

Archive Backup System for OpenVMS

Guide to Operations

Order Number: AA-QHD1Q-TE

Software Version	Archive Backup System for OpenVMS Version V4.0A
Required Operating System	OpenVMS VAX Version 6.2, 7.1, 7.2 or 7.3 & OpenVMS Alpha Version 6.2, 7.1-2, 7.2-1, 7.2-2 or 7.3
Required Software	Media, Device and Management Services for OpenVMS Version V4.0A DECnet (Phase IV) or DECnet-Plus(Phase V) TCP/IP Services for OpenVMS

April 2002

© 2002 Compaq Information Technologies Group, L.P

Compaq, the Compaq logo, OpenVMS, VAX and Tru64 are trademarks of Compaq Information Technologies Group, L.P. in the U.S. and/or other countries. UNIX is a trademark of The Open Group in the U.S. and/or other countries. All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Compaq service tool software, including associated documentation, is the property of and contains confidential technology of Compaq Computer Corporation. Service customer is hereby licensed to use the software only for activities directly relating to the delivery of, and only during the term of, the applicable services delivered by Compaq or its authorized service provider. Customer may not modify or reverse engineer, remove, or transfer the software or make the software or any resultant diagnosis or system management data available to other parties without Compaq's or its authorized service provider's consent. Upon termination of the services, customer will, at Compaq's or its service provider's option, destroy or return the software and associated documentation in its possession.

Printed in the U.S.A.

Contents

Preface	xvii
----------------------	-------------

1 Introduction

2 Overview

2.1 ABS Operational Environment	2-1
2.2 ABS Objects	2-1
2.2.1 Saves	2-2
2.2.2 Restores	2-2
2.2.3 Archives	2-3
2.2.4 Environments	2-4
2.2.5 Selections	2-4
2.2.6 Schedules	2-5
2.3 ABS Catalogs	2-5
2.4 Backup Agent	2-6
2.5 Media, Device and Management Services (MDMS)	2-7
2.6 User Interfaces	2-7
2.7 Scheduler Options	2-7
2.8 MDMS Objects	2-8
2.8.1 Domain	2-8
2.8.2 Drives	2-8
2.8.3 Groups	2-9
2.8.4 Jukeboxes	2-9
2.8.5 Locations	2-10
2.8.6 Magazines	2-10
2.8.7 Media Types	2-10
2.8.8 Nodes	2-11
2.8.9 Pools	2-11
2.8.10 Volumes	2-11
2.9 Getting Started	2-11

3 Saving and Restoring Data

3.1 Archives	3-1
3.1.1 Archive Name	3-2
3.1.2 Archive Type	3-2
3.1.3 Catalog	3-2
3.1.4 Consolidation	3-2
3.1.5 Destination	3-3
3.1.6 Drives	3-3
3.1.7 Expiration Date and Retention Days	3-3
3.1.8 Location	3-4
3.1.9 Maximum Saves	3-4

3.1.10	Media Type	3-4
3.1.11	Pool	3-4
3.1.12	Volume Sets	3-4
3.2	Catalogs	3-4
3.2.1	Catalog Name	3-5
3.2.2	Catalog Node	3-5
3.2.3	Type	3-5
3.2.4	Directory	3-6
3.2.5	Staging	3-6
3.2.6	Catalog Save Entries	3-6
3.2.7	Catalog File Entries	3-7
3.2.8	Improving Catalog Performance	3-8
3.2.8.1	Catalog File Sizes	3-8
3.2.8.2	Catalog File Maintenance	3-8
3.2.8.3	Staging Catalog	3-9
3.3	Cataloging Existing Savesets	3-9
3.4	Environments	3-11
3.4.1	Environment Name	3-11
3.4.2	Action	3-11
3.4.3	Compression	3-11
3.4.4	Data Safety	3-12
3.4.5	Drive Count	3-12
3.4.6	Prologue and Epilogue	3-12
3.4.7	Retry Limit and Interval	3-13
3.4.8	Links Only and Span Filesystems	3-13
3.4.9	Listing Option	3-14
3.4.10	Lock	3-14
3.4.11	Notification	3-14
3.4.12	Profile	3-15
3.5	Saves and Restores	3-15
3.5.1	Save Name or Restore Name	3-16
3.5.2	Archive	3-16
3.5.3	Base Date, Start Date and Skip Time	3-16
3.5.4	Before Date, Since Date and Date Archived (Restore Only)	3-17
3.5.5	Catalog (Restore Only)	3-17
3.5.6	Include, Exclude, Data Type and Source Node	3-18
3.5.7	Delete Interval and Keep	3-19
3.5.8	Destination (Restore Only)	3-20
3.5.9	Environment	3-20
3.5.10	Frequency and Explicit Interval	3-20
3.5.11	Incremental	3-24
3.5.12	Nodes and Groups	3-24
3.5.13	Prologue and Epilogue	3-25
3.5.14	Reschedule	3-26
3.5.15	Selections	3-26
3.5.16	Sequence Option (Saves Only)	3-26
3.6	Selections	3-27
3.6.1	Agent Qualifiers	3-27
3.6.2	Before Date, Since Date and Date Type (Saves Only)	3-27
3.6.3	Conflict Options (Restore Only)	3-28
3.6.4	Include, Exclude, Data Type and Source Node	3-28
3.7	Schedules	3-29

3.7.1	After Schedule	3-30
3.7.2	Command	3-30
3.7.3	Dates, Days and Months	3-31
3.7.4	Include and Exclude	3-32
3.7.5	Times	3-33

4 Media Management

4.1	MDMS Domain Configuration	4-1
4.2	Domain	4-1
4.2.1	ABS Rights	4-2
4.2.2	Application Rights	4-2
4.2.3	Check Access	4-2
4.2.4	Deallocate State	4-2
4.2.5	Default Rights	4-2
4.2.6	Mail Users	4-2
4.2.7	Maximum Scratch Time	4-3
4.2.8	Media Type	4-3
4.2.9	Offsite Location	4-3
4.2.10	Onsite Location	4-3
4.2.11	OPCOM Classes	4-3
4.2.12	Operator Rights	4-3
4.2.13	Protection	4-3
4.2.14	Relaxed Access	4-4
4.2.15	Request ID	4-4
4.2.16	Scheduler Type	4-4
4.2.17	Scratch Time	4-4
4.2.18	SYSPRV	4-4
4.2.19	Transition Time	4-5
4.2.20	User Rights	4-5
4.3	Drives	4-5
4.3.1	Access	4-5
4.3.2	Automatic Reply	4-5
4.3.3	Device	4-5
4.3.4	Disabled	4-6
4.3.5	Drive Number	4-6
4.3.6	Groups	4-6
4.3.7	Jukebox	4-6
4.3.8	Media Types	4-6
4.3.9	Nodes	4-6
4.3.10	Read-Only Media Types	4-6
4.3.11	Shared	4-6
4.3.12	Stacker	4-7
4.3.13	State	4-7
4.3.14	Allocate Drive (DCL Only)	4-7
4.3.15	Deallocate Drive (DCL Only)	4-8
4.3.16	Load Drive	4-8
4.3.17	Unload Drive	4-8
4.4	Groups	4-9
4.4.1	Nodes	4-9
4.5	Jukeboxes	4-9
4.5.1	Access	4-9
4.5.2	ACS ID	4-9

4.5.3	Automatic Reply	4-9
4.5.4	Cap Size	4-10
4.5.5	Control	4-10
4.5.6	Disabled	4-10
4.5.7	Groups	4-10
4.5.8	Library ID	4-10
4.5.9	Location	4-10
4.5.10	LSM ID	4-10
4.5.11	Nodes	4-11
4.5.12	Robot	4-11
4.5.13	Slot Count	4-11
4.5.14	State	4-11
4.5.15	Threshold	4-11
4.5.16	Topology	4-12
4.5.17	Usage	4-12
4.5.18	Inventory Jukebox	4-12
4.6	Locations	4-13
4.6.1	Parent Location	4-14
4.6.2	Spaces	4-14
4.7	Magazines	4-14
4.7.1	Jukebox, Start Slot and Position	4-14
4.7.2	Onsite and Offsite Locations and Dates	4-14
4.7.3	Slot Count	4-15
4.7.4	Spaces	4-15
4.7.5	Move Magazine(s)	4-15
4.8	Media Types	4-16
4.8.1	Capacity	4-16
4.8.2	Compaction	4-16
4.8.3	Density	4-16
4.8.4	Length	4-16
4.9	Node	4-16
4.9.1	Database Server	4-16
4.9.2	Disabled	4-17
4.9.3	OPCOM Class	4-17
4.9.4	Transports and Full Names	4-17
4.10	Pools	4-17
4.10.1	Authorized Users	4-17
4.10.2	Default Users	4-18
4.10.3	Threshold	4-18
4.11	Volumes	4-18
4.11.1	Allocation Fields - Account, Username, UIC and Job	4-20
4.11.2	Allocation and Movement Dates	4-20
4.11.3	History Dates	4-21
4.11.4	State	4-21
4.11.5	Media Types	4-22
4.11.6	Pool	4-22
4.11.7	Previous and Next Volumes	4-22
4.11.8	Placement - Jukebox, Magazine, Locations, Drive	4-23
4.11.9	Formats - Brand, Format, Block Factor, Record Size	4-23
4.11.10	Protection	4-24
4.11.11	Counters	4-24
4.11.12	Allocate Volume	4-24

4.11.13	Allocate Volume(s) by Selection Criteria	4-25
4.11.14	Deallocate Volume	4-25
4.11.15	Bind Volume	4-26
4.11.16	Unbind Volume	4-26
4.11.17	Load Volume	4-27
4.11.18	Unload Volume	4-27
4.11.19	Move Volume(s)	4-27
4.11.20	Initialize Volume(s)	4-28

5 Security

5.1	MDMS Rights	5-2
5.2	Access Control	5-4
5.3	Implementing a Security Strategy	5-5

6 User Interfaces

6.1	Graphical User Interface	6-1
6.1.1	Starting MDMSView	6-2
6.1.1.1	OpenVMS Systems	6-2
6.1.1.2	Windows Systems	6-2
6.1.2	Look and Feel	6-2
6.1.3	Logging In	6-3
6.1.4	Selecting A View	6-3
6.1.5	Creating Objects	6-5
6.1.6	Showing and Modifying Objects	6-6
6.1.7	Deleting Objects	6-8
6.1.8	Viewing Relationships Between Objects	6-8
6.1.9	Performing Operations on Objects	6-9
6.1.10	Running Save And Restore Requests	6-9
6.1.11	Showing Current Operations	6-10
6.1.12	Reporting on Volumes	6-11
6.1.13	Viewing MDMS Audit and Event Logging	6-13
6.1.14	Errors	6-13
6.1.15	Help	6-14
6.2	DCL Interface	6-14
6.2.1	Syntax Overview	6-15
6.2.2	Object Lists	6-16
6.2.3	Qualifier List	6-16
6.2.4	Inherit	6-17
6.2.5	Symbols	6-17
6.2.6	Help and Reference	6-17
6.3	User Interface Restrictions	6-17

7 Preparing For Disaster Recovery

7.1	Disaster Recovery for OpenVMS Systems	7-1
7.1.1	Backup of Your System Disk	7-1
7.1.2	Backup of MDMS\$ROOT	7-3
7.1.3	Backup of ABS\$ROOT	7-4
7.2	Prolog and Epilog Procedure	7-5
7.2.1	Restoring The System Disk	7-7
7.2.2	Restoring Remaining Savesets	7-7

7.3 Non-OpenVMS Systems	7-8
7.4 Thoughts on Save and Restore Procedures	7-8

8 Remote Devices

8.1 The RDF Installation	8-1
8.2 Configuring RDF	8-1
8.3 Using RDF with MDMS	8-2
8.3.1 Starting Up and Shutting Down RDF Software	8-2
8.3.2 The RDSHOW Procedure	8-2
8.3.3 Command Overview	8-2
8.3.4 Showing Your Allocated Remote Devices	8-2
8.3.5 Showing Available Remote Devices on the Server Node	8-2
8.3.6 Showing All Remote Devices Allocated on the RDF Client Node	8-3
8.4 Monitoring and Tuning Network Performance	8-3
8.4.1 DECnet Phase IV	8-3
8.4.2 DECnet-Plus (Phase V)	8-4
8.4.3 Changing Network Parameters	8-4
8.4.4 Changing Network Parameters for DECnet (Phase IV)	8-5
8.4.5 Changing Network Parameters for DECnet-Plus(Phase V)	8-6
8.4.6 Resource Considerations	8-6
8.4.7 Controlling RDF's Effect on the Network	8-8
8.4.8 Surviving Network Failures	8-8
8.5 Controlling Access to RDF Resources	8-9
8.5.1 Allow Specific RDF Clients Access to All Remote Devices	8-9
8.5.2 Allow Specific RDF Clients Access to a Specific Remote Device	8-9
8.5.3 Deny Specific RDF Clients Access to All Remote Devices	8-10
8.5.4 Deny Specific RDF Clients Access to a Specific Remote Device	8-10
8.6 RDserver Inactivity Timer	8-10
8.7 RDF Error Messages	8-11

9 System Backup to Tape for Oracle Databases

9.1 Linking System Backup to Tape with the Oracle Server	9-2
9.1.1 Testing Oracle's Recovery Manager before linking System Backup to Tape	9-2
9.1.2 Authorizing privileges and granting rights to the Oracle server account	9-3
9.1.3 Editing Oracle's Link Option File and Command Procedures	9-3
9.1.3.1 Editing Oracle8i Link Option File and Command Procedures	9-3
9.1.3.2 Editing Oracle9i Link Option file and Command Procedures	9-4
9.1.4 Shutdown the database	9-5
9.1.5 Relinking the ORA_RDBMS: executables	9-5
9.1.6 Startup the database	9-5
9.1.7 Retesting Oracle's Recovery Manager	9-6
9.2 Defining the logical MDMS\$SBT_TRACE_LEVEL	9-6
9.3 Configuring System Backup to Tape in the Archive Backup system	9-6
9.3.1 Creating an ORACLE_DB Catalog	9-7
9.3.2 Creating an Archive	9-7
9.4 Testing the Configuration of SBT	9-9
9.5 Using System Backup to Tape with Oracle's Recovery Manager	9-10
9.5.1 Specifying an Archive	9-11
9.5.2 Specifying a Catalog	9-11
9.5.3 Specifying an I/O Block Size	9-12
9.5.4 Specifying Archives for Duplex Backups	9-12

9.6	Using the Show Catalog Command	9-13
9.7	Using the MDMS Scheduler	9-15
9.8	System Backup to Tape Defaults	9-16
9.8.1	Archive Name	9-16
9.8.2	Catalog Name	9-16
9.8.3	I/O Block Size	9-16
9.8.4	System Backup to Tape Logicals Names	9-16
9.9	System Backup to Tape Restrictions	9-17
9.9.1	Doing Parallel Backups	9-17
9.9.2	Piece Name Length Greater than 254 Characters	9-18
9.9.3	Using RDF Drives with SBT	9-18
9.10	Troubleshooting Tips	9-18
9.10.1	Using the logical MDMS\$SBT_TRACE_LEVEL	9-18
9.10.2	Fatal Internal Error	9-20
9.10.3	Check ORA_DUMP:SBTIO.LOG for Errors	9-20
9.10.4	Using Tape I/O Slaves	9-21

10 Architecture

10.1	The Server Process	10-1
10.1.1	The Database (DB) Server	10-1
10.1.1.1	Database	10-1
10.1.1.2	Becoming a DB Server	10-2
10.1.1.3	Finding another DB Server	10-2
10.1.1.4	Failover of the DB Server	10-2
10.1.1.5	Role of the DB server	10-3
10.1.2	Server Communications	10-3
10.2	Scheduler Interface	10-3
10.2.1	Option INT_QUEUE_MANAGER	10-3
10.2.2	Option EXT_QUEUE_MANAGER	10-4
10.2.3	Option EXT_SCHEDULER	10-4
10.3	Catalogs	10-4
10.3.1	Catalog Sizes	10-5
10.4	Coordinator	10-6
10.4.1	Coordinator Cleanup	10-7
10.4.2	Volume Sets	10-7

11 Troubleshooting

11.1	Save and Restore Requests	11-1
11.1.1	Notification of Save/Restore Completion	11-1
11.1.2	Log Files	11-1
11.1.3	Logical Names	11-1
11.1.4	Alpha Stack Size Logical	11-1
11.1.5	Fast Skip Errors	11-1
11.1.6	Volume Set Locking and Coordinator Cleanup Process	11-2
11.2	Media Management	11-2
11.2.1	Log Files	11-2
11.2.2	OPCOM	11-2
11.2.3	MDMS Requests	11-3
11.2.4	Scheduling Problems	11-3
11.2.4.1	Internal Scheduling	11-3
11.2.4.2	External Scheduling	11-3

11.2.4.3	Scheduler Scheduling	11-3
11.2.5	MDMS Scheduled Activities	11-4
11.3	MDMSView GUI	11-4
11.3.1	Running MDMSView GUI After ABS/MDMS Installation	11-4
11.3.2	Windows Java Path	11-4
11.3.3	MDMSView Log Screen	11-4
11.3.4	MDMSView Command Window	11-4
11.3.5	MDMS\$LOGFILE_*.LOG	11-5
11.4	ABS Catalogs	11-5
11.4.1	Staging Unpack	11-5
11.4.2	Catalog Cleanup	11-5
11.5	Windows and Unix Clients	11-5
11.5.1	Windows Log File	11-5
11.5.2	Windows Quotas	11-6
11.5.3	Permission Denied Errors	11-7
11.5.4	UBS FAILURE	11-7
11.5.5	Considerations for Saving Large Disks on UNIX and NT Clients	11-7
11.5.6	Files Larger than 2gb	11-8
11.6	RDF (Remote Device Facility)	11-9
11.7	Information Required When Reporting Problems	11-9

A Configuration Example

Tables

Table 3-1	Logical Names Available to Environment Prologues and Epilogues	3-13
Table 3-2	Use of Base Date, Start Date and Skip Time	3-17
Table 3-3	Disk, File, Path and Database Specification Formats	3-19
Table 3-4	Logical Names in Save/Restore Prologues and Epilogues	3-25
Table 3-5	Disk, File, Path and Database Specification Formats	3-29
Table 3-6	Date Specifications	3-31
Table 3-7	Day Specifications	3-31
Table 3-8	Month Specifications	3-31
Table 3-9	Combining Dates, Days and Months	3-32
Table 4-1	MDMS Volume State Transitions	4-19
Table 5-1	Examples of Low Level Rights	5-2
Table 5-2	ABS to MDMS Rights Mapping	5-3
Table 5-3	Access Control Allowed Operations	5-4
Table 5-4	Domain Access Control Options	5-4
Table 8-1	How to Change Network Parameters	8-5
Table 9-1	SBT Logical Names	9-16
Table 11-1	Assigning a System Variable for NT Troubleshooting	11-6
Table 11-2	Modifying the Blocking Factor	11-8

Figures

Figure 2–1	ABS Save or Restore Request	2-3
Figure 2–2	ABS Catalogs	2-6
Figure 3–1	Relationships Between ABS Objects	3-16
Figure 3–2	Complex Backup Schedules	3-23
Figure 4–1	Volume State	4-19
Figure 6–1	MDMSView Main Screen.	6-3
Figure 6–2	Object View Screen.	6-5
Figure 6–3	Drive Create Screen	6-6
Figure 6–4	Save Show General Screen	6-7
Figure 6–5	Domain View Showing Expanded Relationships	6-8
Figure 6–6	Load Volume Screen with Queued Dialog Box	6-9
Figure 6–7	Save Log Screen	6-10
Figure 6–8	Show Requests Screen	6-11
Figure 6–9	Report View Selection Criteria Example	6-12
Figure 6–10	Report View Results Screen	6-12
Figure 6–11	Context-Sensitive Help Screen from Show Volume Screen.	6-14

Preface

Intended Audience

This document is intended for storage administrators who are experienced OpenVMS system managers. This document should be used in conjunction with the *Introduction to OpenVMS System Management* manual.

Conventions

The following conventions are used in this document:

Convention	Description
{ }	In format command descriptions, braces indicate required elements.
[]	In format command descriptions, square brackets indicate optional elements of the command syntax. You can omit these elements if you wish to use the default responses.
boldface type	Boldface type in text indicates the first instance of a term defined in the Glossary or defined in text.
<i>italic type</i>	Italic type emphasizes important information, indicates variables, indicates complete titles of manuals, and indicates parameters for system information.
Starting test ...	This type font denotes system response, user input, and examples.
Ctrl/ <i>x</i>	Hold down the key labels Ctrl (Control) and the specified key simultaneously (such as Ctrl/Z).
PF1 <i>x</i>	The key sequence PF1 <i>x</i> instructs you to press and release the PF1 key, and then press and release another key (indicated here by <i>x</i>).
<i>n</i>	A lowercase italic <i>n</i> denotes the generic use of a number. For example, 19 <i>nn</i> indicates a four-digit number in which the last two digits are unknown.
<i>x</i>	A lowercase <i>x</i> denotes the generic use of a letter. For example, <i>xxx</i> indicates any combination of three alphabetic characters.

Related Products

The following related products may be mentioned in this document:

Product	Description
HSM	HSM refers to Hierarchical Storage Management for OpenVMS software.
MDMS	MDMS refers to Media and Device Management Services for OpenVMS software.
OpenVMS	OpenVMS refers to OpenVMS operating system.
SLS	SLS refers to Storage Library System for OpenVMS software.

Associated Documents

The following documents are part of Archive Backup System for OpenVMS documentation set:

- *Archive Backup System for OpenVMS Installation Guide*
- *Archive Backup System for OpenVMS Guide to Operations*
- *Archive Backup System for OpenVMS MDMS Reference Manual*

Introduction

The Archive Backup System for OpenVMS (ABS) is a software product that allows you to save and restore data in a heterogeneous environment. ABS provides you with the ability to perform anything from full system backup operations to user-requested or user-created backup operations. ABS ensures data safety and integrity by providing a secure environment for save and restore operations.

ABS is based on an OpenVMS system environment, and all data is saved to (and restored from) archives on OpenVMS systems. However, ABS supports saving and restoring data that resides on nodes running the UNIX and Windows operating systems, as well as OpenVMS systems.

ABS enables you to implement a backup policy that allows you to save the data through automatic or repetitively scheduled save operations. It also enables you to save data randomly using a one-time-only save operation. ABS allows you to use different scheduler interface options to schedule requests. This feature allows you to customize the scheduling of save or restore requests to your system configuration.

Save and restore operations are accomplished by using two of the objects recognized by ABS, a save request and a restore request. These objects allow you to save data from online to either a offline volume or to another disk and, if necessary, allows you to restore that data to either its original location or to a different output location.

ABS tracks the location of data when saved as a result of an ABS save request. This information is kept in an ABS catalog. Upon request, ABS accesses the catalog to locate or restore the data. Chapter 2 provides an overview of ABS capabilities, and Chapter 3 describes ABS Save and Restore operations, and the associated ABS objects, in more detail.

ABS is integrated with Media, Device and Management Services (MDMS), which performs the following functions on behalf of ABS:

- Database Management Services - MDMS maintains the ABS database objects including saves, restores, archives, environments, catalogs, schedules and selections. Database management services are available within a distributed environment using either TCP/IP or DECnet communication protocols. Chapter 3 describes the ABS objects in detail.
- Media Management Services - MDMS maintains a set of physical and logical objects for management of backup hardware and media. These objects include domain, locations, nodes, groups, jukeboxes, drives, media types, pools, volumes and magazines. Chapter 4 describes Media Management Services in detail.
- Scheduling Services - MDMS provides extensive internal scheduling services for automatically scheduling ABS save and restore requests. Chapter 3 describes Scheduling Services.
- Security Services - MDMS provides flexible security options using rights, privileges and object access control for secure use in a distributed environment. Chapter 5 describes security services.

Introduction

- MDMSview - A graphical user interface that manages all ABS and MDMS objects using a view-based approach for navigation. Views currently supported include Objects, Tasks, Requests, Reports and Domain. Chapter 6 describes the Graphical User Interface.
- DCL Interface to the Database Objects - A comprehensive set of DCL commands to manage all ABS objects, compatible with the interface for MDMS media management objects. Chapter 6 describes DCL operation, with a full reference in the MDMS Reference Guide.

Planning for Disaster Recovery is an important part of any datacenter operation. Chapter 7 offers guidelines on how to plan for disaster recovery with ABS.

Chapter 10 offers an architectural overview of the ABS/MDMS system; you can use this to understand the internal operations of ABS and customize certain operational parameters.

A troubleshooting section has been added in Chapter 11. This chapter describes to how define extended logging options, and offers solutions for some of the more common problems that can occur in an ABS environment.

The appendix offers an example of configuring MDMS.

Overview

This chapter provides an overview of the various components that comprise an ABS/MDMS operational environment, and includes a simple example on how to get started with ABS. The overview discusses the following items:

- ABS Operational Environment
- ABS Objects
- ABS Catalogs
- Backup Agent
- Media, Device and Management Services (MDMS)
- User Interfaces
- Scheduler Options
- MDMS Objects
- Getting Started

This chapter provides an overview on the ABS and MDMS environment. See Chapter 3, *Saving and Restoring Data*, for detailed information on how to save and restore data using ABS. See Chapter 4, *Media Management*, for information about how to configure and maintain the media management environment.

2.1 ABS Operational Environment

ABS operational environment contains the following components:

- ABS objects - ABS objects define physical locations of saved data, the criteria under which save and restore requests are performed, and the save and restore requests themselves. ABS objects are described in Section 2.2.
- ABS catalogs - ABS catalogs are the components of ABS software that contain the history information about ABS save requests. Catalogs contain the records of data saved using ABS. Those records enable you to locate and restore data that was saved using ABS. ABS catalogs are described in Section 2.3.
- Backup agent - A backup agent is the utility that performs the actual data movement operation. For OpenVMS systems, the backup agents are the OpenVMS BACKUP Utility and the RMU Backup Utility. For UNIX and Windows clients, the supported backup agent is `gtar` (tape archiver). ABS uses `gtar` because most UNIX and Windows systems support it. Backup agents are described in Section 2.4.

2.2 ABS Objects

The following sections summarize the objects used by ABS to save and restore data. More detailed information about ABS objects may be found in Chapter 3, *Saving and Restoring Data*.

Overview

2.2 ABS Objects

2.2.1 Saves

A save request defines the data to be saved and executes upon immediate invocation or through an automatic, repetitive schedule. You can create save requests using either the MDMSView GUI or the CLI interface.

A save defines the following criteria:

- The data to back up - you can specify disks, files or databases to back up
- The type of data to back up (VMS files, Oracle Rdb databases or storage areas, UNIX files or Windows files)
- Whether the save is an incremental operation based on a previous save, or otherwise
- When to save the data (base date and frequency)
- Where to save the data (which archive to use)
- The length of time to keep the data (retention period or expiration date)
- Who can access a save request (for data safety)
- What environment to use to execute the save request
- Whether to perform pre- or postprocessing commands

To meet your site's backup requirements, you will need to create save requests that fulfill those requirements.

2.2.2 Restores

A restore request restores data from an archive back to online storage. You can create restore requests using either the MDMSView GUI or the CLI interface. Restore requests can be executed immediately or at a specified time. You can also schedule restores for repeated operations in the same manner as saves.

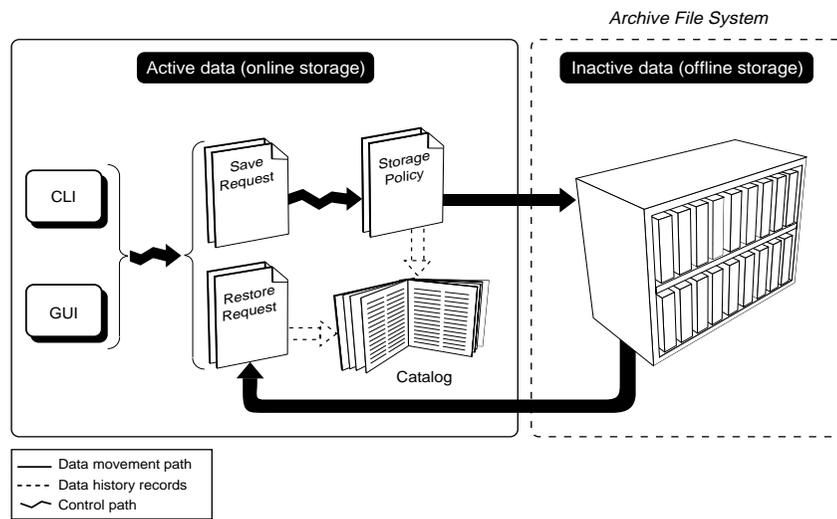
A restore defines the following criteria:

- The data to restore - you can specify disks, files or databases to restore
- The type of data to restore (VMS files, Oracle Rdb databases or storage areas, UNIX files or Windows files)
- Whether the restore is an incremental restore based on a previous restore, or otherwise
- Where to restore the data (optional output location other than the original location)
- Where the data resides (on which archive)
- Who can access a restore request (for data security)
- What environment to use to execute the restore request
- Whether to perform pre- or postprocessing commands

To meet your storage management requirements, you will need to create restore requests that fulfill those requirements.

Figure 2-1 illustrates the path of a save or restore request.

Figure 2–1 ABS Save or Restore Request



A save or restore request is invoked through the GUI or through the CLI (DCL).

IF the request is a . . .	THEN the data is . . .
Save request	Saved from online storage to the archive. An ABS catalog records the location of the saved data.
Restore request	Restored back to online storage. ABS searches the catalog for the location of the data (archive), loads the appropriate volume, and restores the data.

2.2.3 Archives

An archive defines the media type and other characteristics where you can safely store data. Each archive has a unique name and contains a set of archive characteristics. You can simply reference an archive name in a save or restore request rather than a complicated set of characteristics. Archives are designed to be shared among many save or restore requests.

Each archive defines the following:

- The type of archive to use (TAPE or DISK)
- If the archive file system is TAPE, the media type, pool, and location for tape volumes in the archive
- How long to keep the data stored in a particular archive (retention period or expiration date). You can specify two archives for save requests that perform both full and incremental operations (at different times) so that the full and incremental saves can have different retention periods and can reside on different volume sets
- Who is allowed to access the archive (for data safety)
- Who is allowed write data to and read data from the archive (ensures data safety)
- Which catalog contains the information about the data stored in the archive
- How long to use a volume set
- How many save or restore requests can be executed simultaneously

Overview

2.2 ABS Objects

Normally, one archive is associated with both save and restore requests. However, for save requests that perform both full and incremental saves (at different times), you can define two archives: the first for full saves and the second for incremental saves. This allows the full and incremental saves to be performed on different tape volumes with different retention periods.

2.2.4 Environments

An environment object defines the criteria under which save and restore requests are executed. The criteria defined in an environment include:

- Whom to notify when a backup or restore operation has successfully completed (or failed)
- The number of drives to use for the save or restore requests
- Who is allowed access to the environment (for data security)
- Default data safety checks to perform during save or restore operations (such as Full, XOR Redundancy, CRC, or a combination thereof)
- Whether to enable log and listing files.
- How often to retry the save or restore operation before requiring user intervention
- Whether to perform job-wide pre- or post-processing commands
- UNIX compression, file system span, and symbolic link options
- The resulting disposition of the files that are saved
- Locking options

2.2.5 Selections

When you specify a set of disk or file specifications for a save or restore request, you are creating (implicitly or explicitly) a selection object. A selection object contains one or more disk or file specifications, together with additional selection criteria and operational attributes including the following:

- Options to pass to the Backup Agent (agent qualifiers)
- The type of data to be saved (VMS files, Rdb databases and storage areas, UNIX files or Windows files)
- Selection criteria using a combination of before dates and since dates (explicit selection only)
- Specific files to exclude that would otherwise be included in the file specification
- Who is allowed access to the selection (for data security)
- Conflict options (what to do if the file being restored exists)
- For UNIX files and Windows files, the source node on which these files reside

If you specify a set of disk or file specifications as part of the save or restore request, these files are stored in a default selection for that save or restore. You can use a default selection exclusively in your saves and restores as long as the other selection criteria (including data type) are the same for all files in the request. Alternatively, you can create your own selections explicitly using either the MDMSView GUI or the CLI, and associate them with your save and restore requests. Each save and restore can support multiple selections.

2.2.6 Schedules

You can use a variety of ways of scheduling your save and restore requests, including two methods provided by MDMS, or by the use of a third-party scheduler product (see Section 2.7). The schedule object defines on what days and times a save or restore request is run. If you use MDMS scheduling, these schedule objects are executed at the appropriate times and the associated save and restore requests are invoked. If you use a third-party scheduler, the schedule objects are still created, but they do not invoke the associated save or restore requests - that is done by the third-party scheduler. The schedule object is created when you create the associated save or restore request.

For most save and restore requests, you can define a *frequency* of operation, which together with a base date determine the schedule attributes automatically. However, if you use internal MDMS scheduling, you can specify a *custom* schedule, and set attributes for scheduling including the following

- The days of the week you wish a request to run
- The dates of the month you wish a request to run
- The months of the year you wish a request to run
- The times of the day you wish a request to run - a request can run up to 100 times per day
- Specific dates in the next 10 years you wish a request to run, that otherwise would not be run according to the other selection criteria
- Specific dates in the next 10 years you wish the request *not* to run, that otherwise would be run according to the other selection criteria.
- Relate one schedule to another, so that its associated save or restore request runs after the related save or restore request.

If you use a third-party scheduler, you can specify non-standard frequencies by using an *explicit* frequency and interval that is passed to the scheduler, or you can use the scheduler interface directly to manipulate the frequency of the request.

2.3 ABS Catalogs

An ABS catalog consists of a catalog object and the catalog files. The information contained in an ABS catalog object includes:

- The type of catalog (FILES, DISKS, VOLUME_SETS)
- Whether or not to use an intermediate staging file
- Who is allowed to access the catalog (for data safety)
- Who is allowed write data to and read data from the catalog (ensures data safety)

The ABS catalog files contain history information about save requests and can be assigned to one or more archives. Each time a save request is initiated through a particular archive, the save request history is recorded in an ABS catalog associated with the archive.

The information contained in an ABS catalog includes:

- The name of the data that was saved
- The type of data that was saved (OpenVMS Files, Oracle Rdb Database, Oracle Rdb Storage Area, UNIX Files, Microsoft Windows Files, Oracle Database)
- The date and time the data was saved

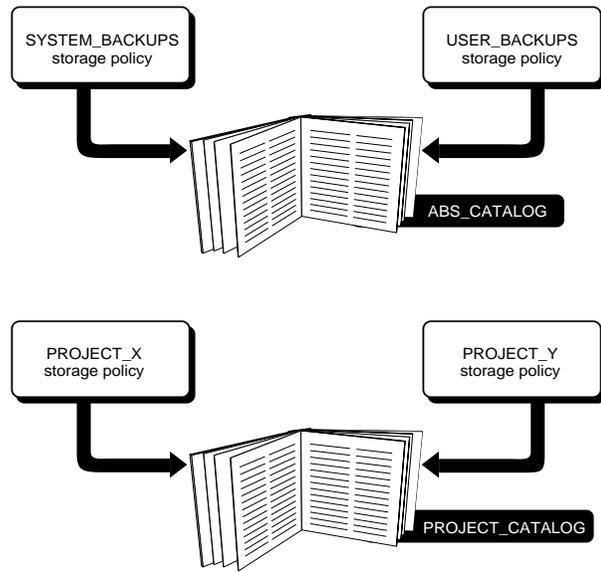
Overview

2.4 Backup Agent

- The save set name where the data is located
- The location of the save set (disk or tape)
- The original location of the data
- The owner of the data

Figure 2–2 shows the relationship between an ABS catalog and an ABS archive.

Figure 2–2 ABS Catalogs



CX05022C

After the installation of ABS is complete, ABS provides a default catalog named `ABS_CATALOG`. By default, this catalog is associated with all archives unless it is changed by the creator of the archive. All ABS catalogs, both the default catalog and user-created catalogs, support lookup and restore capabilities.

ABS catalogs are node specific but in a VMScluster all nodes could share the same catalog files.

2.4 Backup Agent

ABS uses various backup agents to save and restore data. The backup agent is determined by the type of data, such as VMS files, Oracle Rdb databases, Oracle Rdb storage areas, UNIX files, or Windows files. The backup agent is responsible for the actual data movement operation, while ABS is responsible for invoking the correct backup agent and recording the information about the save operation.

ABS supports the following backup agents:

- OpenVMS BACKUP Utility - For OpenVMS files, ABS uses the OpenVMS BACKUP Utility.
- RMU Backup Utility - For Oracle Rdb databases and storage areas, ABS uses the RMU Backup Utility.
- `gtar` (GNU tar) - For UNIX and Microsoft Windows files, ABS uses `gtar` (aka tape archiver or tar).

2.5 Media, Device and Management Services (MDMS)

Media, Device and Management Services (MDMS), a fully-integrated component of ABS, performs several services for ABS, including:

- Database Services - The ABS objects are managed by MDMS databases and are compatible with the MDMS media management databases
- Interfaces - Both the MDMSView GUI and the CLI to all objects are managed by MDMS. The old ABS DCL interface is obsolete, but still supported. The old ABS and MDMS GUIs are not supported.
- Security Services - MDMS manages access rights and privileges to ABS and MDMS objects, including individual access control on all objects. Security is discussed in Chapter 5.
- Media Management Services - MDMS supports a set of objects for the purpose of media management for ABS. Media management services are described in Chapter 4.

2.6 User Interfaces

The interfaces for ABS are provided by MDMS, which performs all database management on behalf of ABS. MDMS provides the following interfaces.

Interface	Description
MDMSView GUI	MDMS provides a graphical user interface (GUI) called MDMSView that allows manipulation of all ABS and MDMS objects in an integrated GUI. MDMSView provides several “views” of accessing ABS and MDMS information, and is usable on OpenVMS Alpha systems (V7.2-1 and later) and any Windows system. See Chapter 6 for a description of MDMSView.
MDMS CLI (DCL)	MDMS also provides a Command Line Interface (CLI), which is the Digital Command Line (DCL) interface, for users who prefer this type of interface, or for users whose OpenVMS systems cannot support the MDMSview GUI. See Chapter 6 for a brief description of the CLI interface. This interface is also described in its entirety in the <i>MDMS Reference Guide</i> .

