Bull

Recovery from Cluster Site Disaster
User's Guide

AIX

Bull

Recovery from Cluster Site Disaster
User's Guide

AIX

Software

January 2002

Trademarks and Acknowledgements

We acknowledge the right of proprietors of trademarks mentioned in this book.

AIX® is a registered trademark of International Business Machines Corporation, and is being used under licence.

UNIX is a registered trademark in the United States of America and other countries licensed exclusively through the Open Group.

About This Book

This manual will help you to install, configure and use the Recovery from Cluster Site Disaster application.

For definitions of abbreviations, see Glossary

Who Should Use This Book

This manual is intended for cluster administrators and anyone who monitors HACMP configuration.

This manual assumes that you are familiar with the HACMP product and your host operating system. You should also know how to perform a basic set of actions in a Windows–type environment, including:

- choosing menu commands,
- moving and resizing windows,
- opening icon windows,
- dragging and dropping icons,
- using mouse controls for your system.

Overview

The manual is organized as follows:

- Overview of Disaster Recovery Solutions,
- Installation,
- Cluster Site Reconfiguration Utility,
- Planning Site Recovery Configuration,
- Recovery from Site Disaster Situation,
- Restore Initial Configuration,
- Importing Volume Groups Using Shared Volumes,
- Glossary.

Related Publications

- Site Monitoring for Disaster Recovery User's Guide, 86 A2 83JX.
- HACMP Concepts and Facilities, 86 A2 54KX.
- HACMP Enhanced Scalability Installation and Administration Guide, Vol. 1, 86 A2 62KX.
- HACMP Enhanced Scalability Installation and Administration Guide, Vol. 2, 86 A2 89KX.
- PowerConsole & Cluster Assistant Setup Guide, 86 A2 81HX.
- EPC & HA Solutions Setup Guide, 86 A2 79HX05.
- OSM SRDF Components for UNIX Product Guide, P/N 200–999–565.
- System Management Design and Development Guidelines – IBM.

Table of Contents

Chapter 1. Overview of Disaster Recovery Solutions

Disaster recovery is a generic name that encompasses various options, with one goal in mind, answering customers number one issues: the recovery of their data or business exploitation in case of major natural disasters. These disasters can be classified from water flooding of an Information System, through fire disaster, up to earthquake ... etc.

Whatever the nature of the disaster, two main types of recovery needs can be identified:

- APPLICATION and DATA recovery: This gives total recovery of the business environment. The customer requires all applications and critical data to his business environment to be recovered and up and running again as fast as possible on the back–up site,

- DATA ONLY recovery: This gives recovery of business critical data only (database, plus a variable part of the data archived for business sake).

Depending on the type of disaster which is being targeted, the customer may add a requirement of distance between the master/backup site for both types of disaster recovery.

distance resource	room	building	campus	geographic	geographic
APPLI + DATA	Bull HA Configuration with DAS (SCSI) or SSA	Bull HA Configuration with DAS (FC) or SSA	Bull HA Configuration with DAS (FC)	To be evaluated on a case per case basis	To be evaluated on a case per case basis
DATA ONLY	Bull ESCALA Systems with DAS (SCSI) or SSA or SRDF/SYMME-TRIX or Oracle replication	BULL ESCALA Systems with DAS (FC) SSA or SRDF/SYMME-TRIX or Oracle replication	BULL ESCALA Systems with DAS (FC) SSA or SRDF/SYMME-TRIX or Oracle replication	BULL ESCALA Systems with DAS SRDF/SYMME-TRIX or Oracle replication (ESCON link)	BULL ESCALA Systems with DAS SRDF/SYMME-TRIX or Oracle replication (ATM link)
Metric distance	< 25 m	< 1000 m	< 10 km	< 60 km limitation due to ESCON link	> 60 km no limitation

Note: Oracle Replication solution is explained in the white paper document available on the Bull web site: **Keep disaster at a distance**

http://www–frec.bull.com –> White papers –> Keep disaster at distance

Bull Disaster Recovery Solutions

There are two Bull disaster recovery solutions depending on the range of the disaster:

- Campus disaster recovery.
- Geographic disaster recovery: inter site monitoring and recovery from cluster site disaster.

Campus Disaster Recovery

The **Campus Disaster Recovery** allows data mirroring on an extended cluster.

Each node in a cluster is defined over two sites, and must belong to the same subnet to respect HA requirement for public network connection even if only Applications and Data Takeover is possible.

Serial networks used for HACMP heartbeat between cluster nodes must be extended using RS232 micro–modems.

If a disaster is diagnosed, HACMP on the surviving node will take over all the resources of the failed node.

For distances over 1 km and under 10 km, Bull delivers a recovery method based on the technology of Fibre–Channel (FC), fiber disks subsystems, optical adapters and fiber links.
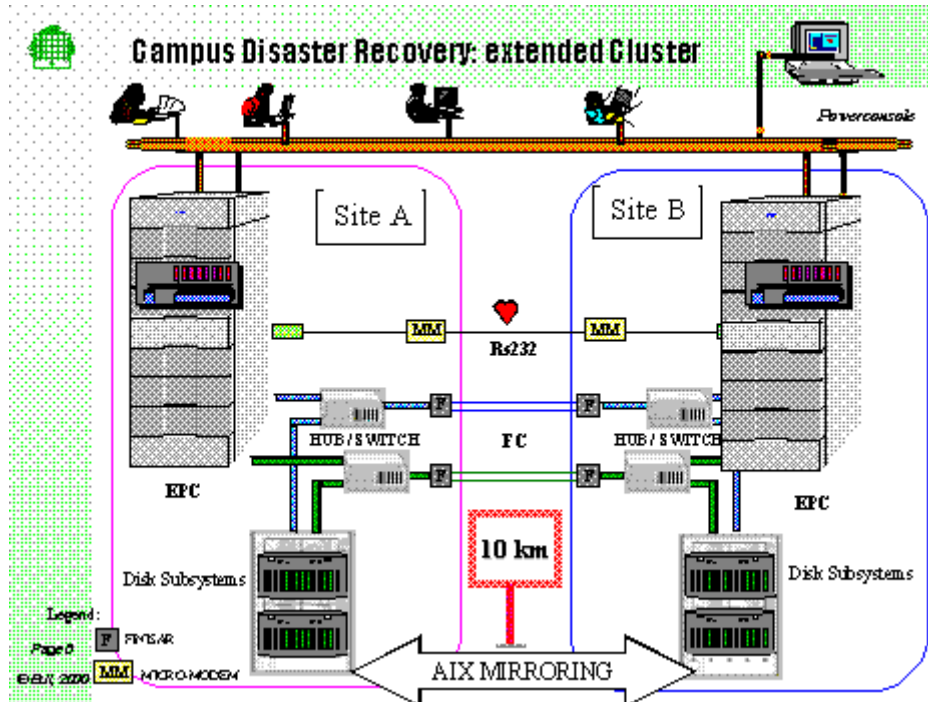

Figure 1. Campus Disaster Recovery

Geographic Disaster Recovery

Geographic recovery is specifically targeted at Disaster Recovery requirements over long or very long distances:

- For distances over 10km, Bull provides DATA Disaster Recovery solutions in partnership with EMC and ORACLE and can answer DATA and APPLICATION Disaster Recovery Solutions on a case by case basis.

