

**EMC ControlCenter
5.2 Service Pack 4**

ADMINISTRATION/USER GUIDE

**P/N 300-002-591
REV A01**

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748 -9103
1-508 -435 -1000
www.emc.com

Copyright © 2001 - 2006 EMC Corporation. All rights reserved.

Printed January 2006

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

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

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

Trademark Information

EMC², EMC, EMC ControlCenter, AlphaStor, ApplicationXtender, Catalog Solution, Celerra, CentraStar, CLARAlert, CLARiiON, ClientPak, Connectrix, Co-StandbyServer, Dantz, Direct Matrix Architecture, DiskXtender, Documentum, EmailXtender, EmailXtract, HighRoad, Legato, Legato NetWorker, Navisphere, OpenScale, PowerPath, RepliStor, ResourcePak, Retrospect, Smarts, SnapShotServer, SnapView /IP, SRDF, Symmetrix, TimeFinder, VisualSAN, VSAM Assist, Xtender, Xtender Solutions, and where information lives are registered trademarks and EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, EMC Storage Administrator, Access Logix, ArchiveXtender, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CLARevent, Codebook Correlation Technology, Common Information Model, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, DiskXtender 2000, EDM, E-Lab, EmailXaminer, Enginuity, eRoom, FarPoint, FLARE, FullTime, Global File Virtualization, Graphic Visualization, InfoMover, Invista, MirrorView, NetWin, NetWorker, OnAlert, Powerlink, PowerSnap, Rainfinity, RecoverPoint, RepliCare, SafeLine, SAN Advisor, SAN Copy, SAN Manager, SDMS, SnapImage, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, UltraPoint, Viewlets, VisualSRM, and WebXtender are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

Preface	xix
----------------------	------------

PART 1 Administering ControlCenter

Chapter 1 Managing ControlCenter Users

Introduction to ControlCenter User Access Management	1-2
Managing User Access to ControlCenter.....	1-2
Working with User Groups.....	1-3
Working with Rules	1-4
Assigning Permissions	1-6
Understanding Groups and Inheritance	1-10
Understanding the ChangeMembership Permission.....	1-11
User Access Management Procedures.....	1-12
Creating ControlCenter Users	1-13
Creating a Domain Account for the ECC Server Service	1-15
Adding a User to ControlCenter	1-16
Managing ControlCenter User Groups	1-17
Creating a New User Group	1-17
Adding a User to a Group.....	1-18
Removing the eccadmin User from ControlCenter	1-19
Creating and Modifying Rules for Users and User Groups.....	1-20
Monitoring and Contacting Console Users	1-22
Monitoring Console Users	1-22
Sending Messages to Console Users.....	1-22

Chapter 2 Administering ControlCenter Data Collection Policies

Data Collection Overview	2-2
Managing Data Collection Policy Definitions and Templates	2-4
Agent Data Collection Policy Reference	2-7

Chapter 3 **Configuring and Managing Alerts and Notifications**

Understanding Alerts and Notifications	3-2
Alert and Notification Life Cycles	3-4
Understanding Metric Types	3-6
Setting Up Your Alert and Notification Strategy	3-7
Controlling Who Creates and Edits Alerts and Notifications	3-7
Deciding Whether to Define Alerts or Notifications	3-7
Defining the Severity Levels	3-8
Establishing Procedures for the Resolution of Alerts	3-9
Refining Alert Definitions	3-10
Keeping Alerts and Notifications	3-10
Creating Alerts and Notifications	3-11
Tips for Setting Alerts and Notifications	3-11
Getting Help	3-11
Gathering Information	3-12
Creating Alert Definitions	3-13
Creating an Alert Definition From a Template	3-13
Creating Alert Definitions in the Edit Thresholds Dialog Box ..	3-14
Copying an Alert Definition	3-14
Testing an Alert Definition	3-14
Controlling Alert Spikes	3-16
Sending Alerts in E-mail or to a Management Framework	3-17
Automating Alert Responses with Autofixes	3-18
Creating an Autofix Definition in the Console	3-18
Passing Alert Information to an Autofix Script	3-19
Creating an Autofix Script on the Host	3-21
Attaching an Autofix to an Alert Definition	3-22
Best Practices for Configuring and Managing Alerts	3-23
Disable Unnecessary or Redundant Alerts	3-23
Set Alert Frequencies to Minimize Processing Impact	3-23
Use Notifications to Reduce Alert Volume	3-24
Create User-Defined Groups to Organize Your Alerts	3-24
Use Management Policies to Notify Appropriate Personnel	3-24
Modify Templates to Facilitate Alert Creation	3-24
Troubleshooting Alerts and Autofixes	3-25
Too Many Alerts Appear in the Console	3-25
Cannot Create or Edit Alerts or Changes Not Saved	3-26
Cannot Clear Alerts	3-26
Alert Does Not Trigger as Expected	3-26
Autofix Does Not Run	3-27
Alert Count Differs Among Users	3-27
Managed Object has Warning Icon but No Alerts	3-28
Alert Created or Modified Date and Time is Incorrect	3-28

Chapter 4 **Maintaining the Repository**

Automatic Tasks	4-2
Backing Up the Repository	4-2
Backing Up Using Third-Party Tools	4-3
Exporting the Repository Backup	4-4
Analyzing Tables	4-4
Rebuilding the Index	4-4
Recompiling Invalid Objects	4-5
Monitoring Tablespace Growth	4-5
Monitoring the Status of Automated Tasks	4-6
Listing Installed ControlCenter Components	4-6
Manual Tasks	4-7
Shutting Down the Repository	4-7
Starting the Repository	4-8
Scanning the Repository Alert Log	4-8
Cleaning Trace Files	4-8
Determining Tablespace Fragmentation	4-8
Determining Which Processes are Currently Running	4-9
Resetting the Repository	4-9
Performing a Media Recovery	4-9
Restoring the Repository	4-10
Importing the Repository Database	4-11
Gathering Data for Remote Diagnostics Assistance	4-11

PART 2 **Using ControlCenter**

Chapter 5 **Using the ControlCenter Console**

Working in the Console	5-2
Using the Console	5-3
Using the Menu Bar	5-4
Using the Taskbar	5-5
Using the Console Toolbar	5-6
Understanding the Information Panel	5-7
Using the Tree Panel	5-8
Using the Target Panel	5-10
Using the Active View	5-11
Using a Table View	5-12
Using Map Views	5-13
Using Special Views	5-14
Splitting a View	5-15
Using the Action Menu	5-15

Console Features	5-16
Creating User-Defined Groups	5-16
Using the Arrange By Feature	5-17
Sorting Multiple Columns of Data	5-17
Using the Drag and Drop Feature	5-18
Using the Drill-Down Feature	5-19
Using the Find Feature	5-20
Using the Hide and Show Columns Feature	5-20
View Preferences	5-21
ControlCenter Online Help	5-23

Chapter 6 Using the EMC ControlCenter Web Console

Working in the Web Console	6-2
Accessing the Web Console	6-2
Web Console Interface	6-3
Menu Bar	6-4
Tree View	6-5
View Bar	6-6
Web Console Views	6-7
Populating Views	6-7
Navigating views	6-7
Common View Actions	6-7
Types of Views	6-8
Web Console Tutorial and Online Help	6-15
Web Console Tutorial	6-15
Web Console Help	6-15
Web Console FAQs	6-16

Chapter 7 Managing Your SAN

EMC ControlCenter SAN Manager Overview	7-2
Managing a SAN	7-2
Features	7-3
Discovery and Monitoring Requirements	7-4
Discovery	7-5
Automatic Discovery	7-5
Connectivity Device Discovery	7-6
Assisted Discovery	7-8
Topology View	7-13
Objects Rendered in Topology View	7-14
ControlCenter Groups in Topology View	7-15

User-Defined Groups in Topology View	7-16
Highlighting Zone and Zone Set Members in Topology View	7-17
Topology View Tools	7-17
Saving Topology Maps (View Preferences).....	7-18
Topology Edit Service (TES).....	7-19
TES and Discovery	7-19
User-Defined Objects	7-19
Associating Unidentified Ports.....	7-20
Viewing the Login History	7-21
Zoning	7-22
Managing Zone Sets.....	7-30
Managing Zones	7-34
Zoning States.....	7-35
EMC Zoning Recommendations	7-37
Monitoring Statistics.....	7-38
Masking.....	7-39
Masking View.....	7-39
Supported Arrays	7-39
Enabling Masking.....	7-40
Masking Actions	7-40
Required Agents	7-40
Symmetrix Masking	7-41
HP StorageWorks HSG80 Masking	7-43
CLARiiON Masking.....	7-44
HP StorageWorks XP 128/1024, HDS 9500V and 9900V Series, SUN StoreEdge 9900V Series Masking.....	7-45
HP StorageWorks XP 48/512 Masking	7-48
Path Details View.....	7-50
Objects to Place Into Path Details View.....	7-50
Table and Graphic Panes	7-51
Viewing Path Details.....	7-51
Troubleshooting Paths	7-52

Chapter 8 **Monitoring Storage with Alerts and Notifications**

Getting the Status of the Storage Environment.....	8-2
Using the At A Glance View	8-3
Using the Alerts View	8-5
Identifying Hosts, Arrays, and Network Components	
Requiring Attention.....	8-8
Responding to Alerts.....	8-9
Getting Help on Triggered Alerts.....	8-10
Getting More Information About an Affected Resource	8-11
Checking the Status of Automatic (Autofix) Responses.....	8-12

Tracking the Progress of Alert Resolution with Notes.....	8-13
Searching Alert Notes.....	8-13
Finding the Alert You Need.....	8-14
Gathering Information for Setting Alerts	8-15

Chapter 9 Monitoring and Analyzing Performance

Performance Monitoring and Analysis Overview	9-2
Performance Monitoring.....	9-2
Performance Analysis.....	9-2
Performance Monitoring Configuration and Startup	9-3
Required Components.....	9-3
Setting Data Collection Policies	9-3
Starting Performance View	9-4
Performance Analysis Configuration and Startup	9-5
Performance Analysis Architecture.....	9-5
Enabling and Editing WLA Policies.....	9-9
Archiving Performance Data.....	9-9
Viewing the Performance Archives and Collections.....	9-13
Accessing the Performance Archives and Collections.....	9-13
Daily and Revolving Performance Analysis	9-14
Host Configuration Information.....	9-15
Symmetrix Array Information	9-20

Chapter 10 Allocating or Deallocating Storage

Storage Provisioning Service Overview.....	10-2
Allocating Storage Overview.....	10-3
Assigning Allocation and Deallocation Permissions	10-4
Gathering Data Using Storage Array Properties	10-5
Gathering Data Using the Free Space View.....	10-6
Displaying Free Space View	10-6
Creating Storage Pools.....	10-8
What You Should Know before Starting.....	10-8
Creating a Storage Pool	10-9
Adding a Device to a Storage Pool.....	10-9
Moving a Logical Device between Pools	10-9
Creating Storage Allocation Policies	10-10
Allocating Storage Using Storage Provisioning Service	10-14
Deallocating Storage Using Storage Provisioning Service	10-16
Deallocation Process Overview.....	10-16
Understanding Deallocation Path Selection.....	10-17
Deallocation Permissions	10-18
Controlling Deallocation Actions through Policies	10-18

Starting the Deallocation Wizard	10-21
Troubleshooting Deallocation	10-22
Verifying Deallocation Actions	10-23
Deferring SPS Tasks for Future Execution	10-24

Chapter 11 Protecting Data

Introduction to Data Protection	11-2
Protecting Data on Symmetrix Storage Arrays	11-3
Configuring Symmetrix Device Mirrors	11-4
Working with Symmetrix Groups	11-5
Local Protection with Symmetrix TimeFinder	11-10
TimeFinder Clones	11-13
EMC Snap	11-14
Remote Protection with Symmetrix SRDF	11-16
Working with EMC Solutions Enabler SYMCLI	11-25
Creating BCV Devices	11-26
Protecting Data on CLARiiON Storage Arrays	11-33
Protecting Data on HP StorageWorks Storage Arrays	11-35
Enabling or Disabling Host Access to Units	11-35
Setting Unit Offset for a Connection	11-36

Chapter 12 Managing Host Storage Resources

Viewing File Systems, Devices, and Their Relationships	12-2
Viewing File Systems and Their Properties	12-3
Viewing Host Devices and Their Properties	12-3
Relating Host Resources to Storage Array Volumes	12-4
Working with Windows and UNIX Files and Directories	12-6
Recapturing UNIX Storage	12-7
Recapturing Windows Storage	12-8
Recapturing Novell Storage	12-9
Managing UNIX Storage to Increase Performance	12-10
Managing Windows Storage to Increase Performance	12-11
Working with MVS Host Resources	12-13
Viewing MVS Host Properties and Relationships	12-13
Viewing Detailed Host Information	12-13
Common Tasks in MVS Hosts	12-15
Recapturing MVS Host Storage	12-16
Managing MVS Storage to Increase Performance	12-18

Chapter 13 Using Reports

Overview of Reports	13-2
Types of Reports	13-2
User-Defined Groups for Reports	13-3
StorageScope Permissions	13-4
StorageScope Permission Types	13-4
Applying StorageScope Permissions	13-4
Working with Reports in the Console	13-6
Launching StorageScope from ControlCenter	13-6
Opening Reports from ControlCenter	13-7
Logging In to StorageScope	13-8

Chapter 14 Incorporating Data from External Sources into Reports

Custom Field Overview	14-2
Planning Overview	14-2
Configuring Custom Fields	14-4
Assigning User Permissions	14-4
Creating the Custom Fields XML Template	14-4
Rules for Creating the Custom Field XML	14-5
Sample Custom Field XML	14-7
Importing the XML File	14-8
Removing Stale Custom Field Names	14-9
Managing Custom Fields	14-11
Viewing Custom Field Properties	14-11
Setting or Modifying Custom Field Names and Values	14-12
Removing Name and Value Pairs	14-13
Using the Custom Field Bulk Load Utility	14-14
Importing Custom Fields Using the Bulk Load Utility	14-14
customFieldsUtil Examples	14-19

Chapter 15 Tuning Symmetrix Performance

Performance Management Overview	15-2
Understanding Optimizer	15-3
Optimizer Process	15-3
Optimizer Devices	15-4
Optimizer Capabilities and Limitations	15-4
Swapping Logical Devices	15-5
Understanding Optimizer Time Windows	15-6

Using Optimizer.....	15-8
Accessing Optimizer	15-8
Monitoring Optimizer Server Status	15-9
Setting Optimizer Parameters.....	15-9
Retrieving Optimizer Logs	15-11
Using the Quality of Service (TimeFinder/SRDF QoS) Tool.....	15-12
Performance Tuning BCV and SRDF Copy	15-12
QoS Performance Settings	15-12
Index	i-1

1-1	Rule Providing a User with SDR Permissions on a Symmetrix Array .	1-3
1-2	Initial Groups	1-4
1-3	Default Authorization Rules	1-5
1-4	New User Dialog Box	1-14
1-5	New User Definition	1-16
1-6	Creating a New User Group	1-17
1-7	Adding a User to a Group	1-18
1-8	Creating a New Rule	1-20
1-9	Naming the New Rule	1-21
2-1	Selecting a Policy to Modify	2-5
2-2	Policy Definition Dialog Box	2-6
3-1	Alerts View	3-2
3-2	At A Glance Health and Performance Views	3-3
3-3	Alert and Notification Life Cycle	3-4
5-1	The ControlCenter Console	5-2
5-2	Basic Console Steps	5-3
5-3	Menu Bar	5-4
5-4	Performing a Task from a Menu	5-5
5-5	The Taskbar	5-5
5-6	Storage Allocation Title Bar	5-5
5-7	Console Toolbar	5-6
5-8	Information Panel	5-7
5-9	Default Tree Folders	5-9
5-10	Expanding and Collapsing Tree Items	5-9
5-11	Navigating in the Target Panel	5-11
5-12	Using the Active View	5-12
5-13	Using the Table View	5-12
5-14	Using Map Views	5-13
5-15	Using Special Views	5-14
5-16	Using the Action Menu	5-15

5-17	Sorting Multiple Columns of Data	5-17
5-18	Drag and Drop Example	5-19
5-19	Use Saved Preferences	5-21
6-1	The Web Console	6-3
6-2	Web Console Default Folders	6-5
6-3	Properties View Table	6-8
6-4	Properties Views Tabbed Tables	6-10
6-5	Properties View Tree Tables	6-10
6-6	Properties View Split Table View	6-11
6-7	Topology View Map	6-12
6-8	Performance View Chart	6-13
7-1	Switched Fibre Channel Fabric in a SAN	7-22
7-2	Active Zoning of a Fabric in a SAN	7-23
7-3	Fibre Channel Port Relationships	7-25
7-4	Zoning Folders Displayed in the Tree Panel	7-30
8-1	At A Glance Host Capacity Chart	8-3
8-2	Viewing Status of Grouped Objects in At A Glance View	8-4
8-3	Positioning the Cursor Over and Clicking At A Glance Icons	8-5
8-4	All Alerts Button	8-5
8-5	The Alerts View, Showing All Active Alerts	8-6
8-6	Selecting the Alert Chart View	8-7
8-7	Alert Chart View	8-7
8-8	Managed Object Status in Tree and View	8-8
8-9	Managed Objects With Alerts	8-8
8-10	Getting Help on a Triggered Alert	8-10
8-11	Getting More Information About an Affected Resource	8-11
9-1	Performance Analysis Architecture	9-6
9-2	Archiving Process for Performance Archives	9-10
9-3	Host-to-Symmetrix Configuration	9-15
9-4	Host Device Map to Non-EMC Disks	9-16
9-5	Host-to-Symmetrix Performance	9-17
9-6	CPU Utilization	9-18
9-7	Host Device Response Times	9-19
9-8	Host Device Response Times Graph	9-20
9-9	Symmetrix Cache Management and Data Flow	9-20
9-10	System Write Pending Count and System Write Pending Limit	9-22
9-11	Host Port – % Utilization	9-23
9-12	All Host Directors – % Utilization	9-24
9-13	Disk Directors – % Utilization	9-25
9-14	All Disks – % Utilization	9-26
9-15	All Disks – % Utilization Ribbon Graph	9-27
9-16	% Utilization for Selected Metrics	9-28
9-17	Disks % Utilization Histogram	9-29

9-18	Disk % Utilization Histogram Drill Down	9-30
9-19	Symmetrix Devices I/Os Per Second	9-31
10-1	Symmetrix Arrays Properties Table	10-5
10-2	Free Space View	10-7
10-3	Storage Pool	10-9
10-4	Allocation Policy Editor Dialog Box	10-10
10-5	Storage Allocation Using the Storage Provisioning Service	10-15
10-6	Storage Deallocation Wizard	10-21
10-7	Execute Later Dialog Box	10-24
10-8	New Task Added to the Task Lists Folder	10-25
10-9	Task List Properties	10-25
11-1	Device Protection Definition Dialog Box	11-4
11-2	Device Group Wizard Dialog Box	11-5
11-3	Create Device Group Dialog Box	11-6
11-4	Associate Device Group Dialog Box	11-7
11-5	Associate BCV Devices Dialog Box	11-8
11-6	BCV Configuration	11-10
11-7	Basic SRDF Configuration	11-16
11-8	Mode Control Dialog Boxes for Device Pairs and Groups	11-20
11-9	BCV/DRV Definition Dialog Box	11-26
11-10	Device Properties	11-27
11-11	Split View Showing Device Properties and TimeFinder	11-28
11-12	SnapView Snapshot Copies	11-34
12-1	Viewing Properties of File Systems	12-3
12-2	Viewing Host Devices and Their Properties	12-4
12-3	Viewing Relationships (Map)	12-5
12-4	Viewing Relationships (Table)	12-5
12-5	Viewing Relationships for an MVS Host (Table and Map)	12-13
12-6	Viewing Host Information in an Agent Window	12-14
12-7	Mapping MVS Volumes to Symmetrix Logical Volumes	12-18
13-1	Types of StorageScope Reports	13-3
13-2	StorageScope Home Page	13-6
13-3	All Arrays Report Example	13-7
13-4	StorageScope Login Page	13-8
14-1	Exported Comma-Delimited File	14-16
14-2	Comma-Delimited File Ready for Import	14-18
15-1	Performance Time Window	15-7
15-2	Optimizer Dialog Box	15-8
15-3	Log Tab	15-11
15-4	Performance Controls	15-12

1-1	Tasks and Permissions	1-7
2-1	Backup Agent DCPs	2-7
2-2	Common Agent DCPs	2-7
2-3	Connectivity Agent DCPs	2-8
2-4	Common Mapping Agent DCPs	2-8
2-5	Database Agent for Oracle DCPs	2-9
2-6	Host Agents for AIX, HP-UX, Linux, and Solaris DCPs	2-9
2-7	Host Agent for Windows DCPs	2-11
2-8	Physical Agent for MVS DCPs	2-11
2-9	EMC Invista Agent	2-12
2-10	Storage Agent DCPs	2-12
3-1	Alert and Notification Life Cycle Stages	3-5
3-2	Defining Alert Severity Levels	3-8
3-3	Gathering Information for File Systems, Directories, and Files	3-12
3-4	Autofix Syntax Requirements and Examples	3-20
3-5	Reducing the Number of Alerts That Display	3-25
6-1	Menu Bar Options	6-4
7-1	Required Agents	7-4
7-2	Effect of Zone Set Actions on Planned Zone Sets Folder	7-33
7-3	Zoning States	7-36
8-1	Autofix Status Icons	8-12
8-2	Criteria for Searching Alert Notes	8-13
8-3	Windows File, Folder, and Volume Alerts	8-14
8-4	UNIX File, Directory, and File System Alerts	8-14
8-5	MVS Disk Alerts	8-14
8-6	Gathering Information for File Systems, Directories, and Files	8-15
9-1	Agent Requirements for Collecting Performance Manager Data	9-8
10-1	Allocation Policy Editor Control Descriptions	10-11
10-2	RAID Level Configuration Options	10-13
10-3	Effects of Deallocation on Managed Objects	10-17

10-4	Deallocation Policy Editor Controls	10-19
11-1	Supported TimeFinder Operations	11-13
11-2	Supported TimeFinder Clone Operations	11-14
11-3	Supported EMC Snap Operations	11-15
11-4	Supported SRDF Operations	11-18
11-5	Advanced SRDF Commands	11-19
11-6	Configuration Modes	11-20
11-7	TimeFinder View Description	11-29
11-8	SRDF View Descriptions	11-30
12-1	Methods of Viewing Host Information	12-2
12-2	Common Tasks on Windows and UNIX Hosts	12-6
12-3	Candidate Files for Space Recovery (UNIX)	12-7
12-4	Exploring Detailed Information about MVS Host Resources	12-14
12-5	Common Tasks on MVS Hosts	12-15
14-1	Relationship of ControlCenter Objects to Custom Field Elements	14-5

As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this guide may not be supported by all revisions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If a product does not function properly or does not function as described in this guide, please contact your EMC representative.

Audience

This guide is part of the EMC ControlCenter documentation set, and is intended for use by system and data storage administrators.

Readers of this guide are expected to be familiar with the following topics:

- ◆ Storage array operation
- ◆ Host operating systems
- ◆ Storage array networks

Detailed information about EMC products is available on the EMC Powerlink website at: [**http://powerlink.emc.com**](http://powerlink.emc.com).

Organization

This guide is organized into two parts as follows:

Part 1, Administering ControlCenter

Chapter 1, *Managing ControlCenter Users*, provides directions for configuring and managing user access to ControlCenter.

Chapter 2, *Administering ControlCenter Data Collection Policies*, provides common tasks involved in administering DCPs.

Chapter 3, *Configuring and Managing Alerts and Notifications*, provides an introduction to alerts and monitoring.

Chapter 4, *Maintaining the Repository*, details the common tasks for maintaining the relational database that holds the current and historical data of both the storage environment and ControlCenter.

Part 2, Using ControlCenter

Chapter 5, *Using the ControlCenter Console*, provides an overview of the ControlCenter Console as well as tips for using online Help.

Chapter 6, *Using the EMC ControlCenter Web Console*, provides an overview of the ControlCenter Web Console.

Chapter 7, *Managing Your SAN*, provides information for managing your SAN.

Chapter 8, *Monitoring Storage with Alerts and Notifications*, discusses how to monitor your storage environment and storage resources using alerts and thresholds.

Chapter 9, *Monitoring and Analyzing Performance*, tells you how to configure, monitor, and perform data analysis on real-time and historical data collections from storage array, switch, host, and database objects.

Chapter 10, *Allocating or Deallocating Storage*, covers the common tasks required to allocate storage in an environment managed with EMC ControlCenter.

Chapter 11, *Protecting Data*, provides procedures for some of the common data protection tasks involving business continuance volumes (BCVs) and remote device mirroring.

Chapter 12, *Managing Host Storage Resources*, provides an introduction to host storage resource management.

Chapter 13, *Using Reports*, provides an introduction to reports.

Chapter 14, *Incorporating Data from External Sources into Reports*, details the ControlCenter custom fields feature that allows you to specify attributes for the managed objects in your storage environment that can later be used in generated reports.

Chapter 15, *Tuning Symmetrix Performance*, tells you how to tune your Symmetrix arrays for optimum performance using Optimizer.

An *Index* is located at the back of this document.

An online glossary of commonly used ControlCenter terms is available through the Console online Help.

Related Documentation

ControlCenter documentation is available in the following locations:

- ◆ **Documentation /Help CD** provided with your EMC ControlCenter installation kit.
- ◆ **EMC Powerlink** (<http://powerlink.emc.com>)
- ◆ **ControlCenter Documentation Library** available from the ControlCenter Console's help menu after installation

Conventions Used in This Guide

EMC uses the following conventions for notes and caution notices.

A note presents information that is important, but not hazard-related.



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.

Where to Get Help

Obtain technical support by calling your local sales office.

For service, call:

United States: (800) 782-4362 (SVC-4EMC)

Canada: (800) 543-4782 (543-4SVC)

Worldwide: (508) 497-7901

and ask for Customer Support.

Sales and Customer Support Contacts

A list of EMC sales locations is available on the EMC home page at:

<http://www.emc.com/contact/>

Additional information about EMC products and services is available on the EMC Powerlink website at:

<http://powerlink.emc.com>

Your Comments

Your suggestions help us improve the accuracy, organization, and overall quality of the user publications. Please send a message to **techpub_comments@emc.com** with your opinions of this guide.

Administering ControlCenter

This section provides administration tasks for EMC ControlCenter Administrators and consists of the following:

- Chapter 1, *Managing ControlCenter Users*
- Chapter 2, *Administering ControlCenter Data Collection Policies*
- Chapter 3, *Configuring and Managing Alerts and Notifications*
- Chapter 4, *Maintaining the Repository*

Managing ControlCenter Users

This chapter provides an introduction to ControlCenter access management as well as procedures for setting up and managing user access to the ControlCenter application.

This chapter contains the following sections:

- ◆ Introduction to ControlCenter User Access Management..... 1-2
- ◆ User Access Management Procedures 1-12
- ◆ Creating ControlCenter Users 1-13
- ◆ Adding a User to ControlCenter..... 1-16
- ◆ Managing ControlCenter User Groups 1-17
- ◆ Removing the eccadmin User from ControlCenter 1-19
- ◆ Creating and Modifying Rules for Users and User Groups 1-20
- ◆ Monitoring and Contacting Console Users..... 1-22

Introduction to ControlCenter User Access Management

This section provides an overview of the concepts behind EMC ControlCenter® access management. Refer to *User Access Management Procedures* on page 1-12 for the procedures to administer user access.

Detailed descriptions of ControlCenter user management procedures are provided in the online Help under **Administering ControlCenter Users**.

This section contains the following topics:

- ◆ *Managing User Access to ControlCenter*, which follows
- ◆ *Working with User Groups* on page 1-3
- ◆ *Working with Rules* on page 1-4
- ◆ *Assigning Permissions* on page 1-6
- ◆ *Understanding Groups and Inheritance* on page 1-10
- ◆ *Understanding the ChangeMembership Permission* on page 1-11

Managing User Access to ControlCenter

Access to the ControlCenter application is controlled through *rules* that grant *permissions* to a single *user* and/or *groups* of users. The permissions determine what actions (commands) a user or group may perform on a given object or group of objects called a *user-defined group*.

Users are not *created* through ControlCenter. Users are initially created on hosts through Windows Administrative Tools (refer to *Creating ControlCenter Users* on page 1-13) or as LDAP users.

- ◆ A user group consists of a set of users to whom you want to apply the same permissions. A user inherits the permissions of the group in which the user is included.
- ◆ A user-defined group contains groups of objects (for example, all Symmetrix® arrays) that you want to provide access to for a specific group of users.
- ◆ A rule controls which users (or user groups) can access which objects (or user-defined groups) and what that user or user group is allowed to do with the user-defined group (permissions).

The terms *user-defined group* and *object group* are both used to represent groups of managed objects in ControlCenter.

Figure 1-1 shows the creation of a rule granting the SDR permission (authorizing the use of Symmetrix Dynamic Reallocation) to a single user (JSmyth) on a single Symmetrix array (S/N 000182402626). The complete procedure for creating a new rule for a user is provided in *Creating and Modifying Rules for Users and User Groups* on page 1-20.




Figure 1-1 Rule Providing a User with SDR Permissions on a Symmetrix Array

Working with User Groups

ControlCenter provides several default user groups that act as examples to help you begin to set up your authorization framework. You can modify or delete these groups (some exceptions apply to ECCAdministrators) to meet your needs.

Default User Groups

ControlCenter provides five initial user groups (Figure 1-2 on page 1-4):

- ◆ **ECCAdministrators** — Access to all objects. This group cannot be deleted or renamed. Members of this group can create users and groups and change permissions. This group initially contains one user, *eccadmin*, that you create during ControlCenter installation.

- ◆ **SAN Manager™** — Access to EMC® Enterprise Storage Network functions.
- ◆ **Symmetrix Configuration Manager** — Access to configuration functions.
- ◆ **Symmetrix Data Protection Manager** — Access to backup and recovery functions.
- ◆ **Symmetrix Performance Manager** — Access to tuning tools such as Optimizer and TimeFinder® SRDF/QoS.




Figure 1-2 Initial Groups

Working with Rules

ControlCenter provides initial authorization rules to help you set up your authorization framework. You can modify or delete these rules (some exceptions apply to ECCAdministrators) to meet your needs.

The basic principles for rule creation are:

- ◆ Each user or user group can have only one rule applied to it.
- ◆ A user may belong to multiple user groups.
- ◆ A rule may mention more than one user or user group.

- ◆ Once a rule is created for a user or a user group, no other rule may be created that applies to that specific user or user group. However, if a user is a member of a user group mentioned in a rule, it is still possible to create a rule for that specific user.

A user can have a rule applied to it individually while at the same time having a rule applied to it as a member of a user group.

Rules are constructed using items under **Administration, Security Management, Authorization Rules** in the ControlCenter Console as shown in Figure 1-3.




Figure 1-3 Default Authorization Rules

Default Authorization Rules

Each entry represents a rule, depicted by a judge's gavel and the rule's name. Each rule may mention one or more users or user groups. The following rules are created by default:

The authorization system maintains a *privileged rule* that is never displayed and cannot be edited. This rule grants members of the ECCAdministrators group full permissions over creating and changing users, user groups, user-defined groups, and authorization rules.

- ◆ **ECCAdministrators Rule** — Provides all access to all objects. ECCAdministrators can create users and groups and change permissions.
- ◆ **SAN Manager Rule** — Provides access to SAN permissions.
- ◆ **Symmetrix Configuration Manager Rule** — Provides access to configuration permissions.
- ◆ **Symmetrix Data Protection Manager Rule** — Provides access to backup and recovery permissions.
- ◆ **Symmetrix Performance Manager Rule** — Provides access to tuning tools such as Optimizer and TimeFinder SRDF®/QoS.
- ◆ **Any User Rule** — Provides access to permissions that you want to apply to all users. During installation, the StorageScope™ user permission is placed in the Any User rule by default. This permission provides view-only access to StorageScope reports (refer to *StorageScope Permissions* on page 13-4).

Assigning Permissions

Authorization involves creating rules that assign permissions to users. Permissions define what a user is allowed to do with a specified object or group of objects in a user-defined group. Users that require similar access permissions can be formed into user groups to simplify management.

The authorization GUI obeys the following principles:

- ◆ A specific user or user group can have only one rule assigned to it. However, a user may be a member of a user group that is mentioned in a rule and still have a rule applied to it that mentions that user specifically.
- ◆ A user may be a member of multiple user groups.
- ◆ A rule may mention more than one user or user group.
- ◆ Objects can be formed into user-defined groups so that rules can apply permissions to an entire group of objects at once.

Descriptions of all ControlCenter permissions are provided in the Console online Help under **Administering ControlCenter Users, Working with authorization rules and permissions, ControlCenter permissions**.

Table 1-1 provides a list of the permissions required to accomplish common high-level ControlCenter tasks.

Table 1-1 Tasks and Permissions

If you want to:	You need to edit this kind of object:	And give it at least this permission:
Manage Agents		
Install, start, stop, and uninstall agents, view agent installation and uninstallation logs, and restart the master agent.	ECC Agent	Agent Management
Manage Users, Groups, Rules, and Objects		
Create, edit, and delete authorized ControlCenter users and user groups.	User Account Manager Data	UserAccountManager
Create, edit, and delete authorized ControlCenter users only.	User Account Manager Data	ManageUsers
Create, edit, and delete user groups only.	User Account Manager Data	ManageUserGroups
View User and group information only.	User Account Manager Data	Read
Create a new rule.	Authorization Data	Authorize (also provides Read privileges)
Read authorization rules.	Authorization Data	Read
Change Membership of User-Defined (Object) Groups		
Change membership of user-defined group.	ObjectGroup	ChangeMembership
Change Permissions on Managed Objects		
Change permissions on ControlCenter managed objects.	Various (refer to the online Help for the specific object)	DatabaseUpdate
Manage the SAN		
Discover and view topology.	N/A	No permissions required
Rename the fabric, enable or disable zoning operations, change default zoning, activate the zone set.	Fabric	Fabric Management
Discover switches, rename a switch, change passwords, delete a switch, clear zone operations, import a zone set.	Switch	Switch Management
Display zones and zone sets.	N/A	No permissions required

Table 1-1 Tasks and Permissions (continued)

If you want to:	You need to edit this kind of object:	And give it at least this permission:
Create, activate, rename, delete, and/or clone zone sets, and remove zones.	ZoneSet	Zoneset commands
Create, delete, rename and/or clone zones, and add or remove ports from zone.	Zone	Zone commands
Manage both zones and zone sets.	Zone/ZoneSet	Zone Administration
Create, delete, modify, and rename zoning policies	Zoning Policy	Zoning Policy Administration
Perform Device Masking Operations on a Storage Array		
Perform CLARiiON® device masking operations.	CLARiiON	CLARiiON Device Masking
Perform device masking on an HP-XP array.	HDS	HDS Device Masking
Perform Symmetrix Storage Device Masking administration (set who can perform SDM operations).	Symmetrix	SDM Administration
Perform Symmetrix storage device masking configuration operations including changing access rights, getting logs, carrying out SID lockdown, and so on.	Symmetrix	SDM Configure
Perform operations on the SDM database including initialize, refresh, synchronize, restore, backup, and edit.	Symmetrix	SDM Database
Perform StorageWorks device masking operations.	StorageWorks	StorageWorks Device Masking
Work with StorageScope Reports		
View reports.	StorageScope Reports	StorageScope User
Create custom report layouts.	StorageScope Reports	StorageScope User
Schedule reports and modify report retention policies.	StorageScope Reports	StorageScope Admin
Run reports in realtime.	StorageScope Reports	StorageScope Admin
View report history.	StorageScope Reports	StorageScope Admin

Table 1-1 Tasks and Permissions (continued)

If you want to:	You need to edit this kind of object:	And give it at least this permission:
View or Edit Alert Definitions		
View and modify alert definitions.	Alert	Edit Alert Definitions
Assign an alert.	Active Alert	Assign Alert
Clear an alert.	Active Alert	Clear Alert
View or Modify Data Collection Policies (DCPs)		
View and modify DCPs.	Data Collection Policy	Edit Data Collection Policy
Map Devices, Configure Arrays, Bind LUNs, and Create RAID Groups for Various Arrays		
Map, unmap, move, and change the addresses of devices connected to front-end ports on HP StorageWorks arrays.	StorageWorks	HP StorageWorks Device Mapping
Expand devices on HP StorageWorks arrays.	StorageWorks	HP StorageWorks Device Modification
Configure CLARiiON arrays.	CLARiiON	CLARiiON Array Configuration
Bind and unbind CLARiiON LUNs.	CLARiiON	CLARiiON LUN Management
Create or delete CLARiiON RAID groups.	CLARiiON	CLARiiON RAID Group
Create metaLUNs for CLARiiON arrays.	CLARiiON	FLARE Fusion
Map, unmap, move, and change the addresses of devices connected to the front-end ports on an HP-XP array.	HDS	HDS Device Mapping
Create logical unit size expansion (LUSE) volumes on an HP-XP array.	HDS	HDS LUSE Management
Allocate or Deallocate Storage Using the Storage Provisioning Services		
Create, modify, or delete storage pools (groups of devices available for allocation) on a specified storage array(s).	Symmetrix CLARiiON StorageWorks	Allocation Administration
Perform storage allocation. You need permissions on both the host and the storage pool.	Host Storage Pool	Allocation Execution

Table 1-1 Tasks and Permissions (continued)

If you want to:	You need to edit this kind of object:	And give it at least this permission:
Save and edit storage allocation tasks on the storage allocation task list. You need permissions on both the host and the storage pool.	Host Storage Pool	Allocation Reservation
Create a deallocation task through Storage Provisioning Services (SPS). You need permissions on both the host and the array.	Host Array	Deallocation Reservation
Execute a deallocation task through Storage Provisioning Services (SPS). You need permissions on both the host and the array.	Host Array	Deallocation Execution

Understanding Groups and Inheritance

A ControlCenter user group is a group of users who share the same authorization characteristics. Objects are formed into user-defined groups in this same way, so that rules may apply permissions to an entire group of objects at once.

One important consequence of user groups and object-defined groups is that permissions are inherited through the group structure. That is, groups may have subgroups, and permissions granted to a group also apply to users in a subgroup of that group. However, user-defined groups of objects have an additional permission that applies to the groups themselves (as opposed to the members) called the *ChangeMembership* permission. This permission restricts the ability to add and delete members from a user-defined group.

Understanding the ChangeMembership Permission on page 1-11 provides an overview of the *ChangeMembership* permission.

Refer to the online Help topic: **Administering ControlCenter users, User management concepts** for details about the differences in managing user groups and user-defined groups.

Understanding the ChangeMembership Permission

The ChangeMembership permission applies to the user-defined group itself, rather than to the objects that are members of the group. This permission controls who is allowed to add or delete objects from the user-defined group.

It is important that you control who can change the membership of user-defined groups used in authorization rules.

ChangeMembership Permission Inheritance

The ChangeMembership permission is inherited through the home of the group. The home of the group is the group in which it was *created* or a group into which it was *moved*.

There is a distinction between *linking* a group to another group, *moving* a group to another group, and *copying* a group.

- ◆ Linking a group to another group (select the group, press CTRL and SHIFT, and drag the group to another group) merely links the group to the new group so that any additions or deletions of objects from the original group are reflected in the new group.

This operation requires ChangeMembership permission on the destination group.

- ◆ Copying the group (select the group, press CTRL, and drag the group to another group) creates a new group with the same members as the old group.

This operation requires ChangeMembership permission on the destination group.

- ◆ Moving a group (by dragging the group from one group to another) changes the home of the group.

This operation requires ChangeMembership permission on both the old and new group.

User Access Management Procedures

This section provides procedures for setting up and maintaining access management at the user level with EMC ControlCenter. Refer to *Introduction to ControlCenter User Access Management* on page 1-2 for an overview of ControlCenter access management concepts.

In general, the following tasks are required to set up user access management:

- ◆ Create eccadmin. You created this user on the ECC Server host through Windows Administrative Tools as a local or domain user or as an LDAP user *before* installing ControlCenter. Refer to the *EMC ControlCenter Planning and Installation Guide, Volume 1* for details.
- ◆ Choose existing LDAP users, or local or domain users on Windows hosts, or create new users through LDAP or Windows Administrative Tools (refer to *Creating ControlCenter Users* on page 1-13).
- ◆ Add new users to ControlCenter from the LDAP users or from Windows users (refer to *Adding a User to ControlCenter* on page 1-16).
- ◆ Create groups of users for efficient user management (refer to *Managing ControlCenter User Groups* on page 1-17).
- ◆ Create or modify the rules that govern the permissions allowed users (or user groups) on specified objects (or user-defined groups).

Creating ControlCenter Users

ControlCenter users in general do not require any special privileges. However, users must exist as LDAP users or as local or domain users on a Windows host. The ControlCenter Server (ECC Server) must be a member of the domain or have an established trust relationship with the domain of the host on which the users were created.

LDAP is supported for user authentication purposes only. LDAP does not control ControlCenter user permissions.

Creating the ControlCenter eccadmin User

During ControlCenter planning and installation, you created a user named `eccadmin` through Windows Administrative Tools on the ECC Server host or as an LDAP user. When the ECC Server starts for the first time, `eccadmin` is added to the ControlCenter Repository as a member of the ECCAdministrators User Group. This user has the authorization to add new users, create and edit user groups, and create and edit authorization rules. Refer to the *EMC ControlCenter Planning and Installation Guide, Volume 1* for more details.



CAUTION

The `eccadmin` account is an anonymous account. Anyone logging into it can perform any function without giving away their identity in log files. For this reason, it is desirable (but not necessary) to limit or eliminate `eccadmin` as a ControlCenter login account. Refer to *Managing ControlCenter User Groups* on page 1-17.

Creating ControlCenter Users through Windows Administrative Tools

You can use current Windows users (Windows 2000 or Windows Server 2003) or LDAP users or you must create new users before you can add them as users to ControlCenter. Create new Windows users as follows:

1. From the **Start** menu, select **Settings, Control Panel**.
2. Open **Administrative Tools**, and then open **Computer Management**.
3. From the Systems Tools folder on the Computer Management dialog box, open the **Local Users and Groups** folder. A view appears displaying the Users and Groups folders.

4. Open the **Users** folder. The local users appear.
5. From the **Action** menu, select **New User**. The New User dialog box appears (Figure 1-4).


The image shows a 'New User' dialog box with a title bar containing a question mark and a close button. The dialog has several text input fields: 'User name:', 'Full name:', 'Description:', 'Password:', and 'Confirm password:'. Below these fields are four unchecked checkboxes: 'User must change password at next logon', 'User cannot change password', 'Password never expires', and 'Account is disabled'. At the bottom right, there are two buttons: 'Create' and 'Close'.

Figure 1-4 **New User Dialog Box**

6. In the New User dialog box, enter the required information (username, full name, description, password, confirm password).



CAUTION

EMC strongly recommends that *all* user accounts assigned as EMC ControlCenter users be given unique, hard-to-guess passwords.

7. Clear (uncheck) the **User Must Change Password at Next Logon** checkbox.
8. Click **Create**.
9. Repeat this procedure for each user requiring access to ControlCenter.

Creating a Domain Account for the ECC Server Service

In some situations you will be unable to add domain users to ControlCenter because the ECC Server service is not a qualified domain user.

If you are unable to add domain users to ControlCenter, ask your Windows Domain Administrator to complete the following steps to create a new domain account for the ECC Server service:

1. On the domain controller, create or select an existing domain account.

This account will be used exclusively as a service account for the ECC Server service.

2. On the ECC Server host, grant this account the *Act as part of operating system* user right.
3. On the ECC Server host, add this account to the local Administrators group.

If steps 2 and 3 are not completed, you will not be able to connect to the server from the ControlCenter Console.

4. Make this account the service account for the ECC Server service.
5. Stop the EMC Web Applications Server, Store, and ECC Server services (in that order).
6. Start the services in the reverse order (ECC Server, Store, Web Applications Server).
7. Bring up the ControlCenter console and add the domain users you want to have access to ControlCenter (refer to *Adding a User to ControlCenter* on page 1-16).

Adding a User to ControlCenter

Once users exist on a Windows host or as LDAP users, you can add them to ControlCenter. Adding a user to ControlCenter allows that user to log on to the ControlCenter application.

Add a user as follows:

1. From the Console tree, right-click **ECC Users** under **Administration, Security Management**, and select **New**. The User Definition (New) dialog box appears (Figure 1-5).
2. Fill in the **Login ID** field in the form *username* for a local user or *domain\username* for a domain login ID (the ControlCenter Server must be a member of the domain or have an established trust relationship with the domain).

The ControlCenter *user ID* in the **Login ID** field is the same as the name you entered in the **Username** field in Windows Administrative Tools shown in Figure 1-4 on page 1-14 or as the LDAP user.

3. Add a description of the login ID if desired.




Figure 1-5 **New User Definition**

You can change the ControlCenter login ID for a user, but it must have already been changed in the **Username** field in Windows Administrative Tools or through LDAP. Once the username is changed on the host or through LDAP, you can change the login ID through ControlCenter.

Once a user is added as a ControlCenter user, you set permissions for rules that control the user's access to monitoring, managing, and controlling objects in the ControlCenter environment. Rules may apply to either individual users or groups of users as explained in the *Managing ControlCenter User Groups* on page 1-17.

Managing ControlCenter User Groups

User groups save administrative time and effort by allowing you to manage rules for groups of users instead of for individuals. This section provides procedures for:

- ◆ *Creating a New User Group*, which follows
- ◆ *Adding a User to a Group* on page 1-18

Refer to the online Help topic: **Administering ControlCenter users, Working with user groups** for detailed procedures for managing user groups.

Creating a New User Group

Create a new user group as follows (Figure 1-6):

1. From the Console tree, right-click **ECC User Groups** under **Administration, Security Management**, and select **New**. The User Group Definition (New) dialog box appears.




Figure 1-6 Creating a New User Group

2. Type the name of the new group and a description, and click **OK**. The new group appears under ECC User Groups.
3. Populate the user group by clicking and dragging users from ECC Users to the new group.

Adding a User to a Group

4. You can also create a user group from within the rule creation dialog box by right-clicking **ECC User Group** and proceeding with steps 1 through 3.

A user may be added to an ECC user group by using the mouse to drag the user from the ECC Users folder to the group.

In the example shown in Figure 1-7, a user named *JSmyth* was added to the Symmetrix Configuration Manager Group.




Figure 1-7 Adding a User to a Group

Removing the eccadmin User from ControlCenter

The `eccadmin` account must be used to initially log on to ControlCenter because no other ControlCenter users exist at that point. However, `eccadmin` is an anonymous account, which means that anyone logging into it can perform any function without revealing their identity in log files.

It is desirable (but not necessary) to limit or eliminate `eccadmin` as a ControlCenter login account.

The `eccadmin` user cannot be removed from the ECCAdministrators group until another user is added to the group (the group cannot be deleted and must always have at least one member). Once you add another user to the ECCAdministrators group, the `eccadmin` account is no longer required and may be removed as follows:

The new ControlCenter user must already be a valid account on a Windows host or as an LDAP user.

1. Add a user to the ECCAdministrators group as outlined in *Managing ControlCenter User Groups* on page 1-17.
2. Log out of the Console and log back in as the user you just added to the ECCAdministrators group.
3. From the Console tree, expand **Administration, Security Management, ECC Users**.
4. Right-click **eccadmin** and select **Delete**.

A message appears asking if you really want to delete the user `eccadmin`.

5. Click **Yes**.

The `eccadmin` account is deleted from ControlCenter.

The corresponding user account for `eccadmin` that exists on the ECC Server host (or as an LDAP user) may be left in place in case there is ever a need to add the `eccadmin` user back into ControlCenter for maintenance purposes.

Creating and Modifying Rules for Users and User Groups

Rules allow you to control what a user or user group can do to specified objects. A user or user group can only have one rule applied to it. You can modify a rule by editing (adding or deleting) the permissions, renaming the rule, or deleting the rule.

Refer to the online Help topic: **Administering ControlCenter users, Working with authorization rules and permissions** for detailed procedures.

You can create a new rule with the following steps:

1. From the Console tree, expand **Administration, Security Management**.
2. Right-click the user or user group for which no rule exists, and select **Authorization, New Rule** (Figure 1-8).




Figure 1-8 Creating a New Rule

If the user or user group is already mentioned in a rule, this option does not appear.

The Rule Definition for User Group (or User, depending on whether you are creating a rule for a user or user group) dialog box appears as shown in Figure 1-9.




Figure 1-9 Naming the New Rule

3. Enter a name for the rule.
4. Under **Choose Objects** (located above the left panel), select either **Groups/Instances** or **Types** depending on the kind of objects you are creating.

If a group of objects, object instance, or object type is selected, the available permissions appear under the middle panel (**Available actions**).

5. Select one or more of the permissions and click > to place them in the rule.
6. Click **OK** to write the new rule to the database.

Once a rule is created containing a user or user group, no additional rules may be created for it. However, existing rules may be edited to remove permissions or grant additional ones.

Monitoring and Contacting Console Users

ControlCenter allows members of the ECCAdministrators group to monitor Console use as well as to send messages to Console users.

Only members of the ECCAdministrators group can monitor Console use or send messages to other Console users. Refer to *Adding a User to a Group* on page 1-18 for the steps for adding a user to a group.

Monitoring Console Users

ControlCenter allows members of the ECCAdministrators group to monitor Console use through the At A Glance view. ControlCenter displays the user IDs that are currently logged on, the name and IP address of the host, and the time they logged on.

You access the status of ControlCenter users as follows:

1. Click the **At A Glance** button in the toolbar to display the At A Glance view, and then click **Show All** to display all of the At A Glance charts.
2. In the ECC Status chart, click the **Consoles** bar.
3. On the Drill Down By toolbar, click **Consoles**. The Consoles view with the current ControlCenter users appears.

Sending Messages to Console Users

ControlCenter allows members of the ECCAdministrators group to send messages to other console users. You can broadcast a message to a specific user or group of users before performing actions that could be disruptive to other users, such as placing a configuration lock on a Symmetrix or stopping a ControlCenter agent.

The message appears immediately in the Consoles of those users, on top of all other windows. They cannot reply to it, nor is it logged.

To send an immediate notification to a user:

1. From the **Consoles** view (that you accessed in the previous steps), right-click the user to whom you want to send the notification and select **Consoles, Send Message**. The Send ECC Console Message dialog box appears.
2. Type the message and click **OK**. The message cannot exceed 200 characters.

Administering ControlCenter Data Collection Policies

This chapter provides a brief overview of the tasks involved in administering data collection policies, and consists of the following sections:

- ◆ Data Collection Overview 2-2
- ◆ Managing Data Collection Policy Definitions and Templates.....2-4
- ◆ Agent Data Collection Policy Reference 2-7

Data Collection Overview

Data collection policies are a formal set of statements used to manage the data collected by most ControlCenter agents. The policies specify the data to collect and the frequency of collection. Each agent has associated predefined collection policies and collection policy templates, which can be managed through ControlCenter Administration.

Data Collection Policy Definitions

Predefined data collection policies are provided automatically with each agent. You can edit, copy, or delete predefined collection policies.

Data Collection Policy Templates

Data collection policy templates provide default values for the creation of new collection policies. ControlCenter provides at least one template for each agent. You can define your own policies by modifying the collection policy templates.

Enabling Data Collection Policies

Agent data collection is performed differently for each agent. Some agents:

- ◆ Start data collection automatically upon startup.
- ◆ Require manual configuration to collect data.
- ◆ Require that data collection policies are defined, assigned, and enabled to manage how and when the data is collected.