ABS provided its own CLI interface in versions prior to version 4. This interface is now deprecated, but is still provided for backward compatibility. The former ABS GUI, however, is not supported.

2.7 Scheduler Options

MDMS allows the use of different scheduler interfaces. By default MDMS uses an internal interface to the OpenVMS Queue Manager to schedule save and restore requests. MDMS supports the following scheduler interfaces:

- INTERNAL (default) - uses an internal interface to OpenVMS Queue Manager
- EXTERNAL - uses DCL commands to interface with the OpenVMS Queue Manager by calling a command procedure
- SCHEDULER - uses DCL commands to interface with the 3rd party scheduler product by calling a command procedure; the pre-V3.0 ABS scheduler DECscheduler V2.1B may be used with this option if you have a license for that product.

Note

The internal queue manager scheduler interface is the only scheduler interface available with the ABS-OMT license.

The scheduler interface is invoked when a save or restore request is created, you can either start the request immediately or define a repetitive schedule.

The scheduler interface is used to:

- Automate and manage ABS jobs that run repeatedly, such as ABS save and even restore requests.
- Capture events through a logging system, so you can generate accounting and historical reports. This may vary depending on the scheduler interface.
- Execute all requests remotely as well as locally - transparently to the user.

2.8 MDMS Objects

This section summarizes the MDMS objects for media management. See Chapter 4, *Media Management*, for more detailed information on MDMS objects.

2.8.1 Domain

The MDMS domain encompasses all objects that are served by a single MDMS database. These include physical resources such as nodes, jukeboxes, drives and volumes, and logical objects such as media types, pools and magazines. The domain also encompasses all the users that access and manage MDMS resources. A domain may encompass a single site location, or can be geographically distributed, linked via Fibre Channel or a wide area network.

The MDMS domain has a single domain object, which contains:

- The default media type, onsite and offsite locations, protections and dates that are assigned to new volumes by default
- The default OPCOM classes assigned to new nodes by default
- The type of scheduler to be used in the domain
- The system users to be notified when volumes are deallocated
- The request ID of the next MDMS request
- The mapping of low-level rights to high-level rights
- The level of access control to be assigned to the domain.

2.8.2 Drives

A drive is a physical resource that can read and write data to tape volumes. Drives may be of one of three types:

- Jukebox - The drive is part of a robot-controlled jukebox, and random-access loading and unloading is performed by the robot
- Stacker - The drive supports the automatic loading of a succession of volumes in sequential access. Once the volumes are exhausted, operator intervention is needed to load new volumes
- Standalone - The drive requires operator intervention for all loads and unloads.

Jukebox drives are associated with a jukebox, and require a drive number identifier if the jukebox is controlled by MRD. Stacker and standalone drives are not associated with a jukebox: this includes drives used in a stacker configuration that are actually in a physical loader.

MDMS supports a drive object for each drive to be managed by MDMS. The drive object includes:

- The OpenVMS device name of the drive (this can be the same or different than the drive name).
- The media types that the drives supports for both read-write and read-only operations
- The nodes and groups with direct access to the drive, including Fibre Channel access
- Flags associated with the drive
- The state of the drive
- Local and/or remote access to the drive
- The jukebox associated with the drive

2.8.3 Groups

The MDMS group object is simply a collection of nodes that have some common association. You may define groups to represent OpenVMS clusters, a set of nodes that can access Fiber Channel devices, or for any purpose whatsoever. Groups can typically be used in all commands that support nodes. It is a convenient way to reference a long list of nodes. In commands that support nodes and groups, it is possible to specify both for the command.

The only attribute that a group has is a list of nodes.

2.8.4 Jukeboxes

In MDMS, a jukebox is a generic term applied to any robot-controlled device that supports automatic loading of volumes into drives. MDMS jukeboxes include:

- Small, single-drive loaders such as the TZ887 or the TLZ9L
- Large, multi-drive libraries with ports, slots and capabilities typically ranging from the tens to the hundreds of volumes, such as the ESL9326
- Very large StorageTek (R) silos that may contains literally thousands of volumes and many tens of drives

A jukebox object is associated with each jukebox, and contains the following fields:

- Control option - controlled by the SCSI-based MRD subsystem, or DCSC for certain silos
- For MRD jukeboxes:
 - The OpenVMS robot name for the jukebox
 - The number of slots in the jukebox
 - The magazine option flag, and optional magazine topology
- For DCSC jukeboxes:
 - The library, ACS and LSM identifiers for the jukebox
 - The CAP sizes for the jukebox
- The location of the jukeboxes
- Access options for local and/or remote access to the jukebox

Overview

2.8 MDMS Objects

- The threshold value for free volumes in the jukebox (before a warning is issued)
- The groups and nodes that have direct access to the jukebox, including access via Fibre Channel
- The state of the jukebox

2.8.5 Locations

A location describes the physical location of other objects, and is used as a selection criterion for allocating drives and volumes, and for placing tape volumes in a specific place. Locations can exist in a hierarchy, and as such are considered compatible locations for allocation purposes if locations share a common root in the hierarchy.

Locations only have two attributes:

- Parent location - The parent location in the hierarchy (a location need not have a parent location)
- Spaces - A range of “spaces” to be used for storing volumes, also optional.

2.8.6 Magazines

A magazine is a logical object that contains a set of volumes that are to be moved as a group. Magazines typically relate to a physical magazine that certain jukeboxes require in order to move volumes in and out of a jukebox (for example, a TZ877 or TLZ9L). However, even for jukeboxes requiring physical magazines, it is not a requirement to configure MDMS magazines if you want to treat the movement of the individual volumes independently.

Magazines contain the following attributes:

- Slot count
- Placement
- Jukebox name, start slot or position
- Onsite and offsite locations and dates

When a volume is in a magazine, its placement and associated locations are those of the magazine. Magazines can be scheduled to move onsite and offsite. In most cases, this means that all the volumes in the magazine are moved onsite or offsite; the physical magazine itself usually stays with the jukebox with a new set of volumes.

The use of magazines is not required.

2.8.7 Media Types

A media type is a logical object that describes certain attributes of tape volume media. Media types are used as a major selection criterion for drive and volume allocation, and are used to match volumes with compatible drives. Media types contain the following attributes:

- Density - A density value or keyword that identifies the density of the media. This value must be one of the keyword values supported by OpenVMS. Density is used in initializing volumes.
- Compaction - A flag indicating whether compaction is desired on volumes. Setting compaction usually results in about twice as much data capacity for a tape volume.
- Capacity - The size of the media in MB (not used by MDMS).
- Length - The length of the media in feet (not used by MDMS).

2.8.8 Nodes

A node is an OpenVMS system in the MDMS domain that is running MDMS. Every node in the domain must have a node definition, which describes the network transports and other information applicable to that system. Node attributes include:

- Location of the node
- OPCOM classes to be used for OPCOM messages on the node
- Supported network transports and transport full names

2.8.9 Pools

A pool is a logical object that contains a set of volumes that can be allocated and used by a set of authorized users. It is one way to separate volumes belonging to different organizations and allowing only users of those organizations to use the volumes. Pool attributes include:

- Authorized users - A list of users in node::username format that are authorized to allocate and use volumes in the pool
- Default users - A list of users in node::username format that are not only authorized to use volumes, but that use volumes from this pool by default.
- Threshold - A minimum value of free volumes in the pool, below which an OPCOM warning message is sent.

A user need only be defined in one of the lists to be able to use volumes in the pool.

The use of pools is not required.

2.8.10 Volumes

A volume is a single piece of tape media that MDMS applications (ABS and HSM) use to store tape-related data. Volumes contains many attributes that are used to describe the type of volume, its placement and location, and dates for scheduling allocation and movement. Volume attributes include:

- Media type and pool for the volume
- Placement and placement objects such as jukebox, slot, location, magazine
- Onsite and offsite locations and scheduled dates
- Allocation state, user and scheduled scratch date
- Formatting information
- Volume protection
- Counters
- Historical information dates

2.9 Getting Started

This section provides a simple example of how to configure a minimal ABS/MDMS domain and create a save and restore request. Although most configurations are more complex than this, it serves to illustrate how to use the MDMS configuration procedure and the default objects provided by ABS.

Before creating save or restore requests, you should first configure the media management environment. This includes the tape volumes, drives, jukeboxes and other media management objects that you may want to use. The recommended way to do this is to run the MDMS configuration command procedure, which offers an online tutorial and help in defining the configura-

Overview

2.9 Getting Started

tion. During execution of this procedure, type “?” to get help on any question, and type “??” to get help and (in many cases) a list of existing objects or possible values for answers to questions. To invoke this procedure:

```
@MDMS$SYSTEM:MDMS$CONFIGURE
```

A complete example of running this procedure is provided in Appendix A.

Having completed the media management configuration, creating a save or restore request in ABS can be very simple if you elect to use the default archives, environments and selection objects. The minimum amount of information you need to specify for a save or restore request is:

- The name of the save or restore.
- The disks or files to be saved.
- The start time of the save.

ABS tries to determine the type of data being saved based on the format of the file specification and assigns by default a relevant archive and environment. So, for example, a save request can be specified and executed in a single DCL command as follows:

```
$ MDMS CREATE SAVE MY_SAVE/INCLUDE=DISK$USER1:[SMITH...]/START
```

This command creates a save called MY_SAVE, includes the file specification DISK\$USER1:[SMITH...] (all files), and starts the save immediately. MDMS determines that this is a save of VMS files based on the file format, and assigns archive SYSTEM_BACKUPS and environment SYSTEM_BACKUPS_ENV, and creates a default selection and schedule. With this save definition, a default frequency of ONE_TIME_ONLY is assigned, and the save is not scheduled for regular execution.

A restore can be defined just as easily. For example, to restore the same files that were saved in MY_SAVE, you can enter the following command:

```
$ MDMS CREATE RESTORE MY_RESTORE/INCLUDE=DISK$USER1:[SMITH...]/START
```

This command creates a restore called MY_RESTORE, includes the file specification DISK\$USER1:[SMITH...] (all files), and starts the restore immediately. MDMS determines that this is a restore of VMS files based on the file format, and assigns archive SYSTEM_BACKUPS and environment SYSTEM_BACKUPS_ENV, and creates a default selection and schedule. With this restore definition, a default frequency of ONE_TIME_ONLY is assigned, and the restore is not scheduled for regular execution.

Since these requests were defined with a frequency of ONE_TIME_ONLY, ABS will automatically delete them after a default interval of 3 days after execution.

Of course, creating the backup environment to backup all data in your production environment will involve more complex definitions, including creating your own archives, environments and in some cases selections and schedules. Chapter 3, *Saving and Restoring Data*, describes all the ABS objects in detail.

Saving and Restoring Data

This chapter expands upon the ABS Overview in Chapter 2 and describes saving and restoring in detail by discussing the ABS objects, and the meanings, possible values and uses for all attributes. For each object, the attributes are listed in alphabetical order for easy reference, but related attributes are discussed together. The attributes are described without specific syntax or instructions on how to manipulate them, but are named according to the qualifiers in the CLI and attributes on the MDMSView GUI screens. For information on the syntax and semantic rules for each object and attribute, refer to the *MDMS Reference Guide*.

All objects have an owner, and optional access control which limits access to the object. Since these attributes are common to all objects, they are described in Chapter 6, *Security*, instead of this chapter.

In addition, ABS supports inheriting attributes from one object to another when creating a new object. For example, if you want to create a new save request SAVE2, but use most of the attributes from another save request SAVE1, you can specify SAVE1 as the inherit attribute when you create SAVE2. From there you can modify SAVE2 to define its unique characteristics. This philosophy applies to all ABS and MDMS objects. You can even inherit restore requests from save requests if you want to restore the same files as were previously saved.

Finally, all objects have a *description* attribute in which you can enter a text string to describe the object. This attribute is not interpreted by either ABS or MDMS, so you can use it for any purpose you see fit. By default, the description is blank.

The following sections discuss all seven ABS objects in detail.

3.1 Archives

Archives define the media type and characteristics about where backup data is stored. Each save and restore uses exactly one archive, except that certain complex save and restores can use two archives (see Section 3.3). You can use a single archive for many different saves and restores by simply referencing the archive in the save and restore request. ABS defines four archives by default, which you can use in your save and restore requests as needed:

- **SYSTEM_BACKUPS** - For system backups that are normally performed by a system administrator at regularly scheduled times
- **USER_BACKUPS** - For backups performed by a non-privileged user to save or restore his or her own data
- **UNIX_BACKUPS** - For backups of UNIX client data, normally performed by a system administrator
- **DISASTER_RECOVERY** - For backups primarily designated for disaster recovery

Although these default archives are provided by ABS, you may modify them as needed to suit your site's operational environment. Alternatively, you can create your own archives and manipulate the attributes as described in the following sections.

Saving and Restoring Data

3.1 Archives

3.1.1 Archive Name

This name is used to reference the archive in save and restore requests. There are no required or ad-hoc conventions for archive names, so they can reflect their usage in your environment. However, there are ad-hoc conventions for environment names based on the archive name, so you should restrict the archive name to 60 characters.

3.1.2 Archive Type

ABS supports two types of archive, which are hopefully self-explanatory:

- **DISK** - The archive data is stored on disk media, which can include optical disk. ABS assumes that all disk media are online and mounted on the OpenVMS system before any save or restore operation is executed. ABS does not perform any load/unload or mount operations on disk archives. When you specify disk archive type, the archive must contain a *destination* attribute indicating the disk and directory location of the archive data.
- **TAPE** - The archive data is stored on tape media, and uses MDMS for media management control of the media. When you specify tape archive type, the archive must contain a media type (defined in MDMS) that defines the type of tape media to be used for the archive. Only a single media type is supported. In addition, the archive may optionally contain a pool specification (indicating a set of volumes reserved to users authorized for the pool) and a location specification (used to allocate a drive).

3.1.3 Catalog

A catalog contains information about what data is stored in the archive and where it is stored. Each archive uses exactly one catalog, although catalogs can be shared among different archives. ABS defines a default catalog called ABS_CATALOG, which is assigned to all archives by default if a different catalog is not specified. If you wish to define a different catalog for an archive, then specify a catalog object name (not its location) in the catalog attribute of the archive. For the archive to be useful, the catalog must be defined as a catalog object in MDMS.

An archive with a name of "DISASTER_RECOVERY" is the only archive allowed to have no catalog associated with it and the save operation is therefore not catalogued (see Chapter 7, *Disaster Recovery*).

3.1.4 Consolidation

ABS supports the concept of consolidation criteria which determine when a volume set should be retired from use in the archive and a new volume set used. ABS supports three types of consolidation criteria, of which none, one, two or all three can be applicable:

- **INTERVAL** - You can specify an interval as a delta time from the creation of the current volume set to the creation of the next volume set. The current volume set is retired if the consolidation interval is exceeded.
- **SAVESETS** - You can specify the maximum number of savesets that should reside on the volume set. If this number would be exceeded, ABS retires the current volume set and allocates a new volume set for the archive. There is an ANSI-imposed maximum of 10000 savesets in a volume set
- **VOLUMES** - A volume set can contain one or more physical tape volumes. You can limit the number of volumes by specifying volumes on the consolidation criteria. If this number would be exceeded, ABS retires the volume set and allocates a new volume set. There is an ANSI-imposed maximum limit of 100 volumes in a volume set.

If you specify multiple consolidation criteria, ABS creates a new volume set when the first of any of the defined criteria are exceeded. The default consolidation criteria is an INTERVAL of 7 days. If no consolidation criteria are specified, then ABS creates a new volume set when the ANSI limits apply, or upon the first error writing to the volume set. This is not recommended as you may create excessively large volume sets, and may have to split a volume set between onsite and offsite (vault) locations. Consolidation criteria are only applicable to an archive type of TAPE.

3.1.5 Destination

If you specified an archive type of DISK, you must enter a destination attribute for the archive, or use the default of ABSS\$ROOT:[000000]. The destination contains the disk and directory location of the data saved in this disk archive. When specifying destination, you should ensure that the specified disk has enough free capacity to handle all data to be saved in this archive. ABS does not monitor the disk for sufficient capacity. ABS clears this attribute if the archive type is TAPE.

3.1.6 Drives

ABS allows you to enter a list of drives that can be used by save and restore operations to and from this archive. This should be used only to restrict the drives that would normally be available for these operations for some reason. Normally, you can let ABS select drives for all operations based on media type and location, and so you do not need to specify the drives in the archive. If you do specify drives, be aware that these drives apply to restores as well as saves. Drives are only applicable to an archive type of TAPE.

3.1.7 Expiration Date and Retention Days

ABS supports two alternative methods of specifying when an archive expires. These are:

- Expiration Date - A date given in OpenVMS absolute time that defines a specific future date that the volume data will expire.
- Retention Days - The number of days following retirement of the volume set that the data will be retained, after which time it will expire.

Either retention days or expiration date may be given, but not both. By default, ABS defines retention days of 365, meaning that volume data is valid for one year after retirement of the volume set.

For an archive type of TAPE it defines the initial scratch date of the tape volume set. Once a volume has transitioned to FREE state and it has been re-used all catalog entries relating to the past usage of this volume are deleted. You can change the expiration of the archive by setting a new scratch date for the volume. Whenever data is added to the volume set a new scratch date will be set if the expiration date extends beyond the old scratch date.

For an archive type of DISK it defines the time at which the on-disk saveset is deleted. At the same time all catalog entries relating to that saveset are also deleted.

Expiration date and retention days are only applicable to an archive type of TAPE.

ABS supports save requests that sometimes perform full backups and sometimes perform incremental backups. Under these circumstances, it is useful to use different volume sets with different retention days or expiration dates for the fulls and the incrementals. To support this, ABS allows you to specify two archives for save requests: the first applies to the full backups, and the second applies to the incremental backups.

Saving and Restoring Data

3.2 Catalogs

3.1.8 Location

A location is an MDMS object that defines the physical location of volumes, drives or juke-boxes. The location is used as one of the selection criteria (along with media type) for allocating a drive to load a scratch volume to extend the archive. If no location is specified for the archive, ABS uses the default onsite location defined in the MDMS domain. This is the default. Location is only applicable to an archive type of TAPE.

3.1.9 Maximum Saves

ABS supports multiple parallel save operations using a single archive, each operating on a different drive and volume set (archive type TAPE). To enable this feature, specify the maximum number of parallel saves that are desired using the maximum saves attribute. Values can range between 1 and 36, with 1 being the default. This attribute also applies to an archive type of DISK, but without the implications of multiple drives and volume sets being allocated.

3.1.10 Media Type

Media type is an MDMS object that describes the type of tape media to be used in the archive. Specify a media type that is defined within MDMS, or use the default domain media type. The media type is used as a mandatory selection criterion (along with optional pool and location) for volumes to be used in the archive. Media type is only applicable to an archive type of TAPE.

3.1.11 Pool

A pool is a logical MDMS object that relates a collection of volumes to a set of authorized users. In this way, you can allocate a collection of volumes to certain users knowing that other users cannot use volumes from the pool. Similarly, you can assign a pool to an archive, so that all volumes used in the archive must be taken from the volumes that are in the pool. You can specify only one pool per archive. If you do not specify a pool, then only volumes that have no pool defined can be used for the archive (this is also known as the *scratch pool*).

3.1.12 Volume Sets

The volume sets attribute indicates which volume sets are currently being used by save requests using the archive. There may be up to the *maximum saves* number of volume sets currently being used. These volume sets are those to which the next save operation will be written to the archive. This attribute is normally maintained by ABS and you should not modify it unless there is a pressing need to remove one or more of the volume sets from the list and let ABS allocate new volume sets. Under no circumstances should you add volumes to the volume set list.

3.2 Catalogs

An ABS catalog contains historical information about ABS save operations. This historical information includes the location of data that was saved using ABS. For this purpose, ABS provides a default catalog named ABS_CATALOG.

Some sites can operate efficiently using only ABS_CATALOG provided by ABS. However, using additional catalogs can improve your ABS operations:

- Speed of record insertion
- Speed of lookup operations
- Segregation of saved data
- Regular catalog file maintenance

3.2.1 Catalog Name

This name is used to reference the catalog. There are no required or ad-hoc conventions for catalog names, so they can reflect their usage in your environment.

3.2.2 Catalog Node

Catalogs are node specific. You have to specify the MDMS node name where the catalog resides. An empty catalog node name means the local node where the command was issued or where the request executes. In a VMScluster, multiple nodes can share the same catalogs on a common disk as long as they have direct access to the catalog files. During a save operation ABS always accesses the catalog on the local node even though a different node name is specified in the archive. During a restore operation the catalog lookup will be performed for the catalog at the specified node. This allows to restore data that has been saved on another node.

3.2.3 Type

ABS supports four types of catalogs, which are hopefully self-explanatory:

- **DISKS** - This catalog type only stores information about save requests performed. No information about individual filenames are stored in the catalog. The size of a DISKS type catalog is drastically smaller than the FILES catalog type but you cannot restore individual files. Save requests using this catalog type must be of type FULL and only specify a disk name. Staging does not apply to these catalogs.

NOTE: You can still use the backup agent outside of ABS to do a selective restore.

Example 3–1 Restoring Individual OpenVMS Files with DISKS Type Catalog

```
$ BACKUP MKA500:01MAY20001234567. /SELECT=[MyDir]MyFile.Dat *
```

restores the file [MyDir]MyFile.Dat from save set 01MAY20001234567 to the current directory.

To view information about saved disks use the /SAVE qualifier with the “SHOW CATALOG” command. The show output lists the data saved and the volume ID and save set name used.

When a save request uses a DISKS type catalog, the following message is displayed in the save request log file:

```
"Filenames will not be catalogued"
```

- **FILES** - The FILES type catalog stores all information about save requests performed and all files saved. It allows individual file lookups and restores.
- **SLS** - The SLS type catalog allows you to create ABS catalogs solely for the purpose of restoring data saved using SLS. These catalogs are only used for restore operations and not for save operations.
- **VOLUME_SETS** - The VOLUME_SETS type catalog stores all information like the FILES type catalog. However, ABS uses individual files for each volume set. Catalog lookups take slightly longer for VOLUME_SETS type catalogs compared to FILES type catalogs. But VOLUME_SETS type catalogs avoid the constant growth of catalog files because the volume set specific file is deleted once the volume set has been re-used. A VOLUME_SETS type catalog cannot be used for DISK archive types.

Saving and Restoring Data

3.2 Catalogs

3.2.4 Directory

By default catalog files are created in ABS\$CATALOG. Without modification ABS\$CATALOG points to ABS\$ROOT:[CATALOG]. When creating a new catalog you can create the catalog files in a different location by specifying a device and directory. You have to update the definition of ABS\$CATALOG logical in ABS\$SYSTARTUP.COM to include the new device and directory solution in form of a search list.

Example 3–2 Adding a New Catalog Location

```
$ CREATE/DIRECTORY DKA100:[ABS_CATALOG]
$ DEFINE/SYSTEM/EXECUTIVE ABS$CATALOG ABS$ROOT:[CATALOG],-
DKA100:[ABS_CATALOGS]
```

This creates a new directory for catalog files and adds it to the ABS\$CATALOG search list. The same definition needs to be set in ABS\$SYSTARTUP.COM.

If a “SHOW CATALOG/FULL” does not display a directory for a catalog it means the catalog’s location is not included in the ABS\$CATALOG search list.

3.2.5 Staging

A catalog that is setup for staging improves the performance of the save operation because the catalog entry for a saved file is first written to a sequential disk file in ABS\$CATALOG. Once the backup operation has completed a separate process moves the entries from the staging catalog file to the final catalog which is specified in the archive associated with the save request.

The final catalog does not contain the information about the save operation until the staging process has completed. If you request a backup operation and immediately look in the final catalog, the entries may not be available, yet. The backup operation and the staging process must complete before the currently saved files can be looked up in the catalog.

You can always modify the staging setting for an existing catalog. The use of staging is highly recommended to improve your overall backup times.

The staging catalog file is created in the first location pointed to by logical name ABS\$CATALOG.

3.2.6 Catalog Save Entries

Save entries contain information about executing or executed save operations:

- Catalog Name - The name of the catalog
- Catalog Node - The name of the MDMS node where the catalog resides
- Date Archived - The date the save operation was performed
- Expiration Date - The original date the entry expires in the catalog (used only for archive type of DISK)
- Source Node - The network name of the node where the saved data was located (UNIX and Windows Files only)
- Include - The include file specification used
- Object Entries - Number of entries added to the catalog
- Archive - The name of the archive or, if the original archive no longer exists the previous archive UID
- Environment - The name of the environment or, if the original environment no longer exists the previous environment UID

- Save - The name of the save or, if the original save no longer exists the previous environment UID
- Save Type - Shows the type of save being performed
 - all files with recording (R) - All files in a full incremental save with final recording of the backup date
 - all files (B) - All files in an incremental selective save
 - all files (S) - All files in a selective save
 - all files (0) - All files in an incremental save
 - increment level n - All files modified between incremental save n and n-1
- Owner - The owner field of the archive being used
- Saveset Format - The format used in the saveset:
 - GTAR - UNIX gtar format
 - NT_GTAR - Windows gtar format
 - RMU_BACKUP - Oracle Rdb/RMU saveset format
 - VMS_BACKUP - OpenVMS BACKUP saveset format
- Archive Type- DISK or TAPE
- Saveset Location -
 - For archive type TAPE the list of volume IDs containing the saveset
 - For archive type DISK the on-disk location of the saveset
- Saveset Name - The filename of the saveset
- Saveset Position - The tape mark offset of the beginning of the saveset on tape
- Status - The ABS status for the save operation
- Severity - The ABS severity level for the save operation

3.2.7 Catalog File Entries

File entries contain information about files which have been saved.

- Catalog Name - The name of the catalog
- Catalog Node - The name of the MDMS node where the catalog resides
- Data Select Type - The format of the entry name
 - RDB_[Vnm_]_DATABASE - An Oracle Rdb database file
 - RDB_[Vnm_]_STORAGE_AREA - An Oracle Rdb storage area
 - UNIX_FILES - UNIX file specification
 - VMS_FILES - OpenVMS file specification
 - VMS_SAVESET - volumeID:saveset specification
 - WINDOWS_FILES - Windows files specification
- Filename - The name of the entry
- Source Node - The network nodename where the entry was located
- Date Archived - The date the entry was saved

Saving and Restoring Data

3.2 Catalogs

- Expiration Date - The original date the entry expires in the catalog (used only for archive type of DISK)
- Creation Date - The date the entry was created on the source node
- Revision Date - The date the entry was last modified on the source node before being saved
- Owner - Owner information of the entry used on the source node
- Saveset Name - Copied from related save entry
- Saveset Location - Copied from related save entry
- Saveset Section -
 - For archive type of TAPE the index into the list of volume IDs indicating the volume which contains the start of the saved entry
 - For archive type of DISK it is always 1
- Save Type - Copied from related save entry
- Status - Copied from related save entry
- Severity - Copied from related save entry

3.2.8 Improving Catalog Performance

Catalog files are RMS index-sequential files and as such need regular maintenance to avoid unnecessary file growth and performance penalties. ABS provides a catalog conversion command procedure (“ABS\$SYSTEM:ABS\$CONVERT_CATALOG.COM”) that improves the target catalog update performance by doing a file-to-file conversion. By converting the target catalogs, you improve catalog update time.

3.2.8.1 Catalog File Sizes

The ABS catalog files will grow as you continue to execute save requests. The sizes depend on the number of files saved and the retention period used. For as long as the retention period has not expired more entries will be added to the catalog. Once the retention period is reached the ABS_CLEAN_CATLG_<node_name> batch job will remove expired entries from the catalog. So the more files you save and the longer you want to keep the archived data the larger the catalog files will become.

Be sure to consider this information when creating catalogs and assigning retention values to your archives. It may be best to create separate catalogs for each archive, if the retention period is different. For example, you may create an archive called MONTHLY_SAVE, with a retention period of one month. Create a catalog called MONTHLY_SAVE to be used by that archive. The catalog size will grow for one month and then maintains its size.

3.2.8.2 Catalog File Maintenance

Run the conversion command procedure for each individual catalog on a regular basis. Catalogs with more frequent delete operations should be converted on a monthly basis. See the logfiles ABS\$SYSTEM:ABS\$CATALOG_CLEANUP.LOG;* for information on catalog file activities. As a rule of thumb the catalog should be converted if more than 10% of its records have been deleted.

Example 3–3 Converting Catalog Files

```
$ @ABS$SYSTEM:ABS$CONVERT_CATALOG MyCatalog
```

This example converts all files for catalog “MyCatalog” by creating new copies of the files in the same directory.

For additional improvement you can also move the target catalogs to a different disk by defining a system level search list logical for ABS\$CATALOG in ABS\$SYSTARTUP.COM. The command procedure also allows you to move the converted files to a different disk or directory.

Example 3–4 Moving Catalog Files to New Location

```
$ @ABS$SYSTEM:ABS$CONVERT_CATALOG MyCatalog DKA100:[ABS_CATALOGS]
```

This example converts the files for catalog “MyCatalog” and places the new files in location “DKA100:[ABS_CATALOGS]”. Once the files have been copied over you can add the new location to the ABS\$CATALOG search list. Rename the old catalog files to *.DAT_OLD and verify that you can lookup information using the new files. Once the new catalog files are used you can delete the old files.

3.2.8.3 Staging Catalog

With staging enabled for a catalog ABS writes the catalog entries to a sequential file during a save operation. The save operation at the end creates a command procedure and executes it in a separate process. This unpack process moves all entries from the staging catalog to the final catalog. If all entries have been moved successfully the command procedure is deleted. If the unpack process failed for some reason you can run the command procedure manually. To do this, find the location and name of the command procedure in the logfile of the save request. Then execute the command procedure on the node where the save request was running.

Example 3–5 Staging Information in Save Log

```
21:21:07      COORD: Staging process PID : 2300143C
21:21:07      COORD: Staging catalog   : ABS$CATALOG:ABS_CATALOG_4.STG;1
21:21:07      COORD: Staging procedure  : ABS$CATALOG:ABS_CATALOG_4_1.COM;1
21:21:07      COORD: Staging logfile   : ABS$LOG:ABS_CATALOG_4.LOG
```

In this example if the command procedure file “ABS\$CATALOG:ABS_CATALOG_4_1.COM” still exists it indicates that the staging unpack process has failed and you can manually execute the command procedure to update the catalog.

Staging files by default are created in the first location pointed to by logical name MDMS\$CATALOG.

3.3 Cataloging Existing Savesets

You may catalog information from existing VMS Backup savesets on tape. This allows you to lookup and restore files from savesets created outside of ABS.

Restrictions:

- The saveset must reside on tape
- Only VMS Backup savesets may be cataloged
- The tape volume must be defined in MDMS and allocated to ABS so that ABS may reference the volume
- A separate catalog and archive should be created for the saveset information

Saving and Restoring Data

3.3 Cataloging Existing Savesets

To catalog the savesets, create a SAVE request with the name of the tape volume and the saveset name, or wildcard, separated by a colon as the selection (include) and a data_type of VMS_SAVESET:

```
MDMS CREATE SAVE mysaveset_catalog/INCLUDE=tap001:mysaveset.sav/DATA_TYPE =  
VMS_SAVESET/ARCHIVE=my_archive/ENVIRONMENT=my_env/START=01-MAY-2002
```

or

```
MDMS CREATE SAVE mysaveset_catalog -  
/INCLUDE=tape001:*/DATA_TYPE=VMS_SAVESET/ARCHIVE=my_archive/ENVIRON-  
MENT=my_env/START=01-MAY-2002
```

ABS will load the tape listed in the include specification, then do a BACKUP/LIST of the contents, loading the information into the ABS catalog defined in the archive. The original date of the saveset will be preserved in the catalog.

Recommended Implementation:

It is recommended that you create a new catalog to store this data. You should also create a new archive to be used by these cataloging operations. This is mainly if you are cataloging copied tapes, where the dates of the original and the copied savesets will be duplicates.

This will allow you to choose to restore from the original or copies by selecting the appropriate archive for the restore request.

For example:

Several ABS save requests were saved on tape ABS000 using the SYSTEM_BACKUPS archive. Saveset Manager (SSM) was used to copy that tape to another tape, TAP000.

Before cataloging the data, do the following:

Create a new catalog called COPIED_TAPES. Create an archive called COPIED_ARCH, which points to the catalog COPIED_TAPES.

Create a save request specifying TAP000:* for the include specification and give it a data_type of VMS_SAVESET.

ABS will execute the request, cataloging the information into the COPIED_TAPES catalog.

To restore the data which is on ABS000 or TAP000, decide which copy you wish to restore and specify the appropriate archive or catalog on the restore request. For example, to restore from the original tapes, specify the SYSTEM_BACKUPS archive. To restore from the copy, specify the COPIED_ARCH archive. The MDMS SHOW CATALOG/FILES command with the /FULL qualifier will show the volumes used for the data.

Note

If the information about the original and copied savesets is put into the same catalog, they will have exactly the same archived data. This could cause confusion when restoring the data because ABS may not choose the tapes you wish to use for the restore. To make it easier to restore, it is recommended to use a separate catalog (as described above).

3.4 Environments

An environment describes the criteria under which save and restore requests execute, and exactly one environment must be associated with each save and restore request. You can use a single environment for many different saves and restores by simply referencing the environment in the save and restore request. ABS defines five environments by default, which you can use in your save and restore requests as needed:

- `SYSTEM_BACKUPS_ENV` - For system backups that are normally performed by a system administrator at regularly scheduled times
- `USER_BACKUPS_ENV` - For backups performed by a non-privileged user to save or restore his or her own data
- `UNIX_BACKUPS_ENV` - For backups of UNIX client data, normally performed by a system administrator
- `DISASTER_RECOVERY_ENV` - For backups primarily designated for disaster recovery
- `DEFAULT_ENV` - Used by default in the event one of the other default environments have been deleted

Although these default environments are provided by ABS, you may modify them as needed to suit your site's operational needs. Alternatively, you can create your own environments and manipulate the attributes as described in the following sections.

3.4.1 Environment Name

This name is used to reference the environment in save and restore requests. There are no required conventions for environment names, but ABS uses an ad-hoc convention that pairs environments and archives. If you specify an archive of name FOO, then by convention there should be a matching environment named FOO_ENV. You can choose to follow or ignore this convention for your site.

3.4.2 Action

The action attribute specifies one of three possible actions to be performed on files saved using this environment. Specify one of the following three actions:

- `RECORD_DATE` - Modify the BACKUP date to reflect the time that this file was backed up; this is the required option if you intend to do incremental backups of this file, and is the default value - supported for files of type VMS_FILES only.
- `NO_CHANGE` - Do not change the online file at all. If this option is specified, you will not be able to perform incremental saves on this file.
- `DELETE_FILE` - This option is used when the backup is intended to be a long-term archive and you wish the file to be removed from the online system. The file is only deleted on a successful save operation.

Although `RECORD_DATE` is supported for VMS_FILES only, it remains the default for all data types, and is simply ignored for the other types.

3.4.3 Compression

ABS supports the following types of compression for UNIX clients:

- No compression (default)
- UNIX Compression
- GZIP Compression

Saving and Restoring Data

3.4 Environments

It is recommended that you either use the default UNIX environment (`UNIX_BACKUPS_ENV`) or a single user-created environment for all your UNIX client saves, using a single type of compression for all UNIX saves. If you mix compression types among your UNIX saves, you should be very careful to assign the appropriate environment on any restore. If you specify the wrong compression option on a restore, then ABS will not be able to restore the data. The default is no compression.

3.4.4 Data Safety

The environment object allows you to specify one or more data safety features to ensure the reliability of the data on your offline tape volumes. You can select one or more of the following options:

CRC - Performs a Cyclic Redundancy check and writes it for each data block on a tape volume. This enables detection of a bad block during a restore operation.

FULL_VERIFY - Rereads all saved data and compares to what is on disk during a save. This option approximately doubles the time for the data copy phase of a save operation.

XOR - If the CRC check detects a bad block during a restore operation, the XOR mechanism allows recovery of the block. This option is applicable only to data type `VMS_FILES`, for which the backup agent is `VMS BACKUP`.

By default, all three options are enabled for maximum data safety.

3.4.5 Drive Count

The drive count specifies the number of tape drives to use for each save or restore using this environment. If there are at least as many drives available as the drive count, that number of drives are allocated for each save and restore request. If not, a reduced number of drives are allocated.

The default and highly recommended value is 1. The number of drives may range from 1 to 32.

3.4.6 Prologue and Epilogue

The prologue and epilogue attributes in the environment allow you to invoke a command procedure before and/or after the entire save or restore request. This allows you to perform pre-processing and post-processing operations around the entire request. Compare the order of environment prologue and epilogue procedures operations to the individual save and restore prologue and epilogue procedures, which are executed before and/or after each file or disk specification in the save or restore request. The order of execution is illustrated below:

- Environment prologue
- Start save or restore request
- First disk/file specification prologue
- First disk/file specification save or restore operation
- First disk/file specification epilogue
- Next disk/file specification prologue
- Next disk/file specification save or restore operation
- Next disk/file specification epilogue
-
- End save or restore request
- Environment epilogue (only on successful completion)

ABS defines logical names that can be used within the prologue or epilogue command procedure. These are defined in the process job table as follows:

Table 3–1 Logical Names Available to Environment Prologues and Epilogues

Logical Name	Meaning
ABS_SAVE_REQUEST_NAME	Name of the save request
ABS_RESTORE_REQUEST_NAME	Name of the restore request
ABS_STORAGE_CLASS	Name of the archive
ABS_EXECUTION_ENVIRONMENT	Name of the environment
ABS_NODE_NAME	Execution node name
ABS_OUTPUT_DEVICE	The name of output device, or devices, used by the save or restore request.

You should enter an OpenVMS command of up to 80 characters in the prologue and/or epilogue strings. For example:

```
@ABS$SYSTEM:ABS_ENV_PROLOGUE.COM
```

By default, there are no prologues or epilogues defined for an environment.

3.4.7 Retry Limit and Interval

The retry limit and retry interval options allows you to specify the number of times and how often ABS should retry a object in a save or restore request before operator intervention is required. Specify the following:

- **Retry Limit** - The number of retries (excluding the first attempt) to attempt before activating the notification options. A value of zero means no retries. You can also specify no limit, which means retries will be performed until the request either succeeds, or is manually stopped.
- **Interval** - The interval between retry attempts, expressed as a delta time. The default retry interval is 15 minutes.