The clustering software HACMP is easily adaptable to local site monitoring, allowing:

- diagnosis of system errors and network errors,

- high availability of data (external disk subsystems) by using data mirroring.

But, this functionality is not well adapted to the inter site availability (heavy infractructure constraints, no data mirroring integration…).

Inter Site Monitoring and Recovery from Cluster Site Disaster

To take into account the inter site availability, some mechanisms are proposed:

- The **Inter Site Monitoring** mechanism.

- The **Recovery from Cluster Site Disaster** mechanism:

   – to define the methods of replication of data and applications from one site to another,

   – to allow the reconfiguration and the restart of activities (applications) on the mirror site when the system administrator makes the decision to proceed with the site takeover,

   – to restore the initial configuration.
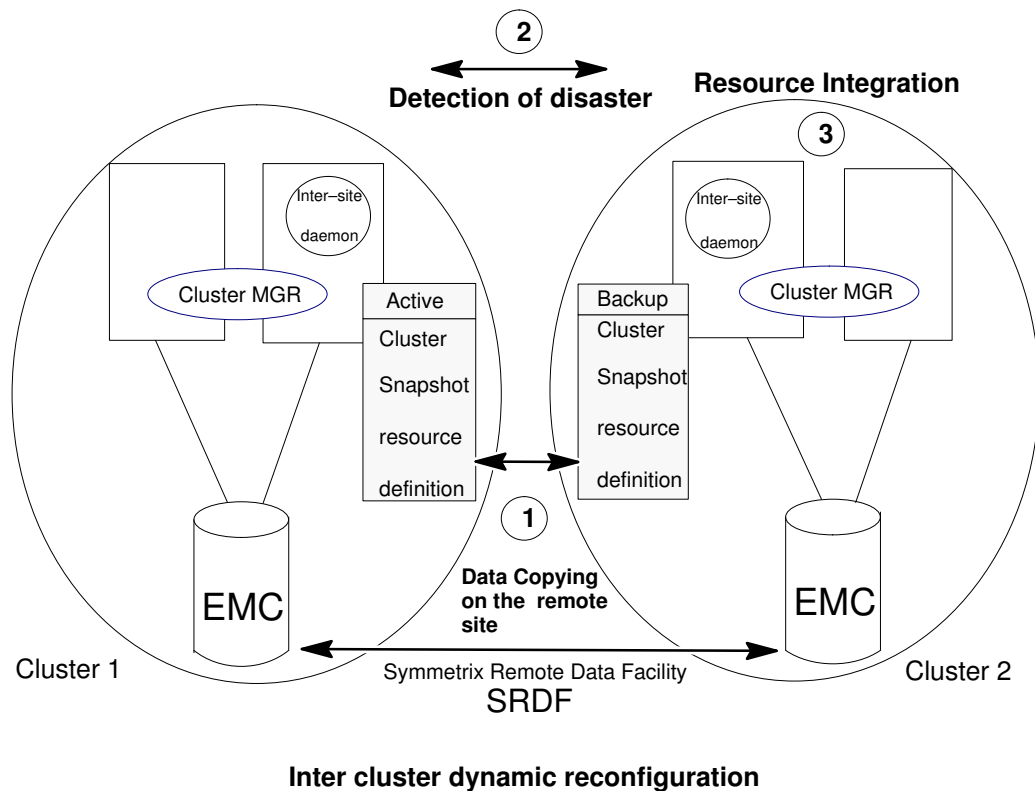

Inter cluster dynamic reconfiguration

Figure 2.   Inter–site monitoring and Recovery from cluster site disaster

The Geographic Disaster Recovery allows data mirroring between clusters on separate sites.

In case of disaster, the failing cluster site can be reconfigured on the surviving cluster site, and the activities can be restarted.

Inter Site Monitoring

The site monitoring functionality allows all the clusters to be monitored to detect a disaster.

A disaster is diagnosed by the monitoring daemon when the site is unreachable.

On each site, a daemon runs on one cluster node. If the cluster node (on which the monitoring daemon is running) is stopped (due to a normal or abnormal reason) using HACMP feature, this daemon is restarted on another cluster node.

To exchange heartbeats, monitoring daemons use a public IP network and a telephone network in case of an IP network failure.

To manage the cluster and to monitor the other site during a maintenance operation on the cluster, a Bull PowerConsole (on which a monitoring daemon is running) is required on each site.

If a disaster is diagnosed by the monitoring daemon, an appropriate alert is sent to the system administrator.

For more information about the **Inter Site Monitoring** facility, refer to "Site Monitoring for Disaster Recovery" User's Guide, 86 A2 83JX.

Recovery from Cluster Site Disaster

When the system administrator makes the decision to proceed with the site takeover, the failing cluster site is automatically reconfigured on the surviving site and the activities are restarted.

The following diagram shows the different components of the geographic disaster recovery infrastructure:

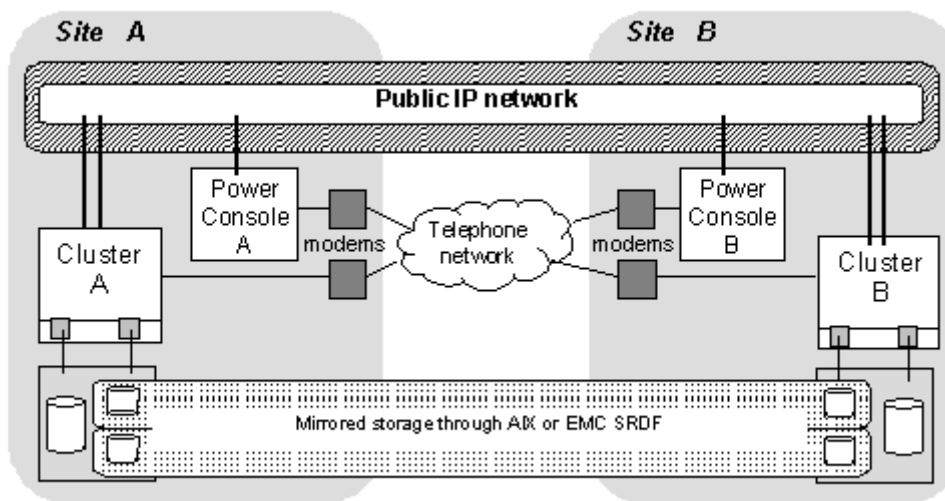
Figure 3.   The geographic disaster recovery infrastructure

Hereafter, a configuration of two sites is shown. An external IP network (for example Internet) links the two LAN's.