Scheduling Data Collection Policies

You can optimize the scheduling of data collection to avoid overtaxing the ECC server with too much data at one time by using the Schedule Discovery utility available on Powerlink™ at <http://powerlink.emc.com>.

Understanding the Schedule Discovery Utility

The Schedule Discovery utility generates a report of recommended times to schedule when discovery of the Host Agents for AIX (MAR), HP-UX (MHR), Solaris (MSR), and Windows (MNR), and Symmetrix SDM Agents (EGZ) should take place to avoid performance issues. Processing too many discovery policies at the same time causes performance issues for the ECC Server. This scheduling utility works for ControlCenter 5.1 and higher. Previous releases of ControlCenter are not supported.

The report generated by the Schedule Discovery utility is provided to help you decide how to configure the Discovery data collection policy (DCP). A DCP is a formal set of statements used to manage the data collected by ControlCenter agents that specifies the data to collect and the schedules for collection.

The Schedule Discovery utility analyzes the existing DCPs and suggests new collection schedules. The new schedules must be input manually by the user and must be both defined and applied to the agent in question. The user should take care that any existing schedules that are set to “Apply this policy to all applicable hosts” should have this attribute turned off, so that only the new schedules are used.

Downloading the Schedule Discovery Utility

On Powerlink, download the *EMC ControlCenter 5.2 Schedule Discovery Utility Technical Notes*. These contain instructions on how to download, install, and use the Schedule Discovery utility.

Managing Data Collection Policy Definitions and Templates

You can define a new data collection policy based upon a predefined policy or policy template, and edit, copy or delete existing collection policies. Once the policy settings are defined or modified, you can apply the policy to specific agents or managed objects.

Note the following additional resources:

- ◆ Refer to the Console online Help for detailed step-by-step procedures for all data collection policy procedures.
- ◆ For environments with many hosts and arrays, refer to the *EMC ControlCenter 5.2 Planning and Installation Guide, Volume 1*.

Policy Management Options

Managing the data collection policies consists of:

- ◆ **Assigning Data Collection Policies** — Each agent is assigned a set of predefined policies and a set of policy templates. You can define new data collection policies from a predefined policy or from a policy template.
- ◆ **Editing Data Collection Policies** — You can edit all settings for an existing data collection policy; however, you can only edit the schedule and properties defined by the data collection policy templates.
- ◆ **Copying Data Collection Policies** — You can use the copy policy function when you want to have more than one data collection policy with similar settings.
- ◆ **Deleting Data Collection Policies** — You can only delete policies in the **Policies Definitions** branch of the Administration tree. Data collection policy templates cannot be deleted.
- ◆ **Viewing Data Collection Policies** — You can create a tabular view of specific data collection policies and template settings.

Accessing Data Collection Templates and Policies

In general, collection policies for a specific agent are accessed as follows (Figure 2-1):

1. From the Console tree panel, expand **Administration, Data Collection Policies, Policy Definitions**, and the specific agent folder (Figure 2-1).





Figure 2-1 Selecting a Policy to Modify

2. Right-click the policy you are modifying or assigning and select **Edit**. The Policy Definition dialog box appears (Figure 2-2 on page 2-6).



The image shows a Windows-style dialog box titled "Policy Definition: Topology Validation_00 (Edit)". It contains several fields and tabs. The "Display Name" field is set to "Topology Validation_00" and the "Unique Descriptor" field is set to "00". A checkbox labeled "Policy Enabled?" is checked. Below these fields are four tabs: "Properties", "Source", "Actions", and "Apply To". The "Properties" tab is selected, showing fields for "Agent" (Fibre Channel Connectivity Agent), "Template Name" (Topology Validation), "Last Modified" (May 2, 2003 4:15:47), and "Who Modified" (Installation). A large text area for "Description" contains the text: "The policy allows the user to monitor the switch topology. If the ECC Repository and the switch connectivity do not match the agent generates an alert." At the bottom right are "OK", "Cancel", and "Help" buttons.

Figure 2-2 Policy Definition Dialog Box

3. Click the **Actions** tab and then the **Edit** button to edit the schedule for this DCP.
4. Click **OK** when you complete your edits to the DCP.
5. If the agent folder does not contain the policy you want to assign, expand the **Policy Templates** folder under **Administration, Data Collection Policies**, and then expand the agent folder to find the appropriate policy template.

Agent Data Collection Policy Reference

The following tables list agents and their data collection policies (DCPs). If an agent is not listed, it means that data collection is managed using another process. Refer to the specific agent overview topic in the online Help for detailed information.

Basic DCPs that are turned on by default when the agent is started are found in the **Policy Definition** folder. If a DCP is not in the Policy Definition folder, it has not yet been defined for this agent and is located in the **Policy Templates** folder. Refer to *Accessing Data Collection Templates and Policies* on page 2-5 for more information.

Table 2-1 Backup Agent DCPs

Backup Agents	DCPs	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
AIX Backup Agent for NetBackup AIX Backup Agent for NetWorker AIX Backup Agent for TSM Generic Backup Agent for EDM™ HP-UX Backup Agent for NetBackup HP-UX Backup Agent for NetWorker HP-UX Backup Agent for TSM Solaris Backup Agent for NetBackup Solaris Backup Agent for NetWorker Solaris Backup Agent for TSM Windows NT Backup Agent for NetBackup Windows NT Backup Agent for NetWorker Windows NT Backup Agent for TSM	Discovery	Policy Definition	Y	Every day at 4 a.m.

Table 2-2 Common Agent DCPs

Common Agents	DCPs	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
WLA Archiver	WLA Retention	Policy Definition	Y	Agent Controlled

Table 2-3 Connectivity Agent DCPs

Connectivity Agents	DCPs	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
Symmetrix SDM Agent	Discovery	Policy Definition	Y	Every 12 hours
	Masking Configuration	Policy Definition	Y	Every 6 hours
Fibre Channel Connectivity Agent	Device Validation	Policy Definition	Y	Every hour
	Fabric Validation	Policy Definition	Y	Every hour
	Historical Data	Policy Definition	N	Every day at 12 a.m.
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 15 minutes
	Discovery Scan	Policy Template	N	Every 30 minutes
	Performance Statistics	Policy Template	N	Every 15 minutes
	WLA Analyst	Policy Template	N	Every 15 minutes

Table 2-4 Common Mapping Agent DCPs

Common Mapping Agent	DCPs	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
Common Mapping Agent	DB2 Host ^{a, b} Informix SqlServer Sybase	Policy Template	N	Every day at 12 a.m.

- a. If you want to perform only monitoring functions on hosts, consider using the Common Mapping Agent and enabling the *Host* DCP instead of running a host agent on each host. Use of Common Mapping Agent improves host performance and reduces overhead, but keep in mind that Common Mapping Agent does not support active commands, real-time explore, or alerts.
- b. You can have either the *Discovery* data collection policy for a given host agent enabled or you can have the *Host* DCP for the Common Mapping Agent enabled, but you cannot have both enabled on the same host at the same time. Note that the Host agent *Discovery* DCP starts by default when the Host Agent is installed and started.

Table 2-5 Database Agent for Oracle DCPs

Database Agents for	DCPs	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
Oracle ^a	WLA Analyst	Policy Definition	N	Every 15 minutes
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Template	N	Every 15 minutes
	Oracle Agent Data Collection	Policy Template	N	Once per day at 6 a.m.

- a. The database agent for Oracle has two additional default settings that must be turned on for the Oracle database instance to be fully discovered. From within the Policy Definition dialog box, click the **Source** tab and set **Collect Configuration** and **Collect Allocation** to **Yes**.

Table 2-6 Host Agents for AIX, HP-UX, Linux, and Solaris DCPs

Host Agents for	Available Data Collection Policies	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
AIX	Discovery	Policy Definition	Y	Every day at 4 a.m.
	Watermarks for File Systems	Policy Definition	Y	Every 15 minutes
	Watermarks for Logical Volumes	Policy Definition	Y	Every 15 minutes
	Watermarks for Volume Groups	Policy Definition	Y	Every 15 minutes
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 2 minutes
	WLA Analyst	Policy Template	N	Every minute
HP-UX	Discovery	Policy Definition	Y	Every day at 4 a.m.
	Watermarks for File Systems	Policy Definition	Y	Every 15 minutes
	Watermarks for Logical Volumes	Policy Definition	Y	Every 15 minutes
	Watermarks for Volume Groups	Policy Definition	Y	Every 15 minutes
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 2 minutes
	WLA Analyst	Policy Template	N	Every minute

Table 2-6 Host Agents for AIX, HP-UX, Linux, and Solaris DCPs (continued)

Host Agents for	Available Data Collection Policies	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
Linux	Discovery	Policy Definition	Y	Every day at 2 a.m.
	Watermarks for File Systems	Policy Definition	Y	Every 15 minutes
	Watermarks for Logical Volumes	Policy Definition	Y	Every 15 minutes
	Watermarks for Volume Groups	Policy Definition	Y	Every 15 minutes
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 2 minutes
	WLA Analyst	Policy Template	N	Every minute
Solaris	Discovery	Policy Definition	Y	Every day at 2 a.m.
	Watermarks for File Systems	Policy Definition	Y	Every 15 minutes
	Watermarks for Logical Volumes	Policy Definition	Y	Every 15 minutes
	Watermarks for Volume Groups	Policy Definition	Y	Every 15 minutes
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 2 minutes
	WLA Analyst	Policy Template	N	Every minute

Table 2-7 Host Agent for Windows DCPs

Host Agent for	Available Data Collection Policies	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
Windows	Discovery	Policy Definition	Y	Every day at 12 a.m.
	Watermarks for File Systems	Policy Definition	Y	Every 15 minutes
	Watermarks for Logical Volumes	Policy Definition	Y	Every 15 minutes
	Watermarks for Volume Groups	Policy Definition	Y	Every 15 minutes
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 2 minutes
	WLA Analyst	Policy Template	N	Every minute
	File Level Reporting	Policy Template	N	Every day at 2 a.m.
	FLS File Level Summary Reporting	Policy Template	N	Every day at 2 a.m.
	FLS File Set	Policy Template	N	Agent Controlled
	FLS File Set Advanced	Policy Template	N	Agent Controlled

Table 2-8 Physical Agent for MVS DCPs

Physical Agents for	Available Data Collection Policies	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
MVS	MMP Data Collection	Policy Definition	Y	Every day at 2 a.m.
	WLA Daily	Policy Definition	N	Every 30 minutes
	WLA Revolving	Policy Definition	N	Every 15 minutes
	WLA Analyst	Policy Template	N	Every 15 minutes

Table 2-9 EMC Invista Agent

Agent	Available Data Collection Policies	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
EMC Invista™ Agent	Discovery	Policy Definition	Y	Every day at 12 a.m.

Table 2-10 Storage Agent DCPs

Storage Agents for	Available Data Collection Policies	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
Centera™	Discovery	Policy Definition	Y	Every day at 12 a.m.
CLARiiON	Discovery	Policy Definition	Y	Every day at 12 a.m.
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 2 minutes
	WLA Analyst	Policy Template	N	Every minute
HDS	Discovery	Policy Definition	Y	Every day at 12 a.m.
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 10 minutes
	WLA Analyst	Policy Template	N	Every 10 minutes
HP StorageWorks	Discovery	Policy Definition	Y	Every day at 12 a.m.
ESS	Discovery	Policy Definition	Y	Every day at 4 a.m.
NAS	Discovery	Policy Definition	Y	Every day at 12 a.m.
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 10 minutes
	WLA Analyst	Policy Template	N	Every 10 minutes

Table 2-10 Storage Agent DCPs (continued)

Storage Agents for	Available Data Collection Policies	Initial Location of DCP (Policy Definition or Policy Template Folder)	Enabled by Default Y/N	Default Schedule
SMI	Discovery for CLARiiON	Policy Definition	Y	Every 8 hours
	Discovery for ESS	Policy Definition	Y	Every 8 hours
	Discovery for HDS	Policy Definition	Y	Every 8 hours
	Discovery for StorageArray	Policy Definition	Y	Every 8 hours
	Discovery for StorageWorks	Policy Definition	Y	Every 8 hours
	Discovery for Symmetrix	Policy Definition	Y	Every 8 hours
	Discovery for TapeLibrary	Policy Definition	Y	Every 8 hours
Symmetrix	Alert Polling	Policy Definition	Y	Every 2 minutes
	BCV/RDF Status	Policy Definition	Y	Every 5 minutes
	CLI Generator	Policy Definition	Y	Every day at 12 a.m.
	Configuration	Policy Definition	Y	Every 10 minutes
	Historical Data	Policy Definition	Y	Every day at 12 a.m.
	Local Discovery	Policy Definition	Y	Every day at 12 a.m.
	Performance Statistics	Policy Definition	Y	Every 2 minutes
	Real-time BCV/RDF Status	Policy Definition	Y	Every minute
	WLA Daily	Policy Definition	N	Every 15 minutes
	WLA Revolving	Policy Definition	N	Every 2 minutes
	Proxy Discovery	Policy Template	N	Every day at 12 a.m.
	WLA Analyst	Policy Template	N	Every 5 minutes

Configuring and Managing Alerts and Notifications

This chapter provides an introduction to alerts and monitoring for the ControlCenter administrator. Administrative tasks involve the mechanics of setting alerts and notifications, directing them to specific personnel, configuring automated responses (autofixes), and troubleshooting alerts that do not trigger as expected. The chapter also describes best practices so that alerts processing provides the maximum benefit without redundancy or system impact.

For user information on how to view and respond to alerts and how to apply alerts to host resources, refer to Chapter 8, *Monitoring Storage with Alerts and Notifications*.

This chapter contains the following sections:

- ◆ Understanding Alerts and Notifications 3-2
- ◆ Setting Up Your Alert and Notification Strategy 3-7
- ◆ Creating Alerts and Notifications.....3-11
- ◆ Creating Alert Definitions 3-13
- ◆ Controlling Alert Spikes.....3-16
- ◆ Sending Alerts in E-mail or to a Management Framework.....3-17
- ◆ Automating Alert Responses with Autofixes 3-18
- ◆ Best Practices for Configuring and Managing Alerts 3-23
- ◆ Troubleshooting Alerts and Autofixes 3-25

Understanding Alerts and Notifications

ControlCenter allows you to monitor hundreds of metrics about your storage environment—one example is the I/O rate of a Symmetrix director. For each metric, you can set values at which you want ControlCenter to notify you—for example, when the I/O rate exceeds 15,000 operations per second. The *EMC ControlCenter Alert Matrix* lists all of the metrics.

You have several options for how ControlCenter notifies you. The most significant choice is whether you want to receive a notification or both an alert and a notification.

Alerts and notifications appear in different parts of the Console; alerts appear in the Alerts view (Figure 3-1), and notifications populate the At A Glance views (Figure 3-2).

	Object Name	Message
	lswk7051	Volume Group /dev/vg00 is at 4% free.
	lswk5041	Volume Group rootvg is at 10% free.
	lswk5041	File System /dev/hd9var is at 15% free.
	lswk5041	File System /dev/hd2 is at 11% free.
	lswk5040	File System /dev/hd3 is at 10% free.
	lswk5040	File System /dev/hd9var is at 2% free.
	lswk5040	Volume Group rootvg is at 4% free.
	lswk5040	File System /dev/hd3 is at 13% free.

Figure 3-1 Alerts View

The Alerts view provides a very detailed look at issues in your environment and offers many tools for tracking the issues to resolution.




Figure 3-2 At A Glance Health and Performance Views

The At A Glance views provide a higher-level perspective by dividing your environment into categories such as storage array performance or host capacity. For the At A Glance views, ControlCenter consolidates related alerts and notifications into charts that indicate the statuses of the various categories (Figure 3-2).

Alert and Notification Life Cycles

Figure 3-3 demonstrates the alert and notification life cycles using a file-size metric. This section also introduces important terminology.

Table 3-1 explains the stages in Figure 3-3.




Figure 3-3 Alert and Notification Life Cycle

Table 3-1 Alert and Notification Life Cycle Stages

Stage 1	ControlCenter can monitor hundreds of metrics about your storage environment, such as file size, the I/O rate of a Symmetrix Director, or the status of a daily backup operation. The online Help and the <i>EMC ControlCenter Alert Matrix</i> provide complete descriptions of each metric.
Stage 2	In the Console, the Administration, Alert Management branch of the tree provides alert templates for each metric that ControlCenter can monitor. The templates are organized by agent. The templates provide default values at which alerts and notifications will trigger. The template also specifies a schedule that determines how often ControlCenter will evaluate alerts and notifications. There is one template for each metric, and you can edit the default values. In addition to trigger values and a schedule, you can specify a management policy, which indicates who should be notified when an alert or notification triggers, and autofix, which is a script that runs automatically when an alert triggers.
Stage 3	You can create multiple alert definitions from a template. An alert definition specifies which object to monitor (for example, the filename for a file-size alert) and on which hosts or storage systems. In the alert definition, you specify whether ControlCenter should send a notification only, or both an alert and a notification when the monitored object exceeds the trigger values. For the trigger values, schedule, and management policy, you can use the template values or modify them as necessary.
Stage 4	When a monitored resource exceeds a trigger value, ControlCenter triggers the alerts and notifications that you have turned on. Alerts are sent to the Alerts view. Notifications are sent to the At A Glance views. In addition, ControlCenter executes the management policy steps (such as sending an alert to an e-mail address or SNMP framework) and runs any autofixes attached to the alert.
Stage 5	If the resource exceeds the next trigger value, ControlCenter updates the display of that alert or notification in the view (for example, changing the color and icon of an alert in the Alerts view, or adding length to the critical bar in an At A Glance view chart and removing length from the warning bar). In addition, the management policy steps and autofixes run again.
Stage 6	When the resource falls back under the trigger values you defined, ControlCenter removes the associated alerts and notifications from the Alerts and At A Glance views. You also can clear alerts and notifications manually from the views. Some alerts and notifications must be cleared manually. After an alert is cleared, it appears in the Alert History view. The alerts appear at the severity level at which they were cleared. For example, if an alert was critical when it was cleared, then it appears as a critical alert in the Alert History view. In the Alerts or Alert History view, right-click the alert and select Alerts, History to view all the changes to the alert, such as severity changes, assignments, and the creation of notes.
Stage 7	Alerts remain in the Alert History view for the period defined in the alert data retention policy. By default, the policy is disabled and alerts remain indefinitely.

Understanding Metric Types

There are two primary types of metrics: state metrics and count metrics. State metrics measure whether a condition is true or false (for example, whether a backup completed successfully). Count metrics measure whether a monitored condition has met a specific numeric value (for example, whether a file system has exceeded a specific size or the rate of a device has fallen below some value).

Within the three primary metric categories (health, capacity, and performance), health metrics are typically state metrics and capacity and performance metrics are count metrics.

After you install and enable an agent, the health metrics for that agent are typically enabled by default and the capacity and performance metrics are disabled.

In addition to state and count metric types, ControlCenter has interval and rate metrics. Interval metrics monitor the number of times an event occurs within a specified time range. For example, if an alert definition has a schedule of 15 minutes and 10 as the trigger value, then the condition must occur 10 times within 15 minutes for an alert to trigger. Rate metrics monitor the number of times a condition occurs every second. For example, if an alert definition specifies 100 as the trigger value, the condition must occur 100 times per second for an alert to trigger. The Alert Definition dialog box indicates the type of each metric.

For more explanation of metric types, refer to *Understanding alert types* in the online Help.

Setting Up Your Alert and Notification Strategy

Develop a strategy for how you want to use alerts and notifications in your data center. Create standards for the configuration and resolution of alerts and communicate them to all ControlCenter users.

Consider the following issues:

- ◆ *Controlling Who Creates and Edits Alerts and Notifications*, which follows
- ◆ *Deciding Whether to Define Alerts or Notifications*, which follows
- ◆ *Defining the Severity Levels* on page 3-8
- ◆ *Establishing Procedures for the Resolution of Alerts* on page 3-9
- ◆ *Refining Alert Definitions* on page 3-10
- ◆ *Keeping Alerts and Notifications* on page 3-10

Controlling Who Creates and Edits Alerts and Notifications

Decide whether you want all users to be able to create and edit alert definitions or whether a single or group of administrators will create and edit all alert definitions.

To create alert definitions, a user must have the Edit Alert Definition permission for the Alert object type. To edit a specific alert definition, the user must have the Edit Alert Definition permission for that alert definition or for the Alert object type. Grant this permission to users who you want to be able to create and edit alert definitions. Allow users to edit alert definitions for which they are responsible, but not other alert definitions.

To clear an alert from the Alerts or Notifications views, a user must have the Clear Alerts permission for the Active Alert object type.

You can also control access to management policies, autofixes, and schedules.

Refer to *User Access Management Procedures* on page 1-12 for specific procedures for assigning permissions.

Deciding Whether to Define Alerts or Notifications

You can define a metric to appear as a notification only or as both a notification and an alert.

Define an alert if you anticipate tracking an issue to resolution. Define a metric as a notification only if you want a metric to

contribute to an overall picture of your environment, but you do not want to track the issue to resolution.

Or, you can configure the metric to appear as both an alert and notification under certain conditions and as a notification only for less severe conditions.

Defining the Severity Levels

For each metric, you can specify five unique trigger values, with each assigned a different severity level. Create standards in your data center for the types of conditions that correspond to the different severity levels. Table 3-2 provides recommendations.

Table 3-2 Defining Alert Severity Levels

Icon in Alerts view	Severity Level	Possible Meaning	Examples
	Fatal	A resource critical to the daily operation of your organization has failed or cannot perform at an acceptable level. The alert requires immediate attention.	A critical volume or file system is out of space. A critical process has failed.
	Critical	A critical resource is failing or its performance is severely degrading. The alert requires immediate attention to ensure the resource can continue to perform or does not fail.	A high disk queue is affecting the performance of a critical application. The memory performance of a critical server is poor.
	Warning	The performance or availability of a resource is nearing an unacceptable range. Monitor the resource carefully and possibly take action.	The free space in a database tablespace is below an acceptable threshold. A Symmetrix subsystem raised an environmental alarm.
	Minor	An abnormal event occurred. The event may indicate current or future problems.	The daily backup of a nonessential file system did not occur. A Symmetrix subsystem issued an error message.
	Information	Use this severity level for informational messages.	The backup of a critical resource completed successfully. A device mapping changed.

Establishing Procedures for the Resolution of Alerts

Acknowledging Alerts

ControlCenter provides several tools in the Alerts view to help track alert resolution. Decide whether and how you will use these tools and communicate alert resolution procedures to all users.

Users can acknowledge a new alert to indicate that someone is working on resolving the alert. In the Alerts view, the alert text changes from bold to plain text and the new alert count in the upper-right corner of the Console decreases. In addition, the name of the user who acknowledged the alert is recorded in the Acknowledged By column of the Alerts view.

Decide who should acknowledge alerts: the ControlCenter or alert administrator, team leaders who assign the alerts to other users, or any user.

If you do not want to use the acknowledge-alert feature, you can hide this column in the Alerts view.

Assigning Alerts

Establish whether you will assign alerts to individual users or user groups and who will assign the alerts (an administrator, team leader, and so on). The name of the user or group appears in the Assigned To column of the Alerts view.

You can change the user or group after you initially assign the alert. One strategy is to assign alerts to a user group and then have the team leader of the group reassign the alert to a team member.

Users can sort or filter the Alerts view by the Assigned To column.

If you do not want to use the alert-assignment feature, you can hide this column in the Alerts view.

Recording Alert Notes

Create standards for how you will use alert notes. You can add a note to an alert when you acknowledge, assign, or clear it or at any other time.

If you use notes to document how you resolve alerts, you can later search the notes to help resolve new alerts of the same type.

Refining Alert Definitions

Creating alert definitions is an iterative process. After an alert triggers, review the alert history to ensure you have defined the alert effectively. This is especially important after you create an alert definition.

Look for the following opportunities to optimize the alert definition:

- ◆ Determine whether trigger values and severity levels are set appropriately. For example, ensure that the trigger values are not set too low, causing the alert to trigger too frequently.
- ◆ Determine whether you can configure some alerts to appear as notifications instead of alerts, thereby reducing the number of alerts that appear in the Alerts view.
- ◆ Notice whether the alert is spiking (frequently exceeding a trigger value and then returning to normal levels). If so, adjust the trigger values or set spike-controlling values that prevent the alert from triggering until it exceeds a trigger value for several consecutive intervals.
- ◆ Notice how quickly users respond to an alert. Determine whether the alert schedule is appropriate considering the user response. For example, if users typically respond to an alert after several days, then a schedule that evaluates the alert every hour may not be necessary.

Keeping Alerts and Notifications

After you clear an alert from the Alerts view or a Notification from the Notifications view, ControlCenter keeps it in the Alert History view. The length of time ControlCenter keeps the alert or notification in the Alert History view depends on the Alert data retention policy. By default, the Alert data retention policy is disabled, and ControlCenter keeps resolved alerts indefinitely.

Refer to *Defining data retention policies* in the online Help for specific procedures on defining data retention policies.

Creating Alerts and Notifications

After you gather the necessary information, set alerts and notifications in the Administration branch of the selection tree. Review the *EMC ControlCenter Alert Matrix*, which lists all the metrics provided by ControlCenter, to determine which metrics you want to monitor.

In addition, determine if any existing schedules and management policies meet the needs of the resources you want to monitor. If they do, you can attach the given schedule and management policy to the alert definition you create. If not, you can create a custom schedule, management policy, or both.

Tips for Setting Alerts and Notifications

Use the following tips to set alerts and notifications:

- ◆ Set more than one alert of a given type.
- ◆ Monitor similar resources together using the same alert, which will have the same schedule and management policy.
- ◆ For resources managed by different teams or individuals, create separate alerts with different management policies.
- ◆ Use separate alerts for resources of different levels of importance.
- ◆ For critical resources, use a more frequent schedule.
- ◆ Use wildcards to restrict alert processing to limited sets of resources. For example, do not set an alert to monitor the size of all files on a system.

Getting Help

When you are creating or modifying an alert definition in the Alert Definition dialog box, click Help to find out more about the alert.

Descriptions of all alerts are also available through the online Help table of contents. To access the Help:

1. Select **EMC ControlCenter Help** from the **Help** menu.
2. On the **Contents** pane of the Help Navigator dialog box, expand **Alert Descriptions** and then expand the node for the desired agent.

After an alert triggers, you can right-click the alert in the Alerts view or Alert History view and select **Alerts, Help** to get help on responding to some specific alerts.

Gathering Information

Table 3-3 describes the information you need to gather to set alerts for file systems, directories, and files. See the online Help for detailed requirements for individual alerts. The *EMC ControlCenter Alert Matrix* lists all the alerts that ControlCenter provides.

Table 3-3 Gathering Information for File Systems, Directories, and Files

Information needed	Description	Instructions	Notes
Hosts	Hosts that need to be monitored	List the hosts you want to monitor.	Other types of alerts may check multiple hosts, but file-system and file alerts are best reserved for a single host.
Source	Resources to be monitored	List the file systems, directories, files, and disks you want to monitor on each host.	Explore hosts for their file systems, directories, important files, and disks.
Conditions	Trigger values and alert severities	For each resource, determine the values that should trigger alerts. - File systems and disks: Determine the threshold free space and percentage free space. - Files and folders: Determine triggers for size, change in size, and percent change in size. Consider multiple thresholds for alerts of increasing severity: warning, critical, and fatal.	To help you determine trigger values, use recent data for resource free space and size. Also, consult the user of the resource.
Schedule	Frequency that the alert conditions are evaluated	For each resource, determine: - How often the alert condition should be checked. - The days of the week on which the alert condition should be checked.	Critical or faster-growing resources should be checked more often (every 5 to 60 minutes). Others should be checked less often to decrease alert processing (every 60 to 360 minutes).
Management policy	Names and e-mail addresses of personnel to notify	Determine whom an alert should notify automatically: - In the Console. - By e-mail. - By page. - In a framework product.	You can limit the display of alerts to the Consoles of administrators with responsibility for the affected systems or applications. You can configure alerts to send e-mail to key personnel at appropriate times.
Autofix	Automated responses to alerts, including predefined or user-defined commands and scripts	Determine an automated action that would help resolve the alert. Assemble any scripts or commands that the alert could issue when triggered.	

Creating Alert Definitions

You can use one of the following procedures to create an alert definition:

- ◆ *Creating an Alert Definition From a Template*, which follows
- ◆ *Creating Alert Definitions in the Edit Thresholds Dialog Box* on page 3-14
- ◆ *Copying an Alert Definition* on page 3-14

Once you create an alert definition, be sure to test it using the procedure outlined in *Testing an Alert Definition* on page 3-14.

As described in *Controlling Who Creates and Edits Alerts and Notifications* on page 3-7, you must have the Edit Alert Definition permission for the Alert object type or for specific alerts to create alert definitions.

Creating an Alert Definition From a Template

To create an alert definition:

1. In the selection tree, expand **Administration**, **Alert Management**, and **Alert Templates**.

The alert templates appear, organized by agent or component.

2. Expand the tree to view the templates for the agent that supplies the metric you want to use.
3. Continue expanding the tree until you reach the alert template.
4. Right-click the template, and then click **New**.

The Alert Definition dialog box appears.

5. On the **Source** tab, specify which resource you want to monitor (for example, a file or a tablespace and table in a database). You must complete all fields, and you can include wildcards. Click **Help** for field definitions. (Some alerts have no fields on this tab.)
6. On the **Conditions** tab, select levels and values at which you want ControlCenter to issue alerts and notifications.
7. On the **Actions** tab, specify the schedule and, optionally, a management policy and autofixes.
8. On the **Apply To** tab, select the objects ControlCenter should monitor, such as a particular storage system, host, file system, or device.

Creating Alert Definitions in the Edit Thresholds Dialog Box

You can also create alert definitions using the Edit Thresholds dialog box. The Edit Thresholds dialog box allows you to create or modify multiple alert definitions at one time. In addition, the dialog box shows the category for each metric. The category determines in which At A Glance view chart a notification will display (for example, Storage System Performance or Host Capacity).

To access the Edit Thresholds dialog box, right-click any object and select **Alert Thresholds, Edit Thresholds**.

Copying an Alert Definition

If similar managed objects have slightly different monitoring needs, you can copy an alert definition and modify the settings as necessary. This prevents having to create a second alert definition from scratch.

To create an alert definition by copying an existing alert definition:

1. In the selection tree, expand **Administration, Alert Management, and Alert Definitions**.
2. Expand the folder for the agent for which you want to create the alert definition.
3. Locate the alert definition you want to copy by expanding the subfolders.
4. Right-click the alert definition and select **Copy As**.

The Alert Definition dialog box appears.

5. Modify the definition as necessary.

Testing an Alert Definition

After you create an alert definition, test it to ensure that it triggers correctly for the desired resources.

To test an alert definition:

1. In the upper-right portion of the Console, view the current number of active alerts.
2. Edit the alert you want to test. The Alert Definition dialog box appears.
3. On the Conditions tab, select **Information**.
4. Specify a trigger value that is guaranteed to trigger the alert. For example, specify a free space percentage of less than 99, or a file size greater than 1.

5. In the **Actions** tab, select a schedule of **Minute_05**. This will cause ControlCenter to check the condition within 5 minutes of editing it. Click **OK**.

Note the following:

- ControlCenter checks alert conditions immediately when you create an alert, but when you edit an alert, it checks conditions at the end of the schedule interval.
 - Some alerts require two iterations of the schedule before they can be evaluated, such as alerts that measure growth over time. You will have to wait longer for these alerts to fire when testing.
6. In the upper-right portion of the Console, click **All Alerts**. Within a few minutes, the alert should trigger and the number of active alerts should increase.
 7. Locate the new alert. In the Alerts view, click the time and date column headings to sort the alert to the top.
 8. Verify the following:
 - The resource that triggered the alert is one that is monitored by the alert you set.
 - The trigger value and severity matches those of the alert you set.

Controlling Alert Spikes

ControlCenter provides a way for you to prevent alerts from triggering when a resource temporarily exceeds a trigger value, called an alert spike. For example, if a user creates a temporary file that causes the free space on a volume to drop below a trigger value, you may want to know about the condition only if the free space remains low for several hours and not when it temporarily dips. You can control these spikes using the Before and After fields in the Alert Definition dialog box and the schedule attached to the alert.

To prevent an alert from triggering when spikes occur:

1. In the selection tree, expand **Administration, Alert Management, and Alerts**.
2. Expand the folder for the agent to which the alert you want to edit belongs.
3. Expand the subfolders to locate the alert.
4. Right-click the alert and select **Edit**. The Alert Definition dialog box appears.
5. Click **Conditions**.
6. In the **Before** boxes, specify the number of consecutive times that the alert conditions must exist before ControlCenter triggers the alert.
7. In the **After** boxes, specify how many times, after an alert has triggered, that the alert must evaluate to false before ControlCenter removes the alert.
8. Click **Actions**.
9. From the **Schedule** list box, select a schedule to indicate how often ControlCenter should evaluate the alert.

Note that you also can access the Alert Definition dialog box from the Alerts view. Right-click the alert that is spiking and select **Edit Definition**. Then modify the Before and After values.

Sending Alerts in E-mail or to a Management Framework

Often, alerts trigger when you are not at the ControlCenter Console. To ensure you receive notification of critical alerts in a timely manner, set up ControlCenter to send alert messages by e-mail or to a management framework like HP OpenView Network Node Manager or Computer Associates Unicenter TNG.

To set up ControlCenter to send events to a management framework or e-mails:

1. To send e-mails, the ControlCenter administrator must configure ControlCenter for SMTP notification. This is typically done during installation of the ECC Server.

To send an event to a management framework, the ControlCenter administrator must configure ControlCenter to send SNMP traps to a management framework during installation and configuration of the ECC Server.

2. Create a management policy.

In the Management Policy Definition dialog box, drag the e-mail icon or SNMP icon into the management policy (and for e-mail, specify the address).

3. Attach the management policy to an alert.

Each time the alert triggers or moves from one severity to another, ControlCenter sends the e-mail or SNMP trap.

For complete instructions on configuring ControlCenter to work with a third-party framework application, refer to *EMC ControlCenter Integration Packages Product Guide*.

Automating Alert Responses with Autofixes

Through autofixes, ControlCenter allows you to specify commands or scripts that should run when an alert triggers.

To execute an autofix on a host in response to an alert, the agent issuing the alert must be running on the host.

Autofixes consist of a unique name, a descriptive name, and the text of the command or script ControlCenter sends to the host when an alert triggers. ControlCenter provides some autofixes; you can also create autofixes from new or existing scripts, batch files, or executables.

Configuring an autofix consists of three primary steps:

1. Creating the autofix definition in the Console.
2. Creating the autofix script on the host.
3. Attaching the autofix to an alert.

You must be a member of the ECCAdministrators user group to create autofix definitions in the Console.

Creating an Autofix Definition in the Console

To create an autofix definition:

1. In the selection tree, expand **Administration, Alert Management, and Autofixes**.
2. Right-click **User** and select **New**. The Autofix Definition dialog box appears.
3. Type a brief but unique name in **Internal Name**. This name cannot contain spaces.
4. Type a brief description in **Display Name**. This field can contain spaces but should not be too long as it appears in many tables throughout the interface.
5. In **Command**, type the syntax of the command or the name of a script or batch file.

When an alert triggers, the host agent uses the text exactly as it appears in this field. You can also pass alert information to your autofix commands. Refer to *Autofix Syntax* on page 3-20 for complete information.

ControlCenter executes an autofix any time an alert or notification to which the autofix is attached triggers or moves from one severity to another, whether the alert or notification increases or decreases in severity.

Passing Alert Information to an Autofix Script

ControlCenter allows you to pass the following information with your autofix command:

- ◆ Metric name
- ◆ Alert severity level
- ◆ Value that caused the alert to trigger
- ◆ Key values related to the alert, such as the name of the resource for which the alert triggered

Use the following syntax to pass this information with your autofix command:

```
your_autofix_command &METRIC &LEVEL &KEY &VALUE
```

&METRIC is the metric name.

&LEVEL is the severity level of the alert, in string format: Fatal, Critical, Warning, Minor, or Information.

&KEY is a value, such as a subsystem ID or filename, that ControlCenter passes when the alert triggers. These are the same values that appear in the alert message in the Alerts view. On UNIX and Windows, if the alert has more than one key, then append a number within brackets to &KEY for each key you want to pass, for example: &KEY[1], &KEY[2], and so on. On MVS, ControlCenter passes the first key only.

&VALUE is the value at which the alert triggered.

The *EMC ControlCenter Alert Matrix* lists the metric names and the messages that ControlCenter issues when alerts trigger. This document is available on the EMC ControlCenter documentation CD-ROM.

Autofix Syntax

Table 3-4 provides requirements and examples for specifying autofixes for UNIX, Windows, and MVS hosts.

Table 3-4 Autofix Syntax Requirements and Examples

Platform	Requirements and Examples
UNIX	Specify the file path and the name of a shell command, shell script, or executable. Ensure that files or programs referenced in a script are in the same directory as the script, or specify the complete path in the script. Examples <pre>/admin/tools/backup/backup.sh &KEY &VALUE /utility/cleanup/fixit.pl</pre>
Windows	Specify the name of a command, script, batch file, or executable. Include <code>cmd.exe /c</code> or <code>cmd /c</code> in front of the autofix string. If the script, batch file, or executable is not included in the Windows path, then specify the full path in the autofix command. Ensure that files or programs referenced in a script are in the same directory as the script, or specify the complete path in the script. Do not launch another GUI application from an autofix command. Doing so may cause the autofix to fail. Examples <pre>cmd.exe /c C:\utilities\cleanup.bat &METRIC &LEVEL &KEY[1] &KEY[2] &VALUE cmd /c C:\backup\delete.wsh cmd /c cp &KEY[1] C:\backup</pre>
MVS	Specify the fully-qualified dataset name of the REXX executable or CLIST. Enclose the autofix within quotes. (ControlCenter submits autofixes as TSO commands.) Place the alert value substitutions outside of the quotes. Separate multiple commands with semicolons. Examples <pre>"system.utility.cleanup" &METRIC &VALUE &KEY "test.batch.restart(rexxfix) "</pre>

Creating an Autofix Script on the Host

Use Perl, Javascript, or another scripting tool to create your autofix script on the host. Place the script in the directory that you specified when you created the autofix command in the Console.

If you plan to use the autofix on multiple hosts, you must place it in the same directory on each host, or create separate autofixes for each host. Before propagating the autofix to multiple hosts, create and test it on one host to make sure it works how you want.

Following is an example of a simple batch file autofix for a Windows host. The autofix accepts four parameters (%1 through %4) and writes them to a log file called `alertinfo.txt`.

```
@echo METRIC NAME: %1 >> alertinfo.txt
@echo LEVEL: %2 >> alertinfo.txt
@echo KEY 1: %3 >> alertinfo.txt
@echo VALUE: %4 >> alertinfo.txt
```

In ControlCenter, the corresponding autofix command might look like this:

```
cmd /c C:\Autofixes\TestAutofix.bat &METRIC &LEVEL &KEY[1] &VALUE
```

After an alert to which the autofix is attached triggers, the `alertinfo.txt` log file would contain content like this:

```
METRIC NAME: MNR.FileSize.File.Size
LEVEL: FATAL
KEY 1: MyMonitoredFile.txt
VALUE: 100000
```

Use a simple autofix such as this to test your autofix scripts before copying them to multiple hosts. Simple scripts are also helpful in parsing the output returned by ControlCenter.

Attaching an Autofix to an Alert Definition

To have an autofix run when an alert triggers, attach the autofix to the alert definition. The autofix runs when the alert first triggers and then each time the alert moves from one severity level to another.

To attach an autofix to an alert:

1. In the selection tree, expand **Administration, Alert Management, and Alert Definitions**.
2. Expand the folder for the agent that contains the alert to which you want to attach the autofix.
3. Expand the subfolders to locate the alert. Right-click the alert and select **Edit Alert**.
4. Click **Actions**.
5. Select the autofix in **Available Autofixes**.
6. Click **Add**.
7. Click **Apply To**.
8. Verify that you want the autofix to run on all the selected hosts or for the selected storage systems. The script, executable, or batch file specified in the autofix must exist in the same place on all hosts.

Best Practices for Configuring and Managing Alerts

Use the following tips to configure and manage a system of alerts and notifications:

- ◆ Disable unnecessary or redundant alerts.
- ◆ Set alert frequencies to minimize processing impact.
- ◆ Use notifications to reduce the volume of alerts.
- ◆ Create user-defined groups to organize your alerts.
- ◆ Use management policies to notify personnel.
- ◆ Modify templates to facilitate alert creation.

Disable Unnecessary or Redundant Alerts

When you install a new ControlCenter agent, a set of alerts and notifications for that agent are enabled by default. Review those alerts and notifications to ensure they are appropriate for your environment. Disable alerts and notifications for metrics that you do not want to monitor, or remove the managed object from the alert definition.

Set Alert Frequencies to Minimize Processing Impact

Each alert definition has an associated schedule that determines how often ControlCenter checks the status of an object. Use the following techniques to minimize the impact of alert and data collection policy processing on the host CPU:

- ◆ Do not configure schedules to run more often than every 5 minutes.
- ◆ Do not assign the same schedule to too many alerts on the same host. As you increase the number of objects an agent is monitoring, extend the schedule to minimize processing impact.
- ◆ Disable alerts that you do not need.
- ◆ Configure schedules to run only on days and during hours when you need notification. For example, configure the schedule to run only during normal business hours.

Consult the *EMC ControlCenter 5.2 Planning and Installation Guide, Volume 1* for duration and CPU usage statistics for data collection policies and alerts.

Use Notifications to Reduce Alert Volume

For any metric, you can enable both notifications and alerts. Notifications populate the At A Glance view, which uses charts to summarize the status of your environment or a subset of your environment, based on your selection. You can then drill down through the charts to identify specific problems. Alerts appear in the Alerts view, which lists problems in a table format.

Often, the number of active alerts in the Alerts view can seem overwhelming. Configure as notifications those metrics that do not require action. Use alerts for more critical problems that require immediate attention. For example, define the Information, Minor, and Warning levels as notifications and the Critical and Fatal levels as both notifications and alerts.

Create User-Defined Groups to Organize Your Alerts

User-defined groups allow you to organize objects (such as hosts and storage systems) into logical groups based, for example, on geographic region, business units, or personnel assignments. In views such as Alerts and At A Glance, you can then show the alerts for a specific group.

Use Management Policies to Notify Appropriate Personnel

If you do not attach a management policy to an alert definition, the associated alerts and notifications appear for all ControlCenter users. Create management policies that send alerts and notifications to the people who need to see them.

Modify Templates to Facilitate Alert Creation

Each metric has an associated template that defines default parameters for new alert definitions. Edit the templates to specify the trigger values, management policy, and schedule that all alert definitions created from that template will have by default. You can then modify the alert definitions as necessary.

Troubleshooting Alerts and Autofixes

This section provides tips for troubleshooting the following alert and autofix problems:

- ◆ Too many alerts appear in Console
- ◆ Cannot create or edit alerts or changes not saved
- ◆ Cannot clear alerts
- ◆ Alert does not trigger as expected
- ◆ Autofix does not run
- ◆ Alert count differs among users
- ◆ Managed object has warning icon but no alerts
- ◆ Alert created/modified date and time is incorrect

Too Many Alerts Appear in the Console

If too many alerts appear in your Alerts view, there are several steps you can take to reduce the display, such as applying management policies and disabling alerts that are not important to you.

Table 3-5 provides suggestions for reducing the number of alerts that display.

Table 3-5 Reducing the Number of Alerts That Display

Action	Description
Disable all alert definitions, and then enable only those critical to you.	Many alert definitions are enabled by default when you install ControlCenter components. Although we identified these alerts as important, they may not be significant in your environment or for every system. Or, the default settings may not be appropriate. Refer to <i>Enabling or disabling multiple alert definitions</i> in the online Help.
Disable unnecessary or redundant alerts.	Instead of disabling all alerts, review all enabled alerts and disable those that are not appropriate for your storage environment. Refer to the description on page 3-23 for more information.
Assign management policies.	If you do not assign a management policy to an alert definition, the alert displays for all users when it triggers. Use management policies to direct alerts to appropriate personnel. For example, assign all of the ECC Server alerts to the ControlCenter administrator. Many alert definitions are preconfigured when you install ControlCenter components. However, these alert definitions do not have management policies attached. Make sure you assign management policies to them.

Table 3-5 Reducing the Number of Alerts That Display (continued)

Action	Description
Clear resolved alerts.	After you address a triggered alert, you can remove it from the Alerts view by clearing it for all users. For a specific procedure, refer to <i>Clearing an alert whose condition has been resolved</i> in the online Help.
Use notifications to reduce the volume of alerts.	Refer to the description on page 3-24 for more information.
Filter the Alerts view.	Filter the view to show a subset of alerts, such as alerts created within the last day or of a specific type, such as Host Performance alerts. Click Filter on the view title bar to set the filter criteria.
View alerts by object or user-defined group.	View the alerts for a specific storage array, host, or other object or group of objects, such as an application group. Refer to <i>Limiting the Active Alerts That Appear</i> on page 8-6 for more information.

Cannot Create or Edit Alerts or Changes Not Saved

If you cannot create an alert, edit an alert, or save changes to an alert, check with the ControlCenter administrator to ensure your user ID has been assigned the necessary permissions.

Cannot Clear Alerts

To clear alerts from the Alerts view and notifications from the Notifications view, you must have the Clear Alert permission for the Active Alert object type. See the ControlCenter administrator to ensure your user ID has the necessary permissions.

Alert Does Not Trigger as Expected

Ensure that:

- ◆ You used correct syntax to specify the alert key (or source). For syntax rules, see the specific alert description in the online Help.
- ◆ You attached a schedule to the alert.
- ◆ Enough time passed for ControlCenter to evaluate the alert. (For example, if you attached a schedule that causes ControlCenter to evaluate the alert every hour, wait for an hour to pass.)
- ◆ The alert is enabled and you selected at least one alert severity level.
- ◆ The alert's management policy is set up to notify your Console.

- ◆ The alert's spike controlling values are configured properly. The Before field in the alert definition indicates how many consecutive times an alert must evaluate to true before ControlCenter triggers the alert.
- ◆ Another user has not cleared or removed the alert (you can review the alert history to find out).
- ◆ There are no additional requirements for the alert. See the alert description in the online Help for requirements.

Autofix Does Not Run

If your autofix command does not run when the alert it is attached to triggers, ensure that:

- ◆ You have used the proper syntax to specify the autofix command.
- ◆ Files or programs referenced in a script are in the same directory as the script or have their complete path specified in the script.
- ◆ On Windows, you have included `cmd.exe /c` in front of the command and you have specified an executable.

Alert Count Differs Among Users

The All Alerts button in the upper-right corner of the Console displays:

- ◆ Total number of alerts.
- ◆ Total number of new alerts (new alerts are alerts that have not been assigned to a user or cleared from the Alerts view).
- ◆ Severity of the alert with the highest severity and the number of alerts at that severity.

These totals may differ among users if:

- ◆ The management policy attached to an alert definition specifies that an alert be sent to one user but not another.
- ◆ The management policy attached to an alert definition has changed since a user logged on to the Console. In this case, ControlCenter applies the management policy and updates the alert count accordingly the next time the user restarts the Console.
- ◆ If a user is running multiple Consoles, the alert count can differ in those Consoles depending on any management policy changes and the last time each Console was restarted.

Managed Object has Warning Icon but No Alerts

If a storage array, host, or other managed object has a Warning icon but no alerts appear for that object in the Alerts or At A Glance views, then the object may have had alerts previously, but the Console has not updated the object's icon. In this case, clear the Warning icon by right-clicking the object and selecting **Refresh Alert Status** to update the alert status.

If, after refreshing the alert status, the problem remains, verify that the managed object does not have alerts for which your user ID is excluded from the management policy.

Make sure that you are included in the list if the management policy specifies that an alert be sent to specific users. If you are not included, you will not see the alerts in your Alerts and At A Glance views even though a Warning icon appears on managed objects for which the alert triggers. (This is the designed behavior.)

Alert Created or Modified Date and Time is Incorrect

When an alert is created (or updated) in the Alerts view on the Console, the time shown in the Alerts view for the creation (or update) of the alert is the local system time for the host running the agent for the managed object of the alert. If the local system time is incorrect, the time shown in the Alerts view will be incorrect. Maintain the correct date and time on your hosts to prevent this problem.

The *Repository* is a relational database that holds the current and historical data of both the storage environment and ControlCenter itself. This data includes configuration details about storage arrays, hosts, databases, statistical data for capacity planning, alerts, and detailed status information about any given device.

This chapter contains the following sections:

◆ Automatic Tasks	4-2
• Backing Up the Repository	4-2
• Backing Up Using Third-Party Tools	4-3
• Exporting the Repository Backup	4-4
• Analyzing Tables	4-4
• Rebuilding the Index	4-4
• Recompiling Invalid Objects.....	4-5
• Monitoring Tablespace Growth.....	4-5
• Monitoring the Status of Automated Tasks.....	4-6
• Listing Installed ControlCenter Components.....	4-6
◆ Manual Tasks	4-7
• Shutting Down the Repository	4-7
• Starting the Repository	4-8
• Scanning the Repository Alert Log.....	4-8
• Cleaning Trace Files	4-8
• Determining Tablespace Fragmentation.....	4-8
• Determining Which Processes are Currently Running.....	4-9
• Resetting the Repository	4-9
• Performing a Media Recovery.....	4-9
• Restoring the Repository	4-10
• Importing the Repository Database.....	4-11
• Gathering Data for Remote Diagnostics Assistance	4-11

Automatic Tasks

The following tasks are scheduled to run automatically at a specific time. The tasks can also be run manually at any time using the procedures provided in each section.

After you complete installation of components on the Repository host, you must reboot the Repository host to enable automatic backup.

Backing Up the Repository

Backup is a database job that runs at 2:00 a.m. for hot backup, which you can expect to complete by 6:00 a.m., everyday. However, you can backup the Repository at any time as follows:

From the **Start** menu, select: **Programs, EMC, EMC Control Center, Repository Maintenance, BackUp Database.**

The hot backup requires that the `RAMBACKUPDIR` environment variable be defined to a directory with enough free space. The variable is defined during install to whatever value is supplied. This can be changed at a later date.



CAUTION

Historical information is purged automatically after 45 days. If you want to keep more data for any trend analysis or other statistics, you can change it from the Console. Make a copy of the backup on a daily basis onto tape or offsite storage to help with recovery, as it is overwritten the next day. Please note that the backup space requirements are approximately the same as the Repository.

Refer to *Restoring the Repository* on page 4-10 for information about restoring the repository in the event the database becomes corrupted.

Backing Up Using Third-Party Tools

If you have a requirement to backup the ControlCenter infrastructure server using third-party software, such as NetBackup, EDM, etc., the following should be considered:

Backups should be scheduled around the following automatic tasks:

- ◆ Hotbackup database job that runs daily at 2:00 a.m.
- ◆ Export Backup that runs daily at 10:00 p.m.
(ramdb_export_db.bat)
- ◆ Both backup jobs are stored in the RAMBBACKUPDIR environment variable directory.

For example, scheduling third-party backup after 3:00 a.m. or 11:00 p.m. (after the automatic tasks end) allows backup of the latest files from RAMBBACKUPDIR. You can determine the exact completion time of hot backup or export backup by looking at the timestamps of %RAMBOH%\Repository\admin\Ramb_Scripts\ramb_hotback.log and %RAMBBACKUPDIR%\exp\rambdb_exp.log.