Each time a retry attempt is made, ABS generates a warning message. If you wish to see the warning messages, select the *warning* option in the *when* attribute for notification.

3.4.8 Links Only and Span Filesystems

For UNIX clients, ABS provides the option to either backup UNIX symbolic links only, or to follow the UNIX symbolic links and backup up the data as well. The default is to backup the symbolic links only (not the data).

In addition, ABS allows you to backup only the root file system (such as the disk the root directory resides on) or an entire file system if the filesystem spans physical devices. The default is *nospan filesystems*, which copies the root file system only.

Both attributes apply to data type UNIX_FILES only.

Saving and Restoring Data

3.4 Environments

3.4.9 Listing Option

The listing option allows you to specify the type of listing generated for save and restore requests using this environment. Specify one of the following options:

- NONE - Does not generate a listing file
- BRIEF - Generates a brief listing file
- FULL - Generates a detailed listing file

If not specified, NONE is the default option.

3.4.10 Lock

ABS allows you to specify what a save request should do when data usage conflicts occur by means of the lock attribute. If you specify *lock*, ABS saves the data even if other applications have the data locked for write access. If you specify *nolock* ABS does not save the data if other applications have the data locked for write - this is the safer approach. If you specify *lock*, it is possible that the data you save is inconsistent, as the other application may be writing to the file during the actual save operation. Use lock with caution. The default is *nolock*.

3.4.11 Notification

ABS uses the notification attributes in the environment to determine who, how and under what circumstances users are notified of ABS events during save and restore operations. ABS supports notification using mail, OPCOM or both. You can specify up to 32 separate notification options in each environment, using the following attributes:

- MAIL - Specifies one or more mail users to be notified on certain types of event; specify one or more OpenVMS mail usernames (including node names as needed).
- OPCOM - Specifies one or more OPCOM classes to be notified on certain types of events - specify one or more OpenVMS OPCOM classes (e.g. TAPES) to be notified - notification will be directed to the (execution) node(s) specified in the save or restore request.
- TYPE - Indicates the level of information given. Select one of the following:
 - BRIEF - Contains only basic information (default)
 - NORMAL - Contains a moderate amount of information
 - FULL - Contains the maximum amount of information
- WHEN - Indicates when the notification should occur. Choose one or more of the following:
 - START - Sends notification at the start of a save or restore request
 - COMPLETE - Sends notification at the completion of a save or restore request with any status (success or failure)
 - WARNING - Sends notification if the request completes with a warning, error or fatal status
 - ERROR - Sends notification if the request completes with an error or fatal status
 - FATAL - Sends notification if the request completes with a fatal status

You associate a TYPE and WHEN for each MAIL or OPCOM option provided. If you do not specify a TYPE and/or WHEN, a notification option acquires a TYPE of BRIEF and a WHEN of COMPLETE.

If you specify no notification options, then by default ABS sends a brief OPCOM message to class TAPES on completion of every request executed under the environment.

3.4.12 Profile

ABS allows you to specify the user name, node name, cluster name, rights and privileges under which save or restore requests in the environment will execute. ABS supports three main options for username:

- **ABS** - This option specifies that all save and restore requests execute in the context of the ABS user (and account). You should not change the cluster, nodes, rights or privileges with this option, otherwise the saves and restores may not execute correctly. This is the default option, and is recommended for all system backup operations. It is also the required option for both UNIX and Windows client operations.
- **<REQUESTER>** - This option (including the angle brackets) instructs ABS to run associated save and restore requests under the user profile of the associated save or restore request. The save and restore user profile (which is not normally displayed and not is changeable) is that of the user who created the save or restore request. This option should be used for user backups. With this option you should not adjust node or cluster, but you can manipulate rights and privileges if the user's normal rights and privileges are not sufficient to run ABS save and restore requests.
- **Other user** - This option instructs ABS to run associated save and restore requests under the profile of a third user (not the save/restore creator or ABS). With this option, you can manipulate rights and privileges if the user's normal rights and privileges are not sufficient to run ABS save and restore requests. In addition, you should also define node and/or cluster to uniquely identify the user in the domain. Wildcard node and cluster names are supported.

It is recommended that you only specify a user profile for user backups. All other backups should run under the default ABS user profile.

3.5 Saves and Restores

The purpose of a save request is to backup data from primary online disk storage to either alternative disk or optical storage, or to tape storage. Saves are typically performed on a regular basis to provide protection in the event of a disk hardware failure, data corruption or deletion, or site disaster. Saves can also be used for archiving data that must be kept for a relatively long time for business purposes, but does not need to be online.

The purpose of a restore request is to return data from tape or alternate disk or optical storage back to primary online storage. In most cases, restores are performed after a disk hardware failure or user file corruption or deletion - these are usually one-time events. However, sometimes it is necessary to bring archived data online, and restores (perhaps scheduled restores) can be used for this purpose also.

ABS models save and restore requests in a similar fashion so in most cases the attributes for one are applicable to the other (exceptions are noted). The main difference is the direction of data transfer between disk and tape storage. As such, we shall discuss saves and restores in a single section.

You create each save or restore with a unique name, and associate a single archive and a single environment with it. Under certain circumstances, you can associate two archives with save requests. In addition, you implicitly create a schedule with each save request, and specify disks or files to save in objects called selections. As such each save or restore request is related to other ABS objects as shown below:

Saving and Restoring Data

3.5 Saves and Restores

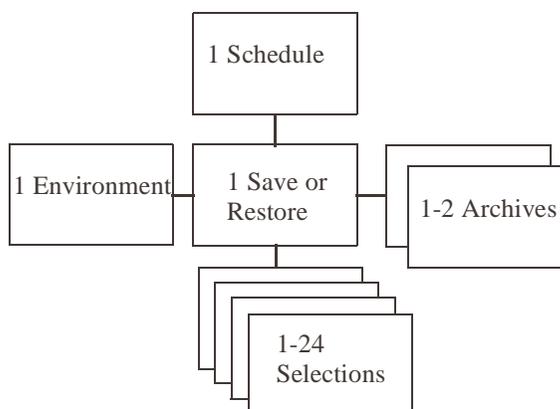


Figure 3–1 Relationships Between ABS Objects

The following sections describes the attributes of save and restore requests.

3.5.1 Save Name or Restore Name

You must assign a unique name to each save and restore request, which is used as the only method of referencing the request. There are no required conventions associated with save and restore names. However, in previous ABS versions, the names could be generated automatically so you might see names that are a combination of the creation date of the request and a generated unique identifier (UID) if you are converting from pre-V4 ABS. For version V4 and later, ABS almost always references saves and restores by name rather than UID, and ABS no longer shows UIDs by default.

3.5.2 Archive

Each save or restore requests is associated with one or two archives, which contain information about where the backed up data is stored. The two archives are for those requests that involve both full and incremental operations: the first archive applies to the full operations and the second applies to the incremental operations. In this way, fulls and incrementals can reside on different volume sets with their own retention days or expiration dates. In other types of request, only one archive is used.

If you do not specify an archive, ABS chooses *SYSTEM_BACKUPS*.

3.5.3 Base Date, Start Date and Skip Time

The base date is the date and time that you wish the request to first execute on a regular basis. The base date is used for two purposes:

- Defining the date and time to be used as a basis for scheduling - all scheduling intervals are based on both the date portion and time portion of the base date, and anniversaries of the base date at intervals defined by the frequency attribute.
- Defining the basis for full versus incremental saves for complex frequencies such as daily-full-weekly, log_2 and log_3. The base date and appropriate anniversaries of the base date define the date of the full saves.

Unless you want to change the scheduling or save type basis for the request, you would not change the base date. As such, the base date will remain a date in the past. Compare this to the *start date*, which specifies the next start date and time for the request. The start date is updated whenever the request is run to reflect the next time it is scheduled, or NONE if it is not scheduled again.

When a request is first created, and you specify only one of the dates, both dates are set (i.e. the next start date is the base date). By default, neither a base date or start date are supplied so the request is not scheduled for execution.

You can use the start date and skip time to request a one-time special, or non-scheduled, execution of the request. For example, assume that the normal scheduled time for a request is 23:00, as specified in the base date. However, you know that this is a particularly busy night and you want to start this request for tonight only at 21:00 instead. You can do this by setting the start date to 21:00. However, when the request is rescheduled, it will be rescheduled to the next iteration of the base date, or 23:00 the same day. You probably do not want this, so to avoid it you can set the start date together with a skip time to avoid running the request twice. The skip time is an exclusion time (expressed as a delta time) from the specified start date in which the request will not be rescheduled: normally you can set this to one day to avoid running the request twice in the same day. The following table shows some examples of base date, start date and skip time definitions based on a daily frequency:

Table 3–2 Use of Base Date, Start Date and Skip Time

Base Date	Start Date	Skip Time	Next Start
23-Aug 23:00	10-Sep 21:00	None	10-Sep 23:00
23-Aug 23:00	10-Sep 21:00	1-00:00:00	11-Sep 23:00
23-Aug-23:00	10-Sep-23:00	2-00:00:00	12-Sep 23:00

When you specify a skip time, ABS saves it in the database until the request is next rescheduled. When the rescheduling takes place, the skip time is applied to the calculation, then cleared from the database. If you set a skip time and do not see it in the request, then it has already been applied to the next start date.

3.5.4 Before Date, Since Date and Date Archived (Restore Only)

When restoring files, you can choose a specific *iteration* of the files based on their archive date - that is, the date that they were saved in the archive. If you know the exact date archived, use the *data archived* attribute. If you know only an approximate date archived, use the *before* or *since* attributes to specify a range of dates. So, for example, if you wish to restore a file as it existed in the first week of January, you can specify a before date of the 8th January (at midnight), or a since date of 1st January (at midnight). When determining appropriate before or since dates, you should probably lookup the files in the catalog before requesting a restore, so that you can specify before and since dates that uniquely identify a single iteration of the file to restore.

The before and since dates in the restore apply only to the archive date of the file. They are not related to the before and since dates in the selection object, which refer to files' online dates that are maintained by OpenVMS.

3.5.5 Catalog (Restore Only)

On a restore, you can specify a catalog name instead of an archive name if you know the name of the catalog from which to locate the restore information, and do not know the name of the archive. Normally, however, you would specify the archive under which the data was saved rather than the catalog.

Saving and Restoring Data

3.5 Saves and Restores

3.5.6 Include, Exclude, Data Type and Source Node

One of the more obvious attributes of a save or restore request are the file names, disk names path names or database names that you wish to save or restore. There are two options for specifying these names in a save or restore request:

- In an INCLUDE specification - You can specify the names directly in the save or restore request in an INCLUDE specification. You can specify multiple disks and/or files in a comma-separated list with the restriction that all disk and file specifications relate to a single data type (for example, VMS files). If you wish to mix file types in a single save or restore request (for example, VMS files and Windows files), then you must use the second option.
- Using SELECTIONS - With this option, you create selection objects directly using the MDMSView GUI or the CLI, specify the appropriate include specifications, then associate the selection object(s) with the save or restore. You can associate multiple selection objects with any save or restore request as long as the total number of disk, file, path or database specifications in all the selection objects does not exceed 24.

When you specify disk, file and database names by including them in the save and restore request, then you are effectively creating a default selection object. This selection object has the same name as the save or restore, with the suffix “_SAVE_SEL_DEL” or “_REST_SEL_DEF” respectively. You can specify the following attributes directly in the save or restore request for inclusion in the default selection:

- INCLUDE - A list of disks, files, paths or databases to include in the save or restore.
- EXCLUDE - A list of files to exclude from the save or restore that would otherwise have been included according to the include specification. This option applies to data type VMS FILES only
- DATA TYPE - The type of data to be saved or restored - select one of the following:
 - VMS Files - Applicable to VMS files. If only a disk is selected, a FULL backup of the entire disk is performed. If directory and file specifications are specified, then a SELECTIVE backup of files is performed.
 - Oracle Rdb Database Options - These options (which are version-number specific) specify that you wish to back up an entire Rdb database using the RMU backup utility. In this case, specify the name of the Rdb database files.
 - Oracle Rdb Storage Area - These options (which are version-number specific) specify that you wish to back up selected storage areas of an Rdb database. In this case specify both the database file name(s) and selected storage areas.
 - UNIX Files - This option applies to saving or restoring UNIX files on a client node. Enter a filesystem name in the format “/usr/...” to the level of granularity you want. With this option you must specify a SOURCE_NODE, which is the name of the UNIX node on which the online data resides.
 - Windows Files - This option applies to saving or restoring Windows files on a client node. Enter a file pathname starting with the device (for example: C:\Windows\...” to the level of granularity you want. With this option you must specify a SOURCE_NODE, which is the name of the Windows node on which the online data resides.
 - VMS Saveset cataloging - This option applies to cataloging existing VMS Backup saveset from tapes. Enter the tape name followed by a colon (:) and the saveset name (or wildcard). See Section 3.3 for more information on this functionality.

- **SOURCE NODE** - This attribute applies to data types UNIX FILES and WINDOWS files only, and specifies the name of the node on which the file data resides. For other data types, the node is specified through the nodes and groups attributes in the request.

The following table shows examples of the appropriate file, disk, path or database names that are valid for each data type:

Table 3–3 Disk, File, Path and Database Specification Formats

Data Type	Examples	Case Sensitive
VMS Files	\$1\$DUA420: (full disk, physical name) DISK\$USER1: (full disk, logical name) DISK\$USER1:[SMITH...]*.*;* (selective) DISK\$USER1:[SMITH]LOGIN.COM;3 (file)	No
Oracle Rdb Databases	DISK2:[USER_RDB]ACCOUNTS.RDB Do not specify a version number.	No
Oracle Rdb Storage Areas	DISK2:[RDB]ACCOUNTS.RDB/INCLUDE =BALANCES (save) DISK2:[RDB]ACCOUNTS.RDB /AREA=BALANCES (restore) Do not specify a version number - the include syntax is required, even from the GUI. If entered from the CLI, you must enclose the specification in quotes.	No
UNIX files	/usr/smith/* If entered from the CLI, you must enclose UNIX specification in quotes.	Yes
Windows files	C:\WINNT\SMITH* If entered from the CLI, you must enclose Windows specifications in quotes.	No
VMS saveset cataloging	tape_name:saveset_name	No

If you prefer to use selection objects directly (which enable you to specify additional selection criteria), then create a selection as shown in Section 3.3.3, then include the selection in the SELECTIONS attribute in the save or restore request. You can include up to 24 selections in a save or restore request with the caveat that a maximum of 24 disk, file or database file specifications (in total) are supported in a single save or restore request.

3.5.7 Delete Interval and Keep

Some saves and most restores are intended to be run only once, and have a frequency of ONE TIME ONLY. With this in mind, ABS automatically deletes such requests after a defined interval after the request has executed. This interval is the delete interval and can be customized for each save and restore request. If not specified, all ONE TIME ONLY requests are deleted approximately 3 days after execution: the actual delete is performed by a daily scheduled activity which runs at a certain time every day. If the frequency is something other than ONE TIME ONLY, ABS does not automatically delete the request. If the delete interval is set to NONE, then the request is deleted the next time the scheduled activity runs after execution of the request.

If you do not wish to have these requests automatically deleted, then set the *keep* attribute. This flags the request to be kept indefinitely and clears the delete interval.

Saving and Restoring Data

3.5 Saves and Restores

3.5.8 Destination (Restore Only)

ABS allows you to restore data to a different disk, directory, file system or pathname from where the data was saved. This is useful if you wish to make additional copies of data from the archive. If you wish to restore to a different location, enter the disk, directory, file system or pathname in the *destination* attribute of the restore. If not specified, the data is restored to the original source location of the data.

3.5.9 Environment

The environment attribute specifies an environment object name for this request. An environment contains attributes relating to how the request is executed. For example, an environment specifies data safety options, notification options and user profile. If not specified, the environment *SYSTEM_BACKUPS_ENV* is selected if available, otherwise *DEFAULT_ENV* is selected.

3.5.10 Frequency and Explicit Interval

ABS supports very flexible options for scheduling save and restore requests, both using the internal MDMS scheduling options and using a third part scheduler. The scheduling options can be divided into three main categories:

- **Standard** - ABS provides a list of standard options that you can specify, and the scheduling information is applied to the schedule object automatically. Standard options are supported by both internal MDMS scheduling and an external scheduler product. Standard options are all those that are neither custom or explicit.
- **Custom** - This option allows you to customize the schedule for the request if the standard options are not sufficient. For example, if you want to run the request every second Sunday in January, April, July and October, then the custom option can do this. You specify *CUSTOM* as the frequency, then modify the schedule object for the request directly. This option is applicable to internal MDMS scheduling only.
- **Explicit** - This option also allows you to customize your schedule, but this time with an external scheduler product. You specify *EXPLICIT* as a frequency, then enter a string into the *EXPLICIT INTERVAL* attribute. This attribute is a string that can be understood by the external scheduler product specifying the desired frequency. Alternatively, you can use the user interface of the external scheduler product to specify the frequency of the request. This option is applicable only to external scheduling options.

Select from one of the following frequencies:

- **ONE TIME ONLY** - Executes the save request one time only according to the option specified for Start Date.

After the save request has been executed and the delete interval (default approximately 3 days) have passed, ABS deletes the job from the database (including any external scheduler database). This is the default frequency if none is specified in the request.

- **ON DEMAND** - This option executes the save request according to the option specified for Start Date. The difference between One Time Only and On Demand is that ABS does not delete the request from the database.

DAILY - Executes a save request once per day according to the option specified for Base Date.

Saving and Restoring Data 3.5 Saves and Restores

- **WEEKLY FULL, DAILY INCREMENTAL (Saves only)** - This option enables you to create a single save request that executes a full backup operation once per week on the day specified in the Base Date, and an incremental backup operation for each subsequent day after the full backup operation is successful. ABS performs the full backup operation on a fixed day of the week during the 7-day cycle.

The Weekly Full/Daily Incremental Process:

For example, if the save request starts the full backup operation on Monday, ABS will always perform the full backup operation on Monday for that particular save request. This happens even if some of the subsequent incremental backup operations fail.

Example A:

Day	Type
Monday	Full
Tuesday	Level 1
Wednesday	Level 2
Thursday	Level 3
Friday	Level 4
Saturday	Level 5
Sunday	Level 6
Monday	Full

If that full backup operation fails, the cycle is repeated until a successful, full backup operation is achieved. ABS considers success and qualified success as a successful completed operation. ABS considers all other status as a failed operation.

Example B:

Day	Date and Time Run	Type	Result
Monday	31-MAR-1997 02:00	Full	Failure
Tuesday	01-APR-1997 02:00	Full	Failure
Wednesday	02-APR-1997 02:00	Full	Success
Thursday	03-APR-1997 02:00	Level 3	Success
Friday	04-APR-1997 02:00	Level 4	Failure
Saturday	05-APR-1997 02:00	Level 5	Success
Sunday	06-APR-1997 02:00	Assume skipping this day using a 3rd party scheduler	
Monday	07-APR-1997 02:00	Full	Success

Saving and Restoring Data

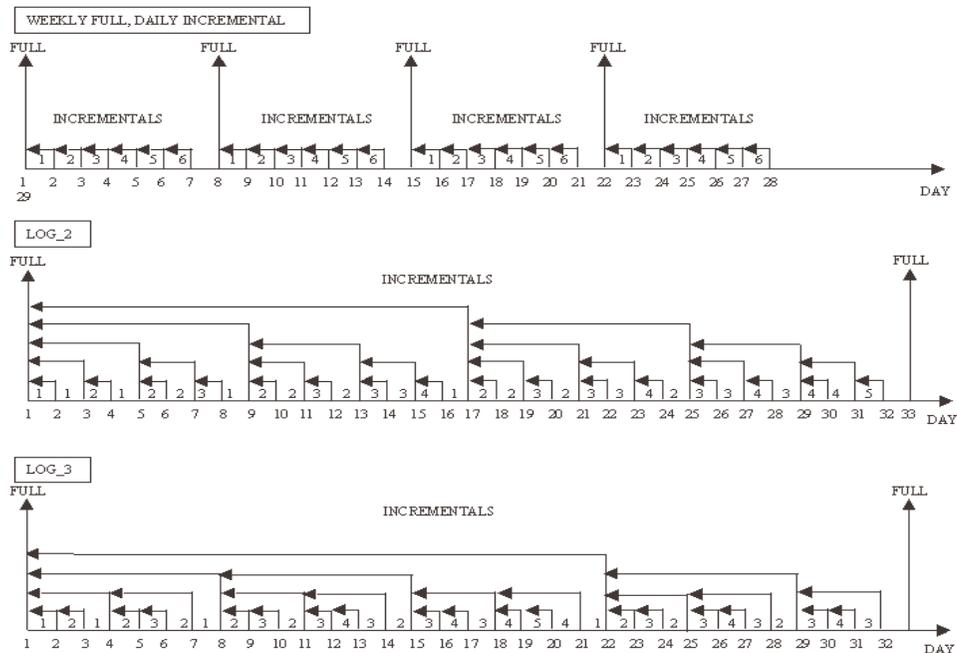
3.5 Saves and Restores

- If you are manually setting up your schedule to skip special days, ABS skips the next level of an incremental backup operation. In Example B, ABS skips Sunday and does not perform the Level 6 incremental backup operation. ABS resumes the full backup operation again on Monday, and the schedule once again repeats itself.
 - Notice also in Example B that ABS repeats the full backup operation until a successful full backup operation is achieved on Wednesday. If one of the incremental backup operations fail, ABS skips to the next level of the incremental backup operations. Unlike repeating the full backup operation, ABS does not repeat the same level of incremental backup operations during the 7-day cycle.
 - In the Example B, the Level 4 incremental backup operation failed on Friday. On Saturday, ABS resumes with a Level 5 incremental backup operation. However, the contents of the incremental backup operations are correct because ABS will back up all new or modified files since the last successful full backup or the last successful lower level incremental backup operation.
 - The save log file will contain the following backup command issued by ABS for Saturday, 05-APR-1997:
 - `$ BACKUP/.../SINCE="03-APR-1997 02:00:00.00"`
 - Because the last successful lower level incremental backup operation was performed on 03-APR-1997, all changes to any file since the date and time specified in the BACKUP command are included in the backup operation.
- WEEKLY - Executes the save request once per week according to the date and time specified for the start time.
 - BIWEEKLY - Executes the save request once every two weeks according to the date and time specified for the start time.
 - MONTHLY - Executes the job the first time on the date and time specified in the start time attribute. Subsequent jobs are scheduled on the first day of each month.
 - QUARTERLY - Executes the job the first time on the date and time specified in the start time attribute. Subsequent jobs are scheduled to execute on the first day of the quarter (3 month period).
 - SEMI_ANNUALLY - Executes the job the first time on the date and time specified in the start time attribute. Subsequent jobs are scheduled to execute on the first day of the month every 6 months and 12 months.
 - ANNUALLY - Executes the job the first time on the date and time specified in the start time attribute. Subsequent jobs are scheduled to execute every 12 months.
 - LOG-2 (Saves only) - ABS executes a full backup operation on day 1, and an incremental backup operation on day 2. On day 3, ABS executes an extended incremental backup operation. An extended incremental backup operation backs up any file modified since the last full or extended incremental backup operation.
 - LOG-3 (Saves only) - ABS executes a full backup operation on day 1, and an incremental backup operation on days 2 and 3. On day 4, ABS executes an extended incremental backup operation. An extended incremental backup operation backs up any file modified since the last full or extended incremental backup operation.

Advantages of Log-*n* backup operations:

Performing Log-*n* backup operations require less restore operations to fully restore a lost or corrupted disk volume. The higher the number of Log-*n*, the less restore operations you need to perform. Log-*n* backup operations are configured on a 32-day schedule, as shown in the examples below:

Figure 3–2 Complex Backup Schedules



- **CUSTOM** - This option allows you to define a specialized frequency by manipulating the associated schedule object directly. In this way you can define more flexible scheduling frequencies than are offered by the standard options. If you specify CUSTOM but do not modify the schedule object, then the default custom frequency is daily. This option applies only if internal MDMS scheduling is enabled (scheduler options INTERNAL and EXTERNAL).
- **EXPLICIT** - This option enables you to submit the save request using a specific scheduler interval. If you select Explicit, you must enter a scheduler time format valid for the scheduler being used in the EXPLICIT INTERVAL attribute. This option applies only if a third-party scheduler is being used (scheduler options SCHEDULER).
- **NEVER** - Never submits the save request and does not call the scheduler to create a job. For example, you may need to create one or more save requests before you determine their schedule. To submit the save request, modify the save request and change the scheduling option.

Depending on the selected scheduling option and the use of a 3rd party scheduler product, the Explicit Interval option allows to specify more flexible intervals. The Explicit Interval is passed as a string to the scheduler in use. Consult your scheduler's manual for more information.

Saving and Restoring Data

3.5 Saves and Restores

3.5.11 Incremental

Every save or restore request can be flagged as an incremental operation or a non-incremental operation. An incremental operation saves or restores files based on a previous operation - either a full operation or another incremental operation. For example, you could define a save request that performs a full disk save on Sunday, and an incremental save request that performs incremental saves on Monday through Saturday. The incremental saves will only save files that have been created or modified since the previous save (whether full or incremental). Restores can be performed in a similar fashion.

By default, saves and restores are not flagged as incremental. If you wish to define an incremental save or restore, then set the incremental attribute.

It is important to point out that if you execute an incremental save 127 times in a row without an intervening FULL save, then the 128th “incremental” save will actually be a full save. This rule actually applies to each individual file, disk, path or database specification within the save request, and as such, it is possible for the various files, disks, paths or databases within a single save request to be backed up at different incremental levels, or have a mixture of fulls and incrementals. As such, it is recommended that you intersperse a non-incremental (full) save at least once a week to avoid unexpected full backups on saves/restores marked incremental and to reduce the restore time required with a large number of incrementals. If you are mixing FULL and INCREMENTAL save requests, use the same catalog for both save requests so that the FULL catalog entry will be found and used as a base for the incrementals.

3.5.12 Nodes and Groups

ABS always performs save and restore operations on an OpenVMS *execution node*, under the control of the ABS coordinator process. Only one execution node actually executes any particular save or restore request at a particular time, but you can specify a list of compatible nodes using either the nodes or groups attributes. At execution time, the node list or group list is scanned in order to determine the execution node, and ABS will attempt to schedule the operation on the first such node. If ABS fails to establish a connection to that node, it will try the next node on the list, and so on until the request is successfully submitted.

For data types VMS files and Oracle databases of all types, the execution node is also the node where the data resides. Therefore, all execution nodes or groups must have access to the data being saved or restored. For data types UNIX files and Windows files, the execution node is the node running the ABS coordinator process, not the UNIX or Windows node on which the data resides: that node is specified by the SOURCE NODE attribute in the save or restore (or selection).

If you wish to enter nodes individually, enter a comma-separated list of nodes in the NODES attribute or select a list of nodes from the GUI. Enter the MDMS node object name (which should be the same as the DECnet Phase IV name if DECnet is running) - do not specify the TCP/IP name or DECnet Phase V fullname.

MDMS supports the notion of groups, whereby you can associate a list of nodes which have something in common (for example, nodes in a cluster) into a group, and simply reference the group name. In this case, you can simply enter one or more group names in the GROUPS attribute.

The NODES attribute and GROUPS attribute are mutually exclusive - you have to choose which one to enter.

If you enter neither nodes nor groups, then ABS enters the node from which the save or restore was created in the NODES attribute.

3.5.13 Prologue and Epilogue

The prologue and epilogue attributes in the save or restore request allow you to invoke a command procedure before and/or after each disk, file, path or database specification in the request. This allows you to perform pre-processing and post-processing operations around individual save or restore iterations. Compare the order of save and restore prologue and epilogue procedures operations to the environment prologue and epilogue procedures, which are executed before and/or after the entire save or restore request. The order of execution is illustrated below:

- Environment prologue
- Start save or restore request
- First disk/file specification prologue
- First disk/file specification save or restore operation
- First disk/file specification epilogue
- Next disk/file specification prologue
- Next disk/file specification save or restore operation
- Next disk/file specification epilogue
-
- End save or restore request
- Environment epilogue (only on successful completion)

ABS defines logical names that can be used within the prologue or epilogue command procedures. Each name is suffixed by “_n”, where n is the iteration number for each include disk, file, path or database specification. The value for n starts at 1 and goes to 24, the maximum number of include specifications supported by ABS. These logical names are defined in the process job table as follows:

Table 3–4 Logical Names in Save/Restore Prologues and Epilogues

Logical Name	Meaning
ABS_OS_OBJECT_SET_n	The include disk, file, path or database name currently being processed
ABS_OS_OBJECT_TYPE_n	The data type for the specification
ABS_OS_DMT_n	The type of operation: FULL, INCREMENTAL, SELECTIVE
ABS_OS_INCREMENTAL_LEVEL_n	For an INCREMENTAL operation, the incremental level being preformed
ABS_OS_VOLUME_SET_n	The volume set being used
ABS_OS_START_RVN_n	Starting relative volume number (RVN) of the volume set for the files being processed. The value is zero if the archive type is DISK,
ABS_OS_LAST_RVN_n	The last relative volume number in the volume set containing this specification. This value is valid for epilogue procedures only, and equates to “Not yet determined” for prologues. The value is zero if the archive type is DISK.

Saving and Restoring Data

3.5 Saves and Restores

Table 3–4 Logical Names in Save/Restore Prologues and Epilogues

Logical Name	Meaning
ABS_OS_START_FILE_POSITION_n	The starting file position of the saveset on the tape volume. This indicates how many tape marks from the beginning of the tape need to be skipped to arrive at the file. The value is zero if the archive type is DISK.
ABS_OS_SAVESET_NAME_n	The name of the saveset being used.
ABS_OS_SAVESET_FORMAT_n	The format of the saveset: either VMS, gtar or RMU.
ABS_OS_STATUS_n	The ABS status of the portion of the request for this specification

3.5.14 Reschedule

The reschedule attribute is used to create a new job with an external scheduler product. Normally, when you create a save or restore request, ABS creates a new scheduler job for the request. If you modify the request, then ABS modifies the existing scheduler job. However, there are circumstances whereby the scheduler job is deleted and needs to be re-created. You can set the reschedule attribute to re-create a new scheduler job for the request. This attribute has no effect when MDMS scheduling is in operation.

3.5.15 Selections

As discussed in section 3.2.3.6, you can specify the files, disks, paths or databases to be included in a save or restore request in one of two ways:

- By using the INCLUDE attribute in the save or restore request, and using a default selection
- By manually creating SELECTIONS, including the files, disks, paths or databases in the selection objects, then associating the selection objects with the save or restore requests.

The SELECTION attribute is how you associate a selection object with a save or restore request. Simply include the selection names as a comma-separated list in the selections attribute. If you wish to have no selections and use the default selection, specify no selections.

3.5.16 Sequence Option (Saves Only)

A save operation involves a data copy phase and a post-processing phase. For archive type TAPE, the post-processing phase does not require the use of a tape drive, so ABS could start on the next data copy phase using the drive before the post-processing phase of the previous operation is complete. This option speeds up the total save operation - if you want to use this option, specify OVERLAPPED as the sequence option. If, on the other hand, you prefer the data copy and post-processing phases to be performed sequentially, enter SEQUENTIAL for the sequence option.

By default, the sequence option is set to SEQUENTIAL.

3.6 Selections

ABS uses selections to hold information about files, disks, paths and databases to be saved or restored. You can elect to specify these names in one of two ways:

- By using the INCLUDE attribute in the save or restore request, and using a default selection
- By manually creating SELECTIONS, including the files, disks, paths or databases in the selection objects, then associating the selection objects with the save or restore requests.

The first option is discussed in Section 3.5.6 as part of the save and restore option. This section discusses the various attributes in the selection object.

The selection object gives you more flexibility to select files based on dates, agent qualifiers for the backup agent, and specifying conflict options on a restore. You can associate up to 24 selections with a given save and restore request, with the caveat that the total number of disk, file, path or database specifications in all selections does not exceed 24.

There are two steps in using selections:

- Creating or modifying a selection object directly by using the MDMSView GUI or the CLI.
- Associating the selection to the associated save and restore request by including it in the SELECTIONS attribute of the request.

The following sections describe attributes in the selection object.

3.6.1 Agent Qualifiers

ABS uses a *backup agent* to perform saves and restores, and the backup agent is dependent on the data type as follows:

- VMS Files - The backup agent is the OpenVMS BACKUP utility.
- Rdb Databases and Rdb Storage Areas - The backup agent is RMU Backup
- UNIX Files and Windows Files - The backup agent is gtar (tape archiver)

Although ABS passes information that you specify in the save, restore and environment to the backup agent, you can pass qualifiers directly to the backup agent using the *agent qualifiers* attribute. Refer to the appropriate backup agent documentation for information on these qualifiers.

3.6.2 Before Date, Since Date and Date Type (Saves Only)

For save requests, you can select files for saving based on the date files were last modified. You can specify either or both of the following:

- Before Date - Any version of the file modified before the specified date
- Since Date - Any version of the file modified after the specified date

If you specify both a before and since date, you are providing a range of dates in which to select files. If a file does not have a revision date (modified date), then ABS uses the creation date instead.

ABS does not yet support the date type attribute, which would allow you to select any one of the four online dates maintained by OpenVMS.

Saving and Restoring Data

3.6 Selections

3.6.3 Conflict Options (Restore Only)

When restoring files, you may find that there are files of the same name already located in the destination directory or original source location. You can specify what ABS should do if it encounters this situation by specifying one of the following conflict options:

- **NEW VERSION** - Restores the data and header and creates a new version of the file - applicable to VMS files only.
- **OVERLAY VERSION** - Overwrites the online version with the archive version of the data, but keeps the online version of the header.
- **REPLACE VERSION** - Deletes the online version of the file, and restores both the header and data from the archive.
- **RETAIN VERSION** - Keeps the online version of the header and data and does not restore the file from the archive.

If not specified, the default is **RETAIN VERSION**.

3.6.4 Include, Exclude, Data Type and Source Node

In exactly the same manner as in save and restore requests, you can specify the following attributes in selection objects directly:

- **INCLUDE** - A list of disks, files, paths or databases to include in the save or restore.
- **EXCLUDE** - A list of files to exclude from the save or restore that would otherwise have been included according to the include specification. This option applies to data type **VMS FILES** only
- **DATA TYPE** - The type of data to be saved or restored - select one of the following:
 - **VMS Files** - Applicable to VMS files. If only a disk is selected, a **FULL** backup of the entire disk is performed. If directory and file specifications are specified, then a **SELECTIVE** backup of files is performed.
 - **Oracle Rdb Database Options** - These options (which are version-number specific) specify that you wish to back up an entire Rdb database using the **RMU** backup utility. In this case, specify the name of the Rdb database files.
 - **Oracle Rdb Storage Area** - These options (which are version-number specific) specify that you wish to back up selected storage areas of an Rdb database. In this case specify both the database file name(s) and selected storage areas.
 - **UNIX Files** - This option applies to saving or restoring UNIX files on a client node. Enter a filesystem name in the format “/usr/...” to the level of granularity you want. With this option you must specify a **SOURCE_NODE**, which is the name of the UNIX node on which the online data resides.
 - **Windows Files** - This options applies to saving or restoring Windows files on a client node. Enter a file pathname starting with the device (for example: C:\Windows\...” to the level of granularity you want. With this option you must specify a **SOURCE_NODE**, which is the name of the Windows node on which the online data resides.
- **SOURCE NODE** - This attribute applies to data types **UNIX FILES** and **WINDOWS** files only, and specifies the name of the node on which the file data resides. For other data types, the node is specified through the nodes and groups attributes in the request.

The following table shows examples of the appropriate file, disk, path or database names that are valid for each data type:

Table 3–5 Disk, File, Path and Database Specification Formats

Data Type	Examples	Case Sensitive
VMS Files	\$1\$DUA420: (full disk, physical name) DISK\$USER1: (full disk, logical name) DISK\$USER1:[SMITH...]*.*;* (selective) DISK\$USER1:[SMITH]LOGIN.COM;3 (file)	No
Oracle Rdb Databases	DISK2:[USER_RDB]ACCOUNTS.RDB Do not specify a version number.	No
Oracle Rdb Storage Areas	DISK2:[RDB]ACCOUNTS.RDB/INCLUDE =BALANCES (saves) DISK2:[RDB]ACCOUNTS.RDB /AREA=BALANCES (restores) Do not specify a version number - the include syntax is required, even from the GUI. If entered from the CLI, you must enclose the specification in quotes.	No
UNIX files	/usr/smith/* If entered from the CLI, you must enclose UNIX specification in quotes.	Yes
Windows files	C:\WINNT\SMITH* If entered from the command line, you must enclose Windows specifications in quotes.	No

3.7 Schedules

ABS supports very flexible options for scheduling save and restore requests, both using the internal MDMS scheduling options and using a third part scheduler. The scheduling options can be divided into three main categories:

- **Standard** - ABS provides a list of standard options that you can specify, and the scheduling information is applied to the schedule object automatically. Standard options are supported by both internal MDMS scheduling and an external scheduler product. Standard options are all those that are neither custom or explicit.
- **Custom** - This option allows you to customize the schedule for the request if the standard options are not sufficient. For example, if you want to run the request every second Sunday in January, April, July and October, then the custom option can do this. You specify CUSTOM as the frequency, then modify the schedule object for the request directly. This option is applicable to internal MDMS scheduling only.
- **Explicit** - This option also allows you to customize your schedule, but this time with an external scheduler product. You specify EXPLICIT as a frequency, then enter a string into the EXPLICIT INTERVAL attribute. This attribute is a string that can be understood by the external scheduler product specifying the desired frequency. Alternatively, you can use the user interface of the external scheduler product to specify the frequency of the request. This option is applicable only to external scheduling options.

Saving and Restoring Data

3.7 Schedules

This section discusses the second option, custom schedules, which are only applicable to internal MDMS scheduling. To use a custom schedule, specify CUSTOM as the frequency on the save and restore request, then modify the attributes of the associated schedule object. The schedule object always has the name of the save and restore request, followed by “_SAVE_SCHED” or “REST_SCHED” respectively.

