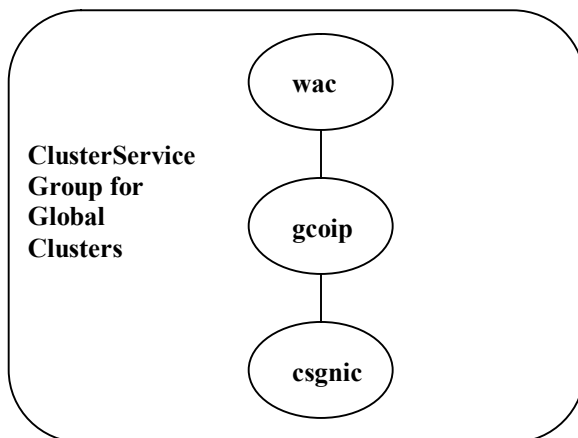
cluster Cluster_248 <
  UserNames = < admin = enoGniNkoJooMwoIn1 >
  ClusterAddress = "10.182.13.98"
  Administrators = < admin >
  HacliUserLevel = COMMANDROOT
  >

remotecluster Cluster_246 <
  ClusterAddress = "10.182.13.96"
  >

heartbeat Icmp <
  ClusterList = < Cluster_246 >
  AYATimeout = 30
  Arguments @Cluster_246 = < "10.182.13.96" >
  >

system thor248 <
  >

system thor249 <
  >
```



```

group ClusterService <
  SystemList = { thor248 = 0, thor249 = 0 }
  AutoStartList = { thor248, thor249 }
  OnlineRetryLimit = 3
  OnlineRetryInterval = 120
  >

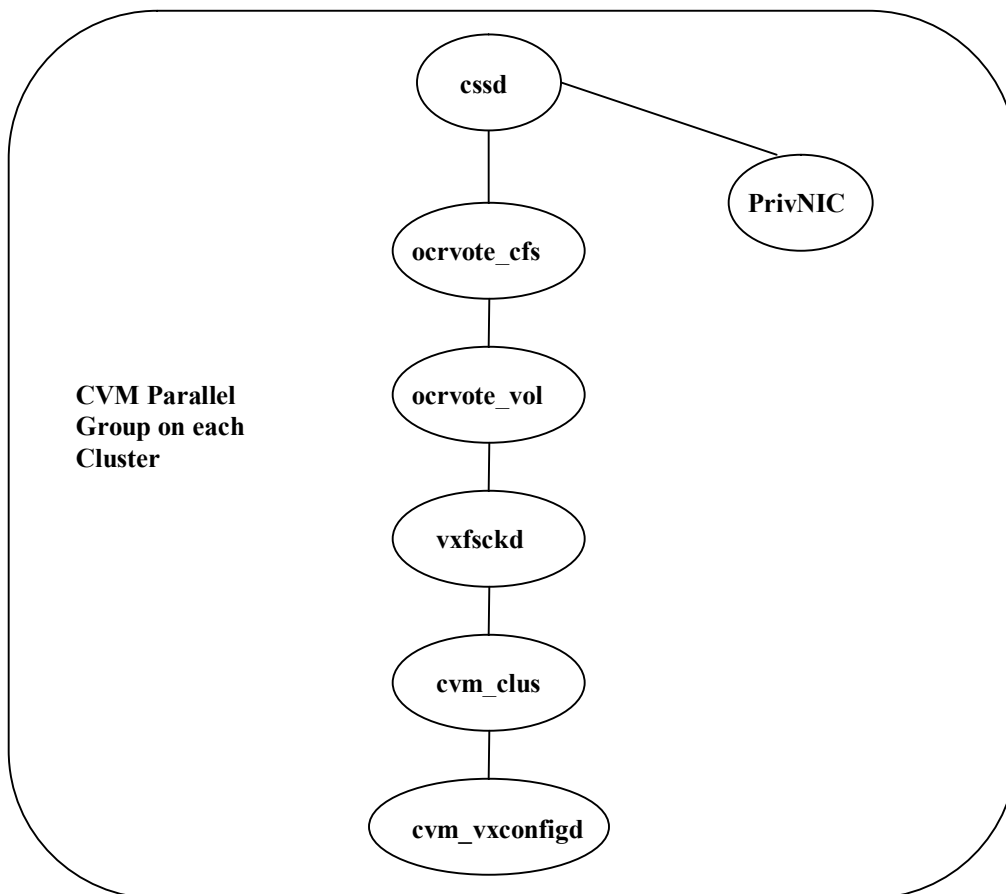
  Application wac <
    StartProgram = "/opt/URTSvcs/bin/wacstart"
    StopProgram = "/opt/URTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/URTSvcs/bin/wac" }
    RestartLimit = 3
  >

  IP gcoip <
    Device = hge0
    Address = "10.182.13.98"
    NetMask = "255.255.240.0"
  >

  NIC csgnic <
    Device = hge0
  >

  gcoip requires csgnic
  wac requires gcoip

```



```

group cvm <
  SystemList = { thor248 = 0, thor249 = 1 }
  AutoFailOver = 0
  Parallel = 1
  Authority = 1
  AutoStartList = { thor248, thor249 }
}