When backing up the database components, the following open files need to be excluded from the backup process:

```
%RAMBOH%\Repository\oradata\rambdb\*.CTL
%RAMBOH%\Repository\oradata\rambdb\*.LOG
%RAMBOH%\Repository\oradata\rambdb\*.DBF
```

This also applies to any third-party monitoring tools, such as BMC Patrol, as the external locks that this tool places on files are intrusive to Oracle (refer to the Oracle website, www.oracle.com, for additional information).

For recovery purposes, when backing up the database components, the following files need to be included in the backup process:

```
%RAMBBACKUPDIR%\hot\*.*
%RAMBBACKUPDIR%\exp\*.*
%RAMBOH%\Repository\oradata\rambdb\archive\*.*
%RAMBOH%\Repository\database\PWD*.*
%RAMBOH%\Repository\database\SP*.*
%RAMBOH%\Repository\admin\rambdb\bdump\*.*
%RAMBOH%\Repository\admin\rambdb\udump\*.*
%RAMBOH%\Repository\admin\rambdb\pfile\*.*
```

Exporting the Repository Backup

Export backup is a database job that runs at 10:00 p.m. every day. However, you can export the repository database any time as follows:

From the **Start** menu, select: **Programs, EMC, EMC Control Center, Repository Maintenance, Export Database.**

In order for the Export Backup to work, the RAMBBACKUPDIR environment variable needs to be defined to a directory which has sufficient free space. This variable is defined during the install to whatever the value is supplied.

You can import this data back into the Repository in the event you need to restore the Repository to a previous point in time (all existing data in the Repository will be overwritten). Refer to *Importing the Repository Database* on page 4-11 for more information.

Analyzing Tables

This process is a scheduled database job to run at 9:00 p.m., every Sunday. However, you can run Analyze Tables anytime with the following batch process:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_analyze_table.bat
```

Error files end in .err and complete output is listed in .log. The files are created in <Install_Root>\Repository\admin\Ramb_scripts.

For example, when you run:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_analyze_table.bat
```

it creates the following files:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_analyze_table.log
<Install_Root>\Repository\admin\Ramb_scripts\ramb_analyze_table.err.
```

Rebuilding the Index

This process is a scheduled database job that runs at 12:05 a.m., every Saturday. However, you can run Rebuilding Index anytime with the following batch process:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_rebuild_index.bat
```

Error files end in .err and complete output is listed in .log. The files are created in <Install_Root>\Repository\admin\Ramb_scripts.

For example, when you run:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_rebuild_index.bat
```

it creates the following files:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_rebuild_index.log  
<Install_Root>\Repository\admin\Ramb_scripts\ramb_rebuild_index.err
```

Recompiling Invalid Objects

This process is a scheduled database job that runs at 9:30 a.m., every Sunday. However, you can run Recompiling invalid objects at any time with the following batch process:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_recomp_invalid.bat
```

Error files end in .err and complete output is listed in .log. The files are created in <Install_Root>\Repository\admin\Ramb_scripts.

For example, when you run:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_recomp_invalid.bat
```

it creates the following files:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_recomp_invalid.log  
<Install_Root>\Repository\admin\Ramb_scripts\ramb_recomp_invalid.err
```

Monitoring Tablespace Growth

This process is a database job scheduled to run at 9:00 p.m., every day. However, you can run Tablespace Growth Monitoring at any time with the following batch process:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_freespace.bat
```

Error files end in .err and complete output is listed in .log. The files are created in <Install_Root>\Repository\admin\Ramb_scripts.

For example, when you run:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_freespace.bat
```

it creates the following files:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_freespace.log  
<Install_Root>\Repository\admin\Ramb_scripts\ramb_freespace.err
```

Monitoring the Status of Automated Tasks

ControlCenter runs a batch process at 8:30 a.m. every day that creates a log file with the status of every automated Repository job (not jobs you ran manually) based on the current information that Oracle has for jobs. The `ramb_jobstatus.log` file provides the name of each job, how many times it has failed (if any), the last successful run time, the next scheduled run time, as well as if the job is broken.

Check the status of the automated Repository jobs by running the following batch process and then viewing the log file:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_jobstatus.bat
```

Listing Installed ControlCenter Components

ControlCenter runs three automated tasks each morning to create log files containing lists of installed ControlCenter components. The log files are identical, except in the way that the lists are ordered.

The log files are:

- ◆ `InstalledECCCompByComp.log` — Display ControlCenter components based on components. Runs at 7 a.m. every day.
- ◆ `InstalledECCCompByVer.log` — Display ControlCenter components based on version. Runs at 7:10 a.m. every day.
- ◆ `InstalledECCCompByHost.log` — Display ControlCenter components based on host. Runs at 7:20 a.m. every day.

You can run these tasks anytime with one of the following batch processes:

```
<Install_Root>\Repository\admin\Ramb_scripts\InstalledECCCompByComp.bat
```

```
<Install_Root>\Repository\admin\Ramb_scripts\InstalledECCCompByVer.bat
```

```
<Install_Root>\Repository\admin\Ramb_scripts\InstalledECCCompByHost.bat
```

Manual Tasks

The following tasks are performed manually and should not be run without contacting your service representative:

- ◆ *Shutting Down the Repository* on page 4-7
- ◆ *Starting the Repository* on page 4-8
- ◆ *Scanning the Repository Alert Log* on page 4-8
- ◆ *Cleaning Trace Files* on page 4-8
- ◆ *Determining Tablespace Fragmentation* on page 4-8
- ◆ *Determining Which Processes are Currently Running* on page 4-9
- ◆ *Resetting the Repository* on page 4-9
- ◆ *Performing a Media Recovery* on page 4-9
- ◆ *Restoring the Repository* on page 4-10
- ◆ *Importing the Repository Database* on page 4-11
- ◆ *Gathering Data for Remote Diagnostics Assistance* on page 4-11

The scripts used to perform these tasks are located in the ControlCenter install root directory:

```
<install_root>\repository\admin\Ramb_Scripts\
```

Shutting Down the Repository

You can shut down the Repository by completing either of the following methods:

- ◆ Run:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_shutdown_db.bat
```

- ◆ On Windows, from the **Start** menu, select **Programs, Administrative Tools, Services, OracleServiceRAMBDB, Stop**.

Starting the Repository

If you shut down the Repository as described in the previous section, you can start the Repository through either one of the following methods:

During the ControlCenter installation, the Repository starts automatically.

- ◆ Run:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_start_db.bat
```

- ◆ On Windows, from the **Start** menu, select **Programs**, **Administrative tools**, **Services**, **OracleServiceRAMBDB**, **Start**.

Scanning the Repository Alert Log

To scan the Repository Alert Log, run the following:

```
findstr /I "ORA- SP2-" <Install_Root>\Repository\RDBMS\trace\rambdbALRT.log
```

Cleaning Trace Files

To clean trace files, complete the following procedure:

Be sure to rename the alert_ramdb.log file (located in the bdump directory) and save it to another location in case you need to check the log files.

1. Delete the old trace files from the following two files:

```
<Install_Root>\Repository\admin\ramdb\bdump
```

```
<Install_Root>\Repository\admin\ramdb\udump
```

2. Delete the old archive log files from:

```
<Install_Root>\Repository\oradata\ramdb\archive
```

Determining Tablespace Fragmentation

Execute the following to run the Tablespace Fragmentation task:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_tbspfrag.bat
```

Determining Which Processes are Currently Running

You can generate a log file listing the status of the processes running on the ControlCenter Repository database with the following task:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_process_info.bat
```

Resetting the Repository



CAUTION

Resetting the Repository will erase all data from the Repository! Contact EMC Customer Support for assistance.

Performing a Media Recovery

If the ECC Server cannot connect to the Repository and if you see ORA-01113 toward the end of the

<Install_Root>\Repository\admin\rambdb\bdump\alert_ramdbd b.log, run the following script to perform a media recovery on the Repository.

The script will not run properly from Terminal Services Client. Run this batch file on the Repository host itself.

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_mediarecovery.bat
```

If this procedure fails, proceed to the next section to restore the Repository.

Restoring the Repository

Use this procedure to restore a corrupted Repository database. You can restore the database to any point in time for which you have the hot backup copy and `archive.log` files (refer to *Backing Up the Repository* on page 4-2).

A two-step process is used to restore the database:

- ◆ **Perform a Media Recovery** — Usually the Repository can be repaired by running a media recovery script as outlined in the previous section.
- ◆ **Restore the Database** — If the media recovery is unsuccessful, then you can run a Repository restore script as detailed in this section.

The script will not run properly from Terminal Services Client. Run this batch file on the Repository host itself.

If you attempted to perform a media recovery and failed, use the steps in this section to restore the Repository.

This procedure requires a hot backup copy and `archive.log` files to restore the database to a specified point in time.

1. Copy the following files into a new folder (for example `Before_Restore_Files`) in case the restore is unsuccessful and you require EMC Support:

```
<Install_root>\Repository\oradata\rambdb\*.DBF
<Install_root>\Repository\oradata\rambdb\*.CTL
<Install_root>\Repository\oradata\rambdb\redo*.LOG
```

2. Copy your good hot backup `*.dbf` and `redo*.log` files to:

```
<Install_root>\Repository\oradata\rambdb
```

3. Run the following script:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_restore.bat
```

The script walks you through the restore procedure.

If this procedure does not work, contact Customer Support.

Importing the Repository Database

In some situations, you may want to import data from a previously exported database (refer to *Exporting the Repository Backup* on page 4-4). The Repository must be functioning properly to import a database. This is not a solution for resolving a corrupted database; use this procedure when you need to restore your Repository to a previously exported point in time.



CAUTION

Importing the Repository database will replace the existing database (including managed objects) with the previously exported database and its managed objects. Be sure to export and/or backup your existing database before attempting this procedure.

Import the Repository database as follows:

1. Export and/or back up the existing Repository database using the procedures outlined in *Exporting the Repository Backup* on page 4-4, and *Backing Up the Repository* on page 4-2.
2. Run the following script to import a previously exported Repository database:

```
<Install_Root>\Repository\admin\Ramb_scripts\ecc_import.bat
```

The script walks you through the import procedure.

Gathering Data for Remote Diagnostics Assistance

EMC may require system and database log files to troubleshoot a remote customer's ControlCenter issue. In this situation, EMC will request that you run the following job, which gathers system and database log information, and then creates a zip file:

```
<Install_Root>\Repository\admin\Ramb_scripts\ramb_RDA.bat
```


This section provides task-based information for EMC ControlCenter users and consists of the following chapters:

- Chapter 5, *Using the ControlCenter Console*
- Chapter 6, *Using the EMC ControlCenter Web Console*
- Chapter 7, *Managing Your SAN*
- Chapter 8, *Monitoring Storage with Alerts and Notifications*
- Chapter 9, *Monitoring and Analyzing Performance*
- Chapter 10, *Allocating or Deallocating Storage*
- Chapter 11, *Protecting Data*
- Chapter 12, *Managing Host Storage Resources*
- Chapter 13, *Using Reports*
- Chapter 14, *Incorporating Data from External Sources into Reports*
- Chapter 15, *Tuning Symmetrix Performance*

Using the ControlCenter Console

This chapter provides an overview of the ControlCenter Console as well as tips for using online Help.

This chapter consist of the following sections:

- ◆ Working in the Console.....5-2
- ◆ Console Features 5-16
- ◆ ControlCenter Online Help 5-23

Working in the Console

The Console is the ControlCenter user interface that allows you to monitor, configure, report on, and manage your SAN.

The Web Console is another ControlCenter user interface, which is primarily used for remote monitoring and reporting.

Before you can begin navigating and performing tasks in the Console, it is important that you understand the different components that make up the Console window.

Figure 5-1 identifies the basic Console components.

A tutorial on basic Console operations is available through the Console by selecting **Help**, **Quick Start Tutorials**, and then selecting **Using the Console**.




Figure 5-1 The ControlCenter Console

The **File**, **Edit**, **View**, and **Help** menus are standard menus that always display in the menu bar. The **Alerts**, **At A Glance**, **Properties**, **Topology**, **Relationship**, and **Performance** views are common views that always display in the toolbar.

Using the Console

Use the Console to perform tasks, such as monitoring, configuring, controlling, tuning, and planning storage across your storage area network (SAN).

Basic Console Steps

Using the Console involves the following basic steps (Figure 5-2):

1. Check the object that you are interested in viewing from the tree panel by clicking the appropriate checkbox.
2. Select a task environment by clicking a task button on the taskbar.
3. Select the view that you want the selected object(s) to appear in from the task button's pull-down menu, or from the toolbar.
4. Use task-associated menus, shortcut menus, and toolbar commands to perform actions on objects on the Console.




Figure 5-2 Basic Console Steps

Using the Menu Bar

Use the Console menu bar to perform (Figure 5-3):

- ◆ Standard file commands, such as **Copy**, from standard menus (the standard menus are File, Edit, View, and Help).
- ◆ Task-associated commands, such as **Zoning**, from task-associated menus (in this example, the tasks associated with Storage Allocation are Configure, Allocate, Zoning, Masking, Discover, Relationship, and Alerts).




Figure 5-3 Menu Bar

The menus that appear in the menu bar depend on the task button you click on the taskbar. When you click a task button, the menu bar changes to display both standard menus and task-associated menus.

Standard Menus

Standard menus are always available on the menu bar regardless of the task you select.

Task-Associated Menus

When you click a task button on the taskbar, the menu bar changes to display menus associated with the selected task.

Performing a Task From a Menu

To perform a task from a menu:

1. Click the appropriate task button on the taskbar.
2. Select the appropriate menu option from the menu bar.
3. From the menus that appear, select a command.

For example, to import the active zone sets, click **Storage Allocation** on the taskbar. The Storage Allocation task menus appear. From the **Zoning** menu, select **Fabric**, **Import Zone Set**, **Import Active**.

In the dialog box that appears, perform the Import Zone Sets task (Figure 5-4).




Figure 5-4 Performing a Task from a Menu

Using the Taskbar

The Console taskbar is shown in Figure 5-5.



Figure 5-5 The Taskbar

Use the Console taskbar to:

- ◆ Change the task environment to display the task-associated menus.
- ◆ Access a list of task-associated views from the task drop-down menus.
- ◆ Open or change a view in the active target panel by selecting from the list of available views.

Changing the Task Environment

To change the task environment, click a task button on the taskbar. The Console title bar updates to identify the current task environment, and the menu bar updates to provide the appropriate task-associated menus.

For example, click Storage Allocation and the title bar shown in Figure 5-6 appears:



Figure 5-6 Storage Allocation Title Bar

Accessing a List of Available Views

Click the task menu to display a list of task-associated views.

Opening or Changing a View

To open or change the view in the active target panel:

- ◆ Click the task menu, to display a list of task-associated views, and then select the view you want to open.
- ◆ Select a common view from the toolbar.
- ◆ Split the current view to open another view of the same type (default setting does not populate the view with objects).

The view opens in the target panel and displays the appropriate view information for objects selected in the tree panel.

Using the Console Toolbar

The Console toolbar is divided to three areas (Figure 5-7):

- ◆ File menu commands
- ◆ Common views
- ◆ Task-associated commands



Figure 5-7 Console Toolbar

File Menu Commands

The following file menu commands are available on the toolbar:

- ◆ Printing
- ◆ Print Preview
- ◆ Export
- ◆ Help

These commands are available regardless of the task that is selected in the taskbar.

Common Views

The following common views are available on the toolbar:

- ◆ Alerts
- ◆ At A Glance
- ◆ Properties
- ◆ Topology
- ◆ Relationship
- ◆ Performance

These views are available regardless of the task that is selected in the taskbar.

Task-Associated Commands

Task-associated commands change depending on the task selected in the taskbar.

Understanding the Information Panel

The information panel displays in the lower panel on the Console. It consists of the following two areas:

- ◆ Hints
- ◆ System and status information



Figure 5-8 Information Panel

The hint area displays dynamic information about navigating the Console. When you position the cursor over an object, it displays information on what you can do to this object. For example, click to expand <folder name>.

The system information area displays the status of objects as you position the cursor over them. It also displays the status of members of that object. Icons within this area represent the status of your system. You can click the icon of interest to view status information.

Using the Tree Panel

The Console tree panel contains a view of all the managed objects available in the monitored system.

Use the Console tree panel to:

- ◆ Access objects in the main tree folders:
 - Expand and collapse tree objects to locate objects of interest
 - Check tree objects, or drag and drop tree objects, to display in a target panel
 - Use the Find feature to quickly locate objects in the tree
 - Create user-defined groups
 - Open Task Lists to view a queue of tasks
- ◆ Perform actions on objects in the tree or in the tree view:
 - Use standard panel toolbar commands
 - Use the Action menu
 - Use the shortcut menu
 - Create user-defined groups
 - Use Task Lists

Accessing Objects in the Main Tree Folders

The tree contains the following default folders that contain the objects that you then select to view, or perform actions on:

Storage Systems — Contains all the available storage and storage devices visible throughout the network.

Storage Administration — Contains the administrative storage objects such as policies and task lists.

Hosts — Contains all the hosts and host devices available throughout the network.

Connectivity — Contains all the connectivity devices available throughout the network.

Database Instances — Contains all the database instances.

Administration — Contains all the agent information and data collection policies used throughout the network.

Status Acknowledged — Provides a location to store managed objects that have a warning or error status. The purpose of this folder is to allow you to acknowledge a status condition without being reminded of it through the parent object.

Task Lists — Provides a means to queue tasks or lists of tasks to execute now or later.




Figure 5-9 Default Tree Folders

Expanding and Collapsing Tree Items

You can expand or collapse any tree item that has a plus (+) or minus (-) sign next to its icon.




Figure 5-10 Expanding and Collapsing Tree Items

Selecting Tree Items

Use any of the methods that follow to select objects from the tree panel.

Check Objects

To view tree items in the active target view, click the checkbox so that a check appears next to the tree item.

Click a checkbox on the far left side of the tree panel to select one or more tree items. Click again to clear the checked item. Notice that the cursor changes to a checkmark for selecting and an eraser for deselecting. Click the eraser at the top of the checkbox column to deselect all checkbox selections. Clicking a folder checks all objects in that folder.

Using Double-Click

You can double-click an object in the tree panel to add that object to the target panel.

Using Right-Click

You can right-click an object in the tree panel to add that object to the target panel. From the shortcut menu, select **Add to View** to add the object to an active target panel view.

Using Highlighting

To perform an action on an object(s) without adding that object to the active target view, select the object(s) in the tree. To select an object, click any tree item name (double-click or drag and drop to make the object appear in the active target panel). Click again to clear the selection. It is possible to make multiple selections by holding the **SHIFT** key when clicking to make contiguous selections and by holding the **CTRL** key when clicking to make non contiguous selections.

Using the Target Panel

Use the Console target panel to:

- ◆ View objects and object associations within the different target panel views:
 - tables
 - maps
 - special views
- ◆ Perform actions against objects, for example, move, copy, delete, associate, group, by using:
 - the Console menu bar
 - shortcut menus
 - Console features
 - toolbar commands
 - Action menus

Navigating in the Target Panel

Figure 5-11 shows different components that make up the target panel and allow you to navigate within the panel.

The active target panel changes depending on the view you select and the objects you select in the tree.




Figure 5-11 Navigating in the Target Panel

Using the Active View

If the title bar of a target panel is your system default color and the background is white, this panel is active. This means it will automatically update its contents as you select items in the tree panel. The background of a current panel is white.

If the title bar and background of a target panel are gray, this panel is not active. Only one panel can be active at a time. A panel that is not active has frozen contents and will not update as you change your checkboxes in the tree panel (Figure 5-12).




Figure 5-12 Using the Active View

Using a Table View

Table views enable you to view object information in a table. For example, checking the Adapters folder under the Host folder in the tree panel, and then selecting the Properties view from the taskbar, generates tabulated properties data for those host adapters in the target panel (Figure 5-13).




Figure 5-13 Using the Table View

You can choose to change the way data displays in the table using the sorting multiple columns, filtering, or hide or show columns features.

You can also select objects in the table, and then perform actions against those objects. To view a list of actions that can be performed against the selected object, you can right-click the object, and a list of actions appear.

Using Map Views

If the active view in the Console target panel displays a Map button on the panel title bar, you can display objects in a map view. The map view is a pictorial rendering of objects and devices in your SAN. The advantage to displaying a map view is that makes it easy to see the relationships between objects at different levels (Figure 5-14).




Figure 5-14 Using Map Views

Using Special Views

Special views are available for certain tasks. For example, checking a **Symmetrix** array in the tree panel, and then selecting **Physical Display** from the **Monitoring** menu, generates a Symmetrix Front and Rear Views display in the target panel.



Figure 5-15 Using Special Views

Splitting a View

To split the current view into two views, select either from the panel title bar:



Splits the view horizontally



Splits the view vertically

In target panels, if you prefer to copy data from the original view into the new split view, select **Copy data** when splitting views from the **File, Console Options** menu. When this option is enabled, each time you split a view, the data is copied to the new view.

Using the Action Menu

The Action menu provides commands that are available to use within target and tree panels. The commands that appear in the Action menu change depending on what is highlighted in the active tree panel or in the active target panel.

You can access the Action menu as a pull-down menu from the panel title bar. If more than one view appears in a panel, an Action menu relevant to that view appears for each.




Figure 5-16 Using the Action Menu

Console Features

This section provides an overview of the following Console features:

- ◆ *Creating User-Defined Groups*, which follows
- ◆ *Using the Arrange By Feature* on page 5-17
- ◆ *Sorting Multiple Columns of Data* on page 5-17
- ◆ *Using the Drag and Drop Feature* on page 5-18
- ◆ *Using the Drill-Down Feature* on page 5-19
- ◆ *Using the Find Feature* on page 5-20
- ◆ *Using the Hide and Show Columns Feature* on page 5-20
- ◆ *View Preferences* on page 5-21

Creating User-Defined Groups

If you frequently access the same objects in the Console tree panel, you can create user-defined groups to organize and group those objects. Creating user-defined groups makes locating objects in the tree panel easy, and allows you to create groups based on your business needs and practices.

You can also create subfolders, or subgroups, within your user-defined groups. You cannot create subgroups under any other folder type. All user-defined groups are distinguished by an orange dot in the folder icon.

The following user-defined tasks are supported:

- ◆ Creating user-defined groups
- ◆ Adding members
- ◆ Moving members
- ◆ Removing members
- ◆ Renaming user-defined groups
- ◆ Arrange By

You must have the appropriate authorization permissions to create or modify user-defined groups.

Create a user-defined group as follows:

1. Right-click in the white space on the left-hand side of the main Console, and a submenu appears.

2. Select **New, Group**. A new group appears in the Console tree.
3. Type a name for the new user-defined group.
4. Populate the new group by dragging members (objects) into it.

You can also create a new group by right-clicking an existing group and selecting **New, Group**. The new group will be a subgroup of the group you right-clicked.

Using the Arrange By Feature

Use Arrange By to arrange objects in a tree folder as you prefer to view and navigate within them. The folder that you select to arrange determines the sort options that appear. For example, if you select to arrange by type, the objects in the folder sort according to the object's associated type. If you select to arrange by group, the objects in the folder are sorted according to group associations.

Sorting Multiple Columns of Data

The sort multiple columns feature allows you to customize the way that data appears in a table. The following sort options are available:

- ◆ Sort the order that the column data appears under each column heading as ascending or descending.
- ◆ Sort the order that data appears in the table by selecting column sort levels.
- ◆ Sort the order that table columns appear by dragging and dropping the columns where you want them.

In Figure 5-17, the columns of interest have been dragged next to each other for easy viewing. The column content is sorted in ascending order. The column sort level is indicated by the triangular shapes in the column headings. The first sort criteria is Memory, the second is Operating System, and the third is Last Update.

This sort indicates that the user is interested in finding operating systems that are running low in memory, and then determining the last time an attempt was made to update the host's data.

Memory ▲	Operating System ▲	Last Update ▲
0 B	Windows NT	Mon May 20 15:04:00 EDT 2002
255.41 MB	Windows NT	Sun Apr 28 13:04:26 EDT 2002
255.55 MB	Windows 2000	Sun Apr 28 13:01:29 EDT 2002
511.55 MB	Windows 2000	Mon Jun 03 15:48:24 EDT 2002

Figure 5-17 Sorting Multiple Columns of Data

You can choose to sort column information using either the **Sort Multiple Columns** dialog box, or by directly manipulating the table within the active view.

The View Preferences Action options allow you to save the sort order for this table view as a named view or as your default view.

Using the Drag and Drop Feature

The drag and drop feature allows you to easily move objects within the Console window for the purposes of viewing and possibly performing actions on those objects.

Use the drag and drop feature when you want to add an object from the:

- ◆ **folder** in the tree panel to another **folder** in the tree panel.
- ◆ **tree** panel to a **view** panel, or **view** panel to the **tree** panel.
- ◆ **tree** panel to a **dialog** box, or the **dialog** box to the **tree** panel.
- ◆ **dialog** box to another **dialog** box (not all dialog boxes support this feature).

You can also drag and drop columns within a table in a table view.

Example

An example of dragging and dropping includes dragging an object from one view and dropping it into another view. To view the properties of an object that is in the Command History table view, select the object, and then drag and drop the object into a Properties view.




Figure 5-18 Drag and Drop Example

If an object cannot be moved to the location you chose, the Console will not allow you to perform the drag and drop feature, indicated by a circle with a line through it.

Using the Drill-Down Feature

Use drill-down views to navigate to more object-specific information for objects appearing in the At A Glance or Free Space views.

The drill-down views and the commands available in a view are determined by both the view selected and the objects currently in that view. The objects are used as target items for the drill-down views. The drill-down view options appear as Drill Down By buttons at the bottom of the view panel. There is no limit to the number of drill-down views you can display.

Drill-down views track the views that are visited. You can revisit views you have drilled to by using the navigation controls.

Using the Find Feature

Use the Find feature to find all occurrences of:

- ◆ Object(s) within a tree, table view, or map or chart view.
- ◆ Word(s), or partial word, in a table view.

Locating Objects in a Tree

In a tree panel, you can search for objects by clicking **Find** in the tree panel title bar. Find searches the database and the tree for all matching object occurrences of the text you enter in **Search for objects named**. You can also search by object type or in user folders by selecting from Search Options.

If you click **Find**, search results are highlighted in the tree one at a time. If you click **Find All**, search results are listed in a dialog box.

Because only objects exist in the tree panel, Find in the Find panel is disabled.

Locating Objects within a Table View

In a table, you can search for objects by clicking **Find**, and then clicking the **Object** radio button. This option only finds objects already loaded in the view. Find searches the selected view for all matching occurrences of the object, and then highlights matching objects in the table.

Locating Text within a Table View

In a table, you can search for words or partial words by clicking **Find**, and then clicking the **Text** radio button in the Find panel. This option finds all matching occurrences of a string of text, and then highlights any table cells containing that string of text in the table.

Locating an Object or Text within a Map or Chart View

In a map view, you can search for objects or text by selecting **Find**, and then selecting either the **Text** or **Object** radio button in the Find panel. This option only finds objects already loaded in the view and highlights any matching objects in the map or the table.

Using the Hide and Show Columns Feature

Use the Hide and Show columns feature to select table columns that you want to appear in a table view.

Use the Sorting Multiple Columns feature to control the order that columns headings appear and sort. Use View Preferences to save your column preferences in a named view that you can later restore.

View Preferences

You can save Console window preferences and view preferences so that the next time you log in, your preferences can be restored. This feature is very helpful if you continuously return to the same Console or view image as a starting point for analysis.

Not all views support view preferences. In these cases, the menu option to save view preferences is not available.




Figure 5-19 Use Saved Preferences

Preference Categories

There are three preference categories:

- ◆ **System defaults** — Console and view preferences preset by EMC as the default preferences (Restoring Console system default preferences).
- ◆ **User defaults** — Console and view preferences that you save, and that are associated with your user id (Saving Console window preferences and Saving default user view preferences).
- ◆ **Named view preferences** — View preferences that you save, and that are associated with a name. Named preferences allow you to create and restore several preferred views (Saving named view preferences).

With the exception of the Topology view, saved preferences do not include object selections in the tree panel, or objects populated in the view panels.

Preferences Console user preferences that you can save and restore include Console window configurations and view panel configurations:

Console Preferences

- ◆ Number of open Console windows
- ◆ Position and size of the Console windows
- ◆ Selected target view panels
- ◆ Window divider locations and orientations for both the tree and target panels

View Preferences

◆ **View**

- View selections
- Filtering options created for a select views
- Object selections displayed in view (Topology view only)
- Active view (when toggling between views)

Map view — Including layout, collapsing or expanding nodes, filters, show overview, zoom factor, node arrangement, and map arrangement

Table view — Including column width, sorting, filtering, ordering, and hide or show columns

- ◆ **Tree panel** — Tree object selections (for Topology view only)
- ◆ **Toolbar panel**
 - Size of the panel
 - Collapsed or expanded state of the toolbar

ControlCenter Online Help

EMC ControlCenter provides a comprehensive Java-based Help system. You access Help from the toolbar by selecting **Help, Topics**. The Help Topic window appears.

Using the Help Contents Tab

Click the **Contents** tab in the Navigator window to display the table-of-contents tree.

Browsing the Tree

Browse the tree by:

- ◆ Double-clicking a closed book to expand the next lower level of topics and books, or
- ◆ Selecting a book and choosing **Expand**, **Expand All**, **Collapse**, or **Collapse All** from the **File** menu.

Displaying a Topic From the Tree

Display a topic from the tree in the Help Topic window by:

- ◆ Double-clicking a topic, or
- ◆ Selecting a topic, then clicking the **Display** icon located in the Help toolbar.

To display a topic in a new window, select a topic, and then click the **Display in New Window** icon located in the Help toolbar.

This chapter provides an overview of the EMC ControlCenter Web Console (Web Console) as well as tips for using online Help.

This chapter consist of the following sections:

- ◆ Working in the Web Console 6-2
- ◆ Web Console Interface 6-3
- ◆ Web Console Views.....6-7
- ◆ Web Console Tutorial and Online Help 6-15
- ◆ Web Console FAQs.....6-16

Working in the Web Console

The EMC ControlCenter Web Console uses data stored in the Repository to monitor your storage-attached network and manage ControlCenter alerts remotely through a Web browser.

The Web Console provides access to the same ControlCenter Repository as the ControlCenter Console. However, the Web Console:

- ◆ Can be accessed without installation on the host.
- ◆ Has an improved interface design for ease of use, without creating a large learning curve for users comfortable with the ControlCenter Console.

The Web Console does not provide all of the administrative and management functions as the ControlCenter Console.

Accessing the Web Console

The following URL is the default for accessing the Web Console:

`http://<hostname>:7070/eccweb`

where *hostname* is the name of the host running the Web Console server.

If your network is configured with Secure Sockets Layer (SSL), you may be required to enter a different URL and port. Check with your system administrator for more information.

Accessing the Web Console with Pop-up Blockers

If pop-up blockers are installed on the host from which you are opening the Web Console, the Web Console may not open in your browser.

To open the Web Console with a pop-up blocker installed on the host, do one of the following:

- ◆ Disable the pop-up blocker program before you launch the Web Console.
- ◆ If available, use the pop-up blocker's unblocking feature while launching the Web Console. Refer to the pop-up blocker's documentation for instructions.

Pop-up blockers may also prevent you from accessing the Web Console online help. Use the same procedures you used to enable access to the Web Console to access the Web Console help.

Web Console Interface

Figure 6-1 describes the different areas in the Web Console interface.

A tutorial covering the differences between the Web Console and the ControlCenter Console and when to use the Web Console is accessible from the Web Console **Help** menu, **Tutorial** option.




Figure 6-1 The Web Console

Menu Bar

Table 6-1 provides a list and description of the menu bar options.

Table 6-1 **Menu Bar Options**

Option	Description
File menu	Provides commands that can be used anytime during your Web Console session. From the File menu you can Print, Export, Launch another Web Console session, and log out of the Web Console.
Alert menu	Provides all the Alert management commands that can be performed from the Web Console. The Alert menu is only functional when the Alert view is open.
Help menu	Provides commands for accessing the Web Console Help, Tutorial, Session information, and About dialog box.
Print	Allows you print the tree or the currently open view. If no view is open the tree is printed by default.
Export	Allows you to export data that displays within an open view or tree view, as a .csv, jpeg, or HTML file.
Logout	Logs you out of the Web Console.
New Alerts	Allows you to open the alerts view with all of the most recently issued alerts.

Tree View

The Web Console tree view contains a view of all the managed objects available in the monitored system.

Tree Folders

The tree contains the following default folders:




Figure 6-2 Web Console Default Folders

- ◆ **Storage Systems** — Contains all the available storage and storage devices visible throughout the network.
- ◆ **Hosts** — Contains all the hosts and host devices available throughout the network.
- ◆ **Connectivity** — Contains all the connectivity devices available throughout the network.

In the Web Console, **unidentified ports** have been split into two groups of ports, unassigned and unidentified. Unassigned ports are not associated with a fabric. Unidentified ports are associated with a fabric.

The ports are organized in different locations in the Connectivity folder. Unassigned ports are located in the **Connectivity>Unassigned Ports** folder.

Unidentified ports are located in the tree view in the **Connectivity>Fabric>fabricname>Unidentified Ports** folder or the **Connectivity>Fabric>Ciscophysicalfabricname>VSANname>Unidentified Ports** folder.

- ◆ **Databases** — Contains all the database instances.
- ◆ **Agents** — Contains a list of all the types of agents installed in the ControlCenter environment.

Storage Administration and Administration folders are not available in the Web Console. Storage administration and other administration tasks cannot be performed through the Web Console. Use the ControlCenter Console to access the folders and perform the associated tasks.

Tree View Actions

You can perform the following actions from the Web Console tree:

- ◆ Expand and collapse tree objects to locate objects of interest.
- ◆ Select a folder or tree object to include in a view.

In the Web Console, only one folder or object can be selected in the tree at a time.

- ◆ Right-click a folder and select the **Arrange by** option to arrange the folder objects by name or type.
- ◆ Right-click a folder or object in the tree and select **Show**, then a view, to open a view for the selected folder or object.

If an item is not available for a view, it will not be displayed in the view after it is selected.

- ◆ Right-click an object and select **Element Manager** to launch the vendor Web management application or site.
- ◆ Use the Search feature to quickly locate objects in the tree.
- ◆ View user-defined groups.

The Web Console presents user-defined groups that were created in the ControlCenter Console. User-defined groups cannot be created from the Web Console.

View Bar

The view bar provides buttons to open the views available in the Web Console and a previous and next button for navigating between the views.

In the Web Console, the Command History option is provided in the view bar. In the ControlCenter Console, Command History is accessed through the Administration tasks' drop-down list of views.

Web Console Views

Web Console views are used to display the Repository data. Different views are provided for monitoring the SAN and alert management.

Populating Views

By default, the Web Console opens with the Properties view displayed unpopulated.

- ◆ Populate the Properties view by clicking a folder or object in the tree.
- ◆ Change a view by clicking another view in the view bar. The view changes, but still contains the data for the last object selected in the tree.
- ◆ Change the content of a view by selecting another object or folder in the tree.

Navigating views

Multiple views cannot be displayed simultaneously in the Web Console. Only one view can be displayed at a time. However, you can use the previous and next buttons in the view bar to navigate through different views or you can open another browser to look at multiple views at the same time.

There is no way to clear a view. A view or the data in the view can be changed, but once a view is populated all the views that follow are populated with the data for the object most recently selected in the tree.

Common View Actions

A shortcut menu is available with all the Web Console views. The shortcut menu provides access to common view actions that are directed at the object selected in the view. The view shortcut menu provides the following options:

The following options are not enabled for every view.

- ◆ **Show** — Changes to another view for the object selected in the view.
- ◆ **Element Manager** — Launches the vendor Web management application or site.
- ◆ **Arrange by** — Arranges the contained objects by name. This option is only enabled for map views.

- ◆ **Acknowledge** — Only enabled for the Alert view.
- ◆ **Alerts** — Only enabled for the Alert view.

Types of Views

Table views

There are three types of views: tables, maps, and charts.

Properties, Alerts, and Command History are presented in tables.

The Web Console provides different layout options for the table views. Each layout presents different information for a specific table view. For example, the Alerts view has a Basic and Details layout. The Basic layout provides a subset of the columns provided in the Details layout. In the Properties view, layouts are used to display different sets of data for the object selected in the tree. For example, when looking at the properties for all Symmetrix arrays, the Basic layout displays the general system information about each Symmetrix array in the folder (refer to the *Properties View Table* below), while the Symmetrix Allocation layout provides the storage allocation details for each Symmetrix array.

EMC ControlCenter (losbe138:7070) - Microsoft Internet Explorer provided by EMC

EMC ControlCenter V 5.2.0.0.0

File Alerts Help Print Export Logout Alerts Loading...

Properties Relationship Command History Alerts Topology Performance

Search

Storage Systems (by Type)

- CELESTRA CNS
- CLARiiON
- ESS
- Hitachi-based
- HP StorageWorks
- HP XP
- Network Appliance
- SMI Discovared
- Symmetrix

Hosts (by Type & OS)

Connectivity

Databases (by Type)

Agents (by Name)

Linux Hosts

Test Root Group

Database user_group

Properties - Symmetrix Layout Basic

Array	S/N	Type	Model	Vendor	Configured(GB)
000000006223	000000006223	Symmetrix	DMX2000P	EMC	274.7
000182601265	000182601265	Symmetrix		EMC	
000183502217	000183502217	Symmetrix	3830	EMC	224.6
000183600358	000183600358	Symmetrix		EMC	
000184600314	000184600314	Symmetrix	8130	EMC	201.3
000187900611	000187900611	Symmetrix	DMX800	EMC	3,554.2
002806000215 (test)	002806000215	Symmetrix		EMC	

->Sent:1399,received:85950

Local intranet

Figure 6-3 Properties View Table

Customizing Tables

Tables can be customized in the following ways:

Table preferences cannot be saved in the Web Console. Changes to the table are lost once you open a new table.

- ◆ Rearrange the order of the columns.

Drag and drop the column to the position in the table where you want it.

- ◆ Sort the table by values in a particular column.

Click a column header. An up arrow appears in the column header the first time you click it. The up arrow indicates that the current sort is descending. Columns are first sorted numerically and then alphabetically.


Click the column again. A down arrow appears in the column header. The down arrow indicates that the current sort is ascending.

The sort feature is not available for all table columns. To find out if a column supports the sort feature, hold the cursor over the column header. If the cursor changes to a pointing hand, then the column can be sorted.

Properties Views Tables

The Web Console provides different types of Properties views tables for easy navigation and investigation of a system's components and resources. The following types of tables are provided in the Properties views:


- ◆ **Tables** — Present the object properties in a single table. It contains no tabs or trees and requires no navigation.
- ◆ **Tabbed tables** — Present the properties of the object selected in the tree and provide tabs to tables that display properties of associated objects. Tab tables may contain tables and tree tables.



Properties - Symmetrix 000000006223	
Array	000000006223
S/N	000000006223
Type	Symm6
Model	DMX2000P
Vendor	EMC
Configured(GB)	274.78
Unconfigured(GB)	1772.89
Raw(GB)	2187.72
Local Replica Allocated(GB)	NC
Allocated(GB)	

Figure 6-4 Properties Views Tabbed Tables

- ◆ **Tree tables** — Present the object properties and a drill-down feature for more information. A tree table is identified by the expandable objects contained in the table.



Properties - Symmetrix 000000006223						
General	Associated Groups (1)	Host Directors (7)	Mapped Devices (196)	Unmapped Devices (179)		
	Status	# Ports	# Devices	WWN	Type	
FA-14B	Online	2 NC				
Port 0			NC	50060480000613dd	SymFibrePort VCM Enabled	
Port 1			NC	50060480000613fd	SymFibrePort Unique WWN	
FA-3D	Online	2 NC				
Port 0			NC	50060488000613d2	SymFibrePort VCM Enabled	
Port 1			NC	50060488000613f2	SymFibrePort VCM Enabled	
FA-3C	Online	2 NC				
FA-14C	Online	2 NC				
FA-14A	Online	2 NC				

Figure 6-5 Properties View Tree Tables

- ◆ **Detail table views** — Present the details of an object selected in a table below the table. A detail view is identified by a bar below the table, which is used to split the view in the right pane. If there is no bar below the table, then no detail view is available for the objects in the table.

Properties - Symmetrix 000000006223

General Associated Groups (1) Host Directors (7) Mapped Devices (196) Unmapped Devices (179)

	Status	# Ports	# Devices	WWN	Type	
FA-14B	Online	2	NC			
Port 0			NC	50060480000613dd	SymFibrePort VCM Enabled	
Port 1			NC	50060480000613fd	SymFibrePort Unique WWN	
FA-3D	Online	2	NC			
Port 0			NC	50060488000613d2	SymFibrePort VCM Enabled	
Port 1			NC	50060488000613f2	SymFibrePort VCM Enabled	
FA-3C	Online	2	NC			
FA-14C	Online	2	NC			

Devices for: FA-3D > Port 0

Device	Name	Status	Service State	Size(GB)	Configuration	# of Map
000	000	Write Disabled	Normal	0.02	2-Way Mir	
029	029	Ready	Normal	0.003	Unprotected	
02A	02A	Ready	Normal	0.003	Unprotected	
02B	02B	Ready	Normal	0.003	Unprotected	

Figure 6-6 Properties View Split Table View

Map Views Relationships and Topology are displayed in maps.

The map view is a pictorial rendering of objects and devices in your SAN. Map views makes it easy to see the relationships between objects at different levels.




Figure 6-7 Topology View Map

Chart Views

The last 24 hours of Symmetrix array or Fibre Channel Connectivity device performance is displayed in charts in the Web Console Performance view. The charts provide a graphical picture of the system performance. You define the content of the charts by selecting objects and statistics to include.

There is no Performance table in the Web Console Performance view. Use the Time Marker to find a metric value at a specific time.




Figure 6-8 Performance View Chart

In addition to creating performance charts, the Web Console Performance view allows you to:

- ◆ Select multiple objects to include in the chart.

If you select a folder in the tree, a list of the folder objects opens in the Performance view. Use the CTRL key to select multiple objects to chart.

The Performance view is the only view in the Web Console that supports multiple selection of objects.

- ◆ Get context-sensitive help for the selected metric.

Click the question mark “?” next to the metric list and a description of the selected metric opens.

- ◆ Specify a time range to chart.

If you select a time other than real time, the chart displays the time range of 3 hours before the time and 1 hour after. For example, if you selected 3 p.m., the chart X-axis would be from 12:00 p.m. - 4:00 p.m.

- ◆ Set a maximum value for the Y-axis.

Click the Max Limit button and enter the maximum value for the Y-axis. The Y-axis resizes to the maximum value entered.

- ◆ Remove a line from the chart.

From the chart legend, click the eraser next to the line to remove from the chart.

- ◆ Find a metric value at a specific time

Slide the chart Time Marker to a specific time on the X-axis. The value is displayed in the chart legend Value column.

- ◆ Scale a line value

From the chart legend Scale column, select a value by which you are scaling the line. The line moves to a location in the graph that represents how you've scaled.

The actual values of the metric are the values displayed in the chart multiplied by the scale factor in the corresponding row in the legend table.

Web Console Tutorial and Online Help

The Web Console is provided with a tutorial and help system.

Web Console Tutorial

The tutorial is accessed from the **Help** menu, **Web Console Tutorial** option. It discusses how to access the Web Console, the differences between the Web Console and the ControlCenter Console, and when to use the Web Console.

The Web Console tutorial requires the Flash plug-in for your browser.

Web Console Help

The help is a WebHelp system opened in your system browser. It supports the Web navigation tools provided with your browser as well as providing Contents, Index, and Search panes for easy navigation.

Access the help from the:

- ◆ **Help** menu, **Contents**, **Index**, and **Search** options to review general information about the Web Console or to access the help Contents, Index, Search, and Glossary.
- ◆ View title bar, click the question mark “?” to open a context-sensitive help topic that provides information about the open view.
- ◆ Alerts view, right-click a specific alert in the table, and select **Alerts, Help**. A context-sensitive help topic opens for the selected alert. The alert help topics describe how to respond to the alert in addition to other pertinent information about the selected alert.
- ◆ Dialog box, click the **Help** button to open a context-sensitive help topic for the dialog box. The dialog box help provides details about the options provided in the dialog box.
- ◆ **How Do I...** links within the help system to open another topic that provides needed information or explains how to perform a task. Click a link in the **How Do I...** column in a help topic.

The “No Context-sensitive Help Available” topic opens if a context-sensitive help topic is missing or is linked incorrectly to the application. Please follow the steps described in the topic to get the information you need, or contact Customer Support about the issue.

Web Console FAQs

This topic provides answers to the most frequently asked questions.

What functionality is provided with the Web Console?

This version of the Web Console provides partial monitoring and alert management functionality.

What are the browser requirements for using the Web Console?

The Web Console can run on the following browsers with JRE 1.4.2 installed as a plug-in:

- ◆ Internet Explorer Version 5.5 and higher
- ◆ Netscape 6.2.3 and higher

When I try to launch the Web Console from my browser or shortcut, nothing happens?

This will happen if you have a pop-up blocker enabled on your browser. Disable the pop-up blocker or use the blocker's unblocking mechanism while accessing the Web Console.

You may have to do the same when accessing the help from the Web Console.

When I try to access the Web Console help or tutorial nothing happens?

This will happen if you have a pop-up blocker enabled on your browser. Disable the pop-up blocker or use the blocker's unblocking mechanism while accessing the Web Console help.

How do I access my licensed ControlCenter applications such as TimeFinder, SRDF, Path Details, StorageScope, SDM, or Workload Analyzer?

The applications are accessed from the ControlCenter Console and are not accessible through this version of the Web Console.

Why is the alert status not displayed on the objects in the tree?

The functionality is not supported in this version of the Web Console. However, the object alert status is included in the tooltip that appears when you position the cursor over the object in the tree.

Why does the alert status button in the menu bar not show the status of the new alerts?

The functionality is not supported in this version of the Web Console.

Can I select an object in the view and go to another view?

Yes, right-click any managed object in a view and select the view you want to open with the data for the selected object.

A managed object is identified in a Web Console view by an icon with the object name.

Can I select multiple objects in a view and open another view?

No, you can only open a view for one object or folder at a time.

How do I look at multiple views?

The Web Console does not support opening multiple views in a single Web Console session. Only one view can be displayed at a time. However, you can use the previous and next buttons in the view bar to navigate through different views or you can open another browser to look at multiple views at the same time.

How do I clear a view?

There is no way to clear a view. A view or the data in the view can be changed, but once a view is populated, all the views that follow are populated with the data for the object most recently selected in the tree.

How do I save a view?

All populated views are saved for as long as the Web Console session is running. Use the previous and next arrows to locate a saved view. No views are saved after the Web Console session ends.

How do I create my own layouts?

You can not create custom layouts. The functionality is not supported in this version of the Web Console.

Why can I not show all the commands at once in Command History?

The functionality is not supported in this version of the Web Console. The Web Console Command History view displays the commands for the selected object.

Managing Your SAN

This chapter provides information for managing your SAN and consists of the following sections:

◆ EMC ControlCenter SAN Manager Overview	7-2
◆ Discovery and Monitoring Requirements	7-4
◆ Discovery.....	7-5
◆ Topology View	7-13
◆ Topology Edit Service (TES)	7-19
◆ Viewing the Login History	7-21
◆ Zoning.....	7-22
◆ Monitoring Statistics.....	7-38
◆ Masking	7-39
◆ Path Details View	7-50

EMC ControlCenter SAN Manager Overview

The primary purpose of a storage area network (SAN) is the transfer of data between computer systems and storage elements and among storage elements. A SAN consists of a communication infrastructure, which provides physical connections (both Fibre Channel and iSCSI), and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. A SAN typically consists of storage elements, storage devices, computer systems, and/or appliances, plus all control software, communicating over a network.

Managing a SAN

As storage networks increase in size, it can become challenging to manage the network environment. Some of the problems businesses face deploying and maintaining these networks are:

- ◆ Managing access to storage
- ◆ Zoning the network (managed differently by each switch vendor)
- ◆ Managing hardware and software products from multiple vendors
- ◆ Lack of trained IT staff and storage administrators to manage SAN environments
- ◆ Limited tools for measuring and analyzing the productivity of storage networks

Through SAN Manager, EMC offers a comprehensive set of management software tools that support the entire switched fabric infrastructure.

Features

EMC SAN Manager, a component of EMC ControlCenter, provides SAN discovery, topology representation, monitoring, performance, and active LUN masking features.

Key attributes in the EMC SAN Manager include:

- ◆ Discovery and monitoring of storage networks and their components to detect and respond to error conditions.
- ◆ Assisted Discovery feature to assist ControlCenter agents that do not perform automated discovery.
- ◆ Ability to manage and zone McDATA/Connectrix®, Brocade, Computer Network Technology, Corp. (CNT), and Cisco switches and fabrics, as well as inter operability fabrics, including the ability to import active and inactive zoning during discovery.
- ◆ Thresholds for connectivity device port performance monitoring, event notification, policy based management, and auditing the history of user actions.
- ◆ Option to launch the Brocade and Cisco fabric managers, McDATA Connectrix Manager, and EMC Invista instances.
- ◆ Ability to run ControlCenter with a Brocade Secure FabricOS and McDATA Enterprise Fabric Mode/SANtegrity fabric.
- ◆ Topology view for displaying a pictorial view of the entire SAN, including user-defined objects and groups, zone/zoneset highlighting, and toggling on and off links belonging to a selected object.
- ◆ Masking, a feature that manages host port access privileges to logical volumes in storage arrays.
- ◆ Enhanced user security, including the ability to create user-defined roles and privileges.
- ◆ Command line interfaces for storage device masking on Windows, Windows 2000, Solaris, HP, and AIX.

Discovery and Monitoring Requirements

The following agents must be installed and running before you can use ControlCenter to discover, monitor, or set up devices in the SAN.

Table 7-1 Required Agents

Agent	Comments
Master Agent	Manages all agents. Typically installed on each host during ControlCenter installation. Cannot be installed through the Console.
Fibre Channel Connectivity Agent	Discovers and monitors connectivity of Fibre Channel switches and other connectivity devices. Can be installed on any host connected to the ECC Server, with no more than one per host. Typically, one instance of the Fibre Channel Connectivity Agent is required within each subnet managed by the ECC Server.
EMC Invista Agent	Monitors EMC Invista instances. Run one agent on each client host.
Host Agents	Helps manage a host's storage activities. Install one or more through the Console.
Storage Agents	Manage storage arrays. Install one or more on your computer through the Console.
Symmetrix SDM Agent	Monitors volume-access control in Symmetrix arrays. Install on your computer through the Console.
Database Agent for Oracle	Monitors and manages Oracle database instances.
Database Agent for DB2	Monitors database instances (not the ECC Repository). Install on hosts where your databases are installed.
Common Mapping Agent	Discovers and monitors Informix, Sybase, SQL Server, and DB2/UDB databases.
Backup Agent for EDM	Monitors EMC Data Manager (EDM).
Backup Agent for Legato® NetWorker	Monitors Legato NetWorker.
Backup Agent for NetBackup	Monitors NetBackup.
Vendor Agents	Refer to documentation provided by device vendors.

It may be necessary for the ControlCenter administrator to edit the data collection policies associated with some of these agents before the agent will function in the storage network.

Discovery

Once you install and start the various ControlCenter components (Repository, Store, Console, ECC Server, Master Agent, and ControlCenter Agents), you must *discover* the topology of the SAN before you can begin to view, monitor, and manage it.

Discovery is the process of using ControlCenter agents to find, identify, and represent the various elements in your SAN and the relationships among them. These elements include hosts and storage arrays and their ports and devices; connectivity devices and their ports and links; file system and database objects; and relational elements such as fabrics, zones and zone sets, groups, and so on.

Different ControlCenter agents discover different elements of the SAN in different ways. Some agents discover devices automatically; others require user involvement. Discovery by these varied agents can occur in any sequence, but the results may vary depending on that sequence.

This section describes the following types of discovery:

- ◆ Automatic Discovery
- ◆ Connectivity Device Discovery
- ◆ Assisted Discovery

All of the steps for logging in to ControlCenter, installing agents, discovering topology, etc., are covered extensively in the *EMC ControlCenter 5.2 Planning and Installation Guide, Volume 1* and/or the ControlCenter online Help.

Automatic Discovery

Once they are installed and running, ControlCenter's host agents and the Storage Agent for Symmetrix discover hosts and Symmetrix storage arrays automatically, as long as their Discovery data collection policies are enabled. (All such policies are enabled by default except the Storage Agent for Symmetrix Proxy discovery policy, which enables discovery of Symmetrix arrays on hosts other than the one on which the agent is running.) Once an agent discovers a device, ControlCenter automatically continues to monitor it.

The Console online Help provides detailed information on configuring agents for discovering and monitoring objects in a SAN. Refer to the following online Help sections for more information:

- ◆ **Agent Overviews** — Refer to *ControlCenter agents* under *Introducing EMC ControlCenter*
- ◆ **Agent Administration** — Refer to *Installing ControlCenter components and agents*
- ◆ **Agent Data Collection Policies** — Refer to *Data Collection Policy descriptions*

Connectivity Device Discovery

Connectivity device discovery is performed by the Fibre Channel Connectivity Agent (FCC Agent). Some or all of the following information is discovered during connectivity device discovery, depending on the options selected in the Search for Connectivity Devices dialog box:

- ◆ Switch properties — hardware revision, software revision, user-assigned information, and so on
- ◆ Switch topology information — switch ports, links, information about neighboring ports
- ◆ Fabric information — including other switches in the same fabric
- ◆ Zoning information
- ◆ Nicknames, if any
- ◆ Connectrix AP-7420B and Brocade SilkWorm AP7420 Multiprotocol Router — limited discovery, alerts and visualization. Active zoning control, statistics, or fabric discovery are not supported.
- ◆ EMC Invista — properties information
- ◆ Other connectivity devices — basic device properties

You can discover connectivity devices at any time by selecting Discover, Connectivity in the Console Monitoring task menu bar, launching the Search for Connectivity Devices dialog box.

Some switches cannot be discovered without additional connectivity information supplied by the user. Refer to *Assisted Discovery* on page 7-8.

The FCC Agent performs the following operations after discovery:

1. Monitors the condition (status, attributes, and configuration) of all discovered connectivity devices and ports, generating alerts when the condition of a device, port, or fabric changes, and updating the console with topology information such as changes to neighbors, links, and fabric configurations. Active zoning configurations can also be monitored and updated if this option is set in the Fabric Validation data collection policy.
2. Monitors and updates information about the connection settings of connectivity devices you have discovered—management IP addresses, usernames, passwords, and so on—and displays this information in the Console. (To update this information when it changes, you must rediscover the device.)
3. Collects performance, frame flow, error, and operational statistics for connectivity device ports, generating alerts when specified thresholds are reached.

Monitoring Connectivity Device Port Statistics

ControlCenter supports the following methods for monitoring connectivity device port performance:

- ◆ Using Performance view
- ◆ Using performance alerts

Using Performance View

You can monitor the statistical data from connectivity device ports by viewing collected data in the Performance view. Real-time or historical data can be viewed in charts or tabular format. (You can display tables for multiple ports at the same time, but charts for only one port at a time.) To access this information, click Performance in the Console toolbar to open Performance view in the target panel. Then, in the tree panel, click the checkbox beside the ports you want to monitor. The data appears in the target panel.

Using Performance Alerts

Set connectivity device port alert thresholds for each monitored statistic in the Fibre Channel Connectivity Agent user-configurable alerts. Refer to the online Help topic *Fibre Channel Connectivity Agent alerts* under *Alert descriptions, Fibre Channel Connectivity Agent* for a list of connectivity device port statistical alerts.

To set a switch statistical alert:

1. In the tree panel, expand **Administration, Alert Management, Alert Templates, Fibre Channel Connectivity Agent, alerts, stats**.
2. Right-click an alert in the **connectivityport** folder and select **New**. The Alert Definition dialog box opens, in which you configure the alert and assign it to a particular switch.

Assisted Discovery

Assisted Discovery allows the storage administrator to guide ControlCenter agents to discover objects that cannot be found by other discovery methods. Some agents discover objects automatically, but need assistance in special instances, while some cannot discover any objects without guidance from a ControlCenter administrator. The following agents require Assisted Discovery to discover the following objects:

- ◆ **Storage Agent for Centera** — EMC Centera Content Addressed Storage
- ◆ **Storage Agent for NAS** — EMC Celerra® Network Servers, Network Appliance Filers
- ◆ **Storage Agent for CLARiiON** — EMC CLARiiON storage arrays
- ◆ **Storage Agent for ESS** — IBM ESS arrays
- ◆ **Storage Agent for HDS** — HDS 9500V and 9900V series arrays, HP StorageWorks XP arrays, and SUN StoreEdge 9900V series arrays
- ◆ **Storage Agent for HP StorageWorks** — StorageWorks HSG80 storage arrays
- ◆ **Storage Agent for SMI** — SMI-enabled storage arrays and tape libraries
- ◆ **EMC Invista Agent** - EMC Invista instances
- ◆ **Database Agent for Oracle** — Oracle database objects on the host on which the agent is running, or on a host accessible from that host
- ◆ **Common Mapping Agent** — Database objects and host objects

Assisted discovery does not apply to connectivity devices.

Discovering Objects

Assisted Discovery consists of four dialog boxes that allow you to initiate discovery, monitor the results of discovery, and modify the set of agents that have permission to manage discovered objects. The discovery settings you enter into the *Discover Other Objects* dialog box generate discover commands sent to the ECC Server to allow discovery of one or more objects, and provide access to objects when adding or removing agents with permission to manage them. Discovered objects appear in the relevant directories in the tree panel.

The same agent type is typically installed on multiple hosts. You must specify which instance of the agent (host-based) you wish to execute the discovery command.

To discover objects using the agents that require Assisted Discovery:

1. Make sure that all agents to receive a command through Assisted Discovery are running.
2. Select **Monitoring, Discover, Assisted ...**

The *Discover Other Objects* dialog box appears. Use this dialog box to select the appropriate panel for the object type you want to discover and further select (if required) the specific device model(s).

3. Click **Apply** or **OK**.

Reviewing Assisted Discovery Results

You can review the results of all discover commands sent from all Console sessions since the ControlCenter infrastructure was installed, plus all commands (discover and nondiscover) sent by the current local Console session in the *Review Discovery Results* dialog box. Stored discovery settings for commands that discovered no objects can be deleted in this dialog box while discover commands for other objects are in progress.

To review the results of Assisted Discovery operations:

1. Select **Monitoring, Discover, Review Progress...**

The *Review Discovery Results* dialog box appears.

2. Select a command or discovery settings row in the Result Summaries table. The command or discovery settings details appear in the Assistance Result Details panel.

Deleting an Object and Preventing Rediscovery

In some cases, you delete an object from the Console with the expectation that it will eventually be rediscovered, but in other cases you may want to prevent rediscovery of the object. You prevent rediscovery of an object by deleting the stored discovery settings that were specified by the user during the original discovery command.

You can delete discovery settings only if they are not currently used to grant permission to manage any object. If the discovery settings are currently being used, you must first remove management permission for all objects discovered by these settings as outlined in step 3 below.

Note: Be sure you want to permanently delete an object before using this procedure.

To delete an object and prevent rediscovery:

1. In the Console Monitoring task menu bar, select **Discover, Review Progress**. The Review Discovery Results dialog box appears.
2. In the dialog box, select the row that contains stored discovery settings of the object you want to delete. If the value in the # Objects column is zero, skip to step 7. If the value in the # Objects column is greater than zero, close the dialog box and proceed to step 3.
3. Right-click the object in tree panel and select **Discover, Remove Permission to Manage**. The Select Discovery Settings dialog box appears.
4. Select the discovery settings you checked in the previous step (you may have to repeat this procedure for more than one agent type) and click **Remove Permission to Manage**. The Remove Permission to Manage dialog box appears.
5. Select all of the active, enabled agents and click **OK** to remove the agent(s) from those with permission to manage the object.
6. Return to the Review Discovery Results dialog box. Select the stored discovery settings you want to delete and see if the value in the # Objects column is zero. If so, proceed to step 7. If the value in the # Objects column is greater than zero, refresh the Review Discovery Results dialog box by closing and reopening it.

If this does not reduce the object count to zero, look at the list of object canonical names in the details panel; these names tell you which discovered object is still able to be managed by some agent. Then, return to step 3 and again invoke **Remove Permission to Manage** for the specified object.

Note: Before the value in the # Objects column can change to zero, the Remove Permission to Manage command executed in Step 4 must finish running and the Review Discovery Results dialog box must be refreshed.

7. When the value in the # Objects column for the stored discovery settings is zero, select the row and click the **Delete** button, which is now enabled. The stored discovery settings are deleted and can no longer be sent to any agent.
8. Stop the agents that discovered the objects you want to delete. See Starting and stopping agents for more information on stopping and starting agents.
9. Delete the objects in question from the Console and restart the agents, if desired.

Note: To avoid resource impact to the server, delete one managed object at a time. Wait until you receive a delete completed notification before deleting the second object.

Granting and Denying Permission to Manage

The *Add Permission to Manage <object>* and *Remove Permission Managing <object>* dialog boxes allow you to add or remove agents to or from the set of agents that have permission to manage an object.

Once an object has been discovered by an agent through Assisted Discovery, that agent has permission to manage the object using the discovery settings you provided for discovery. But multiple active agents can potentially manage a discovered object. When more than one agent has permission to manage an object, the ECC Server selects one of them for execution of management policies such as data collection. (If no active agents have permission to manage a discovered object, the ECC Server cannot execute management policies on the object, but it remains in the Repository.)

When you grant an agent permission to manage using stored discovery settings, it can manage not only the object or objects that were discovered when those settings were used in a discover command, but also any additional discovered objects to which those settings provide access.

To grant or deny one or more agents permission to manage an object discovered through Assisted Discovery:

1. Right-click the object in the tree panel and select **Discover, Add (Remove) Permission to Manage...**

The *Select Discovery Settings* dialog box appears.

2. Select the agent type to which you want to grant or deny permission to manage the object. If there are multiple stored discovery settings available, select the settings you want to use. Check the Discovery Settings Detail panel to make sure you have the settings you want, and click **Add (Remove) Permission to Manage**. The *Add (Remove) Permission to Manage* dialog box appears.
3. Select one or more of the active agents listed in the middle portion of the dialog box to grant or deny permission to.
4. Click **OK**. ControlCenter sends a command to the ECC Server to grant or deny the selected agent(s) permission to manage the selected object. You can check the status of the command in the *Review Discovery Results* dialog box. Note that if you deny all active agents permission, the ECC Server cannot execute management policies on the object, but it remains in the Repository until you delete it.

Topology View

In Topology view, you create topology maps, pictorial renderings of your SAN. When objects are checked in the tree panel with Topology view open in the target panel, those objects, the objects to which they are connected, and the connectivity relationships between them, are displayed in the map.

Objects must be discovered before they can appear in Topology view. Refer to *Discovery* on page 7-5 for more information.

Topology view is updated in much the same way as the tree panel. Object status, relationships, and alerts displayed in the tree are also displayed in Topology view. When an object is expanded so that it is displayed in both the tree panel and in the view, selecting that object in one place also selects it in the other. Most operations performed in the ControlCenter tree panel can also be performed in Topology view.

Topology maps can be saved and stored for use at a later time. Saved topology maps are referred to as saved Topology view preferences (refer to *Saving Topology Maps (View Preferences)* on page 7-18).

Objects cannot be deleted from the Repository through Topology view. However, when an object displayed in the view is deleted from ControlCenter, the object and its children are removed from the view, along with all associated links.

Objects can be dragged from Topology view into other ControlCenter views (for example, Properties view, Masking view, Path Details view, Performance view, and so on) to display comprehensive object data and perform storage management operations.

To view objects in Topology view:

1. Click **Topology** in the ControlCenter toolbar.
2. Drag objects from anywhere in the Console into the view.

Objects Rendered in Topology View

When an object is placed into Topology view, the elements to which it is connected appear in the view with connecting links drawn.

In the case of a host, connectivity to other hosts is not displayed.

Topology view displays any object that ControlCenter discovers, including, but not limited to the following:

Host objects

- ◆ hosts
- ◆ HBAs
- ◆ ports, etc.
- ◆ Fujitsu-Siemens BS2000 mainframe

Storage objects for the following arrays:

- ◆ Symmetrix
- ◆ CLARiiON
- ◆ Celerra
- ◆ Network-attached storage (NAS)
- ◆ StorageWorks HSG80, XP
- ◆ HDS
- ◆ SUN
- ◆ IBM ESS
- ◆ SMI-based

Connectivity device objects

- ◆ switches
 - McDATA/Connectrix
 - Cisco
 - Brocade
 - CNT
- ◆ hubs
- ◆ bridges/gateways
- ◆ extenders
- ◆ fabrics/VSANs
- ◆ unidentified ports

Miscellaneous objects

- ◆ adapters
- ◆ ports
- ◆ links
- ◆ user-defined objects and groups

Fabrics and Cisco VSANs

When you check a fabric or a Cisco VSAN, or check objects in the tree panel under a physical fabric including a switch, switch ports, zone sets, zones, or VSANs, the entire physical fabric is displayed in Topology view. In the case of a Cisco VSAN, all members (switch ports) of the VSAN are highlighted in green.

Links

Containers objects can be expanded in Topology view. The final rendition in a map contains ports and physical links. As you drill down, logical links between elements are displayed. Physical links (links to ports) are always displayed. Logical links (links to containers) can be toggled on and off using the **Show/Hide Links** button, which can be useful when displaying complex configurations.

Positioning the cursor over a link changes the color of the link to blue, displays the status of the link in the status area, and elicits a tool tip that displays the endpoints of the link. Physical links display the link name and both end-port names. If one or both endpoints of a link are not fully expanded, the link is logical, and the name of the parent container appears as a logical endpoint in the tool tip.

A legend palette can be launched to identify the icons displayed in Topology view.

ControlCenter Groups in Topology View

ControlCenter objects are grouped together in the topology map, typically in folders. When collapsed, the parent container is displayed with a (+) sign, indicating that it can be expanded to reveal its member objects. The following ControlCenter groups appear in the indicated columns (by default) of the topology map:

- ◆ **Hosts** appear in the left column of the map.
- ◆ **Unidentified ports** share the left column of the map with hosts.
- ◆ **Storage arrays** appear in the right column of the map.
- ◆ **Connectivity devices** appear in the center column of the map.
- ◆ **User-defined groups** appear in the center column of the topology map ahead of connectivity devices and fabrics.
- ◆ **Fabrics** share the center column of the map with connectivity devices. If the connectivity device is a switch, the fabric associated with the switch is shown as the parent of the switch. When expanded, the fabric is displayed as a colored box that contains the switches and switch ports making up the fabric. Fabrics are named by the WWN of the primary switch.

Arranging Groups in Topology View

Hosts, storage, and unidentified ports can be arranged by name in the topology map. Right-click the parent folder and select **Arrange By, Name**.

This feature is not available for user-defined groups, connectivity devices or fabrics.

There must be a minimum of four devices in a group in order to arrange its members by name. The name of the group is determined by the first and last elements in the group. In addition, hosts can be arranged by **Type** and unidentified ports can be arranged by **Vendor**. The default criteria for arranging groups is **None**.

User-Defined Groups in Topology View

User-defined groups created in the tree panel can be selected to display in the map. Group members can be highlighted and arranged by name. User-defined groups are displayed in the middle column of the topology map.

Links between members of a user-defined group are always displayed in the map. If a member is connected to objects outside the group, the external links are displayed when the group is selected.

Expand user-defined group folders by clicking the plus icon beside the group folder. This action replaces the folder icon with the members of the group. Group members can be subgroups or member objects. When a user-defined group is expanded to reveal its member objects, those objects are moved to the appropriate columns in the map—hosts and unidentified ports in the left column, connectivity devices in the middle column and storage in the right column.

A **combo box** appears in the upper-left corner of the map and contains the names of user-defined groups that have been expanded in the topology map. Groups and subgroups in the map can be expanded and collapsed as well as highlighted through the combo box. When you select an expanded group in the combo box, its members are highlighted in the map. If you select a collapsed group in the combo box, only the group folder is highlighted in the map. A **Clear Highlight** button at the bottom of the combo box allows you to clear the highlights from the selected group's members in the map. Collapsing a group in the combo box hides all of its members in the map and disables the (-) sign next to its subgroups in the combo box. Groups in the combo box with a grayed icon are not visible in the map, and its descendants are removed from the map.

Highlighting Zone and Zone Set Members in Topology View

If you check a zone or zone set in the tree panel, all of its members are brought into the map and highlighted. In addition, the neighbors connected to the zone and zone set members are brought into the map with connecting links drawn.

When a zone set or a zone is selected in the tree, all of its members (or parent containers, if collapsed) that are currently displayed in the map are highlighted. The links associated with those members are also highlighted. Interswitch links are not highlighted.

If you select multiple zones or zone sets, or a mixture of both, all of the members of the selected objects are highlighted.

Topology View Tools

ControlCenter provides a suite of tools and features that create flexibility in using the topology map. Click **Action** in the topology map title bar and select **Tools Palette**. The following tools can be accessed from the Tools Palette:

- ◆ **Show Neighbors/SAN** — Allows you to display only the neighbors of a selected object.
- ◆ **Hide/Show Links** — Toggles on and off logical links in the map.
- ◆ **Find** — Helps to locate objects in the map.
- ◆ **Filter** — The following filters can be used to clear the view of unnecessary objects when performing SAN operations:
 - **Show Used Ports** (default) — Displays only ports connected to the selected object.
 - **Show Used + Available Ports** — Displays ports connected to the selected object, and any available ports on the objects connected to the selected object.
 - **Show All Ports** — Displays the entire SAN.
- ◆ **Legend Palette** — Identifies the icons displayed in the map.
- ◆ **Expand All** and **Collapse All** — On the right-click context menu lets you expand and collapse all containers in the map.
- ◆ **Print** or **Export** — Prints or exports a topology map.
- ◆ **Tools Palette** — Contains the following set of map tools:
 - **Move** — Move container objects around in the map.
 - **Link** — Toggle on and off links associated with an object.
 - **Zoom box** — Draw a box around any section of the map to magnify it.

- **Zoom in** — Point and click to magnify any section of the map.
- **Zoom out** — Point and click to reduce the map magnification.
- **Default size** — Reset the map to the default (100%) setting.
- **Overview** — Open a small window depicting the entire topology map with a selection box drawn around the section of the map that is currently magnified.

Saving Topology Maps (View Preferences)

A **Save As** feature on the Topology view **Action menu** allows you to save topology maps as Topology view preferences. Topology maps can be saved, opened, and removed by selecting options from the View Preferences submenu. A dialog box appears containing a list of all your saved views (maps).

The following functions can be saved as Topology view preferences:

- ◆ Layout
- ◆ Collapsing/Expanding objects
- ◆ Hide/Show links
- ◆ Overview setting
- ◆ Zoom factor
- ◆ Arrangement of ControlCenter groups (by name, type or vendor)
- ◆ Arrangement of user-defined groups (by name)

All other Topology view preferences conform to the generic ControlCenter preferences.

The first time you launch Topology view it opens empty. If you saved Topology view preferences in previous Topology view sessions, the last saved configuration opens in Topology view. These saved sessions are referred to as topology maps. Saved topology maps are also referred to as saved or named Topology view preferences.

ControlCenter supports running multiple Topology view sessions, allowing you to simultaneously work on multiple instances of a topology map, or view different maps at the same time.

Inactive links are not displayed in Topology view.

Topology view does not display SRDF relationships. The SRDF view provides detailed information in tabular form, or use the Relationship view and Agent Relationship view for a map depiction of a SRDF relationships.

Topology Edit Service (TES)

ControlCenter's Topology Edit Service allows you to manually depict elements in the topology that cannot be discovered by the agents; for example, those objects that do not have software-based management interfaces; or a switch or a storage array that is not yet installed, but which you want to configure in the Console. You can depict these objects in the topology by providing some basic object properties through the **Create/Associate wizard**. User-defined objects and their information are entered in the Repository and persist just as discovered information does.

The Create/Associate wizard is used to:

- ◆ Create user-defined objects
- ◆ Associate unidentified ports

TES and Discovery

User-defined objects may be created and then later discovered to have properties that do not match those that a user had defined earlier. To avoid inconsistencies, you must enter valid properties when you create a user-defined object. ControlCenter uses these properties to correlate a user-defined object with a discovered object. If a user-defined object exists with the same properties as a discovered object, ControlCenter overwrites inconsistent properties and notifies the user via an alert.

User-Defined Objects

The following user-defined objects can be created in the topology with the TES Create/Associate wizard:

- ◆ **Hosts:** Windows, UNIX, mainframe, generic
- ◆ **Storage systems:**
 - CLARiiON arrays
 - HDS arrays
 - HP StorageWorks HSG80 arrays
 - HP StorageWorks XP arrays
 - ESS arrays
 - Celerra attached storage
- ◆ **Connectivity devices:** switches, hubs, bridges, gateways, extenders.
- ◆ **Fibre Channel ports:** host, storage, connectivity device, and unidentified ports.

- ◆ **Adapters:** Created automatically by TES when associating unidentified ports with managed objects such as hosts, switches, and storage arrays. You cannot directly create a user-defined adapter with TES.

To create a user-defined storage array or host:

- ◆ Right-click the **Storage Systems** or **Hosts** in the ControlCenter tree panel and select **New**.

To create a user-defined storage array port:

- ◆ Right-click a storage array in the ControlCenter tree and select **New**.

Refer to the ControlCenter online Help to create other user-defined containers and ports.

Associating Unidentified Ports

TES can be used to manually associate unidentified ports with containers such as hosts, storage systems, and connectivity devices. You can associate a port with a discovered container as well as a user-defined container. For example, after an unidentified port has been associated with a host, the port can then be granted access to storage devices.

To associate an unidentified port with a user-defined object:

1. In the ControlCenter tree, expand the **Connectivity** folder, and then expand the **Unidentified ports** folder.
2. Right-click an unidentified port in the tree, and select **Topology, Associate Port**. The Associate wizard displays with the port WWN and vendor.

Although unidentified iSCSI initiators can appear in the Unidentified Ports folder, these objects can neither be associated with parent containers, nor can they be user-created.

Unassociating Unidentified Ports

After an unidentified port has been associated with a container through the TES Associate wizard, this operation can be undone by right-clicking the port again, and selecting **Topology, Unassociate**. This undo feature becomes unavailable if, in the meantime, a ControlCenter agent has discovered the associated port.

Viewing the Login History

The Login History Viewer displays information from the login history tables of all VCM-enabled Symmetrix arrays. Login history tables store current and historical login information for all Fibre Channel adapters in a Symmetrix array. When a host HBA logs in to a Symmetrix array, a record is created in the login history table. Each FA port on a Symmetrix array contains its own row in the login history table.

You can use the information in the Login History Viewer to:

- ◆ Verify connections between hosts and Symmetrix arrays.
- ◆ Track configuration changes. For example, you can view a list of host bus adapters that were once connected to an FA, and see if they were connected through a switch.

The Login History Viewer can be minimized, maximized, or placed behind the main Console window, and can run multiple sessions.

You must have the Symmetrix SDM Agent installed and running on the host connected to Symmetrix arrays in order to read the current login information.

To view the Login History, click the **Monitoring** button in the Console taskbar, and then select **Topology, View Login History** on the task-associated menu bar above it. This launches the *Login History Viewer*.

Zoning

This section provides an overview of zoning and the zoning tasks performed through ControlCenter.

A fabric is a group of Fibre Channel switches that are connected to each other and to the end ports of hosts and storage devices, such that data can be transmitted between any connected host end port and any connected storage device end port. These connected end ports are part of, or members of, the fabric, along with its switch ports.

Figure 7-1 shows a switched Fibre Channel fabric with a SAN.




Figure 7-1 Switched Fibre Channel Fabric in a SAN

ControlCenter supports Brocade, Computer Network Technology Corp. (CNT), McDATA/Connectrix, and Cisco fabrics, as well as interoperability (heterogeneous) fabrics containing some combination of Brocade and McDATA/Connectrix switches. Cisco physical fabrics are partitioned into multiple logical fabrics, or VSANs (Virtual SANs); the VSANs sometimes require procedures different from those used for other fabrics.

Fabrics and Zoning

The ability of all connected devices to communicate with each other makes switched fabrics very powerful. Their complexity, however, makes it necessary to partition fabrics into subsets of connected logical devices through zoning.

A zone is a group of devices within a fabric that you want to communicate with each other—for example, a host, a storage array, and the switch that connects them. You create a zone by grouping the end ports of the devices involved, which are most often host bus adapters (HBAs) and host directors (also known as front end directors), or the switch ports physically connected to those end ports. As long as these devices are connected through the fabric, you can place them together into a zone. Because zones are created by grouping ports, not devices and switches, a single host can communicate with multiple storage devices using multiple zones, and vice versa.

Figure 7-2 shows the previously illustrated SAN with three active zones, the first (blue) connecting a host and a Symmetrix array, and the second (red) and third (green) connecting another host with two different Symmetrix arrays, respectively.




Figure 7-2 Active Zoning of a Fabric in a SAN

A *zone set* is a collection of zones that can be activated together, partitioning a fabric into zones. Although a fabric can have any number of zone sets associated with it, only one of these zone sets is active at any time. It is this active zone set that determines which of the devices connected to the fabric can communicate with each other.

Like zones, zone sets can be created, modified (by changing which zones are included), renamed, and deleted. Zone sets can be copied from one fabric to another, and can be copied under a different name within a fabric. Zones associated with a fabric can be included in any number of zone sets and copied (under a new name) within a zone set, and can also be copied from one fabric to another.

Fibre Channel Port Types

Fibre Channel standards use the term *node* to describe any device connected to one or more other devices over Fibre Channel. Each node has at least one port that connects (directly or through a switch or fabric) to other ports on other nodes. These ports are called *end ports*. In general, the ports of a switch are called *switch ports*, but further distinction is made depending on what the ports are connected to. Types of ports used in zoning in ControlCenter include the following:

- ◆ **N_Port** — A port that connects a node to a fabric or to another node; roughly synonymous with end port. A node's N_Port connects to a fabric's F_Port or to another node's N_Port. For example, an HBA's N_Port connects to the F_Port of a switch that is part of a fabric, or to an N_Port of a host director. An N_Port handles creation, detection, and flow of message units to and from the connected array.
- ◆ **F_Port** — A port on a switch that connects to an N_Port, thereby connecting the node the N_Port is on to the fabric.
- ◆ **E_Port** — A port on a switch in a fabric that connects to another E_Port on the same or on a different switch.

The link joining a pair of E_Ports is called an inter-switch link (ISL). E_Ports carry frames originating from the node ports as well as frames originating from within the fabric. Frames generated within the fabric provide control, management, and support of the fabric.

The relationships of N_Ports, F_Ports, and E_Ports are shown in Figure 7-3.




Figure 7-3 Fibre Channel Port Relationships

Cisco switches also support a trunking E_Port, or TE_Port, which connect two switches. Unlike E_Ports, TE_Ports can be members of more than one logical fabric (VSAN).

Zoning Types

When you create a zone in ControlCenter, you select the ports to be included in the zone, which in turn determines which devices are included. You have a choice, however, of methods by which you can specify these ports. These differing methods of specifying the ports in a zone are called *zoning types*.

Switch Port Zoning

Switch port zoning creates zones by grouping switch ports rather than end ports. The end ports (and therefore devices) to be zoned are determined by their physical connection to the specified switch ports.

A zone or zone set using switch port zoning contains only switch ports. Benefits of using switch port zoning include simplicity and convenience in certain environments, because zoning more closely reflects the physical configuration of the equipment and HBAs and host directors can in some instances be replaced without affecting zoning configurations.

EMC does not recommend using switch port zoning exclusively, however, for the following reasons:

- ◆ Management of ISL congestion by relocation of high-traffic end port pairs to a common switch cannot be handled automatically.
- ◆ Switch port replacement and the use of spare ports require manual changes to the zone configuration.
- ◆ When a switch's domain ID changes (for example, when you reconfigure a set of independent switches to form a fabric), the zone configuration may become invalid, increasing the chance of data corruption.
- ◆ Interoperability fabrics do not support switch port zoning.

End Port Zoning

End port zoning creates zones by grouping end ports (and thus the nodes they are on) rather than switch ports. The switch ports to be zoned are determined by their physical connection to the specified end ports. The two main advantages to using end port zoning include:

- ◆ Flexibility, because the identification of zone members does not change when their connections to switch ports are rearranged. Connections between fabrics and adapters can be changed affecting zoning configurations.
- ◆ Dynamic fabric reconfiguration, because management of ISL congestion by relocation of high-traffic end port pairs to a common switch can be handled automatically, without affecting device driver configurations, switch-zoning configurations, or storage device configurations.

However, disadvantages of end port zoning can include greater complexity, diminished flexibility, and diminished security.

End port zoning uniquely identifies the end ports involved through either their WWNs or their FCIDs (when members of Cisco fabrics).

In end port WWN zoning, each end port is identified by its World Wide Name (WWN), a unique identifier combining the device it is part of, and its port number. WWNs are factory-set on HBAs, and are generated on host directors in managed storage objects (for example, a Symmetrix array).

A WWN consists of eight hexadecimal numbers separated by colons. For example:

10:00:08:00:88:44:50:ef

Cisco FCID Zoning

Cisco VSANs also allow a variant of end port zoning in which each end port is identified by its Fibre Channel identifier (FCID), a unique logical attribute of each N_Port connected to a Fibre Channel fabric. ControlCenter provides limited support for FCID zoning.

You can import zones containing FCID members from Cisco fabrics, and you can activate zone sets containing such zones on Cisco VSANs. But you cannot change a member from WWN end port zoning to FCID zoning, even within a zone on a Cisco VSAN; you cannot add a port zoned by FCID to another zone (new or existing) unless its WWN is known (in which case it is converted to WWN zoning); and you cannot drag and drop a zone containing ports zoned by FCID from a Cisco VSAN to a non-Cisco fabric unless their

WWNs are known (in which case they are converted to WWN zoning).

If you choose to maintain zoning containing FCID members for a Cisco VSAN, ensure that N_Port FCIDs are locked on the switches. If they are not, rebooting a switch reassigns the FCIDs, making zoning out of date.

Mixed Zoning

You can create or edit a zone to contain a mixture of switch port zoning and end port zoning, as long as a zoning policy is not applied. More commonly, mixed zoning is employed by combining zones using different types of zoning in a single zone set. Under some circumstances, this may combine some of the advantages of each form of zoning.

Zoning in ControlCenter and on the Fabric

Third-party switch and fabric management tools allow fabrics and their switches to be zoned independently of ControlCenter. A fabric's active zone set as represented in the ControlCenter Repository can differ from the active zone set on the fabric itself. Both the fabric and ControlCenter can also store inactive zones and zone sets.

Zone Set Activation and Import

You can change the active zone set for a fabric in ControlCenter whenever you want by activating one of the inactive zone sets associated with the fabric. This automatically activates the new zone set on the fabric, overwriting the current active zone set on the fabric.

The active zone set on the fabric can also be imported into ControlCenter by one of several methods, overwriting the current active zone set in ControlCenter. Inactive zone sets can also be imported from the fabric.

Because the active zone sets may differ, you may lose zoning configurations when you take action to import or activate and therefore overwrite one or the other. Each of these operations allows you to specify ControlCenter's behavior when differences in the active zone sets are detected: cancel the operation, or continue and allow one active zone set or the other to be overwritten. Whichever option you choose, an alert is generated.

Comparing Zone Sets

When activating a zone set, you can compare the configurations of the zone set about to be activated with the current active zone set in the ECC Repository. The configuration differences are displayed on the Activate Zone Set dialog box. After clicking Execute on this Activate Zone Set dialog box, a confirmation dialog box appears prompting you to choose one of the following options:

- ◆ **Cancel operation** — Cancels the activate zone set operation if the active zone set on the fabric is different from the active zone set in ControlCenter.
- ◆ **Continue and overwrite** — Activates the new zone set even if the zone set active on the fabric is different from the ControlCenter active zone set. The new zone set overwrites the zone set that is currently active on the fabric without saving a copy.

Refer to the following online Help topics:

- ◆ *Activating zone sets*
- ◆ *Comparing a zone set with the active zone set*
- ◆ *Comparing any two zone sets (in the same fabric)*

Zoning Folders in the Tree Panel

ControlCenter stores active and inactive zone sets and zones in a series of folders located under each fabric or VSAN in the Console tree:

Active Zone Set Folder

Contains the current active zone set as represented in the ControlCenter Repository. The active zone set in the folder expands to display the zones it contains, and the zones expand to display their port members. Active zones and zone sets are indicated by green icons.

An active zone or zone set in this folder can be copied to the Planned Zones or Planned Zone Sets folder. When you right-click the active zone set and select Edit Zone/Zone Set, a Copy of Active instance opens in the Manage Zone Set dialog box. After you change the Copy of Active instance, it is saved with the tag Modified Copy of Active in the Planned Zones/Zone Sets folder.

When an inactive zone set is activated in ControlCenter, or when the active zone set is imported from the fabric, the zone set in this folder is replaced and saved in the Planned Zone Sets folder with the tag Last Active appended to its name.

Planned Zone Sets Folder

Contains the inactive zone sets that belong to the fabric. Zone sets created within ControlCenter appear in this folder. Zone sets expand to display the zones they contain, and the zones expand to display their port members. Planned zone sets and their members are indicated by gray icons, and can be copied, edited, dragged, and dropped.

A copy of the active zone set and its members appears in this folder with a green icon and the tag Copy of Active. If this zone set is changed it is tagged Modified Copy of Active. The previous active zone set is saved in this folder with the tag Last Active.

Inactive zone sets imported from a fabric are imported to this folder. Existing zoning elements are not overwritten during the import.

Planned Zones Folder

Contains inactive zones that are not associated with a zone set. Zones created within ControlCenter appear in this folder. Zones expand to display the ports that comprise them. Planned zones are indicated by gray icons, and can be copied, edited, dragged, and dropped.

Inactive zones imported from a fabric are imported to this folder. Existing zones are not overwritten during the import.

Switches Folder

The Zones and Zone Sets folders under each switch within the Switches folder contain inactive zoning imported from the fabric. These zones and zone sets cannot be modified, and must be copied to the Planned Zone Sets folder before any changes are made. Folders remaining from the previous import are replaced during each import.

The Imported Zone Sets folder in the Connectivity directory in the ControlCenter tree panel contains zone sets that are imported into ControlCenter from the EMC ESN Manager application using the ESN Migration Wizard. This folder is dynamic, appearing only as needed.

Figure 7-4 shows an example of these zoning folders in the tree panel.




Figure 7-4 Zoning Folders Displayed in the Tree Panel

Zoning Policies

ControlCenter lets you create zoning policies to apply to new zones as they are created. Zoning policies specify characteristics of new zones—such as zoning type, maximum number of host and storage ports that can be included, and a zone naming expression—and help to ensure that these characteristics are the same for all zones created within a fabric, all zones created by or for certain groups of users, and so on. Once a zone is created, however, the zoning policy that was applied during creation no longer has any effect on it.

Managing Zone Sets

Managing zone sets includes creating and modifying zone sets, activating and deactivating zone sets on fabrics, and importing zone sets. Zone sets are associated with a single fabric only; however, a zone set can be copied and associated with another fabric. Most zone set management operations can be performed through the *Manage Zone Set* dialog box.

To launch the dialog box, right-click a fabric or a zone set in a **Planned Zone Sets** folder and select **Manage Zone Set...** (Using **Edit** on a zone set displays the similar *Edit Zone Set* dialog box.)

Zone Set Operations

The following zone set management operations can also be performed by right-clicking a fabric, zone or zone set in the Console tree panel and selecting a command from the context-menu:

- ◆ To create a new zone set, right-click a fabric or a **Planned Zone Sets** folder, and select **New, Zone Set**. The *New Zone Set* dialog box appears.
- ◆ To add a zone to a zone set, right-click one or more zones in a **Planned Zones** folder and select **Zoning, Add To Zone Set**. The *Add To Zone Set* dialog box appears.
- ◆ To remove a zone from a zone set, right-click a zone in a **Planned Zone Sets** folder, and select **Zoning, Remove From Zone Set**. The zone is removed from the zone set.
- ◆ To rename a zone set, right-click a zone set in the **Planned Zone Sets** folder and select **Rename**. A cursor appears in the zone set name field with the text selected.
- ◆ To copy a zone set, right-click a zone set in the **Planned Zone Sets** folder, and select **Copy As**. The *Duplicate Name* dialog box appears; enter a unique name.
- ◆ To delete a zone set, right-click a zone set in the **Planned Zone Sets** folder, and select **Delete**. A delete confirmation dialog box appears.
- ◆ To activate a zone set, right-click a zone set in the **Planned Zone Sets** folder, and select **Zoning, Activate Zone Set**. A confirmation dialog box appears. You also can perform this operation on the active zone set in the **Active Zone Set** folder to reactivate it.
- ◆ To deactivate zoning, expand the Connectivity and Fabrics folders, right-click on a fabric and select **Zoning, Deactivate Zoning**. A confirmation dialog box appears.
- ◆ To import all the zone sets associated with a fabric, right-click a fabric in the tree panel and select **Zoning, Import Zone Sets, Import All**. A confirmation dialog box appears.
- ◆ To import only the active zone set on a fabric, right-click a fabric in the tree panel and select **Zoning, Import Zone Sets, Import Active**. A confirmation dialog box appears.

Active Zone Set Synchronization

Third-party switch and fabric management tools allow a fabric to be zoned independently of ControlCenter. As a result, a fabric's active zone set as represented in the ControlCenter Repository can differ from the active zone set on the fabric itself.

When the active zone set is imported from a fabric, the fabric's active zone set in the Repository is overwritten. The active zone set on a fabric can be imported into ControlCenter, overwriting the active zone set in ControlCenter, under the following circumstances:

- ◆ When a switch in the fabric is discovered and the Import active zone set option is selected during connectivity device discovery (refer *Connectivity Device Discovery*).
- ◆ During FCC Agent polls of the fabric, if one of the necessary Import Active Zone Set options of the Fabric Validation data collection policy is enabled.
- ◆ When you choose to import by right-clicking a fabric in the tree panel and selecting Zoning, Import Zone Sets.

ControlCenter does not support either zoning aliases or zone member types other than switch port, end port by WWN, or end port by FCID in zones imported from Cisco fabrics (refer to *Zoning Types*). When ControlCenter encounters a zone alias or unsupported member type during zone set import, it alerts you with an error dialog.

The active zone set on the fabric can be overwritten or removed by the following operations:

- ◆ When you activate an inactive zone set, or reactivate the current active zone set, in ControlCenter, that zone set is activated on the fabric as well.
- ◆ When you deactivate the active zone set in ControlCenter (refer to *Deactivating Zoning* on page 7-35), the active zone set on the fabric is removed and the fabric has no active zoning.

If the active zone set on the fabric has been changed using a third-party tool, the active zone set on the fabric may be different from (out of sync with) the active zone set in ControlCenter. Therefore, importing the active zone set from the fabric, activating a zone set in ControlCenter, or reactivating or deactivating the active zone set in ControlCenter can cause you to lose zoning configurations that exist only in one or the other of the active zone sets.

Before you import the active zone set, activate or reactivate a zone set, or deactivate the active zone set, the confirmation dialog box gives you the choice of two actions should the two active zone sets turn out to be different: canceling the operation, or continuing and overwriting one zone set or the other.

The Search for Connectivity Devices dialog box also includes settings by which you can choose one of these options. The Source panel of the Fabric Validation data collection policy template lets you specify whether the policy should never import the active zone set, compare the active zone sets and import only if they are the same, or import even if they are different, overwriting the active zone set in the Repository.

In any of these situations, if an import, activation, or deactivation is canceled because you selected the appropriate option, an *Active Zone Sets Out of Synch* alert is generated. If an import, activation, or deactivation continues and overwrites one zone set or the other, an *Active Zone Sets Synchronized* alert is generated.

Because importing, activating, and deactivating zone sets affects the tagged zone sets in the Planned Zone Sets folder, you may want to make and keep copies of these zone sets (Table 7-2).

Table 7-2 Effect of Zone Set Actions on Planned Zone Sets Folder

Action	Changes to Planned Zone Sets folder
Importing active zone set	Replaces Last Active with copy of previous active zone set and Copy/Modified Copy of Active with copy of imported zone set
Activating new zone set	Replaces Last Active with copy of previous active zone set and Copy/Modified Copy of Active with copy of activated zone set
Deactivating active zone set	Replaces Last Active with copy of deactivated zone set and deletes Copy/Modified Copy of Active.

Use **Copy As** to make copies of these zone sets with different names.

Editing the Active Zone Set

You cannot edit the active zone set directly. However, you can edit the Copy of Active zone set in the Planned Zone Set folder. After you have modified it the zone set is saved as Modified Copy of Active, and you can select it in the Planned Zone Sets folder and activate it.

Managing Zones

Like zone sets, zones are associated with one fabric only, but can be copied and associated with other fabrics. A zone cannot be activated on a switch or fabric by itself, but only as part of a zone set.

Zones contains ports (host, storage, switch, or unidentified) as members. Ports can be added to or removed from zones. When a port is not associated with a zone, it appears in the Unzoned Ports folder under a fabric in the tree panel.

ControlCenter's drag-and-drop feature makes it easy to manage zones. You can drag a port into a zone, and drag a zone into a zone set, which in turn can be dragged into a fabric or a switch.

Most zone management operations can be performed through the Manage Zone dialog box. To launch the Manage Zone dialog box, Right-click a zone the Planned Zones folder or a zone within it and select **Manage Zone....** (Right-clicking a zone and selecting **Edit** opens the similar *Edit Zone* dialog box.)

The following zone management operations can also be performed by right-clicking a fabric, zone or zone set in the Console tree panel and selecting a command from the context-menu:

- ◆ To create a new zone, right-click a fabric, or a Planned Zones folder, and select **New, Zone**. The *New Zone* dialog box appears.
- ◆ To add a port to a zone, right-click one or more ports in the tree panel and select **Zoning, Add To Zone**. The *Add To Zone* dialog box appears.
- ◆ To remove a port from a zone, right-click a zone in a Planned Zones folder, and select **Zoning, Remove From Zone**. The port is removed from the zone.