3.7.1 After Schedule

With ABS custom scheduling, you can actually define one schedule to execute after another schedule has completed. For example, if you want SAVE2 to execute immediately after SAVE1 completes, you can modify SAVE2’s schedule object and setting its AFTER SCHEDULE attribute to SAVE1’s schedule object. In this case:

```
SAVE2_SAVE_SCHED:
```

```
After Schedule: SAVE1_SAVE_SCHED
```

If you specify an after schedule and only want the associated request to execute after the after schedule (and not at any other time), then do not specify any other date or time attributes in the schedule. If on the other hand you want the associated request to execute at regular times AND after the specified after schedule, then you can associate date and time attributes to the schedule.

With after schedule, you can also define conditions upon which the schedule will run after the other schedule. The conditions are stored in an attribute called after schedule when. Select from one of the following:

- ALL - Always run the schedule after the dependent schedule completion
- SUCCESS - Run the schedule if the dependent save or restore completed with a successful status
- WARNING - Run the schedule if the dependent save or restore completed with a Warning, Error or Fatal status
- ERROR - Run the schedule if the dependent save or restore completed with an Error or Fatal Status
- FATAL - Run the schedule if the dependent save or restore completed with a fatal status
- NONE - Never run the schedule (can be used as a temporary placeholder)

If an after schedule name is defined, but no conditions are specified, the default condition is ALL. To remove the after schedule dependency, specify *no after schedule*.

3.7.2 Command

For ABS save and restore commands, the command to run a schedule and execute the associated save and restore request is:

```
MDMS RUN SCHEDULE schedule_name
```

You should not modify this command line, unless you know how to activate an ABS request in some other way.

For non-ABS save or restore requests, this command line can be any command that can be submitted to the OpenVMS CLI.

3.7.3 Dates, Days and Months

ABS supplies three attributes in the schedule object by which you can specify on what days you want the schedule to be regularly executed. These are:

- Dates - The dates of the month you want the schedule to execute
- Days - The days of the week you want the schedule to execute
- Months - The months of the year you want the schedule to execute

You can specify the actual dates in the month that you want the schedule to run by number. Here are some examples:

If you don't specify a date attribute, the default is every day of the month.

Table 3–6 Date Specifications

Dates	Explanation
1	First day of month
1-7	First week of month
1-7, 15-21	First and third week of month
1-31	Every day of month (default)

You can specify the actual day in the week that you want the schedule to run by name. Here are some examples:

Table 3–7 Day Specifications

Dates	Explanation
SUN	Sunday Only
MON-FRI	Monday through Friday Only
MON, WED, FRI	Monday, Wednesday and Friday Only
FRI-MON, WED	Friday, Saturday, Sunday, Wednesday

If you don't specify a day attribute, the default is every day of the week.

Finally, you can specify the actual months in the year that you want the schedule to run by name. Here are some examples:

Table 3–8 Month Specifications

Dates	Explanation
MAR	March Only
APR-SEP	April through September Only
JAN, APR, JUL, OCT	January, April, July, October Only

Saving and Restoring Data

3.7 Schedules

Table 3–8 Month Specifications

Dates	Explanation
JAN-DEC	All months

If you don't specify a month attribute, the default is every month of the year.

The dates, days and months attributes work together so that all must qualify for the schedule to be run. Therefore if you specify days SUN, but months of JAN, JUL only, then the schedule only runs on Sundays in January and July.

The following table shows some examples of how the days, dates and months attributes work together to produce custom schedules

Table 3–9 Combining Dates, Days and Months

Custom Schedule	Dates	Days	Months
First sunday of every month	1-7	SUN	JAN-DEC
First day of the quarter	1	SUN-SAT	JAN, APR, JUL, OCT
First and third saturdays of month	1-7, 15-21	SAT	JAN-DEC
First of month, every four months	1	SUN-SAT	FEB, JUN, OCT
Weekdays only	1-31	MON-FRI	JAN-DEC
Summer weekends only	1-31	SAT-SUN	JUN-SEP

If there are schedules that cannot be accommodated by this scheme, then you can use the INCLUDE and EXCLUDE attributes as explained below.

3.7.4 Include and Exclude

Although the days, dates and months attributes can produce a very flexible scheduling scheme, there may be specific days that you want to include or exclude regardless of the regular schedule. You can do this using the following attributes:

- **INCLUDE** - Include specific dates that otherwise may not be included using the days, dates and months attributes
- **EXCLUDE** - Exclude specific dates that otherwise may be included using the days, dates and months attributes

The dates are specified in the standard OpenVMS format DD-MMM-YYYY, and can range from the current date to up to 10 years in the future. Only dates may be specified, not times. Specification of include and exclude dates override the regular schedule as determined by the dates, days and months attributes.

You can also use the include and exclude attributes to augment the days, dates and months in situations that they do not cover what you want. For example, to run on the last day of every month, you can specify DATES 31, DAYS MON-SUN and MONTHS JAN-DEC, then specifically include 28-Feb, 30-Apr, 30-Jun, 30-Sep, 30-Nov.

3.7.5 Times

ABS allows you to specify times that you wish your schedule to run. Normally a schedule runs only once per day, but ABS allows you the flexibility to specify up to 100 times per day for a schedule to run. Simply specify times in the times attribute as a comma-separated list. Be careful to not specify so many times that the schedule executions overlap each other.

Media Management

This chapter expands on the MDMS object summary given in Chapter 2, and describes all the MDMS objects in detail, including the object attributes and operations that can be performed on the objects.

Before going into details on each object, however, the use of the MDMS\$CONFIGURE.COM procedure is recommended to configure your MDMS domain and the objects in it. In many cases this should take care of your entire initial configuration.

4.1 MDMS Domain Configuration

If you are configuring your MDMS domain (including all objects in the domain) for the first time, Compaq recommends that you use the MDMS\$CONFIGURE.COM command procedure. This procedure prompts you for most MDMS objects, including domain, drives, jukeboxes, media types, locations and volumes, and establishes relationships between the objects. The goal is to allow complete configuration of simple to moderately complex sites without having to read the manual.

The configuration procedure offers extensive help, and contains much of the information contained in this chapter. Help is offered in a tutorial form if you answer “No” to “Have you used this procedure before”. In addition, for each question asked, you can enter “?” to have help on that question displayed. Furthermore, if you type “??” to a question, not only will the help be displayed, but in most cases a list of possible options is also displayed.

This procedure is also useful when adding additional resources to an existing MDMS configuration. To invoke this procedure, enter:

```
@MDMS$SYSTEM:MDMS$CONFIGURE.COM
```

and just follow the questions and help.

A complete example of running the procedure is shown in Appendix A.

4.2 Domain

The MDMS domain encompasses all objects that are served by a single MDMS database, and all users that utilize those objects. A domain can range from a single OpenVMS cluster and its backup requirements, to multi-site configurations that may share resources over a wide area network or through Fibre Channel connections. An OpenVMS system running MDMS is considered a node within the MDMS domain, and MDMS server processes within a domain can communicate with one another.

The MDMS domain object is created at initial installation, and cannot be deleted. Its main focus is to maintain domain-wide attributes and defaults, and these attributes are described in the following sections.

Media Management

4.2 Domain

4.2.1 ABS Rights

The domain attribute `ABS_RIGHTS` controls whether a user having certain pre-V4.0 ABS rights can map these to MDMS rights for security purposes (see Chapter 5, *Security*, for more information about rights). Setting the attribute allows the mapping, and setting the attribute to false disallows the mapping.

4.2.2 Application Rights

The right `MDMS_APPLICATION_RIGHTS` is a high-level right that maps to a set of low level rights suitable for MDMS applications (for example, ABS and HSM). Normally these rights should not be changed, or at least not reduced from the default settings otherwise ABS and HSM may not function correctly. You may add rights to application rights if you have your own MDMS applications or command procedures. The ABS and `MDMS$SERVER` accounts should have `MDMS_APPLICATION_RIGHTS` granted in the User Authorization File.

4.2.3 Check Access

The check access attribute determines if access controls are checked in the domain. MDMS uses two forms of security: Rights and Access Control. Rights checking is a task-oriented form of security and is always performed. However, access control is an object-oriented form of security and can be optionally enabled or disabled with this attribute. Setting Check Access enables access control checking. Clearing Check Access disables access control checking even if there are objects with access control entries.

4.2.4 Deallocate State

When a volume is deallocated after its data has expired, it may go into one of two states. The transition state is an interim state that the volume goes into after deallocation, but it is not eligible to be used again until a period of time called the transition time expires. This is a safety feature that allows you to examine whether the data has legitimately expired, and if not to retain the volume (put back to the allocated state). If you do not wish this feature, you can disable the transition state and allow volume to return directly to the free state, where it is eligible for immediate allocation and initialization for new data. The domain deallocate state is applied to all volumes that are automatically deallocated by MDMS. When manually deallocating volumes, you can override the domain deallocate state with a state on the deallocate operation itself.

4.2.5 Default Rights

The MDMS default rights attribute maps a set of MDMS low-level rights to all users in the domain. This allows you to give all users a limited set of rights to access MDMS objects and perform operations, without having to expressly modify their accounts. Be aware that default rights are applied to all users on all nodes in the domain, so granting such rights should be carefully reviewed. By default, MDMS maps no rights to the default rights.

4.2.6 Mail Users

When MDMS deallocates volumes based on their scratch date (an operation that is performed once per day), it sends a mail message indicating which volumes were deallocated to the set of users defined in the mail users attributes. You should enter a list of users in the format `node::username`. Every user in the list will receive the deallocate volume mail messages. This mail address is also used when the ABS catalog unpack process encounters an error.

4.2.7 Maximum Scratch Time

The maximum scratch time is the maximum scratch time that can be applied to any volume when it is allocated. The scratch time is the period of time that you wish the volume to stay allocated because its data is still valid. The maximum scratch time imposes a maximum limit and overrides the volume's scratch time if it exceeds the maximum. For HSM, the maximum scratch time should be set to zero (unlimited), as HSM volumes' data remains valid until it is repacked. For ABS uses, this value should be set to the longest period of time you wish to retain any volume.

4.2.8 Media Type

The domain media type attribute is the media type that is applied to new volumes and drives by default when they are created. In a simple configuration, you may only have a single media type, so specifying it in the domain allows you to not have to specify it when creating individual drives and volumes. It may also be applied as a default to ABS archives. You may always override the domain default media type with a specific media type when you create or modify drives and volumes.

4.2.9 Offsite Location

The domain offsite location attribute is applied by default to the offsite location field of new volumes when they are created. The offsite location is an MDMS location that is used for secure storage of the volumes in case of a disaster. You can always override the domain default offsite location when you create or modify volumes.

4.2.10 Onsite Location

The domain onsite location attribute is applied by default to the onsite location field of new volumes when they are created. The onsite location is an MDMS location that is used for storage of the volumes when they are onsite, or quickly accessible to jukeboxes and drives. You can always override the domain default onsite location when you create or modify volumes.

4.2.11 OPCOM Classes

The domain OPCOM classes attribute contains the default OPCOM classes that are applied to new node objects by default when they are created. OPCOM classes are classes of users whose terminals are enabled to receive certain OPCOM classes. You can override the domain default OPCOM classes with specific classes on a per-node basis when you create or modify a node.

4.2.12 Operator Rights

The right MDMS_OPERATOR_RIGHTS is a high-level right that maps to a set of low level rights suitable for operators managing the domain. The default set of operator rights allow for normal operator activities such as loading and unloading volumes into drives, showing any object or operations, and moving volumes offsite and onsite. However, you can add or remove low level rights to/from the operator rights as you wish.

4.2.13 Protection

The domain protection attributes defines the default protection applied to new volumes when they are created. This protection is used by MDMS when it initializes volumes, and writes the protection on the magnetic tape volume itself. You can always override the domain default protection by specifying the protection specifically when creating or modifying a volume.

Media Management

4.2 Domain

4.2.14 Relaxed Access

The relaxed access attribute controls the security when a user or application tries to access an object without any access control entries, and access control checking is enabled. If relaxed access is set, such access is granted. If relaxed access is clear, such access is denied. The relaxed access attribute is ignored if the check access attribute is clear.

4.2.15 Request ID

MDMS uses sequentially increasing request identifiers for each request received by the MDMS database server, and this attribute displays the ID of the next request. If this ID is becoming very large, you can reset it to zero or one (or indeed any value) if you wish. The request ID automatically resets to one when it reaches one million.

4.2.16 Scheduler Type

MDMS performs scheduling operations on behalf of itself and ABS. For ABS scheduling, you can choose a scheduler type that best meets your needs, as follows:

- *Internal* - The default internal scheduler type uses MDMS schedule objects and OpenVMS batch queues. This option should be sufficient for most sites as the schedule object supports many custom scheduling options.
- *External* - This option uses MDMS schedule objects and OpenVMS batch queue, but the scheduling is submitted through a command procedure. You can use this option if you have a need to modify the command procedure to perform site-specific operations.
- *Scheduler* - This option uses an external scheduler product via command procedures. ABS supplies a template scheduler command procedure that you can modify to access your own scheduler product. You can also use this option to invoke the pre-V3.0 ABS DECScheduler V2.1B, as long as you have a license for that product.

MDMS-initiated scheduled operations such as MDMS\$MOVE_VOLUMES always use the internal MDMS scheduler.

4.2.17 Scratch Time

The domain default scratch time is the default scratch time applied to new volumes when they are created. Scratch time indicates how long a volume is to remain allocated (that is, how long its data is valid and needs to be kept). You can override the domain volume scratch time when you create, modify or allocate individual volumes. For HSM volumes, the scratch time should be set to zero (unlimited), since HSM data remains valid until a volume is repacked.

4.2.18 SYSPRV

MDMS uses user account rights as one mechanism for security within the domain. MDMS allows you to control whether the OpenVMS privilege SYSPRV can map to the ultimate MDMS right MDMS_ALL_RIGHTS. If you set the SYSPRV attribute, users with SYSPRV are assigned MDMS_ALL_RIGHTS, which means they can perform any operation subject to access control checks. Clearing SYSPRV gives users with SYSPRV no special rights.

Note

If you wish to use the SYSPRV attribute from the MDMSView GUI, the user's authorization file must have SYSPRV defined as a privilege and a default privilege. Having SETPRV is not sufficient as there is no way to set the SYSPRV privilege from the GUI.

4.2.19 Transition Time

The domain default transition time is applied to volumes by default when they are deallocated into the transition state. The transition time determines how long the volumes remain in the transition state before moving to the free state. This attribute is used alongside the deallocation state attribute, which determines the default state that volumes are deallocated into. You can override the domain default transition time when you create, modify, or deallocate a volume.

4.2.20 User Rights

The right MDMS_USER_RIGHTS is a high-level right that maps to a set of low level rights suitable for non-privileged users that perform ABS or HSM operations. The default set of user rights allow for user activities such as creating and manipulating their own volumes and loading and unloading those volumes into drives, showing their volumes. However, you can add or remove low level rights to/from the user rights as you wish.

4.3 Drives

A drive is a physical resource that can read and write data to tape volumes. Drives can be standalone requiring operator intervention for loading and unloading, in a stacker configuration that allows limited automatic sequential loading of volumes, or in a jukebox which provides full random-access automatic loading. Drives are named in MDMS using a unique name across the domain; it may or may not be the same as the OpenVMS device name, as these may not be unique across the domain.

The following sections describe the attributes of a drive.

4.3.1 Access

The access attribute controls whether the drive may be used from local access, remote access or both. Local access includes direct SCSI access, access via a controller such as the HSJ70, access via TMSCP, or access via Fibre Channel, and does not require use of the Remote Device Facility (RDF). Remote access is via a DECnet network requiring RDF. You can set the access to one of the following:

- All - Allows both local and remote access (default)
- Local - Allows only local access (as defined above)
- Remote - Allows only remote access using RDF

4.3.2 Automatic Reply

Automatic reply is the capability of polling hardware to determine if an operator-assist action has completed. For example, if MDMS requests that an operator load a volume into a drive, MDMS can poll the drive to see if the volume was loaded, and if so complete the OPCOM request without an operator reply. Set automatic reply to enable this feature, and clear to require an operator response. Please note that some operations cannot be polled and always require an operator reply. The OPCOM message itself clearly indicates if a reply is needed or automatic replies are enabled.

4.3.3 Device

The device attribute is the OpenVMS device name for the drive. In many cases you can set up the drive name to be the OpenVMS device name, and this is the default when you create a drive. However, the drive name must be unique within the domain, and since the domain can consist of multiple clusters there may be duplicate device names across the domain. In this case you must use different drive names from the OpenVMS device names. Also, you can specify simple or descriptive drive names which are used for most commands, and hide the OpenVMS device in the device name attribute.

Media Management

4.3 Drives

4.3.4 Disabled

By default, drives are enabled, meaning that they can be used by MDMS and its applications. However, you may wish to disable a drive from use because it may need repair or be used for some other application. Set the disabled flag to disabled the drive, and clear the flag to enable the drive.

4.3.5 Drive Number

If the drive is in a robotically-controlled jukebox, and the jukebox is controlled by MRD, you must set the drive number to the relative drive number in the jukebox used by MRD. Drives in jukeboxes are numbered from 0 to n, according to the SCSI addresses of the drives. Refer to the jukebox documentation on how to specify the relative drive number.

4.3.6 Groups

The groups attribute contains a list of groups containing nodes that have direct access to the drive. Direct access includes direct-SCSI access, access via a controller such as an HSJ70, access via TMSCP, and access via Fibre Channel. You can specify as many groups as you wish, in addition to nodes that may not be in a group.

4.3.7 Jukebox

If the drive is in a jukebox, you must specify which jukebox using the jukebox attribute. Enter a valid jukebox name from an MDMS-defined jukebox. If there is no jukebox, MDMS treats the drive as a standalone drive or as a stacker.

4.3.8 Media Types

A drive must support one or more media types in order for volumes to be used on the drive. In the media type attribute, specify one or more MDMS-defined media types that this drive can both read and write. If you wish, you can restrict the media types to a subset that you wish this drive to handle, and not all the media types it could physically handle. In this way, you can restrict the drive's usage somewhat.

4.3.9 Nodes

The nodes attribute contains a list of nodes that have direct access to the drive. Direct access includes direct-SCSI access, access via a controller such as an HSJ70, access via TMSCP, and access via Fibre Channel. You can specify as many nodes as you wish, in addition to groups of nodes in the groups attribute.

4.3.10 Read-Only Media Types

In addition to media types that a drive can read and write, a drive may support one or more additional media types that it can only read. In the read-only media type attribute, specify one or more MDMS-defined media types that this drive can only read. This allows this drive to be used when the application operation is read-only (for example, HSM unshelves or ABS restores). Do not duplicate a media type in both the media type list and read-only media type list.

4.3.11 Shared

You can designate whether a drive is to be used by MDMS applications and users only, or by non-MDMS users. If the drive is not shared, the MDMS server process allocates the drive on all clusters to prevent non-MDMS users and applications from allocating it. However, when an MDMS user attempts to allocate the drive, MDMS will deallocate it and allow the allocation. Set the shared attribute if you wish to share the drive with non-MDMS users, and clear if you wish to restrict usage to MDMS users. ABS users who do their own user backups are considered MDMS users, as are all system backups and HSM shelving/unshelving users.

4.3.12 Stacker

Certain types of drive can be configured as a stacker, which allows a limited automatic sequential loading capability of a set of volumes. Such drives may physically reside in a loader or have specialized hardware that allows stacker capabilities. If you wish the drive to support the stacker loading capability, set this attribute and make sure the jukebox attribute does not contain a jukebox name. If you wish the drive to operate as a jukebox or standalone drive, clear this attribute.

4.3.13 State

The drive state field determines the load state of the drive. The drive can be in one of four states:

- *Empty* - There is no volume in the drive
- *Full* - There is a volume in the drive
- *Loading* - A volume is being loaded into the drive
- *Unloading* - A volume is being unloaded from the drive

This is a protected field that is normally handled by MDMS. Only modify this field if you know that there are no outstanding requests and the new state reflects the actual state of the drive.

4.3.14 Allocate Drive (DCL Only)

You allocate a drive so that you can use it for reading and writing data to a volume. If you allocate a drive, your process ID and node is stored in the MDMS database, and the drive is allocated in OpenVMS for your process. Because the MDMSView GUI does not operate in a process context, it is not possible to allocate drives from the GUI.

You can either allocate a drive by name, or you can specify selection criteria to be used for MDMS to select an available drive for you and allocate it. The allocation selection criteria include:

- *Media Type* - Select a drive with the specified media type
- *Location* - Used with media type, select a drive in the specified location
- *Jukebox* - Used with media type, select a drive in the specified jukebox
- *Group* - Used with media type, select a drive that is supported by a node in the group
- *Node* - Used with media type, select a drive that is supported by the node
- *Volume* - Select a drive that is compatible with the specified volume (media type and placement)

You can also specify the following options when allocating a drive:

- *Assist* - A flag indicating whether you wish operator assistance if a drive cannot be allocated. Set if you wish assistance, and clear if you wish to use the retry limit and intervals to automatically retry (that is, wait for drives to become available).
- *Define* - Use define to set a logical name for the drive. The logical name evaluates to both the MDMS Drive Name and the OpenVMS device name, and can be used in either MDMS or other DCL commands.
- *Retry Limit and Interval* - If you wish the allocate to retry if there are no available drives, set the retry limit and interval, and specify noassist.
- *Preferred* - If you allocated a drive for a specific volume, you can set preferred to request that the same drive that the volume was last loaded is the preferred drive. If you clear preferred, this forces MDMS to perform a round-robin allocation of the drives.
- *Reply* - You can specify a symbol to receive an operator's reply message.

Media Management

4.3 Drives

- *Nowrite* - You can specify that the drive only has to be compatible for read-only media types, as the desired operation will only read from the drive.

4.3.15 Deallocate Drive (DCL Only)

If you allocated a drive using the DCL “Allocate Drive” command, you should deallocate the drive when you are finished using it, otherwise the drive will remain allocated until your process exits.

Simply issue a deallocate drive and specify the drive name or the logical name obtained from the define option in “Allocate Drive”.

4.3.16 Load Drive

MDMS supports two ways to load volumes into drives:

- *Load Drive* - This loads a scratch volume into a drive via operator intervention or by stacker operation. As such, this option is only for standalone and stacker controlled drives.
- *Load Volume* - This loads a specific volume into a drive, and can apply to all types of drives.

This section discusses the load drive option. The load volume option is discussed under volumes.

The “Load Drive” operation requests either that a scratch volume (in the free state) be loaded into the drive, or the next volume in the stacker is loaded into the drive. In either case, the volume ID of the volume is not known until the load completes, and MDMS reads the magnetic tape label to determine the volume.

The loaded volumes may or may not already be defined in the MDMS database. You can choose to create volume records by setting the “Create” flag, and optionally providing attributes to apply to the volume as follows:

- *Inherit volume ID* - This is the most comprehensive option as it allows the new volume to inherit all non-protected fields from the specified volume.
- *Media type* - Assign this media type to the volume. If you use inherit and media type, the specified media type overrides the inherit media type
- *Pool* - Assign this volume to the specified pool. If you use inherit and pool, the specified pool overrides the inherit pool.

When issuing the load drive request, you can specify whether the load is for read/write (almost always the case) or read-only, and whether operator assistance is required.

You can also specify an alternative message for the operator. This is included in the OPCOM message instead of the normal MDMS operator message. Use of an alternative message is not recommended.

When initiating a load from the DCL, you can choose a synchronous operation (default) or an asynchronous operation using the /NOWAIT qualifier. From MDMSView, a load is always asynchronous, so that you can continue performing other tasks.

4.3.17 Unload Drive

Unlike the load drive operation, the unload drive can be applied to any type of drive at any time. What it does is simply unload the current volume in the drive, and so you can use this when you don't know which volume is in the drive. Alternatively, you can use the unload volume operation if you know the volume ID in the drive.

The only option for unload drive is to request operator assistance if needed.

When initiating an unload from the DCL, you can choose a synchronous operation (default) or an asynchronous operation using the /NOWAIT qualifier. From MDMSView, an unload is always asynchronous, so that you can continue performing other tasks.

4.4 Groups

The group object is a logical object that is simply a list of nodes that have something in common. Groups can be used to represent an OpenVMS cluster, a collection of nodes that have access to a device, or for any other purpose. A node may appear in any number of groups. Groups can be specified instead of, or in addition to nodes in drive, jukebox, save and restore objects, and can be used interchangeably with nodes in pool authorization and access control definitions.

Groups contain only one attribute.

4.4.1 Nodes

The list of nodes that comprise the group. Nodes must be OpenVMS nodes that are defined in the MDMS database. You should not use groups for non-OpenVMS nodes (for example, ABS UNIX or Windows clients).

4.5 Jukeboxes

In MDMS, a jukebox is a generic term applied to any robot-controlled device that supports automatic loading of volumes into drives. Jukeboxes include small, single-drive loaders, large multi-drive libraries and very large silos containing thousand of volumes. In general MDMS does not make distinctions among the types of jukeboxes, except for the software subsystem used to control them. MDMS supports both the Media Robot Device (MRD) subsystem for SCSI-controlled robots, and the Digital Cartridge Server Component (DCSC) subsystem for certain silos.

The next sections describe the jukebox attributes.

4.5.1 Access

The access attribute controls whether the jukebox may be used from local access, remote access or both. Local access includes direct SCSI access, access via a controller such as the HSJ70, or access via Fibre Channel, and does not require use of the Remote Device Facility (RDF). Remote access is via a DECnet network requiring RDF. You can set the access to one of the following:

- All - Allows both local and remote access (default)
- Local - Allows only local access (as defined above)
- Remote - Allows only remote access using RDF

4.5.2 ACS ID

For DCSC-controlled jukeboxes, the ACS identifier specifies the Automated Cartridge System Identifier. Each MDMS jukebox maps to one Library Storage Module (LSM), and requires the specification of the Library, ACS and LSM identifiers.

4.5.3 Automatic Reply

Automatic reply is a capability of polling hardware to determine if an operator-assist action has completed. For example, if MDMS requests that an operator move a volume into a port, MDMS can poll the port to see if the volume is there, and if so complete the OPCOM request without an operator reply. Set automatic reply to enable this feature, and clear to require an operator response. Please note that some operations cannot be polled and always require an operator reply. The OPCOM message itself clearly indicates if a reply is needed or automatic replies are enabled.

Media Management

4.5 Jukeboxes

4.5.4 Cap Size

For DCSC-controlled jukeboxes equipped with Cartridge Access Points (CAPs), this attribute specifies the number of cells for each CAP. The first number is the size for CAP 0, the second for CAP 1, and so on. If a size is not specified, a default value of 40 is used. Specifying a cap size optimizes the movement of volumes to and from the jukebox by filling the CAP to capacity for each move operation.

4.5.5 Control

The control attribute determines the software subsystem that performs robotic actions in the jukebox. The control may be one of the following:

- *MRD* (Media Robot Device) - The default control uses SCSI commands to control the robot in the jukebox. When you specify MRD, you should also specify slot count, robot device name and a flag as to whether the jukebox supports magazines.
- *DCSC* (Digital Cartridge Server Component) - MDMS uses the DCSC subsystem to control the device. When you specify DCSC, you should also specify library ID, ACS ID, LSM ID and CAP sizes. DCSC is used for certain large silo devices only.

4.5.6 Disabled

By default, jukeboxes are enabled, meaning that they can be used by MDMS and its applications. However, you may wish to disable a jukebox from use because it may need repair or be used for some other application. Set the disabled flag to disabled the jukebox, and clear the flag to enable the jukebox.

4.5.7 Groups

The groups attribute contains a list of groups containing nodes that have direct access to the jukebox. Direct access includes direct-SCSI access, access via a controller such as an HSJ70, and access via Fibre Channel. TMSCP access is not supported for jukeboxes. You can specify as many groups as you wish, in addition to nodes that may not be in a group.

4.5.8 Library ID

For DCSC-controlled jukeboxes, the Library identifier specifies the library that this jukebox is in. Each MDMS jukebox maps to one Library Storage Module (LSM), and requires the specification of the Library, ACS and LSM identifiers.

4.5.9 Location

The location attribute specifies the physical location of the jukebox. Location can be used as a selection criterion for selecting volumes and drives. Specify an MDMS-defined location for the jukebox. This location may be the same as, or different from, the onsite location that volumes are stored in when not in a jukebox. If different, moves from the jukebox to the onsite location and vice versa will be done in two phases: jukebox to jukebox location, then jukebox location to onsite location, and vice versa.

4.5.10 LSM ID

For DCSC-controlled jukeboxes, the Library Storage Module (LSM) identifier specifies the LSM that comprises this jukebox. Each MDMS jukebox maps to one Library Storage Module (LSM), and requires the specification of the Library, ACS and LSM identifiers.

4.5.11 Nodes

The nodes attribute contains a list of nodes that have direct access to the jukebox. Direct access includes direct-SCSI access, access via a controller such as an HSJ70, and access via Fibre Channel. TMSCP access to jukeboxes is not supported. You can specify as many nodes as you wish, in addition to groups of nodes in the groups attribute.

4.5.12 Robot

For MRD-controlled jukeboxes, the robot name is the OpenVMS device name of the robot device. Robot names normally fall into one of several formats:

- GKx0 or GKxn01 for direct-connect SCSI
- \$n\$DUAnnn for access via an HSJ-type controller
- \$2\$GGnx for Fibre Channel access

If the jukebox is controlled by direct connect SCSI (first option), the device must be first loaded on the system with one of the following DCL commands:

```
Alpha - $ MCR SYSMAN IO CONNECT GKxxx/NOADAPTER/DRIVER=SYS$GKDRIVER.EXE
VAX - $ MCR SYSGEN CONNECT GKxxx/NOADAPTER/DRIVER=GKDRIVER
```

and the device name must begin with GK.

4.5.13 Slot Count

For MRD jukeboxes, the slot count is simply the number of slots (which can contain volumes) in the jukebox. Volumes reside in numbered slots when they are not in a drive. Slots are numbered from 0 to (slot count - 1). Filling in this field is optional: MDMS calculates the slot count by polling the jukebox firmware.

4.5.14 State

The state attribute is a protected field that describes the current state of the jukebox. A jukebox can be in one of three states:

- *Available* - Available for use, and not currently performing an operation
- *In-Use* - Currently performing a robot operation: robot operations occur sequentially; any new operation requested while the robot is in-use is queued
- *Unavailable* - The robot is unavailable for use for some reason

This field is normally maintained by MDMS, so you should not modify it unless a problem has occurred that needs manual cleanup (for example, the robot is stuck in the in-use state when it is clear that it is not in-use).

4.5.15 Threshold

MDMS provides the capability of monitoring the number of free volumes in a jukebox. A free volume is one that is available for allocation and writing new data. Many users would like to maintain a minimum number of free volumes in a jukebox to handle tape writing needs for some period of time. You can specify a threshold value of free volumes, below which an OPCOM message is issued that asks an operator move some more free volumes into the jukebox. In addition, the color status of the jukebox in MDMSView changes to yellow if the number of free volumes falls below the threshold, and to red if there are no free volumes in the jukebox. If you wish to disable threshold OPCOM messages and color status, set the threshold value to 0.

Media Management

4.5 Jukeboxes

4.5.16 Topology

The topology attribute specifies the physical configuration of a certain type of jukebox when it is being used with magazines. Topology is only useful when all of the following conditions are true:

- The jukebox is controlled by MRD
- The jukebox is in the TL820 class that allows you to open the jukebox door and insert entire magazines
- The jukebox is configured with towers, faces and levels

You specify the topology of the jukebox so that you can move magazines into and out of the jukebox by specifying a position rather than a start slot.

For each tower in the jukebox, you specify the number of faces in the tower, the number of levels in each face, and the number of slots in each level. For TL820-class jukeboxes, the typical values for each tower are 8 faces, 2 or 3 levels per face and 11 slots per level. The associated magazine contains 11 slots and fits into a position specified by tower, face and level. Other jukeboxes may vary.

4.5.17 Usage

The usage attribute determines whether this jukebox is set up to use magazines, and has two values:

- *Magazine* - The jukebox is configured to use magazines
- *Nomagazine* - The jukebox is not configured to use magazines

You should only set usage to magazine if you plan to use MDMS magazine objects and move all the volumes in the magazines together. An alternative is to move individual volumes separately, even if they reside in a physical magazine; in this case set usage to nomagazine.

4.5.18 Inventory Jukebox

MDMS provides the capability to inventory jukeboxes, and “discover” volumes in them and optionally create volumes in the MDMS database to match what was discovered. With this feature, you can simply place new volumes in the jukebox and let MDMS create the associated volume records with attributes that you can specify.

There are two types of inventory:

- Inventory using a vision system, which polls the jukebox’s firmware to locate volumes; this option is available for most larger library and silo type jukeboxes, and this operation takes only a few seconds to a few minutes depending on the size of the jukebox.
- Physical inventory, which actually loads volumes into drives to read volume labels. This is the only kind of inventory available for small loader-type jukeboxes that lack a vision system. This option is also available for larger jukeboxes, but is not recommended as it takes a considerable amount of time.

You can inventory whole jukeboxes, or specify a volume range or slot range, as follows:

- *Volume range* is supported for DCSC-controlled jukeboxes and MRD-based jukeboxes that have a vision system. Specify a range of volumes such as ABC001-ABC024. Up to 1000 volumes can be specified in a single range. When specifying a volume range, only those volumes are inventoried; other volumes in the jukebox are not.
- *Slot range* is available only for MRD-controlled jukeboxes, and can be applied to either vision or non-vision varieties. With slot range, only the specified slots are inventoried; other slots are not.

While inventorying jukeboxes, MDMS can find volumes that are defined and in the jukebox, that are not defined but are in the jukebox, and that are defined but missing from the jukebox. MDMS provides several options to handle undefined and missing volumes.

If you set the “Create” flag during an inventory, MDMS will create a volume record for each undefined volume it finds in the jukebox. You can specify in advance certain attributes to be applied to this volume record:

- *Inherit volume ID* - This is the most comprehensive option as it allows the new volume to inherit all non-protected fields from the specified volume. You normally use a volume known to be in the jukebox as the inherit volume ID.
- *Media type* - Assign this media type to the volume. If you use inherit and media type, the specified media type overrides the inherit media type
- *Preinitialized* - If you set this flag, the volume will be set to the free state and is immediately available for use. If you clear this flag, the volume will be set to the uninitialized state, and needs to be initialized prior to use. You should set or clear this flag depending on whether the volume is already physically initialized.

If you do not set the “Create” flag, then MDMS will not create new volume records for undefined volumes it finds.

Conversely, you can also define what to do if a volume that should be in the jukebox (according to the database) is found not to be in the jukebox. There are three options that you can apply using the “Missing” attribute:

- *Delete* - Delete the volume from the database; this is not normally what you would want to do because in most cases the volume is simply in another location and you probably want to keep it.
- *Ignore* - Do not change the database; this will probably leave the database in an inconsistent state, but you may prefer to perform the changes manually.
- *Move* - This is the default option, and changes the database to flag that the volume is in the volume’s onsite location.

When initiating an inventory from the DCL, you can choose a synchronous operation (default) or an asynchronous operation using the /NOWAIT qualifier. From MDMSView, an inventory is always asynchronous, so that you can continue performing other tasks.

4.6 Locations

A location is an MDMS object that describes the physical location other objects. Nodes, jukeboxes, magazines, volumes and archives can all have locations associated with them. Locations are used for volume and drive allocation selection criteria, and for placing volumes and magazines in known labelled locations.

Locations can be hierarchical, and locations in hierarchy that have a common source are considered compatible locations. For example, locations SHELF1 and SHELF2 are compatible if they have a common parent location such as ROOM2. Compatible locations are used when allocating drives and volumes using selection criteria, so you should only define hierarchies to the extent that you wish compatible locations. Locations that extend beyond a room or floor are generally not considered compatible, so you should not normally build location hierarchies beyond that level. Locations can also contain “spaces”, that are normally labelled areas in a location that volumes and magazines can be placed in an onsite location. If a volume or magazine contains a space definition, this is output in OPCOM messages so that operator can easily locate a volume or magazine when needed.

Locations contain two attributes, as defined in the following sections.

Media Management

4.7 Magazines

4.6.1 Parent Location

The parent location is an MDMS location object which is the next level up on the location hierarchy. For example, a location SHELF1 might have a parent location ROOM2, indicating that SHELF1 is in ROOM2. You should define a parent location only if you wish all locations belonging to the parent (including the parent itself) to be compatible when selecting volumes and drives. For example, in a hierarchy of SHELF1 and SHELF2 in ROOM2, volumes in any of the three locations would match a request to allocate a volume from ROOM2. Do not use the location hierarchy for other purposes.

4.6.2 Spaces

Locations can contain spaces, that are used in OPCOM messages when volumes and magazines are being moved from one place to another. Enter a range of spaces in an alphanumeric range separated by a dash. Examples of space ranges are 1-10, A-Z, AAA001-AAA099, 10A-10Z.

4.7 Magazines

A magazine is an MDMS object that contains a set of volumes that are planned to be moved together as a group. It can also relate to physical magazines that some jukeboxes (most notably small loaders) require to move volumes into and out of the jukebox. Magazines can be moved into and out of MRD-controlled jukeboxes with all their volumes at once.

However, just because a jukebox requires a physical magazine does not necessarily mean that you must use MDMS magazines. The physical magazine jukebox can be handled without magazines, and volumes are moved individually as far as MDMS is concerned. The choice should depend on whether you wish the volumes to move independently (don't use magazines) or as a group together (use magazines).

Magazines are not supported for DCSC-controlled jukeboxes. Magazines have the following attributes.

4.7.1 Jukebox, Start Slot and Position

The jukebox name contains the name of the jukebox if the magazine is in a jukebox. When in a jukebox, a magazine can optionally have a start slot or position, as follows:

- In a single-drive loader jukebox, only one magazine can be loaded at a time. In this case, the start slot is always zero, and the number of slots in the jukebox becomes the number of slots in the magazine.
- In larger, TL820-type jukeboxes, the magazine can be placed in many different places. If you have associated a topology with the jukebox, you can place the magazine in a "Position", specified by a tower, face and level specification. This is easier to physically locate in such jukeboxes than the alternative, which is a start slot designation. OPCOM messages for Move Magazine operations will state either position or start slot depending on whether a topology was specified.

All three fields are protected and normally managed by MDMS when a "Move Magazine" operation occurs. Only manipulate these fields if an error occurs and you need to recover the database to a consistent state.