Application cssd <
  Critical = 0
  StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
  StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline stopcrs"
  CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
  MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
  OnlineRetryLimit = 20
}

CFMount ocrvote_cfs <
  Critical = 0
  MountPoint @thor248 = "/ocrvote"
  MountPoint @thor249 = "/ocrvote"
  BlockDevice @thor248 = "/dev/ux/dsk/ocrvote_dg_thor248_thor249/ocrvote_vol"
  BlockDevice @thor249 = "/dev/ux/dsk/ocrvote_dg_thor248_thor249/ocrvote_vol"
}

CFSfsckd vxfscd <
}

CUMCluster cvm_clus <
  CUMClustName = Cluster_248
  CUMNodeId = { thor248 = 0, thor249 = 1 }
  CUMTransport = gab
  CUMTimeout = 200
}

CUMVolDg ocrvote_vol <
  Critical = 0
  CUMDiskGroup @thor248 = ocrvote_dg_thor248_thor249
  CUMDiskGroup @thor249 = ocrvote_dg_thor248_thor249
  CUMVolume @thor248 = { ocrvote_vol }
  CUMVolume @thor249 = { ocrvote_vol }
  CUMActivation @thor248 = sw
  CUMActivation @thor249 = sw
}

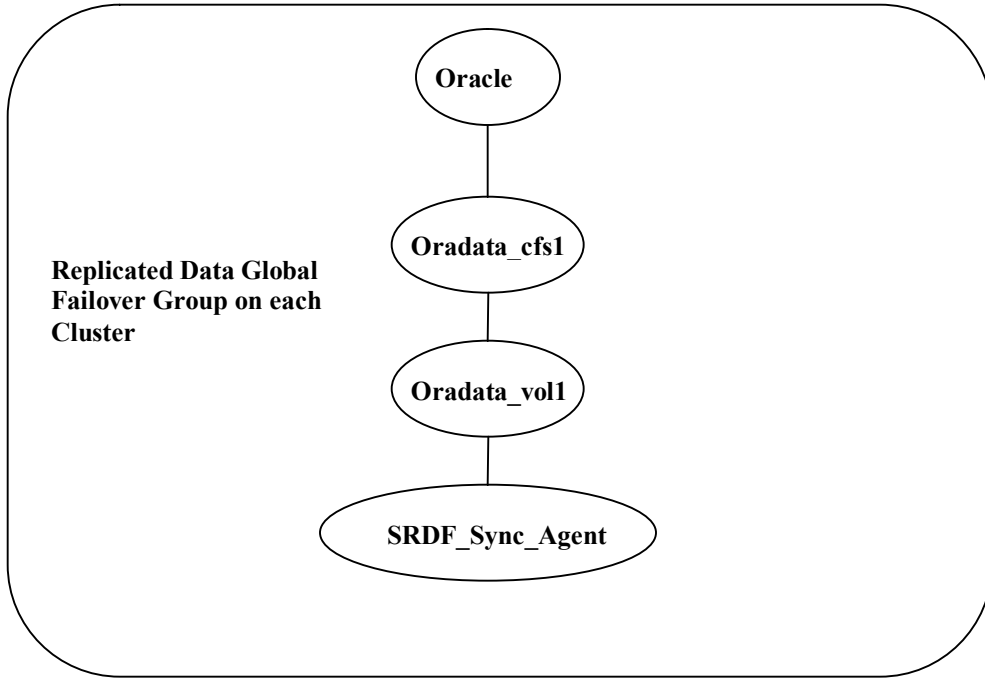
CUMvxconfigd cvm_vxconfigd <
  Critical = 0
  CUMVxconfigdArgs = { syslog }
}

PrivNIC ora_priv <
  Device @thor248 = { bge1 = 0, bge2 = 1 }
  Device @thor249 = { bge1 = 0, bge2 = 1 }
  Address @thor248 = "1.1.1.248"
  Address @thor249 = "1.1.1.249"
  NetMask = "255.255.255.0"
}

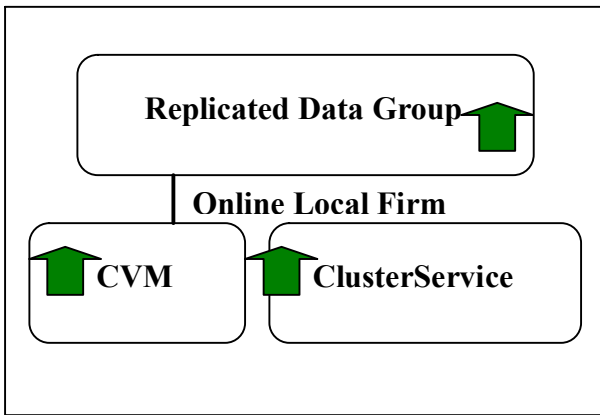
cssd requires ocrvote_cfs
cssd requires ora_priv
cvm_clus requires cvm_vxconfigd
ocrvote_cfs requires ocrvote_vol
ocrvote_vol requires vxfscd
vxfscd requires cvm_clus

```

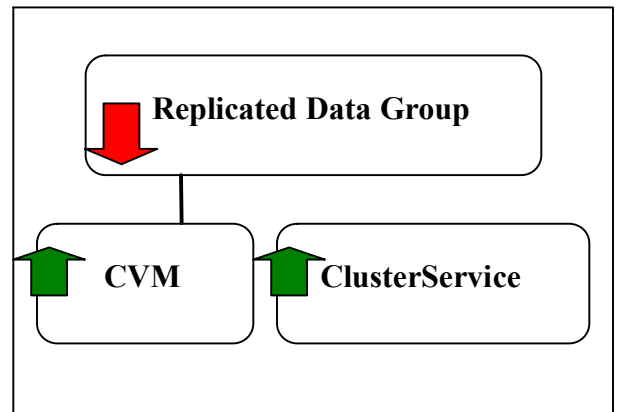
4. Install and configure the SRDF Agent. For Solaris, agent binaries are available as a package set which the 'pkgadd' command can install.
5. Configure the VCS Replicated Data Global Failover Group on each cluster. This group is online only on one cluster. 'hawizard srdf' can configure the SRDF agent.



**Primary Cluster**



**Remote Cluster**



= Online      = Offline

The secondary cluster main.cf contents are similar to primary. Note the Replicated Data Group is offline on the secondary until a primary failover occurs.

The following example displays part of the secondary cluster main.cf contents:

```
include "types.cf"
include "CFSTypes.cf"
include "CUMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "SRDFTypes.cf"

cluster Cluster_246 <
  UserNames = < admin = bk1DkfKh1G11Jt1Fki >
  ClusterAddress = "10.182.13.96"
  Administrators = < admin >
  HacliUserLevel = COMMANDROOT
>

remotecluster Cluster_248 <
  ClusterAddress = "10.182.13.98"
>

heartbeat Icmp <
  ClusterList = < Cluster_248 >
  AYATimeout = 30
  Arguments @Cluster_248 = < "10.182.13.98" >
>

system thor246 <
>

system thor247 <
>
```


## Failure and Recovery Scenarios

### Intra-Cluster Failover and Failback

#### Scenario A:

##### Failover – Node A is down

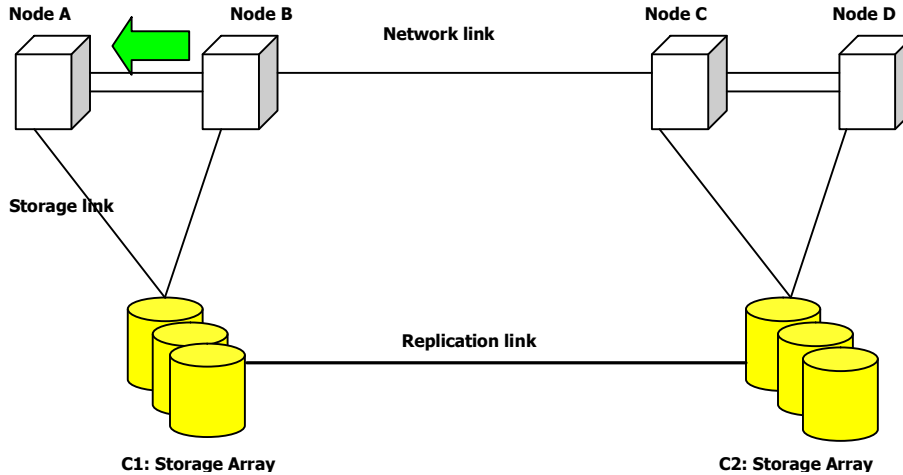
- Service Groups (Failover) from Node A failover to Node B.
- Parallel Service Groups like Oracle DB group remain online on Node B.



##### Failback – Node A is up

- Service Groups (Failover) from Node B can manually failback to Node A.
- Parallel Service Groups like Oracle DB group are online on Node A and Node B.
- No replication roles change.
- C2 remains write-disabled.






## Inter-Cluster Failover and Failback

### Scenario B:

#### Failover – Node A and Node B are down


- Global Service Groups failover to Cluster 2 from Cluster 1. The failover policy can be automatic or manual (default).
- The SRDF Agent reverses replication roles.
- Storage array on Cluster 1 is write-disabled.
- Storage array on Cluster 2 is write-enabled.



#### Failback – Node A and Node B are up

- When the primary nodes are back online, Global Service Groups can manually failback to Cluster 1
- The SRDF Agent reverses replication roles.


- Storage array on Cluster 1 is write-enabled.
- Storage array on Cluster 2 is write-disabled.



### Scenario C:

#### Failover – Storage Links are down on the Primary

- When the Storage link on just one node on the Primary is down, resources become offline on that node based on the service group and resource dependencies.
- When both the Storage links are down on the primary cluster, Global Service Groups failover to Cluster 2 from Cluster 1.