The EMC Symmetrix subsystems are connected by an SRDF link. Each Cluster site can communicate with the other one and its PowerConsole through a phone link.
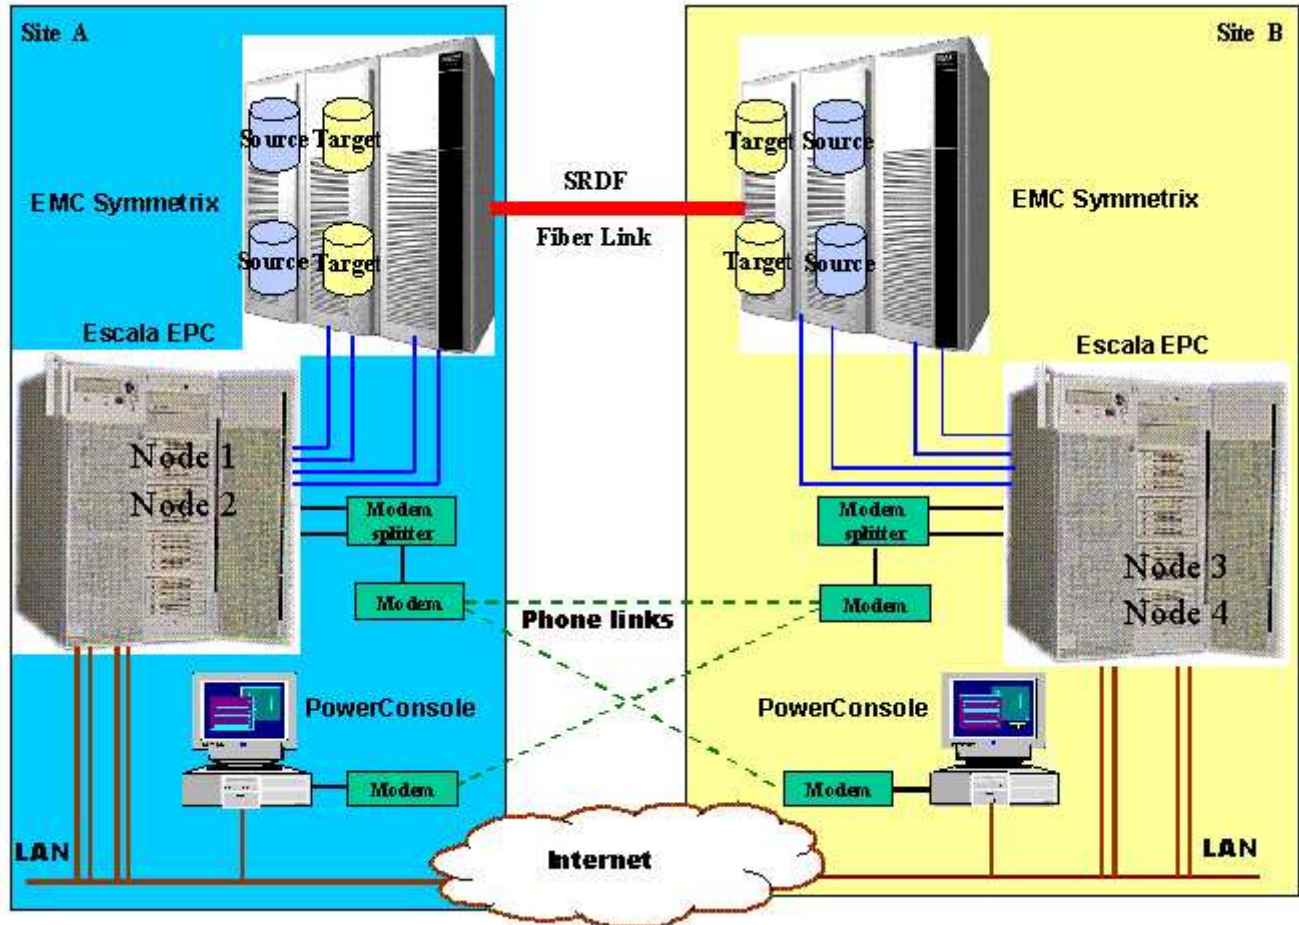

Figure 4.   Geographic disaster recovery: distant clusters

Overview of Recovery from Cluster Site Disaster

The Recovery from Cluster Site Disaster allows the system administrator:

* to define the method of data replication,

* to reconfigure and restart all activities from the failing site to the surviving site,

* to restore the initial configuration of the cluster site.

In normal mode, the HACMP clusters ensure the local monitoring of their resources.

With the inter–site recovery, the HACMP clusters ensure the reconfiguration and the dynamic restart of the failing cluster site applications.

Data Replication

The Data Replication depends of the level of data lost that can be acceptable.

If in case of disaster, the restart of the application with one day of data lost is acceptable, the backup / restore solution can be considered.

If no data lost is acceptable, two solutions can be considered:

* A solution with the **Symmetrix Remote Data Facility** of the EMC Company ensures the data replication between remote sites.

   To minimize the degradation of performance due to distance and telecommunication delays, EMC recommends using the semi–synchronous mode for the extended distance connections (> 60 km).

   In normal mode, the file used by a cluster site belongs to a local copy (source) in an active state, and a remote copy (target) is locked.

   In case of disaster, some scripts delivered with the SRDF facility allow to unlock the target copy and to make available the target copy in an active state.

   In the same way, to come back to the initial configuration (after the failing site was repaired) others scripts are delivered with the SRDF facility allowing to update the data copy.

* Another solution is to use the **AIX Mirroring Facility** that ensures the data replication between remote sites.

Cluster Site Reconfiguration

To reconfigure the failing cluster site resources, the surviving cluster site takes into account the last saved snapshot of the failing site.

From the snapshot:

* A checking is done to ensure the compatibility with the configuration of the surviving cluster site.

* The resources to configure are extracted from the snapshot.

* The system administrator can choose the location of the failing cluster resources on the surviving cluster nodes, and the reconfiguration of the failing cluster resources will be applied.

Restore Initial Configuration

Once the failing cluster site was repaired, this service is used to integrate again the initial configuration of the surviving Cluster Site.

Note: The remote site belongs to a different network. The remote IP addresses of the HACMP configuration will be migrated on the surviving cluster site and aliased on the network adapter of the surviving cluster site. It will be necessary to verify the address accessibility in the context of the customers, via a network solution ex: Network Address Translation, or via the update of the network route).

Chapter 2. Software Installation

Prerequisites

- AIX 4.3.3 minimum

- HACMP/ES 4.3.1 minimum

- Bullcluster 4.3.3 minimum with following minimum fileset release:

  – *bullcluster.basic_soft.ha_cmp* 4.3.3.7

  – *bullcluster.classistant.Dt* 4.3.3.7

  – *bullcluster.classistant.help.en_US* 4.3.3.6

Software Delivery

The Site Monitoring for geographic Disaster Recovery software is available on the Site Monitoring 4.3.3 CD–ROM.

Software Installation

The prerequisites and Site Monitoring for geographic Disaster Recovery are AIX LPPs which can be installed with install (or SMIT install) as follows:

1. `smit install_all`

   – Select either *bullcluster.bundles.sitemon.cons* on the PowerConsole, or *bullcluster.bundles.sitemon.node* on the nodes.

2. `smit install_bundle`

   – Select either *sitemon.cons* on the PowerConsole, or *sitemon.node* on the nodes.

Chapter 3. Cluster Site Reconfiguration Utility

Reconfiguration Utility Purpose

The Cluster Site Reconfiguration utility is used to:

- integrate resources from the failed cluster to the surviving cluster after a disaster,

- execute some checking even if no disaster is detected.

Its interface allows to create new resource groups or to add resources from the failing cluster to the existing resource groups on the surviving cluster.

Inputs

The inputs of the Cluster Site Reconfiguration utility are:

- the snapshot of the remote cluster,

- the configuration of the local cluster.

Script

A script is generated from this snapshot to add or modify HACMP resource group(s) on the local cluster.

This script is then executed on the local cluster to take into account all the resources.

Usage

Before a Disaster

If the Cluster Site Reconfiguration utility is used while no disaster is detected, the system administrator can:

1. Transfer the snapshot from the remote cluster to the local cluster (from time to time, at a date, ...), (see page 4-1).