4.7.2 Onsite and Offsite Locations and Dates

When not in a jukebox, a magazine may be either in an onsite or offsite location. An onsite location is one where the magazine can be quickly accessed and moved into a jukebox, which is also onsite. An offsite location is meant to be a secure location in the case of disaster recovery, and generally does not have local access to a jukebox. However, nothing in MDMS precludes the possibility of offsite locations having their own jukeboxes.

Each magazine should have an onsite and offsite location defined, so that operators know where the magazine is physically located. They use these locations, the jukebox name and the placement to determine where a jukebox is at a certain time. Both onsite and offsite locations should be MDMS-defined location objects.

Together with the offsite and onsite locations, you can associate an offsite and onsite date. These dates represent the date the magazine is due to be moved offsite or onsite respectively. Typically, magazines are moved offsite while their volumes' data is still valid and needs to be protected in a secure location. When the volumes' data expires, the magazine should be scheduled to be brought onsite, so that the newly-freed volumes can be used for other purposes.

If an offsite and/or onsite date is specified, MDMS initiates the movement of the magazines at some point on the scheduled date automatically. This is performed by the "Move Magazine" scheduled operation, which by default runs at 1:00 am each day. Operators will see OPCOM messages to move the magazines to either the onsite or offsite location.

If you do not wish to have MDMS move magazines automatically, either remove the onsite and offsite dates from the magazine, or disable the scheduled "Move Magazine" activity by assigning a zero time to its schedule object "MDMS\$MOVE_MAGAZINES".

4.7.3 Slot Count

The slot count specifies how many slots are in the magazine. Unlike jukeboxes, this value is required to make magazines work properly.

4.7.4 Spaces

While in an onsite location, the magazine can occupy a space, which is a labelled part of a location that uniquely identifies where the magazine is. A space can be designed to handle a single volume, but since magazines hold multiple volumes, multiple spaces can also be assigned. Enter either a space or a range of spaces for the magazine.

4.7.5 Move Magazine(s)

The supported way to move magazines from one place to another is to use the "Move Magazine" operation. You can move magazines on demand by issuing this operation, or you can let MDMS automatically move magazines according to pre-defined onsite or offsite dates (this is called a "scheduled" move). You can also force an early scheduled move if you want it to occur before the time that MDMS would initiate the move. Moving magazines into jukeboxes must always be performed manually.

When initiating a "Move Magazine", you can choose a destination for the magazine if the move is not a scheduled move. The destination can be one of three types of places:

- *Jukebox* - You wish to move the magazine and all of its volumes into a jukebox; you would then specify the jukebox name. If the jukebox is a large TL820-type jukebox, you must also specify the "Position", using tower, face and level, or start slot for the magazine.
- *Onsite location* - You wish to move the magazine to a location that is onsite to the computer hardware that normally uses it. You would then specify the onsite location name, and optionally one or more spaces that the magazine (or volumes from the magazine) will occupy.
- *Offsite location* - You wish to move the magazine to an offsite location for safety in case of a disaster. Specify an offsite location name.

If you wish to force a scheduled move, you can select "Scheduled". In most cases, the destination is predefined, so you don't need to specify it. However, you can specify an alternative destination for the scheduled move if you wish by specifying a destination as outlined above.

Media Management

4.8 Media Types

Finally, you can specify if you need operator assistance. This is recommended with “Move Magazine” as magazines cannot be moved without human intervention. Only if you plan to do the physical move yourself or you manually let someone know would you disable operator assistance.

4.8 Media Types

MDMS uses media type objects to hold information about the type of media that volumes and drives can support. MDMS uses media type as a major selection criterion for allocating volumes and drives, and volumes can only be loaded into drives with compatible media types.

Media types contain four attributes, as defined in the following sections.

4.8.1 Capacity

The capacity attribute indicates the capacity of the media in MB. This field is not used by ABS or HSM, but is used by the obsolete product “Sequential Media Filesystem” (SMF).

4.8.2 Compaction

This important field indicates whether you wish the tape to be written with firmware compaction. Enabling compaction usually doubles the capacity of the tape, so this is a desirable option which is set by default. Clear the attribute if you do not wish compaction.

4.8.3 Density

This field indicates the density of the tape that you desire. Many types of tape media (especially DLT tapes) support multiple densities, and certain types of drive can either read and write a certain density, or just read some densities. As such, you can define many media types with different densities that can be assigned to volumes and drives.

MDMS uses the density field when initializing volumes, so the density must be a valid OpenVMS density for the version of the operating system being used. Issue a “HELP INITIALIZE /DENSITY” command to determine the valid densities on the platform.

4.8.4 Length

The length field is used for information purposes only. If your media comes in various lengths, you can differentiate between types by using the length field. Specify an integer value that has meaning to your operators.

4.9 Node

An MDMS node is an OpenVMS system that is running MDMS. All nodes running MDMS must have a node object defined in the database for MDMS to work properly. The node name must be one of the following, in this order: the SCSNODE name, the DECnet Phase IV name, the host name of the IP node, or the DECnet Phase V name. If none of these are defined then MDMS\$SERVER should be the node name.

Nodes contain attributes as outlined in the following sections.

4.9.1 Database Server

MDMS operates as a group of co-operating processes running on multiple nodes in multiple clusters in an MDMS domain. One of these MDMS processes is known as the “Database Server”, and it actually controls all MDMS operations in the domain. Although only one node is the database server at any one time, you should enable multiple nodes to be *possible* database servers in case the actual database server node fails. In this way, failover is supported.

A database server must have direct access to the database files located in MDMS\$DATABASE_LOCATION. Direct access, access via MSCP, and access via Fibre Channel are all considered local access. Access via a network protocol or DFS are not considered local access. It is recommended that you enable at least 3 nodes as potential database servers to ensure failover capabilities.

4.9.2 Disabled

Set to disable the node as an MDMS node. Clear to enable the node as an MDMS node.

4.9.3 OPCOM Class

You can specify the OPCOM classes to be used by MDMS for operator messages on this node. By default, the domain default OPCOM classes are used, but you can override this on a node-by-node basis. Specify one or more of the standard OpenVMS OPCOM classes - messages are directed to all login sessions with these OPCOM classes enabled.

4.9.4 Transports and Full Names

You can define which network transports are defined for this node. There are four choices:

- *DECnet* - The DECnet transport is used
- *TCPIP* - The TCP/IP transport is used, and the TCP/IP full name is specified
- *DECnet, TCPIP* - The DECnet and TCP/IP transports can be used, with DECnet preferred
- *TCPIP, DECnet* - The TCP/IP and DECnet transports can be used, with TCP/IP preferred

If you identify TCP/IP as a supported transport, you must define the TCP/IP fullname in the TCP/IP fullname field. These fullnames are normally in the format “node.loc.org.ext”. For example, SLOPER.CXO.CPQCORP.COM

If you identify DECnet as a transport, you need to specify a DECnet full name only if you are using DECnet-Plus (Phase V). In this case, enter the full name, which is normally in a format such as LOCAL:.node. If you are running DECnet Phase IV, do not specify a DECnet full name. The node’s node name is used.

4.10 Pools

A pool is a logical MDMS object that associates a set of volumes with a set of users that are authorized to use those volumes. Every volume can be assigned one pool, for which we say that the volume is in the pool. The pool is then assigned a set of users that are authorized to use the volumes in the pool. If a volume does not have a pool specified, then it is said to belong to the “scratch pool” for which no authorization is required.

Pools have three attributes that are discussed in the following sections.

4.10.1 Authorized Users

You can specify a list of authorized users for the pool, as a comma-separated list of users. Each user should be specified as node::username or group::username, where both the node/group and username portions can contain wildcard characters (*%). To authorize everyone, you can specify *::*. To authorize everyone on a node you can specify nodename::*. Everyone in the authorized user list is allowed to allocate volumes in the pool. Other users require MDMS_ALL_RIGHTS or MDMS_ALLOCATE_ALL rights.

Media Management

4.11 Volumes

4.10.2 Default Users

Default users are authorized like the authorized users, but in addition are assigned this pool as their default pool. In this case, if they attempt to allocate a volume and don't specify a pool, they will allocate a volume from this pool. A particular user need only appear in one list: they do not need to be listed in both lists to be an authorized user to their default pool.

4.10.3 Threshold

Pools are useful for dividing volumes between groups or organizations, but they are only useful if there are free volumes in the pool. MDMS provides the capability of monitoring the number of free volumes in a pool. A free volume is one that is available for allocation and writing new data. Many users would like to maintain a minimum number of free volumes in a pool to handle tape writing needs for some period of time. You can specify a threshold value of free volumes, below which an OPCOM message is issued that asks an operator add some more free volumes to the pool. In addition, the color status of the pool in MDMSView changes to yellow if the number of free volumes falls below the threshold, and to red if there are no free volumes in the pool. If you wish to disable threshold OPCOM messages and color status, set the threshold value to 0.

4.11 Volumes

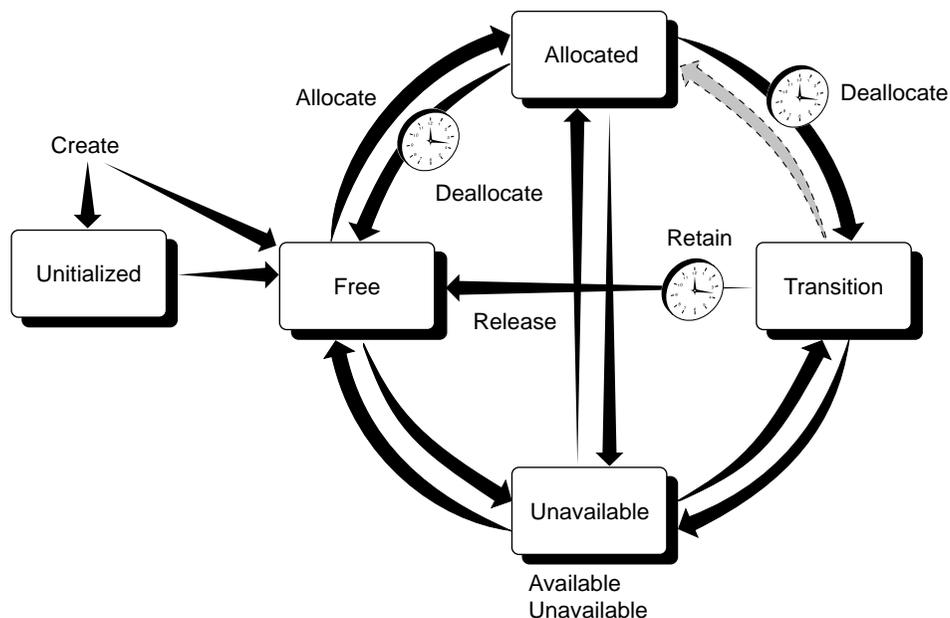
A volume is a physical piece of tape media that contains (or will contain) data written by MDMS applications (ABS or HSM), or user applications. Volumes have many attributes concerning their placement, allocation status, life-cycle dates, protection attributes and many other things.

Volume records can be created manually with a "Create Volume" operation, or automatically by MDMS with "Inventory Jukebox" and "Load Drive" operations. The MDMS\$CONFIGURE command procedure can also be used to create volumes.

Once a volume is created it acquires a state. This state determines how the volume may be used at any time, and to an extent where the volume should be placed.

The following figure illustrates the life cycle of volumes, and the following table indicates how a volume transitions from one state to another.

Figure 4–1 Volume State



CXO6756A

Each row describes an operation with current and new volume states, commands and GUI actions that cause volumes to change states, and if applicable, the volume attributes that MDMS uses to cause volumes to change states. Descriptions following the table explain important aspects of each operation.

Table 4–1 MDMS Volume State Transitions

Current State	Transition to New State	New State
Blank	MDMS CREATE VOLUME Volume Create	UNINITIALIZED
Blank	MDMS CREATE VOLUME/PREINIT	FREE
UNINITIALIZED	MDMS INITIALIZE VOLUME Volume Initialize	FREE
FREE	MDMS INITIALIZE VOLUME Volume Initialize	FREE
FREE	MDMS ALLOCATE VOLUME Volume Allocate	ALLOCATED
ALLOCATED	MDMS DEALLOCATE VOLUME Volume Deallocate or automatically on the volume scratch date	TRANSITION
ALLOCATED	MDMS DEALLOCATE VOLUME Volume Deallocate or automatically on the volume scratch date	FREE

Media Management

4.11 Volumes

Table 4–1 MDMS Volume State Transitions

Current State	Transition to New State	New State
TRANSITION	MDMS SET VOLUME /RELEASE Volume Release or automatically on the volume transition time	FREE
Any State	MDMS SET VOLUME /UNAVAILABLE Volume Unavailable	UNAVAILABLE
UNAVAILABLE	MDMS SET VOLUME /AVAILABLE Volume Available	Previous State
UNINITIALIZED	MDMS DELETE VOLUME Volume Delete	BLANK
FREE	MDMS DELETE VOLUME Volume Delete	BLANK

The following sections describes all the volume attributes in detail, followed by operations that you can perform on volumes.

4.11.1 Allocation Fields - Account, Username, UIC and Job

The account, username and UIC fields are filled in automatically when a volume is allocated, and reflect the calling user or specified user during the allocate. The username is a valid OpenVMS username on the client system performing the allocate, and the account and UIC is from the user's entry in the system Authorization (UAF) file.

These fields are normally maintained by MDMS and are protected fields. You should not modify these fields unless the volume is deallocated. MDMS maintains the Account, Username and UIC in the volume even after the volume is deallocated, so that you can "retain" the volume back to the allocated state in case of accidental deallocation.

The job name field is not used by ABS, HSM or MDMS.

4.11.2 Allocation and Movement Dates

There are several dates that maintain or control allocation and movement dates for volumes. These are as follows:

- Allocation Date - This is the date that the volume was last allocated using the "Allocate Volume" function. This field is protected and maintained by MDMS and should not normally be manually changed.
- Scratch date - This is the date the volume is due to be deallocated. MDMS will automatically deallocate the volume on the scratch date, but you can manually deallocate the volume before the scratch date as needed.
- Deallocation Date - This is the date the volume is actually deallocated. The volume may go into either the transition state or the free state depending on whether there is a transition time on the volume. This field is protected and maintained by MDMS and should not normally be manually changed.
- Onsite Date - This is the date the volume is due to be moved onsite from an offsite location. If this date is specified, MDMS automatically generates a "Move Volume" operation to move the volume onsite. Clear this field if you do not wish MDMS to automatically move the volume onsite.

- **Offsite Date** - This is the date the volume is due to be moved offsite. If this date is specified, MDMS automatically generates a “Move Volume” operation to move the volume offsite. Clear this field if you do not wish MDMS to automatically move the volume offsite.
- **Transition Time** - The transition time indicates that the volume is to enter the transition state when it is deallocated and remain in this state until the transition time has expired. In the transition state, the volume cannot be allocated for use. When the transition time expires, the volume enters the free state and may be re-used.

If an offsite and/or onsite date is specified, MDMS initiates the movement of the volumes at some point on the scheduled date automatically. This is performed by the “Move Volumes” scheduled operation, which by default runs at 1:00 am each day. Operators will see OPCOM messages to move the volumes to either the onsite or offsite location.

If you do not wish to have MDMS move volumes automatically, either remove the onsite and offsite dates from the volume, or disable the scheduled “Move Volumes” activity by assigning a zero time to its schedule object “MDMS\$MOVE_VOLUMES”.

4.11.3 History Dates

The history dates are maintained by MDMS, but are for information purposes only. MDMS does not use these dates to perform any operations. The following history dates are maintained:

- **Creation Date** - The date the volume was created in the database. This field is protected and maintained by MDMS and should not normally be manually changed.
- **Initialize Date** - The date the volume was last initialized. This field is protected and maintained by MDMS and should not normally be manually changed.
- **Freed Date** - This is the date the volume was last freed, either directly on deallocation, or upon expiration of the transition time. This field is protected and maintained by MDMS and should not normally be manually changed.
- **Last Access Date** - The date the volume was last loaded (and presumably accessed). This field is protected and maintained by MDMS and should not normally be manually changed.
- **Cleaned Date** - If the volume is a cleaning volume, MDMS updates the cleaned date to reflect the date that the volume was last used for cleaning. Otherwise it is set to the creation date.
- **Purchase Date** - The date the volume was purchased. MDMS makes this the same values as the creation date, but you can adjust this if needed.

4.11.4 State

The state field indicates where in a volume’s life cycle the volume exists. The state field itself is protected, and you should not normally adjust it unless an error occurs. However, you can “Update State” using certain keywords, which checks for validity and results in a consistent database state.

A volume can be in one of the following states, which are shown in normal life-cycle order:

- *Uninitialized* - The default state when a volume is created. This state indicates that the volume needs to be initialized prior to use, and cannot be allocated until then.
- *Free* - When a volume is initialized, and after it has been freed, the volume is in the free state. This means that the volume’s data (if any) is no longer valid and can be used to write new data. This is the only state from which a user can allocate a new volume for use.
- *Allocated* - After a volume is allocated, it enters the allocated state. It remains in this state until the scratch date is reached. MDMS automatically deallocates the volume when the

Media Management

4.11 Volumes

scratch date is reached, and it transitions to either the transition state (if there is a transition time on the volume) or the free state.

- *Transition* - If a volume is deallocated to the transition state, it remains in this state until the transition time expires. At this point, the volume re-enters the free state.
- *Unavailable* - This state is used by MDMS if it detects a problem with the volume. For example, if MDMS cannot read the label on this volume during a load, it puts it in the unavailable state. MDMS remembers the previous state (the “Available State”), so that when it comes out of the unavailable state, it goes back to its previous state.

A picture showing the normal state transitions is provided at the top of the volumes section.

While changing the state directly is not recommended, there are several options for changing state that are supported:

- *Available* - This changes the state from the unavailable state to the volume’s previous state
- *Unavailable* - This changes the current state to unavailable, and remembers the volume’s previous state. The volume cannot be used in this state. You should set this if you believe the volume is corrupted or broken and cannot be used.
- *Release* - If the volume is in the Transition State and you have verified that its data has expired, you can “Release” it to the free state immediately.
- *Retain* - If the volume is in the Transition State and you have verify that its data has NOT expired and is still useful, you can “Retain” the volume back to the allocated state. The existing allocated user, UIC and account are maintained. If the volume was in a volume set, the volume set is re-created.
- *Preinitialize* - If you know that the volume is already initialized, and the volume is in the Uninitialized state, this changes it to the Free state. It does not change the state if the volume is in any other state.

4.11.5 Media Types

A volume’s media types define the type of media for the volume, and what potential compaction or density options the volume can support. As such, before a volume is initialized, it can potentially support many media types. However, once a volume is initialized, MDMS uses the density and compaction attributes from a media type to physically write the tape. As such, a volume should only support one media type at and after the first initialization.

If the volume is in the Uninitialized state, select one or more MDMS-defined media types for the volume. If the volume is in any other state, select a single media type. If no media type is specified, the domain default media type is used.

4.11.6 Pool

A pool contains a collection of volumes that can be used by a set of authorized users. To insert a volume into a pool, simply specify a pool name in the volume’s pool field. If not defined, the volume is placed in the “scratch pool”, and it can be allocated by any user. If the volume is in the free state, the number of free volumes in the pool is incremented.

4.11.7 Previous and Next Volumes

These read-only fields indicate if a volume is in a volume set, and what the previous and next volumes are in the set, relative to this volume. A volume set is created when a tape write operation reaches end-of-tape and a new tape is required to complete the operation. ABS and HSM bind the next volume to the current volume, and create a volume set.

These fields are manipulated by “Bind Volume” and “Unbind Volume” operations, both manually and under control of MDMS applications.

4.11.8 Placement - Jukebox, Magazine, Locations, Drive

The placement fields of a volume indicate where the volume resides, and where it should reside when moved to an onsite or offsite locations. The placement attributes include the following:

- *Placement* - The current placement of the volume - options can be:
 - *Drive* - The volume is in a drive, indicated by the drive field
 - *Jukebox* - The volume is in a jukebox, indicated by the jukebox field and the slot field
 - *Magazine* - The volume is in a magazine, indicated by the magazine field and slot field
 - *Offsite* - The volume is in an offsite location, indicated by the offsite location field
 - *Onsite* - The volume is in an onsite location, indicated by the onsite location field, with optional space field
 - *Moving* - The volume is moving between one place and another

Placement is a protected field managed by MDMS. You should not change placement unless error recovery is needed.

- *Drive* - The name of the drive containing the volume. This field may contain a value even if the volume is not currently in a drive. The drive is a protected field managed by MDMS. You should not change drive unless error recovery is needed.
- *Jukebox* - The name of the jukebox containing the volume. The jukebox is a protected field managed by MDMS. You should not change jukebox unless error recovery is needed. The slot field indicates the jukebox slot the volume is in, and is filled in even if the volume is actually in a drive.
- *Magazine* - The name of the magazine containing the volume. The magazine is a protected field managed by MDMS. You should not change placement unless error recovery is needed. The slot field indicates the magazine slot the volume is in (this may or may not be the same as the jukebox slot). When the volume is in a magazine, its onsite and offsite location and date fields are invalid, as the magazine's onsite and offsite location and dates are used instead.
- *Offsite Location* - The designated offsite location for the volume (not valid if the volume is in a magazine)
- *Onsite Location* - The designated onsite location for the volume (not valid if the volume is in a magazine). The Space field indicates which space in the onsite location the volume is in or would be in if the placement is onsite.

4.11.9 Formats - Brand, Format, Block Factor, Record Size

The format fields are not used by ABS, HSM or MDMS, but can be used to document certain characteristics of the volume and its data format. The fields are as follows:

- *Brand* - The manufacturer of the volume - string
- *Format* - The record format used on the tape volume. Options are:
 - ASCII
 - BACKUP
 - EBCDIC
 - NONE
 - RMUBACKUP

Media Management

4.11 Volumes

- *Record Size* - An integer
- *Block Factor* - An integer

4.11.10 Protection

The protection field provides System, Owner, Group and World access protection for the volume. This protection is written to the volume when it is initialized, and provides protection from unauthorized use and re-initialization. The standard protection is:

```
SYSTEM(R, W) OWNER (R, W) GROUP (R) WORLD (None)
```

If protection is not set for the volume, the domain default protection is used.

4.11.11 Counters

MDMS provides three counters for volumes, as follows:

- *Mount Count* - This is a count of the number of times the volume is loaded - maintained by MDMS and incremented every time MDMS loads this volume
- *Error Count* - Not maintained by MDMS - set this field any integer you wish
- *Times Cleaned* - If the volume is a cleaning volume, this value is incremented each time the volume is loaded and used for cleaning. Otherwise it is set to 0.

4.11.12 Allocate Volume

You allocate volumes so that you can use them for writing new data. Allocating a volume places it into the Allocated state, and assigns the calling user (or specified user), UIC, and account in the allocation fields. This effectively reserves the volume to the user. The volume remains allocated to the user and unavailable for other use until the scratch date is reached, or unless the volume is manually deallocated.

When allocating a volume, you may specify the user for which you are allocating the volume (for example, ABS). If you do not specify a user, then you as the calling user are placed in the allocation fields.

Also, during allocation, you can change the following fields in the MDMS database to reflect the format to be used on the tape:

- *Format* - The record format used on the tape volume. Options are:
 - ASCII
 - BACKUP
 - EBCDIC
 - NONE
 - RMUBACKUP
- *Record Size* - An integer
- *Block Factor* - An integer
- *Scratch Date* - The date when the volume's data becomes obsolete and the volume should be deallocated - MDMS will automatically deallocate the volume at this time.
- *Transition Time* - When the volume is deallocated, the volume should go into the Transition State and remain in this state until the transition time expires, after which it will go into the Free State. If not specified, the volume goes into the Free State immediately on deallocation.

4.11.13 Allocate Volume(s) by Selection Criteria

Instead of allocating a volume by name, you can specify selection criteria to be used for MDMS to select a free volume for you and allocate it. You can also allocate a volume set by specifying a count of volumes to allocate. The allocation selection criteria include:

- *Media Type* - Select a volume with the specified media type
- *Location* - Used with media type, select a volume in the specified location
- *Jukebox* - Used with media type, select a volume in the specified jukebox
- *Pool* - Select a volume in the specified pool
- *Like Volume* - Select a volume like the specified volume (with the same media type, pool and placement)
- *Bind Volume* - Select a volume like the specified volume (with the same media type, pool and placement) and bind the new volume to the specified volume in a volume set

If you specify a volume count of more than one, then that many volumes will be allocated and placed in a volume set. If you also use the “Bind Volume” selection option, the new volume set is bound to the specified volume set.

You can also specify that you wish to change certain attributes of the volume as follows:

- *Format* - The record format used on the tape volume. Options are:
 - ASCII
 - BACKUP
 - EBCDIC
 - NONE
 - RMUBACKUP
- *Record Size* - An integer
- *Block Factor* - An integer
- *Scratch Date* - The date when the volume’s data becomes obsolete and the volume should be deallocated - MDMS will automatically deallocate the volume at this time.
- *Transition Time* - When the volume is deallocated, the volume should go into the Transition State and remain in this state until the transition time expires, after which it will go into the Free State. If not specified, the volume goes into the Free State immediately on deallocation.

4.11.14 Deallocate Volume

MDMS normally deallocates volumes when their scratch date expires. However, you can deallocate volumes manually in order to free them up earlier than planned. You can deallocate your own volumes, or with the appropriate rights deallocate volumes allocated to other users.

If the volume is in a volume set, the volume is also unbound from the volume set.

The following options are available when you deallocate a volume:

- *Deallocation State* - You can specify if the volume goes into the transition state or the free state on deallocation. The transition state disallows allocation until the transition time expires. You should make sure a transition time is specified, otherwise the domain default transition time is used. If you select the free state, the volume immediately goes into the free state.

Media Management

4.11 Volumes

- *Transition Time* - If the deallocation state is set to Transition, this is the length of time the volume remains in the transition state. If not specified, the volume's existing transition time is used, or the domain default transition time is used.
- *User Name* - If the volume is allocated to a user other than yourself, you must specify that user name for the deallocation to occur. You need MDMS_DEALLOCATE_ALL for this option.
- *Deallocate Volume Set* - If the volume is in a volume set, the entire volume set is deallocated by default. You can avoid this by deallocating only the single volume clearing the volume set attribute. Note that the volume set is still unbound at the deallocated volume.

4.11.15 Bind Volume

Binding volumes is the way to create volume sets, by binding one volume (or volume set) to another volume (or volume set). Normally, MDMS applications such as ABS and HSM perform automatic binding when they reach end-of-tape. However, it is sometimes necessary to perform manual binding. For example, if a volume set has been accidentally deallocated but is still needed, you may need to manually bind the set together (although the retain feature does this quite well).

There are only two options when binding a volume set:

- *Bind Volume ID* - The volume or volume set you wish to bind the current volume to. The current volume is always bound to the end of the specified volume set. Note that the allocated user of the volume set must match the allocated user of the current volume for the bind to be successful.
- *User Name* - If this volume is allocated to a different user than yourself, you must specify that user name. This requires the MDMS_BIND_ALL right.

When you bind a new volume to a volume or volume set, the new volume acquires the following attributes of the volume set:

- Onsite Date
- Offsite Date
- Scratch Date

The next and previous volumes are also updated appropriately.

4.11.16 Unbind Volume

Unbinding a volume removes the volume from the volume set without deallocating it. When unbinding a volume you can choose whether to unbind the entire volume set, or break the volume set at the point of the unbind. You can also unbind on behalf of the allocated user.

There are only two options for unbind:

- *User Name* - If this volume is allocated to a different user than yourself, you must specify that user name. This requires the MDMS_UNBIND_ALL right.
- *Unbind Volume Set* - Set this flag if you wish to unbind the entire volume set (that is, none of the volumes will be in a volume set anymore). Clear the flag if you wish to unbind at the point of the current volume (that is, the volumes before and the volumes after will remain in two separate volume sets).

4.11.17 Load Volume

MDMS supports two ways to load volumes into drives:

- *Load Drive* - This loads a scratch volume into a drive via operator intervention or by stacker operation. As such, this option is only for standalone and stacker controlled drives
- *Load Volume* - This loads a specific volume into a drive, and can apply to all types of drive.

This section discusses the load volume option. The load drive option is discussed under drives.

When loading a specific volume, you normally need to specify the drive in which to load the volume, unless a drive has been specifically allocated for a volume (via DCL only). Select a drive with a compatible media type for the volume.

If you are loading a volume into a jukebox drive, and the volume is not in the jukebox, you can specify an automatic “Move Volume” request to move the volume into the jukebox if desired. If you do not specify this option, and the volume is not in the jukebox, the operation will fail.

Another option is to request MDMS to check the volume label. This is normally a good idea as there can be mismatches between the volume’s magnetic label and its bar code label. If the labels do not match, the load fails. If you do not set the label check flag, the load may succeed but the label may be wrong. Use this option with caution.

When issuing the load volume request, you can specify whether the load is for read/write or read-only, and whether operator assistance is required.

You can also specify an alternative message for the operator. This is included in the OPCOM message instead of the normal MDMS operator message. Use of an alternative message is not recommended.

4.11.18 Unload Volume

You can unload a specific volume from a drive by issuing the “Unload Volume” operation. Unlike the “Unload Drive” operation which unloads any volume from the drive, the “Unload Volume” function checks the label on the volume on the drive before unloading it. If the label can be read and does not match the specified volume, the unload fails.

There is only one option for unload volume - operator assistance. This is recommended unless you are personally monitoring the unload operation.

4.11.19 Move Volume(s)

The supported way to move volumes from one place to another is to use the “Move Volume” operation. You can move volumes on demand by issuing this operation, or you can let MDMS automatically move volumes according to pre-defined onsite or offsite dates (this is called a “scheduled” move). You can also force an early scheduled move if you want it to occur before the time that MDMS would initiate the move. Moving volumes into jukeboxes or magazines must always be performed manually.

When initiating a “Move Volume”, you can choose a destination for the volume if the move is not a scheduled move. The destination can be one of four types of places:

- *Jukebox* - You wish to move the volume into a jukebox; you would then specify the jukebox name. You can also specify slot: this is required for small loader jukeboxes unless you perform an inventory afterwards. For vision-equipped jukeboxes, MDMS can determine an appropriate slot for the volume automatically.
- *Onsite location* - You wish to move the volume to a location that is onsite to the computer hardware that normally uses it. You would then specify the onsite location name, and optionally a space that the volume will occupy.

Media Management

4.11 Volumes

- *Offsite location* - You wish to move the volume to an offsite location for safety in case of a disaster. Specify an offsite location name.
- *Magazine* - You wish to move the volume into a magazine, and specify a magazine slot for the volume.

If you wish to force a scheduled move, you can select “Scheduled”. In most cases, the destination is predefined, so you don’t need to specify it. However, you can specify an alternative destination for the scheduled move if you wish by specifying a destination as outlined above.

Finally, you can specify if you need operator assistance. This is recommended with “Move Volume” because human intervention is necessary to move volumes. Only if you plan to do the physical move yourself or you manually let someone know would you disable operator assistance.

4.11.20 Initialize Volume(s)

MDMS supports initialization of volumes to make them available for use. Initializing a volume consists of writing an ANSI label on the volume, and applying compaction and density attributes and the volume protection field in the label. The volume is then free to be written. If the volume was in the Uninitialized state, it will now change to the Free state. All volumes need to be initialized at least once before ABS and HSM can allocate and use them.

Volumes that are already written need to be initialized again if you wish to use the whole volume for writing again. Both ABS and MDMS initialize volumes on every allocation.

When initializing volumes, you can specify four options:

- *Media Type* - If the volume does not have a media type specified, or has more than one media type specified, this is the time to specify a single media type for the volume. This is because the initialize instantiates the density and compaction attributes of the media type when writing to the volume.
- *Operator Assistance* - Recommended if a problem occurs during loading/unloading during the initialization.
- *Overwrite* - If set, this indicates that you wish the volume label to be written regardless of the label currently on the tape. If clear, the initialize will not take place if there is a different label on the tape.

Security

The security model used by ABS and MDMS is designed to provide flexibility in both the level of security and ease-of-use. ABS uses the MDMS security model, which is based on two main elements:

- *Rights* - The assignment of individual rights to particular users or classes of users that allow them to perform specific operations across the domain. Rights allow users to perform operations on all objects or certain object classes across the domain. This is a task-based form of security.
- *Access Control* - The assignment of access control is on a per-object basis, and allows specific users to perform specific types of operations on the object. This is an object-based form of security.

In addition, you can assign your MDMS domain one of three levels of access-control based security as follows:

- *No Access Control* - As the name implies, MDMS and ABS perform no access control based checking, even if individual objects have access control entries defined. However, rights continue to be checked.
- *Loose Access Control* - This option supports access control checking on objects, but only on those objects that have at least one access control entry. If there is at least one entry, access to the object is restricted to users with access control entries supporting the requested access. With objects with no access control entries, access to the object is implicitly granted.
- *Tight Access Control* - Designed for secure environments, this option supports access control checking on all objects. If there is at least one access control entry on an object, access to the object is restricted to users with access control entries supporting the requested access. With objects with no access control entries, access to the object is implicitly denied. This basically requires that all objects have appropriate access controls to be defined for the object to be used. Certain domain users may access normally inaccessible objects to prevent accidental lock-out due to insufficient access controls.

In general, the security model requires that both rights and access control are applied to users wanting to perform operations. In other words, having the “super” right MDMS_ALL_RIGHTS does not necessarily mean that you can do anything - any access control restrictions must also be satisfied.

This chapter discusses the security model in more detail.

5.1 MDMS Rights

MDMS controls users' operations with process rights granted to users and applications through low-level and high-level rights. High-level rights are simply a list of low-level rights assigned to selected classes of users as follows:

- *All users* - The MDMS Default rights are applied to all users, even though those users may not have any MDMS rights defined in their user authorization file. By default, MDMS does not assign any rights to the default rights, but you can change this.
- *MDMS Users* - The MDMS_USER right contains a set of rights typically granted to non-privileged MDMS users that require access to tape drives to perform their own backups and restores, or their own file shelving.
- *MDMS Operators* - The MDMS_OPERATOR right contains a set of rights typically needed for operations management of ABS, HSM and MDMS. This option contains more rights than are typically assigned to a non-privileged user, and allows such actions as creating volumes, loading volumes and drives, and inventoring jukeboxes.
- *MDMS Applications* - The MDMS_APPLICATION right contains a set of rights typically needed for MDMS applications - ABS and HSM. You should not modify these rights, as they may cause ABS and HSM to fail.
- *MDMS Administrators* - The MDMS_ALL_RIGHTS low-level right allows a user to perform any operation across the domain.

MDMS assigns defaults to all the high-level rights, you can modify high level rights to contain any list of low-level rights you wish.

To increase flexibility, you can also assign individual users a combination of low-level rights and high-level rights as needed.

The MDMS rights grant permission to perform certain kinds of *operations* across the domain, rather than restrict access to specific objects. The low-level rights typically are named in the following manner:

MDMS_operation_scope

where "operation" is typically an MDMS DCL command verb such as Allocate or Set. The "scope" may restrict the operation to a certain group of objects. Four common scopes are:

- *All* - Allows the operation on all objects. This is the most powerful scope.
- *Pool* - A volume-specific scope that allows operations on volumes to which you have authorization to the pool to which the volume belongs.
- *Own* - Allows you to perform operations on objects that you own.
- *Volume* - Allows the operation on all volumes.

The following table shows several examples of how the low-level rights work:

Table 5-1 Examples of Low Level Rights

Right	Explanation
MDMS_ALLOCATE_ALL	Can allocate any drive or volume
MDMS_SET_VOLUME	Can modify any volume's attributes
MDMS_SET_POOL	Can modify a volume's attributes, if the volume is in a pool for which you have authorization

Table 5–1 Examples of Low Level Rights

Right	Explanation
MDMS_DELETE_OWN	Can delete any object that you own
MDMS_SHOW_ALL	Can show any object
MDMS_SET_ALL	Can modify the attributes of any object

Refer to the *MDMS Reference Guide* for a complete list of MDMS low-level rights, and the default mapping of low-level rights to high-level rights.

In previous versions, ABS had a set of rights of its own, and you could map ABS rights to MDMS rights. For backward compatibility purposes, this mapping is still supported as shown in the following table:

Table 5–2 ABS to MDMS Rights Mapping

ABS Right	MDMS Rights Granted
ABS_BYPASS	MDMS_ALL_RIGHTS
ABS_BACKUP_JOB	MDMS_HR_USER
ABS_SHOW_ALL	MDMS_SHOW_ALL
ABS_CREATE_EXECUTION_ENV ABS_CREATE_STORAGE_CLASS	MDMS_CREATE_ALL MDMS_SET_ALL MDMS_SHOW_ALL

The mapping of ABS to MDMS rights is optional, and is controlled by the “ABS Rights” attribute in the domain. If you enable this attribute, the ABS to MDMS rights mapping is supported.

Finally, you can optionally enable the OpenVMS privilege SYSPRV to map to MDMS_ALL_RIGHTS. This makes it convenient for system managers to gain all needed rights by simply turning on SYSPRV. You can control this option using the “SYSPRV” attribute in the domain. If you enable this attribute, the SYSPRV mapping is supported.

Note.

If you wish to use the SYSPRV attribute from the MDMSView GUI, the user’s authorization file must have SYSPRV defined as a privilege and a default privilege. Having SETPRV is not sufficient as there is no way to set the SYSPRV privilege from the GUI.

Having access rights alone does not necessarily mean that you can perform all operations granted by those rights. Access control checks (if any) are applied in addition to rights to determine the final access to an object.

5.2 Access Control

Access control complements the MDMS rights access by granting object-based control over operations. You can assign up to 32 access control entries on any MDMS object, and define the types of access that the user in the entry is granted. There are seven kinds of access that users can be granted as shown in the following table:

Table 5–3 Access Control Allowed Operations

Allowed Access	Explanation
CONTROL	The user may modify the object's access control
EXECUTE	The user may perform operations on the object
DELETE	The user may delete the object
READ	The user may perform restore requests using this object (ABS only)
SET	The user may modify the attributes of this object
SHOW	The user may show this object
WRITE	The user may perform save requests using this object (ABS only)

You can manipulate access control from MDMSView using the Access tab on an object's Show screen. From the DCL, you can use the /ACCESS qualifier. In either case, the user name specification should include both node name and user name in the format:

```
node::username
```

From either interface, wildcards are supported in both the node and username portions of the specification. For example:

```
HOUST%::SMITH* allows users whose name begins with SMITH access from HOUST%
```

```
JUNGLE::* allows all users access from node JUNGLE
```

```
::SYSTEM allows all users named SYSTEM from all nodes
```

```
SYS001::JAMES allows user JAMES from node SYS001 only
```

If an access control entry matches a requesting user, only the access that is granted in the entry is granted to the user. Allowances that are not specifically listed are not granted.