- The SRDF Agent reverses replication roles.
- Storage array on Cluster 1 is write-disabled.
- Storage array on Cluster 2 is write-enabled.



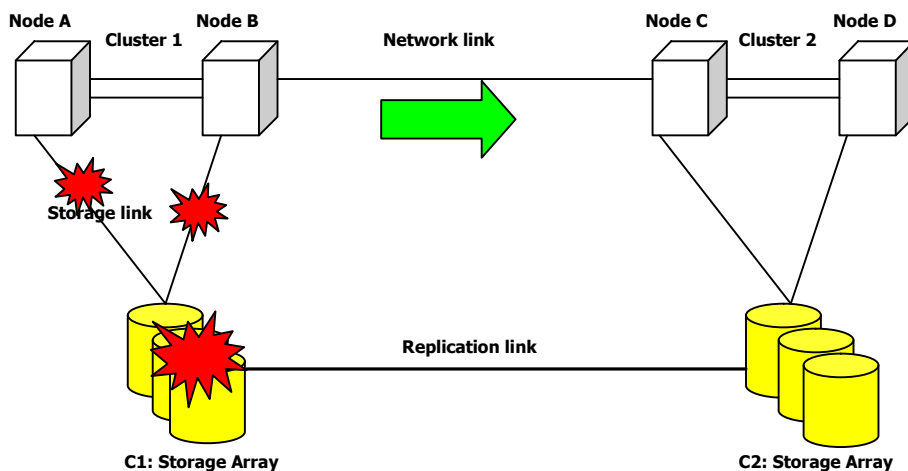
### Failback – Storage links are up on the Primary

- When the storage links are back online, Global Service Groups can manually failback to Cluster 1
- The SRDF Agent can reverse replication roles.
- Storage array on Cluster 1 is write-enabled.
- Storage array on Cluster 2 is write-disabled.

### Scenario D:

#### Failover – Primary storage array loss

- When the Primary Storage both the Storage links are down on the primary cluster, Global Service Groups failover to Cluster 2 from Cluster 1.
- Storage array on Cluster 2 is write-enabled.




#### Failback – Storage array and storage links are up on the Primary

- When the Storage array and links are back online, manually restore the primary site.
- Re-sync the Storage in the Primary from the secondary Array.
- Failback service groups from secondary to primary.
- Storage array on Cluster 1 is write-enabled.
- Storage array on Cluster 2 is write-disabled.

### Scenario E:

#### Loss – Replication link is down

- There is no failover.




### Restore – Replication link is up

- There is an attempt to auto-resynchronize the data in the secondary storage if the array supports it.

### Scenario F:

#### Loss: Network link is down

- GCO detects the failure and prompts for any manual failover by the administrator.



### Restore – Network link is up


- No action is required.

### Scenario G:

#### Loss – Replication and network links are down.

- No communication to the secondary.
- SRDF Agent attempts to enable C2 Storage Array on the secondary.

- Both the Arrays are write-enabled.



### Restore – Replication and network links are up

- Manual restoration is required to re-synchronize both the primary and the secondary.

## Symmetrix Remote Data Facility (SRDF)

EMC SRDF maintains a mirror data image across Symmetrix arrays which physically reside at separate sites. Symmetrix devices are designated as either a source or a target and SRDF synchronizes data between the source and the target devices. SRDF transfers data using one of the following protocols