2. Verify this snapshot (see page 4-2).

3. Generate the script (only to verify what the script do) (see page 4-7).

After a Disaster

If the Cluster Site Reconfiguration utility is used after a disaster (in this case, the remote cluster is the failed cluster and the local cluster is the surviving cluster), the system administrator:

1. Checks the configuration of the shared SRDF volumes, if the SRDF mechanism of the EMC is used (see page A-1).

2. Verifies the snapshot from the failed cluster (see page 4-2).

   Not mandatory: the verification is performed when the script is executed.

3. Generates the script (see page 4-7),

4. Applies this generated script (see page 5-1).

To have a correct behavior:

- All the volume groups of each cluster site have to be imported on each node of the other cluster site.

- The transfer of the snapshot of the remote cluster and the checking phase of this snapshot is recommended.

Chapter 4. Planning Site Recovery Configuration

The following operations must be planned:

- Transfer the snapshot from one site to another, page 4-1.

- Verify snapshot, page 4-2.

- Validate AIX Mirroring (if needed), page 4-5.

- Generate Reconfiguration script, page 4-7.

- Import Volume Groups, page 4-8.

Transfer the Snapshot from one Site to another

To transfer the snapshot from one site to another, the following solutions can be used:

- Magnetic support exchanged between site

  Snapshots are executed periodically on each site and saved on a magnetic support, which is sent to the other site, and restored on one Cluster node. This solution is fully manual.

- Network

  The two sites are assumed to communicate through the network. Snapshots are executed periodically and sent to one cluster node of the other site using the network. This operation can be done manually or with a cron command that executes the snapshot and transfers it.

- shared storage

  Snapshots on both sides are executed periodically and stored on a shared storage. This can be executed manually or with a cron command that executes the snapshot and transfers it.

The snapshot files (<snapshot name>.odm and <snapshot name>. bullodm) must be transferred from the *usr/sbin/cluster/snapshots/bull* directory to the *usr/sbin/bullcluster/disaster/snapshots* directory.

Verify Snapshot

The operations described below allow the system administrator to verify the snapshot of the remote cluster site and the current configuration on the local site.

Verify Snapshot with ClusterAssistant

From the Disaster Recovery folder of the ClusterAssistant Application Manager, double–click the "**Site Recovery**" icon.

Figure 5.   ClusterAssistant Application Manager: Disaster Recovery folder
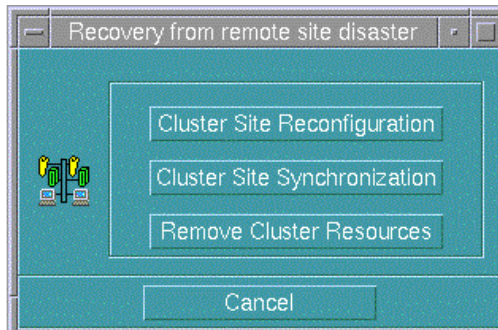
The following window is displayed:

Figure 6.   "Recovery from remote site disaster" window

Click on the "**Cluster Site Reconfiguration**" button. The following window is displayed:
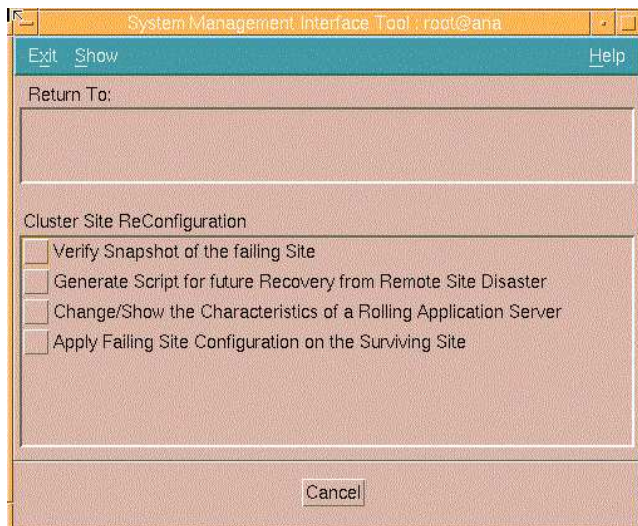

Figure 7.   smit menu: "Cluster Site Reconfiguration" window

Select the "**Verify Snapshot of the failing Site**" smit menu. The following screen is displayed. Select the name of the snapshot to be verified, from the list:
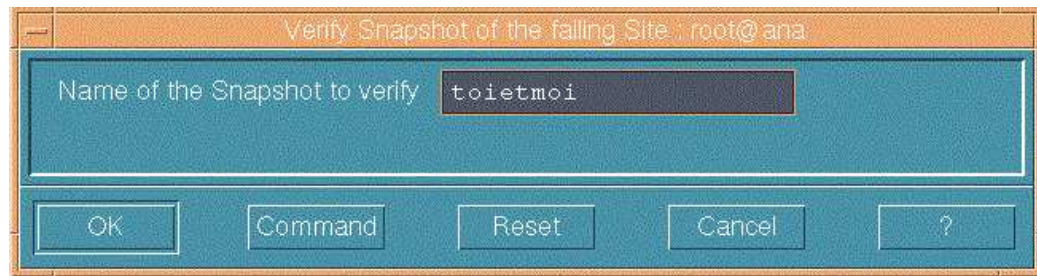

Figure 8.   Verify snapshot window

Click "**OK**" to validate.

Verify Snapshot with smit

Enter the smit command:

```
smit disaster
```

The following smit screen is displayed:

Figure 9.   smit menu: disaster window

Select the "**Recovery from Remote Site Disaster**" smit menu.

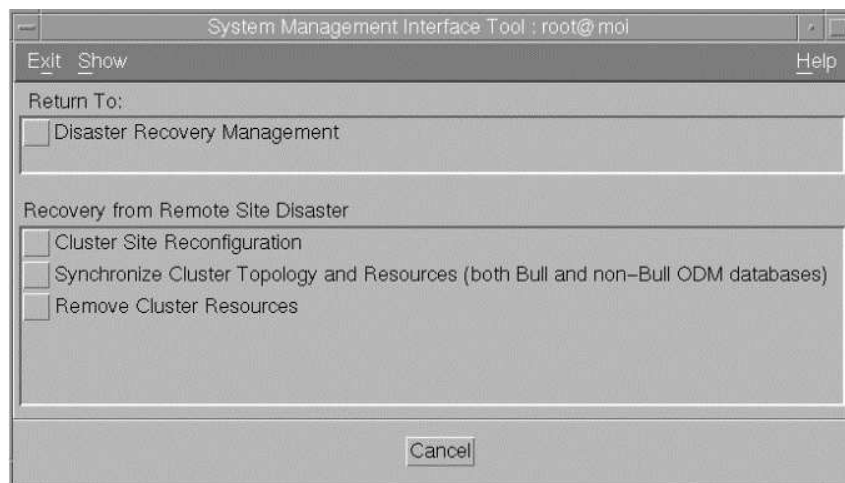The following smit screen is displayed:

Figure 10. smit menu: Recovery from Remote Site Disaster

Select "**Cluster Site Reconfiguration**" smit menu:
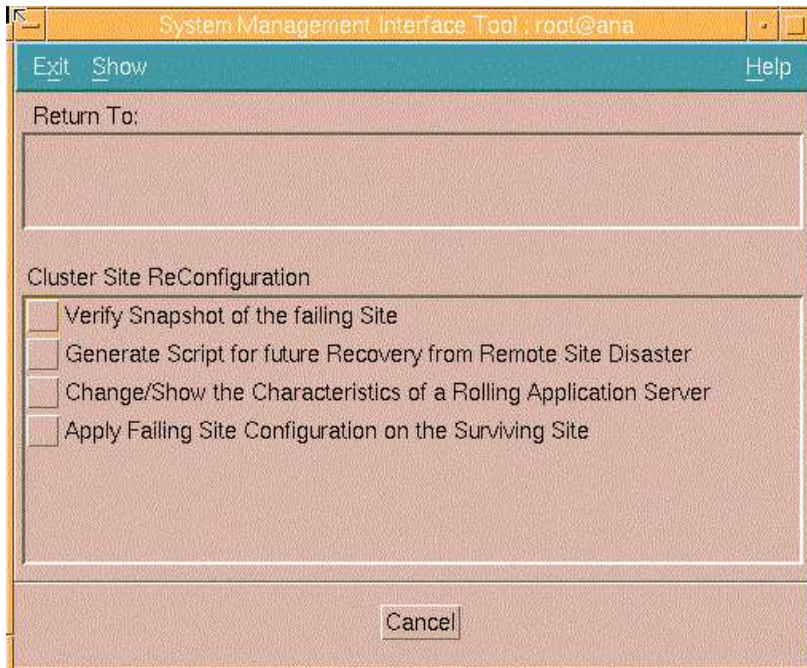
The following smit screen is displayed:

Figure 11. smit menu: "Cluster Site Reconfiguration" window

Selecting the "**Verify Snapshot of the failing Site**" smit menu. The following menu is displayed. Select from the list the name of the snapshot to be verified:
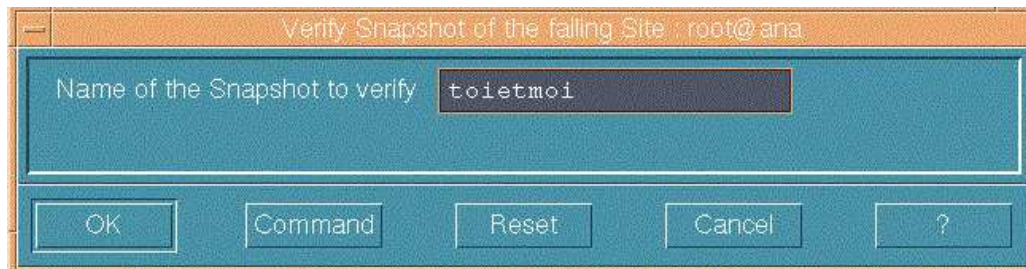

Figure 12. Verify snapshot window

Click "**OK**" to validate.

Validate AIX Mirroring (if needed)

To mirror storage, if you do not use the SRDF mechanism, the following configuration is mandatory to validate the AIX mirroring facility.

Using AIX Mirroring with ClusterAssistant

From the "**Disaster Recovery**" folder of the ClusterAssistant main folder (see figure 5), double–click on the "**Campus Disaster Recovery**" button.
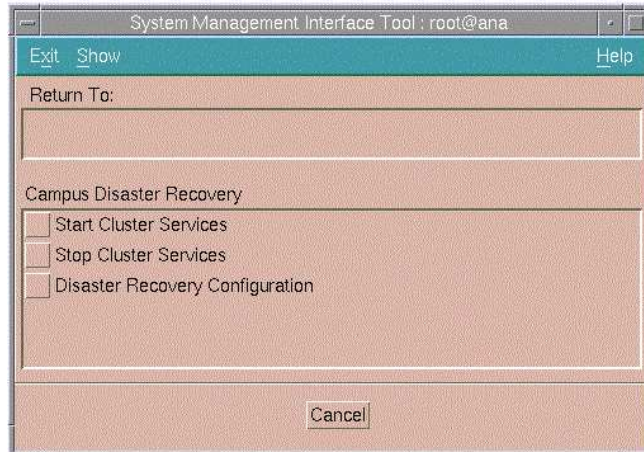
The following screen is displayed:

Figure 13. ClusterAssistant menu: Campus Disaster Recovery

Select the "**Disaster Recovery Configuration**" smit menu, and then activate the AIX mirroring facility:

Figure 14. ClusterAssistant menu: Activate Disaster Recovery

Click "**OK**" to validate.

Using AIX Mirroring with smit

Enter the smit command:

```
smit disaster
```

The following smit screen is displayed:

Figure 15. smit menu: Disaster Recovery Managenent window

Select the '**Campus Disaster Recovery**' smit menu. The following smit screen is displayed:
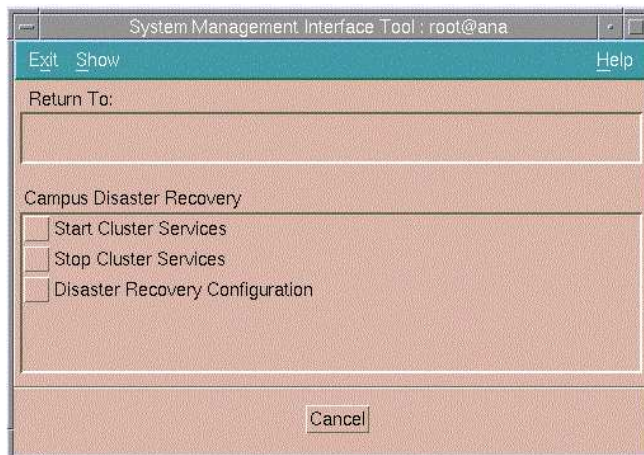

Figure 16. smit menu: Campus Disaster Recovery

Select the "**Disaster Recovery Configuration**" smit menu, and then activate the AIX mirroring facility.
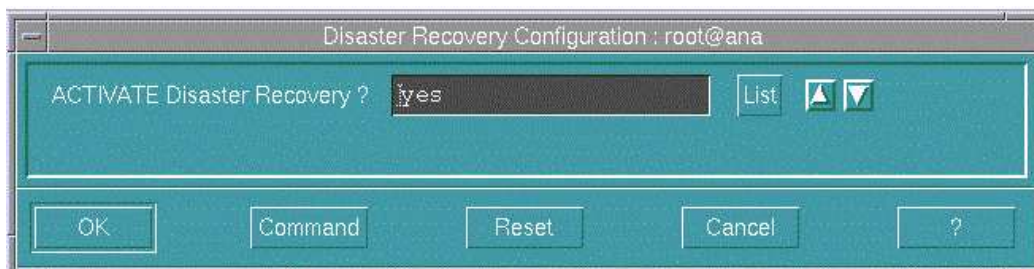

Figure 17. smit: Activate Disaster Recovery

Click "**OK**" to validate.

Generate Reconfiguration Script

This section explains how to generate the reconfiguration script to process the snapshot of the failed cluster and to add or modify HACMP resource group(s) on the surviving cluster.

According to the given list of dead node names on the failed Site (retrieved from the snapshot), the system administrator must enter the list of running node names in the right order, as described in the following example:

Example:

Two–site configuration with two nodes on each site:

• Site A

Node_1 (with all of its owned resources).

Node_2 (with all of its owned resources).

• Site B

Node_3 (with all of its owned resources).