- ◆ To rename a zone, right-click a zone in a Planned Zones folder, and select **Rename**. A cursor appears in the zone name field with the text selected.
- ◆ To copy a zone, right-click a zone in a Planned Zones folder, and select **Copy As**. The *Duplicate Name* dialog box appears.
- ◆ To delete a zone, right-click a zone in a Planned Zones folder, and select **Delete**. A delete confirmation dialog box appears.

Editing Active Zones

You cannot edit zones in the active zone set directly. However, you can edit the Copy of Active zone set in the Planned Zone Set folder. After you have modified it the zone is saved in the Modified Copy of Active zone set, and you can select it in the Planned Zone Sets folder and activate it.

Zoning States

Different types of fabrics can be in different zoning states depending on whether zoning is activated (there is an active zone set) and whether default zoning is enabled or disabled.

Deactivating Zoning

The active zone set can be deactivated on a fabric without activating another zone set on the fabric. The behavior of a fabric depends on the type of fabric and whether default zoning is enabled or disabled.

Default Zoning

The default zone consists of all unzoned ports. When default zoning is enabled, all unzoned ports can see each other. When default zoning is disabled, only ports in the same active zone can see each other.

- ◆ To enable default zoning, right-click a fabric in the tree panel and select **Zoning, Enable Default Zoning**.
- ◆ To disable default zoning, right-click a fabric in the tree panel and select **Zoning, Disable Default Zoning**.

Default zoning applies to McDATA/Connectrix, Cisco, and CNT fabrics, but not to Brocade fabrics, or to interoperability fabrics containing Brocade switches.

Default zoning on a logical fabric (Cisco VSAN) can be inconsistent as well as enabled or disabled. This is due to inconsistencies in the default zoning of the switches within the logical fabric (Cisco switches can be part of more than one logical fabric). To correct this, you can change the state of default zoning either from ControlCenter or by using Cisco Fabric Manager.

You can enable or disable default zoning on a fabric whether or not a zone set is activated on the fabric. When you disable default zoning on fabrics with no active zone set, no device can see any other device.

Table 7-3 shows the zoning states that result from the various possible combinations of active/inactive zoning and enabled/disabled default zoning.

Table 7-3 Zoning States

Type of Fabric	Possible Zoning States
Brocade	zoning disabled — No zone set is active; all devices see all other devices. zoning enabled — One zone set is active; all devices in the active zone set can see other devices in the same zone.
Cisco logical fabric	no active zone set and default zone disabled — no device sees any other device. no active zone set and default zone enabled — all devices see all other devices. zone set active and default zone disabled — only devices in the same zone in the active zone set see each other.
CNT	no active zone set and default zone disabled — no device sees any other device. no active zone set and default zone enabled — all devices see all other devices. zone set active and default zone disabled — only devices in the same zone in the active zone set see each other. zone set active and default zone enabled — all devices in the active zone set can see other devices in the same zone, and all unzoned devices see each other, but zoned devices see neither devices in other zones nor unzoned devices, and unzoned devices do not see zoned devices.
McDATA/Connectrix	no active zone set and default zone disabled — no device sees any other device. no active zone set and default zone enabled — all devices see all other devices. zone set active and default zone disabled — only devices in the same zone in the active zone set see each other. zone set active and default zone enabled — all devices in the active zone set can see other devices in the same zone, and all unzoned devices see each other, but zoned devices see neither devices in other zones nor unzoned devices, and unzoned devices do not see zoned devices.
Interoperability	If fabric does not contain Brocade switches, see McDATA and CNT row. If fabric does contain Brocade switches, see Brocade row.

EMC Zoning Recommendations

End Port WWN Zoning

EMC recommends the following zoning practices.

Use end port zoning where possible. Zones with switch port zoning do not automatically adjust to physical link changes that may take place during switch maintenance and repair. Refer to the *Converting to switch port, end port, and mixed zoning* online Help topic.

Zoning policies can be applied to new zones to determine (among other things) the type of zoning used, although they do not affect existing zones. The development and application of a set of zoning policies is a useful way to ensure that all zones created within ControlCenter are consistent with their purposes and use the appropriate type of zoning.

Planned Zoning Directories

Perform all ControlCenter zoning procedures in the following directories under each fabric in the ControlCenter tree panel:

- ◆ **Planned Zone Sets**
- ◆ **Planned Zones**

Monitoring Statistics

In addition to the Console-wide alert mechanisms, you can monitor the performance statistics of Symmetrix arrays and Fibre Channel connectivity device ports. Symmetrix statistics monitoring can be done on an individual port basis as well as for the entire array. Statistics cannot be collected for entire switches, however, but only for individual switch ports.

A variety of statistical data is monitored including, but not limited to, the following:

- ◆ Errors and failures
- ◆ Buffer-to-buffer credit events
- ◆ Byte counts
- ◆ Frames composition and counts
- ◆ Link sets
- ◆ Offline sequences

Masking

Masking allows you to restrict host access to a defined set of logical devices on a given storage array. Masking is also referred to as LUN masking, LUN security, and storage device masking.

Masking View

Masking view, which appears in the Console target panel, provides a series of filters and a pictorial or tabular interface used to isolate logical storage devices to which you want to grant host access. Right-clicking one or more storage devices in the view presents menu options to add or remove host access. Host access is then configured through a dialog box or a wizard interface. To use Masking view:

1. Click the drop-down arrow beside **Storage Allocation** in the Console taskbar, and select **Masking**. Masking view opens in the ControlCenter target panel.
2. Select an object in the tree panel to populate Masking view.
3. In the upper Masking view panel, use the selection filters to choose host ports (WWNs) to configure for storage access.
4. Use the interdependent device filters to display combinations of logical devices in the lower panel of Masking view.
5. Right-click one or more storage devices (LUNs) to access the shortcut menu that contains masking commands to grant or remove (host) access.

CLARiiON arrays do not use Masking view. Perform CLARiiON masking operations on objects in the ControlCenter tree.

Supported Arrays

ControlCenter supports masking for the following storage arrays:

- ◆ EMC Symmetrix
- ◆ EMC CLARiiON
- ◆ HP StorageWorks HSG80
- ◆ HP StorageWorks XP 48/512
- ◆ HP StorageWorks XP 128/1024
- ◆ HDS 9500V Series and 9900V Series
- ◆ SUN StoreEdge 9900V series

Enabling Masking

Masking must be enabled on a storage array or storage port before masking operations can be performed.

If masking is not enabled on a storage array/port, a host can see all the logical devices (LUNs) mapped to a storage port to which it is physically connected. After masking has been enabled on a storage array/port, you can configure host access to selected logical devices through ControlCenter's masking functionality.

Masking Actions

After you create specific masking *actions* in the various masking dialog boxes and wizards, the actions are sent to the ECC Server for execution. The new masking configurations appear in ControlCenter after the masking actions have finished executing on the ECC Server.

Required Agents

The following ControlCenter agents must be running on a host in order for masking operations to be performed:

Symmetrix Storage Arrays

- ◆ Storage Agent for Symmetrix
- ◆ Symmetrix SDM Agent

CLARiiON Storage Arrays

- ◆ Storage Agent for CLARiiON

StorageWorks HSG80 storage arrays

- ◆ Storage Agent for StorageWorks

HDS 9500V / 9900V Series Arrays, HP StorageWorks XP 48/512 Arrays, HP StorageWorks XP 128/1024 Arrays, and SUN StoreEdge 9900V Series Arrays

- ◆ Storage Agent for HDS

Symmetrix Masking

Enabling Symmetrix Masking

Masking is enabled on a Symmetrix storage port by setting the VCM port flag. This state is called VCM-enabled, and applies to both Fibre Channel and iSCSI. Only VCM-enabled storage ports and adapters can be masked and only VCM-enabled storage ports appear in ControlCenter's Masking view and masking dialog boxes. Refer to *Symmetrix Manager: Setting port flags* in the online Help for more information. For additional information about enabling and disabling masking on Symmetrix FA ports, refer to the *Solutions Enabler Symmetrix Device Masking CLI Product Guide*.

When performing masking operations on Symmetrix arrays, an Execute Now / Execute Later dialog box prompts you to add the masking actions to a task list. Through the Execute Now / Execute Later dialog box the task list is sent to the ECC Server for immediate execution or execution at a later time.

VCM Database

The Symmetrix Volume Configuration Management (VCM) database stores access configurations that are used to grant host access to logical devices in a Symmetrix storage array.

Masking operations performed on Symmetrix storage devices result in modifications to the VCM database in the Symmetrix array. The VCM database can be backed up, restored, initialized, and activated. The Symmetrix SDM Agent must be running in order to perform VCM database operations (except deleting backup files).

The following VCMDb maintenance operations can be performed through the ControlCenter Console:

- ◆ **Activating the VCM database** — Refreshes VCMDb changes (masking modifications and VCMDb maintenance operations) to all the Fibre Channel adapters on the Symmetrix array. Runs the Make Active command.
- ◆ **Initializing the VCM database** — Initialize the VCM database on a Symmetrix array in order to perform masking operations on the logical devices in that array.

Initialize the VCM database only once, before performing masking operations on that array. Subsequent initialization of the VCM database removes all access configurations from the storage array. After initializing the VCM database for the first time, initialize it again *only* when you want to clear all host access from all the logical devices on the storage array.

- ◆ **Backing up the VCM database** — Saves a copy of the access configurations currently active on a Symmetrix storage array.
- ◆ **Editing VCM database backup files** — Lets you change the name of an existing backup file and edit the comments in the backup file.
- ◆ **Restoring VCM database backup files** — Restores a backed up access configuration to a Symmetrix storage array.
- ◆ **Deleting VCM database backup files** — Requires no agent.

Make Active Command

The Symmetrix *Make Active* command activates changes made to the VCM database to all the Fibre Channel adapters on a Symmetrix array. For example, masking modifications made to one host become visible to all the hosts that are connected to the Symmetrix storage array. Running a Make Active command is synonymous with activating the VCM database.

Run the Make Active command after performing any of the following Symmetrix operations:

- ◆ Masking operations
- ◆ Setting SID Lock
- ◆ Setting Volume Visibility
- ◆ Initializing the VCM database
- ◆ Restoring the VCM database

Use one of the following methods to run the Make Active command:

- ◆ Select the **Activate the VC MDB** option on Symmetrix masking and VC MDB maintenance dialog boxes before clicking **Execute**. This causes ControlCenter to add a Make Active command to the bottom of the task list. After the actions on the task list execute, ControlCenter runs the Make Active command. The Activate the VC MDB option only appears on a dialog box when the VC MDB is being modified.
- ◆ Manually run a Make Active command on the Symmetrix array.
- ◆ Manually add a Make Active command to the bottom of the same task list that contains the actions created in a dialog box session. The Make Active command must not begin running until all Symmetrix masking and VC MDB maintenance actions have finished executing.

Symmetrix Masking Operations

The following masking operations can be performed on Symmetrix storage logical devices:

- ◆ **Viewing masking configurations** — Display existing device masking configurations.
- ◆ **Modifying masking configurations** — Adds (or removes) host access to storage logical devices.
- ◆ **Replacing masking configurations** — Swaps storage device access between two host ports.
- ◆ **Clearing masking configurations** — Removes storage device access configurations from hosts and storage arrays by deleting the WWNs of host ports and storage array ports from the Symmetrix VCM database. This operation also removes SID lock settings and volume visibility settings.
- ◆ **Setting SID lock** — Restricts host access to Symmetrix storage devices by adding switch source ID information to the VCM database, preventing WWN spoofing when multiple hosts are connected to the same storage port.
- ◆ **Setting volume visibility** — Allows certain HP and Linux hosts to view noncontiguous devices in Symmetrix arrays.

To view and modify host access configurations:

- ◆ Right-click one or more hosts, host ports, HBAs, or unidentified ports in the ControlCenter tree panel and select **Masking, Modify Masking Configurations**.

HP StorageWorks HSG80 Masking

When performing masking operations on HP StorageWorks HSG80 arrays, an Execute Now/Execute Later dialog box prompts you to add the masking actions to a task list. Through the Execute Now/Execute Later dialog box the task list is sent to the ECC Server for immediate execution or execution at a later time.

StorageWorks HSG80 Masking Operations

The following masking operations can be performed on StorageWorks HSG80 storage logical devices:

- ◆ **Viewing masking configurations** — Displays existing masking configurations.

- ◆ **Modifying masking configurations** — Adds (or removes) host access to storage logical devices.
- ◆ **Clearing masking configurations** — Removes storage access rights from hosts and storage arrays by removing the WWNs of the host and storage ports.

CLARiiON Masking

Masking CLARiiON storage arrays is performed in user-defined storage groups created on the CLARiiON array itself. Host access is granted and removed by adding and removing LUNs to and from the storage group, and then binding and unbinding hosts to the same storage group.

Masking operations for CLARiiON storage arrays can be performed through the Navisphere® CLI. For more information, refer to the *EMC Navisphere Command Line Interface (CLI) Guide* available on the EMC Powerlink website.

Enabling CLARiiON Masking

Masking is enabled or disabled on each EMC CLARiiON storage array by enabling or disabling the Access Logix™ software. Make sure that Access Logix is enabled on all CLARiiON storage arrays for which you want to control host access.

Storage Group Wizard

Storage group management actions are typically created in the CLARiiON Storage Group Configuration wizard, assigned to a task list and sent to the ECC Server for execution.

To launch the CLARiiON Storage Group Configuration wizard, right-click a CLARiiON storage object or subdirectory and select commands such as the following from the popup menu:

- ◆ New
- ◆ Edit
- ◆ Add To Storage Group
- ◆ Bind/Unbind Hosts
- ◆ Remove From Storage Group

CLARiiON Snapshots

A snapshot is a point-in-time picture of storage data that resides on a CLARiiON storage array. It is also referred to as a snapshot copy. Snapshots can be used for testing, backup, or protecting data. Snapshots typically behave like logical unit numbers (LUNs) and

reside in the Storage Systems/CLARiiON directory structure in the ControlCenter tree panel:

For more information, refer to the help topic *CLARiiON: Using EMC SnapView*.

CLARiiON Masking Operations

The following masking operations can be performed on CLARiiON storage logical devices:

- ◆ Viewing storage group configurations
- ◆ Creating storage groups
- ◆ Adding LUNs to storage groups
- ◆ Removing LUNs from storage groups
- ◆ Binding and unbinding hosts to storage groups
- ◆ Editing storage groups
- ◆ Renaming storage groups
- ◆ Deleting storage groups

HP StorageWorks XP 128/1024, HDS 9500V and 9900V Series, SUN StoreEdge 9900V Series Masking

Whenever *HP StorageWorks XP 1024* appears in the Help, it applies to *HDS 9500V and 9900V series* and to *SUN StoreEdge 9900V series* arrays.

Granting host access is performed by adding host ports (WWNs) and LUNs to the same host group on the storage port, so as to grant storage access to all the host ports in the host group at once. Similarly, host access is removed by removing the host ports and/or LUNs from the host group.

LUN addresses are assigned to LUNs in the same step as access rights are granted to host ports. Adding access rights is the same as adding the devices to the host groups.

Enabling LUN Security

LUN security must be enabled on a storage port in order to perform LUN masking operations. This can be done through in ControlCenter. Refer to *Enabling LUN security on storage ports* in the online Help.

Host Groups

User-defined host groups are used to grant multiple host port access to storage logical devices in a single masking operation, and are typically used to allocate storage to servers in a clustered environment. Host groups are created on storage ports and contain the following folders:

- ◆ WWNs folder — Contains host ports granted access all the LUNs in the LUNs folder.
- ◆ LUNs folder — Contains LUNs that are accessed by all the host ports in the WWNs folder.

HDS 9500V 9900V series arrays refer to host groups as storage domains.

Preconditions for Creating Host Groups

- ◆ The front end storage port must be Fibre Channel.
- ◆ LUN security must be enabled on the storage port.
- ◆ Maximum of 128 host groups per port.
- ◆ Host group names and host port WWNs must:
 - Be unique on the storage port.
 - Have a maximum length of eight characters and not contain special characters.
- ◆ Host port nicknames are optional and apply only to HP.
- ◆ Once you specify a host mode, the *Available Host Ports* table in the New Host Group dialog box shows only host ports belonging to a host whose operating system is compatible with that host mode. Host ports not discovered through the ControlCenter Host Agent have an unknown operating system. Such host ports are compatible with all host modes.
 - Maximum of 255 host ports per storage port.
 - Maximum of 255 host ports per host group.

Host mode

The host mode is set on the host group, and not on the controller or storage port. The host mode setting on the host group must be compatible with the host platform. All the host ports in a host group must have the same or compatible host mode. A single storage port can have host groups set to different host modes. The host mode is set on the host group when it is created or through an edit host group properties operation.

Nicknames

For convenience, host ports are given user-defined nicknames. A host port can have different nicknames in different host groups.

Masking Operations

The following masking operations can be performed on HP StorageWorks XP 128/1024 arrays:

- ◆ Granting access rights
- ◆ Removing access from selected host groups
- ◆ Removing access from all host groups
- ◆ Enabling LUN security on storage ports
- ◆ Disabling LUN security on storage ports
- ◆ Creating host groups
- ◆ Viewing host group members
- ◆ Editing host group properties
- ◆ Editing host group members
- ◆ Deleting host groups
- ◆ Creating host ports
- ◆ Renaming host ports
- ◆ Setting the host mode on array ports
- ◆ Mapping devices to array ports
- ◆ Unmapping devices from array ports
- ◆ Creating LUSE volumes
- ◆ Deleting LUSE volumes

HP StorageWorks XP 48/512 Masking

Enabling LUN Security

LUN security must be enabled on a storage port in order to perform LUN masking operations. Refer to *Enabling LUN security on storage ports* in the online Help for information on how this can be done with ControlCenter.

HBA WWNs

A host port is also referred to as an HBA WWN. HBA WWNs must be registered with a storage port in order to be granted access to the logical devices mapped to storage port. After registering with a storage port, an HBA WWN can be found in the WWNs and WWN Groups folders in the storage port directory structure. Refer to *Registering host ports with storage ports* in the online Help.

Registered HBA WWNs can be grouped together in storage port WWN groups. A WWN group contains a minimum of two registered host ports. WWN groups are used to grant multiple host port access to LUNs and LUN groups in a single masking operation. WWN groups are typically used to allocate storage to servers in a clustered environment. Refer to *Creating WWN groups* in the online Help.

Masking can also be performed for individual host ports, exclusive of a WWN group. Refer to *Masking registered host ports* in the online Help.

For convenience, registered host ports are given user-defined nicknames. A registered host port can have different nicknames in different WWN groups.

Host Mode

Host mode is set at the storage port level. In order for a host port to be registered with a storage port, the host mode setting on the storage port must be compatible with the host platform. All host HBA WWNs and WWN groups defined for a storage port must be of the same host platform type. Refer to *Setting the host mode on array ports* in the online Help.

LUNs and LUN Groups

Masking can be performed on a single LUN or a group of LUNs. Mapped storage devices can be grouped to form LUN groups. A LUN group must contain at least two LUNs. LUN groups are used to grant or remove host port (WWN) access to multiple logical devices in a single masking operation.

Masking Operations

The following masking operations can be performed on HP StorageWorks 48/512 arrays:

- ◆ Masking registered host ports
- ◆ Registering host ports with storage ports
- ◆ Unregistering host ports with a storage port
- ◆ Creating WWN groups
- ◆ Editing WWN groups
- ◆ Deleting WWN groups
- ◆ Renaming WWN Groups
- ◆ Renaming registered host ports
- ◆ Creating LUN groups
- ◆ Editing LUN groups
- ◆ Deleting LUN groups
- ◆ Renaming LUN groups
- ◆ Enabling LUN security on storage ports
- ◆ Disabling LUN security on storage ports
- ◆ Removing LUN security from LUNs
- ◆ Removing LUN security from LUN groups

Path Details View

Path Details view opens in the Console target panel and displays host-to-storage path information in both tabular and graphic form. It is a diagnostic tool used to:

- ◆ Troubleshoot paths.
- ◆ Identify the host and storage devices at the endpoints of paths.
- ◆ View the host, connectivity, and storage elements in a path.
- ◆ View the zoning status of a path.

In order for a path to appear in Path Details view, host access to a storage array must have been granted host access through the ControlCenter's masking utility.

Objects to Place Into Path Details View

The following objects can be selected in the ControlCenter tree panel, or dragged into Path Details view.

Host Objects

- ◆ host
- ◆ host device
- ◆ host adapter
- ◆ host Fibre Channel port
- ◆ unidentified port
- ◆ PowerPath® device
- ◆ PowerPath Volume Manager device

Storage Objects

- ◆ storage array
- ◆ storage adapter
- ◆ storage port
- ◆ controller
- ◆ storage logical device
- ◆ LUN
- ◆ storage group
- ◆ WWN
- ◆ WWN group

You can select an entire container to appear in Path Details view. However, it can take several minutes to load all the information into the view. To minimize the delay, select lower-level objects.

Table and Graphic Panes

Selection Pane

A selection pane at the top of Path Details view is used to narrow the number of objects displayed in the view. After you make selections and click the Show Devices button, the Path Details table populates with the respective Path Details. You can minimize and maximize the selection area by clicking the button in the upper right corner of the selection area.

Table Pane

When path objects are brought into Path Details view, they appear in the table pane, in the Path Details table. Each row in the table lists the details of one path. When one or more rows are selected in the Path Details table, the physical and logical SAN elements in the paths are graphically displayed in the lower graphics pane in the view.

Graphic Pane

The elements of a path are pictorially displayed in the graphics pane when the path is selected in the Path Details table. The object icons in the graphic pane can be expanded and contracted to reveal the physical and logical elements of the path. However, objects cannot be moved around in the view. The absence of a line between two ports on the path depicts a break in the physical connectivity of a path.

Viewing Path Details

To display the details of a host-to-storage path in Path Details view:

1. Click the drop-down arrow beside Storage Allocation in the Console taskbar and select **Path Details**.
2. In the ControlCenter tree panel, check one or more path objects into Path Details view. The Path Details table displays the details.
3. Narrow your search using the filters in the selection area, if applicable, and click **Show Devices**.
4. Select one or more paths (rows) in the Path Details table. A graphical depiction of the paths selected appears in the graphic pane.

Troubleshooting Paths

For functioning I/O to occur along a path, the following conditions must be met:

- ◆ **Agents** — ControlCenter agents must have discovered all the objects in the path.
- ◆ **Physical connectivity** — The host must be physically connected to the storage array, either directly, or through a fabric, indicated by a **Yes** value in the *Is Connected?* column of the Path Details table.
- ◆ **Storage device mapping** — The storage logical device must be mapped to the storage port, indicated by a **Yes** value in the *Is Mapped?* column of the Path Details table.
- ◆ **Zoning** — The ports in the path must be in the same zone, and in the active zone set, indicated by a **Yes** value in the *Is Zoned?* column of the Path Details table. A **Yes** value may also indicate that default zoning is enabled.
- ◆ **Masking** — The host port must have been granted access to the storage device indicated by a **Yes** value in the *Has Access Rights?* column of the Path Details table. If masking is disabled on the storage port/array, all connected hosts can access all the LUNs mapped to the storage port.

Using the Path Details Table

The following four columns of the Path Details table indicate vital areas where path I/O can break down. A **Yes** value is required in all four columns for I/O to occur on a path.

- ◆ **Is Connected?** — A **No** value in this column indicates that the physical connectivity on the path is down. I/O cannot occur on the path. One of the following physical cables, or a hardware device, may be down:
 - Host-to-switch cable
 - Switch-to-switch cable
 - Switch-to-storage array cable

Solution: Plug in the cable or troubleshoot the hardware failure.
- ◆ **Is Mapped?** — A **No** value in this column indicates that a storage logical device is not mapped to the storage port. I/O cannot occur on the path.

Solution: Map the logical device to the storage port.

◆ **Is Zoned?**

- A **No** value in this column indicates one of the following:
 - The ports in the path are not members of the same zone. I/O cannot occur on the path. **Solution:** Add the ports in the path to the same zone, then activate the zone set containing the zone.
 - The ports in the path are not in the active zone set. I/O cannot occur on the path. **Solution:** Activate the zone set that contains the zone with all the ports in the path.
- A **Not Applic** value indicates that the fabric has not been discovered.

- ◆ **Has Access Rights?** — A **No** value indicates that the host has not been granted access to the storage device. I/O cannot occur on the path. **Solution:** Grant the host access to the storage logical volume.

When the **Host Device** column of the Path Details table contains no value (is null), it indicates one of the following:

- ◆ The host agent may not have discovered the host device. **Solution:** Start the Host Agent.
- ◆ The host may have to scan for I/O paths. **Solution:** Perform a Rescan I/O operation on the host, or reboot the host.

Monitoring Storage with Alerts and Notifications

This chapter discusses how to monitor your storage environment and storage resources using alerts and notifications. The chapter is intended for the storage administrator and provides advice on finding the metrics you need to monitor your environment, checking the status of your environment using the alert views, and responding to alerts.

The chapter assumes that the ControlCenter administrator will create all alerts and notifications and manage the alert environment. For information about creating and managing alerts and notifications, refer to Chapter 3, *Configuring and Managing Alerts and Notifications*.

This chapter contains the following sections:

- ◆ Getting the Status of the Storage Environment 8-2
- ◆ Identifying Hosts, Arrays, and Network Components
Requiring Attention.....8-8
- ◆ Responding to Alerts.....8-9
- ◆ Tracking the Progress of Alert Resolution with Notes 8-13
- ◆ Finding the Alert You Need 8-14
- ◆ Gathering Information for Setting Alerts 8-15

Getting the Status of the Storage Environment

ControlCenter allows you to monitor hundreds of metrics about your storage environment—one example is the I/O rate of a Symmetrix director. For each metric, you can set values at which you want ControlCenter to notify you—for example, when the I/O rate exceeds 15,000 operations per second. When a metric exceeds a trigger value, it can appear either in the At A Glance view as a notification or in both the At A Glance view and the Alerts view.

- ◆ At A Glance views — Provide a higher-level perspective by dividing your environment into categories such as storage array performance or host capacity. For the At A Glance views, ControlCenter consolidates notifications into charts that indicate the statuses of the various categories, such as storage-system performance or host capacity.
- ◆ Alerts view — Provides a very detailed look at issues in your environment and offers many tools for tracking the issues to resolution.

Many alerts and notifications are enabled by default when the ControlCenter administrator installs an agent. You can work with the ControlCenter administrator to tailor these alerts and notifications to the needs of your datacenter. In addition, you can identify other metrics you want the ControlCenter administrator to configure for you.

After an alert or notification triggers, it appears in one of the alert views. For alerts of higher severity, ControlCenter adds a graphic to the icon of the object for which the alert triggered. The graphic appears throughout the Console.

Using the At A Glance View

To get an overall picture of the status of your environment, look at the At A Glance views. The following example demonstrates the host capacity chart in the At A Glance view (Figure 8-1).




Figure 8-1 At A Glance Host Capacity Chart

The Host folder is selected in the Console tree. As a result, the host capacity chart shows the notifications for all hosts. In this case, the chart shows that there are eight capacity notifications for all hosts.

Some options in this view:

- ◆ Show the members of the Host folder to see which hosts have triggered the notifications.
- ◆ Click the **Hosts** bar in the chart, and then click **Notifications** to see the details of the notifications.
- ◆ Click the **Hosts** bar, and then click **Capacity** to show the Free Space View, which provides a more detailed view of the hosts for which the notifications triggered.

Refer to *Using the Drill-Down Feature* on page 5-19 for more information about drilling down to other views from the At A Glance view.

Viewing Status of Grouped Objects in At A Glance

When you create and define your own groups of objects in the tree, you can view the status of the grouped resources in the At A Glance view.

Figure 8-2 shows health and capacity conditions for the Accounting Group. This hypothetical user-defined group includes hosts and a storage array used by the accounting department. You can expand and contract the group to show conditions for each member of the group (under Health) or for the entire group (as shown under Capacity).




Figure 8-2 Viewing Status of Grouped Objects in At A Glance View

For more information about using groups in the At A Glance views, refer to the online Help. Click the Help button on the view title bar to access specific Help for the view.

Accessing the At A Glance View

There are several ways to access the At A Glance view. The standard method is to click the **At A Glance** button on the toolbar.

You can also access the At A Glance view by clicking the Capacity and Performance icons in the lower-right corner of the Console. If you click the Performance icon (Figure 8-3), the At A Glance view opens with the Performance charts visible. If you click the Capacity icon, the At A Glance view opens with the Capacity charts visible.

Note that as you position the cursor over these icons, ControlCenter shows the current number of notifications at each severity level (also Figure 8-3).



Figure 8-3

Positioning the Cursor Over and Clicking At A Glance Icons

Using the Alerts View

Typically, you use alerts for more serious issues. In the Alerts view, ControlCenter provides several tools for tracking alerts to resolution, such as the ability to assign alerts and to record notes.

The All Alerts button, in the upper-right corner of the Console, is the quickest way to show the Alerts view. In addition, it lists the number of new alerts, the highest severity of any alert, and the total number of alerts at that severity (Figure 8-4).




Figure 8-4

All Alerts Button

Click the button to view all active alerts. The Alerts view appears (Figure 8-5).

Object Name	Message	Created
000183600408	Power Subsystem Error 24V...	06/07/2002 18:18:00
000183600408	Power Subsystem Error 24V...	06/07/2002 18:10:00
000183600408	Power Subsystem Error - S...	06/07/2002 18:06:00
000183600408	Power Subsystem Error 24V...	06/07/2002 18:00:00
000183600408	Power Subsystem Error 24V...	06/07/2002 17:38:00
000184600314	The Service Processor coul...	06/12/2002 14:02:00
001	RAID device not ready - 0001...	06/07/2002 15:42:00
01E	RAID device not ready - 0001...	06/07/2002 15:18:00
03D	RAID device not ready - 0001...	06/07/2002 17:56:00
03D	Volume not ready - 00018360...	06/07/2002 16:20:00
Agent=LOSBE105...	Symmetrix Agent agent has ...	06/12/2002 10:27:04
Agent=LOSBE105...	MGA agent has become inac...	06/12/2002 10:27:04
Agent=WANG124...	Storage Agent for CLARiON ...	06/07/2002 13:45:10

Figure 8-5 The Alerts View, Showing All Active Alerts

The color and icon indicate the status of each alert. In the fourth column, the number one icon represents the highest severity, Fatal, and the number five indicates the lowest, Information.

Alerts in bold text are new. After you acknowledge a new alert (by right-clicking it and selecting **Acknowledge**), the alert appears in plain text, and the icon in the first column dims.

The second column indicates if any notes are attached to the alert. An icon appears if there are notes. Right-click the alert and select **Alerts, Notes, Note** to view or add notes.

The third column indicates the status of any autofixes attached to the alert, either pending, completed, or failed. To get detailed status information, right-click an alert and select **Alerts, Autofixes**.

Limiting the Active Alerts That Appear

Often, you may find it easier to view subsets of active alerts rather than all active alerts at once. You may want to view the active alerts for a particular storage system or host, or a logical grouping of storage systems or hosts.

View the alerts for a particular object by opening the Alerts view and then double-clicking the object. ControlCenter adds to the view the alerts belonging to the selected object and its child objects.

Viewing Active Alerts in a Chart

You may find it easier to view a chart of active alerts for an object. To view the active alerts as a series of charts, click **Chart** in the Alerts view title bar (Figure 8-6).




Figure 8-6 Selecting the Alert Chart View

The Alert Chart view provides bar charts showing for each selected object the total number of alerts at each severity. The lower half of the view lists the alert details in a table. To see the alert details for a particular object, click the object name in its chart. To see the details of alerts of a particular severity only, click the corresponding bar in the chart (Figure 8-7).




Figure 8-7 Alert Chart View

Identifying Hosts, Arrays, and Network Components Requiring Attention

ControlCenter uses icons to indicate the status of storage systems, hosts, and network components. The status icons appear wherever the object appears. Figure 8-8 shows a tree view in which a storage system has one or more alerts with a severity of Warning and a Properties view in which a host has one or more Warning alerts.




Figure 8-8 Managed Object Status in Tree and View

If the managed object has multiple alerts, ControlCenter displays the icon for the alert with the highest severity. ControlCenter does not display an icon for objects that have Information alerts, only Minor, Warning, Critical, and Fatal.

As mentioned in the previous section, you can display the storage system or host in the Alerts view to see the alert details.

ControlCenter uses a red arrow that points down and to the right when a child object of a managed object has an alert. You can follow the trail of arrows to discover which child object has the alert. Figure 8-9 shows three hosts that have child objects with alerts, perhaps file systems with capacity alerts (left). If a managed object has an alert and one of its child objects has an alert, both the alert icon and down-arrow icon appear, as shown for host losav132 (right).




Figure 8-9 Managed Objects With Alerts

Responding to Alerts

When you respond to alerts, use the following general steps.

Responding to Alerts

To respond to alerts:

1. View active alerts. Click **All Alerts** in the upper-right portion of the screen. The Alerts view appears.
2. Find the alert in the Alerts view. Sort alerts and search alert messages.
 - Click the column headings to sort alerts by date and time, device or object, and severity.
 - Click **Find** to search for a specific storage resource in the Alerts view.
3. Get a description of the triggered alert. Right-click a triggered alert and select **Alerts, Help** for help in responding to the triggered alert. In the Help topic, click the alert name for a full description.
4. Find the affected resources. Navigate the selection tree or use the **Find** command to search the selection tree for the affected resource.
5. Fix the problem.
6. Tune the alert. If necessary, edit the alert to modify trigger values, the schedule, or the monitored resources, if this would make the alert more useful in discovering problems.

Verifying a Triggered Alert

Many alerts have similar names and alert messages, and slight but important differences in functionality. It is important to know exactly which alert triggered.

To verify the exact alert that triggered:

1. Right-click the triggered alert and select **Edit Definition**.
The Alert Definition dialog box appears.
2. Read the alert description.
3. If you are still not sure, click **Help** in the Alert Definition dialog box. A Help topic appears with a description of the alert, including the alert name.

Getting Help on Triggered Alerts

You can also get help when responding to triggered alerts. To access alert response Help, right-click an alert in the Alerts view, and select **Alerts, Help**, as shown in Figure 8-10.




Figure 8-10 Getting Help on a Triggered Alert

The Help system displays a topic with a statement of the problem, normally with immediate and long-term suggestions for resolving the condition and preventing it in the future. Further procedures are often available as well.

Getting More Information About an Affected Resource

To get more information on a resource that has an alert, right-click the resource in the Alerts view and select **Properties**, as shown in Figure 8-11.




Figure 8-11 Getting More Information About an Affected Resource

Note that the menu options vary depending where you click. You must right-click the resource in the Object Name column to get the menu options related specifically to the resource.

In addition to the Properties view, you can get information about an object by opening another view and dragging the resource to that view. For example, use the Relationship view to find out how a host relates to storage systems, databases, and other resources.

Checking the Status of Automatic (Autofix) Responses

ControlCenter allows you to run a command or script on a host when an alert triggers. You specify the command or script in an autofix. After an alert triggers, you can find out the status of the autofix in the Alerts view. The third column in the view indicates the autofix status.

Table 8-1

Autofix Status Icons

Icon	Status
	Autofix completed successfully.
	Autofix failed.
	Autofix has not yet completed.

If there are no autofixes attached to the alert, then no icon appears.

If there are multiple autofixes attached to the alert and one fails, then the Failed icon appears, even if other autofixes have succeeded.

To view the states of all autofixes, right-click the alert and select **Alerts, Autofixes**. The Autofixes dialog box shows the start and end time for each autofix attached to the alert.

The autofix status only indicates whether the associated script completed. It does not indicate whether the autofix corrected the issue it was intended to address.

Tracking the Progress of Alert Resolution with Notes

To better track the actions you have taken to resolve an alert, you can document your actions by attaching a note to an alert. You can continue to update the alert notes to create a log of your actions.

To create, edit, or view a note for an alert:

1. Display the active alerts.
2. In the Alerts view, right-click the alert and select **Alerts, Notes, Note**. The View/Add Notes dialog box appears.
3. To add a note, type in the **Add text for new note** box.
4. Click **OK** to save your changes.

Searching Alert Notes

When you receive a new alert, you can search the notes attached to older alerts to help resolve the new alert.

To search the alert notes:

1. Display the Active Alerts or Alert History view.
2. Right-click an alert and select **Alerts, Notes, View Related Notes**.
3. On the Search Notes dialog box, use the following criteria to search for notes. You can also use a combination of these criteria.

Table 8-2 Criteria for Searching Alert Notes

Type of Notes to Find	Action
Notes associated with alerts that derived from the same alert definition as the selected alert.	Select With same definition .
Notes associated with alerts that triggered for the same object (host, storage array, device, and so on) as the selected alert. These notes may help resolve the selected alert or may reveal trends with an object.	Select With same source object .
Notes associated with alerts that have the same template as the selected alert. These notes may provide clues for how to resolve the selected alert.	Select With same template .
Notes that were created within a date range.	Select Search Between and specify a date range.
Notes that contain a text string.	Type the string in Containing text .

4. Click **Search**. Notes matching your search criteria appear in Notes. Click any note to display the full text in Selected Note.

Finding the Alert You Need

ControlCenter alerts monitor hosts, storage systems, network components (such as switches), backup applications, and database management systems. This section lists commonly used host alerts.

Finding Alerts Using the Online Help

The online Help provides descriptions of all EMC ControlCenter alerts. Search the Help for topics beginning with the word *Monitoring*. These topics describe alerts from a task-based perspective.

Commonly Used Host Alerts

Tables 8-3 through 8-5 describe commonly used alerts for monitoring files, file systems, and disks.

Table 8-3

Windows File, Folder, and Volume Alerts

Task	Alerts
Monitoring the free space in a Windows volume	Percent Free Megabytes Free
Monitoring the number of files in a Windows folder	File Count File Count Change File Count Percent Change
Monitoring Windows file and folder size	File Size (KB) Size File Size Change (KB) Percent File Size Change

Table 8-4

UNIX File, Directory, and File System Alerts

Task	Alerts
Monitoring the free space in a ufs file system	MB Free Disk Space
Monitoring UNIX file and directory size	Percent Free Disk Space

Table 8-5

MVS Disk Alerts

Task	Alerts
Monitoring free space on a disk	% Free Space
Monitoring disk and tape units for operator intervention	Intervention
Monitoring fragmentation of volumes	Fragmentation

Gathering Information for Setting Alerts

Gather the following information you need to set alerts. Provide this information to the ControlCenter administrator when you request an alert. See the online Help for requirements for individual alerts.

Table 8-6 **Gathering Information for File Systems, Directories, and Files**

Information Needed	Description	Instructions	Notes
Hosts	Hosts that need to be monitored	List the hosts you want to monitor.	Other types of alerts may check multiple hosts, but file system and file alerts are best reserved for a single host.
Source	Resources to be monitored	List the file systems, directories, files, and disks you want to monitor on each host.	Explore hosts for their file systems, directories, important files, and disks.
Conditions	Trigger values and alert severities	For each resource, determine the values that should trigger alerts. • File systems and disks: Determine the triggers for free space and percentage free space. • Files and folders: Determine triggers for size, change in size, and percent change in size. Consider multiple trigger values for alerts of increasing severity: warning, critical, and fatal.	To help you determine trigger values, use recent data for resource free space and size. Also, consult the user of the resource.

Table 8-6 Gathering Information for File Systems, Directories, and Files

Information Needed	Description	Instructions	Notes
Schedule	Frequency that the alert conditions are evaluated	For each resource, determine: - How often the alert condition should be checked. - The days of the week on which the alert condition should be checked.	Critical or faster-growing resources should be checked more often (every 5 to 60 minutes). Others should be checked less often to decrease alert processing (every 60 to 360 minutes).
Management policy	Names and e-mail addresses of personnel to notify	Determine whom an alert should notify automatically: - In the Console - By e-mail - By page - In a framework product	You can limit the display of alerts to the Consoles of administrators with responsibility for the affected systems or applications. You can configure alerts to send e-mail to key personnel at appropriate times.
Autofix	Automated responses to alerts, including predefined or user-defined commands and scripts	Determine an automated action that would help resolve the alert. Assemble any scripts or commands that the alert could issue when triggered.	

Monitoring and Analyzing Performance

EMC ControlCenter provides the ability to quickly generate performance and configuration views of data collected by individual agents. This chapter describes the tools available, the configuration requirements, and the process for monitoring and analyzing performance through EMC ControlCenter.

- ◆ Performance Monitoring and Analysis Overview 9-2
- ◆ Performance Monitoring Configuration and Startup.....9-3
- ◆ Performance Analysis Configuration and Startup.....9-5
- ◆ Viewing the Performance Archives and Collections 9-13
- ◆ Daily and Revolving Performance Analysis 9-14

Performance Monitoring and Analysis Overview

ControlCenter provides you with the capability to monitor and analyze the performance of storage arrays, host systems, databases, and Fibre Channel switches within your SAN.

Performance Monitoring

Performance monitoring is provided for Symmetrix arrays and Fibre Channel switches. Performance monitoring is accessed through the EMC ControlCenter Console (ControlCenter Console) or the EMC ControlCenter Web Console (Web Console).

Refer to *Performance Monitoring Configuration and Startup* on page 9-3 for guidelines on configuring, starting, and using the ControlCenter performance monitoring tool.

Performance Analysis

ControlCenter allows you to analyze the performance of:

- ◆ Symmetrix storage arrays
- ◆ CLARiiON storage arrays
- ◆ Celerra systems
- ◆ HDS arrays
- ◆ Host systems including Windows, HP-UX, IBM-AIX, Linux, Sun Solaris, and MVS OS/390, and z/OS systems
- ◆ Oracle databases
- ◆ Fibre Channel connectivity devices

Performance analysis requires an EMC ControlCenter Performance Manager license.

This chapter provides the following sections outlining performance analysis configuration and use:

- ◆ *Performance Analysis Configuration and Startup* on page 9-5
- ◆ *Viewing the Performance Archives and Collections* on page 9-13
- ◆ *Daily and Revolving Performance Analysis* on page 9-14

Performance Monitoring Configuration and Startup

Performance monitoring of Symmetrix storage arrays and Fibre Channel Connectivity devices within the SAN is provided by the **Performance** view, located in the Console. For each object, real-time data can be displayed in chart or table form.

Required Components

The following components are required for performance monitoring:

- ◆ ControlCenter Console or Web Console (installed and running)
- ◆ Storage Agent for Symmetrix (to collect Symmetrix data)
- ◆ Fibre Channel Connectivity Agent (to collect switch data)
- ◆ Performance statistic data collection policy

Refer to the *ControlCenter 5.2 Planning and Installation Guide, Volumes 1 and 2* for procedures for installing and configuring ControlCenter components.

Setting Data Collection Policies

Once the required agents are installed, you can schedule how often performance statistics are collected for a specific Symmetrix array or connectivity device by modifying the Performance Statistics data collection policies (DCPs).

The Performance Statistic DCPs are enabled by default to collect data every 2 minutes. Users can define which Symmetrix arrays or connectivity devices the policies are applied to and can modify the polling schedule.

DCPs are edited from the ControlCenter Console tree panel by expanding **Administration, Data Collection Policies, Policy Definitions**, and then the agent folder for the appropriate object. Right-click the Performance Statistics DCP and select **Edit**.

Refer to the online Help accessed through the Console for more information.

Starting Performance View

To display real-time performance metrics for an object, from a Console:

1. Click **Performance**.
2. Select an object within the tree, such as components within a Symmetrix array.

Only Symmetrix arrays and Connectivity devices are available for Performance view. In the ControlCenter Console, objects that are not available for Performance view have a grayed out checkbox. In the Web Console, no data is populated to the view pane after selecting an unavailable object in the tree.

Performance metrics cannot be retrieved from a remote Symmetrix array. (The text "Remote" will appear in the row next to the row label.)

The Performance view opens. Click Help (?) in the view title bar for information about using the Performance view in the Console.

Performance Analysis Configuration and Startup

Performance analysis is done through EMC ControlCenter Performance Manager.

Before beginning performance analysis, the appropriate agents must be installed, configured, and started, and one or more WLA DCPs must be enabled for each data provider.


Once the agents are configured and running and data collection is started, the performance analysis process is as follows:

1. Data is collected by agents.
2. Agents transfer the data to the Repository and the Workload Analyzer (WLA) Archiver.
3. The Repository updates the configuration data.
4. WLA Archiver processes and appends the data in the performance archives.
5. The historical data is viewed from Performance Manager.

Performance Analysis Architecture

Figure 9-1 on page 9-6 shows the components and necessary architecture required for performance analysis.

Note that in Figure 9-1, the WLA Archiver can be installed on the same host as the ControlCenter infrastructure. Refer to the *EMC ControlCenter 5.2 Planning and Installation Guide, Volume 1* for details.



CC-000183

Figure 9-1 Performance Analysis Architecture

The components required for performance analysis are:

- ◆ The ControlCenter infrastructure
- ◆ ControlCenter Console
- ◆ The WLA Archiver
- ◆ Performance Manager

The Performance Manager is comprised of two components:

- Performance Manager
- Performance Manager Automated reports

If you intend to use the Performance Manager Automated reports, the Performance Manager server must be started.

- ◆ At least one of the following agents (based on the objects for which you are collecting and analyzing data):

There are certain steps that must be performed during installation of some of the agents in order to collect performance statistics. Refer to the following section, *Agent Requirements*, for more details.

- Storage Agent for Symmetrix Agent to gather Symmetrix statistics.
- Storage Agent for CLARiiON to gather CLARiiON statistics.
- Storage Agent for NAS to gather Celerra (NAS) statistics.
- Storage Agent for HDS to gather HDS statistics.
- Database Agent for Oracle to gather Oracle database statistics.
- Host agent to gather statistics for the applicable host on which the agent is running.
 - Host Agent for Windows
 - Host Agent for Solaris
 - Host Agent for HP-UX
 - Host Agent for AIX
 - Host Agent for LINUX
 - Physical Agent for MVS
- Fibre Channel Connectivity Agent to gather statistics about the switches in the SAN.

Agent Requirements

The installation and configuration requirements for each agent are included in the *EMC ControlCenter 5.2 Planning and Installation Guide, Volume 1*.

Table 9-1 on page 9-8 outlines the agents that have specific requirements for collecting Performance Manager statistics.

The following information is also included throughout *EMC ControlCenter 5.2 Planning and Installation Guides*.

Table 9-1 Agent Requirements for Collecting Performance Manager Data

Agent	Performance Manager Data Collection Requirements
Physical Agent for MVS	• EMC ResourcePak® Base for zOS must be installed and running on the host running the Physical Agent for MVS. • The RMF/SMF record types must be enabled. (Refer to the zOS documentation for instructions for enabling record types.)
Storage Agent for CLARiiON	• Navisphere CLI must be installed and running on the host running the Storage Agent for CLARiiON. • The SP Statistics logging must be enabled through Navisphere Manager. (Refer to the Navisphere Manager documentation for more information.)
Storage Agent for HDS — running on a Windows or Solaris host	• HiCommand Tuning Manager (HTM) and the HTM agents are configured and running. • Configure the HTM RAID Agent to collect statistics every minute using the HTM client. • JDK/JRE V1.4.2 or higher must be installed on the host running the Storage Agent for HDS. • After installing the Storage Agent for HDS, enter the HTM server and port information into the ADA Setting dialog box in the ControlCenter Console. This ensures that the HTM server and port are recognized during discovery. • Ensure that the HDS agent host can communicate with the HTM server.
	Storage Agent for HDS running on Solaris: • Performance statistics can only be collected from Solaris systems running on a SPARC host. • JDK/JRE V1.4.2 or higher is the default JRE. • System Variable: JAVA_HOME should be properly set to point to the JDK/JRE install directory.

Enabling and Editing WLA Policies

At a minimum, you must enable the WLA policies for performance data collection to begin.

Each agent that is used to collect performance statistics is installed with a set of predefined DCPs for managing the data:

- ◆ **WLA Daily** — Manages the data collected for performance analysis. The daily data is collected in intervals defined in the policy. The daily collection is the “standard” collection that makes up the performance archives.

When setting WLA Daily DCP collection intervals, valid intervals are factors of 60 minutes (for example, 15, 20, or 30 minutes). Examples of intervals that are not valid are 7, 13, or 45 minutes.

- ◆ **WLA Revolving** — Manages the continuous performance data collections. The amount of data contained in a revolving collection is determined by the defined window size. Once the duration is reached, the oldest interval is removed and the most recent data is appended to the collection.
- ◆ **WLA Analyst** — Allows you to create collections that contain information to perform analysis for a specific purpose during a defined time period. This collection is useful for troubleshooting issues in detail that occurred in the recent past.


At the same time you enable the policy, you can edit it to suit your performance analysis needs, and check the status of a running collection. For instructions on how to perform these procedures, refer to *Setting Data Collection Policies* on page 9-3.

Enabling or disabling the WLA Daily or WLA Revolving data collection policy changes the status for all objects to which the collections are assigned.

Once the data collection policies are enabled, statistical and configuration data is collected by the agents. The performance data collected by the agents must be archived before it can be viewed.

Archiving Performance Data

Performance data collected by agents is transferred to the WLA Archiver. The WLA Archiver converts the revolving and analyst data into Revolving and Analyst collections accessible from Performance Manager. The WLA Archiver converts the interval data into Performance Archives (see Figure 9-2 on page 9-10).



*In the event that data collection stops before the time specified as the end of a day, the WLA Archiver will wait a specified time period (90 minutes by default) for another sample before processing the data.

Figure 9-2 Archiving Process for Performance Archives

Once data archiving completes, the following performance data collections and archives are available for analysis through Performance Manager:

- ◆ Performance archives
- ◆ Revolving collections
- ◆ Analyst collections

Performance Archives

The performance archives are automatically processed from the data collected by the agents and managed by the WLA daily DCP.

Performance archives contain records of measured metrics and the time the measurement was taken. Each record is represented by one point on a graph created in Performance Manager.

There are four types of Performance Archives:

- ◆ Interval Archives

Interval archives are a direct conversion of the data collected by the agents and managed by the WLA daily DCP. If the agent was set to collect performance data every 15 minutes for 24 hours, then the Interval archive would contain 96 records, one record for each 15 minute collection.

- ◆ Daily Archives

The Daily archives have one record for each hour of the day. The maximum number of records in a Daily archive is 24, however, the number of hours that makes up the day is defined in the WLA Daily DCP schedule. Each record in the Daily archive represents one hour of the day. The record represents the average of the data collected for the WLA Daily DCP for that hour.

For example, if the WLA Daily DCP was set to collect data every 15 minutes, each record in the Daily archive would be the average of the four values collected for that hour.

- ◆ Weekly Archives

The Weekly archives have one record for each hour of the day. Each hour is the average of that hour for one week.

For example, the record for 9:00 a.m. to 10:00 a.m. for the Weekly archive is the average of the 9:00 a.m. to 10:00 a.m. record for each day of the week. The maximum number of records in a Weekly archive is 24, however, the number of hours is dependent on the number of hours defined as a day in the WLA Daily policy schedule.

- ◆ Monthly Archives

The Monthly archives have one record for each hour of the day. Each hour is the average of that hour for one month. For example, the record for 9:00 a.m. to 10:00 a.m. in a Monthly archive is the average of the 9:00 a.m. to 10:00 a.m. record for each day of the month.

The maximum number of records in a Monthly archive is 24, however the number of hours is dependent on the number of hours defined as a day in the WLA Daily policy schedule. Also, the weekdays that are included in the Monthly archive are specified in the WLA Daily policy schedule.

For example, you may have set your WLA Daily policy to only collected data from 9:00 a.m. to 5:00 p.m. on Monday through Friday, therefore the Monthly archive would only contain statistics for those days and times.

Revolving Collection

The Revolving data collection is a continuous collection with characteristics of an Interval archive that contains data for a defined time window. The data stored in the Revolving data collection time window is continuously updated with current data.

For example, a Revolving data collection set up to collect data at 2-minute intervals for a 120-minute time window will always contain 60 records. If the collection began at 9:00 a.m. the first collection cycle is complete at 11:00 a.m. At 11:02 a.m., the 9:00 a.m. record is deleted, and the collection consists of data collected from 9:02 a.m. to 11:02 a.m. This process continues until the collection is stopped, removed, or edited.

The content of the Revolving collections is defined from the WLA Revolving policy.

Analyst Collection

The Analyst data collection is a unique collection created for a specific purpose.

For example, suppose you are testing a new application on Saturday from 9:00 a.m. to 12:00 p.m. and you want to see how this affects system performance. You can create an Analyst data collection specifically Saturday from 9:00 a.m. to 12:00 p.m., while your other collections continue to run on schedule.

The content of the Analyst collections is defined from the WLA Analyst policy.

Viewing the Performance Archives and Collections

Performance Manager is provided with EMC ControlCenter for performance analysis. The Performance Manager provides two options for viewing the performance archives and collections:

- ◆ **Performance Manager** — Provides access to all the objects and statistics available in a data set and allows you to chart the information as selected.
- ◆ **Performance Manager Automated Reports** — Are HTML performance reports presented in a browser. The report content is defined through automation jobs. The Automation Job Scheduler is accessed from the Performance Manager.

Accessing the Performance Archives and Collections

Performance Manager with Repository Connection

How data is accessed from Performance Manager depends on whether the Performance Manager host is connected to the Repository.

When the Performance Manager host is connected to a Repository, the connection to the archives is automated.

The Repository maintains the location of the performance archives and collections. Once Performance Manager is started, it retrieves the location of the performance archives and collections from the Repository.

Revolving data stays on the agent host until it is manually retrieved from the Policies view of the ControlCenter Console. (Refer to *Setting Data Collection Policies* on page 9-3 for further instructions.)

Stand-alone

When the Performance Manager host is not connected to a Repository, you must manually point to the data collection location.

If the Performance Manager host is stand-alone, then the automated jobs scheduled to run will not work.

Instructions for manually pointing to the data collection location are provided in the Performance Manager online **Help, Contents, Accessing data collections** book, **Defining the location of data collections**.

Daily and Revolving Performance Analysis

Daily and revolving performance analysis is done from Performance Manager.

Performance Manager is accessed from the Windows **Start** menu, **Programs, EMC, EMC ControlCenter, Performance Manager**.

When starting performance analysis, it is important to understand your system environment, such as:

- ◆ The type of workload running on the Symmetrix array.
- ◆ The number and types of hosts.
- ◆ The number and types of applications running.

Performance Manager allows you to view system performance from the host to the Symmetrix device level. The following section reviews how you can view:

- ◆ Host information by looking at the host's CPU utilization and host response time, as well as relating host to Symmetrix array performance.
- ◆ Symmetrix array information by looking at cache management, balance and limits of the front end and back end, and disk contention.

While doing performance analysis, look for:

- ◆ Load balancing across Symmetrix components.
- ◆ The performance limits of the components.
- ◆ Trends in performance.
 - Daily and business workload patterns (including repeating patterns or issues).
 - Longer-term trends for performance planning.

Host Configuration Information

Because many users take a host perspective of performance, it is important to understand that Performance Manager collects open systems, host performance information and can associate this host information to Symmetrix information.

By collecting and displaying both host and Symmetrix information Performance Manager allows you to diagnose issues faster and more easily tune the interaction between host and storage arrays.

One way to look at host to Symmetrix information is to look at the host-to-Symmetrix configuration. After opening a host or Symmetrix array data set from the Data Selection dialog box, click the **Config** tab. Double-click a config file. A Configuration table opens that shows the host devices with their corresponding Symmetrix device, and the rest of the Symmetrix configuration (Figure 9-3).

Host\ID: 172.23.172.151 (Daily 8/14/2001, 0:00-24:00)

Vendor	Port	Director	Host Device	Symm Device	Meta V	Protection	P. Group	Disk	DA	DA Port	Pos.
EMC.0001841	148-01	FIBRE-14b	\\PHYSICALDRIVE1	0x008		DATA	None	02a 0x00	02a	D	0
						MIRR2	None	01b 0xC1	01b	C	1
			\\PHYSICALDRIVE10	0x02F		DATA	None	02a 0xC0	02a	C	0
						MIRR2	None	01b 0xD1	01b	D	1
			\\PHYSICALDRIVE11	0x043		DATA	None	02a 0x00	02a	D	0
						MIRR3	None	01b 0xC1	01b	C	1
			\\PHYSICALDRIVE12	0x044		DATA	None	01a 0x00	01a	D	0
						MIRR3	None	02b 0xC1	02b	C	1
			\\PHYSICALDRIVE13	0x045		DATA	None	02b 0xC0	02b	C	0
						MIRR3	None	01a 0xD1	01a	D	1
			\\PHYSICALDRIVE14	0x046		DATA	None	01b 0xC0	01b	C	0
						MIRR3	None	02a 0x01	02a	D	1
			\\PHYSICALDRIVE15	0x047		DATA	None	02a 0xC0	02a	C	0
						MIRR3	None	01b 0xD1	01b	D	1
			\\PHYSICALDRIVE16	0x048		DATA	None	01a 0xC0	01a	C	0
						MIRR3	None	02b 0xD1	02b	D	1
			\\PHYSICALDRIVE17	0x049		DATA	None	02b 0xD1	02b	D	1
						MIRR3	None	01a 0xC0	01a	C	0
			\\PHYSICALDRIVE18	0x04A		DATA	None	01b 0xD1	01b	D	1
						MIRR3	None	02a 0xC0	02a	C	0
			\\PHYSICALDRIVE19	0x04B		DATA	None	02a 0xD1	02a	D	1
						MIRR3	None	01b 0xC0	01b	C	0
			\\PHYSICALDRIVE2	0x00C		DATA	None	01a 0x00	01a	D	0
						MIRR2	None	02b 0xC1	02b	C	1
			\\PHYSICALDRIVE20	0x04C		DATA	None	01a 0xD1	01a	D	1
			MIRR3	None	02b 0xC0	02b	C	0			
\\PHYSICALDRIVE21	0x05B		DATA	None	01b 0xC0	01b	C	0			
\\PHYSICALDRIVE22	0x05C		DATA	None	02a 0xD1	02a	D	1			
\\PHYSICALDRIVE23	0x05D		DATA	None	02a 0xC0	02a	C	0			
\\PHYSICALDRIVE24	0x05E		DATA	None	01b 0xD1	01b	D	1			
\\PHYSICALDRIVE25	0x05F		DATA	None	01a 0xC0	01a	C	0			

Figure 9-3 Host-to-Symmetrix Configuration

Because data can be collected from the host, there is visibility into disk devices that may not reside on a Symmetrix array, such as devices on another array or devices local to the host (Figure 9-4 on page 9-16).

Vendor	Port	Director	Host Device	Symm Device	Meta Vol.	Protection	P. Group	Disk
EMC:000183400005	12B-1	12b	c2t2d2	0x44F				
			c2t4d2	0x66A				
			c2t4d1	0x669				
			c2t4d0	0x668				
			c2t2d1	0x44B				
			c2t4d3	0x66B				
			c2t4d5	0x66D				
			c2t4d4	0x66C				
			c2t4d7	0x66F				
			c2t4d6	0x66E				
			c2t2d0	0x443				
			c2t1d7	0x442				
			c2t1d1	0x43C				
			c2t1d6	0x441				
			c2t1d0	0x43B				
			c2t0d2	0x7B9				
			c2t0d1	0x7B8				
			c2t1d2	0x43D				
			c2t1d5	0x440				
			c2t1d3	0x43E				
EMC:509650965096	15B-0	15b	c2t0d0	0x7B7				
			c2t1d4	0x43F				
			c3t0d5	0x006				
			c3t0d3	0x02C				
			c3t0d4	0x004				
			c3t0d2	0x02D				
			c3t0d7	0x005				
			c3t0d0	0x02E				
FUJITSU:000000000000 IBM:000000000000	01A-0	01a	c0t1d0	0x000				
			c0t0d0	0x000				

Figure 9-4 Host Device Map to Non-EMC Disks

Viewing Host-to-Symmetrix Configuration

Use the host-to-Symmetrix configuration table to select host and Symmetrix devices to analyze.

For example, three graphs that you could create are (see Figure 9-5):

The procedure for creating the following graphs is described in *Creating Histograms of Host-to-Symmetrix Performance* on page 9-17.

- ◆ Symmetrix Devices I/Os per sec
This graph represents the “how many” aspect of performance.
- ◆ (host) Physical Devices KB/s
This graph shows the corresponding host devices and their throughput in KB/s, which represents the “how much” aspect of performance.
- ◆ (host) Physical Devices response times (ms)

This graph displays the response times as seen by those host volumes in milliseconds, which represents the “how fast” aspect of performance. In this case, the longer responses are seen by the volumes doing large throughput (larger blocks) as opposed to volumes doing a high number of I/Os (and a smaller block size).



Figure 9-5 Host-to-Symmetrix Performance

Creating Histograms of Host-to-Symmetrix Performance

After creating a host-to-Symmetrix configuration table:

1. Select the Symmetrix devices you would like to look at.
2. Click the **Set Selection Objects** button to add the devices to the Metrics panes.

3. Select **i/os per sec** from the list of metrics, and click the **Histogram** button.
4. Go back to the configuration table and select the host devices associated to the previously selected Symmetrix devices.
5. Add the devices to the Metrics panes.
6. Select **Kbytes transferred per sec** from the metrics list and click the **Histogram** button.
7. Using the same host devices, select **response times** from the metrics list and click the **Histogram** button.

Host CPU Utilization and Response Time

One of the host vital signs in Performance Manager is CPU utilization. Simply double-click **CPU Utilization** in the Views panel to open the vital sign (Figure 9-6).




Figure 9-6 CPU Utilization

In Figure 9-6, there is a spike of high %iio wait CPU. A spike such as this should lead you to investigate the issue and time period more closely.

You may want to look at the host response time. Figure 9-7 shows that a response time has been selected and the host physical devices have been sorted by average response time in the Metrics panes. This allows you to highlight devices that may be causing a problem. In this case, one device is showing very long response time while other devices show no activity.




Figure 9-7 Host Device Response Times

From the Metrics panes (Figure 9-7), graphs can be created for host devices. The graph in Figure 9-8 on page 9-20 shows the response time for two host devices over the course of a day. One host device response time has a very high response time (generally, response times above 10-15 ms should be investigated).




Figure 9-8 Host Device Response Times Graph

Symmetrix Array Information

In addition to open-system host information, with Performance Manager you can analyze Symmetrix array performance. When analyzing Symmetrix performance, it is important to remember the Symmetrix layout and cache management principles. Figure 9-9 shows a simplified view of the Symmetrix layout.




Figure 9-9 Symmetrix Cache Management and Data Flow

Cache Management and Data Flow

In the Symmetrix array, all reads and writes go through cache.

When a read is requested by a host, the information may be in the Symmetrix cache (a cache hit) or the Symmetrix array may need to go to disk to retrieve the information (a cache miss). For a cache hit, the information moves from cache through the host director and out the port to the host. For a cache miss, the disk director must read the information from disk and transfer it to cache before it can be read out to the host.

When a write operation is performed by a host system, data goes through the Symmetrix port and host director and is written to a slot in cache. An acknowledgement is sent back to the host. The cache slot is then marked as write pending — to be written to the disk at some later time. If the cache is full with data that has not yet been written out to the disk, the write operation cannot be completed due to a lack of available cache slots. The I/O will have to wait until slots get written to the disk and become available.

The system write pending count and the system write pending limit are two metrics you can look at through Performance Manager (Figure 9-10 on page 9-22).

- ◆ System Write Pending Count is a dynamic value that is captured at the time the statistics are taken. It is the total number of slots in the cache that are writes pending destage (not yet written to disk). A good value for this parameter depends on the cache size.
- ◆ System Write Pending Limit equals 80 percent of the user available cache. The System Write Pending Limit is the maximum number of cache slots that can be used for writes pending in the system.

The Symmetrix is designed to favor reads and therefore always reserves 20 percent of the available cache slots for read activity.

When the write pending count reaches the write pending limit, all writes for all Symmetrix devices are deferred until some of the write pending slots are destaged to disk. It is this condition that can result in deferred writes (write misses).



Figure 9-10 System Write Pending Count and System Write Pending Limit

System Component Limits and Balance

In addition to cache management, you can view the performance of ports, directors, and disks. The following vital signs use the utilization metrics to show the balance and limits of Symmetrix components. (Refer to Figure 9-9 for a view of the Symmetrix layout.)

Host Port - % Utilization

The Symmetrix component closest to the host is the front-end port.

This graph shows both the balance between the ports and how close each port is to its performance limit (Figure 9-11).




Figure 9-11 Host Port - % Utilization

In Figure 9-11, one port is nearing 100 percent utilized (close to its maximum throughput), and many of the ports are close to zero percent.

In situations like this, you should ask:

- ◆ Can we spread out the load across more of the ports (and relieve the load on our busy port)?
- ◆ Why are so many ports under utilized?

All Host Directors - % Utilization

After the port, the front-end director is the next component in the Symmetrix array.

The graph in Figure 9-12 shows the utilization of the directors (how close to the performance I/O limit).




Figure 9-12 All Host Directors - % Utilization

Things to note on this graph:

- ◆ Four directors are essentially unused.
- ◆ Three directors show identical activity, indicating a tool such as PowerPath in use.
- ◆ The director on the end has the highest utilization, but the level is below 14 percent, indicating that the directors are not near their performance limits. You may want to look into the balance of the directors, but since the overall level is low, time may be better spent investigating issues elsewhere.

Disk Directors – % Utilization

The third component with % utilization information is the disk director, or back-end director. This vital sign shows activity of the disk directors (Figure 9-13).




Figure 9-13 Disk Directors – % Utilization

In Figure 9-13, two groupings can be seen, with none of the directors near zero percent utilized. Perhaps the activity could be better balanced, but like the previous example of front-end directors (Figure 9-12 on page 9-24), the overall level of activity is not approaching the performance limits.

Understanding the storage environment is important to interpreting a graph like this. Is there striping in this environment? Are there multiple applications running? These things will help in understanding Performance Manager graphs.

All Disks - % Utilization

The final component with % utilization information is the physical Symmetrix disk. This % utilization is related to the I/O activity of the disk, and not the amount of information stored on the disk.




Figure 9-14 All Disks - % Utilization

This graph shows a number of disks with very high activity (peaks approaching 90 percent utilization), and a number of disks with essentially no activity.

Using the Graph Wizard, you can display the same data in a ribbon graph (Figure 9-15).




Figure 9-15 All Disks - % Utilization Ribbon Graph

This format allows us to better see the distribution of workload and to gauge how many disks have very low utilization. In this graph, some disks are very heavily utilized and many disks are inactive. In situations like this, you should ask:

- ◆ Can we spread the activity across more of the disks by changing the data placement with tools such as Optimizer (refer to Chapter 15, *Tuning Symmetrix Performance*)?
- ◆ Why are so many disks under utilized?

Disk Contention

This section describes how to investigate if disk contention is due to Symmetrix device placement.

Performance Manager allows you to identify Symmetrix device placement issues by first determining if the disks are busy.

The graph in Figure 9-16 shows the performance % utilization of selected disks (activity to the disks over the period of a day).




Figure 9-16 % Utilization for Selected Metrics

In Figure 9-16, the graph shows that there are peaks approaching 70 percent, so we know that at least some of the disks are very busy.

Another way to look at the information is to create a histogram of % utilization for all of the disks (Figure 9-17). This shows the average % utilization for each disk (average of the entire day).




Figure 9-17 Disks % Utilization Histogram

The graph in Figure 9-17 shows that several disks have low average utilization.

By viewing the information in a histogram, the “drill-down” function can be used. Right-click a bar in the histogram to select related metrics for creating another graph (Figure 9-18).




Figure 9-18 Disk % Utilization Histogram Drill Down

Note: Positioning the cursor over a bar shows the disk and the value. For example, in Figure 9-18, disk 01A xC0 has an average utilization of 37.39 percent.

By selecting **Devices, IOs per sec** we see a display of the average I/O activity for all of the Symmetrix devices located on the physical disk (Figure 9-19).




Figure 9-19 Symmetrix Devices I/Os Per Second

The graph in Figure 9-19 can give an indication of which Symmetrix devices you may want to move to another disk.

EMC ControlCenter Symmetrix Optimizer will automatically analyze Symmetrix device placement and move active volumes across the back end to improve performance.

Allocating or Deallocating Storage

This chapter covers the common tasks required to allocate or deallocate storage in an environment managed with EMC ControlCenter.

Allocating and deallocating storage resources requires an in-depth knowledge of storage architecture and device management. EMC strongly recommends that you attend EMC technical training before attempting the procedures outlined in this section.

You need ControlCenter authorization privileges on Host, Storage Array, Storage Pool, and Deallocation Policy objects to complete the procedures outlined in this chapter. Refer to Chapter 1, *Managing ControlCenter Users*, or the online Help for more information.

This chapter consists of the following sections:

- ◆ Storage Provisioning Service Overview 10-2
- ◆ Allocating Storage Overview 10-3
- ◆ Assigning Allocation and Deallocation Permissions 10-4
- ◆ Gathering Data Using Storage Array Properties 10-5
- ◆ Gathering Data Using the Free Space View 10-6
- ◆ Creating Storage Pools 10-8
- ◆ Creating Storage Allocation Policies 10-10
- ◆ Allocating Storage Using Storage Provisioning Service 10-14
- ◆ Deallocating Storage Using Storage Provisioning Service 10-16
- ◆ Deferring SPS Tasks for Future Execution 10-24

Storage Provisioning Service Overview

You can use the Storage Provisioning Service (SPS) to automate the task of adding new storage to a host or host cluster based on your requirements or removing storage that has been allocated to a host. The Storage Provisioning Service:

- ◆ Supports user allocation requests.
- ◆ Searches for devices and paths based on best-fit rules and displays the matches.
- ◆ Assists you in selecting storage allocation and deallocation options.
- ◆ Monitors execution of your requests and displays status and warns of problems.
- ◆ Includes Symmetrix Disk Reallocation (SDR) functionality.

Consider using the Storage Provisioning Service when:

- ◆ A new host is added.
- ◆ A new storage array (Symmetrix, CLARiiON, or StorageWorks) is added.
- ◆ A host requires more storage (for example, to extend a database or file system).
- ◆ A host or host device no longer needs storage that has been allocated to it.
- ◆ A migration of storage is required (for example, for the purpose of consolidation or expansion).

Refer to the online Help if you need to allocate or deallocate storage on a storage array using individual commands, instead of the Storage Provisioning Service. To access commands for manually configuring storage, click **Storage Allocation** on the taskbar, and then use the **Configure** menu on the menu bar.

The first half of this chapter discusses how to allocate storage using the Storage Provisioning Service. The second half of the chapter discusses deallocation.

Allocating Storage Overview

Storage Provisioning Services automates your storage provisioning tasks but before you can efficiently add storage or reallocate existing storage, you need to determine how storage is currently allocated on your storage network.

Once you determine the current state of your storage, you then set up storage pools that contain the devices available for allocation and create allocation policies that simplify and formalize repetitive allocation tasks.

The following sections outline the flow of the storage allocation process:

- ◆ *Assigning Allocation and Deallocation Permissions* on page 10-4
- ◆ *Gathering Data Using Storage Array Properties* on page 10-5
- ◆ *Gathering Data Using the Free Space View* on page 10-6
- ◆ *Creating Storage Pools* on page 10-8
- ◆ *Creating Storage Allocation Policies* on page 10-10
- ◆ *Allocating Storage Using Storage Provisioning Service* on page 10-14

Assigning Allocation and Deallocation Permissions

Allocation and Deallocation permissions are based on storage pools and hosts. By giving a user permissions for a specific storage pool, you allow that user to create and execute allocation tasks that draw from that storage pool. In addition, users must have permissions for the hosts involved in allocation/deallocation tasks. The following permissions apply to allocation and deallocation tasks:

To manage storage pools, users need the following permission:

- ◆ Allocation Administration — Allows you to create a storage pool, populate a pool with devices, remove devices from a pool, and delete a pool.

To create and execute allocation tasks, users must have the following permissions for *both* the storage pool and hosts involved in the task:

- ◆ Allocation Reservation — Allows you to create and save allocation tasks only.
- ◆ Allocation Execution — Allows you to execute allocation tasks.

To create a deallocation policy, users must have the following permission:

- ◆ Deallocation Administration — Allows you to create a deallocation policy.

To execute a deallocation task, a user must have the following permissions for both the storage pool and the hosts involved in the task:

- ◆ Deallocation execution — Allows you to perform a deallocation task such as unmapping storage devices from the front end of the array or dissolving metadevices.
- ◆ Deallocation reservation — Allows you to create a deallocation task.

Refer to *Assigning Permissions* on page 1-6 for more information about how these permissions are assigned and used.

Use the features outlined in the following two sections to gather information about your current storage resources.

Gathering Data Using Storage Array Properties

Each of the storage arrays displayed in the ControlCenter Console is an object with data collected and stored in the Repository. This data can then be used in the Properties view to determine the storage capacity and how it is being used.

Finding Unallocated Storage

Find unallocated storage on a Symmetrix array as follows:

1. In the tree panel, expand the **Storage Systems** folder.
2. If the Storage Systems folder is not currently arranged by array type, right-click the folder and select **Arrange By, Type**.
3. Select the checkbox next to the **Symmetrix arrays** folder.

The Properties table appears (Figure 10-1).

1 Properties - Symmetrix Arrays									Action	Filter	Find							
Symmetrix	S/N	Managed	Model	Configured Capacity	Unconfigured Capacity	Raw Capacity	Total Standard	Total BCV										
000000006223	000000006223	Direct Attached ...	DMX2000P	274.78 GB	1,772.89 GB	410.84 GB	205.28 GB	31.40										
000182601265	000182601265	Direct Attached ...	5300	421.63 GB	104.11 GB	421.63 GB	210.65 GB	168.69										
000183600358	000183600358	Direct Attached ...	5630	421.64 GB	121.05 GB	421.65 GB	168.37 GB	168.69										
000184600314	000184600314	Direct Attached ...	8430	201.31 GB	751.11 GB	334.30 GB	111.44 GB	48.03										
000187900611	000187900611	Direct Attached ...	DMX800	3,554.21 GB	120.45 GB	3,580.99 GB	3,502.53 GB	9.38										
				4,873.57 GB	2,869.61 GB	5,169.40 GB	4,198.27 GB	426.20										

Figure 10-1 Symmetrix Arrays Properties Table

4. Use the Unconfigured Capacity column to find a Symmetrix array with unconfigured storage space.
5. In the tree panel, right-click the **Mapped Devices** folder for the Symmetrix array and select **Properties**. The properties for all mapped devices on that Symmetrix array appear.
6. Right-click the **Unmapped Devices** folder for the Symmetrix array, and select **Properties**. The properties for all mapped and unmapped devices appear in a new table.
7. Double-click the **Allocated** column heading to sort the devices. You can now search the properties table for devices not allocated by SAN Manager that meet your storage requirements.

This step assumes that SAN Manager is running on your network. If not, you need to know which mapped devices are currently in use.

Gathering Data Using the Free Space View

Free Space view shows free space and allocated space for storage arrays related to a selected host or host cluster, with options to get detailed information for HBAs, ports, devices, and pools.

The following storage arrays are supported:

- ◆ EMC CLARiiON
- ◆ EMC Symmetrix
- ◆ HDS
- ◆ HP StorageWorks EMA with HSG80 controllers
- ◆ HP StorageWorks XP
- ◆ IBM ESS

You can use Free Space view to address the following questions:

- ◆ How much storage has been used by a host, where is it, and how much storage has not yet been used by the host?
- ◆ How much storage has been mapped to a host-accessible port, but has not been allocated?
- ◆ How much storage is available to the host from a specific storage array?
- ◆ How much storage has been used by databases, tablespaces, file systems, and other logical objects?
- ◆ Are there enough LUN slots available on the storage ports?

Displaying Free Space View

To display the Free Space view:

1. Select one of the supported storage arrays, a host, or host storage object, such as a file system or storage pool.
2. Click the **Storage Allocation** task button.
3. Select **Free Space** from the Storage Allocation task menu. The Free Space view appears (Figure 10-2).
4. After the display populates, you can select other objects or drill down to narrow the field of inspection.




Figure 10-2 Free Space View

There are many views of storage arrays, their components, and hosts available from within Free Space view.

This should be your primary tool in determining available storage resources within your environment.

Refer to the Console online Help topic *Free Space view: Overview* for more details.

Creating Storage Pools

Storage pools are collections of logical devices you can draw upon when allocating storage with the Storage Provisioning Service.


Some characteristics of storage pools are:

- ◆ Names are unique.
- ◆ You can nest storage pools within other storage pools to create a hierarchy. For example, you might create a New York pool and then Symmetrix and StorageWorks pools within it. If you create a pool to contain other pools, you cannot also add logical devices to that pool.
- ◆ A logical device can exist in only one storage pool.
- ◆ You can move logical devices between pools.
- ◆ After you allocate a logical device, it remains in its storage pool. In addition, you can add previously allocated devices to a storage pool to help track device ownership, for example.
- ◆ You can add all of the available devices in a storage array to a storage pool at once.
- ◆ A pool can contain devices from multiple arrays of different vendors.

What You Should Know before Starting

The Storage Provisioning Service allocates storage from the storage pools you create; therefore, you must create and populate storage pools before using the service.

The storage pool architecture has been designed with your entire enterprise storage environment in mind. This allows you to tailor the creation of storage pools to reflect your organization. For example, in Figure 10-3, the Boston facility has created two heterogeneous storage pools intended to serve database and file service operations that are stored on a variety of storage platforms. The New York facility has applications that are exclusively associated with a single storage platform, so they decided to have storage pools dedicated to each array type.



CC-000076

Figure 10-3 Storage Pool

Creating a Storage Pool

To create a storage pool:

1. In the tree panel, expand the **Storage Administration** folder.
2. Right-click the **Storage Pools** folder and select **New, Storage Pool**.

A new folder appears in the tree below the Storage Pools folder. The name of the new folder is highlighted and editable.

3. Type a folder name.

Adding a Device to a Storage Pool

Add devices to a storage pool by dragging the devices from elsewhere in the Console into the pool.

You can also make one storage pool a subset of another pool by dragging it on top of the pool in the Console tree.

A storage pool can contain either logical devices or other storage pools, but not both.

Moving a Logical Device between Pools

You can drag a device from one storage pool to another. This operation is a move, not a copy. A device can only belong to one pool at a time.

Creating Storage Allocation Policies

Allocation policies are used by the Storage Provisioning Service to simplify and formalize repetitive allocation tasks. In a policy, you can specify attributes to be applied to storage allocated through the Storage Provisioning Service, such as:

- ◆ Replica class (remote or local replicas).
- ◆ Storage pools from which to draw devices.
- ◆ Whether to restrict device selection to ports that are already zoned.

Allocation policies allow senior administrators to define parameters for allocations performed by others.

To create and edit allocation policies:

1. In the tree panel, expand **Storage Administration**.
2. Right-click the **Allocation Policies** folder, and select **New, Allocation Policy**.

The Allocation Policy Editor dialog box appears (Figure 10-4).

Allocation Policy Editor - New Policy

Policy Name:

Replica Class:

☐ Use this Policy as my default

☐ Add to SymAPI Device Group

☐ Allow devices from multiple arrays

Description:

Storage Element's Attributes:

Storage Element	Storage Pool	Storage Type	Raid Level	Port Balancing	#Paths	Mapped Device Only	Zoned Storage Only
Primary Device		Symmetrix	RAID_1	Host Based	1	☒	☒
Remote Replice		Symmetrix	RAID_1	Host Based	1	☒	☒

All cells except left most column are editable and have choices. Click on cell to display options.

Figure 10-4 Allocation Policy Editor Dialog Box

You can also create a new policy from within the Allocation Wizard. Table 10-1 describes the Allocation Policy Editor fields.

Table 10-1 Allocation Policy Editor Control Descriptions

Control	Description
Policy Name	User-specified name of the policy. You can edit this after it has been assigned.
Use This Policy as My Default	Make this your default allocation policy in the Storage Provisioning Service. You can select a different policy when the Storage Provisioning Service starts.
Enforce homogeneity of host OS on storage ports	Select this option if you want hosts that are running the same operating system (OS) to access storage through a fixed set of front-end ports. With this option selected, SPS will filter out the following types of ports during path selection: - Storage ports with allocated storage to multiple host operating system types - Storage ports with allocated storage to a host running an OS that is different from the OS of the host being allocated SPS only considers the operating system type, not the operating system version.
Pause between zoning changes	Pause the execution of a task list between zoning changes. If a task list contains multiple zone set activations, SPS pauses the execution of the task list after each zone set activation to allow you to confirm the zoning changes before changes are made to the next zone set. In the tree and Task List Viewer, a Paused Task List icon indicates that the task list is paused. To resume the storage allocation task after you confirm the zoning changes, right click the paused task list and select Resume Task List.
Replica Class	Indicates the type of replica device to be associated with the primary storage device. Click the Select Class button to display a dialog box that allows you specify a different replica class.
Description	User-specified string describing the policy.
Add to SymmAPI™ Device Group	Add newly allocated devices, which are allocated using this policy, to an existing SymAPI (Solutions Enabler) device group. This attribute only applies to Symmetrix arrays. This attribute is selected automatically if the replica class is Local, but the attribute is optional if the replica class is Primary Only or Remote. You must create the device group and prepare it to receive devices before you run the Allocation wizard.
Meta or non-meta, existing only	Choose either meta or standard devices.
Allow devices from multiple arrays	Allow the allocated storage to include devices from multiple arrays of the same type.

Table 10-1 Allocation Policy Editor Control Descriptions (continued)

Control	Description
Meta devices only	Choose metadevices only.
If not found, create striped metas	If the Allocation wizard cannot find metadevices that meet the parameters of the request, create striped metadevices.
Storage Element's Attributes	An editable summary table indicating the attributes of the primary and replica devices. The following attributes are shown: • Storage Element — Row-level indicator of the replica type for the device attributes to follow. • Storage Pool — Storage pool from which devices should be taken. • Storage Type — Type of storage devices to choose from within the storage pool. Only Symmetrix devices are allowed when the policy specifies local or remote replicas. • RAID Level — Type of RAID configuration associated with the device. See RAID levels later in this section for more information. • # of Paths — Number of paths to create between the host and the storage. • Mapped Device Only <Symmetrix only> — Whether to select Symmetrix devices that are already mapped (no SDR activity allowed). • Create New Storage Group <CLARiiON only> — Whether to create a new storage group if a group does not already exist for a host. Allocated devices are added to the new group. • Zoned Storage Only — Whether to select storage that is currently zoned to the host. This usually means a storage array that is in use or preconfigured for use. This supports the requirement to keep applications inside one storage array. If you do not specify this, SPS can choose any storage in the pool, even if that storage array has never been zoned to (used by) the host before. • Disable Host Actions — Specify this option to prevent any host actions during the allocation, such as rescanning, PowerPath operations, or actions on volume groups, logical volumes, and file systems.

RAID Levels For EMC Symmetrix, EMC CLARiiON, or HP StorageWorks, the RAID Level column in the Allocation Policy Editor provides the option to allocate the appropriate RAID configuration for the new storage (Table 10-2).

Table 10-2 RAID Level Configuration Options

RAID Level	EMC Symmetrix 5xxx (Enginuity™ 5668)	EMC Symmetrix DMX™ ^a	EMC CLARiiON	HP StorageWorks EMA with HSG80 Controllers
RAID_0	BCV or R1	BCV or R1	Supported	Supported
RAID_1	Supported	Supported	Supported	Supported
RAID_1/0	Supported RAID 1/0 devices are striped metadevices with mirrors. Create using the Symmetrix Manager.	Supported RAID 1/0 devices are striped metadevices with mirrors. Create using the Symmetrix Manager.	Supported RAID 1/0 devices are mirrored and striped LUNs. Create using the Storage Agent for CLARiiON.	Supported RAID 1/0 devices are striped LUNs. Create using native configuration tools.
RAID_5	Supported	Supported RAID 5 3+1 RAID 5 7+1	Supported	Supported
Parity Protection	Supported (RAID-S or RAID-R, depending on Symmetrix configuration)	Supported (Parity RAID 3+1 or Parity RAID 7+1, depending on Symmetrix configuration)		

a. DMX-specific RAID configurations require Enginuity 5669 or higher.

Guidelines for Creating RAID Devices

When a RAID device is required, SPS searches for this storage type in the storage pool specified in the allocation policy. If there are no devices matching this storage type in the specified storage pool, then the allocation will fail.

To avoid allocation failures of this type, you must first create devices of the desired type on each array. The devices should then be placed in the appropriate storage pool before executing an allocation policy that uses RAID storage.

A good practice is to create a pool containing multiple devices of this type of storage for allocation. Another good practice is to name the pools according to the type of storage they contain.

Allocating Storage Using Storage Provisioning Service

After you create a storage pool and set up an allocation policy, you can begin to use the Storage Provisioning Service as outlined in the flowchart shown in Figure 10-5.

The following steps provide an overview of the tasks involved in using the service:

1. **Select an initial object in the Console** — Select a host, file system, logical volume, or volume group.
2. **Start the service** — Start the Allocation Wizard by selecting **Allocate** from the **Allocation** menu or by right-clicking. The initial dialog box varies in appearance depending on the type of object selected.
3. **Review the policy** — Review the proposed allocation policy to be used with this allocation. You can edit an allocation policy from within SPS. This step does not apply if you selected a host.
4. **Select a host to see the replicas** — If the primary device receiving allocation has protection, then you can specify which hosts will be able to see the replicas.
5. **Configure metadevices** — Approve or edit the configuration parameters if the file system to be extended is on a concatenated metadevice.
6. **Configure path details** — Approve or edit the proposed path details, which include port, director, zone, LUN, and other characteristics.
7. **Review the allocation task** — Review the proposed final configuration, including the host for replication, general storage specifications, and path details.
8. **Execute immediately or later** — Add the task to a Task List in the tree panel and start it immediately, or defer the task to a later time (refer to *Deferring SPS Tasks for Future Execution* on page 10-24).




Figure 10-5 Storage Allocation Using the Storage Provisioning Service

Deallocating Storage Using Storage Provisioning Service

Deallocating storage with the Storage Provision Service removes the relationship between hosts or unidentified ports and the storage devices to which they have access.

You can only deallocate Fibre Channel Fabric connected paths; you cannot use the Deallocation wizard to deallocate storage connected through SCSI or Fibre Channel Arbitrated Loop.

Deallocation Process Overview

The deallocation process through the Storage Provisioning Service includes these steps:

1. Assign deallocation permissions — To deallocate storage, you must have permissions to create deallocation policies and to perform deallocation actions on the hosts and arrays involved in the deallocation.
2. Create deallocation policies — These policies define the actions to be performed during a deallocation task, such as rediscovering a host, unmapping devices from the front end of a storage array, and dissolving metadevices.
3. Select an initial object in the Console — Select a host, host device, host port, unidentified port, or storage device.
4. Start the service — Select Deallocate from the Allocation menu or by right-clicking and selecting **Allocation, Deallocate**.
5. Confirm the objects and select a policy — Add or remove objects from the deallocation task and select a policy to use.
6. Select the paths to be deallocated — Confirm the paths from host devices to storage devices that will be deallocated and add related paths if necessary.
7. Review the deallocation task — Review the final proposed deallocations.
8. Execute immediately or later — Add the task to a task list in the tree panel, and start it immediately or defer to a later time.

The remainder of this section discusses certain aspects of the deallocation process in more detail.

Understanding Deallocation Path Selection

You can start a deallocation task by selecting one of the objects in Table 10-3. The Storage Provisioning Service presents all the paths from that object to its storage devices or, if you have selected a storage device, the paths back to host devices and other WWN ports. You can then choose which paths to deallocate.

Table 10-3 Effects of Deallocation on Managed Objects

Selected Object	Objects Affected by Deallocation	Paths Presented for Deallocation