- Synchronous (SRDF/S)
- Asynchronous (SRDF/A)
- Adaptive Copy ( for bulk data transfer )

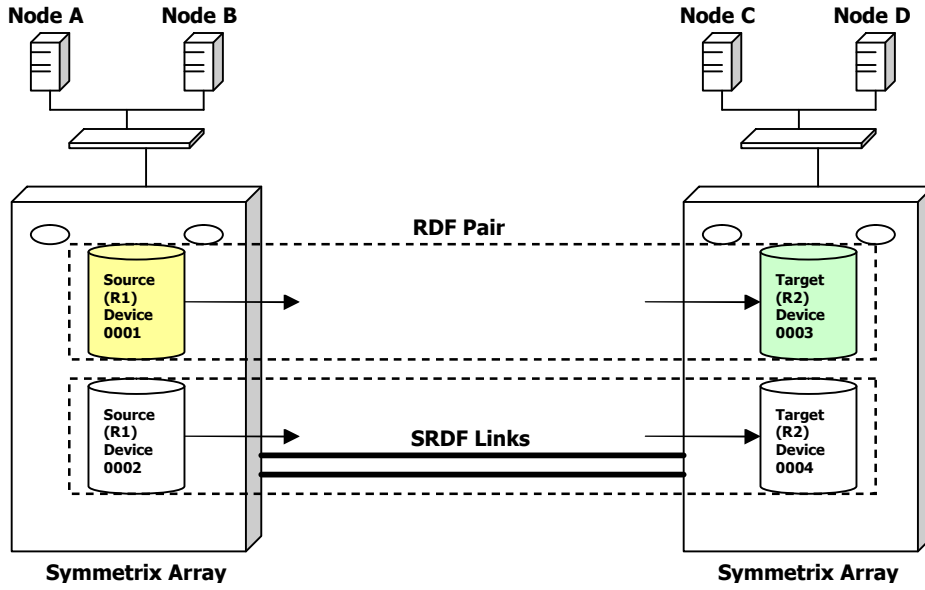
### SRDF Links

Symmetrix arrays can link together with Fibre Channel. Symmetrix DMX models can support mixed combinations of port type connectivity, including Fibre Channel, ESCON, FICON, Gig-E or iSCSI. Using T1/T3 or E1/E3 links, the distance between the target and the source can extend to 37.5 miles distance with Synchronous Replication.

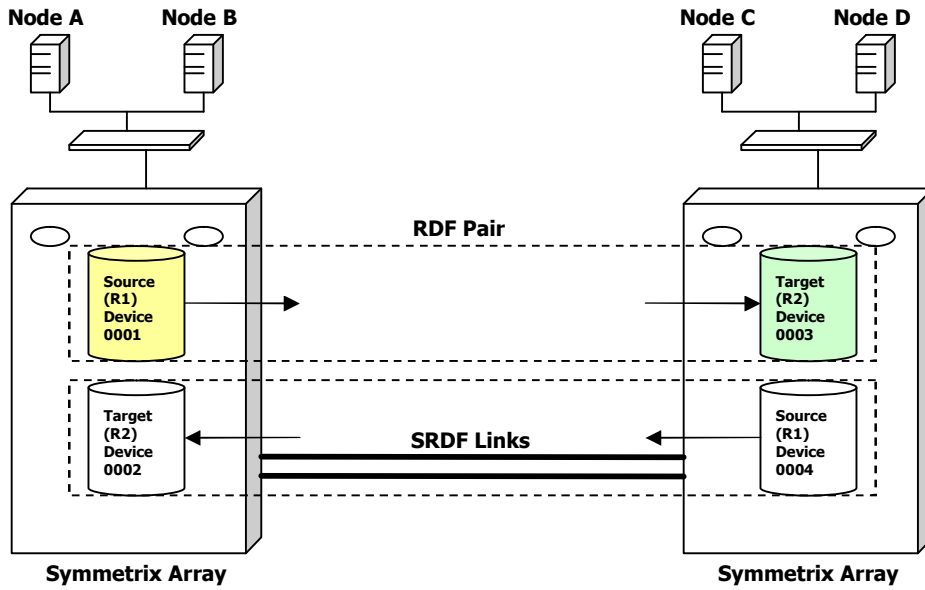
### Unidirectional / Bidirectional Data Transfer

Through configuration, SRDF can transfer data from one storage site to another site unidirectionally or bi-directionally. With unidirectional transfers, all source devices reside in a local Symmetrix array and all target devices reside in a remote Symmetrix array. Data transfers over the SRDF link from source devices to target devices. With bidirectional configurations, both the source and the target devices reside in each Symmetrix array as master and mirror copies.

### SRDF Unidirectional Configuration

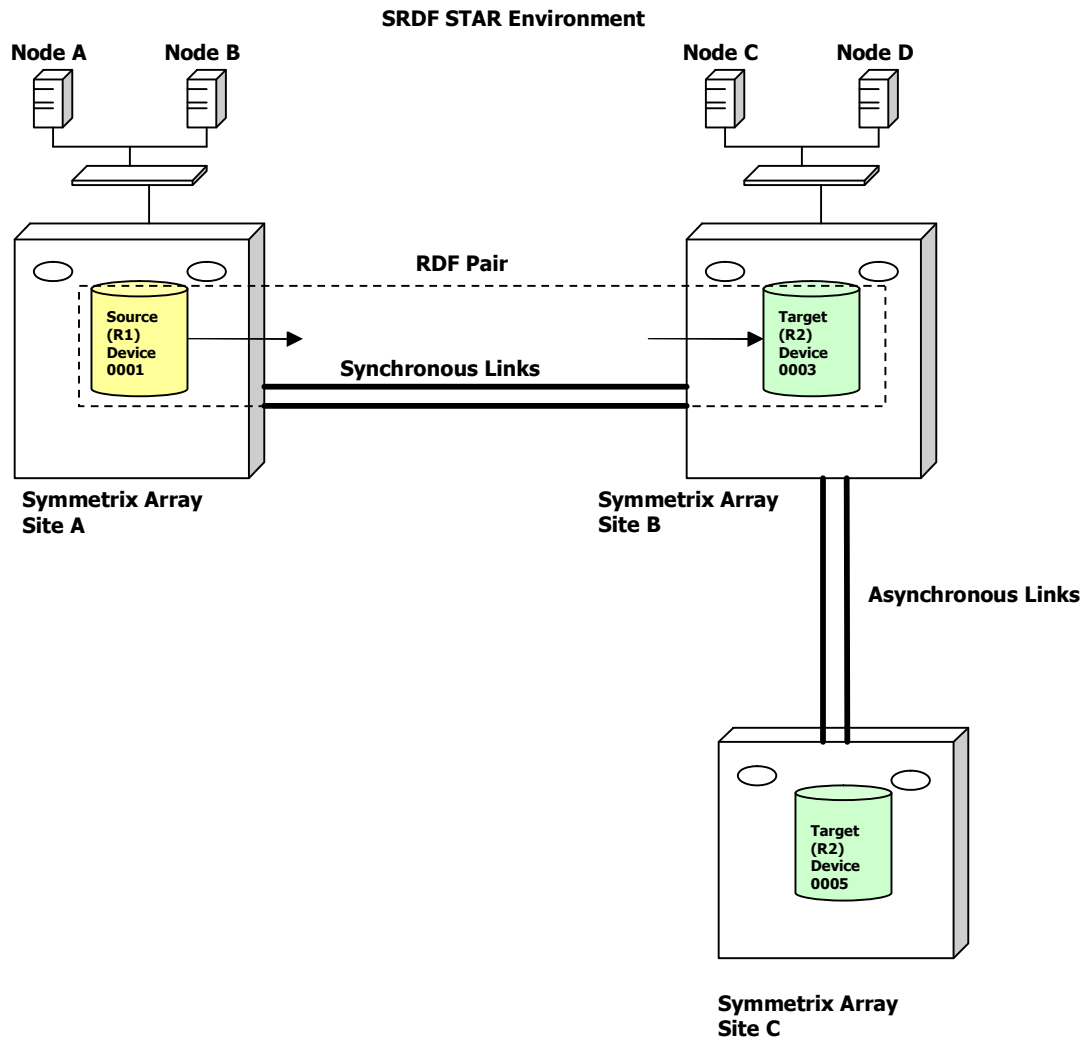


### SRDF Multidirectional Configuration



## SRDF Topologies

Symmetrix supports other topologies such as Multi-Hop and SRDF/Star. In a Multi-Hop topology, three Symmetrix sites connect using RDF links. The third RDF site backups the remote RDF site. In SRDF/Star environments, the source site (Site A) is synchronously mirrored on a remote site (Site B) and asynchronously to another remote site (Site C).



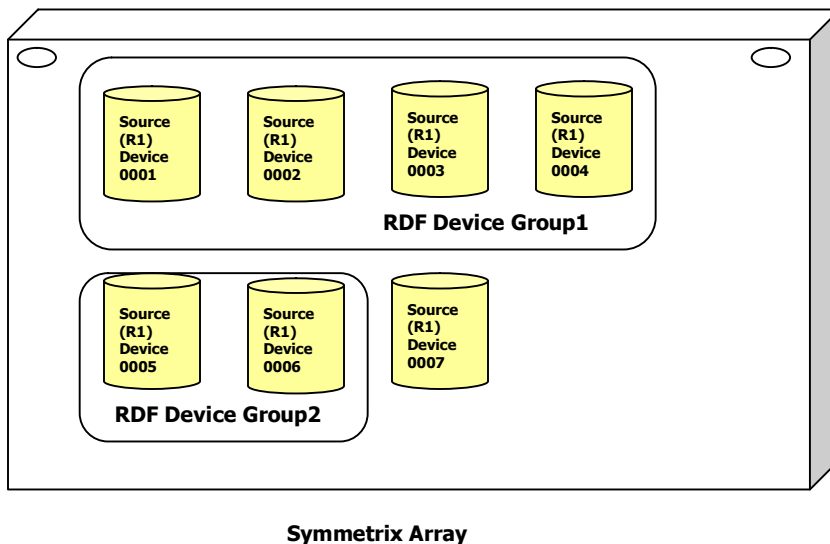
## SRDF Device Types and Groups

Individual Symmetrix devices are either assigned as a source (R1 device) or a target device (R2 device) to synchronize and maintain remote mirroring activities. If the source device fails, target device data can be accessed by the local host. Once the source device is replaced, it can be resynchronized. An RDF pair consists of a source device (R1) and the corresponding target device (R2).