Access control checks are optionally performed depending on attributes that you can set in the domain. The following table explains the settings:

Table 5–4 Domain Access Control Options

Check Access	Relaxed Access	Explanation
Clear	Clear	No access control checking is done
Clear	Set	No access control checking is done
Set	Clear	Access control is checked; if there are no access control entries, access is denied.

Table 5–4 Domain Access Control Options

Check Access	Relaxed Access	Explanation
Set	Set	Access control is checked; if there are no access control entries, access is granted.

Because of the nature of access control, it is possible to set up access control on an object so that no-one can access the object (even to restore its access control to a usable value). As such, MDMS provides three “escape” mechanisms to allow certain individuals to access the object even if not normally allowed through the access control mechanism:

- The owner of an object always has full access to the object. You can disable this feature by clearing the owner field in the object.
- Any user that is listed in the access control list of the domain has the same level of access to all objects.
- Finally, the user designated by “Last Updated By” in the domain has full access to all objects. This is the user who last modified the domain object, and so is assumed to be a trusted individual with recent access to the object.

To help in determining who the authorized domain users are, the SHOW DOMAIN operation does not use access control checking, so that anyone with MDMS_SHOW_ALL rights can show the domain.

Note

Access control checking is in addition to MDMS rights checking; both must be validated for access to be granted. In addition, if access control checking is disabled in the domain, MDMS rights checking is still performed for all operations.

5.3 Implementing a Security Strategy

Before assigning rights and access control entries to specific users you, as the system administrator, should carefully review your MDMS and ABS domain and determine what kind of access to allow your users.

The MDMS domain is a key determining factor as to what level of security you should implement. The MDMS includes all locations, nodes, jukeboxes, drives, volumes and other MDMS objects that are served by a single MDMS database. Implicit in this statement are that all users, operators and system managers on nodes in the domain are also part of the MDMS domain and need to be granted appropriate access to the domain resources.

Another key issue is what kind of security to the MDMS domain resources, including backup tape volumes, jukeboxes and drives, do you wish to assign to the domain users. Some possible scenarios with suggestions are shown below:

- Your domain consists of a single site that is managed by a single organization in a relatively free environment: MDMS high-level rights assigned to specific users are probably all that is necessary here. This is the simplest form of security to maintain.
- Your domain consists of a limited number of sites managed by a single organization in a secure environment: Since management of the domain is still under a single organization, a combination of high-level and low-level rights MDMS rights and limited access control checking may be appropriate. Access control entries on volumes and archives might be appropriate to specifically limit who can access data. Loose access control is recommended

Security

5.3 Implementing a Security Strategy

so objects without access control entries can be accessed. This level of security requires a moderate amount of maintenance.

- Your domain needs to be very secure, or your domain is geographically distributed or managed by multiple organizations that do not wish to interfere with each other's resources. In this case, tight access control with access control entries on every object may be required. This allows each organization to maintain their own resources (volumes, pools, saves, restores and so on), while sharing common resources such as nodes, jukeboxes and drives. An alternative to a distributed domain is to have multiple domains, but resources such as jukeboxes cannot be shared across domains. This level of security requires a substantial amount of maintenance.

Compaq recommends that you begin your security setup by assigning MDMS rights to users, and determining the high-level to low-level mappings carefully. Once these are assigned, assign various users high-level rights based on their function. For certain users whose access needs are not cleanly defined as "User" or "Operator", assign additional needed low-level rights to those users.

Compaq also recommends that you disable access control checking in the domain until all of the following are complete:

- You have installed the product(s), including any conversions from previous versions or previous products such as SLS.
- You have configured your domain.
- You have utilized the product(s) successfully in a production environment. You can perform ABS saves and restores, or HSM shelving and unshelving, successfully.
- You have analyzed your security requirements and determined that access controls on individual objects are required.

You may be concerned that MDMS enforces both access control and MDMS rights in order to access objects. Why can't MDMS_ALL_RIGHTS override all access controls? The answer to this is that MDMS_ALL_RIGHTS can be granted to anyone with SYSPRV privilege on any node in the MDMS domain. As the domain is a distributed object, potentially available to multiple organizations, you may not want privileged users in the domain but outside of your organization accessing your resources. As such, even users with MDMS_ALL_RIGHTS should be subject to access control checking.

However, you can enable domain-wide "super users" by defining them with full access control access to the domain. You should limit this access to trusted users across the domain. As these users have the same level of access to all objects as they do the domain, if they are also granted MDMS_ALL_RIGHTS, then they can perform any operation on any object in the domain.

User Interfaces

ABS and MDMS support two distinct user interfaces, as follows:

- A Graphical User Interface that combines both ABS and MDMS functions in a single GUI, and which you can run on OpenVMS systems and Windows PCs.
- A DCL interface, which now exclusively uses MDMS commands. The old ABS DCL interface is still available for backward compatibility, but will not be enhanced any further.

Both interfaces are designed to be full-function, so the choice of which interface to use is strictly your preference. It is not necessary to switch between interfaces to perform routine management tasks.

6.1 Graphical User Interface

MDMS provides a graphical user interface called *MDMSView*, which provides several views that you can use to manage your MDMS domain. MDMSView provides support for both media management and (if you have an ABS license) the Archive Backup System. MDMSView is designed to be the preferred interface to ABS and MDMS, with the goal of supporting most, if not all, of the regular management tasks. MDMSView supersedes all previous graphical interfaces for both ABS and MDMS.

MDMSView provides several views into the management of MDMS objects and requests, including ABS objects managed by MDMS. In V4.0, a limited number of views have been implemented, but many more are planned for future releases. MDMSView currently supports the following views:

- *Domain View* - With this view, you can see the relationship between objects. For example, under a specific location, you can see the nodes, (child) locations and jukeboxes in that location. At the next level, you can, for example, see the drives in the jukebox. On selecting a specific object, you can then examine and optionally change its attributes.
- *Event View* - This view allows you to examine the MDMS event and audit logfile, using a variety of selection criteria.
- *Object View* - Similar to the domain view, but the navigation is by object class and is not hierarchical. For example, all 17 objects classes are listed, and all objects in those classes are displayed. You can then select an object to manipulate.
- *Report View* - This view allows you to generate reports on a class of object using selection criteria and attribute display options. Currently, the report view supports only volumes.
- *Request View* - This view allows you to examine current activities in the MDMS database server. A request summary and detailed request information is available, with a single click refresh.
- *Task View* - While both the domain and object view allow manipulation on a single object at a time, the task view allows you to perform operations on multiple objects at once, or use

User Interfaces

6.1 Graphical User Interface

selection criteria to allocate objects. For example, you can create, show, delete and modify multiple objects (of the same type) in one operation.

- *Event View* - This view allows you to look at the MDMS\$LOG:MDMS\$LOGFILE_<node>.LOG file. It also allows you to reset the log, starting a new file, and set the logfilter (MDMS\$LOGFILTER logical).

Each view is provided in a tab from the main screen, and you can be working in several views at the same time, although only one is visible at a time. When switching from one view tab to another, the contents of the tab you are leaving are retained, and you can return to it at any time.

6.1.1 Starting MDMSView

6.1.1.1 OpenVMS Systems

MDMSView is installed at installation time on OpenVMS systems. Please refer to the Installation Guide for instructions on how to install MDMSView and Java on OpenVMS systems.

Once the installation is complete, the following commands are required to activate the GUI:

```
$ RUN YOUR VERSION SPECIFIC JAVA SETUP.COM FILE
```

```
$ SET DISPLAY/CREATE/NODE=nodename/TRANSPORT=TCPIP
```

```
$ MDMS/INTERFACE=GUI
```

where *nodename* is the TCP/IP node name of the system on which the MDMSView display is to appear. Although the GUI itself must run on an Alpha System V7.2-1 and higher, using Java 1.2 or higher, the MDMSView display can be redirected to any OpenVMS system, including VAX systems and those running OpenVMS versions less than V7.2-1.

6.1.1.2 Windows Systems

A SETUP.EXE package is also installed on OpenVMS systems for use on Microsoft Windows (R) PCs. This file may then be transported to any Microsoft Windows PC and executed. The SETUP.EXE will install MDMSView at a default location of C:\MDMSView, although alternative locations are possible. Once the PC installation is complete, you can execute MDMSView by clicking on the mdmsview.bat file in that directory. This batch file may need to be edited to include the machine and/or version specific directory of java.exe. Also, if you prefer not to use the default C:\MDMSView directory for the GUI files, you will need to edit those directories in the batch file.

6.1.2 Look and Feel

Once MDMSView is started, it will come up with the default look and feel for the system. For OpenVMS systems, this is the Java/Metal look and feel. For Windows systems, this is the Windows look and feel. You can adjust the look and feel to your taste by using the View menu as follows:

- OpenVMS systems: View>Java Look and Feel or View>Motif Look and Feel
- Windows systems: View>Java Look and Feel, View>Motif Look and Feel or View>Windows look and feel

Changing the look and feel requires a new login, so it's a good idea to change this before logging in. The value is saved in the MDMSView initialization file, and is used on all subsequent invocations from this location.

6.1.3 Logging In

Once MDMSView is started and the look and feel is set, you need to log into an OpenVMS system, even if you are running on an OpenVMS system already. You can log into any OpenVMS node in the MDMS domain, as long as it supports TCP/IP communication. Logging in requires three fields, as follows:

- Node name: TCP/IP name, address or node name alias indicating the OpenVMS node that you wish to log into. This node must be running MDMS.
- User name: A valid OpenVMS user name on the selected node.
- Password: The password associated with the user account on the selected node.

If there is a login failure for any reason, the node name and user name are retained for subsequent retries, but the password must always be re-entered.

After a successful login, the login screen disappears and the MDMSView splash screen is displayed.



6.1.4 Selecting A View

The next step is to select a view depending on what you want to do. Here are some tasks that you might wish to perform, and the associated view(s) that support them:

- Configure the MDMS domain - Domain view, object view or task view
- Create new objects - Domain view, object view or task view
- Modify object attributes - Domain view, object view or task view
- See relationships between objects - Domain view
- Delete objects - Domain view or object view

User Interfaces

6.1 Graphical User Interface

- Observe current MDMS operations - Request view
- Look at MDMS audit log entries and events - Event view
- Generate volume reports - Report view
- Create multiple objects - Task view
- Allocate volumes based on selection criteria - Task view
- Initialize a set of volumes - Task view
- Run and save or restore request - Domain view or object view
- View, reset, or add filters to the MDMS logfile.

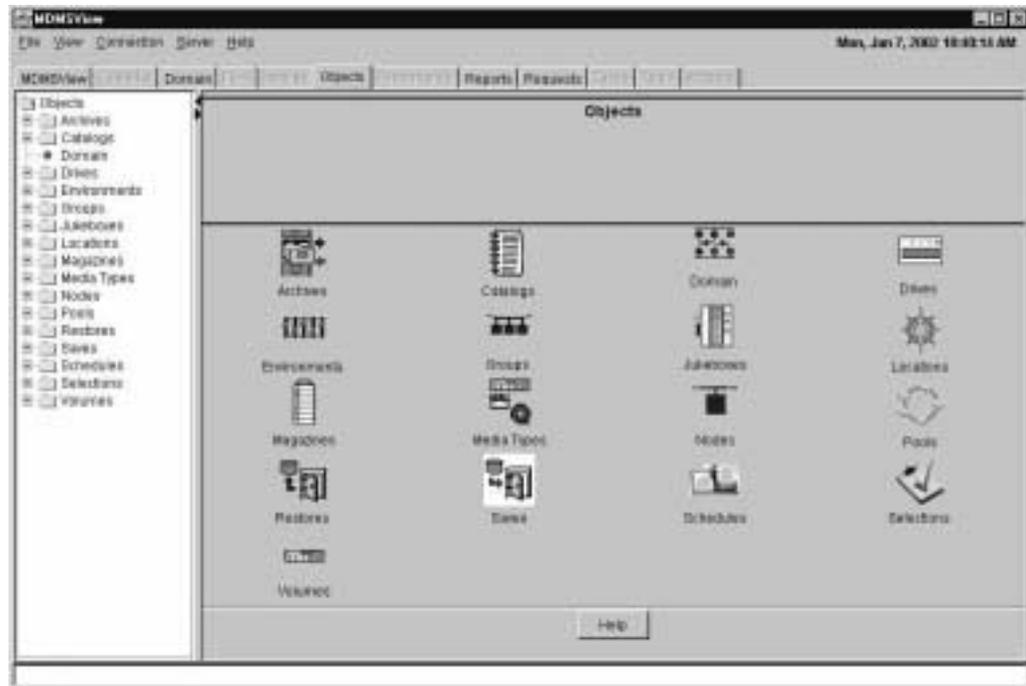
The domain view and object view produce attribute and operation screens that work on one object at a time. The task view produces screens that can operate on multiple objects, but restrict the display of attributes to those that are common across the objects. The request view is a specialized view that allows you to show current requests (as a whole or in detail), and allows you to delete requests as needed. The report view is a specialized view that generates customized volume reports.

All view displays are divided into two parts:

- A left screen containing tree nodes for navigation purposes. The structure of the nodes are view-specific, but the general concept is that there is a level for object classes (for example, Jukebox or Drive), and under the class is a list of relevant objects (e.g. JUKE_1, DRIVE_1). You can expand or contract any node (except for leaf nodes) in a manner similar to Windows explorer. If you click on a class name, the associated list of objects are displayed on the right side of the screen. If you click on an object name, the object's attributes and operations screens are displayed on the right.
- A right screen which contains the object attributes, request information or report that you wish to view. When clicking a class name from the left side, the objects in that class are displayed as icons on the right side. You can double-click on any object icon to bring up the object's attributes and operations screens. In the request view, you can refresh the whole request display or an individual request display by clicking the refresh button.

While resizing the MDMS view screens is not supported, you can choose to view only the left or right screens by using the arrows at the top of the division between the left and right screens. Clicking on the left arrow eliminates the left screen, and clicking on the right arrow eliminates the right screen. To restore the dual screens, click on the opposite arrow.

Figure 6–2 Object View Screen



6.1.5 Creating Objects

If you wish to create a new object, you can choose the Domain, Object or Task Views to accomplish this. The Domain and Object Views create objects one at a time, while the Task View can create multiple objects.

To create an object, use one of the following methods:

- Click on a class name (e.g. Jukebox) on the left screen, and the class object's icons are displayed on the right screen. On the right screen press the "Create" button to display a create screen.
- From the object view only, click on "Object", then double-click on one of the class icons that are displayed. On the right screen press the "Create" button to display a create screen.
- From the left screen, right-click on a class name, and a popup menu appears. Click on "Create" to display a create screen for that class.
- From the task view, expand the Create task and click on one of the class names that appear.
- From the task view, right click on the create task, and a popup menu appears. Click on the appropriate class name.

Once a create screen appears, (except for catalogs) you are prompted for two pieces of information:

- A name for the new object or objects
- An inherit object

The domain and object views allow creation of only one object at a time, whereas the task view allows a comma-separated list of new objects (and also ranges in the case of volumes). Depending on the view, enter the name or names of the new objects you wish to create.

The inherit object allows you to copy most of the attributes from the inherit object to the object being created. If you wish to specify an inherit object, use the combo box to select the existing

User Interfaces

6.1 Graphical User Interface

inherit object. This must be the same type of object, except in the case of restores, in which case you can inherit from either a restore or a save object.

After clicking create, the new object attribute and operations screens appear, which you can then modify to your liking. In the task view, this screen modifies all the newly created objects.

Figure 6–3 Drive Create Screen



6.1.6 Showing and Modifying Objects

For objects that already exist, you can use the Domain View, Object View or Task View to show and optionally modify objects, or to perform operations on them.

To view an object, use one of the following methods:

- From the Domain or Object Views, from the left screen, expand a class name, and click on an object name.
- From the Domain or Object View, click on a class name from the left to bring up the class object icons on the right screen, then double click on an object icon.
- From the Task View, right click on the Show task, and a popup menu appears, then click on an appropriate class and object.

When an object is selected, its attributes and operations are displayed in a two-dimensional tab screen as follows:

- Vertical tabs on the right side of the screen contain the Show and any operations associated with the object. Many objects just have a Show tab, but some (for example, volumes) have a whole list of operational tabs such as load, unload and so on. You can switch between the tabs by simply clicking on them.
- For the Show screen, there are also horizontal tabs that display related attributes about the object. Many simple objects have only a General tab that shows all attributes. Other attributes have General and Advanced tabs, if there is not enough room on one tab. Other tabs include:
 - A Show Access tab, which shows the access controls on the object. This is in a common format for all objects. If your site does not use access controls, you can disable these tabs using the view menu: View>No Access Control Tabs

User Interfaces

6.1 Graphical User Interface

- The Show screen for Jukeboxes and Magazines also has a Contents tab that shows the current contents of the drives and slots in the jukebox, and the slots in a magazine.
- Saves and Restores have a selections tab, that shows all selections for the save or restore, and a log tab that displays the latest version of the associated log file.

If you select the Show screen and wish to modify attributes, use the tool tip text for help on any field. Select appropriate values (from all the show tabs as needed), then click on Set. This sends the currently displayed values from all tabs to the MDMS server. If you just wish to view the object's attributes without modification, click on Cancel after viewing the attributes. This returns you to the object class screen.

MDMSView supports switching from one object to another during displaying of values. For objects that appear in combo boxes or lists, you can view related objects without losing the context of the current object. Each combo box or list attribute supports two methods of viewing, selecting and creating objects:

- Click on a small button to the right of the combo box or list to receive a popup menu for the field
- Right-click on the combo box or list and receive the same popup menu for the field

From the menu, there are the following options:

- Show - To show the selected object
- Create - To create a new object
- Reset - To go back to the previously selected objects
- Clear - To clear the selection
- Add and Remove (list only) - To add and remove an object by name
- List all (list only), lists all the objects

If you select Show or Create, you will go to an appropriate screen. When you then complete your operation on that object, you will come back to the original object.

Figure 6–4 Save Show General Screen

The screenshot shows the 'SAVE SHOW GENERAL' window. At the top left, there's a 'SAVE' icon and the text 'SAVE1'. The status bar at the top right shows 'Status: NONE'. Below that, 'Job Number: 0', 'Last Start Date:', and 'Next Start Date:' are visible. The main area has a 'Description' field with 'PDC2-SMITH' and an 'Owner' field. Below this are tabs for 'General', 'Advanced', 'Selections', 'Log', and 'Access'. The 'General' tab is active, showing a 'Data Select Type' dropdown set to 'NONE', a 'Source Node' field, and a 'Default Selection' field with the value 'USER:[SMITH]*.*'. There are 'Include' and 'Exclude' buttons. The 'Schedule' section has 'Base Date', 'Start Date', 'Bkp Time' (00:00:00), 'Explicit Interval', and 'Delete Interval' (00:00:00). The 'Attributes' section has 'Archive 1' (USER_SAVE), 'Archive 2' (NONE+), and 'Environment' (USER_SAVE_ENV). There are also 'Groups' and 'Execution Nodes' sections. At the bottom, there are buttons for 'Delete', 'Unprotect', 'Run', 'Set', 'Cancel', and 'Help'.

User Interfaces

6.1 Graphical User Interface

6.1.7 Deleting Objects

You can delete objects from the Domain, Object and Task Views. To delete an object, perform one of the following:

- Display the object as discussed in the previous section, then click the Delete button at the bottom of the screen.
- Right-click on the object name from the left screen, then select Delete from the pop-up menu.
- From the task view, select the Delete task, then select the object class, then select the object names from the list on the delete screen.

A request to delete an object will always bring up a Delete dialog box for confirmation of the delete. You can confirm “OK” or “Cancel” from here.

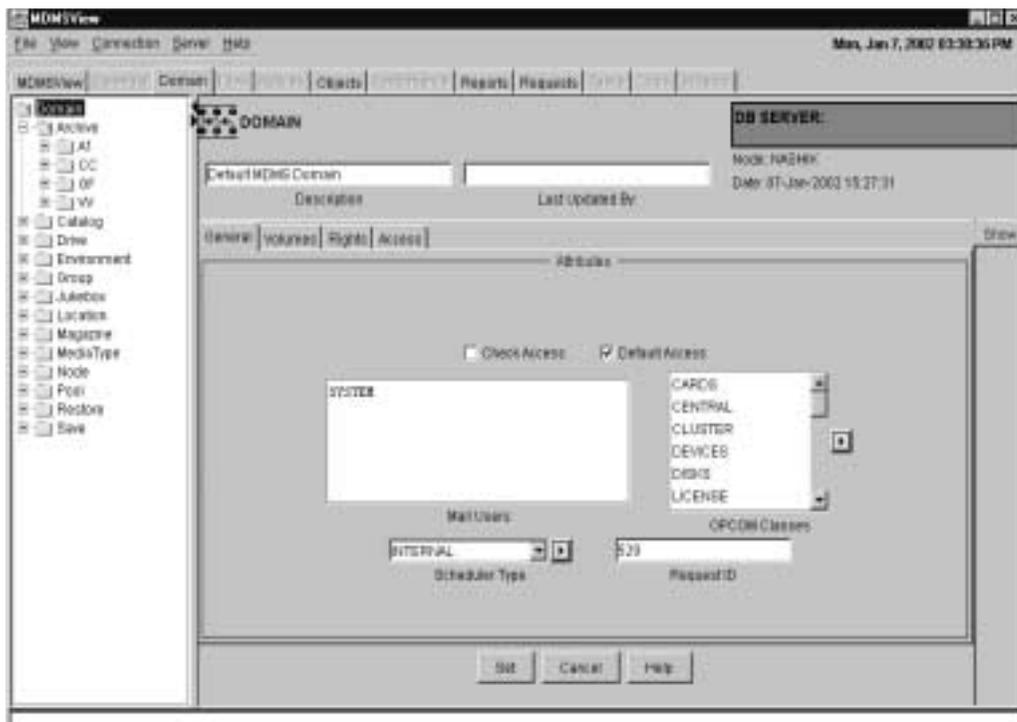
6.1.8 Viewing Relationships Between Objects

The Domain view provides a way to view the hierarchical structure of the MDMS domain. The left side of the screen provides an object-class-object... hierarchy of objects belonging to other objects, or objects contained in other objects. The left side of the screen displays most of the object classes which contain other objects (the exceptions: selections, schedules and volumes, which have no sub-objects). You can begin the hierarchical navigation at any level, and all sub-levels can be displayed.

For example, starting at jukebox, you can view all objects that reside in a jukebox: Drives, Magazines and Volumes. If you then click on Drives, you will see all drives in this jukebox. If you then select a drive and click on it, you can see the volume in the drive.

If your domain is sufficiently complex, you might want to expand the left side of the screen by using the right arrow between the left and right screen. You can then view the entire hierarchy of the domain.

Figure 6–5 Domain View Showing Expanded Relationships



6.1.9 Performing Operations on Objects

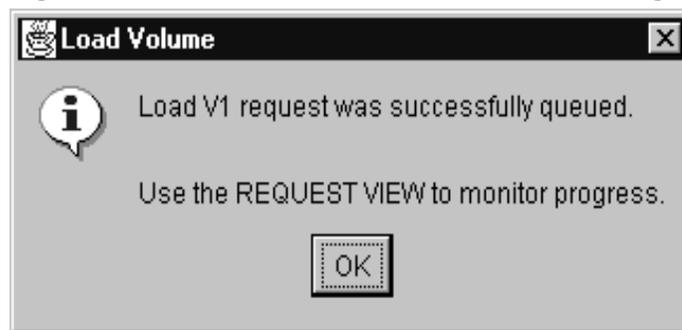
If you wish to perform an operation on an object (for example, to load a volume into a drive), you should first display the object's attributes and operations screens. Then select the desired operation tab, on the right side of the screen. For example, to load a volume, show the volume then click on the Load tab.

The load tab is called an operations tab, and they all follows the same basic concepts. You enter options concerning the operation (for example, operator assistance), then press the appropriate operation button on the bottom left of the screen. This button is always labelled with the appropriate operation (for example, Load).

MDMS has the capability of performing long-running operations synchronously or asynchronously. However, in MDMSView, long-running operations are always submitted asynchronously and control is returned to the user. Asynchronous operations show a dialog box that states that the operation has been queued for processing, but has not yet completed. If you perform an operation that does not result in the dialog box, then you can safely assume it has been completed synchronously.

If you receive a "queued" dialog box, it does not necessarily mean that the operation was fully validated. If you want to check on the status of the operation, use the Request View to monitor the request's progress.

Figure 6–6 Load Volume Screen with Queued Dialog Box



6.1.10 Running Save And Restore Requests

MDMS treats saves and restores in the same manner as other objects that it manages. You can create new saves and restores in the same way that you create other objects, then select a start time for them to run. Clicking Set will schedule the save or restore to run at the requested start date and time.

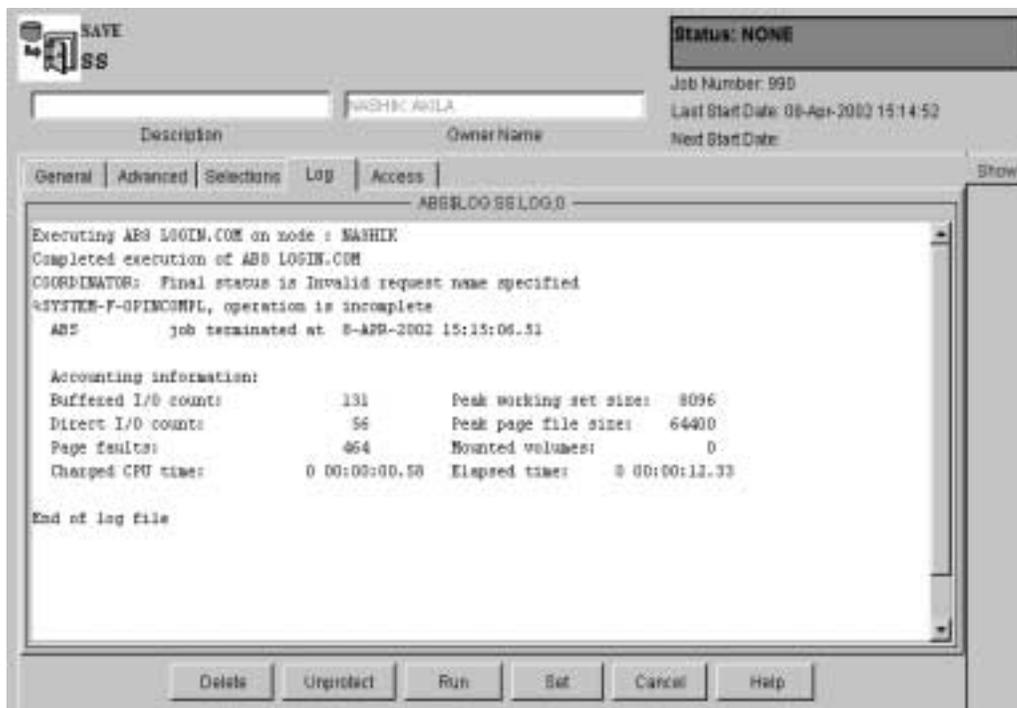
From MDMSView, however, there is an additional mechanism to run a save or restore. If you wish to run the request immediately, press the "Run" button at the bottom of the Show screen. This initiates an immediate run of the save or restore.

Once you run a save and restore request, you can monitor its progress by pressing the "Log" tab for the save or restore. This tab provides an up-to-date display of the request's log file.

User Interfaces

6.1 Graphical User Interface

Figure 6–7 Save Log Screen



6.1.11 Showing Current Operations

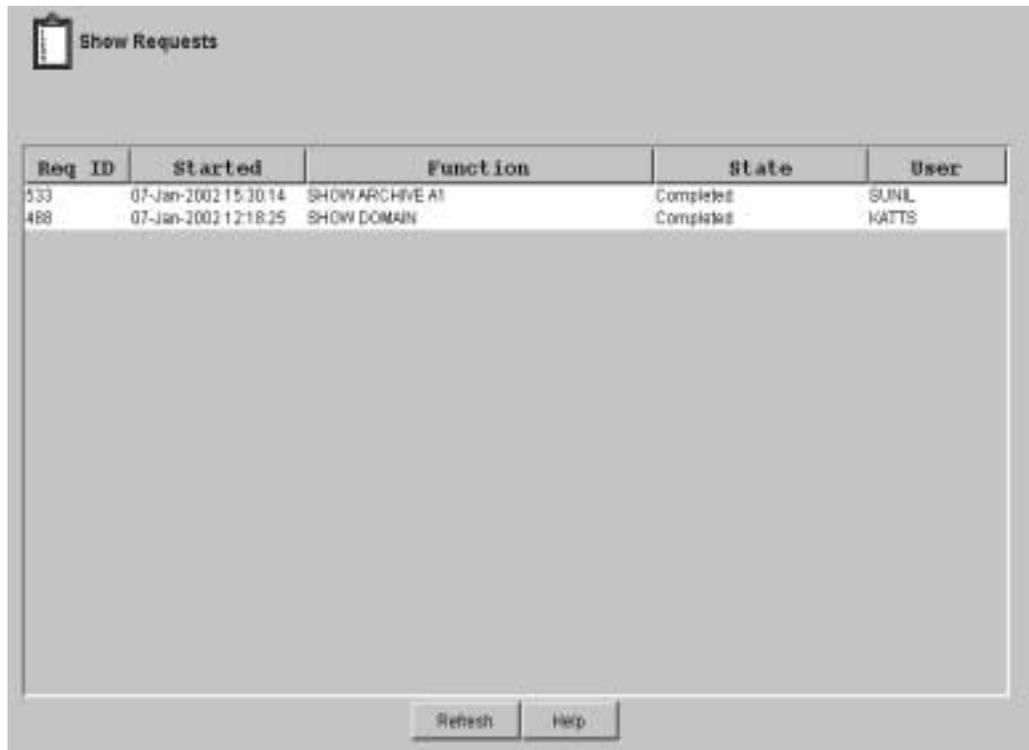
The Request View provides a monitoring capability for all current MDMS operations. You can display all current requests by clicking on Show Requests - this results in a table of requests being displayed. This includes all current requests, and some recently-completed requests.

You can also expand the requests on the left side of the screen and click on a specific request for detailed information about the request. Or you can right-click on the request number on the left screen and select Show.

If you feel that a request is not working correctly, or for any reason you wish to delete the request, you can click on delete from the detailed request screen, or select a request number on the left screen, right-click and select delete from the popup menu.

As with other deletes, a dialog box will appear to confirm the delete of the request.

Figure 6–8 Show Requests Screen



6.1.12 Reporting on Volumes

The Report View provides the capability of generating custom reports on volumes. With this view, you can choose attributes that can be displayed and/or used as selection criteria for volumes.

To select an attribute for display, simply click on the attribute and then press the right arrow button to move it to the display screen. The attributes are displayed in the report in the order selected. If you change your mind or wish to re-order the attributes, select an attribute on the display screen and press the left arrow button to deselect it.

If you wish to use an attribute as a selection criterion, click on the attribute, then click on “Use for Selection”. This will enable a field below (either a text field or combo box) to allow you to enter a selection.

You may display any number of fields and use any number of selection criteria to customize the report. When your selections are ready, you can generate the report by clicking on “Generate”. You can see the resultant report in the “Report Results” tab.

If you wish to save this report, enter a report title in the text field at the bottom of the screen and click on save. The report is saved to the following locations:

- OpenVMS Systems:
 - `sys$common:[mdms.gui.vms]Report_year_month_day_hour_minut_second.txt`
- Windows Systems:
 - `C:\MDMSView\Report_year_month_day_hour_minute_second.txt`

For example, a report file name is: `Report_2001_12_17_8_35_17.txt`

User Interfaces

6.1 Graphical User Interface

Once the results screen is displayed, you can sort the report using any field by clicking on the field's header. You can reverse-sort the same field by clicking on the field header again.

Figure 6–9 Report View Selection Criteria Example

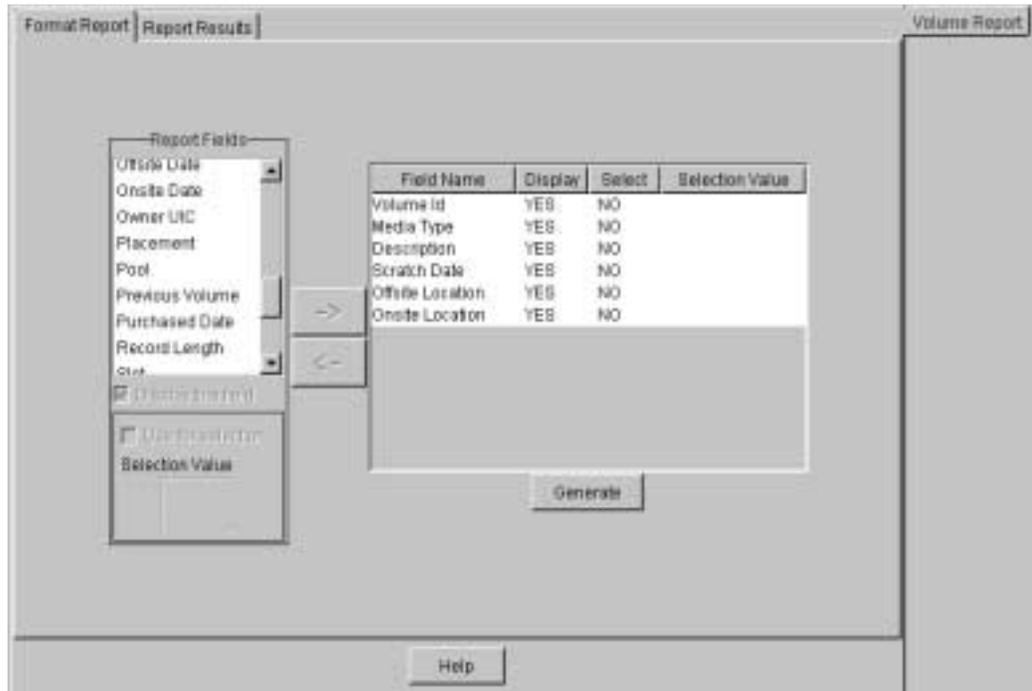
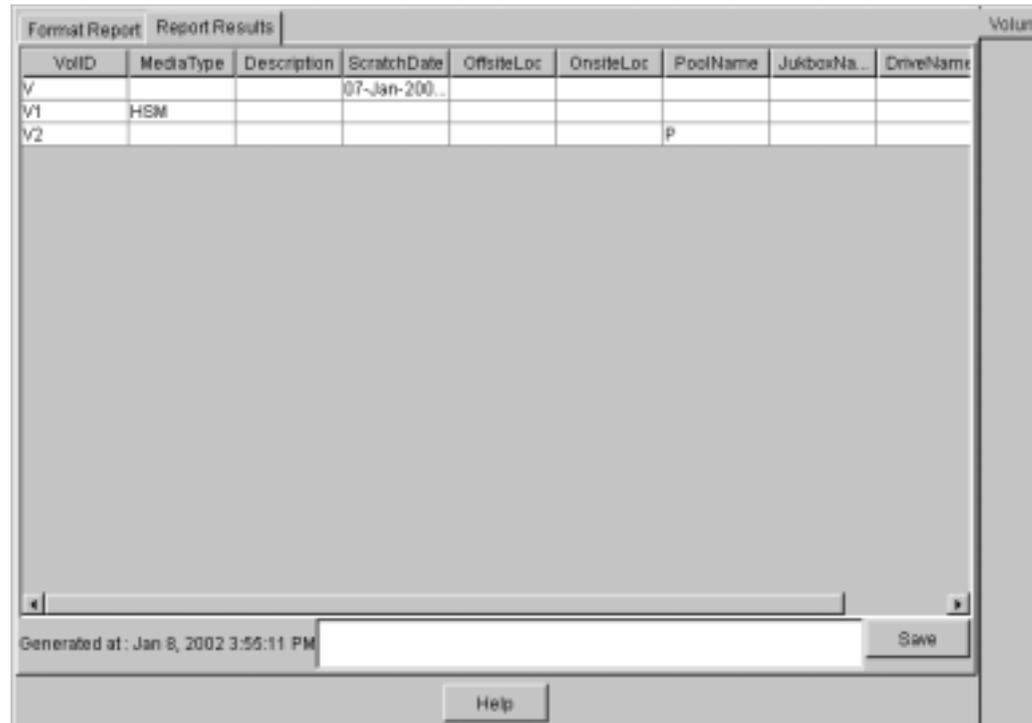


Figure 6–10 Report View Results Screen



6.1.13 Viewing MDMS Audit and Event Logging

To examine past operations in MDMS, you can use the event view to view the MDMS audit and event logfile. There are five pre-configured options and a fully flexible custom option to allow you to select what you wish to see from the MDMS logfile. The five pre-configured options all apply to the MDMS Database Server logfile and show all operations (auditing and events) for the following amounts of time before the current time:

- The last minute
- The last 10 minutes
- The last hour
- The last 24 hours
- The last 72 hours

If you wish to see the logfile using other selection criteria, you can use the “Custom” setting. By clicking on “Custom”, a selection screen appears that allows you to select the entries to be displayed as follows:

- *Node selection:* You can choose the default of the DB server (which contains the most complete information), or select a specific client node. Note that request IDs are not supported on client nodes, and nor is selection by low and high request IDs.
- *Selection Options:* You can select a range of entries in the logfile to display by one of:
 - Elapsed time in minutes (default of 60 minutes) -OR-
 - Before and/or since dates (specified as an absolute time) -OR-
 - Low and high request IDs (for DB server only)
- *Severity Options:* For audit completion entries, you can select that only entries of a certain combinations of completion status are displayed. You can select one or more of:
 - Success (S)
 - Informational (I)
 - Warning (W)
 - Error (E)
 - Fatal (F)

After entering the selection criteria, you click on the Show button to display. Depending on the size of the log file, this operation may take several seconds to complete. You may want to regularly reset your log files to avoid long response times. The code has been written to scan previous versions of log files if the date and or request selections are not in the latest log file.