Node_4 (with all of its owned resources).

If the given list of dead node names on the failed site (Site A) is:

Node_1   Node_2

The system administrator must indicate:

• Node name by which all the resources belonging to the Node_1 must have to be taken over.

• Node name by which all the resources belonging to the Node_2 must have to be taken over.

If the system administrator indicates:

Node_3   Node_4

Then

– All the resources previously belonging to the failed cluster node: Site A Node_1 will be reconfigured on the cluster node of the surviving site: Site B Node_3.

– All the resources previously belonging to the failed cluster node: Site A Node_2 will be reconfigured on the cluster node of the surviving site: Site B Node_4.
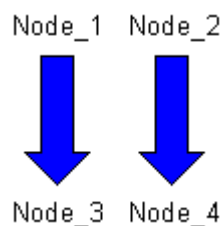

Figure 18. Reconfiguration example

To generate the configuration script of the failed site, proceed as follows:

From the "**Cluster Site Reconfiguration**" window (see figure 11). Click on "**Generate Script for future Recovery from Remote Site Disaster**" button.

On the new screen opened, a dialog box prompts you to:

1. Select, from the list, the name of the snapshot.

2. Enter a list of running node names in the right order, corresponding to the list of dead node names on the failed site.

3. Enter or select from the list the name of the script to generate.

In the example below the script "genscript" is generated with the following rules:

- All the resources previously belonging to the node *moi* would be processed to be configured later on the node *ana*, using the "**Apply Failing Site Configuration on the Surviving Site**" menu (see page 5-1).

- All the resources previously belonging to the node *toi* would be processed to be configured later on the node *purna*, using the "**Apply Failing Site Configuration on the Surviving Site**" menu (see page 5-1).
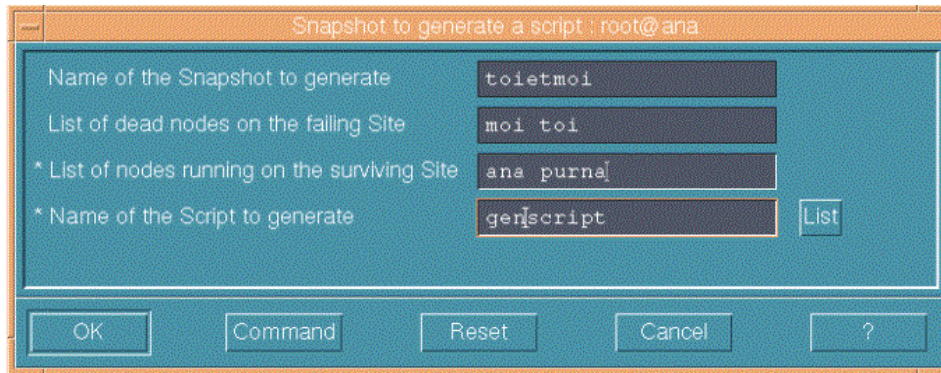

Figure 19. Generate a script

Importing Volume Groups

All the volume groups of each cluster site have to be imported on each node of the other cluster site.

To have more information to import volume groups using the shared SRDF volumes, refer to the Appendix A.

Chapter 5. Recovery from Disaster Situation

When the system administrator decides to proceed with the site takeover (e.g. recover the failing site resources by the surviving site), he has to apply the generated script and to synchronize the cluster site topology.

Note: with **SRDF facility** only, before applying the failing site configuration on the surviving site, the system administrator must:

- Switch data processing from source (R1) side to target (R2) side, entering the following line command:

```
symrdf -g DgName failover
```
       `DgName` = Device Group Name

- After the successful completion of this command, the system administrator can apply the generated script.

If no script has been generated from the snapshot, you have to generate a reconfiguration script before to apply it.

For more information, refer to section "Generate Reconfiguration Script" page 4-7.

The **Recovery from Disaster Situation** is realized performing the following steps:

- **Step 1**: apply the Failing Site Configuration on the Surviving Site

- **Step 2**: Synchronize the Cluster Site Topology

Step 1. Apply the Failing Site Configuration on the Surviving Site

The reconfiguration script must be applied on the surviving cluster to take into account:

- all the resources previously belonging to the failed cluster.

- all the resources owned by the surviving cluster.

From the "**Disaster Recovery**" folder of the ClusterAssistant Application Manager, double–click the "**Site Recovery**" icon (see figure 5 page 4-2).

The "**Recovery from Remote Site Disaster**" window is displayed.

Click on the "**Cluster Site Reconfiguration**" button (see figure 6 page 4-2).

Select "**Apply Failing Site Configuration on the Surviving Site**" from the new window displayed.

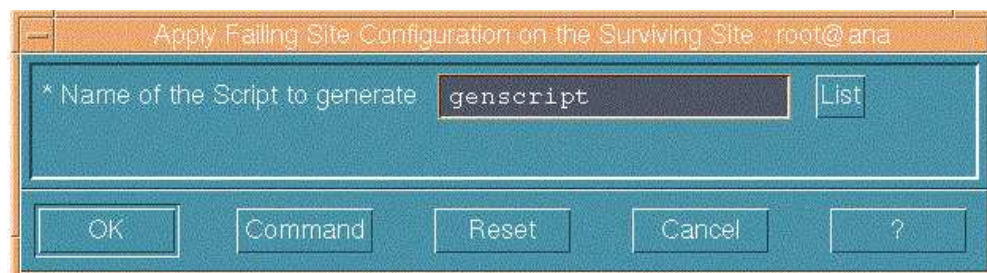

Figure 20. Apply failing site configuration on the surviving site window

Enter the name of the script to be applied or select it using the "**list**" button.

Click **OK** to validate.

To restart the activities of the failing site on the surviving site (dynamic reconfiguration), the surviving cluster must now be synchronized. Go to Step 2.

Step 2. Synchronize the Cluster Site Topology

To synchronize the cluster site topology, you can use either the ClusterAssistant Application or smit.

- Using the ClusterAssistant Application:

  From the Disaster Recovery folder of the ClusterAssistant Application Manager, double–click the "**Site Recovery**" icon (see figure 5 page 4-2).

  The "**Recovery from Remote Site Disaster**" window is displayed.

  Click on the "**Cluster Site Synchronization**" button (see figure 6 page 4-2).

- Using smit:

  Enter smit command:

  ```
  smit disaster
  ```

  The "**Disaster Recovery Management**" screen opens (see figure 9 page 4-3).

  Select the "**Recovery from Remote Site Disaster**" smit menu.

  Click on the "**Synchronize Cluster Topology and Resources**" button.