An RDF device group is a user-defined group consisting of a set of RDF devices belonging to a single Symmetrix array. The RDF device group could be of REGULAR, RDF1 (source) or RDF2 (target) types and is defined during the time of the group creation. Each group may contain various device lists for standard, BCV, virtual (VDEV), and remote devices.

By default, a device cannot belong to more than one device group. But, you can modify the default behavior to allow the device to belong to multiple groups. Using device group, you can execute group commands applicable to all devices in the group. For example, the `'symrdf -g RAC establish -full'` command initiates a full replication of all devices in the RDF group RAC. All the data from the source devices are copied over the target devices and replication is established.



SRDF provides capabilities to create composite, dynamic and consistent RDF groups. A composite group can have multiple Symmetrix arrays and RDF groups. A consistency group preserves the integrity and dependent write consistency of a database distributed across multiple arrays. Consistency groups prevent dependent I/O from getting out of sync, thus ensuring the integrity and consistency of the remote data. By enabling consistency protection to a device group, if data cannot be copied from the R1 to R2 for a particular device then all the devices in



the group is made Not Ready on the link to preserve target data consistency. This feature helps databases to safely recover at the target location.

The SRDF devices and links maintain a (Ready, Not Ready or Write Disabled) state and based on the SRDF control operations the status of their state changes.

## SRDF Setup and Management

Symmetrix command line interface (SYMCLI) can be used to configure and manage SRDF. The following examples describe some of the basic commands. For further information refer to the Symmetrix SRDF documentation.

### To display all the Symmetrix array that are reachable through SRDF links

symcfg list