Host	All host devices connected to storage devices through Fibre Channel	All paths from host devices to storage devices and all paths from those devices to other host devices.
	All Fibre Channel ports on host	All paths from WWNs of ports to storage devices, based on masking.
Host Fibre Channel adapter	All host ports on adapter	All paths from WWNs of ports to storage devices, based on masking and all paths from those devices to other host devices.
Host Fibre Channel port	Host port	All paths from host devices to storage devices through that port and all paths from those devices to other host devices. All paths from WWN of port to storage devices, based on masking.
Unidentified port	Unidentified port	All paths from WWN of port to storage devices, based on masking and all paths from those devices to other host devices.

Table 10-3 Effects of Deallocation on Managed Objects (continued)

Selected Object	Objects Affected by Deallocation	Paths Presented for Deallocation
Host device	Host device	All paths from host device to its storage device and all paths from that device to other host devices.
Storage device	Storage device	All paths from known host devices to storage device All paths from WWNs of ports to storage device, based on masking.
User-defined group	All hosts, unidentified ports, host devices, and storage devices in the group	Refer to previous rows to see which paths are presented for each object type.

Deallocation Permissions

To perform deallocation actions, you must have the following permissions:

- ◆ **Deallocation Administration** — Allows you to create a deallocation policy.
- ◆ **Deallocation Reservation** — Allows you to create a deallocation policy. You need the Deallocation Reservation permission for each host and array involved in the deallocation or for the Host and Array object types.
- ◆ **Deallocation Execution** — Allows you to execute a deallocation task. You need the Deallocation Execution permission for each host and array involved in the deallocation task or for the Host and Array object types.

See your ControlCenter administrator to get the necessary permissions.

Controlling Deallocation Actions through Policies

The actions performed during a deallocation task are determined by a policy that you define. Some of the optional actions include:

- ◆ Rediscovering a host before and after the deallocation.

- ◆ Unmapping storage devices from the front end of the array (for example, removing a LUN from a storage group on a CLARiiON array).
- ◆ Dissolving metadevices or metaLUNs.

Storage administrators can create different policies for storage with different needs or to control deallocation tasks performed by junior administrators.

Deallocation Policy Permissions

To create a deallocation policy, you must have the Deallocation Administration permission for the Deallocation Policy object type. To edit a policy, you must have the permission for the specific policy or for the type.

See your ControlCenter administrator to get the necessary permissions.

Creating a Deallocation Policy

To create a deallocation policy:

1. In the tree panel, expand **Storage Administration**.
2. Right-click **Deallocation Policies** and select **New**.

The Deallocation Policy Editor appears. Table 10-4 describes the Deallocation Policy Editor controls.

You can also create deallocation policies by clicking **New** on the first screen of the Deallocation wizard.

Table 10-4 Deallocation Policy Editor Controls

Control	Description
Policy Name	Policy Name. This field is disabled if you accessed this dialog box by clicking the Edit button in the Deallocation wizard.
Use this policy as system default	The Storage Provisioning Service has one default deallocation policy. The first time you use the Deallocation wizard, this system-wide default policy appears in the Deallocation Policy field. If you choose another policy in the Deallocation wizard, that policy becomes your default policy. However, the system-wide default still appears for other users the first time they use the Deallocation wizard. Select this checkbox to make this policy the system-wide default.
Host Actions	
Disable host actions	Do not perform deallocation actions on hosts that use this policy.

Table 10-4 Deallocation Policy Editor Controls (continued)

Control	Description
Before Executing Tasks: Rediscover host and fail or error	Before performing the deallocation task, rediscover the host and update the Repository. The Deallocation wizard indicates which host devices are currently in use by host file systems, databases, device groups, and so on. If a deallocation task encounters storage that is still in use by a host, the task fails. Note: If you retry a task list, SPS does not check again whether host devices are in use. If the task list succeeds this time, you may lose data on host file systems, databases, and so on, that use deallocated devices.
After Executing Tasks: Rediscover host	Rediscover the host and update the Repository after completing the deallocation task. This keeps the Repository synchronized with changes that result from your deallocation task and prevents you and other users from seeing and acting on out-of-date data.
Array Actions	
Remove masking access rights	Remove host access to storage ports for deallocated devices (for CLARiiON arrays, remove LUNs from storage groups). This action is always performed.
Unmap storage device from front end	Unmap devices from the array front end ports that you specify in the Deallocation wizard. This action applies to Symmetrix devices only. This task is only performed if the device is no longer part of any path.
Remove from Symmetrix device group	If a device or replica belongs to a device group, remove it from the group. This task is only performed if the device is no longer part of any path.
Delete device (CLARiiON) or dissolve meta device if not SRDF (for Symmetrix, requires unmap)	For CLARiiON arrays, delete deallocated devices. For Symmetrix arrays, dissolve metadevices after unmapping them (assuming the metadvice is not an SRDF device). This task is only performed if a device is no longer part of any path.
Deallocate replicas if primary deallocated	When a primary device on a Symmetrix is deallocated, remove any access rights a host has to local or remote replicas of that device. Note: SPS treats a BCV as a replica, even after the BCV has been split from its standard device.

Starting the Deallocation Wizard

To deallocate storage, in the selection tree, right-click any of the objects described in Table 10-3 and select **Allocation, Deallocate**.

The Storage Deallocation Wizard appears (Figure 10-6).




Figure 10-6 Storage Deallocation Wizard

You can select multiple objects for deallocation, but generally, it is easier to select one at a time.

For detailed descriptions and assistance using the wizard, access the online Help. There is a Help button on each screen in the wizard.

Troubleshooting Deallocation

If the Allocation, Deallocate menu command is disabled, look for an explanation in the Hint Area in the lower-left corner of the Console window.

Possible reasons include:

- ◆ The Deallocation wizard does not support the selected object. Refer to Table 10-3 on page 10-17 to see which objects are supported.
- ◆ The selected object is not part of any Fibre Channel connected paths.
- ◆ Access to the selected object is not managed through masking access rights, or masking is disabled.
- ◆ You selected a storage device that is not mapped to any storage ports.
- ◆ You selected a storage device that is not exposed to any host ports through masking.

To determine whether a host has paths that can be deallocated, drag the host to a Path Details view. For a storage array device, use the Masking view.

In the Paths to Deallocate screen of the wizard, if a checkbox is disabled, position the cursor over the checkbox to find out why.

If the disabled checkbox *is not* selected, typically this indicates that you do not have the Deallocate Execution permission for the host or array.

If the disabled checkbox *is* selected, typically this indicates that the path has been selected for deallocation in another task, but the task has not yet executed or the Repository has not been updated.

Verifying Deallocation Actions

After you create a deallocation task, always verify the task actions in the Task List properties view, which appears automatically after you create the task. (For more information, refer to *Deferring SPS Tasks for Future Execution* on page 10-24.)

Some devices may be unmapped from ports that you did not explicitly select and that are not shown on the Paths to Deallocate screen of the Deallocation wizard. These additional unmappings occur when:

- ◆ You have selected the **Unmap storage device from front end** option in the allocation policy.
- ◆ A device is not visible to any hosts through the storage port.
- ◆ You are completely deallocating the device from all other storage ports.

For example, some primary devices may have replicas that do not have masking records and, therefore, are not visible to hosts. A list of these replicas appears after you click Next on the Paths to Deallocate screen of the Deallocation wizard.

If you do not want to deallocate those replicas, edit the deallocation policy and clear the **Deallocate replicas if primary deallocated** option. You can edit the policy from the first and last screens of the Deallocation wizard.

If you selected the **Delete device (CLARiiON) or dissolve meta device if not SRDF (for Symmetrix, requires unmap)** option in the Deallocation policy, metadevices that are completely deallocated (all paths to the device are removed) are dissolved. SRDF (R1 and R2) devices are not dissolved.

If **Deallocate replicas if primary deallocated** is also selected, local replica metadevices that are completely deallocated are dissolved. However, a replica that does not have any paths is only deallocated if its primary device is completely deallocated.

To identify the primary device associated with a replica, drag the replica to an SRDF or TimeFinder view.

Deferring SPS Tasks for Future Execution

The ControlCenter Task List provides a means to queue tasks or lists of tasks for future execution.

Tasks are added to the visible queue in the Task List once they are created.

You may want to defer a task because:

- ◆ The task requires technical review or permissions
- ◆ You need to delay an action to accommodate shift changes (some allocation tasks may only be performed on a certain shift)
- ◆ The task requires budgetary approval or internal customer review

Task Lists differ from batch files in that they execute one time only; they are not a repeated action like a backup batch file operation. A Task List may contain many individual tasks, such as storage allocation to a variety of file systems.

A task can be created by Storage Provisioning Service as part of the final review in the last dialog box.

Tasks do not run automatically. All tasks are manually started, including deferred tasks.




Figure 10-7 Execute Later Dialog Box

After a task has been created inside one of these applications, it is added to the Task Lists folder in the tree panel (Figure 10-8).




Figure 10-8 New Task Added to the Task Lists Folder

Tasks remain in the tree panel until they have been deleted. The properties of scheduled tasks can be displayed in the Properties view.

You can monitor a task in the target panel, observing each command execution. A checkbox in the last dialog box in the Storage Provisioning Service triggers the Properties view for that Task List. Use the view to inspect a task during or after execution.

The Status column is dynamically updated during task execution.

Select **Action's Details** to inspect the general properties of an allocation task, including existing and new size.

Select **Action's Commands** to display the command history of the task.




Figure 10-9 Task List Properties

This chapter provides a description of ControlCenter supported applications available to protect your data through local and remote mirroring of the devices on a storage array.

The terms *volumes* and *devices* are used synonymously in this chapter.

This chapter consists of the following sections:

- ◆ Introduction to Data Protection 11-2
- ◆ Protecting Data on Symmetrix Storage Arrays..... 11-3
- ◆ Protecting Data on CLARiiON Storage Arrays 11-33
- ◆ Protecting Data on HP StorageWorks Storage Arrays..... 11-35

Introduction to Data Protection

ControlCenter allows you to perform various tasks for protecting data on supported storage arrays. Depending on your array, you can perform business continuity functions from the ControlCenter Console, such as creating and managing business continuity volumes on CLARiiON, Symmetrix, and HDS arrays, or performing remote data protection on Symmetrix storage arrays using SRDF.

This chapter discusses:

- ◆ *Protecting Data on Symmetrix Storage Arrays* on page 11-3
- ◆ *Protecting Data on CLARiiON Storage Arrays* on page 11-33
- ◆ *Protecting Data on HP StorageWorks Storage Arrays* on page 11-35

Protecting Data on Symmetrix Storage Arrays

Symmetrix storage arrays are capable of providing both local (through Symmetrix TimeFinder BCVs) and remote (through Symmetrix SRDF) data protection. You can configure and manage TimeFinder and SRDF devices through EMC ControlCenter.

TimeFinder and SRDF concepts and capabilities are covered in detail in the *EMC Symmetrix Remote Data Facility Product Guide*, and the *EMC TimeFinder OS/390 and z/OS Product Set Product Guide*. Both of these guides are available on the EMC technical publications document CD provided with the ControlCenter 5.2 kit.

This section provides guidelines for completing the following data protection tasks on a Symmetrix storage array:

- ◆ *Configuring Symmetrix Device Mirrors* on page 11-4
- ◆ *Working with Symmetrix Groups* on page 11-5
- ◆ *Local Protection with Symmetrix TimeFinder* on page 11-10
- ◆ *TimeFinder Clones* on page 11-13
- ◆ *EMC Snap* on page 11-14
- ◆ *Remote Protection with Symmetrix SRDF* on page 11-16
- ◆ *Working with EMC Solutions Enabler SYMCLI* on page 11-25
- ◆ *Creating BCV Devices* on page 11-26

Configuring Symmetrix Device Mirrors

You can add or remove mirrors for Symmetrix devices with the Device Protection Definition dialog box accessed from the Console taskbar through **Storage Allocation, Configure, Symmetrix**, or through the shortcut menu (Figure 11-1).




Figure 11-1 Device Protection Definition Dialog Box

The Device Protection Definition dialog box appears with a list of devices on the Symmetrix array. There are four columns, each of which can be sorted by clicking on the column header. This display of devices has been prefiltered to prevent the selection of inappropriate devices.

You can select a device from the list and click **Unprotected** to propose removing mirroring for that device. You can click **Mirror** to propose mirroring for unprotected devices. Note that once you redefine a device, the type is displayed in blue italics. Illegal operations are flagged in a subsequent dialog box.

Repeat this operation as required for other devices.

If you create an unprotected device, it should be used only as a BCV or an SRDF device, not as a standard (STD) device.

A device must be in a ready state before protection is added or removed.

Working with Symmetrix Groups

Symmetrix device groups are used by TimeFinder and SRDF to support operations on large numbers of devices. The procedures outlined in this section are accessed through the shortcut menu by selecting **Data Protection, Device Groups**.

ControlCenter uses a Device Group Wizard to automate many of the tasks involved in creating groups for use by SRDF and TimeFinder. this section provides steps for using the wizard to complete the following tasks:

- ◆ Create a device group after you have selected your devices.
- ◆ Create an empty device group from the Device Group folder for future use.

Creating a Device Group After Selecting the Devices

The normal sequence of steps includes the following:

1. Select the devices (all of the same type) from anywhere in the Console. You may find it easiest to open the appropriate folder in the Symmetrix tree and select the devices from there.
2. Use the shortcut menu and select **Data Protection, Device Groups, Create**. The Device Group Wizard dialog box appears (Figure 11-2).




Figure 11-2 Device Group Wizard Dialog Box

3. The Device Group Wizard dialog box reflects your earlier choices for the Symmetrix containing the devices and their device type. If you selected multiple device types, it asks you to specify the **Device group type**. Once you select the device group type, click **Next** to proceed to the Create Device Group dialog box (Figure 11-3).
4. Specify the name of the new group (or existing group name if you are editing values) and the name of the host.
5. Once you specify the name of the group:
 - Click **Next** to continue with device association and proceed to step 6
 - or
 - Click **Finish** to let the wizard make the remaining decisions, and proceed to the final review in step.




Figure 11-3 Create Device Group Dialog Box

After you click Next in the Create dialog box, the Associate Device Group dialog box appears (Figure 11-4).

6. Specify the replica type (BCV, BRBCV, RBCV) and if a Gatekeeper is to be associated with this device group, by selecting the related line in the Association dialog box and click **Next**.




Figure 11-4 Associate Device Group Dialog Box

An association dialog box appears for the replica type you selected as shown in the associate BCV devices example in Figure 11-5 on page 11-8.

7. If you specified a replica type, or gatekeeper, and there were not enough compatible devices to cover the group, the dialog box describes which devices are missing and which are available. Add the missing devices if necessary and click **Next** to associate any other device types or gatekeepers.




Figure 11-5 Associate BCV Devices Dialog Box

8. Once you complete the device association, click **Next** or **Finish** to proceed to the final review.

A message appears stating that ControlCenter is acquiring a lock on the Symapi database.

9. Click **OK**.

The Review dialog box appears.

10. Review your selections and click either **Back** to return to a previous dialog box to revise your selections, or **Finish** to create the group.

Creating an Empty Device Group

The normal sequence of steps includes the following:

1. Either select the Symmetrix array containing the devices, the host on which the device group will reside, or the Device Group folder.
2. From the shortcut menu, select **Data Protection, Device Groups, Create**.

3. The Device Group Wizard dialog box shows the relationship between the selected host and Symmetrix arrays, or between the selected Symmetrix array and hosts, and offers a choice of available **Device types**. Select the device type and click **Next** to proceed to the Create Device Group dialog box.
4. Use the Create Device Group dialog box to specify the name of the new group. Click **Next** or **Finish** to proceed to the final review.
5. A message appears stating that ControlCenter is acquiring a lock on the Symapi database. Click **OK**.
6. Review all the configuration specifications summarized in the Review dialog box. Click **Finish** to create the group.

Notes

- ◆ If you are using GNS to maintain device groups on multiple hosts, be aware that even though data is refreshed every 15 minutes, it is possible for a device group to be renamed or updated from a host other than the one currently in use for device group operations.
- ◆ If additional gatekeeper devices are required, they must be created from SYMCLI.
- ◆ To perform these operations, one of the following must be true:
 - The Storage Agent for Symmetrix must be running on the host.
 - The SYMAPI server must be running on the host and a remote discovery policy is enabled for that host.
- ◆ The Device Group Wizard implements a lock on the associated host's SYMAPI database (`symapi_db.bin`). Once the lock has been set, do not leave a Device Group Wizard session open for extended periods, as this will prevent ControlCenter from updating the Repository with any configuration changes for any Symmetrix arrays being managed through the host, as well as preventing the discovery of new Symmetrix arrays.

RA Groups

RA groups, which can be found in the Symmetrix portion of the tree panel, in the RA Groups folder, can also be used to perform SRDF group operations. RA groups contain all the logical devices associated with a particular remote adaptor.

Local Protection with Symmetrix TimeFinder

This section provides guidelines and an overview of the common tasks used to create and manage business continuance volumes (BCVs) on a Symmetrix storage array. In addition, TimeFinder clones are discussed (available with *Solutions Enabler Version 5.2* and *Engenuity Version 5568* or higher).

TimeFinder Overview

An understanding of TimeFinder BCVs is necessary to fully use ControlCenter for data protection. BCV devices are Symmetrix devices that are specially configured in the Symmetrix array to be dynamic mirrors. Each BCV device has its own host address, and is configured as a stand-alone Symmetrix device. Figure 11-6 illustrates a basic business continuance process using *Establish* and *Split*, two of the EMC TimeFinder operations that can be performed on BCV devices.




Figure 11-6 BCV Configuration

A business continuance sequence first involves *establishing* the BCV device as an additional mirror of a standard Symmetrix device. Once the BCV is established as a mirror of the standard device, it is not accessible through its original device address.

The BCV device may later be separated, or *split*, from the standard Symmetrix device with which it was previously paired. It then becomes available for backup or other host processes through its original device address.

Once host processes on the BCV device are complete, the BCV may again be mirrored to a standard Symmetrix device to acquire new data for other business continuance processes or updating the standard device with the data from the completed business continuance processes.

TimeFinder Devices

TimeFinder uses two basic device types:

- ◆ **Standard (STD) devices** — Devices configured for normal I/O operations on the Symmetrix array.
- ◆ **BCV devices** — Symmetrix devices designed for dynamic mirroring. Such devices have additional attributes that allow them to independently support host applications and processes.

Once a BCV device is established as a mirror of a standard device, those two devices together are referred to as a *BCV pair*.

A BCV device and the standard device it is paired with both reside in the same Symmetrix array.

In addition to the STD and BCV devices, TimeFinder also supports two other device configurations:

- ◆ **RBCV devices** — The level of data replication for the normal SRDF configuration can be increased by treating the SRDF R2 device in the remote Symmetrix as the primary part of a TimeFinder pair and pairing a BCV device to the SRDF R2 device. A BCV device in this configuration is called an RBCV.
- ◆ **BRBCV devices** — The normal TimeFinder configuration has the BCV device in the local Symmetrix array. This local BCV device can simultaneously operate as the R1 device for an SRDF pair with the R2 device in a remote Symmetrix array. This remote R2 device can then operate as the primary device in a TimeFinder pair within that remote Symmetrix. The BCV device in the remote Symmetrix configuration is called a BRBCV.

TimeFinder manages the relationship between a standard storage device (STD) and separately addressable mirrored volumes (BCVs) within the local Symmetrix array.

These mirrored volumes contain a copy of the data while the original device is online for regular I/O operation. After the mirror image is established, you can split it from the standard device, manipulate the data (back it up or perform applications testing), and later reestablish the mirror image with the standard device.

Before using TimeFinder, some of the disks in the Symmetrix array must be configured (through the Device Type Definition command and dialog box) as special disks known as business continuance volumes (BCVs). Each BCV device has its own host address, and is configured as a stand-alone device. You then pair the BCV with a standard Symmetrix volume.

The TimeFinder Process

BCV devices are created by modifying existing STD devices. The devices must be unmapped and unavailable to the host before they can be modified.

1. Define a device as a BCV using Symmetrix Manager.
2. Establish a new pair relationship between an STD device and the BCV device. Once the BCV is established as a replica of the standard device, it is not accessible through its original device address.
3. Split the BCV device from the standard device with which it was previously paired. After a split, the BCV device has valid data and is available for backup or other host processes through its original device address.

Once host processes on the BCV device are complete, the BCV may again be mirrored to a standard device (either the same device to which it was previously attached or a different device). It can then acquire new data for other BCV processes or update the standard device with any new data from the completed BCV processes.

Most TimeFinder operations can be performed on pairs of devices and host-based groups of devices.

TimeFinder Commands The following table summarizes the TimeFinder operations supported in ControlCenter.

Table 11-1 Supported TimeFinder Operations

Operation	Description
Attach Device Group	Attaching BCV pairs is the process of defining a specific BCV device as the preferred BCV device. This pairing eliminates the need to specify a device for each subsequent establish and restore operation. To access this command, select the device group, right-click, and select Data Protection, TimeFinder .
Establish Concurrent	Starting with the Engenuity 5568, you can establish two BCV devices as concurrent mirrors of a single standard device all within the same symmir command line. This relationship is known as a concurrent BCV pair. This feature allows you or an application script to instantly generate two synchronized copies of the standard data.
Cancel Deltamark Session	The Cancel Deltamark Session command demotes the preferred device to a normal BCV device, allowing a different BCV device to be used during a restore or an establish operation.
Detach Device Group	Detaching BCV pairs means to detach a BCV device as the preferred BCV device to pair with the standard device the next time a full establish or full restore is issued without specifying a BCV device to pair with the standard device. To access this command, select the device group, right-click, and select Data Protection, TimeFinder .
Establish	Establishing BCV pairs is the process of copying data from the STD device to the BCV device until both devices are identical. Once they contain exactly the same data, normal TimeFinder operations can commence.
Establish New Pairs	Establishing new BCV pairs is the process of pairing an STD device with a BCV device. This process copies the entire contents of the standard device to the BCV.
Restore	The Restore command copies the contents of the BCV device to the standard (STD) device. You can optionally copy to the STD device only that data that has been written to the BCV device while it was split from the STD device.
Split	Splitting a BCV pair makes each device available to hosts through their separate device addresses.

TimeFinder Clones

The TimeFinder clone feature allows you to make copies of data simultaneously on multiple target devices from a single source device. The data is available to a target's host instantly. You can copy data from a single source device to as many as 16 target devices. A source device can be either a Symmetrix standard device or a TimeFinder BCV device. A target device can also be either an STD or BCV device designated to be a clone.

Unlike a BCV mirror copy, which must be completely synchronized with its source and then split to access the data, the clone copy activation makes data on the clone immediately accessible to its host, even while copying is occurring in the background.

The following table summarizes the TimeFinder clone operations supported in ControlCenter.

Table 11-2 Supported TimeFinder Clone Operations

Operation	Description
Create Clone Copy	Create the relationship between the source and target devices.
Activate Clone Copy	Activate the copy operations.
Recreate	Once a clone session has been created and executed, it can be re-created later in order to incrementally copy changed data to the target device.
Set Mode	Change the mode after the creation or re-creation of a Clone session. The supported modes are Copy, No Copy (the default), and Pre-Copy.
Terminate Clone Copy	End the relationship between the source and target devices.

EMC Snap

You can use EMC Snap™ to perform virtual copying operations that provide a space-saving method of creating instant, point-in-time copies of logical volumes. Snapping to a virtual device (VDEV) creates the appearance of copying volumes by simply copying the original data from changed tracks and the pointers to that data.

Snap operations can be used to create a single instantaneous copy, or used to create multiple copies of the same logical volume at different points in time.

The EMC Snap operation uses two types of devices: VDEV and SAVE. A VDEV device contains pointers to the changed data, while a SAVE device holds the actual data that has been changed.

The following table summarizes the EMC Snap operations supported in ControlCenter.

Table 11-3 Supported EMC Snap Operations

Operation	Description
Create a Snap Session	Create a Snap session by selecting a Symmetrix array or any suitable device under it. Suitable source devices include standard devices, split BCV devices and metadevices. The target device must be a virtual device (VDEV).
Activate a Snap Session	Activating multiple snap pairs with a specific source device means that the Symmetrix system retains changed track information about those virtual devices while these copy sessions exist. By activating various snap pairs over a period of time, you can capture progressive historical snapshots of the specified source device.
Restore Virtual Device Data	Restore data on the VDEV back to the STD device or BCV device from which it originated.
Terminate a Snap Session	A copy session for a snap pair originates when you issue the Create command and ends when you issue the Terminate command. A device cannot participate in two copy sessions simultaneously unless it is a source device that has multiple target devices. To pair an existing target device with a different source device, you need to first terminate the target's original copy session.
Restore Anywhere	Restore data on the VDEV to STD devices or BCV devices <i>other than</i> from which it originated.

Remote Protection with Symmetrix SRDF

Basic SRDF Configuration

SRDF is a business continuance and data replication solution that maintains multiple, real-time copies of logical volume data in more than one physical Symmetrix array location.

Figure 11-7 shows a basic SRDF configuration consisting of a production site and a remote backup site. At the production site, a local host connects to Symmetrix A. At the remote backup site, a second host connects to Symmetrix B. The Symmetrix arrays communicate through SRDF links.

The volumes located on the SRDF-enabled Symmetrix array that are not involved in the SRDF process are referred to as *local volumes* in Figure 11-7.




Figure 11-7 Basic SRDF Configuration

SRDF Volume Types

SRDF designates Symmetrix devices as either *source* volumes, or *target* volumes.

Source Volumes

Source volumes contain production data that is mirrored in a different Symmetrix array. Source volumes are also referred to as *R1* volumes. Updates to a source volume are automatically copied to a target volume in the other Symmetrix array.

A source volume can be paired with a BCV to provide an additional working copy of the data at the same location.

Target Volumes

Target volumes contain a mirrored copy of data from a source volume. Target volumes are also referred to as *R2* volumes. A target volume can also be paired with a BCV to provide an additional working copy of the data at the same location.

Physical Configuration

Before you can use SRDF, the local and remote Symmetrix arrays must each be set up with at least two Remote Link Directors (RLDs) through which the two arrays are linked. The Symmetrix array being mirrored is designated as the source (*R1*); the Symmetrix array maintaining the remote mirror is designated as the target (*R2*). Data is transferred across the SRDF link from the source to the target array. SRDF arrays can be up to 12,000 miles apart.

Before you can perform the SRDF operations described in this section, SRDF must be installed and configured by an EMC representative. This configuration includes the mapping between the *R1* and *R2* devices.

Displaying SRDF Information in Table View

SRDF devices can be found in many locations in the tree panel, but the most efficient path is **Storage Systems, Symmetrix, SymmetrixID, SRDF**. After you select one or more SRDF devices, you can click Properties to display a full range of information about them in the table view, as shown in the SRDF Properties table.

Accessing the SRDF Command Menu

Perform an SRDF operation by selecting the appropriate Symmetrix, SRDF devices, RA group, or device group in the tree panel and selecting **Data Protection, SRDF** from the shortcut menu.

The following SRDF operations are supported in ControlCenter.

Table 11-4 Supported SRDF Operations

Operation	Description
Establish	Establishing (or Copy Source to Target) resumes links and copies data between source (R1) devices and target (R2) devices. Establish will propagate any updates made to the R1 devices while the links were suspended, bringing the R2 devices up to date and completely overwriting the content of the R2 devices with the source device content.
Split	Splitting is an SRDF control operation that suspends SRDF link traffic and read/write enables the R2 device to its local host. The split causes the R2 device to provide an additional copy to the local host.
Restore	Restore operations copy data from target (R2) to source (R1). This operation is useful if, for example, you performed application testing on the R2 devices, production processing was halted on the R1 devices, the testing was successful, and you want to keep the updates.
Suspend Link	Suspending a link breaks all link paths between the selected source (R1) and target (R2) devices, preventing data transfer to R2 volumes. The Suspend Link operation is directed to the Symmetrix array containing the R1 devices. The link can be suspended only if there are no invalid tracks for the source (R1) volumes, and no invalid tracks for the R2 volumes seen on either the source or target volumes.
Resume Link	Resuming links resumes data transfers between the selected source (R1) and target (R2) devices, allowing data transfer to R2 devices. The Resume Link operation is run against the Symmetrix array containing the R1 devices. The link can be resumed only if the R2 device is write-enabled and there are no invalid tracks for the R1 devices on the R2 devices.
Fail Over	The Failover operation causes the target (R2) devices to take over read/write operations for source (R1) devices. This operation halts all I/O activity to the Symmetrix array containing the R1 devices; this will write-disable the R1 devices. This operation is typically performed when you need to transfer I/O operation from the R1 devices to the R2 devices.
Fail Back	This topic describes the steps to execute a source takeover (failback) on devices in the Symmetrix array to which the host is attached.
Mode Control	ControlCenter allows you to change the SRDF mode of operation and attributes for selected device groups.
Create Dynamic Pair	Dynamic pairs are used to create an ad hoc SRDF relationship between two Symmetrix SRDF devices.

Table 11-4 Supported SRDF Operations (continued)

Operation	Description
Delete Dynamic Pair	Dynamic pairs are used to create an ad hoc SRDF relationship between two Symmetrix SRDF devices. This command breaks an existing pair apart.
Update Source	Update Source (R1) operations update the source device with the changes from the target (R2) device while the target device is still operational with its local host(s). This operation is required if you have previously performed a failover operation from the R1 volume to the R2 volume.
Swap	Swap the R1 and R2 devices that compose an SRDF pair. You can swap R1 and R2 devices for data recovery purposes, when the remote site has a greater degree of data integrity.

Advanced SRDF Commands

ControlCenter provides a set of advanced SRDF commands that allow you to explicitly control the availability or state of SRDF devices. You access the commands through the Console: **Data Protection, SRDF, Advanced** or through the shortcut menu.

These commands require the Advanced SRDF permission level (refer to *Assigning Permissions* on page 1-6) and should only be used by advanced SRDF users. Improper use can compromise data integrity.

Table 11-5 Advanced SRDF Commands

Operation	Description
R2 NR if Invalids	Sets R2 devices to Not Ready if invalid tracks exist.
R2 RW Disable	Blocks read and write operations between the target (R2) devices to their local host. This option enables you to set a device to the Not Ready state on the R2 side by making the device Not Ready on the RA.
SRDF Ready Device	Sets an R1 or R2 device to the Ready state for its local host.
SRDF Not Ready Device	Sets an R1 or R2 device to the Not Ready state for its local host. No data will be transmitted over the SRDF links, but I/O may continue for the R1 device.
SRDF RW Enable	Enables read and write operations on either the R1 or R2 device for the local host.
SRDF Write Disable	Write-disables operations on either the R1 or R2 device to the local host.
SRDF Refresh	Refreshes the contents of either the R1 or R2 device from the other side of the pair.
RDFA Enable	Enables consistency protection on the specified RDF-Asynchronous backed devices.
RDFA Disable	Disables consistency protection the specified RDF-Asynchronous backed devices.

SRDF Modes of Operation

There are a variety of options to control the overall mode of operation for SRDF. Modes are changed in the SRDF Mode Control dialog boxes (Figure 11-8 and Table 11-6). These modes are selectable based upon distance, performance, and speed of recovery requirements.




Figure 11-8 Mode Control Dialog Boxes for Device Pairs and Groups

Table 11-6 Configuration Modes

Mode or Option	Description
Force	Overrides some of the normal checking for SRDF operations. For the establish operation, an R1 device will not be processed if an appropriate R2 device cannot be found. For the restore operation, this means that it will skip, but not reject, devices in the group that are NEVER ESTABLISHED, or are not properly paired and split with R2s associated with the group. This will also allow devices that are SPLIT BEFORE RESTORE to be restored. Note: Exercise extreme caution when using this option as it may result in data loss if used improperly.
All	Specifies that the SRDF control operation is for both standard SRDF and BCV SRDF devices that belong to the same RA group.
Remote BCV	Specifies that the SRDF control operation is intended for remotely associated BCV devices (connected through the local STD device) that are configured as BCV devices in the SRDF hop 2 link. The remote devices are known as RBCV devices.

Table 11-6 Configuration Modes (continued)

Mode or Option	Description
Bypass Lock	Bypasses Symmetrix exclusive locks for the local and/or remote arrays during SRDF operations. Use this only if you are sure that no other SRDF operation is in progress to the local and/or remote Symmetrix arrays.
BCV	Specifies that the SRDF control operation is intended for one or more R1 BCV devices in the SRDF hop 1 link that are associated with the group.
BCV Remote BCV	Specifies that the operation will be targeted at the remote side of any R1 BCVs associated with the group and BCV Remote BCV (BRBCV) device(s) that are associated with the device group. Only applies to device group operations.
Mode: Sync	Notifies host of successful I/O only after target (R2) device signals success.
Mode: SemiSync	Notifies host of successful I/O after source (R1) device signals success.
Mode: Asynchronous	Allows a group of SRDF R1 devices to group their I/Os into a cycle and periodically destage the data to the R2 side. Only supported for device groups and pairs of Symmetrix DMX.
Domino Effect	Causes an R1 device to become not ready to its host, and all I/O activity ceases with that device if the R2 volume fails or a link failure occurs. When the fault condition is corrected, you must manually make the master device ready. This feature ensures that a remotely mirrored pair is always synchronized.
Adaptive Copy: Disk	Transfers data from the R1 device to the R2 device and does not wait for confirmation. This mode is intended to be a temporary SRDF operating state and is designed for situations requiring the transfer of large amounts of data without loss of performance.
Adaptive Copy: Change Skew	Modifies the adaptive copy skew threshold. When the skew threshold is exceeded, the remotely mirrored pair operates in the predetermined SRDF state (synchronous or semi-synchronous). As soon as the number of invalid tracks drop below this value, the remotely mirrored pair reverts back to the adaptive copy-write pending mode. The skew value is configured at the device level and may be set to a value between 0 and 65,534 tracks. For devices larger than a 2 GB capacity drive, a value of 65,535 can be specified to target all the tracks of any given drive.
Adaptive Copy: Write Pending	Transfers data from the R1 device to the R2 device and does not wait for confirmation. This mode is ideal for situations when a large amount of data must be transferred to remote devices and performance must not be compromised at the local site.

Asynchronous Mode

SRDF Asynchronous (also known as SRDF/A) mode allows a group of SRDF R1 devices to group their I/Os into a cycle and periodically destage the data to the R2 side. Use this feature when the cost of immediate replication is not required for certain classes of data.

- ◆ SRDF/A is supported only on pairs of Symmetrix DMX Series systems.
- ◆ All operations must be performed on either an RA group or a device group which contains all devices under that RA group. There are no controls for individual device pairs.
- ◆ You can only configure one SRDF/A group per DMX system.
- ◆ Dynamic SRDF is not supported for SRDF/A-backed devices.
- ◆ Concurrent SRDF devices that have an SRDF/A-backed mirror are not supported.

The Mode Control dialog box allows you to specify the Asynchronous mode, and the Device Group wizard supports the creation of groups of SRDF/A pairs. Use RDFA Enable to enable consistency protection on SRDF/A-backed devices, and RDFA Disable commands to disable consistency protection.

All other relevant SRDF/A operations are supported in the normal SRDF dialog boxes.

Concurrent SRDF

In an SRDF configuration, a single source (R1) device can be concurrently remotely mirrored to two target (R2) devices. This feature is known as concurrent SRDF and is supported with ESCON and Fibre Channel interfaces. Concurrent SRDF is valuable for duplicate restarts or disaster recovery, or for increased flexibility in data mobility and migrating applications.

Concurrent SRDF is supported in Enginuity levels 5567 and higher.

Operating Mode Restrictions

Concurrent SRDF supports each of the two remote target devices operating independently (but concurrently) in any of the following SRDF modes:

- ◆ Synchronous
- ◆ Semi-synchronous
- ◆ Adaptive copy-disk mode
- ◆ Adaptive copy-write pending mode
- ◆ Asynchronous mode (Enginuity 5671 and higher)

While the modes for the two remote target devices can be the same or different, the following restrictions apply:

- ◆ Each of the two concurrent mirrors must belong to a different RDF (RA) group.
- ◆ You cannot have one mirror in synchronous and the other in semi-synchronous mode.
- ◆ You can only enable concurrent SRDF for asynchronous mode with Engenuity 5671 and higher.

Remote Data Copy

All SRDF operations that emanate outward from the R1 device can be performed on concurrent SRDF configurations. Failback, restore, and R1 update operations cannot be performed concurrently, as data cannot be copied from two R2 devices to a single R1 device.

To resolve this problem, you must use the Remote Data Copy option. Use this option when you want to restore data to the R1 device and to any other concurrent R2 devices.

(This implies that only a single R2 device has the correct data.) The data is first copied from the specified R2 device to the R1 device; then, when the concurrent link is ready, data will also be copied to the concurrent SRDF R2 mirror(s). This option is available with the failback, restore, and update operations.

The Remote Data Copy option sets the state of the concurrent link to Write Disabled or Not Ready.

Device Groups and RA Groups

When concurrent SRDF is enabled in the Symmetrix array, a device group can contain up to two RA groups. BCV, RBCV, and BRBCV devices can be added from either RA group in the device group, but not from both. For example, you can create a standard SRDF1 device group and add a device from SRDF group 1 to it, followed by a device from SRDF group 2, followed by a concurrent SRDF device from SRDF groups 1 and 2. If you add a third RA group, a failure is returned.

To create a device group for the concurrent SRDF devices and initially synchronize (establish) the devices across the concurrent SRDF:

1. Create an R1/R2 device group by selecting the SRDF folder within a Symmetrix folder.
2. Add all devices to the group.
3. Establish the concurrent group.

SRDF: Consistency Groups

A consistency group is a group of Symmetrix SRDF devices specially configured to act in unison to maintain the integrity of a database distributed across multiple SRDF arrays. Consistency groups maintain coherency for an SRDF configuration by monitoring data propagation from the source (R1) devices in a consistency group to their corresponding target (R2) devices.

Consistency groups require PowerPath 2.1 or higher.

When a typical DBMS application updates a database, it first writes to the disk containing a log, it then writes the data to the actual database datafiles, and finally writes to the log volume to indicate these write I/Os (log, database) are related, and each I/O is not issued until the prior I/O has successfully completed.

Even in a remote disk copy environment, data consistency cannot be ensured if one of these I/Os was remotely mirrored, but its predecessor was not remotely mirrored.

This could occur, for example, in a rolling disaster where there is a communication loss that affects only a portion of the disk controllers that are performing the remote copy function.

SRDF-established consistency groups can prevent this situation by using the PowerPath pseudo-device driver to intercept any I/O to a disk device that cannot communicate to its remote mirror. The consistency protocol is to then suspend the remote mirroring for all devices defined to the consistency group before the intercepted I/O and return control to the application. In this way, consistency groups prevent dependent I/O from getting out of sync, thus ensuring the integrity and consistency of the data at the remote site.

Creating Groups and Adding Devices

You must use the EMC Solutions Enabler SYMCLI to create consistency groups and add members.

The ControlCenter Console provides the capability to monitor the status of these groups.

Working with EMC Solutions Enabler SYMCLI

The EMC Solutions Enabler SYMCLI is a specialized set of commands invoked from the host operating-system command line (shell). The Solutions Enabler SYMCLI is used in single command-line entries and scripts to perform control operations on Symmetrix devices and data objects, as well as to monitor device configuration and status.

The SYMCLI used in conjunction with the ControlCenter GUI can provide you with flexibility in configuring and managing data protection.

Many data protection operations performed through ControlCenter can also be performed through the SYMCLI. Refer to the EMC Solutions Enabler SYMCLI documentation for specific details.

For example, the SYMCLI can be used to create R1/R2 devices, or to automate complex data protection processes through the use of scripts. You can use ControlCenter to view device groups created in the SYMCLI, determine the location of device groups, monitor a SYMCLI script, or test the individual script processes in a simple case before applying the SYMCLI script to the production environment.

The ability to use both the SYMCLI and ControlCenter to manage and configure data protection requires care in executing operations on devices in production environments.

For example, you could change the state of an operation with ControlCenter but forget to reflect that change in a script that you run through the SYMCLI, causing the scripted operation to fail.

Be aware of the possibility of conflict between ControlCenter operations and the SYMCLI scripted commands you may have running on the same devices.

Another possible conflict is created if you attempt ControlCenter operations on individual devices within a group. Operations should be performed on all devices within a group. If necessary, the device can be removed from the group before making any changes to it.

EMC recommends that you perform ControlCenter operations on device groups rather than individual devices within the device groups.

Creating BCV Devices

BCV devices are created by modifying existing STD devices. The devices must be unmapped and unavailable to the host before they can be modified.

Use the following procedure to create BCV devices:

1. From the Storage folder in the tree panel, go to the Unmapped Devices folder for the Symmetrix array to which you are adding BCVs. Right-click and select **Configure, Device Type Definition**. The BCV/DRV Definition dialog box appears (Figure 11-9).




Figure 11-9 BCV/DRV Definition Dialog Box

Each column can be sorted by clicking on the column header.

2. Select a device from the Device column.
3. Click **BCV** to redefine the drive. Note that once you reassign a drive, the type is displayed in italics. Illegal operations are flagged in the Status panel.
4. Repeat this operation as required.
5. The configuration process commences when you click **OK**.

Determining Device Size

Determine the size of the STD devices in the device group as follows:

1. From the Hosts folder in the tree panel, expand the attached host for the Symmetrix devices you are adding.
2. From the Device Groups folder, expand the device group and right-click the **Standard Devices** folder. The properties of the STD devices in that device group are displayed in the right panel.
3. Record the sizes of the STD devices and leave the properties view open for use in future steps.

Finding Available BCVs

Find BCVs that have the correct size and are connected to the correct host as follows:

1. From the Storage folder in the tree panel, expand the Symmetrix array containing the STD devices.
2. Expand the **Mapped Devices** folder, right-click the **BCV Devices** folder (if present), and select **Properties**. The mapped BCV device properties are displayed in the Properties view (Figure 11-10).




Figure 11-10 Device Properties

3. Expand the **Unmapped Devices** folder, right-click the **BCV Devices** folder, and select **Properties**. Now both the mapped and unmapped BCV device properties display in the Properties view.
4. Sort the BCV devices (by clicking the column heading) for size and then port to determine which BCVs are the correct size and which are connected to the correct host.

Adding BCVs to the Device Group

Add BCVs to the device group with the following steps:

1. Split the view in the right-hand panel and open a TimeFinder view (select **Data Protection**, **TimeFinder** from the toolbar). You now have a device properties view in the top panel and a TimeFinder view in the lower panel (Figure 11-11).




Figure 11-11 Split View Showing Device Properties and TimeFinder

2. Drag each BCV device into the TimeFinder view to determine the state of the device. Verify that the device is not already paired with another device.

3. Once you determine that a BCV is the correct size, can be seen by the correct host, and is not paired with another device, drag the device icon into the BCV folder under the correct device group.
4. Click **Yes** at the Add devices to a device group confirmation dialog box.
5. Repeat steps 2 through 4 until you finish adding BCV devices to the device group.

Monitoring TimeFinder Operations

Display a dynamically updated table listing all the BCV devices and their characteristics as follows:

1. Select a Symmetrix array from the tree panel.
2. Select **Data Protection, TimeFinder** from the toolbar. The TimeFinder view appears in the information panel.

The following table describes the columns in the TimeFinder view.

Table 11-7 TimeFinder View Description

Column Heading	Description
Symmetrix	Serial number of the Symmetrix array.
STD Host - Dev Group	Name of the device group containing the STD device.
STD	Device ID of the STD device.
BCV	Device ID of the BCV device.
BCV Host - Dev Group	Name of the device group containing the BCV device.
State	Current state of the BCV pair.
Last Action	Date and time of the last TimeFinder operation (passed from the SymmAPI, and stored in the Repository).
STD inv trks	Number of invalid tracks on the STD device.
BCV inv trks	Number of invalid tracks on the BCV device.
MBs left	Remaining storage capacity, measured in megabytes.
MBs/sec	Data flow rate between the BCV pair, measured in megabytes per second.
Time left (sec)	Time remaining to complete a data transfer.
CBCV	Specifies if configured for concurrent BCV.

Table 11-7 TimeFinder View Description (continued)

Column Heading	Description
STD QoS	QoS setting for a STD device (1-10).
BCV QoS	QoS setting for a BCV device (1-10).
BCV Config	Configuration of a BCV device.

The creation of the R1/R2 mirrored devices is done either during Symmetrix configuration or through the EMC Solutions Enabler SYMCLI.

Monitor SRDF Operations

To display dynamic data about the status of SRDF devices, select the appropriate Symmetrix array, and then select **SRDF** from the Data Protection menu.

The SRDF view displays all relevant information relating to SRDF configuration and operational status.

Table 11-8 SRDF View Descriptions

Column Heading	Description
R1 Host - Dev Grp	Name of the SRDF device group on the R1 Symmetrix array.
R1 Symmetrix	Serial number of the Symmetrix array housing the R1 device.
R1	ID of the R1 device.
R2	ID of the R2 device.
R2 Symmetrix	Serial number of the Symmetrix array housing the R2 device.
R2 Host - Dev Grp	Name of the SRDF device group on the R2 Symmetrix array.
R1 State	State of the R1 device. Possible values are: READY, NOT READY, WRITE DISABLED, NA, and MIXED.
R2 State	State of the R2 device. Possible values are: READY, NOT READY, WRITE DISABLED, NA, and MIXED.
Pair State	State of the pair. Possible values are: INVALID, SYNC_IN_PROG, SYNCHRONIZED, SPLIT, SUSPENDED, FAILED_OVER, PARTITIONED, R1_UPDATED, R1_UPDINPROG, and MIXED.
Link Status	Status of the link. Possible values are: Ready, Not Ready.
Rem inv on R1	Number of remote invalid tracks on the R1 device.

Table 11-8 SRDF View Descriptions (continued)

Column Heading	Description
Rem inv in R2	Number of remote invalid tracks on the R2 device.
Loc Inv on R1	Number of R1 invalid tracks on the R1 side.
Loc Inv on R2	Number of R2 invalid tracks on the R2 side.
R1 RA Grp	I/O of the R1 RA group.
R2 RA Grp	I/O of the R2 RA group.
Uncommitted Tracks	Number of tracks that have not been committed to the R2 device.
R2 Consistent	Indicates whether the data on the R2 side of an asynchronous SRDF pair is in a consistent state. A value of NA is displayed for non-asynchronous SRDF pairs.
Consistency Protection	Indicates if consistency protection is enabled.
R1/R2 Time Lag	Number of seconds that the data on the R2 is behind the data on the R1.
Async SRDF	Yes or No indicates support for Asynchronous SRDF. The RA group must also be in asynchronous mode before asynchronous operations can be performed.
Mode	Level of synchronization between R1 and R2 devices. Possible values are: SYNCHRONOUS, SEMI-SYNCHRONOUS, ADAPTIVE_COPY, and MIXED.
Domino	State of the Domino attribute which forces data on the R1 and R2 devices to be synchronized. Possible values are: ENABLED, DISABLED, and MIXED.
Ad. Copy	Adaptive copy-disk mode and adaptive copy-write pending mode Possible values are: • Enabled: WP Mode • Enabled: Disk Mode • Mixed • Disabled
AC Skew	Number of invalid tracks allowed when in adaptive copy mode. Possible values range from 0 to 65535.
Link Status	Status of the link. Possible values are: READY, NOT READY, WRITE DISABLED, NA, and MIXED.
R1 SA Status	SA status of R1. Possible values are: READY, WRITE DISABLED, and NA (if there is no front-end director).
R2 SA Status	SA status of R2. Possible values are: READY, WRITE DISABLED, and NA (if there is no front-end director).

Table 11-8 SRDF View Descriptions (continued)

Column Heading	Description
RA Status	Status of the remote link director. Possible values are: READY, NOT READY, WRITE DISABLED, NA, and MIXED.
Dev SRDF Status	Status of the SRDF device. Possible values are: READY, NOT READY, WRITE DISABLED, NA, and MIXED.
CRDF	Specifies concurrent RDF or not.
DyRDF	Specifies dynamic RDF capable or not.
R1 QoS	QoS setting for the R1 device (1-10).
R2 QoS	QoS setting for the R2 device (1-10).
R1 Config	Configuration of the R1 device.
R2 Config	Configuration of the R2 device.
Async SRDF:	Yes or No. indicates support for Asynchronous SRDF. The RA group must also be in asynchronous mode before asynchronous operations can be performed.
Uncommitted tracks:	Number of tracks that have not been committed to the R2 device.
R2 Consistent:	Indicates whether the data on the R2 side of an asynchronous SRDF pair is in a consistent state. A value of NA is displayed for non-asynchronous SRDF pairs.
Consistency Protection:	Indicates if consistency protection is enabled.
R1/R2 Time Lag:	Number of seconds that the data on the R2 is behind the data on the R1.

Protecting Data on CLARiiON Storage Arrays

ControlCenter lets you use EMC SnapView™ software to protect data on EMC CLARiiON FC4700 and CX series subsystems. EMC SnapView takes snapshot copies of production data and stores the copies on the subsystem. You can use these copies for testing purposes, or you can back them up, without affecting the production data that resides on the production host.

Creating a SnapView snapshot involves:

1. Identifying the LUNs of which you want to take a snapshot.
2. Creating the snapshot, either through the ControlCenter Console or with the Navisphere Manager or CLI.
3. Adding the snapshot to a storage group, other than the storage group that contains the target LUN (the LUN you snapped), so you can access it from the hosts connected to that storage group.
4. Making the snapshot active (called starting a session). You must use the Navisphere Manager or CLI to start or stop a session, which you can do from the production host or another host that can see the snapshot.

When you start a SnapView session I/O to the target LUN continues. SnapView monitors for any changes to the target LUN. If the LUN changes, for example the production host writes to it, SnapView copies the original block (snapshot session or snapshot copy) into a private area on the array, called the SnapView Save Area or snapshot cache, before modifying the original data. For any changed block, the copy happens only once when the block is first modified. You can have up to eight active copies per LUN. See Figure 11-12 on page 11-34.

5. Use the Navisphere Manager or CLI to give attached hosts access to the snapshot cache, for accessing the snapshot copies.
6. Use the snapshot copies as you want.
7. Use the Navisphere Manager or CLI to stop the snapshot session.

You can perform most SnapView functions through the ControlCenter Console, but to start or stop a snapshot session, or make the snapshot copies available to other hosts, you use the Navisphere CLI. See your Navisphere and EMC SnapView documentation for full details about using SnapView with CLARiiON subsystems. For more information on CLARiiON Storage Groups, refer to *Masking* on page 7-39.




Figure 11-12 SnapView Snapshot Copies

Protecting Data on HP StorageWorks Storage Arrays

ControlCenter allows you to control host access to units (devices) on HP StorageWorks Enterprise Modular Arrays (EMAs), which are sometimes referred to by their controller model (ControlCenter currently supports EMA array with HSG80 controllers). You can also set the unit offset, which specifies the range of units the connected hosts can access.

Enabling or Disabling Host Access to Units

To enable or disable host access to units on the array:

1. In the tree panel, expand the **Storage Systems** folder.
2. Locate the StorageWorks array on which you want to configure unit access.
3. Right-click the array and select **Explore, Storage Agent for HP StorageWorks**. A new window appears with the array(s) in the selection tree on the left.
4. Right-click the StorageWorks array on which you want to configure unit access and select **Units**. The Units icon appears below the array.
5. Expand the **Units** icon to display the units on the array.
6. Continue to *Enable* or *Disable*.

Enable

7. Right-click a unit and select **Enable Access Connection**. The Enable Access Path dialog box appears.
8. For **Select Connection**, select the connection you want to connect to the unit.
Select **ALL** if you want to enable access to all connected hosts.
9. Continue to step 12.

Disable

10. Right-click a unit and select **Disable Access Connection**. The Disable Access Path dialog box appears.
11. For **Select Connection**, select the connection you want to disconnect from the unit.
Select **ALL** if you want to disable access for all connected hosts.
12. Click **OK** to apply changes and exit.

Setting Unit Offset for a Connection

To set the unit offset for hosts connected to the array:

1. In the tree panel, expand the **Storage Systems** folder.
2. Locate the StorageWorks array on which you want to set the unit offset for a host connection.
3. Right-click the array and select **Explore, Storage Agent for HP StorageWorks**. A new window appears with the array(s) in the selection tree on the left.
4. Right-click the StorageWorks array on which you want to set the unit offset for a host connection and select **Units**. The Units icon appears below the array.
5. Expand the **Units** icon.
6. Right-click a unit and select **Access Connections**. The Access Connections icon appears.
7. Expand the **Access Connections** icons to display the current host connection(s) to that unit.
8. Right-click the host connection and select **Set Unit Offset**. The Set Connection Unit Offset dialog box appears. The host connection appears next to **Connection**.
9. For **New Unit Offset**, enter the new offset value for the connection.
10. Click **OK** to apply the new offset and exit.

Managing Host Storage Resources

This chapter provides an introduction to host storage resource management. Host storage resource management includes viewing and exploring host resources, identifying free space, and proactively managing growth and performance-related storage problems.

This chapter only briefly mentions extending file systems. For a complete discussion of host storage allocation, refer to Chapter 10, *Allocating or Deallocating Storage*.

This chapter contains the following sections:

- ◆ Viewing File Systems, Devices, and Their Relationships..... 12-2
- ◆ Working with Windows and UNIX Files and Directories..... 12-6
- ◆ Working with MVS Host Resources 12-13

Viewing File Systems, Devices, and Their Relationships

The EMC ControlCenter Console presents considerable information about hosts in the tree view, plus the Properties and Relationship views. At the highest level, you can easily view:

- ◆ Lists of file systems, host devices, and volume groups (open systems) and volumes (MVS).
- ◆ Operating system, service pack, and other properties.
- ◆ Size, free space, and other status information.
- ◆ End-to-end configuration information from the host to the physical storage.

ControlCenter regularly gathers updated information on file system and device properties and configurations. The ControlCenter administrator regulates how often updated information is gathered.

Beyond the high-level information about your host resources, you can gather detailed information about files and directories (open systems) and data sets (MVS) on demand. You can issue commands against resources to view and change their configuration. You can also gather storage and performance data in reports.

Table 12-1 Methods of Viewing Host Information

Type of host information	Where to view the data	How the data is collected	Where to read in this chapter
Windows and UNIX hosts and their file systems and devices	• Tree view • Properties view • Relationships view	High-level host data, including basic properties, is collected by data collection policies. Host files and directories and other detail information is collected when you click the shortcut menu command for it.	Continue reading the current section.
MVS hosts and their devices	• <i>High-level info (hosts and volumes):</i> Tree view, Properties view, Relationships view • <i>Detail info (Data sets and detailed reports):</i> The agent window for the desired agent	Hosts and their devices are collected by data collection policies. Most detail information is collected on an internal schedule by the agent and shown when you click the menu command for it.	For Properties and Relationship views, continue reading the current section. For data sets and MVS reports, refer to <i>Working with MVS Host Resources</i> on page 12-13.

Viewing File Systems and Their Properties

To view Windows and UNIX file systems and their properties:

1. In the tree, expand **Hosts**, the desired host, and **File Systems**. A list of file systems appears in the tree.
2. For more information, right-click **File Systems**, and then select **Properties**. A table appears in the target panel (as shown in Figure 12-1).




Figure 12-1 Viewing Properties of File Systems

3. To find out how recent the data is, scroll to the far right in the table and read the Last Update column.

Viewing Host Devices and Their Properties

To view host devices (host “physical disks”) and their properties:

1. In the tree panel, expand **Hosts**, the desired host, and **Host Devices**. A list of devices appears in the tree.
2. Right-click **Host Devices**, and then select **Properties**. A table appears in the target panel (as shown in Figure 12-2).
3. To find out how recent the data is, scroll to the far right in the table and read the Last Update column.

Figure 12-2 shows the host devices displayed in the tree and a target panel.




Figure 12-2 Viewing Host Devices and Their Properties

Relating Host Resources to Storage Array Volumes

You can view the relationships among host resources (file systems, volumes, devices, MVS volumes) and the storage array devices on which they reside. You can switch between a map view and a table view. The following example shows this procedure for Windows and UNIX.

To view a map of file systems and the devices on which they reside:

1. On the toolbar, click **Relationship**. The active target panel becomes a Relationship view.
2. In the tree, right-click **Hosts** and select **Arrange By, Type**. Folders for each host type appear.
3. Expand the **Windows Hosts** folder and double-click a host.
A map appears in the target panel, showing the relationship between the host file systems and host devices.
4. In the tree panel, double-click a host of another type, for example, a UNIX host. The relationships for that host appear in the same view as the first host (Figure 12-3).




Figure 12-3 Viewing Relationships (Map)

- On the Relationship view title bar, click **Table**.

The file system and disk relationship changes to a tabular format (as shown in Figure 12-4).

Relationship - Filesystems - /sales,E:\							
File System	Volume Group	Logical Volume	Host Device	Director	Port	Device	
/sales	salesdg	salesvol	/dev/rdisk/c2t2d18s2	FA-3A	0	03C	
E:\	rootvg	E:\	\\.PHYSICALDRIVE1				

Figure 12-4 Viewing Relationships (Table)

Working with Windows and UNIX Files and Directories

In the current release, Windows and UNIX functions previously performed in a separate agent window are performed in the main window of the Console.

The table gets you started performing some common tasks on open systems hosts.

Table 12-2 Common Tasks on Windows and UNIX Hosts

Task	Steps
Explore Windows and UNIX hosts	In the tree, expand the host and the subfolders beneath it. Select the file systems, devices, and other resources you desire, and add them to a Properties view using any method.
Extend a file system	In the tree or any view, right-click the file system and select Allocation, Extend Filesystem .
Mount a file system	In the tree or any view, right-click the file system, and then select Host, Mount .
List all files and directories	In the tree, expand Files and Directories under the host.
Back up a UNIX file system with <code>tar</code>	Right-click the UNIX file system, and then select Host, Backup .
Perform most storage-related commands on Windows and UNIX hosts	Right-click the host, file system, device, or other resource you desire, and then select Host and then the command you desire.
Explore and manage Windows users, groups, services, Registry, and other non storage entities	In the tree, expand System Information under the Windows host.
Explore and manage UNIX users, groups, page space, processes, and other non storage entities	In the tree, expand System Information under the UNIX host.

The following tasks use the Host Agents for AIX, HP-UX, or Solaris to recapture inefficient storage.

Recapturing UNIX Storage

Recovering Disk Space From Obsolete Files, Log Files, and Temporary Files

On Windows and UNIX systems, you can quickly identify and delete unneeded files, freeing storage space for more urgent uses.

To identify and remove obsolete files of large size:

1. Right-click a UNIX file system anywhere in the Console, and then select **Host, Recover Disk Space, Obsolete Large Files**. The Find Obsolete Large Files dialog box appears.

On a large host or on a host with large NFS mounted directories, searching for files can take many minutes.

2. Complete the dialog box.
 - Specify the minimum file size of the files to be listed.
 - Specify the minimum age in days of the files to be listed.
 Click **OK**.
3. In the resulting display, sort the largest files to the top by clicking the Size column heading.
4. Identify files you no longer need.
5. Right-click an unneeded file and select **Remove**. Repeat for all files you want to delete.
6. Use a similar procedure for other unneeded files. Table 12-3 shows other types of files you can quickly identify and remove.

Table 12-3 Candidate Files for Space Recovery (UNIX)

Type of File	Action in Agent Tree
Log files	Right-click the file system, and then select Host, Recover Disk Space, Recover Disk Space, LogFiles .
Core dump files	Right-click the file system, and then select the Host, Recover Disk Space, Core Dump Files .
Temporary files	Right-click the file system, and then select Host, Recover Disk Space, Recover Disk Space, tmp Files .

Recapturing Windows Storage

The following tasks use the Host Agent for Windows to recapture inefficient storage.

Locating Large Files and Their Owners

To locate large files and their owners:

1. Right-click the drive letter and select **Host, Search**.
2. In the Search dialog box, specify a start directory and the size range of the files.

Click **OK**. The Search Results dialog box appears.

3. To sort by size, click the Size column heading.
4. Right-click each file and select **Edit**. The Properties dialog box shows the owner of the file.

Compressing Files, File Systems, and Partitions

You can search for large files and folders and compress them. When you activate compression, Windows compresses the files. When future users access or save compressed files, Windows uncompresses them at read time and compresses them at write time.

To compress all files written to a folder or file system:

1. Explore the host and find the folder or file system.
2. Right-click the file or folder and select **Host, Compress**.

Automatically Backing Up and Clearing Event Logs

Operating system event logs can consume large amounts of space. You can configure alerts and autofixes to back up event logs and clear them automatically.

Alerts

- ◆ Application Event Log Size (KB) Limit
- ◆ Security Event Log Size (KB) Limit
- ◆ System Event Log Size (KB) Limit

Autofixes

- ◆ **Execute Clear The Event Log** automatically clears the event log when the alert triggers.
- ◆ **Execute Backup and Clear The Event Log** first backs up the log, and then clears it.

Refer to the online Help for the backup file naming convention and system requirements for these alerts.

Recapturing Novell Storage

The following tasks use the Host Agent for Novell to recapture storage on Novell NetWare Servers.

Locating and Deleting Large Log Files

NetWare Servers can contain many log files, which can grow large.

1. Explore a NetWare Server, and then right-click and select **Administration, Log File Explore**. The log files appear.
2. Sort the log files by size.
3. Right-click a log file to view its contents, head, or tail, and then to delete it if desired.

For the log files on a single volume rather than on a server, use a similar procedure. In step 1, explore the volume you desire.

Compressing Files

To compress a file (up to 256 MB in size) on a NetWare Server:

1. Explore and select the file to compress.
2. Right-click the file and select **Compress Immediate**.

Purging Deleted Files

Novell retains deleted files, allowing users quick recovery from inadvertent mistakes. Purge older deleted files to free disk space. To purge a deleted file:

1. Explore and select the deleted file that you want to purge.
2. Right-click the deleted file and select **Purge**.

Migrating Files to Offline Media

You can migrate old or rarely accessed files to offline media. To migrate a file already configured for Data Migration:

1. Explore and select the file that you want to migrate to secondary storage.
2. Right-click the file and select **Migrate**.

Monitoring Novell Space

ControlCenter allows you to monitor your Novell NetWare 4.x and NetWare 5.x servers through the following alerts that you set to warn you of impending problems with your NetWare servers:

- ◆ Deleted File Space Threshold Alert
- ◆ Large File Alert
- ◆ Space Usage Alert
- ◆ User Quota Alert

- ◆ Volume Percent Free Space Alert (Novell)
- ◆ Volume Free Space Alert (Novell)

Refer to the *Alerts* topic in the ControlCenter online Help for the concepts and procedures required to set these alerts.

Managing UNIX Storage to Increase Performance

Use ControlCenter to manage storage-related performance on UNIX hosts, monitor swap space, and add more swap space if necessary. Also, ControlCenter allows you to explore and manage processes.

Monitoring Swap Space

Use the following alerts to monitor swap space.

- ◆ Swap Space Percent Free alert
- ◆ Swap Space Megabytes Free alert

If the alert triggers at a critical or fatal severity, add page spaces to improve performance. Also, explore processes and change priority or kill certain processes.

Adding Page Spaces to Improve Performance

Add file-system page spaces to an HP-UX or Solaris host to boost the amount of virtual memory available to it. While this decreases storage space on the host, it can improve overall performance.

To create a new page space for an HP-UX or Solaris file system:

1. In the tree, expand **Hosts**, the host you want to explore, **System Information**, and then **Page Space**. A list of the host's device page spaces appears in the tree.
2. Right-click the **Page Space** folder and select **Host, Create Page Space for File Type**.

The Create a File System Paging Space dialog box (HP-UX or Solaris) appears, which prompts you to specify the location and size for the new page space.

Generally, to determine the amount of page space required for a host, double its RAM size. Heavily used hosts will require even more paging space.

Exploring Processes

You can view the processes running on a UNIX host. Then, identify and kill processes that are taxing server memory resources, kill unnecessary processes, or identify the resources a process consumes.

To explore processes:

1. In the tree, expand **Hosts**, the host you want to explore, **System Information**, and then **Processes**. A list of the host's processes appears in the tree.
2. Right-click an individual process and select **Properties** to see the UNIX Processes view in the target panel. The view lists detailed information about each process selected.

Also, you can right-click a process to:

- ◆ Learn the resources consumed by the process.
- ◆ Change the priority of the process.
- ◆ Kill the process.

Managing Windows Storage to Increase Performance

Manage Windows performance by creating performance baselines, and monitoring and fixing bottlenecks.

Creating Performance Baselines

Use the recordings feature to log performance data. You can then examine the logs to create baselines of normal performance behavior.

To record performance statistics:

1. Right-click the host, and then select **Others, Agents, Host Agent for Windows, Setup, Recordings**. The Recordings dialog box appears.
2. Complete the Recordings dialog box with the sampling interval and the days you want sampling to occur for each resource: paging file, physical storage, logical storage, cache, server, process, operating system, and memory. Click **OK**.
3. After a day, check the proper execution of the recordings.
 - ControlCenter writes the recording information to the `\hdata` subdirectory of the agent's working directory on the Windows host.
 - The recording files are comma-delimited, plain-text files.
 - The naming convention for the files is:
`MNR_categoryrecorder.csv`, where *category* is the name of one of the statistical categories, such as memory or cache.
4. After a few more days of processing, use the output files to analyze the performance data.

Detecting and Fixing Bottlenecks

The first step in diagnosing a disk bottleneck is ensuring that a memory bottleneck is not the real cause.

To find the cause of a memory bottleneck, use snapshots that show an instantaneous view of performance objects such as logical disks, memory, and paging files, along with their relevant performance counters. Excessive I/O to the paging file could mean a memory bottleneck.

To take a snapshot of memory and paging files:

1. Right-click the host, and then select **Performance, Memory Snapshot**. The Memory Snapshot dialog box appears.
2. Use the counters on the **Memory** and **Paging File** tabs to identify a memory bottleneck, if any.

Monitoring Physical Disks for Bottlenecks

Use the following alerts to monitor bottlenecks on Windows disks.

- ◆ **Physical Disk Queue Length alert** — A key indicator of a disk bottleneck is a high disk queue, which indicates that a disk is not handling I/O requests sufficiently. The Physical Disk Queue Length alert watches this key indicator.
- ◆ **Physical Disk Average Transfer Rate alert** — This alert measures how long it takes to read from or write to a physical disk. To determine the transfer rate, the agent watches the physical disk object's Avg. Disk sec/Transfer counter. If you create a baseline for this counter, you can set this alert to trigger when disk performance is poor compared to the baseline.

Working with MVS Host Resources

MVS host resources are visible in the tree view, where you can select them for use in Properties and Relationship views. Also, for a wide range of additional data and reports specific to MVS systems, use the agent windows for logical, physical, SMS, and HSM storage. ControlCenter agents allow direct access to host configuration data, status information, operational commands, and a number of reports.

Viewing MVS Host Properties and Relationships

You can view MVS properties and relationships using the methods shown at the beginning of this chapter.

Figure 12-5 shows a relationship view for MVS. In this view, data columns not applicable to MVS are not shown.




Figure 12-5 Viewing Relationships for an MVS Host (Table and Map)

Viewing Detailed Host Information

For detailed and up-to-date information about MVS host storage resources, issue commands directly to agents through the agent windows. Most MVS results appear in their own dialog box. Agent data is collected when you issue an **Explore** command.

You can explore and control MVS host resources using the following agent windows in the Console:

- ◆ Logical Agent for MVS window
- ◆ Physical Agent for MVS window
- ◆ Host Agent for MVS SMS window
- ◆ Host Agent for MVS HSM window

To explore one of the agents: In the tree, right-click the host, and then select the agent name and click **Explore**.

For more detailed information about host file systems and physical devices, use Table 12-4.

Table 12-4 Exploring Detailed Information about MVS Host Resources

To explore	In the	Right-click the host, then the agent, then select
MVS high-level qualifiers	Logical Agent for MVS	Logical, Explore by DSN
MVS Open Edition file systems	Logical Agent for MVS	Logical, Explore Open Edition
MVS physical devices	Physical Agent for MVS	Explore

After you perform these commands, right-click any device to explore it further. Data appears to the right as shown in Figure 12-6.




Figure 12-6 Viewing Host Information in an Agent Window

Common Tasks in MVS Hosts

Use the table for quick reference to perform tasks for MVS hosts.

Table 12-5 Common Tasks on MVS Hosts

Task	Action	Additional steps
Explore MVS hosts and their volumes, SMS groups, and network interfaces	In the tree, expand the host and the subfolders beneath it.	Select the file systems, devices, and other resources you desire, and add them to a Properties view using any method.
List data sets using wildcards	Right-click the host and select Logical Agent for MVS, Logical, Explore by DSN .	Specify partially or fully-qualified dataset names you want to locate. Example: DSN.*.*
View SMS groups	Expand the host, and then expand SMS Groups .	To perform operations, explore the Host Agent for MVS SMS and view its separate agent window.
View volumes	Expand the host, and then expand Volumes - All or Volumes - non-SMS .	To perform operations, explore the Host Agent for MVS SMS and view its separate agent window.
Vary a volume online	Right-click the host and select Physical Agent for MVS, Physical, Units, Disk, Offline .	In the agent window, expand Units Disk Offline . Right-click the desired disk and select Vary Online .
Vary a volume offline	Right-click the host and select Physical Agent for MVS, Physical, Units, Disk, Online - by Volume .	In the agent window, expand Units Disk Online by Volume . Right-click the desired disk and select Vary Offline .
Initialize one or more volumes	Right-click the host and select Physical Agent for MVS, Physical, Units, Disk, Offline .	In the agent window, expand Offline Volumes , right-click an offline DASD volume and select Initialize Volume . Complete the dialog box.
Define your application data sets for use with DASD space reports	Right-click the host and select Logical Agent for MVS, Setup, Application IDs .	In the Create/Edit Application IDs dialog box, click Create . Then specify data sets, account codes, or other defining characteristics of application storage resources.
Identify users and jobs using excessive DASD space	Right-click the host and select Logical Agent for MVS, Logical, Reports, DASD Space Utilization, Application Space Overview .	Pick one of the application IDs you previously defined to run a report showing the space utilization of jobs and users for that application. Click Generate Report .
Identify users and jobs consuming DASD space at an excessive rate	Right-click the host and select Logical Agent for MVS, Logical, Reports, DASD Space Activity .	Specify the jobs, users, dataset names, and account codes whose storage consumption you want to view.
Perform most storage-related commands on MVS hosts	Right-click the host you desire, then select the agent name and the command you desire.	

Recapturing MVS Host Storage

One of the most beneficial tasks you can perform for your hosts is to recapture storage. This section briefly describes numerous methods for doing this on MVS hosts.

The following tasks (and the reports they describe) use the Logical Agent for MVS to recapture inefficient storage.

You can identify who is using the space and how quickly they are using it. To start, you define the data sets that an application uses, and then run reports to identify the jobs and users using large amounts of space or consuming space at a quick rate.

Defining Applications for Space Reporting

MVS space reporting uses “application IDs” to define the data sets and other resources that characterize an application. Application IDs allow you to generate the same report repeatedly without repeating the data entry for the parameters of the space reports.

To create an application ID to save for running space usage reports:

1. Expand **Hosts** in the selection tree. Right-click the preferred MVS host. Select **Logical Agent for MVS, Setup, Application IDs**. The Create/Edit Application IDs dialog box appears.
2. Click **Create**. The Application ID Resources dialog box appears.
3. Type the criteria that define the application data. For field descriptions, click **Help**.

Be sure to specify a new name for the Application ID.

4. Click **OK**.
5. Return to the Create/Edit Application IDs dialog box.
6. Click **Refresh**. Verify that the newly created application ID appears in the list.

See the online Help for copying and deleting application IDs.

Identifying Jobs and Users With Too Much Space

You can often trace space problems to a few jobs and users with excessive space. Use the DASD Space Utilization report to find out the names of these jobs and users.

To run an application space usage report:

1. Expand Hosts in the selection tree. Right-click a preferred MVS server, and select **Logical Agent for MVS, Logical, Reports, DASD Space Utilization, Application Space Overview**.

The DASD Space Utilization - Select Application IDs dialog box appears.

2. Select an application ID, and click **Generate Report**.

The output for the report appears, displaying with space usage statistics about that application.

Identifying Jobs and Users With High Rates of Consumption

When a volume is filling quickly, you can find out the jobs or users that have consumed the space and the amount of space they have consumed in a given time. Use the DASD Space Activity report to find out the names of these jobs and users.

To run an application space consumption report:

1. Right-click the host and select **Logical Agent for MVS, Logical, Reports, DASD Space Activity**.
2. Specify the data sets, volumes, users, jobs, and other criteria for which you want a report.
3. Specify the time period whose space consumption you want to analyze. Complete any other desired fields and click **OK**.
4. In the report output, click the column headings to sort jobs or users to the top that have consumed the most space.

Managing MVS Storage to Increase Performance

Performance problems frequently originate in storage configuration and utilization. You can use the main Console window and the host agents to solve many storage-related performance issues.

Gathering MVS Host Performance Data

You can analyze MVS host performance with Performance Manager. Refer to Chapter 9, *Monitoring and Analyzing Performance*.

Mapping MVS Volumes to Disk Array Storage

When you analyze performance of MVS volumes, you can troubleshoot problems by associating the volumes with the storage on which they reside.

To associate MVS volumes with their disk array counterparts:

1. Expand **Hosts**, then the MVS host, then **Volumes - All**.
2. In the tree, check the poorly performing MVS volumes.
3. Drag the volumes into a Relationship view.

Figure 12-7 shows how to map MVS volumes to their counterpart volumes on the Symmetrix array.




Figure 12-7 Mapping MVS Volumes to Symmetrix Logical Volumes

Identifying Orbiting Data Sets in HSM

Configuration problems in HSM can cause some data sets to be migrated, then recalled, and then migrated again, in a cycle that decreases application performance, impairs timely access to data, and interferes with other I/O processing. To correct the problem, identify so-called *orbiting* data sets and assign them a management class.

To identify orbiting data sets:

1. Right-click the host, and then select **Host Agent for MVS HSM, Console, Recall, DSN Recall Report**.
2. In the Get Report Date dialog box, specify the date range for the report and the number of recalls that you consider excessive for that interval.

The HSM Recall Orbiter report appears in a dialog box. Any data set recalled the specified number of times during the date range appears in the report.

3. Right-click the data set for additional commands, such as viewing the recall history and altering the management class.

This chapter provides an introduction to StorageScope and StorageScope reports. It contains the following:

- ◆ Overview of Reports..... 13-2
- ◆ StorageScope Permissions 13-4
- ◆ Working with Reports in the Console..... 13-6
- ◆ Logging In to StorageScope..... 13-8

Overview of Reports

StorageScope, an integral part of the EMC ControlCenter family of storage management software, enables capacity planning, chargeback, and efficient asset management. With StorageScope's historical Storage Resource Management (SRM) reports, you can:

- ◆ Identify under utilized or inefficient utilization of storage assets.
- ◆ Facilitate billing and chargeback operations by location, line of business, or application.
- ◆ Plan capacity across your entire infrastructure.
- ◆ Summarize your multi vendor storage capacity and usage in application and business contexts.
- ◆ Track historical metrics to predict growth.

StorageScope also provides a variety of management and customization tools that allow you to create custom report layouts and custom graphs, export report data into CSV or XML format, print reports in HTML format or PDF, and schedule reports to run automatically on a daily or weekly basis.

Refer to the StorageScope online Help and the *EMC ControlCenter StorageScope Reference Guide* for detailed descriptions of the StorageScope user interface, reports, and procedures for using and administering StorageScope.

Types of Reports

StorageScope provides a variety of reports that show both the capacity and utilization of managed objects in ControlCenter. These managed objects include storage systems, hosts, NAS, devices, and fabrics. The group reports present an overview or summary of data into which you can drill down to get data about individual objects. In addition, the File Summaries reports show file age and size distribution data across the entire enterprise by file set, host, and file system. Figure 13-1 on page 13-3 shows the various types of reports available through StorageScope.

Where do I find reports I need?
Reports by Object
Arrays
Connectivity
Hosts
Backup
NAS
File Summaries
CAS Clusters
Combined
Reports by Use
Storage Utilization
Capacity Planning
Asset Management
Chargeback
Troubleshooting/Configuration
Backup/Restore
All Reports
All Reports
All Custom Layouts

Figure 13-1 Types of StorageScope Reports

User-Defined Groups for Reports

A user-defined group contains groups of objects (for example, all NAS file servers at a physical location or all hosts being utilized by a specific department). You can use the ControlCenter grouping functionality to group objects to produce reports that meet specific business needs.

For example, you could group all Symmetrix arrays whose storage is being utilized by your accounting department, in order to plan for additional capacity or to track utilization.

Refer to *Creating User-Defined Groups* on page 5-16 for specific information about creating user-defined groups.

StorageScope Permissions

To access StorageScope, you must have StorageScope permissions. The StorageScope functions available to you depend on the type of StorageScope permissions assigned to you.

StorageScope Permission Types

StorageScope provides two permission types: user and administrator.

- ◆ **User** — StorageScope users can view, print, and export reports, and create custom report layouts and graphs.
- ◆ **Administrator** — In addition to performing all the tasks that a StorageScope user can perform, StorageScope administrators can schedule reports, run reports in real time, modify report retention policies, and view report history.

Applying StorageScope Permissions

When you install ControlCenter for the first time, an Any User authorization rule is created. *StorageScope user* (read-only) permissions are placed in this rule. Other permissions that the ControlCenter administrator wants all users to have can also be placed in this rule.

By default, the Any User Rule is automatically applied to a user when that user is added to ControlCenter. Therefore, all ControlCenter users have read-only access to StorageScope.

The Any User Rule does not count as an assigned rule, so individual users can have the Any User Rule and one other rule assigned to them.

If you do not want all users to have this access, do the following:

1. Delete the StorageScope user permissions from the Any User Rule.
2. Create a group and add those users that you want to have StorageScope user permissions.
3. Select the group and create a rule with StorageScope user permissions.

When you create the rule, select the **Types** radio button and select **StorageScope Reports** from the list of available object types. Then, select *StorageScope user* from the list of available actions.

StorageScope administrator permissions must be explicitly assigned to a user or group of users. You can create a StorageScope authorization rule and apply it to specific users or a group of users in the same way you would any other type of ControlCenter permissions.

Refer to Chapter 1, *Managing ControlCenter Users*, for detailed information about and procedures for creating users, authorization rules, and user groups.

Working with Reports in the Console

The following sections describe how to launch StorageScope and how to display specific reports from ControlCenter.

Launching StorageScope from ControlCenter

To launch StorageScope from the ControlCenter Console:

1. Click **ECC Administration** in the task bar.
2. From the **Reports** menu, select **Launch StorageScope**.
StorageScope launches in a browser window (Figure 13-2).




Figure 13-2 StorageScope Home Page

Opening Reports from ControlCenter

You cannot open any of the StorageScope File Summary reports directly from the ControlCenter Console. You must Launch StorageScope first.

To open a report from the ControlCenter Console:

1. Right-click a managed object anywhere in the Console.

You can right-click an object in the tree view or in any other view except Performance view.

2. Select **StorageScope Reports**. The Select Report dialog box appears.
3. Select the report to open from the list of reports.
4. Click **OK**. The report opens in a Web browser window. From this window, you can access other reports and perform all the StorageScope tasks for which you are authorized.

Figure 13-3 provides an example of the All Arrays report.

Arrays - Accounting		Pages 1 of 1		Rows 1-21 of 21		About this page ?
Array ^	Array Type	Raw - Total (GB)	Configured - Usable (GB)	Unconfigured - Total (GB)	Allocated - Usable - Total (GB)	Unallocated - Usable - Total (GB)
000183400005	Symmetrix	5,569.48	1,193.70	3,355.47	399.34	794.36
000183500601	Symmetrix	3,189.62	1,079.94	1,441.32	453.21	626.73
000184600023	Symmetrix	2,187.72	1,994.30	100.11	1,267.93	726.37
000184799999	Symmetrix	17,228.33	2,475.49	12,633.03	1,042.38	1,433.11
000185400279	Symmetrix	1,628.10	925.87	228.42	293.70	632.17
000185400295	Symmetrix	1,628.10	888.83	331.82	67.74	821.10
000185400331	Symmetrix	1,628.10	1,040.38	96.63	290.99	749.38
000187400019	Symmetrix	19,689.52	7,550.60	8,204.62	3,748.17	3,790.09
000187700037	Symmetrix	38,832.12	12,013.58	24,289.34	6,326.40	5,631.37
000187900038	Symmetrix	8,203.97	2,147.69	4,198.35	36.70	2,088.88
000187900110	Symmetrix	8,067.24	2,012.99	4,679.00	127.88	1,862.71
000187940103	Symmetrix	7,930.50	1,871.99	4,728.44	111.58	1,739.34
000187940284	Symmetrix	5,873.73	646.87	4,868.87	114.52	532.35
000187940311	Symmetrix	6,152.98	2,095.24	2,607.42	358.93	1,722.94
000187940318	Symmetrix	6,152.98	1,006.64	4,393.94	300.84	665.79
30471	HPXP	0.00	4,633.06	0.00	2,957.57	1,670.12
APM00024401348	Clariion	999.58	514.21	0.00	59.64	192.39
APM00024600161	Clariion	599.75	441.39	0.00	166.78	181.50
APM00024600162	Clariion	666.39	447.60	0.00	433.18	12.01
APM00034701854	Clariion	6,015.59	3,589.84	133.68	864.29	750.46
WRE00022401063	Clariion	0.00	0.00	0.00	0.00	0.00
Totals (21):		142,243.80	48,570.19	76,290.46	19,421.76	26,623.18

Figure 13-3 All Arrays Report Example

Logging In to StorageScope

You can also log in to StorageScope through your browser. You log in using your ControlCenter username and password.

Refer to the *EMC ControlCenter Support Matrix* on the Powerlink™ website at <http://powerlink.EMC.com> for supported browsers and versions.

To log in to StorageScope through a browser:

1. Open a browser window.
2. Enter the URL of the StorageScope server in the following format:
http://<storagescope hostname>:<storagescope port number>
 - *storagescope hostname* is the name of the host on which the StorageScope server is installed.
 - *storagescope port number* is the number of the port on the StorageScope host through which you communicate with StorageScope. The default port number is **8080**.
3. When the StorageScope login page (Figure 13-4) appears, enter your ControlCenter username and password, and then click **Login**. The StorageScope home page appears.




Figure 13-4 StorageScope Login Page

Incorporating Data from External Sources into Reports

This chapter provides an introduction to configuring ControlCenter custom fields. Custom fields allow an administrator to specify attributes for the managed objects in their storage environment that can later be used in generated reports. This chapter contains the following:

- ◆ Custom Field Overview 14-2
- ◆ Configuring Custom Fields 14-4
- ◆ Managing Custom Fields 14-11
- ◆ Using the Custom Field Bulk Load Utility 14-14

Custom Field Overview

ControlCenter allows you to associate managed object types with custom field names and values that are relevant to your business needs. Once objects are “tagged” with this information, you can then generate reports for such things as asset management or chargeback.

For example, users could tag objects with asset management information such as location, vendor, lease date, and cost. Or, they might choose to collect chargeback information from objects by tagging them with project, department, and client names. Users can then generate reports about the objects using ControlCenter StorageScope or some other reporting tool.

A brief overview of StorageScope features and operation is provided in Chapter 13, *Using Reports*. Refer to the StorageScope online Help and the *EMC ControlCenter StorageScope Reference Guide* for detailed descriptions of the StorageScope user interface, reports, and procedures for using and administering StorageScope.

Planning Overview

The custom field feature of ControlCenter requires planning and configuration before it can be used effectively. Planning your asset management or chargeback strategy is beyond the scope of this guide, but you should have your business model in place before creating the template that associates your custom field names with ControlCenter predefined managed object types.

Custom fields are valid for the following ControlCenter objects:

- ◆ Array
- ◆ Array Device
- ◆ Disk
- ◆ Array Port
- ◆ Database Instance
- ◆ File System, NAS File System
- ◆ Host (including Data Movers and NAS Servers)
- ◆ Host Device
- ◆ Host HBA

- ◆ Logical Volume
- ◆ Volume Group
- ◆ NAS Container
- ◆ Switch
- ◆ Fabric
- ◆ Switch Port

Configuration Overview

The following steps are required to set up and configure the custom field feature on ControlCenter.

1. Assign appropriate permissions to users.

Specific user permissions are required to administer and use custom fields. There are two permissions that must be set based on the type of access required by an administrator/user. Refer to *Assigning User Permissions* on page 14-4.

2. Create the custom fields template as an XML file.

Once planning is complete, the administrator creates custom name and value pairs using an XML template defined by ControlCenter. This XML file provides the relationship between ControlCenter object types and the custom terms developed by the business. Refer to *Creating the Custom Fields XML Template* on page 14-4.

3. Import the custom field XML file into ControlCenter.

The custom fields XML file is imported into the ECC Repository using the Import Custom Fields menu option. Refer to *Importing the XML File* on page 14-8.

4. Modify the custom fields XML template if necessary, and then reimport it.

If changes need to be made to the custom field definitions, the administrator modifies the template, and then reimports it. Refer to *Creating the Custom Fields XML Template* on page 14-4.

Configuring Custom Fields

Complete the following tasks to configure ControlCenter custom fields:

- ◆ Assign user permissions
- ◆ Create custom field XML template
- ◆ Import the custom field XML template

The following sections describe each task.

Assigning User Permissions

ControlCenter provides two user permissions for custom fields:

- ◆ Custom Field Administrator — Allows a user to import the custom field XML template, and assign, disassociate (remove the association between an object and custom name/value pairs), and modify custom fields. This permission also allows a user to remove stale custom fields.
- ◆ Custom Field User — Allows a user to assign, disassociate, and modify custom fields.

Refer to *Assigning Permissions* on page 1-6 for additional information on assigning user permissions. If an administrator/user does not have the appropriate permissions, the custom fields menu items are unavailable.

Creating the Custom Fields XML Template

The custom field XML template provides the relationship between ControlCenter managed object types and your custom field values. The custom fields XML template may be created anywhere as long as it is located on a host accessible by the ControlCenter Console.

Information pertinent to the custom fields is specified in two distinct sections of the custom fields template:

- ◆ Definitions section — Specifies the custom field names and the allowable custom field values.
- ◆ Relationship section — Defines the relationship between ControlCenter managed object types and custom field names.

Rules for Creating the Custom Field XML

The contents of the XML file are described as follows:

```
<UDF_settings> ::= <UDF_name> {<UDF_name>} <TYPE_TO_UDF>
                  {<TYPE_TO_UDF>}
<UDF_name> ::= {<UDF_value>}
<UDF_value> ::= restricted ASCII set of characters
<TYPE_TO_UDF> ::= <UDF>{<UDF>}
<UDF> ::= Custom field names specified by the "name" attribute of the
          <UDF_name> xml tag
```

The following general rules apply to creation of the custom field XML file:

- ◆ Specify only one <UDF_settings> XML tag in the file.
- ◆ The <UDF_settings> element should contain at least one <UDF_name> element and at least one <TYPE_TO_UDF> element.
 - The Definitions section consists of a collection of <UDF_name> elements.
 - The Relationship section consists of a collection of <TYPE_TO_UDF> elements.
- ◆ A <UDF_name> element can contain zero or more <UDF_value> elements. The <UDF_name> contains the value for the custom fields name. Values specified for this attribute should be unique.

If the <UDF_name> has no <UDF_value> specified, it is considered a *free-form* element. At runtime, the user can provide the value for this type of element through the ControlCenter Console.

- ◆ There should be at least one <TYPE_TO_UDF> element defined. One attribute name should be associated for each <TYPE_TO_UDF> element. These values are restricted to the types shown in Table 14-1.

Table 14-1 Relationship of ControlCenter Objects to Custom Field Elements

Object Type	Type Name Specified in XML File
Array	Array
Array Device	ArrayDevice
Disk	ArrayDisk
Array Port	ArrayPort

Table 14-1 Relationship of ControlCenter Objects to Custom Field Elements

Object Type	Type Name Specified in XML File
Database Instance	DatabaselInstance
File System, NAS File System	FileSystem
Host, NAS Server, Data Mover	Host
Host Device	HostDevice
Host HBA	HostHBA
Logical Volume	LogicalVolume
NAS Container	NASContainer
Switch	Switch
Fabric	SwitchFabric
Switch Port	SwitchPort
Volume Group	VolumeGroup

5. The <TYPE_TO_UDF> element should contain at least one <UDF> element.
6. The <UDF> element should contain only the values that pertain to the name attribute of the <UDF_name> as defined in the Definitions section of the XML. These values correspond to a set of custom field names for a given managed object type.

StorageScope restricts the number of custom fields for a given type to 10. Although the user is allowed to specify more than 10 names in the XML file, only 10 are processed.

Sample Custom Field XML

This section contains a simple custom field XML example to assist in developing your own based on your business needs.

The following example has three custom fields templates. The custom field names are Location, Account Number, and Cost. Location and Account Number have a corresponding set of allowable custom field values, but Cost is a free form template that allows the user to assign the value at run time through the Console Custom Field Set/Modify dialog box. Refer to *Setting or Modifying Custom Field Names and Values* on page 14-12 for information on assigning custom fields values.

The Relationship section contains two types of Managed Objects: Host and VolumeGroup.

In this example, the administrator associates the Account Number and Cost templates with the Host object type, and Location and Account Number with VolumeGroup. Cost has no values specified, thus is a *free form* template.