A dialog box appears, allowing the administrator to set the parameters:

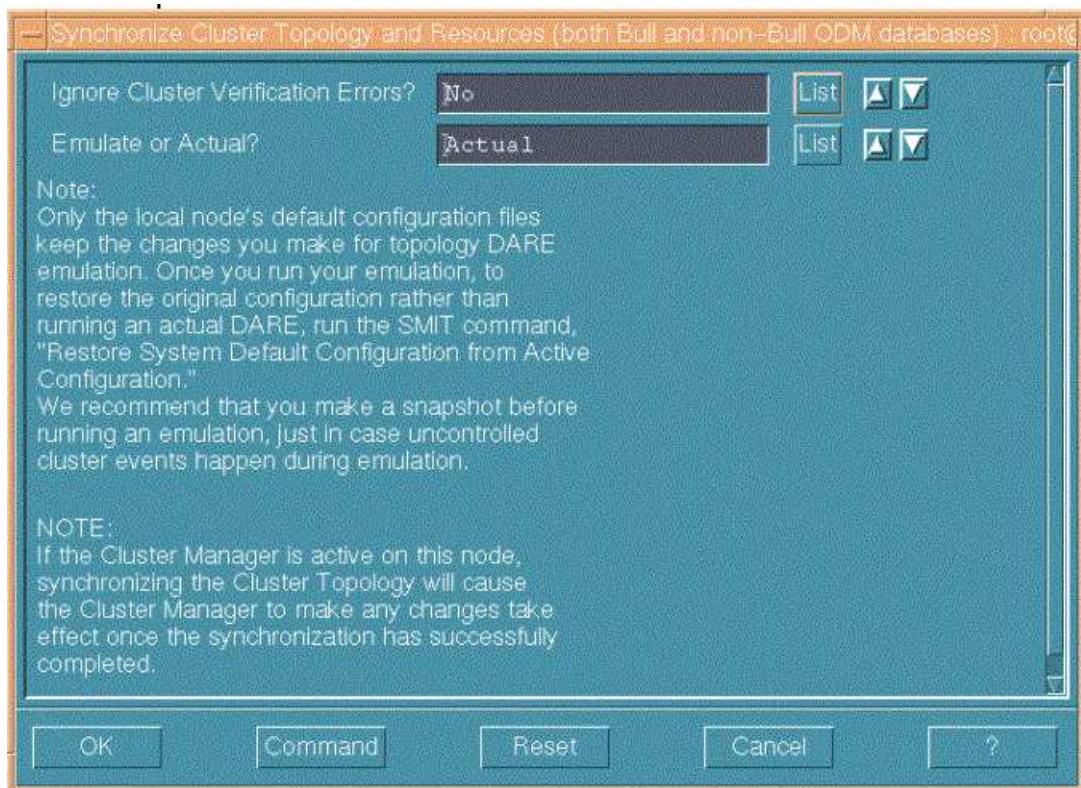- Ignore the cluster verification errors.
- Use the emulate mode.

Figure 21. Synchronize cluster topology window

Click **OK** to validate.

Chapter 6. Restore Initial Configuration

This service is used after a disaster to restore the initial configuration of the Cluster Site.

This operation consists, from the survival site, in removing the cluster resources of the failing site.

The restoration is realized performing the following steps:

- **Step 1**: Restore the Iniatial Configuration on the Surviving Site
- **Step 2**: Synchronize the Cluster Site Topology

Step 1. Restore the Initial Configuration (using the SRDF or the AIX Mirroring Facilities)

1. Restore the Initial Site Configuration of the surviving cluster site.

   From the Disaster Recovery folder of the ClusterAssistant Application Manager, double–click the "**Site Recovery**" icon (see figure 5 page 4-2).

   From the new screen opened, click on the "**Remove Cluster Resources**" button (see figure 6 page 4-2).
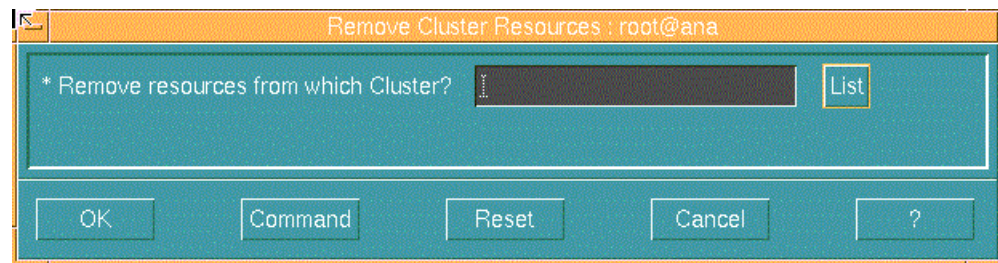
   Select the cluster.

Figure 22. Remove cluster resources window

2. To restart the activities of the surviving site (dynamic reconfiguration), the surviving cluster must now be synchronized. Go to Step 2.

Step 2. Synchronize the Cluster Site Topology

To synchronize the cluster site topology, you can use either the ClusterAssistant Application or smit.

- Using the ClusterAssistant Application:

  From the Disaster Recovery folder of the ClusterAssistant Application Manager, double–click the "**Site Recovery**" icon.

  From the new window opened, click on the "**Cluster Site Synchronization**" button.

- Using smit:

  Enter smit command:

  ```
  smit disaster
  ```

From the new screen opened, click on the "**Synchronize Cluster Topology and Resources**" button.

Then a dialog box appears, allowing the administrator to:

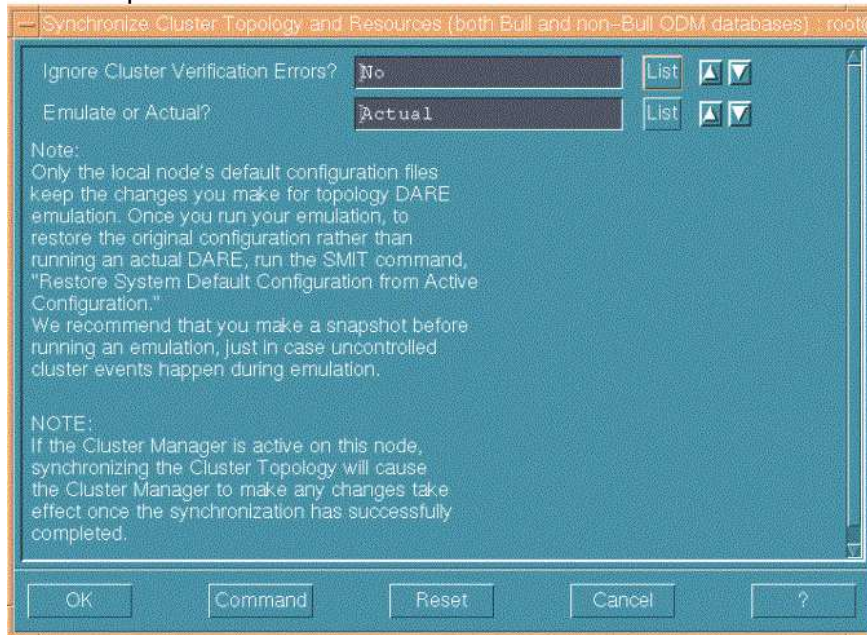- Ignore the cluster verification errors.
- Use the emulate mode.

Figure 23. Synchronize cluster topology window

If you use the **AIX mirroring facility**, proceed as follows:

On the surviving cluster site, execute the following script to synchronize all the disks of each cluster site:

`/usr/sbin/bullcluster/nativeha/nativeha_nodeup_rmtcmp_pre`

If you use the **SRDF facility**, proceed as follows:

1. To ensure that the source volumes are fully synchronized, and that the takeover is accomplished smoothly, the system administrator must start an update of the source (R1) side after a failover and while the target (R2) side may still be operational to its local host(s). This command makes sure that the source and target volume is within a pre–specified level of synchronization. This operation implies that write operations with the specified target (R2) volumes will be stopped.

   This command can be run only after a target takeover and prior to a source takeover.