The Refresh button at the bottom of the screen refreshes whatever selection is currently on the screen. The Cancel button allows you to enter a new selection.

6.1.14 Errors

MDMSView can report two types of errors:

- Those generated by MDMSView itself: these typically appear in a dialog box and in the status bar at the bottom of the screen. These errors normally explain an illegal user operation.
- Those generated by the MDMS server on the log-in node. These errors appear in a dialog box with the standard MDMS DCL syntax. These errors are documented in the *MDMS Reference Guide*.

User Interfaces

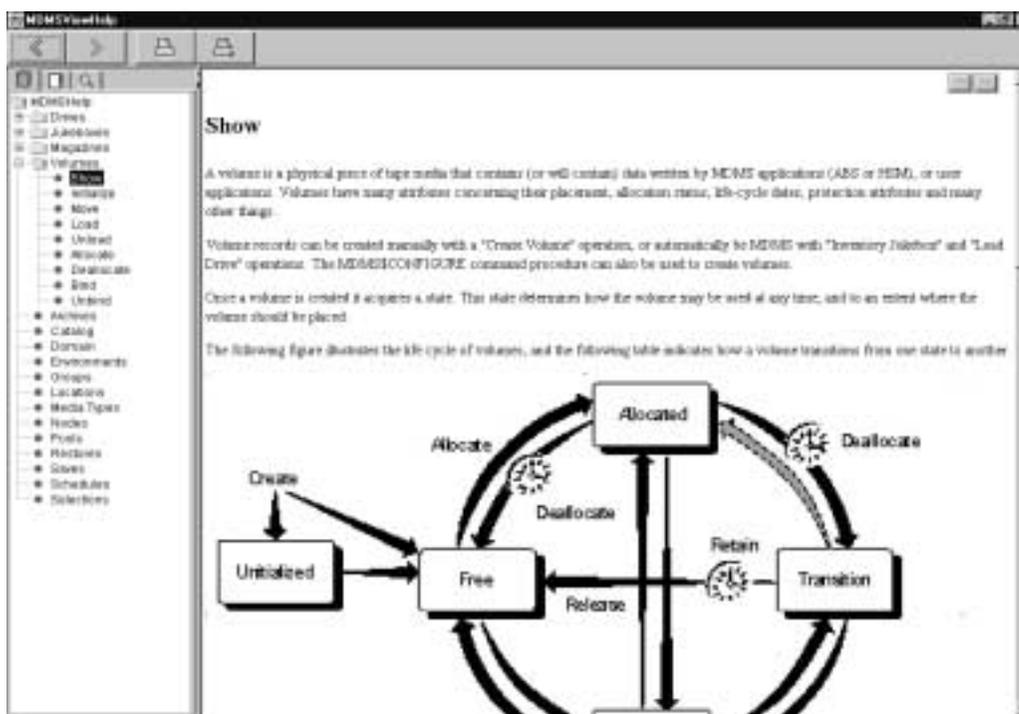
6.2 DCL Interface

6.1.15 Help

MDMSView provides three types of help:

- Tool-tip Help for every field on every screen. To obtain tool-tip help on a field, simply position the cursor on that field. The help appears near the field within one second and remains on the screen for 4 seconds.
- Screen-sensitive Help. For every screen in the Domain, Object or Task Views there is a “Help” button at the bottom right of each screen. If you press this “Help” button, a help screen pops-up with information about the screen you from which you pressed the button. The help information displayed is derived from this manual, the *ABS Guide to Operations*.
- Finally, there is a “Help” pull-down menu from the main screen. This provides the same type of Help as the “Help” button, but starts from the beginning of the manual. You can use the left-screen navigation or a search capability to find what you are interested in.

Figure 6–11 Context-Sensitive Help Screen from Show Volume Screen



6.2 DCL Interface

MDMS provides a DCL command line interface in addition to MDMSView. Some people prefer a command line interface, and it can also be used for automated command procedures. With this release, the entire command line interface is supported within MDMS, which maintains the database for both media management and ABS objects.

In previous releases, there was an ABS DCL interface that supported the ABS objects. This interface is now deprecated, but still works for backward compatibility. If you have command procedures that use this interface, they will still work. However, this interface will not be enhanced, so a migration to the MDMS DCL verbs is recommended for the long term.

6.2.1 Syntax Overview

The MDMS DCL interface uses a consistent syntax for virtually all commands in the format:

```
$ MDMS VERB OBJECT_KEYWORD OBJECT_NAME /QUALIFIERS
```

The verb is a simple action word, and may be one of the following:

- ALLOCATE
- BIND
- CREATE
- DEALLOCATE
- DELETE
- INITIALIZE
- INVENTORY
- LOAD
- SET
- SHOW
- UNBIND
- UNLOAD

The object keyword is the object class name that the verb is to operate on. In MDMS, the object keyword cannot be omitted. MDMS supports the following object keywords:

- ARCHIVE (ABS object)
- DOMAIN
- DRIVE
- ENVIRONMENT (ABS object)
- GROUP
- JUKEBOX
- LOCATION
- MAGAZINE
- MEDIA_TYPE
- NODE
- POOL
- RESTORE (ABS object)
- SAVE (ABS object)
- SERVER
- SCHEDULE
- SELECTION (ABS object)
- VERSION
- VOLUME

User Interfaces

6.2 DCL Interface

Following the object keyword, you should enter an object name. This must be the name of an already-existing object unless the verb is “Create”, in which case the object must not already exist.

The qualifiers for all commands are non-positional and may appear anywhere in the command line.

There are two exceptions to the general command syntax, as follows:

- The Move verb takes two arguments. The first is the object name as normal, and the second is a destination object name. The destination object name is not preceded by an object keyword. An example of a Move command is:

```
MDMS MOVE VOLUME TLZ234 TLZ_JUKE/SLOT=4
```

- The Report verb, which takes a variable number of arguments. This verb uses the syntax of the old SLS Storage Report Volume command. Since the Report verb does not operate on any specific object, the first argument is always the keyword “Volume”, and the other argument is a comma-separated list of display and/or selection attributes. For example:

```
$ MDMS REPORT VOLUME VOLUME,STATE=ALLOCATED,SCRATCH_DATE,PLACEMENT,PLACNAME
```

6.2.2 Object Lists

With this release of MDMS, all of the following commands accept a list of objects, so that you can operate on multiple objects in a single command:

- CREATE
- DELETE
- SET
- SHOW

If you specify an attribute in a CREATE or SET command and use an object list, then that attribute value is applied to all objects in the list.

6.2.3 Qualifier List

Certain qualifiers accept a list of attributes, and the list can be applied in one of three ways using an appropriate qualifier:

- /ADD - The specified value or list is added to any pre-existing list; this is the default option if you do not specify a qualifier
- /REMOVE - The specified value or list is removed from any pre-existing list
- /REPLACE - The specified value or list replaces any pre-existing list.

Consider the following examples:

```
MDMS CREATE GROUP COLORADO/NODES=(DENVER, SPRINGS, PUEBLO)
```

The group Colorado contains nodes Denver, Springs and Pueblo

```
MDMS SET GROUP COLORADO/NODE=ASPEN
```

The group Colorado now contains nodes Denver, Springs, Pueblo and Aspen. With no list qualifier specified, /ADD is applied by default.

```
MDMS SET GROUP COLORADO/NODE=ASPEN/REPLACE
```

The group Colorado now contains only node Aspen.

6.2.4 Inherit

All MDMS objects now accept the /INHERIT qualifier on Create. This allows you to create new objects and inherit most attributes of an existing object. This provides an easy way to “clone” objects, then apply the any differences in individual commands. It saves the effort of typing in all the attributes once a prototype has been established. In general, only non-protected fields of objects can be inherited.

In addition, the object list capability allows you to clone multiple objects in a single command. For example:

```
MDMS CREATE DRIVE DRIVE_2, DRIVE_3, DRIVE_4/INHERIT=DRIVE_1
```

This command creates three drives and applies all non-protected attributes of DRIVE_1 to the three new drives.

6.2.5 Symbols

MDMS now supports symbols on all objects which command procedures can read and process. To use symbols, enter a Show command for a single object (symbols are not supported for object lists). The symbols are generally in the format “MDMS_INQ_qualifier”, where “qualifier” is almost always the associated qualifier name for the attribute. The list of symbols for each show command is documented for that command, and is also available in DCL help.

When you issue a Show/Symbols, the show output is not displayed by default. If you wish to see the output as well, use Show/Symbols/Output.

6.2.6 Help and Reference

MDMS supports the normal DCL help mechanisms, as follows:

```
$ MDMS HELP [VERB] [KEYWORD] [/QUALIFIER]
```

```
$ HELP MDMS [VERB] [KEYWORD] [/QUALIFIER]
```

In addition, you can request help on any error message, for example:

```
MDMS HELP MESSAGE NOSUCHOBJECT
```

You can request help on any MDMS logical name, for example:

```
MDMS HELP LOGICAL MDMS$LOGFILTER
```

Finally, you can locate the mapping of the old (pre-version 4.0) ABS commands to the MDMS equivalent, for example:

```
MDMS HELP MAPPING CREATE ARCHIVE
```

The *MDMS Reference Guide* fully documents all DCL commands and qualifiers.

6.3 User Interface Restrictions

MDMSView and the MDMS DCL supports operations on Archive Backup System (ABS) objects only if an ABS or SLS license is loaded on the system. The ABS objects are:

- Archives
- Environments
- Restores
- Saves
- Selections

User Interfaces

6.3 User Interface Restrictions

MDMS supports operations on the other media management objects if the system only has a Hierarchical Storage Management (HSM) license installed, or with an ABS or SLS license.

In addition, if the ABS license is the restricted OMT license, the following operations are not supported:

- Creation of archives or environments
- Support of the Remote Device Facility
- Support of DCSC-controlled jukeboxes
- Support of external scheduler products
- Save/restore frequencies other than Daily_Full_Weekly, On_Demand and One_Time_Only

Preparing For Disaster Recovery

In case of a disaster you may need to restore all or part of the on-disk data of your computing environment. Basically you need a bootable system disk and a complete ABS/MDMS environment to restore all the rest of your data. This chapter explains the task to get you ABS environment up-and-running from scratch. The procedure differs slightly between OpenVMS and non-OpenVMS systems.

7.1 Disaster Recovery for OpenVMS Systems

To recover from a total loss of your online data you need the following items for recovery:

- a. An image copy of your system disk
- b. A copy of your MDMS\$ROOT including the database files
- c. A copy of your ABS\$ROOT including all catalog files
- d. A copy of files of any other product required by ABS, such as your 3rd party scheduler product

In all cases you need to keep the information about the saves in a safe place. This information is in the ABS save log and includes:

- a. The volume IDs of the tapes used
- b. The name of the savesets created
- c. The source path of the files being saved

Note

It is important to note the full pathname of the original location of the files.

You can print out this information from the epilog procedure of the environment.

7.1.1 Backup of Your System Disk

The easiest way to save your system disk is by using an ABS SAVE object like this:

Example 7–1 Save Object for Disaster Recovery of System Disk

```
Save: SYSTEM_DISK
Description: System Disk Backup for Recovery
Access Control: NONE
Owner: BONFYR::FROEHLIN
Archive: DISASTER_RECOVERY
Base Date: NONE
Delete Interval: NONE
Environment: DISASTER_RECOVERY_ENV
Epilogue:
```

Preparing For Disaster Recovery

7.1 Disaster Recovery for OpenVMS Systems

```
Execution Nodes: BONFYR
Explicit Interval:
Frequency: DAILY
Groups:
Incremental: NO
Job Number: 0
Prologue:
Schedule: SYSTEM_DISK_SAVE_SCHED
Sequence Option: SEQUENTIAL
Skip Time: NONE
Start Date: NONE
Transaction Status:
Selections: SYSTEM_DISK_SAVE_SEL_DEF

Default Selection -
- Data Select Type: VMS_FILES
  - Include: $1$DUA300:
  - Exclude:
- Source Node:
```

This SAVE uses the standard archive of DISASTER_RECOVERY and the standard environment of DISASTER_RECOVERY_ENV which comes with ABS. If these objects do not exist on your system run the ABS database initialization program:

```
$ RUN SYS$SYSTEM:ABS$DB_INIT
```

This program adds all the missing default ABS objects to the MDMS database.

Saving an OpenVMS system disk online produces many errors for files open for write by the operating system and layered products. Even though, the image backup produced can be used to restore a bootable system disk. The problem comes when executing the site-specific SYSTARTUP_VMS.COM. For example when starting the OpenVMS Queue Manager the command could hang because the Queue Manager files had been saved in an inconsistent state. There are three ways to avoid these kind of problems.

- Do a standalone backup of your system disk.

For Alpha systems see the section "Backing Up the System Disk" in the Appendix of the "Alpha Upgrade and Installation Manual" in the OpenVMS Documentation.

For VAX systems see the chapter "Using BACKUP" in the "System Manager's Manual" in the OpenVMS Documentation.

- Shutdown components of your system until all critical files are closed before starting the backup of your system disk. To find out which files are open for write use the following method:

```
$ BACKUP/IMAGE/IGNORE=INTERLOCK SYS$SYSDEVICE: NLA0:DUMMY.SAV/SAVE
```

Before the backup shutdown all components for which BACKUP reported:

```
%BACKUP-W-ACCONFLICT, SYS$SYSDEVICE:[SYS0.SYSCOMMON.SYSEXE]QMAN$MASTER.DAT;1 is open for write by another user
```

Shutdown of these components can be done in the prolog procedure in environment DISASTER_RECOVER_ENV. The same components can be automatically restarted in the epilog procedure.

- Ignore any error messages during the save operation. After restoring your system disk boot into conversational boot and rename your SYSTARTUP_VMS.COM and SYLOGICALS.COM to prevent any startup of extra components or layered products and reboot.

7.1.2 Backup of MDMS\$ROOT

Backing up MDMS\$ROOT with ABS will always find the MDMS database files open for write. This cannot be avoided. To copy the contents of these files in a consistent way online copies should be made prior to starting the save request. You can use the standard MDMS command procedure to do this:

```
$ @MDMS$SYSTEM:MDMS$COPY_DB_FILES
```

This command procedure uses DCL CONVERT/SHARE to create file copies with a file extension of “*.DAT_COPY”. This can be automatically done by executing this command procedure in the prolog of the environment. All files with that extension can then be automatically deleted in the epilog of the environment.

See ABS\$SYSTEM:ABS\$DISASTER_RECOVERY.TEMPLATE for an example.

If MDMS\$ROOT is not located on your system disk or you want a separate save operation, you can use a separate SAVE object like this:

Example 7–2 Save Object for Disaster Recovery of MDMS\$ROOT

```
Save: DISASTER_RECOVERY
Description:
Access Control: BONFYR::ABS (READ, WRITE, EXECUTE, DELETE, SET,
SHOW,
CONTROL)
Owner: BONFYR::ABS
Archive: DISASTER_RECOVERY
Base Date: NONE
Delete Interval: NONE
Environment: DISASTER_RECOVERY_ENV
Epilogue: @ABS$SYSTEM:ABS$DISASTER_RECOVERY.COM EPILOG
Execution Nodes: BONFYR
Explicit Interval:
Frequency: ON_DEMAND
Groups:
Incremental: NO
Job Number: 0
Prologue: @ABS$SYSTEM:ABS$DISASTER_RECOVERY.COM PROLOG
Schedule: DISASTER_RECOVERY_SAVE_SCHED
Sequence Option: SEQUENTIAL
Skip Time: NONE
Start Date: NONE
Transaction Status:
Selections: DISASTER_RECOVERY_SAVE_SEL_DEF

Default Selection -
- Data Select Type: VMS_FILES
  - Include: MDMS$ROOT:[000000...]*.*;*
  - Exclude: [...]*.LOG;*,[...]*_DB.DAT;*
- Source Node:
```

This save request excludes all the files open for write by MDMS and therefore does not create any error messages in the save log file.

If you want you can combine the save of the MDMS\$ROOT and the ABS\$ROOT into one save object.

Preparing For Disaster Recovery

7.1 Disaster Recovery for OpenVMS Systems

7.1.3 Backup of ABS\$ROOT

Backing up ABS\$ROOT with ABS will always find the ABS log files open for write because a save request always has a catalog file open for write. Not all catalog files might be accessible through ABS\$ROOT. For example if you have created a search list for ABS\$CATALOG and extensions to ABS\$CATALOG point to other directories and/or disk devices.

If ABS\$ROOT is not located on your system disk or you want a separate save operation, you can use a separate SAVE object like this:

Example 7-3 Save Object for Disaster Recovery of ABS\$ROOT

```
Save: DISASTER_RECOVERY
Description:
Access Control: BONFYR::ABS (READ, WRITE, EXECUTE, DELETE, SET,
SHOW,
CONTROL)
Owner: BONFYR::ABS
Archive: DISASTER_RECOVERY
Base Date: NONE
Delete Interval: NONE
Environment: DISASTER_RECOVERY_ENV
Epilogue: @ABS$SYSTEM:ABS$DISASTER_RECOVERY EPILOG
Execution Nodes: BONFYR
Explicit Interval:
Frequency: DAILY
Groups:
Incremental: NO
Job Number: 0
Prologue: @ABS$SYSTEM:ABS$DISASTER_RECOVERY PROLOG
Schedule: DISASTER_RECOVERY_SAVE_SCHED
Sequence Option: SEQUENTIAL
Skip Time: NONE
Start Date: NONE
Transaction Status:
Selections: DISASTER_RECOVERY_SAVE_SEL_DEF

Default Selection -
- Data Select Type: VMS_FILES
- Include: ABS$ROOT:[000000...]*.*;*
- Exclude: [... ]COORD_CLEANUP.DAT;*,[...]*.LOG;
- Source Node:
```

If this object does not exist on your system run the ABS database initialization program:

```
$ RUN SYS$SYSTEM:ABS$DB_INIT
```

This program adds all the missing default ABS objects to the MDMS database.

This save request excludes all the files open for write by ABS like the current logfile and the database file used by the coordinator cleanup process (“ABS\$COORD_CLEAN”). There should not be any error message in the save log file.

You have to make sure that you use the DISASTER_RECOVERY archived with an empty catalog name defined. The disaster recovery archive is the only archive which allows no catalog name. Otherwise you get open and verify errors for the catalog being used.

If you have an extended ABS\$CATALOG search list, then you have to include the extra entries in the include specification as well. By default the catalog subdirectory is included under ABS\$ROOT.

If you want, you can combine the save of the MDMS\$ROOT and the ABS\$ROOT into one save object.

7.2 Prolog and Epilog Procedure

To use ABS\$SYSTEM:ABS\$DISASTER_RECOVERY.TEMPLATE, you should rename it to ABS\$SYSTEM:ABS\$DISASTER_RECOVERY.COM and use it as a prolog and epilogue for the save(s). To automatically prepare the system for a save operation you can use the prolog and epilog feature in the environment object being used. The following example shows you how to use one procedure for both purposes.

Example 7-4 ABS\$SYSTEM:ABS\$DISASTER_RECOVERY.TEMPLATE.

```

$ !
$ ! Abstract:
$ !     This command file is used for saving ABS and MDMS information
$ !     for later disaster recovery. The procedure is used for prolog
$ !     as well as epilog procedures in a SAVE operation.
$ !
$ !
$ ! INPUT:
$ !
$ !     P1 = "" - no operation
$ !     = "PROLOG" - prepares a disaster recovery save operation
$ !                 by making online copies of MDMS database files
$ !     = "EPILOG" - does cleanup of save operation by deleting
$ !                 copies of files created during prolog
$ !                 - lists information about restoring the data
$ !
$ !-----
$ !
$ !
$! Start:
$!
$ SET NOON
$ IF P1.EQS."PROLOG" THEN GOTO Prolog
$ IF P1.EQS."EPILOG" THEN GOTO Epilog
$ EXIT
$!
$ Prolog:
$!
$ WRITE SYS$OUTPUT "Disaster Recovery Prolog"
$ WRITE SYS$OUTPUT "."
$ IF F$SEARCH("MDMS$DATABASE_LOCATION:MDMS$DOMAIN_DB.DAT").NES." "
$ THEN
$   WRITE SYS$OUTPUT "Creating online copies of MDMS database
files..."
$   WRITE SYS$OUTPUT "."
$   @MDMS$SYSTEM:MDMS$COPY_DB_FILES
$ ENDIF
$ EXIT
$!
$ Epilog:
$!
$ WRITE SYS$OUTPUT "."
$ IF
F$SEARCH("MDMS$DATABASE_LOCATION:MDMS$DOMAIN_DB.DAT_COPY").NES." "
$ THEN
$   WRITE SYS$OUTPUT "Deleting copies of MDMS database files..."
$   WRITE SYS$OUTPUT "."
$   DELETE/NOLOG MDMS$DATABASE_LOCATION:MDMS$*_DB.DAT_COPY;*
$ ENDIF
$ WRITE SYS$OUTPUT "BACKUP restore commands:"
$ WRITE SYS$OUTPUT "."
$ nmax = 'F$TRNLNM("ABS_OS_OBJECT_NUMBER")'
$ n     = 1

```

Preparing For Disaster Recovery

7.2 Prolog and Epilog Procedure

```

$!
$ NextObject:
$!
$ VolSetLog = "ABS_OS_VOLUME_SET_'n'"
$ VolRVNLog = "ABS_OS_START_RVN_'n'"
$ ObjectLog = "ABS_OS_OBJECT_SET_'n'"
$ SavsetLog = "ABS_OS_SAVESET_NAME_'n'"
$ CALL GetVolumeList "'F$TRNLNM(VolSetLog)'" 'F$TRNLNM(VolRVNLog)'
VolumeList
$ Destination = "'F$TRNLNM(ObjectLog)'"
$ Destination = F$EXTRACT(0,F$LOCATE(":",Destination),Destination)
$ Destination = "'F$TRNLNM(Destination)'"
$ IF F$LOCATE(".",Destination).NE.F$LENGTH(Destination)
$ THEN
$   Destination = Destination + "[...]" - "]"["
$ ELSE
$   Destination = Destination + "[*...]"
$ ENDIF
$ WRITE SYS$OUTPUT "      $ BACKUP/OVERLAY/EXACT_ORDER/NOASSIST -"
$ WRITE SYS$OUTPUT "    _$  tape:", "'F$TRNLNM(Savset-
Log)'/LABEL=(",VolumeList,
$ WRITE SYS$OUTPUT "    _$  "'Destination'"
$ WRITE SYS$OUTPUT "."
$ n = n + 1
$ IF n.LE.nmax THEN GOTO NextObject
$ WRITE SYS$OUTPUT "."
$ WRITE SYS$OUTPUT "After restoring the savesets rename the MDMS
database"
$ WRITE SYS$OUTPUT "files from "MDMS$_DB.DAT_COPY" to
"* .DAT"/NEW_VERSION."
$ WRITE SYS$OUTPUT "."
$ EXIT
$!
$ GetVolumeList: SUBROUTINE
$!
$ VolumeID   = "'P1'"
$ RVN        = 'P2'
$ 'P3'       == ""
$ VolumeRVN  = 1
$!
$ NextVolume:
$!
$ MDMS SHOW VOLUME 'VolumeID'/SYMBOL
$ IF VolumeRVN.GE.RVN
$ THEN
$   IF 'P3'.NES."" THEN 'P3' == 'P3' + ","
$   'P3' == 'P3' + "'MDMS_INQ_VOLUME_ID'"
$ ENDIF
$ IF "'MDMS_INQ_NEXT_VOLUME'.EQS."" THEN GOTO EndVolumeList
$ VolumeID = "'MDMS_INQ_NEXT_VOLUME'"
$ VolumeRVN = VolumeRVN + 1
$ GOTO NextVolume
$!
$ EndVolumeList:
$!
$ EXIT

```

The example procedure creates copies of the MDMS database files in the prolog phase. This allows to save the files in a consistent state. After a restore from the saveset the files need to be renamed to their original names.

For convenience the procedure prints out the backup commands needed to restore the data using information in logical names defined by ABS during the save operation.

7.2.1 Restoring The System Disk

To restore your system disk you need to use Standalone BACKUP.

- For Alpha systems see the section "Backing Up the System Disk" in the Appendix of the "Alpha Upgrade and Installation Manual" in the OpenVMS Documentation.
- For VAX systems see the chapter "Using BACKUP" in the "System Manager's Manual" in the OpenVMS Documentation.

Use the information from the ABS save log to specify the parameters for the BACKUP command line:

- a. /LABEL=(volume_1,volume_2,...volume_n) - the volume IDs of the tapes being used
- b. The saveset name
- c. The target disk
- d. /IMAGE/NOASSIST qualifiers

Example 7-5 BACKUP Command to Restore the System Disk

```
$ BACKUP/IMAGE MKA500:24DEC20012359590./LABEL=(GKF011,GKF022) -  
_DGA100:/NOASSIST
```

This restores an image of your system disk in saveset "24DEC20012359590." on tape volumes "GKF011" and "GKF022" to disk device "DGA100".

After a successful restore, boot from your restored system disk. If your system does not boot all the way through you may have to disable the execution of your "SYSTARTUP_VMS.COM" command procedure by using a conversational boot and renaming the file.

7.2.2 Restoring Remaining Savesets

Once your system is up-and-running you can restore other save sets necessary to complete the disaster recovery: Make sure that all of these components or products are shut down before you restore the individual files. Use the following restore order:

1. First, any other product required by ABS, such as your 3rd party scheduler data if it has been saved separately. You should startup the component or product just restored.
2. Restore MDMS\$ROOT if it has been saved separately. After the restore rename the "MDMS\$*_DB.DAT_COPY" files to "*_DB.DAT". You can startup MDMS now.
3. Restore ABS\$ROOT if it has been saved separately. Restore any other save used to save the catalog files which are located outside of ABS\$ROOT. After the restore you can startup ABS.

Use the information from the ABS save log to specify the parameters for the BACKUP command lines:

- a. /LABEL=(volume_1,volume_2,...volume_n) - the volume IDs of the tapes being used
- b. The saveset name
- c. The target disk
- d. /IMAGE/NOASSIST qualifiers

Preparing For Disaster Recovery

7.3 Non-OpenVMS Systems

Example 7–6 BACKUP Command to Restore ABS\$ROOT

```
$ BACKUP/NOASSIST/OVERLAY -  
$_MKA500:25DEC20010101010./LABEL=(GKF033,GKF044)-  
$_DGA100:[VMS$COMMON.ABS.*...]/LOG
```

Because ABS has not been started up the ABS\$ROOT logical is not available yet. This restores the ABS\$ROOT files in saveset “24DEC20012359590.” on tape volumes “GKF033” and “GKF044” to disk location “DGA100:[VMS\$COMMON.ABS...]”. This assumes that you ABS\$ROOT logical was defined as a concealed device name of “DGA100:[VMS\$COMMON.ABS.]”.

Note

It is important to note the full pathname when you save these components or products.

7.3 Non-OpenVMS Systems

ABS cannot restore a bootable system disk of a non-OpenVMS system. Therefore you need to be able to save and restore the system disk locally. Once you have the system disk restored and booted the system you have to install the ABS client software for that platform. Once the ABS client software has been installed you can use ABS on your OpenVMS system to restore data to the client node.

7.4 Thoughts on Save and Restore Procedures

When it comes to setup procedures on how to save and restore files for disaster recovery there is a variety of possibilities depending on your configuration and other system activities.

You do not need to have an up-to-date copy of your system disk to restore your ABS environment. You could start with a fresh installation of OpenVMS. Install ABS and products required to run ABS (e.g. a 3rd party scheduler). While ABS is shutdown restore the MDM\$ROOT and ABS\$ROOT and other required components. Startup ABS to restore all the rest of your data.

Or in a VMSCluster with more than one system disk you may be able to restore all your data online using ABS from another node in the cluster.

You can keep a printout of the ABS save log in a safe place. This allows you to restore the data for files on OpenVMS systems using OpenVMS BACKUP. You need to keep the volume IDs, the name of the save sets and the include specifications used in the save operation.

Typically you do not want to keep multiple copies of your disaster recovery saves. You may want to keep 2 copies. So, if you are doing daily disaster recovery saves the archive expiration should be set to 2 days.

You should use non-incremental saves for the disaster recovery. This allows for an easy restore in case of an emergency. You can use incremental saves, but on a restore you have to do all the incremental restores on your own until you have ABS fully up-and-running

Note

And a final word: Make sure that you have a clear procedure on how to do a disaster recovery. Test your disaster recovery procedure!

Remote Devices

This chapter explains how to configure and manage remote devices using the Remote Device Facility (RDF). RDF is used for devices remotely connected over a wide-area network, and DECnet is still a requirement for access to these remote devices. RDF is not required for devices connected remotely via Fibre Channel, as these are considered local devices.

8.1 The RDF Installation

When you install ABS (non-standard installation) or MDMS, you are asked whether you want to install the RDF software. With the ABS standard installation, the RDF client and server software is installed by default.

During the installation you place the RDF client software on the nodes with disks you want to access for ABS or HSM. You place the RDF server software on the systems to which the tape devices (jukeboxes and drives) are connected. This means that when using RDF, you serve the tape device to the systems with the client disks.

All of the files for RDF are placed in `SY$COMMON:[MDMS.TTI_RDF]` for your system. There are separate locations for VAX or Alpha.

Note

RDF is not available if you are running ABS/MDMS with the ABS-OMT license.

8.2 Configuring RDF

After installing RDF you should check the `TTI_RDEV:CONFIG_nodename.DAT` file to make sure it has correct entries.

This file:

- is located on the RDF server node with the tape device
- is created initially during installation
- is a text file
- includes the definition of each device accessible by the RDF software. This definition consists of a physical device name and an RDF characteristic name.

Example:

```
Device $1$MIA0 MIA0
```

Verify:

Check this file to make sure that all RDF characteristic names are unique to this node.

Remote Devices

8.3 Using RDF with MDMS

8.3 Using RDF with MDMS

The following sections describe how to use RDF with MDMS.

8.3.1 Starting Up and Shutting Down RDF Software

Starting up RDF software:

RDF software is automatically started up along with then MDMS software when you enter the following command:

```
$ @SYS$STARTUP:MDMS$STARTUP
```

Shutting down RDF software:

To shut down the RDF software, enter the following command:

```
$ @SYS$STARTUP:MDMS$SHUTDOWN
```

8.3.2 The RDSHOW Procedure

Required privileges:

The following privileges are required to execute the RDSHOW procedure: NETMBX, TMP-MBX.

In addition, the following privileges are required to show information on remote devices allocated by other processes: SYSPRV, WORLD.

8.3.3 Command Overview

You can run the RDSHOW procedure any time after the MDMS software has been started. RDF software is automatically started at this time.

Use the following procedures:

```
$ @TTI_RDEV:RDSHOW CLIENT
$ @TTI_RDEV:RDSHOW SERVER node_name
$ @TTI_RDEV:RDSHOW DEVICES
```

node_name is the node name of any node on which the RDF server software is running.

8.3.4 Showing Your Allocated Remote Devices

To show remote devices that you have allocated, enter the following command from the RDF client node:

```
$ @TTI_RDEV:RDSHOW CLIENT
```

Result:

```
RDALLOCATED devices for pid 20200294, user DJ, on node OMAHA::
Local logical      Rmt node  Remote device
TAPE01             MIAMI::   MIAMI$MUCO
```

DJ is the user name and OMAHA is the current RDF client node.

8.3.5 Showing Available Remote Devices on the Server Node

The RDSHOW SERVER procedure shows the available devices on a specific SERVER node. To execute this procedure, enter the following command from any RDF client or RDF server node:

```
$ @TTI_RDEV:RDSHOW SERVER MIAMI
```

MIAMI is the name of the server node whose devices you want shown.

Result:

```
Available devices on node MIAMI::
Name          Status Characteristics/Comments
MIAMI$MSA0    in use  msa0
...by pid 20200246, user CATHY (local)
MIAMI$MUA0    in use  mua0
...by pid 202001B6, user CATHY, on node OMAHA::
MIAMI$MUB0    -free-  mub0
MIAMI$MUC0    in use  muc0
...by pid 2020014C, user DJ, on node OMAHA::
```

This RDSHOW SERVER command shows any available devices on the server node MIAMI, including any device characteristics. In addition, each allocated device shows the process PID, username, and RDF client node name.

The text (local) is shown if the device is locally allocated.

8.3.6 Showing All Remote Devices Allocated on the RDF Client Node

To show all allocated remote devices on an RDF client node, enter the following command from the RDF client node:

```
$ @TTI_RDEV:RDSHOW DEVICES
```

Result:

```
Devices RDALLOCATED on node OMAHA::
RDdevice  Rmt node  Remote device  User name      PID
RDEVA0:   MIAMI::    MIAMI$MUC0    DJ             2020014C
RDEVB0:   MIAMI::    MIAMI$MUA0    CATHY         202001B6
```

This command shows all allocated devices on the RDF client node OMAHA. Use this command to determine which devices are allocated on which nodes.

8.4 Monitoring and Tuning Network Performance

This section describes network issues that are especially important when working with remote devices.

8.4.1 DECnet Phase IV

The Network Control Program (NCP) is used to change various network parameters. RDF (and the rest of your network as a whole) benefits from changing two NCP parameters on all nodes in your network. These parameters are:

- PIPELINE QUOTA
- LINE RECEIVE BUFFERS

Pipeline quota

The pipeline quota is used to send data packets at an even rate. It can be tuned for specific network configurations. For example, in an Ethernet network, the number of packet buffers represented by the pipeline quota can be calculated as approximately:

```
buffers = pipeline_quota / 1498
```

Default:

The default pipeline quota is 10000. At this value, only six packets can be sent before acknowledgment of a packet from the receiving node is required. The sending node stops after the sixth packet is sent if an acknowledgment is not received.

Remote Devices

8.4 Monitoring and Tuning Network Performance

Recommendation:

The PIPELINE QUOTA can be increased to 45,000 allowing 30 packets to be sent before a packet is acknowledged (in an Ethernet network). However, performance improvements have not been verified for values higher than 23,000. It is important to know that increasing the value of PIPELINE QUOTA improves the performance of RDF, but may negatively impact performance of other applications running concurrently with RDF.

Line receive buffers

Similar to the pipeline quota, line receive buffers are used to receive data at a constant rate.

Default:

The default setting for the number of line receive buffers is 6.

Recommendation:

The number of line receive buffers can be increased to 30 allowing 30 packets to be received at a time. However, performance improvements have not been verified for values greater than 15 and as stated above, tuning changes may improve RDF performance while negatively impacting other applications running on the system.

8.4.2 DECnet-Plus (Phase V)

As stated in DECnet-Plus(Phase V), (DECnet/OSI V6.1) Release Notes, a pipeline quota is not used directly. Users may influence packet transmission rates by adjusting the values for the transport's characteristics MAXIMUM TRANSPORT CONNECTIONS, MAXIMUM RECEIVE BUFFERS, and MAXIMUM WINDOW. The value for the transmit quota is determined by MAXIMUM RECEIVE BUFFERS divided by Actual TRANSPORT CONNECTIONS.

This will be used for the transmit window, unless MAXIMUM WINDOW is less than this quota. In that case, MAXIMUM WINDOW will be used for the transmitter window.

The DECnet-Plus defaults (MAXIMUM TRANSPORT CONNECTIONS = 200 and MAXIMUM RECEIVE BUFFERS = 4000) produce a MAXIMUM WINDOW of 20. Decreasing MAXIMUM TRANSPORT CONNECTIONS with a corresponding increase of MAXIMUM WINDOW may improve RDF performance, but also may negatively impact other applications running on the system.

8.4.3 Changing Network Parameters

This section describes how to change the network parameters for DECnet Phase IV and DECnet-PLUS.

8.4.4 Changing Network Parameters for DECnet (Phase IV)

The pipeline quota is an NCP executor parameter. The line receive buffers setting is an NCP line parameter.

The following procedure shows how to display and change these parameters in the permanent DECnet database. These changes should be made on each node of the network.

Table 8–1 How to Change Network Parameters

Step	Action
1	Enter: <pre>\$ run sys\$system:NCP NCP>show executor characteristics</pre> Result: <pre>Node Permanent Characteristics as of 24-MAY-1991 10:10:58 Executor node = 20.1 (DENVER) Management version = V4.0.0 . . . Pipeline quota = 10000</pre>
2	Enter: <pre>NCP>define executor pipeline quota 45000 NCP>show known lines</pre> Result: <pre>Known line Volatile Summary as of 24-MAY-1991 10:11:13 Line State SVA-0 on</pre>
3	Enter: <pre>NCP>show line sva-0 characteristics</pre> Result: <pre>Line Permanent Characteristics as of 24-MAY-1991 10:11:31 Line = SVA-0 Receive buffers = 6 <-- value to change Controller = normal Protocol = Ethernet Service timer = 4000 Hardware address = 08-00-2B-0D-D0-5F Device buffer size = 1498</pre>
4	Enter: <pre>NCP>define line sva-0 receive buffers 30 NCP>exit</pre>

Requirement:

For the changed parameters to take effect, the node must be rebooted or DECnet must be shut down.

Remote Devices

8.4 Monitoring and Tuning Network Performance

8.4.5 Changing Network Parameters for DECnet-Plus(Phase V)

The Network Control Language (NCL) is used to change DECnet-Plus network parameters. The transport parameters MAXIMUM RECEIVE BUFFERS, MAXIMUM TRANSPORT CONNECTIONS and MAXIMUM WINDOW can be adjusted by using NCL's SET OSI TRANSPORT command. For example:

```
NCL> SET OSI TRANSPORT MAXIMUM RECEIVE BUFFERS = 4000      !default value
NCL> SET OSI TRANSPORT MAXIMUM TRANSPORT CONNECTIONS = 200 !default value
NCL> SET OSI TRANSPORT MAXIMUM WINDOWS = 20                !default value
```

To make the parameter change permanent, add the NCL command(s) to the SYSS\$MANAGER:NET\$OSI_TRANSPORT_STARTUP.NCL file. Refer to the DENET-Plus (DECnet/OSI) Network Management manual for detailed information.

8.4.6 Resource Considerations

Changing the default values of line receive buffers and the pipeline quota to the values of 30 and 45000 consumes less than 140 pages of nonpaged dynamic memory.