```
thor248# symcfg list
```

S Y M M E T R I X						
SymmID	Attachment	Model	Mcode Version	Cache Size (MB)	Num Phys Devices	Num Symm Devices
000187400527	Local	DMX1000S	5671	16384	264	1709
000187900584	Remote	DMX800	5671	8192	0	846
000190100422	Remote	DMX3-24	5771	24576	0	1597

```
thor248#
```

### To list the device group by name and type

symdg list

```
thor248# symdg list
```

D E V I C E      G R O U P S								
Name	Type	Valid	Symmetrix ID	Devs	GKs	Number of BCUs	Number of UDEUs	TGTs
RAC_async	RDF2	Yes	000187400527	6	0	0	0	0
RAC	RDF1	Yes	000187400527	6	0	0	0	0
RAC-nonrep	REGULAR	Yes	000187400527	5	0	0	0	0

```
thor248#
```

### To display SRDF devices that are visible to the host

symrdf list pd

```
thor248# symrdf list pd
Symmetrix ID: 000187400527

-----
Local Device View
-----
Sym   RDF   STATUS   MODES   R1 Inv   R2 Inv   RDF   S T A T E S
Dev  RDev  Typ:G   SA RA LNK  MDA   Tracks  Tracks Dev  RDev  Pair
-----
0042 0024  R1:11   RW RW RW   S..   0       0 RW  WD   Synchronized
0049 002B  R2:11   RW WD RW   S..   0       0 WD  RW   Synchronized
004A 002C  R2:11   RW WD RW   S..   0       0 WD  RW   Synchronized
004B 002D  R2:11   RW WD RW   S..   0       0 WD  RW   Synchronized
004C 002E  R2:11   RW WD RW   S..   0       0 WD  RW   Synchronized
004D 002F  R2:11   RW WD RW   S..   0       0 WD  RW   Synchronized
004E 0030  R1:11   RW RW RW   S..   0       0 RW  WD   Synchronized
004F 0031  R1:11   RW RW RW   S..   0       0 RW  WD   Synchronized
0050 0032  R1:11   RW RW RW   S..   0       0 RW  WD   Synchronized
0051 0033  R1:11   RW RW RW   S..   0       0 RW  WD   Synchronized
0052 0034  R1:11   RW RW RW   S..   0       0 RW  WD   Synchronized
0053 0035  R2:8    RW RW NR   A..   2       0 RW  RW   Split
0054 0036  R2:8    RW RW NR   A..   2       0 RW  RW   Split
0055 0037  R1:12   RW RW RW   S..   0       0 RW  WD   Synchronized
0056 0038  R1:12   RW RW RW   S..   0       0 RW  WD   Synchronized
0057 0039  R1:12   RW RW RW   S..   0       0 RW  WD   Synchronized
```

To view the SRDF information on all the devices in a device group

```
symrdf -g device_group query
```

```
thor248# symrdf -g RAC query
Device Group (DG) Name      : RAC
DG's Type                  : RDF1
DG's Symmetrix ID          : 000187400527

-----
Source (R1) View           Target (R2) View           MODES
-----
Standard   ST           LI           ST
Logical    A           N           A
Device     I   R1 Inv   R2 Inv   K   I   R1 Inv   R2 Inv
Dev        E   Tracks  Tracks  S   Dev  E   Tracks  Tracks  MDA   RDF Pair
-----
DEU001  0042  RW        0        0 RW 0024  WD        0        0 S..   Synchronized
DEU002  004E  RW        0        0 RW 0030  WD        0        0 S..   Synchronized
DEU003  004F  RW        0        0 RW 0031  WD        0        0 S..   Synchronized
DEU004  0050  RW        0        0 RW 0032  WD        0        0 S..   Synchronized
DEU005  0051  RW        0        0 RW 0033  WD        0        0 S..   Synchronized
DEU006  0052  RW        0        0 RW 0034  WD        0        0 S..   Synchronized

Total
Track(s)           0        0        0        0
MB(s)             0.0      0.0      0.0      0.0

Legend for MODES:
M(code of Operation): A = Async, S = Sync, E = Semi-sync, C = Adaptive Copy
D(omino)             : X = Enabled, . = Disabled
A(daptive Copy)      : D = Disk Mode, W = WP Mode, . = ACp off
```

To create an RDF Device Group of type RDF1

```
symdmg create device_group -type device_type
```

```
thor248# symdmg create RAC -type RDF1
thor248# symrdf -g RAC query
Cannot perform the operation because no devices have been added to the group
```

## To add devices into an RDF Device Group

```
symld -g device_group add dev device_id
```

```
thor248# symld -g RAC add dev 0042
thor248# symld -g RAC add dev 004E
thor248# symld -g RAC add dev 004F
thor248# symld -g RAC add dev 0050
thor248# symld -g RAC add dev 0051
```

To create an SRDF pair

```
symrdf createpair -sid array_id -file devicefile -type rdf1 -rdfg grpnum -establish
```

Array\_id is the local Symmetrix Array Id (Example: 000187400527)

The last four digits are sufficient in most cases. (Example: 0527)

The grpnum is the RDF (RA) group number. 'symrdf list pd' command displays the RDF Type and Group number. In the above example, the RDF group number is 11.

A devicefile contains device pairs as mentioned in the below format.