```
<?xml version="1.0" encoding="ASCII"?>

<UDF_settings>
  <UDF_name name="Location">
    <UDF_value>Hopkinton</UDF_value>
    <UDF_value>Tokyo</UDF_value>
  </UDF_name>

  <UDF_name name="Account Number">
    <UDF_value>111-22</UDF_value>
    <UDF_value>222-33</UDF_value>
  </UDF_name>

  <UDF_name name="Cost">
  </UDF_name>

  <TYPE_TO_UDF name="Host">
    <UDF>Account Number</UDF>
    <UDF>Cost</UDF>
  </TYPE_TO_UDF>

  <TYPE_TO_UDF name="VolumeGroup">
    <UDF>Account Number</UDF>
    <UDF>Location</UDF>
  </TYPE_TO_UDF>
</UDF_settings>
```

In summary: the administrator defined three custom fields named Location, Account Number, and Cost. The templates were applied to two types of managed objects: Host and VolumeGroup. Cost was not associated with any values and is a free form template that can be assigned values through the Custom Field Set/Modify dialog box.

Once a file is created, you import it into ControlCenter as outlined in the next section.

Importing the XML File

Import the appropriate custom field XML file as follows:

You must have Custom Field Administrator permission to complete this task (refer to *Assigning User Permissions* on page 14-4).

1. From the ControlCenter Console, click the **Administration** task button, and then click **Reports**.
2. Select **Import Custom Fields**. The Import Custom Fields dialog box appears.

If the Import Custom Fields dialog box does not appear, you do not have Custom Fields Administrator permission.

3. Browse or enter the pathname where the XML file resides, and then click **Apply**.

As the import proceeds, the Progress Area displays each step (Parsing, Importing, Import Complete). After successfully validating the template entries in the new XML file, the Validation area displays import statistics.

For example, the Validation area might list that a previous field named *Accounts* existed in the Repository, and that the field does not exist in the latest import of the XML file. It will also list any new field names that are imported with the XML file, but did not previously exist in the Repository.

If the import is unsuccessful, an error message appears.

4. Click **Finish** to close the dialog box.

Once you import the XML file, you can start using the ControlCenter custom field feature to tag specific objects.

If this was a reimport in which you modified your original XML template, you may choose to remove the *stale* custom fields (those fields that you removed from your current XML file but that ControlCenter retains in the Repository). Refer to *Removing Stale Custom Field Names* on page 14-9 for details.

Proceed to *Managing Custom Fields* on page 14-11 after successfully importing the XML file.

Removing Stale Custom Field Names

When importing a modified XML file, ControlCenter checks for differences between the custom names/value pairs and managed object type definitions that currently exist in the Repository, to those being imported from a modified XML file. If a name existed in the Repository, but does not exist in the current XML template, the Repository entry is marked as *stale*.

To delete stale custom field names for all instances of a managed object type from the Repository:

You must have Custom Field Administrator permission to complete this task (refer to *Assigning User Permissions* on page 14-4).

1. From the ControlCenter Console, click the **Administration** task button, and then click **Reports**.
2. Select **Remove Stale Custom Fields**. The Remove Stale Custom Fields dialog box appears.

If the Remove Stale Custom Fields option does not appear, you do not have Custom Fields Administrator permission.

3. Select the name/value pairs to delete by doing any of the following:
 - Highlight contiguous or noncontiguous rows, and then click **Select Highlighted**.
 - Select one or more checkboxes.
 - Select all field name and value pairs by clicking **Select All**.

4. To remove the selected name/value pairs from the Repository, click **Apply**. The dialog box refreshes and remains open, allowing you to continue making additional changes. Or, click **OK** to remove the name/value pairs and close the dialog box.
5. If you choose to cancel all selections, but want to keep the dialog box open, click **Deselect Highlighted** or **Deselect All**. If you want to cancel all selections and exit the dialog box, click **Cancel**.

The only values that appear in the Remove Stale Custom Fields dialog are those custom field values that were assigned to a managed object but then were consequently made stale.

For example, a user has a custom field named *City* with all the cities in Massachusetts as values, but then imports a new XML template file that does not use the *City* custom field name. If the value *Boston* was never assigned to any managed object and is made stale, it will not appear in the Remove Stale Custom Fields dialog. However, the Import Custom Fields dialog will report that the value *Boston* was made stale (along with all of the other values associated with the *City* custom fields name).

Managing Custom Fields

The template you created in *Creating the Custom Fields XML Template* on page 14-4 establishes the relationship between specific ControlCenter object types and custom field names and values. Once the template is imported, you can use ControlCenter to associate these custom field names and values to specific ControlCenter objects (as long as that object type was defined in your template). This section provides procedures for the following custom field operations:

- ◆ *Viewing Custom Field Properties* on this page
- ◆ *Setting or Modifying Custom Field Names and Values* on page 14-12
- ◆ *Removing Name and Value Pairs* on page 14-13

Viewing Custom Field Properties

You can view the custom field values associated with managed objects in ControlCenter as follows:

1. From the ControlCenter Console, click the **ECC Administration** task button, and then select **Custom Fields** from the ECC Administration menu.

The Custom Fields view opens.

2. Drag the object(s) from the Console tree into the Custom Fields view.

A table appears listing the names of the objects, the custom field names, and the values associated with the custom fields. If the object has no custom field values associated with it, the table displays *Not Set*.

Setting or Modifying Custom Field Names and Values

Use the following procedure to set or modify the custom field values that are associated with specific managed objects:

1. Ensure that you are logged in as a user with Custom Field User permissions.
2. In the Console tree, select and then right-click the object or objects (of the same type) that you want to associate to a custom field name and value (multiple custom field name/value pairs may be set at once).
3. Select **Set/Modify Custom Fields**. All custom defined names that are valid for the selected objects appear in the Set/Modify dialog box.
 - You may choose multiple objects from the tree and associate those selected objects to the same custom field name and value pairs. When choosing multiple objects from the Console tree, the objects must be of the same object type (i.e., multiple hosts or multiple arrays).
 - Selected items with existing assigned values appear in the Selected Items dialog box.
 - The Custom Fields with Static Values dialog box displays all custom names and value pairs that are currently available.
 - The Custom Fields with Free Form values box lists names to which free form values may be assigned.
4. Select the appropriate name/value pairs from the Custom Fields with Static Values field, or if the values section had been left blank in the template, enter a new value in the Custom Fields with Free Form Values field.
5. Click **Apply** to continue setting associations between objects and name/value pairs appearing in the Selected Items box, or if you are done, click **OK** to close the dialog box.

Removing Name and Value Pairs

To remove an association between an object(s) and that object(s) custom name/ value pair(s):

Removing an association does not delete name/value pairs from the Repository. Only the association, or link, to objects are removed from the Repository.

1. Right-click an object(s) in the Console tree panel and select **Remove Custom Fields**. All custom field names and value pairs that are associated to the selected object(s), appear in the Remove Custom Fields dialog box.

You can remove object associations to name/value pairs by doing any of the following:

- Select contiguous or noncontiguous rows, and then click **Select Highlighted**.
 - Select one or more checkboxes.
 - Select all field name and value pairs by clicking **Select All**.
2. To remove selected associations from the Repository, click **Apply**. The dialog box refreshes and remains open, allowing you to continue making additional changes. Or, click **OK** to remove the associations and close the dialog box.

If you choose to cancel all your selections but want to keep the dialog box open, click **Deselect Highlighted** or **Deselect All**. If you want to cancel all selections and exit the dialog box, click **Cancel**.

Using the Custom Field Bulk Load Utility

ControlCenter provides a utility to aid in importing custom field values and/or definitions into the ControlCenter Repository. This utility allows you to:

- ◆ Migrate existing custom field values and definitions into ControlCenter.
- ◆ Import custom field values and definitions for different object types all at once.
- ◆ Use CLI commands instead of GUI commands in instances where large volumes of data need to be imported.
- ◆ Export custom field managed object types.

The import of custom fields *values* is done from a comma-delimited file. The import of custom fields *definitions* is done from an XML file (refer to *Creating the Custom Fields XML Template* on page 14-4). The export of *object identifiers* (types) from ControlCenter is done to a comma-delimited file.

Importing Custom Fields Using the Bulk Load Utility

This section provides an outline of the steps required to import custom field values and definitions using the bulk load utility. The following tasks are required:

- ◆ Access the Custom Field Utility
- ◆ Create the custom field XML template
- ◆ Import the template into the database
- ◆ Export managed objects
- ◆ Associate exported objects with custom field names
- ◆ Import custom field-to-managed object assignments into ControlCenter

Accessing the Custom Fields Utility

The custom fields utility is located on the ECC Server host. Access the utility as follows:

1. Open a DOS prompt on the ECC Server host.
2. Change directories to:

```
C:\><ECC_INSTALL_ROOT>\Tools\utils\CustomFields\
```

Custom Fields Help

Use the following command to access custom fields help:

```
CustomFieldsUtil -h
```

Creating the Custom Field Definition XML Template

A custom field definition XML template is required to provide the relationship between ControlCenter managed object types and the custom field values you will load using the bulk load utility.

The custom field definition XML template is created using the same procedure as for the XML template you wrote in *Creating the Custom Fields XML Template* on page 14-4.

Note that once the template has been imported via either the console or through the Custom Fields Utility, it does not need to be imported again, unless the XML template has been modified.

This template should contain the definitions for the custom fields you are importing through the bulk load utility as well as any definitions previously created and stored in the Repository. If a name existed in the Repository, but does not exist in the current XML template, the Repository entry is marked as *stale* (refer to *Removing Stale Custom Field Names* on page 14-9).

Importing Definitions

This command imports custom field definitions specified in the definitions file (XML format). If `-r` is specified, all stale entries as a result of a new definition are removed from the Repository. Otherwise, all stale instances remain in the Repository and are marked as "stale." Refer to *Removing Stale Custom Field Names* on page 14-9 for more information.

Use the following command to import custom field definitions from the definitions XML file:

```
CustomFieldsUtil -defImport <definitions_file_path> -r
```

The results of the definitions import are provided in an output file that is created upon completion of the import at `<definitions_file_path>`. The file is named:

```
<name_of_config_file>_xml_def_import_<date>_<time>.out
```

This file can be opened with a text editor and contains the same information as in the Validation Results area of the Import Custom Fields dialog (Refer to *Importing the XML File* on page 14-8).

Exporting Object Identifiers

This command exports ControlCenter managed object identifiers (such as Host, Array, ArrayDevice, etc.) to the specified file (requires the full path). All supported types are exported unless explicit further qualification of the desired types is specified. A complete list of supported types is provided in Table 14-1 on page 14-5.

Use the following command to export object identifiers to a specified location:

```
CustomFieldsUtil -export <managed_Objects_File_Path>
[-a] "type1 type2..."
```

If the **-a** option is specified, the identifiers are appended to an existing file.

The information is exported into a comma-delimited file in the following format:

```
moType,moId,moCanonical
```

where:

moType — The custom fields supported managed object type. A list of object types is provided in Table 14-1 on page 14-5.

moId — The unique internal database identifier for the managed object.

moCanonical — The managed object canonical name.

In the following example, the user wants to export all of the object identifiers for the objects using the Host and Volume Group object types into a comma-delimited file named `sample_export.csv`:

```
customFieldsUtil -export sample_export.csv "Host VolumeGroup"
```

The exported file appears as follows:

	A	B	C
1	Host	8503	Host=lumper062.bdcdevdr.bdcdev.bdc.com
2	Host	21105	Host=USENSLAMOD1E.eng.bdc.com
3	Host	58845	Host=mgandhi200013d.eng.bdc.com
4	VolumeGroup	11019	Host=lumper062.bdcdevdr.bdcdev.bdc.com>NtVolumeGroup=rootvg
5	VolumeGroup	19281	Host=USENSLAMOD1E.eng.bdc.com>NtVolumeGroup=rootvg

Figure 14-1 Exported Comma-Delimited File

The export file contains the object identifiers for three hosts (lumper062, USENLAMOD1E, and mgandhi200013d), and two volume groups (rootvg on host lumper062, and rootvg on host USENLAMOD1E).

Associating Managed Objects with Custom Field Names

Once managed object identifiers are exported from ControlCenter, you associate them with your existing custom fields values and definitions. You can then import the associated managed objects with their custom field values and definitions back into ControlCenter to complete the migration.

The process you use to associate the ControlCenter object identifiers with your custom fields must create a comma-delimited file for the custom field values as shown in *Creating the Comma-Delimited Import File* which follows.

Creating the Comma-Delimited Import File

The bulk load utility uses a comma-delimited file to import the managed object and associated values back into ControlCenter.

The custom fields import file obeys the following comma-delimited file format:

moType,moId,moCanonical,customFieldName,customFieldValue

where:

moType — The Custom Fields supported managed object type. A list of object types is provided in Table 14-1 on page 14-5.

moId — The unique internal database identifier for the managed object.

moCanonical — The managed object canonical name. This field is optional.

customFieldName — The custom field name.

customFieldValue — The value of the custom field to insert/update for the specified custom field name. If not specified, designates that this custom field should be deleted for the specified managed object.

Figure 14-2 shows an association between the object identifiers exported in the previous section and a user's custom field names (Account Number and Location) and values.

	A	B	C	D	E
1	Host	8503	Host=lumper062.bdcdevdr.bdcdev.bdc.cor	Account Number	222-33
2	Host	21105	Host=USENSLAMOD1E.eng.bdc.com	Account Number	333-44
3	Host	58845	Host=mgandhi200013d.eng.bdc.com	Account Number	444-55
4	VolumeGroup	11019	Host=lumper062.bdcdevdr.bdcdev.bdc.cor	Account Number	222-33
5	VolumeGroup	19281	Host=USENSLAMOD1E.eng.bdc.com>Nt	Location	Atlanta

Figure 14-2 Comma-Delimited File Ready for Import

In the following example the administrator is adding the custom field *Location* with the value *Hopkinton* to a Symmetrix device, and in the second line is removing that custom field from a host by simply not specifying a value after the custom field name.

```
ArrayDevice,23011,Symmetrix=00187900611>Device=5F5,
    Location,Hopkinton
Host,53215,Host=ZZUR0D20ENG.LSS.EMC.COM,Location
```

Once a file with the changes you want to import is created, you need to import it into ControlCenter using the procedure outlined in *Importing Values into ControlCenter* which follows.

Importing Values into ControlCenter

This command imports the custom field values specified in the comma-delimited format values file. A log file of the import progress and failures is produced in the file:

```
<valuesFilePath>\<FileName>_import_<date>_<time>.out
```

Use the following command to import the values from the file you built in *Creating the Comma-Delimited Import File* on page 14-17:

```
CustomFieldsUtil -import <values_File_Path>
```

customFieldsUtil Examples

Use the following examples to gain a better understanding of the custom fields utility.

Importing Values

Import the values specified in `c:\temp\mycustomfields.csv` and produce a log file of progress and failures in

`C:\temp\myCustomFields.csv_import_date_time.out:`

```
customFieldsUtil -import -values c:\temp\myCustomFields.csv
```

Importing Definitions

Import the definitions in the specified XML file without removing stale instances:

```
customFieldsUtil -defImport c:\temp\customFieldsDef.xml
```

Exporting Object Identifiers

Export all ControlCenter managed object identifiers that support custom fields to the specified path:

```
customFieldsUtil -export c:\temp\allObjectIds.csv
```

Export all managed objects of type Host and ArrayDisk to the specified path:

```
customFieldsUtil -export c:\temp\allObjectIds.csv "Host ArrayDisk"
```


Tuning Symmetrix Performance

This chapter tells you how to tune your Symmetrix systems for optimum performance and contains the following sections:

- ◆ Performance Management Overview 15-2
- ◆ Understanding Optimizer 15-3
- ◆ Using Optimizer..... 15-8
- ◆ Retrieving Optimizer Logs 15-11
- ◆ Using the Quality of Service (TimeFinder/SRDF QoS) Tool ... 15-12

Performance Management Overview

Performance management initially involves gathering real-time and historical performance data about your Symmetrix arrays, hosts, and Oracle databases. You collect and analyze this data using the tools described in Chapter 9, *Monitoring and Analyzing Performance*.

If you determine that there are “hot spots” (disks with a higher volume of I/O than others) on a specific Symmetrix array, you can use the *Optimizer* tool outlined in this chapter to balance the load among selected physical drives for better performance.

Optimizer is designed to operate in the background, detecting physical drive performance problems, and automatically swapping logical drives to balance the load. Optimizer should be active at all times, guarding the back-end performance of your array.

Performance management also includes the use of Quality of Service tools to increase the response time for TimeFinder BCV or SRDF copy operations on selected devices to improve the overall performance of the Symmetrix array. Refer to *Using the Quality of Service (TimeFinder/SRDF QoS) Tool* on page 15-12 for details.

The functionality described in this chapter is not available at all Symmetrix microcode levels. For additional information, contact an EMC representative.

The ControlCenter Optimizer tool performs the following:

- ◆ Analyzes Symmetrix logical device activity
- ◆ Determines which logical devices to swap
- ◆ Swaps the logical devices and their data

This process is explained in detail in *Understanding Optimizer* on page 15-3.

In general, the following steps are required to set up Optimizer:

- ◆ *Setting the Startup Mode* on page 15-9
- ◆ *Specifying the Workload Analysis Period* on page 15-9
- ◆ *Setting the Swap Mode* on page 15-10
- ◆ *Specifying the Swap Settings* on page 15-10
- ◆ *Setting Device Attributes* on page 15-10
- ◆ *Configuring Time Windows* on page 15-10

Understanding Optimizer

Optimizer is a tool that performs self-tuning of Symmetrix data configurations from the Symmetrix service processor by:

1. Collecting and analyzing samples of Symmetrix back-end logical device activity.
2. Determining which logical devices should have their physical locations moved to enhance Symmetrix performance.
3. Swapping logical devices and their data using internal Dynamic Reallocation Volumes (DRVs) to hold customer data while reconfiguring (on a device-to-device basis) the devices chosen for optimization.

A related tool, Quality of Service, allows you to set priority levels limiting TimeFinder and SRDF copy operations. Refer to *Using the Quality of Service (TimeFinder/SRDF QoS) Tool* on page 15-12.

Optimizer Process

Once configured, the Optimizer runs continuously in the background, analyzing the performance of logical and physical devices, and performing swaps as dictated by the internal algorithms.

There are a variety of manual actions that you can take, including:

- ◆ Configuring time windows for when performance should be analyzed.
- ◆ Configuring time windows for when swaps may take place.
- ◆ Manually creating your own swap lists.
- ◆ Rolling back swaps to an earlier point in time.
- ◆ Viewing graphs of performance.

Optimizer Devices

The following types of devices are used by Optimizer for logical device swapping:

- ◆ Logical devices
- ◆ DRV's

Logical Devices

A logical device is a unit of storage. In a Symmetrix array, you can define multiple logical devices on a single physical disk device.

Logical devices must be locally mirrored to work with Optimizer.

DRVs

A DRV is a Dynamic Reallocation Volume. It is a *nonuser-addressable* unit of storage used by Optimizer to hold customer data while reconfiguration (on a hypervolume granularity) is executed.

Optimizer Capabilities and Limitations

Optimizer provides the following capabilities:

- ◆ Improve subsystem performance
- ◆ Reduce response time during peak I/O loads
- ◆ Swap process is transparent to software applications
- ◆ Swap process is transparent to hosts (open systems or mainframes)
- ◆ Support mixed-drive environments.

Optimizer has the following performance limitations:

- ◆ Seeks to optimize from the perspective of the Symmetrix internal disk access only (Symmetrix back-end). It does not attempt to optimize across the entire array.
- ◆ Works more efficiently with arrays that display changes in patterns of behavior over medium-to-longer periods of time.
- ◆ Uses forecasting methods that assume past behavior predicts future behavior.
- ◆ Seeks to optimize performance over long periods of time.

Swapping Logical Devices

This section describes the requirements and process for swapping logical devices.

Logical Device Swapping Requirements

To swap logical devices, the devices must:

- ◆ Be the same size.
- ◆ Be the same emulation.
- ◆ Reside on the same system bus pair (applies only to Symmetrix 4xxx and 5xxx).
- ◆ Not be configured as BCV devices nor as AS400 devices.
- ◆ Not be configured as RAID devices, except for RAID5 (requires Enginuity 5x70.5).

You must have DRVs configured and available on your Symmetrix array. DRVs are configured using Symmetrix Manager. Refer to the Console online Help topic **Symmetrix Manager: Defining the device type** devices for more information.

Logical Device Swapping Process

Logical device swapping relies on the availability of DRV devices. DRVs are Symmetrix devices that are specially configured in the Symmetrix array to hold data while logical devices are being swapped. The DRV enables the data to remain protected and fully accessible during the swap process.

Optimizer initiates logical device swapping only after it has evaluated the logical device activity statistics and developed a reconfiguration plan.

The logical device swapping process has five main steps:

1. A pair of logical devices is identified for swapping.
2. Each logical device to be swapped is associated with a DRV and data from the logical device is copied to the designated DRV.
3. The address locations of the logical devices are swapped.
4. Data is copied back from the DRVs to the designated logical devices in their new locations.
5. The DRVs are split from the logical devices.

Understanding Optimizer Time Windows

The time window feature allows you to configure aspects of Optimizer's behavior. A time window is a period in time during which an aspect of Optimizer's behavior is controlled.

Time windows have the following characteristics:

- ◆ An effective time range, consisting of the following:
 - A start time and end time (may be infinite in the start direction, end direction, or both) in increments of 30 minutes, and aligned on the half-hour
 - A periodicity (daily, weekly, or ranging)
- ◆ A type of behavior and whether to include or exclude the behavior during the time window

In Figure 15-1 on page 15-7, the lower table is an editable summary of either the performance or swap time windows.

The type of list is controlled by the menu on the left side. The position of the time windows in the table assigns them a priority, with the first row having the highest priority. If multiple time windows have time ranges that overlap each other, the higher-listed time window overrides the others. The order of time windows in the list resolves conflicts between overlapping time windows.

Conflict resolution only applies to time windows of the same type.

Types of Time Windows

Optimizer provides two types of time windows:

- ◆ Performance
- ◆ Swap

Performance Time Windows

Performance Time Windows allow you to specify date and time ranges (past or future) when performance samples taken will be included in or excluded from Optimizer analysis (Figure 15-1). Optimizer includes all performance samples in its analysis by default.

Swap Time Windows

Swap Time Windows allow you to specify date and time ranges when Optimizer should or should not start swap activity. By default, Optimizer performs no swap activity.




Figure 15-1 Performance Time Window

Using Optimizer

Optimizer is a client/server architecture application. On the server side, the Optimizer runs on the Symmetrix service processor, and directly connects to the Symmetrix array. The Symmetrix array must be locally attached to the host before you can use Optimizer client.

Accessing Optimizer

Access Optimizer from the ControlCenter Console:

1. From the **Storage** folder in the tree panel, expand the Symmetrix folder, and right-click the appropriate Symmetrix array.
2. Select **Optimizer, Setting** from the shortcut menu (or highlight the device and select **Optimizer, Settings** from the toolbar). The Optimizer dialog box appears (Figure 15-2).

Optimizer shares a locking mechanism with Symmetrix Manager so that only one configuration process is running on the service processor at a time. If an error message occurs stating that the option is locked, wait for the application to release the lock and try again.




Figure 15-2 Optimizer Dialog Box

Monitoring Optimizer Server Status

The **Parameters** tab (Figure 15-2) allows you to monitor Optimizer status and set Optimizer parameters.

The **Optimizer Status** field allows you to monitor Optimizer operating status. It includes the following information:

- ◆ **State:** Displays Optimizer's current state. Valid values are:
 - **Running**
 - **Not Running**
 - **User Defined Swap Session**
 - **Error**
- ◆ **Activity:** Displays Optimizer's current activity. Valid values are:
 - **Idle** — Optimizer is between operations.
 - **Fetching Statistics** — Optimizer is collecting and analyzing samples of Symmetrix logical device activity.
 - **Performance Analysis** — (i.e., generating swap suggestions) Optimizer is determining which logical devices to swap.
 - **Swapping** — Optimizer is swapping logical devices.

The **Client Status** field indicates if this client holds the **Lock** (can change the parameters). The **Mode** is **Read** (if the client does not hold the lock) or **Read-Write** (if the client holds the lock).

Setting Optimizer Parameters

Be aware of the following when setting Optimizer parameters:

- ◆ Parameters can only be modified while Optimizer is stopped.
- ◆ The changes made to the parameters are not activated until you click **Apply** or **OK**. If you do not click **Apply** or **OK**, the settings remain as they were before the changes were made. Click **Cancel** if you do not want to save your changes.

Setting the Startup Mode

Select either the **Manual** or **Automatic** start option. The Automatic option starts Optimizer automatically when the service processor in the Symmetrix array is started.

Specifying the Workload Analysis Period

Specify the **Workload Analysis Period** (the amount of workload sampling that Optimizer should maintain for analysis). For example, if your I/O load characteristics repeat every week, you should set this parameter to 1 week (default), since this characterizes your typical workload.

Set the **Initial Period**, which specifies the minimum amount of workload sampling that Optimizer should complete before analyzing the performance samples for the first time. This enables Optimizer to start analysis and swap activities before the entire Workload Analysis Period has elapsed. This parameter is specified in hour/day/week units with a maximum of the Workload Analysis Period specified previously.

Setting the Swap Mode

Set the **Swap Mode** to either **Automatic** to have Optimizer perform swaps without user permission, or **User Approved** to have Optimizer prompt you for approval before each swap operation.

Specifying the Swap Settings

Set the **Swap Settings** as follows:

1. Specify the **Maximum Swaps per Day** in a 24-hour period starting at midnight. This is only applicable in Automatic mode.
2. Specify the **Maximum Simultaneous Swaps** (maximum number of devices to move simultaneously). This setting gives you better control over Optimizer moves and array performance.

Setting Device Attributes

Device attributes allows you to give priority to specific logical devices during optimization. The priority settings are:

- ◆ **High Priority** — Assign this device the highest priority because it contains crucial data. Optimizer attempts to achieve the best performance for this device without sacrificing the performance of other devices in this high-priority group.
- ◆ **Normal Priority** — This device is eligible for swap, but assign it a normal priority.
- ◆ **No Swap** — Do not swap.

Configuring Time Windows

A **time window** is a period of time during which an aspect of Optimizer behavior is controlled. You can configure the following time windows:

- ◆ **Performance Time Windows** — Specify which samples, based on date and time ranges (past or future), will be included in or excluded from Optimizer analysis.
- ◆ **Swap Time Windows** — Specify date and time ranges for Optimizer to start swap activity.

Specific procedures for time windows are described in the online Help.

Retrieving Optimizer Logs

The log feature allows you to retrieve current and past Optimizer state information. Optimizer logs contain operational information (All Activity), as well as specific information on swap activity and errors (Error Log) it may have encountered. The log appears as read-only text.

Log data is kept in memory, not files, by Optimizer. You can use the **Export** button to save the current log as a file.

From the Optimizer dialog box, click the **Log** tab (Figure 15-3).




Figure 15-3 Log Tab

Using the Quality of Service (TimeFinder/SRDF QoS) Tool

The Quality of Service (QoS) tool provides flexibility in managing a Symmetrix array's performance. By imposing a delay before BCV or SRDF copy operations on selected devices, you can increase the overall performance of other Symmetrix devices.

Performance Tuning BCV and SRDF Copy

You can specify BCV or SRDF QoS performance settings for each device in your Symmetrix array. This setting, 0 (default) to 10 (lowest copy rate), affects the service that Symmetrix provides to copy operations associated with the device type.

A QoS interface lets you view and change the current settings for each device. You choose the appropriate settings for each device based on its priority within your storage array, and apply the settings through the interface. QoS settings can be changed at any time to adjust for changes in your array I/O requirements.

QoS Performance Settings

The default QoS performance setting of 0 allows the device to receive full service as resources are available within the array. Applying an incremental value from 1 to 10 introduces a delay time before each track is copied between standard and BCV devices or SRDF R1 and R2 devices. Figure 15-4 shows the 1-10 scale, which represents a nonlinear scale from 1 millisecond to 1 second.




Figure 15-4 Performance Controls

A

- active alerts view 8-5
- agents, host 12-13
- Alert Log, scanning 4-8
- alerts 3-2
 - assigning 3-9
 - autofixes 3-18
 - best practices 3-23
 - creating 3-11
 - finding 8-14
 - notes 8-13
 - resolution 3-9
 - responding to 8-9
 - send to e-mail and pager 3-17
 - severity level 3-8
 - spikes 3-16
 - templates 3-13
 - testing 3-14
 - triggered 8-9
 - troubleshooting 3-25
- Any User rule 1-6
- At A Glance view 3-3, 8-3
- authorization rules
 - Any User rule 1-6
 - ECCAdministrators rule 1-6
 - SAN Manager rule 1-6
 - StorageScope 13-4
 - Symmetrix Configuration Manager rule 1-6
 - Symmetrix Data Protection Manager rule 1-6
 - Symmetrix Performance Manager rule 1-6
- autofix 3-12, 8-16
 - creating definition 3-18
 - status 8-12

- troubleshooting 3-25
- automated alert responses 3-12, 8-16

B

- backing up the Repository 4-2
- BCV configuration 11-10
- BCV devices, creating 11-12, 11-26
- BCVs 11-11
- bottlenecks, Windows 12-11
- BRBCV devices 11-11

C

- ChangeMembership permission 1-11
- CLARiiON
 - data protection 11-33
 - free space 10-6
 - Navisphere Manager 11-33
- CLI 11-25
- collection policies
 - accessing 2-5
 - managing 2-4
- combination agent data collection policy 2-8
- comments xxii
- common agent data collection policy 2-7
- compression 12-9
 - Novell storage 12-9
 - Windows storage 12-8
- configuring device mirrors 11-4
- connectivity agent data collection policy 2-8
- Console
 - information panel 5-7
 - menu bar 5-4
 - target panel 5-10

- taskbar 5-5
- toolbar 5-6
- tree panel 5-8
- using 5-3
- creating, user-defined groups 5-16
- Customer Support xxii

D

- data collection 9-9
 - policy templates 2-2
 - predefined policy 2-2
 - WLA Analyzer Archiver 9-9
- data collection policies
 - accessing 2-5
 - combination agent 2-8
 - common agent 2-7
 - connectivity agent 2-8
 - database agent 2-9
 - host agent 2-9, 2-11
 - managing 2-4
 - physical agent 2-11
 - storage agent 2-12
- data protection, CLARiiON 11-33
- database agent data collection policy 2-9
- definitions, creating 3-13
- device mirrors, configuring 11-4
- devices
 - BRBCV 11-11
 - host devices, properties 12-3
 - RBCV 11-11
 - STD 11-11
- discovery
 - agents 7-4
 - requirements 7-4
 - topology map editing 7-19
- disk contention 9-28
- disks, monitoring 8-14
- documentation, related xxi

E

- eccadmin 1-13
- ECCAdministrators 1-3
- ECCAdministrators rule 1-6
- e-mail, alert notification 3-12, 8-16

F

- file systems
 - monitoring 8-14
 - page spaces 12-10
 - properties 12-3
 - relationships with devices 12-4
- files
 - large 12-8
 - monitoring 8-14
 - obsolete 12-7
 - properties 12-8
- fragmentation, monitoring 8-14
- free space 10-6
- frequency, alert checking 3-12, 8-16

G

- graph, points 9-11
- groups and inheritance 1-10

H

- help xxii, 5-23
- high-level qualifiers, exploring 12-14
- host agent data collection policy 2-9, 2-11
- host CPU utilization and response time 9-18
- HSM orbiting data sets 12-19

I

- inheritance
 - ChangeMembership 1-11
- initialization parameters
 - performance time windows 15-10
 - swap time windows 15-10

J

- jobs using excessive space 12-17

L

- limitations on performance 15-4
- load balancing 9-14
- lock mechanism, Optimizer 15-8
- logical volume swapping 15-5
- login history table 7-21
- logs, size of 12-8

M

- management
 - policy 3-12, 8-16
 - user access 1-12
- migration 12-9
- mirrors, configuring 11-4
- MVS 12-13
 - exploring host storage resources 12-14
 - increasing performance 12-18
 - orbiting data sets 12-19
 - recapturing inefficient storage 12-16

N

- Navisphere Manager 11-33
- NetWare Server 12-9
- new user, creation 1-13
- Novell NetWare storage 12-9

O

- online help 5-23
- Optimizer
 - capabilities 15-4
 - Dynamic Reallocation Volume (DRV) 15-4
 - lock 15-8
 - logical volume swapping 15-5
 - logical volumes (devices) 15-4
 - overview 15-3
 - performance limitations 15-4
 - Quality of Service 15-3
 - setting parameters 15-9
 - swap process 15-3, 15-4
 - Symmetrix service processor 15-3
 - time windows 15-6

P

- page spaces 12-10
- path details 7-50
- performance
 - archives 9-11
 - data collection, historical and revolving 9-10
 - initialization parameter 15-10
 - points defined 9-11
 - setting QoS 15-12
 - statistics 7-38
 - time windows 15-7

- tuning 15-12
 - Windows 12-11
- permissions
 - assigning 1-6
 - ChangeMembership 1-11
 - tasks 1-7
- physical agent data collection policy 2-11
- Powerlink website xxii
- purging files 12-9

Q

- Quality of Service 15-3, 15-12

R

- RA groups 11-9
- RAID 10-13
- RBCV devices 11-11
- real-time performance analysis 9-3
- recompiling invalid objects in Repository 4-5
- Repository
 - analyzing tables 4-4
 - backing up 4-2
 - cleaning trace files 4-8
 - export backup 4-4
 - monitoring tablespace growth 4-5
 - overview 4-1
 - recompiling invalid objects 4-5
 - scanning alert logs 4-8
 - shutting down 4-7
 - starting 4-8
 - tablespace fragmentation 4-8
- requirements for swapping 15-5
- revolving performance data collection 9-10
- rule, create new 1-20

S

- SAN Manager 1-4
- SAN Manager rule 1-6
- schedules, alert 3-12, 8-16
- Service xxii
- setting parameters, Optimizer 15-9
- Solutions Enabler SYMCLI 11-25
- sorting columns 5-17
- space, rapid consumption of 12-17
- SPS 10-14

- scheduling tasks 10-24
 - storage policies 10-10
 - storage pool 10-8
 - TaskList 10-25
 - SRDF
 - basic configuration 11-16
 - monitoring 11-30
 - SRDF QoS 15-12
 - starting the Repository 4-8
 - status, autofix 8-12
 - STD devices 11-11
 - determine size 11-27
 - storage
 - compression of Novell NetWare storage 12-9
 - compression of Windows files 12-8
 - storage agent data collection policy 2-12
 - storage policies 10-10
 - storage pools 10-8
 - Storage Provisioning Service 10-14
 - scheduling tasks 10-24
 - storage policies 10-10
 - storage pools 10-8
 - TaskList 10-25
 - StorageScope
 - authorization rules 13-4
 - launching 13-6
 - logging in 13-8
 - opening reports 13-7
 - overview 13-2
 - permissions, types of 13-4
 - reports, types of 13-2
 - StorageWorks
 - free space 10-6
 - swap process 15-5
 - swap space 12-10
 - swap time windows 15-7
 - swap time windows initialization parameter 15-10
 - swapping requirements 15-5
 - SYMCLI 11-25
 - Symmetrix
 - Configuration Manager 1-4
 - Configuration Manager rule 1-6
 - Data Protection Manager 1-4
 - Data Protection Manager rule 1-6
 - free space 10-6
 - Performance Manager 1-4
 - SDM agent 7-21
 - system information 9-20
 - unallocated storage 10-5
 - Symmetrix Performance Manager rule 1-6
- ## T
- tables
 - Relationship view 12-5
 - sorting columns 5-17
 - tablespace
 - fragmentation in Repository 4-8
 - growth monitoring 4-5
 - Technical Support xxii
 - threshold, alert 3-12
 - thresholds 3-2
 - time windows
 - conflict resolution 15-6
 - performance 15-10
 - swap 15-10
 - TimeFinder
 - BCV devices 11-11
 - BRBCV devices 11-11
 - clone commands 11-14
 - commands 11-13
 - components 11-11
 - monitor operations 11-29
 - RBCV devices 11-11
 - STD devices 11-11
 - TimeFinder/SRDF QoS 15-12
 - topology map, edit 7-19
 - trace files, cleaning 4-8
 - tree panel 5-8
 - trigger values, alert 3-12, 3-14, 8-15
 - triggered alerts 8-9
- ## U
- unallocated storage 10-5
 - unidentified ports, Web console 6-5
 - user
 - access management overview 1-12
 - change ID 1-16
 - create new 1-13
 - user groups
 - add user 1-18
 - create new 1-17
 - default 1-3

- ECCAdministrators 1-3
- inheritance
 - of permissions 1-10
- SAN Manager 1-4
- Symmetrix Configuration Manager 1-4
- Symmetrix Data Protection Manager 1-4
- Symmetrix Performance Manager 1-4
- user-defined groups, creating 5-16
- users using excessive space 12-17

V

- views
 - active alerts 8-5
 - At A Glance 3-3, 8-3

- path details 7-50
- preferences 5-21
- volume types and states 11-17

W

- Web console
 - unassigned ports 6-5
 - unidentified ports 6-5
- website, Powerlink xxii
- wildcards, using with alerts 3-11
- Windows 2000 users 1-13
- Windows storage compression 12-8
- WLA Daily 9-9
- WLA Revolving 9-9