   On the surviving cluster, enter the following command:

   `Symrdf –g DgName update`

2. To switch data processing from target (R2) side to source (R1) side with source data re–synchronization.

   If the target (R2) is operational, write disables the device(s) on the target side to its/their local host(s), and resumes I/O traffic on the RDF links. Then it writes enables the device(s) on the source (R1) side to its/their local host(s).

   On the surviving cluster, enter the following command: Symrdf –g DgName failback

3. On the repaired site, the system administrator have to initiate an establish on the target (R2) side to refresh R2 (target) data from R1 (source) data:

   `Symrdf –g DgName establish`

4. The system administrator can restart HACMP on the repaired site.

Appendix A. Importing Volume Groups Using Shared Volumes

This appendix is a help to check the configuration of the shared SRDF volumes. For more information, refer to the EMC documentation.

1. **Check if Symmetrix Device groups exists**, executing the following command:

        symdg list

   If the answer is " No Symmetrix device groups were found ", create Symmetrix Device groups. Go to to step 2: "Create the Symmetrix Device Groups".

   Otherwise go to step 8: "Check the state of the Symmetrix Device Groups".

2. **Create the Symmetrix Device Groups**

   On the local host:

   – Create a Symmetrix device group of type RDF1:

        symdg create DgName –type RDF1

        DgName: Device Group name

   – Add a physical device belonging to the specified Symmetrix device group:

        symld –g DgName add pd PdevName

        PdevName: Physical Device Name

   On the remote host:

   – Create a Symmetrix device group of type RDF2:

        symdg create DgName –type RDF2

   – Add a physical device belonging to the specified Symmetrix device group:

        symld –g DgName add pd PdevName

   Note: If you use EMC PowerPath devices with RDF devices, you can also use the **symcg** command to perform operations on a Symmetrix RDF consistency group (for more information, refer to the EMC documentation).

3. **Add Gatekeepers** (optional)

   Define Gatekeepers:

        symgate define pd PdevName

   Associate the Gatekeepers to the Symmetrix Device Group:

        symgate –g DgName associate pd PdevName

4. **Check Information about Device Group**

   ·   Symdg show DgName

5. **Suspend I/O traffic on the RDF links**

   On the local host, stop mirroring between source (R1) and target (R2) devices:

        symrdf –g DgName split

6. **Imports AIX Volume Group on the remote host**

   Remove Symmetrix devices:

        rmdev –l PdevName

Make Symmetrix devices available :

```
mkdev -l PdevName
```

Import Volume Groups:

```
importvg -y VgName [-v MajorNumber] PdevName
```

Change Volume group characteristics (auto–varyon)

```
chvg -an -Qn VgName
```

Varyoff Volume Group

```
varyoffvg VgName
```

7. **Synchronize the SRDF mirrored pairs**

   Initiate a data copy from source (R1) side to target (R2) side.

   ```
   symrdf -g DgName establish [-full]
   ```

8. **Check the state of the Symmetrix Device Groups**

   ```
   symrdf -g DgName query
   ```

   If the State is "Split" (Failed over), switch data processing from the source (R1) side to the target (R2) side:

   ```
   symrdf -g DgName failover
   ```

   After the successful completion of this previous command, apply the generated script (refer to "Applying Failing Site Configuration on the Surviving Site", on page 5-1).

   If the State is "Failed Back", switch data processing from the target (R2) side to the source (R1) side:

   ```
   symrdf -g DgName failback
   ```

   If the State is "SyncInProg", the synchronization is in progress, so wait until the state becomes "Synchronized".

9. **Verify the state of the RDF device pairs in a device group**:

   ```
   symrdf  -g DgName verify
   ```

   The State of the SRDF mirrored pairs would be synchronized.

Glossary

This glossary contains abbreviations, key–words and phrases that can be found in this documentation.

AIX
Advanced Interactive Executive

API
Application Programming Interface

ATF
Application Transparent Failover

CDE
Common Desktop Environment

DAS
Disk Array Subsystem

EPC
Escala PowerCluster

GUI
Graphical User Interface

HA
High Availability/ Highly Available

HACMP
High Availability Cluster Multi–Processing

HACMP/ES
High Availability Cluster Multi–Processing
Enhanced Scalability

HAS
The High Availability Solution Project

LPP
Licensed Program Product

ODM
Object Data Manager

PTF
Program Temporary Fix

RDF
Remote Data Facility

SMIT
System Management Interface Tool

SRDF
Symmetrix Remote Data Facility

SSA
Serial Storage Architecture:

INDEX

Vos remarques sur ce document / Technical publication remark form

Titre / **Title :**   Bull   Recovery from Cluster Site Disaster User's Guide

Nº Reférence / **Reference Nº :**   86 A2 86JX 00

Daté / **Dated :**   January 2002

ERREURS DETECTEES / ERRORS IN PUBLICATION

AMELIORATIONS SUGGEREES / SUGGESTIONS FOR IMPROVEMENT TO PUBLICATION

Vos remarques et suggestions seront examinées attentivement.
Si vous désirez une réponse écrite, veuillez indiquer ci-après votre adresse postale complète.

Your comments will be promptly investigated by qualified technical personnel and action will be taken as required.
If you require a written reply, please furnish your complete mailing address below.

NOM / NAME : _____   Date : _____

SOCIETE / COMPANY : _____

ADRESSE / ADDRESS : _____

Remettez cet imprimé à un responsable BULL ou envoyez-le directement à :

Please give this technical publication remark form to your BULL representative or mail to:

BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

Technical Publications Ordering Form
Bon de Commande de Documents Techniques

To order additional publications, please fill up a copy of this form and send it via mail to:
Pour commander des documents techniques, remplissez une copie de ce formulaire et envoyez-la à :

BULL CEDOC
ATTN / **Mr. L. CHERUBIN**
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

Phone / Téléphone :   +33 (0) 2 41 73 63 96
FAX / Télécopie   +33 (0) 2 41 73 60 19
E–Mail / Courrier Electronique :   srv.Cedoc@franp.bull.fr

Or visit our web sites at: / Ou visitez nos sites web à:
http://www.logistics.bull.net/cedoc
`http://www-frec.bull.com`   `http://www.bull.com`

CEDOC Reference # Nº Référence CEDOC	Qty Qté	CEDOC Reference # Nº Référence CEDOC	Qty Qté	CEDOC Reference # Nº Référence CEDOC	Qty Qté
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	

[_ _] :   **no revision number means latest revision** / pas de numéro de révision signifie révision la plus récente

NOM / NAME : _____   Date : _____

SOCIETE / COMPANY : _____

ADRESSE / ADDRESS : _____

_____

PHONE / TELEPHONE : _____   FAX : _____

E–MAIL : _____

For Bull Subsidiaries / Pour les Filiales Bull :

Identification: _____

For Bull Affiliated Customers / Pour les Clients Affiliés Bull :

Customer Code / Code Client : _____

For Bull Internal Customers / Pour les Clients Internes Bull :

Budgetary Section / Section Budgétaire : _____

For Others / Pour les Autres :

Please ask your Bull representative. / Merci de demander à votre contact Bull.

Bull

BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

ORDER REFERENCE
86 A2 86JX 00

Bull

AIX

Recovery from
Cluster Site
Disaster
User's Guide

86 A2 86JX 00

AIX

Recovery from
Cluster Site
Disaster
User's Guide

86 A2 86JX 00

AIX

Recovery from
Cluster Site
Disaster
User's Guide

86 A2 86JX 00