```
0042 0024
```

```
004E 0030
```

```
004F 0031
```

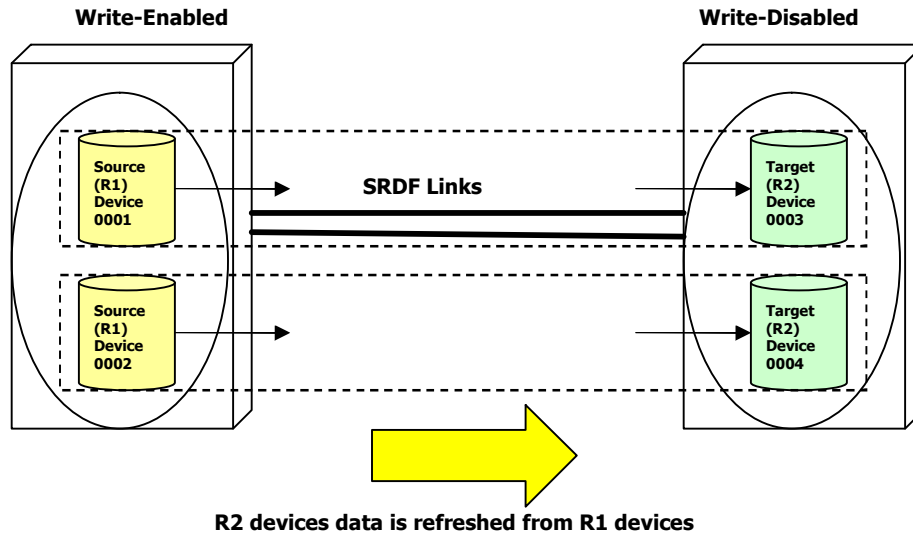
```
0050 0032
```

```
0051 0033
```

```
0052 0034
```

## To perform establish on RDF pairs

When a full/incremental establish is initiated, the target devices(R2) are write disabled and all the tracks on the target devices are refreshed by the corresponding R1 devices.



`symrdf -g device_group establish -full` (For a full establish)


`symrdf -g device_group establish -full DEV001` (For a full establish on one RDF pair)

`symrdf -g device_group establish` (for an incremental establish)

`symrdf -g device_group establish DEV001 DEV002` (for an incremental establish on two RDF pairs)


### Other SRDF Control Operations

Using SRDF control commands, the replication process can be halted and the target device data can be accessed for reading and writing. With a full restore command, the entire contents of the target devices are copied to the source devices.



**When a restore command is issued, R2 devices data is copied to R1 devices**

With a failover command, the source devices are write-disabled and the target devices are write-enabled. The target devices are available for write access. With a failback command, the R2 data changes are copied back to the R1 data and the roles are switched back. The source (R1) devices are write-enabled and the target (R2) devices are write-disabled.



**When Failover command is issued on SRDF devices**

## Conclusion

Storage Foundation for Oracle RAC with SRDF simplifies an organization's ability to recover mission-critical applications. The accelerated recovery time provides the ability to recover in a well-defined, highly-automated manner. Using synchronous replication, the failovers can be tested switching the remote cluster as the primary cluster without loss of transactional data. Using VCS SRDF Agent, the replication direction reverses very easily.

## About the Author

**Naveen Williams** is a Symantec Principal SQA Engineer specializing in Oracle Database Administration. Before joining Symantec, he had worked at Sun Microsystems for 6 years as a Senior Oracle DBA for Production support. He has a Masters Degree in Computer Science from Northwestern Polytechnic University, California.

## About Symantec Technology Network (STN)

**Symantec Technology Network (STN)** is Symantec Corporation's technical information generation and dissemination organization. It distributes a free monthly technical newsletter that discusses timely technology events to 120,000 email subscribers across the globe. STN also publishes technical data storage and security articles each month for large enterprise and Small and Medium Business (SMB) readers, as well as hosts a variety of blogs and product discussion forums discussing Symantec product tips and insights. To subscribe to STN's Technical Newsletter and review other STN materials, please visit STN at:

<http://www.symantec.com/enterprise/stn/index.jsp>

**W. David Schwaderer** is STN's Editor-In-Chief. He has authored ten technical books, and six commercial software programs. His eleventh book titled *Innovation Survival* is nearing completion. He has a Masters Degree in Applied Mathematics from the California Institute of Technology and an MBA from the University of Southern California. He lectures at Stanford in Silicon Valley on the subject of innovation.

**Russ G. Gregg** is STN's Video Content Producer. He is presently focusing on producing state-of-the-art digital video for the Web and small form-factor, portable video devices such as Apple Inc.'s iPhone™ and Video iPod®. Russ has a Bachelor of Arts in English Literature, Texas from A&M University.

## **About Symantec**

Symantec is the world leader in providing solutions to help individuals and enterprises assure the security, availability, and integrity of their information. Headquartered in Cupertino, Calif., Symantec has operations in more than 40 countries. More information is available at [www.symantec.com](http://www.symantec.com).

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 800 745 6054.

Symantec Corporation  
World Headquarters  
20330 Stevens Creek Boulevard  
Cupertino, CA 95014 USA  
408 517 8000  
800 721 3934  
[www.symantec.com](http://www.symantec.com)

Symantec and the Symantec logo are U.S. registered trademarks of Symantec Corporation. Other brands and products are trademarks of their respective holder/s. Any technical information that is made available by Symantec Corporation is the copyrighted work of Symantec Corporation and is owned by Symantec Corporation. NO WARRANTY. The technical information is being delivered to you as-is and Symantec Corporation makes no warranty as to its accuracy or use. Any use of the technical documentation or the information contained herein is at the risk of the user. Copyright © 2007 Symantec Corporation. All rights reserved. 09/04 10318317