In addition, you may need to increase the number of large request packets (LRPs) and raise the default value of NETACP BYTLM.

Large request packets

LRPs are used by DECnet to send and receive messages. The number of LRPs is governed by the SYSGEN parameters LRPCOUNT and LRPCOUNTV.

Recommendation:

A minimum of 30 free LRPs is recommended during peak times. Show these parameters and the number of free LRPs by entering the following DCL command:

```
$ SHOW MEMORY/POOL/FULL
```

Result:

```
System Memory Resources on 24-JUN-1991 08:13:57.66
Large Packet (LRP) Lookaside List Packets Bytes
Current Total Size                36      59328
Initial Size (LRPCOUNT)           25      41200
Maximum Size (LRPCOUNTV)          200     329600
Free Space                         20      32960
```

In the LRP lookaside list, this system has:

- Current Total Size of 36

The SYSGEN parameter LRPCOUNT (LRP Count) has been set to 25. The Current Size is not the same as the Initial Size. This means that OpenVMS software has to allocate more LRPs. This causes system performance degradation while OpenVMS is expanding the LRP lookaside list.

The LRPCOUNT should have been raised to at least 36 so OpenVMS does not have to allocate more LRPs.

Recommendation:

Raise the LRPCOUNT parameter to a minimum of 50. Because the LRPCOUNT parameter is set to only 25, the LRPCOUNT parameter is raised on this system even if the current size was also 25.

8.4 Monitoring and Tuning Network Performance

- Free Space is 20

This is below the recommended free space amount of 30. This also indicates that LRPCOUNT should be raised. Raising LRPCOUNT to 50 (when there are currently 36 LRPs) has the effect of adding 14 LRPs. Fourteen plus the 20 free space equals over 30. This means that the recommended value of 30 free space LRPs is met after LRPCOUNT is set to 50.

- The SYSGEN parameter LRPCOUNTV (LRP count virtual) has been set to 200.

The LRPCOUNTV parameter should be at least four times LRPCOUNT. Raising LRPCOUNT may mean that LRPCOUNTV has to be raised. In this case, LRPCOUNTV does not have to be raised because 200 is exactly four times 50 (the new LRPCOUNT value).

Make changes to LRPCOUNT or LRPCOUNTV in both:

- SYSGEN (using CURRENT)
- SYS\$SYSTEM:MODPARAMS.DAT file (for when AUTOGEN is run with REBOOT)

Example: Changing LRPCOUNT to 50 in SYSGEN

```
Username: SYSTEM
Password: (the system password)
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> USE CURRENT
SYSGEN> SH LRPCOUNT
Parameter Name           Current   Default   Minimum   Maximum
LRPCOUNT                 25       4         0         4096
SYSGEN> SET LRPCOUNT 50
SYSGEN> WRITE CURRENT
SYSGEN> SH LRPCOUNT
Parameter Name           Current   Default   Minimum   Maximum
LRPCOUNT                 50       4         0         4096
```

Requirement:

After making changes to SYSGEN, reboot your system so the changes take effect.

Example: Changing the LRPCOUNT for AUTOGEN

Add the following line to MODPARAMS.DAT:

```
$ MIN_LRPCOUNT = 50      ! ADDED {the date} {your initials}
```

Result:

This ensures that when AUTOGEN runs, LRPCOUNT is not set below 50.

NETACP BYTLM

The default value of NETACP is a BYTLM setting of 65,535. Including overhead, this is enough for only 25 to 30 line receive buffers. This default BYTLM may not be enough.

Recommendation:

Increase the value of NETACP BYTLM to 110,000.

How to increase NETACP BYTLM:

Before starting DECnet, define the logical NETACP\$BUFFER_LIMIT by entering:

```
$ DEFINE/SYSTEM/NOLOG NETACP$BUFFER_LIMIT 110000
$ @SYS$MANAGER:STARTNET.COM
```

Remote Devices

8.4 Monitoring and Tuning Network Performance

8.4.7 Controlling RDF's Effect on the Network

By default, RDF tries to perform I/O requests as fast as possible. In some cases, this can cause the network to slow down. Reducing the network bandwidth used by RDF allows more of the network to become available to other processes.

The RDF logical names that control this are:

```
RDEV_WRITE_GROUP_SIZE
RDEV_WRITE_GROUP_DELAY
```

Default:

The default values for these logical names is zero. The following example shows how to define these logical names on the RDF client node:

```
$ DEFINE/SYSTEM RDEV_WRITE_GROUP_SIZE 30
$ DEFINE/SYSTEM RDEV_WRITE_GROUP_DELAY 1
```

Further reduction:

To further reduce bandwidth, the RDEV_WRITE_GROUP_DELAY logical can be increased to two (2) or three (3).

Note

Reducing the bandwidth used by RDF causes slower transfers of RDF's data across the network.

8.4.8 Surviving Network Failures

Remote Device Facility (RDF) can survive network failures of up to 15 minutes long. If the network comes back within the 15 minutes allotted time, the RDCLIENT continues processing WITHOUT ANY INTERRUPTION OR DATA LOSS. When a network link drops while RDF is active, after 10 seconds, RDF creates a new network link, synchronizes I/Os between the RDCLIENT and RDSERVER, and continues processing.

The following example shows how you can test the RDF's ability to survive a network failure. (This example assumes that you have both the RDSERVER and RDCLIENT processes running.)

```
$ @tti_rdev:rdallocate tti::mua0:
RDF - Remote Device Facility (Version 4.1) - RDALLOCATE Procedure
Copyright (c) 1990, 1996 Touch Technologies, Inc.
Device TTI::TTI$MUA0 ALLOCATED, use TAPE01 to reference it
$ backup/rewind/log/ignore=label sys$library:*.* tape01:test
```

from a second session:

```
$ run sys$system:NCP
NCP> show known links

Known Link Volatile Summary as of 13-MAR-1996 14:07:38
Link      Node          PID      Process      Remote link  Remote user
24593    20.4 (JR)      2040111C MARI_11C_5   8244        CTERM
16790    20.3 (FAST)    20400C3A -rdclient-   16791      tti_rdevSRV
24579    20.6 (CHEERS)  20400113 REMACP       8223        SAMMY
24585    20.6 (CHEERS)  20400113 REMACP       8224        ANDERSON
NCP> disconnect link 16790
.
.
.
```

Remote Devices

8.5 Controlling Access to RDF Resources

Backup pauses momentarily before resuming. Sensing the network disconnect, RDF creates a new `-rdclient-` link. Verify this by entering the following command:

```
NCP> show known links
Known Link Volatile Summary as of 13-MAR-1996 16:07:00

Link      Node          PID      Process      Remote link  Remote user
24593    20.4 (JR)      2040111C MARI_11C_5    8244         CTERM
24579    20.6 (CHEERS)  20400113 REMACP        8223         SAMMY
24585    20.6 (CHEERS)  20400113 REMACP        8224         ANDERSON
24600    20.3 (FAST)    20400C3A -rdclient-    24601        tti_rdevSRV

NCP> exit
```

8.5 Controlling Access to RDF Resources

The RDF Security Access feature allows storage administrators to control which remote devices are allowed to be accessed by RDF client nodes.

8.5.1 Allow Specific RDF Clients Access to All Remote Devices

You can allow specific RDF client nodes access to all remote devices.

Example:

For example, if the server node is MIAMI and access to all remote devices is granted only to RDF client nodes OMAHA and DENVER, then do the following:

1. Edit `TTI_RDEV:CONFIG_MIAMI.DAT`
2. Before the first device designation line, insert the `/ALLOW` qualifier

```
Edit TTI_RDEV:CONFIG_MIAMI.DAT
CLIENT/ALLOW=(OMAHA,DENVER)
DEVICE $1$MUA0: MUA0, TK50
DEVICE MSA0: TU80, 1600bpi
```

OMAHA and DENVER (the specific RDF CLIENT nodes) are allowed access to all remote devices (MUA0, TU80) on the server node MIAMI.

Requirements:

If there is more than one RDF client node being allowed access, separate the node names by commas.

8.5.2 Allow Specific RDF Clients Access to a Specific Remote Device

You can allow specific RDF client nodes access to a *specific* remote device.

Example:

If the server node is MIAMI and access to MUA0 is allowed by RDF client nodes OMAHA and DENVER, then do the following:

1. Edit `TTI_RDEV:CONFIG_MIAMI.DAT`
2. Find the device designation line (for example, `DEVICE 1MUA0:`)
3. At the end of the device designation line, add the `/ALLOW` qualifier:

```
$ Edit TTI_RDEV:CONFIG_MIAMI.DAT
DEVICE $1$MUA0: MUA0, TK50/ALLOW=(OMAHA,DENVER)
DEVICE MSA0: TU80, 1600bpi
```

OMAHA and DENVER (the specific RDF client nodes) are allowed access only to device MUA0. In this situation, OMAHA is not allowed to access device TU80.

Remote Devices

8.6 RDserver Inactivity Timer

8.5.3 Deny Specific RDF Clients Access to All Remote Devices

You can deny access from specific RDF client nodes to all remote devices. For example, if the server node is MIAMI and you want to deny access to all remote devices from RDF client nodes OMAHA and DENVER, do the following:

1. Edit TTI_RDEV:CONFIG_MIAMI.DAT
2. Before the first device designation line, insert the /DENY qualifier:

```
$ Edit TTI_RDEV:CONFIG_MIAMI.DAT
CLIENT/DENY=( OMAHA ,DENVER )
DEVICE $1$MUA0:  MUA0, TK50
DEVICE MSA0:      TU80, 16700bpi
```

OMAHA and DENVER are the specific RDF client nodes denied access to all the remote devices (MUA0, TU80) on the server node MIAMI.

8.5.4 Deny Specific RDF Clients Access to a Specific Remote Device

You can deny specific client nodes access to a specific remote device.

Example:

If the server node is MIAMI and you want to deny access to MUA0 from RDF client nodes OMAHA and DENVER, do the following:

1. Edit TTI_RDEV:CONFIG_MIAMI.DAT
2. Find the device designation line (for example, DEVICE \$1\$MUA0:)
3. At the end of the device designation line, add the /DENY qualifier:

```
$ Edit TTI_RDEV:CONFIG_MIAMI.DAT
DEVICE $1$MUA0:  MUA0, TK50/DENY=( OMAHA ,DENVER )
DEVICE MSA0:      TU80, 16700bpi
```

OMAHA and DENVER RDF client nodes are denied access to device MUA0 on the server node MIAMI.

8.6 RDserver Inactivity Timer

One of the features of RDF is the RDserver Inactivity Timer. This feature gives system managers more control over rdallocated devices.

The purpose of the RDserver Inactivity Timer is to rddeallocate any rdallocated device if NO I/O activity to the rdallocated device has occurred within a predetermined length of time. When the RDserver Inactivity Timer expires, the server process drops the link to the client node and deallocates the physical device on the server node. On the client side, the client process deallocates the RDEVn0 device.

The default value for the RDserver Inactivity Timer is 3 hours.

The RDserver Inactivity Timer default value can be manually set by defining a system wide logical on the RDserver node prior to rdallocating on the rdclient node. The logical name is RDEV_SERVER_INACTIVITY_TIMEOUT.

To manually set the timeout value:

```
$ DEFINE/SYSTEM RDEV_SERVER_INACTIVITY_TIMEOUT seconds
```

For example, to set the RDserver Inactivity Timer to 10 hours, you would execute the following command on the RDserver node:

```
$ DEFINE/SYSTEM RDEV_SERVER_INACTIVITY_TIMEOUT 36000
```

8.7 RDF Error Messages

CLIDENY	Access from this CLIENT to the SERVER is not allowed. Check for "CLIENT/ALLOW" in the RDserver's configuration file.
CLIENTSBUSY	All 16 pseudo-devices are already in use.
DEVDENY	Client is not allowed to the Device or to the Node. This error message is dependent on the "CLIENT/ALLOW", "/ALLOW" or "CLIENT/DENY", "/DENY" qualifiers in the configuration file. Verify that the configuration file qualifier is used appropriately.
EMPTYCFG	The RDserver's configuration file has no valid devices or they are all commented out.
LINKABORT	The connection to the device was aborted. For some reason the connection was interrupted and the remote device could not be found. Check the configuration file as well as the remote device.
NOCLIENT	The RDdriver was not loaded. Most commonly the RDCLIENT_STARTUP.COM file was not executed for this node.
NOREMOTE	This is a RDF status message. The remote device could not be found. Verify the configuration file as well as the status of the remote device.
SERVERTMO	The RDserver did not respond to the request. Most commonly the RDSERVER_STARTUP.COM file was not executed on the server node. Or, the server has too many connections already to reply in time to your request.

System Backup to Tape for Oracle Databases

This chapter describes the System Backup to Tape (SBT) for Oracle databases feature of Archive Backup System (ABS). You can use this feature of the Archive Backup System and Media, Device and Management Services (MDMS) to back up Oracle databases directly to tape using Oracle's Recovery Manager.

As of this writing, the following versions of Oracle databases are supported:

- Oracle8i
- Oracle9i

In the rest of this section we use Oracle to refer to either Oracle8i or Oracle9i. If there is something that is specific to a release of Oracle, we will specify that release.

This section does not cover all aspects of configuring ABS /MDMS. This section only covers what you need to do to use SBT in the ABS/MDMS domain. Before configuring and using SBT, you must configure the following MDMS objects:

- Media
- Location
- Domain
- Node
- Jukebox
- Tape drives
- Pool
- Tape volumes

If you have been using ABS/MDMS you will already have your domain configured. If this is your first installation of ABS/MDMS, be sure to configure the above objects before proceeding with this section.

This section is presented in two portions. The first portion of this section reads like a tutorial in configuring and using SBT. You should read through this portion to see what is involved to configure and use SBT with Oracle's Recovery Manager. The second portion describes things like defaults, logicals, and troubleshooting.

The following topics are covered in this section:

1. Linking SBT with the Oracle server
2. Defining the logical MDMS\$SBT_TRACE_LEVEL
3. Configuring ABS

System Backup to Tape for Oracle Databases

9.1 Linking System Backup to Tape with the Oracle Server

4. Testing the configuration of SBT
5. Using SBT with Oracle's Recovery Manager
6. Using the show catalog command
7. Using the MDMS scheduler
8. System Backup to Tape defaults
9. System Backup to Tape logicals names
10. System Backup to Tape Restrictions
11. Troubleshooting tips

9.1 Linking System Backup to Tape with the Oracle Server

Before you can use SBT, you must link the SYS\$SHARE:MDMS\$SBTSHR_MA64.EXE shareable image with the Oracle server. This is a one time procedure. After performing this procedure, you can install a new release of ABS/MDMS and not have to relink the Oracle server to the new release of SBT.

Note

This linking of SBT to the Oracle server is not the same as described in the Oracle installation guide. The procedure in the Oracle installation guide does not use a shareable image and has you shutdown the database and relink with the vendors product with each new release. With the SBT shareable image and this procedure, you only have to shutdown the database and link one time.

This section takes you through the procedure to link the SYS\$SHARE:MDMS\$SBTSHR_MA64.EXE shareable image with the Oracle server:

1. Testing Oracle's Recovery Manager
2. Authorizing privileges and Granting rights to the Oracle server account
3. Editing Oracle's link option file and command procedures
4. Shutdown the database
5. Relinking the ORA_RDBMS: executables
6. Startup the database
7. Retesting Oracle's Recovery Manager

9.1.1 Testing Oracle's Recovery Manager before linking System Backup to Tape

Before doing the configuration for SBT, you should make sure that Oracle's Recovery Manager is setup and you are able to access it. If you have been using Oracle's Recovery Manager to write to disk, then you are ready to switch to SBT. If you have not been using Oracle's Recovery Manager, you should try a backup of a tablespace as shown in Example 9–1.

Example 9–1 Oracle's Recovery Manager Backup of System Tablespace to Disk

```
RMAN> run
2> {
3>   allocate channel d1 type disk;
4>   backup tablespace system;
5>   release channel d1;
6> }
```

System Backup to Tape for Oracle Databases

9.1 Linking System Backup to Tape with the Oracle Server

Example 9–1 created the file ORA_DB:02D9MBV4_1_1.;1 on my system. Of course your file name will be different. If everything works then you are ready to link Oracle to SBT. If this step did not work, then your Oracle Recovery Manager is not setup correctly. You need to correct this before proceeding.

9.1.2 Authorizing privileges and granting rights to the Oracle server account

Before using SBT, you must authorize the VOLPRO privilege and grant the MDMS_APPLICATION identifier to the Oracle server database administrator account. The VOLPRO privilege allows the Oracle server to mount volumes that belong to ABS. All volumes belong to ABS. The MDMS_APPLICATION allows the Oracle server to use the objects in the MDMS database.

The following example shows the commands that modify the privileges and grant the right to an account called ORACLE9I:

```
$ MCR AUTHORIZE
UAF> MODIFY ORACLE9I/PRIVILEGES=VOLPRO
UAF> GRANT/ID MDMS_APPLICATION ORACLE9I
UAF> EXIT
$
```

Note

Be sure to logout and log back in so that the privileges and rights take effect.

9.1.3 Editing Oracle's Link Option File and Command Procedures

In order to link the SBT shareable image, you must change Oracle's link option file and command procedures. This section covers what you need to do for Oracle8i and Oracle9i.

9.1.3.1 Editing Oracle8i Link Option File and Command Procedures

In order to link the SBTshareable image, you must change three of Oracle's files. Before editing these three files, we suggest that you make a copy of the file. In each file, you may need to comment out the line ora_rman_mml_64/lib and/or add the line SYSSHARE:MDMS\$SBTSHR_MA64.EXE/SHARE. The following are the three files and an example of each file with the line to comment out and/or the line to be added commented:

- ORA_UTIL:RDBMS_RMAN_NOSHARE.OPT as shown in Example 9–2.
- ORA_RDBMS:LORACLE_64.COM as shown in Example 9–3.
- ORA_UTIL:LOUTL.COM as shown in Example 9–4

Example 9–2 Edited Oracle8i ORA_UTIL:RDBMS_RMAN_NOSHARE.OPT file

```
!
!
! rdbms libraries
ora_olb:libvsn8/lib
! ora_rman_mml/lib          COMMENT OUT THIS LINE
ora_olb:libwtc8/lib
ora_olb:libclient8/lib
ora_olb:libcommon8/lib
ora_olb:libgeneric8/lib
ora_olb:libclient8/lib
ora_olb:libcommon8/lib
ora_olb:libgeneric8/lib
```

System Backup to Tape for Oracle Databases

9.1 Linking System Backup to Tape with the Oracle Server

Example 9–3 Edited Oracle8i ORA_RDBMS:LORACLE_64.COM

```
ora_olb:libclient8_64/lib/incl=(kgu),-
'rdbmslib$$'-
'plsqllib$$'-
'rdbmslib$$'-
! ora_rman_mml_64/lib,-                COMMENT OUT THIS LINE
ora_olb:libnro8_64/lib,-
'network$$'-
ora_olb:libtrace8_64/lib,-
'oracore$$'-
'cart64$$'-
ora_olb:libslax8_64/lib,-
'utl$$'-
'oracore$$'-
sys$input/options
sys$share:mdms$sbtsshr_ma64.exe/share, -    !!! ADDED THIS LINE
```

Example 9–4 Edited Oracle8i ORA_UTIL.LOUTL.COM

```
$nonSharedLink:
$ 'loutl_link_cmd$$'/alpha/nouserlibrary'dotrace$$' 'map$$' 'mapex-
tra$$' 'image$$'=
'filename$$' 'switch$$' 'userlink$$'/sysexe -
'p2',-
ora_olb:libclient8/lib,-
ora_olb:libsql8/lib,-
'ocis$$'-
'fastupi$$'-
'network$$'-
'rdbmslib_noshare$$'-
'oracore$$'-
'network$$'-
'rdbmslib_noshare$$'-
'otrancelib$$'-
'oracore$$'-
'rdbmslib_noshare$$'-
'oracore$$'-
'useroption$$'-
sys$input/opt
sys$share:mdms$sbtsshr_ma64.exe/share, -    !!! ADDED THIS LINE
sys$share:decc$shr/share
! Temporary: fixup readonly attributes between compiler
versions.
psect_attr = $readonly$,pic,shr
```

9.1.3.2 Editing Oracle9i Link Option file and Command Procedures

In order to link the SBT shareable image, you must change Oracle's link option file and command procedures. The three files require you to comment out one line and add one line in each file. Before editing these three files, we suggest that you make a copy of the file. In each file, you need to comment out `ora_rman_mml_64/lib` and add the line `SYSSHARE:MDMS$SBTSHR_MA64.EXE/SHARE`. The following are the three files and an example of each file with the line to comment out and the line to add commented:

- ORA_RDBMS:RDBMS_RMAN_NOSHARE_64.OPT as shown in Example 9–5
- ORA_RDBMS:LORACLE_64.COM as shown in Example 9–6
- ORA_RDBMS:LSHRCLIENT_64.COM as shown in Example 9–7

System Backup to Tape for Oracle Databases

9.1 Linking System Backup to Tape with the Oracle Server

Example 9–5 Edited ORA_RDBMS:RDBMS_RMAN_NOSHARE_64.OPT file

```
!  
!  
! rdbms libraries  
ora_olb:libvsn9/lib  
ora_olb:libobk/lib  
! ora_rman_mml_64/lib COMMENT OUT THIS LINE  
sys$share:mdms$sbtshr_ma64.exe/share !!! ADDED THIS LINE  
ora_olb:libwtc9/lib  
ora_olb:libclient9/lib  
ora_olb:libcommon9/lib  
ora_olb:libgeneric9/lib  
ora_olb:libclient9/lib  
ora_olb:libcommon9/lib  
ora_olb:libgeneric9/lib
```

Example 9–6 Editing ORA_RDBMS:LORACLE_64.COM

```
ora_olb:libobk_64/lib,-  
! ora_rman_mml_64/lib,- COMMENT OUT THIS LINE  
ora_olb:libnro9_64/lib,-  
'network$$'-  
ora_olb:libtrace9_64/lib,-  
'oracore$$'-  
'cart64$$'-  
ora_olb:libslax9_64/lib,-  
'utl$$'-  
'oracore$$'-  
sys$input/options  
sys$share:mdms$sbtshr_ma64.exe/share, - !!! ADDED THIS LINE  
sys$share:decc$shr/share
```

Example 9–7 Editing ORA_RDBMS:LSHRCLIENT_64.COM

```
o$$:libobk/lib,-  
! ora_rman_mml_64/lib,- COMMENT OUT THIS LINE  
o$$:libnro9_64/lib,-  
'network$$'-  
'oracore$$'-  
'rdbmslib$$'-  
'oracore$$'-  
'network$$'-  
'rdbmslib$$'-  
'otrancelib$$'-  
'oracore$$'-  
'plsqli$$'-  
'slax$$'-  
'utl$$'-  
'oracore$$'-  
'rdbms2$$'-  
o$$:libcore9_objlib_64/lib/include=(sscoreed),-  
sys$input/opt  
sys$share:mdms$sbtshr_ma64.exe/share, - !!! ADDED THIS LINE  
sys$share:decc$shr/share
```

9.1.4 Shutdown the database

Before relinking the database, you should shutdown the database.

9.1.5 Relinking the ORA_RDBMS: executables

Now that you have prepared the files for relinking, you must relink the ORA_RDBMS: executables. Invoke ORA_INSTALL:ORACLEINS and select RDBMS for rebuild.

9.1.6 Startup the database

Now that you have relinked the Oracle server, you should startup up the database.

System Backup to Tape for Oracle Databases

9.2 Defining the logical MDMS\$SBT_TRACE_LEVEL

9.1.7 Retesting Oracle's Recovery Manager

Before proceeding, you should retest Oracle's Recovery Manager as you did in Section 9.1.1.

9.2 Defining the logical MDMS\$SBT_TRACE_LEVEL

The logical MDMS\$SBT_TRACE_LEVEL allows you to define how much tracing of SBT you want to appear in your trace file. The logical is defined in SYS\$STARTUP:MDMS\$SYSTARTUP.COM. However, if you had previous versions of ABS/MDMS on your system, it may not be in SYS\$STARTUP:MDMS\$SYSTARTUP.COM. You can check using the following command:

```
$ SEARCH SYS$STARTUP:MDMS$SYSTARTUP.COM MDMS$SBT_TRACE_LEVEL
%SEARCH-I-NOMATCHES, no strings matched
```

If the search command did not find it, you need to edit SYS\$STARTUP:MDMS\$SYSTARTUP.TEMPLATE and pull the following code out of it and put the code in SYS\$STARTUP:MDMS\$SYSTARTUP.COM.

```
$ !
$ !
$ !   The MDMS$SBT_TRACE_LEVEL log name controls what is written to the
$ !   Oracle trace file for SBT.  The trace level can be controlled by
$ !   this logical separately from the trace level in the Oracle parameters.
$ !
$ TR_ERROR      = %X00000000 ! Always trace errors, cannot be changed
$ TR_SBTENTRY   = %X00000001 ! Entry and exit of oracle called SBT functions
$ TR_SBTPARAM   = %X00000002 ! Trace oracle called SBT functions parameters
$ TR_SBTREENTRY = %X00000004 ! Trace of SBTREAD/SBTWRITE entry
$ TR_SBTREPARAM = %X00000008 ! Trace of parameters for SBTREAD/SBTWRITE
$ TR_GENINFO    = %X00000010 ! Trace general information like backup file
name
$ TR_MEDINFO    = %X00000020 ! Trace media movement information
$ TR_TAPSTAT    = %X00000040 ! Trace tape/disk transfer stats
$ TR_COMENTRY   = %X00000080 ! Entry and exit for common functions
$ TR_COMPARAM   = %X00000100 ! Trace parameters for common functions
$ TR_MEDENTRY   = %X00000200 ! Entry and exit for media functions
$ TR_MEDPARAM   = %X00000400 ! Trace parameters for media functions
$ TR_CATENTRY   = %X00000800 ! Entry and exit for catalog functions
$ TR_CATPARAM   = %X00001000 ! Trace parameters for catalog functions
$ TR_TAPENTRY   = %X00002000 ! Entry and exit for VMSTAPE functions
$ TR_TAPPARAM   = %X00004000 ! Trace parameters for VMSTAPE functions
$ TR_VOLENTRY   = %X00008000 ! Entry and exit for VOLSET functions
$ TR_VOLPARAM   = %x00010000 ! Trace parameters for VOLSET functions
$ tracefilter = TR_GENINFO .OR. TR_MEDINFO
$ DEFINE/SYSTEM/NOLOG MDMS$SBT_TRACE_LEVEL 'tracefilter'
$!
```

After editing SYS\$STARTUP:MDMS\$SYSTARTUP.COM, be sure to execute it so the logical is defined. By default the general information and media movement information are traced in the trace file ORA_DUMP:SBTIO.LOG. Refer to Section 9.10.1 for more information about the logical MDMS\$SBT_TRACE_LEVEL.

9.3 Configuring System Backup to Tape in the Archive Backup system

Before you can use SBT, you must configure a catalog and an archive in ABS/MDMS. This section describes how to create catalogs and archives. Catalogs store information about what information was backed up. Archives allow you to implement your storage policies.

This section shows you how to create the default catalog (ORACLE_DB) and the default archive (ORACLE_DB_ARCHIVE). By creating these two default objects, you can test SBT easier as described in Section 9.4.

System Backup to Tape for Oracle Databases

9.3 Configuring System Backup to Tape in the Archive Backup system

9.3.1 Creating an ORACLE_DB Catalog

Before you can use the SBT feature, you must create a catalog. The catalog stores the tape volume, saveset name, and piece name. The catalog allows SBT to lookup the tape volume for a restore. You should only create one oracle_db type catalog that is accessible to Oracle's Recovery Manager(s). The catalog is created in ABS\$CATALOG:. The catalog must be created in the ABS\$CATALOG: directory that is local to all nodes that will access the catalog.

Use the following command to create a catalog named ORACLE_DB:

```
$ MDMS CREATE CATALOG ORACLE_DB /TYPE=ORACLE_DB
```

Note

Do not use the /NODE qualifier when creating the catalog. This version of SBT can only access catalogs that are local to the node.

When doing an Oracle Recovery Manager restore, allocateForMaint, or validate command, you must specify the catalog you want to use. However, the default catalog name for SBT is ORACLE_DB. Therefore, if you do not specify a catalog, it will lookup information in the ORACLE_DB catalog on the local node.

Refer to Section 9.6 for information on how to access the ORACLE_DB catalog.

9.3.2 Creating an Archive

An archive defines the tape volumes and archive attributes where you can safely store data. Each archive has a unique name and contains a set of archive characteristics. You can have as many archives as you want to implement your storage policies. This section describes how to create an archive and what attributes pertain to SBT. You should refer to the command reference guide for information about creating archives and their attributes. You may want to create an archive that keeps tape volumes for a year and another that keeps tape volumes for 35 days.

Note

When using SBT you do not use an environment, save, or restore object.

The following command creates an archive named ORACLE_DB_ARCHIVE:

```
$ MDMS CREATE ARCHIVE ORACLE_DB_ARCHIVE -  
/ARCHIVE_TYPE=TAPE -  
/CATALOG=(NAME=ORACLE_DB) -  
/MAXIMUM_SAVES=36 -  
/MEDIA_TYPE=DLT_III -  
/POOL=DB_BACKUP_POOL -  
/RETENTION_DAYS=35  
$ MDMS SHOW ARCHIVE ORACLE_DB_ARCHIVE
```

```
Archive: ORACLE_DB_ARCHIVE 1  
Description:  
Access Control: NONE  
Owner: MOE::ORACLE9I  
Archive Type: TAPE 2  
Catalog -  
- Name: ORACLE_DB 3  
- Nodes:  
Consolidation -  
- Interval: 0007 00:00:00 4  
- Savesets: 0  
- Volumes: 0
```

System Backup to Tape for Oracle Databases

9.3 Configuring System Backup to Tape in the Archive Backup system

```
Destination:
  Drives: 5
Expiration Date: NONE
  Location: CR_2 6
Maximum Saves: 36 7
  Media Type: DLT_III 8
  Pool: DB_BACKUP_POOL 9
Retention Days: 35 10
Volume Sets:
```

The following describes the different attributes of the archive:

1. **Archive:** ORACLE_DB_ARCHIVE is the name of the archive created. This is the name you must specify in Oracle's Recovery Manager allocate command. See Section 9.5.1 on how to specify the archive name in the allocate command. If you do not specify an archive in Oracle's Recovery Manager commands, ORACLE_DB_ARCHIVE is used as the default. See Section 9.8.1 for more information.
2. **Archive Type:** this specifies that the backup will be archived to tape. This version does not support archive to disk.
3. **Catalog Name:** you need to specify which catalog used by this archive. If you create more than one catalog, you must have a different archive for each catalog. However, all archives can use the same catalog. See Section 9.8.2 for more information.
4. **Consolidation Interval:** specifies the consolidation interval for the volume sets. The volumes sets are stored in the Volume Sets: attribute. In this example, after 7 days a new volume set is started. The default consolidation interval is 7 days.
5. **Drives:** specifies the tape drives that you want to use when using this archive. This limits the tape drives that you can use. Unless you need to limit which tape drives to use, I suggest you do not. See Section 9.9.1 for restrictions in using this attribute.
6. **Location:** this location came from the MDMS domain object. SBT uses it in selecting a tape volume and tape drive. Your tape volume, jukebox, and node must also have this location attribute. If you do not need a location to specify the volumes you want to allocate, specify /NOLOCATION.
7. **Maximum Saves:** this attribute specifies how many backups can use the volume sets in this archive at a time. This works great for ABS saves, however, there is a difference when used with SBT. We suggest that you set it at 36 which is the maximum. You should control the number of backups using Oracle's Recovery Manager and not this attribute. See Section 9.9.1 for restrictions in using this attribute.
8. **Media Type:** this is the media type that SBT uses to allocate tape volumes and tape drives. You must have a media type.
9. **Pool:** this is the pool that SBT uses to allocate tape volumes along with location and media type. If you do not need a pool to specify the volumes you want to allocate, specify /NOPOOL.
10. **Retention Days:** this specifies the number of days that the volume will be retained before being scratched. If you want different retention days for different backups, you can use a different archive with a different retention days specified.

Now that you have created a catalog and archive, you are ready to test that everything is configured correctly. The next section shows how to test the configuration.

9.4 Testing the Configuration of SBT

Now that you have linked SBT with the Oracle server and created a catalog and archive, you are ready to test the SBT configuration. Supplied with SBT is a test program, SYSSYS-TEM:MDMS\$SBTTEST_MA64.EXE, that allows you to test the SBT interface with ABS/MDMS. If you have been using ABS/MDMS you may want to skip this section. It just gives you confidence that ABS/MDMS is setup correctly.

Using sbttest is described in the Oracle documentation. However, their executable will not work with the SYSSHARE:MDMS\$SBTSHR_MA64.EXE. Oracle supplied the sbttest code and I compiled and linked it to work with SYSSHARE:MDMS\$SBTSHR_MA64.EXE and supplied it for your use as SYSSYSTEM:MDMS\$SBTTEST_MA64.EXE.

You must have a catalog and archive defined to use sbttest. In Section 9.3 you created a catalog and archive. You will use these for the test. The following commands show how to use sbttest:

```

$ SBTTEST := $SYSSYSTEM:MDMS$SBTTEST_MA64.EXE 1
$ DEFINE MDMS$SBT_ARCHIVE ORACLE_DB_ARCHIVE 2
$ DEFINE MDMS$SBT_CATALOG ORACLE_DB 3
$ SBTTEST TESTFILE -TRACE SBTTEST.TRC 4
MM software supports SBT API version 2.0
MM software is version 4.0.0.0

sbtinit, vendor description string="System Backup to Tape V4.0 (436)" 5
sbtinit successful
sbtinit2 successful
sbtbackup successful
sbtwrite2 successful, wrote 100 blocks
sbtinfo2, SBTBFINFO_NAME=testfile
sbtinfo2, SBTBFINFO_METHOD=stream
sbtinfo2, SBTBFINFO_COMMENT=Onsite: Description for AIF078
sbtinfo2, SBTBFINFO_CREATIME=Fri Dec 7 05:11:07 2001
sbtinfo2, SBTBFINFO_EXPTIME=Sat Dec 7 05:11:07 2002
sbtinfo2, SBTBFINFO_SHARE=single user
sbtinfo2, SBTBFINFO_ORDER=sequential access
sbtinfo2, SBTBFINFO_LABEL=AIF078
sbtinfo2 successful
sbtrestore successful
file was created by this program; seed=1007752246,
blk_size=16384, blk_count=100
sbtread2 successful, read 100 buffers
sbtclose2 successful
sbtremove2(remove_after) successful, remove "testfile"
sbtend successful
*** The SBT API test was successful ***

```

The following describes the commands in the above example:

1. Define the sbttest symbol. By default the sbttest command is pointing to Oracle's ORA_RDBMS:SBTTEST.EXE. However, this executable will not work with SBT.
2. Define the MDMS\$SBT_ARCHIVE logical. This logical points to the archive that you created in Section 9.3.2. In this case, we would not have to define this logical because it is the default. If you did not create an ORACLE_DB_ARCHIVE archive, you would have specified the archive here.
3. Define the MDMS\$SBT_CATALOG logical. This logical points to the catalog that you created in Section 9.3.1. In this case, we would not have to define this logical because it is the default. If you did not create the default catalog, ORACLE_DB, you would have to specify it here.

System Backup to Tape for Oracle Databases

9.5 Using System Backup to Tape with Oracle's Recovery Manager

4. Run `sbttest` using the `sbttest` command. The parameter `TESTFILE` is a made up name that gets put in the catalog during the backup and then is used for the restore. The data stored and retrieved is 100 blocks of 16384 characters per block. By specifying the `-TRACE` flag the file `SBTTEST.TRC` is written (see below).
5. Vendor description string shows that you are using the `SYSS$SHARE:MDMS$SBTSHR_MA64.EXE` shareable image.

The `-TRACE` flag generates the `SBTTEST.TRC` trace file. The trace file should look like the following example:

```
SBT-00001DB2 12/14/01 10:12:30 Using archive ORACLE_DB_ARCHIVE
SBT-00001DB2 12/14/01 10:12:31 Starting backup of testfile for DB: sbtdb
SBT-00001DB2 12/14/01 10:12:31 Using catalog ORACLE_DB
SBT-00001DB2 12/14/01 10:12:31 Attempting to allocate volume set BEB026
SBT-00001DB2 12/14/01 10:12:33 Allocated drive: TLZ88D Device: MOE$MKC200:
SBT-00001DB2 12/14/01 10:12:33 Drive is in jukebox TLZ88J
SBT-00001DB2 12/14/01 10:12:37 Loading/mounting volume BEB026 on drive TLZ88D
SBT-00001DB2 12/14/01 10:12:37 Loading volume BEB026 on drive TLZ88D
SBT-00001DB2 12/14/01 10:12:45 Mounting volume BEB026 on device MOE$MKC200:
SBT-00001DB2 12/14/01 10:12:53 Skipping 9 tapemarks to end of tape
SBT-00001DB2 12/14/01 10:13:02 Ready to write to saveset 2001121410125267.
on volume BEB026
SBT-00001DB2 12/14/01 10:13:07 Using catalog ORACLE_DB
SBT-00001DB2 12/14/01 10:13:07 Finished writing saveset 2001121410125267.
on volume BEB026
SBT-00001DB2 12/14/01 10:13:09 Starting restore of testfile
SBT-00001DB2 12/14/01 10:13:09 Using catalog ORACLE_DB
SBT-00001DB2 12/14/01 10:13:10 Dismounting volume set member: BEB026 RVN 1
SBT-00001DB2 12/14/01 10:13:10 Deallocating drive TLZ88D
SBT-00001DB2 12/14/01 10:13:11 Allocated drive: TLZ88D Device: MOE$MKC200:
SBT-00001DB2 12/14/01 10:13:11 Drive is in jukebox TLZ88J
SBT-00001DB2 12/14/01 10:13:20 Loading/mounting volume BEB026 on drive TLZ88D
SBT-00001DB2 12/14/01 10:13:20 Loading volume BEB026 on drive TLZ88D
SBT-00001DB2 12/14/01 10:13:29 Mounting volume BEB026 on device _RDEVA0:
SBT-00001DB2 12/14/01 10:13:37 Skipping 9 tapemarks to beginning of saveset
SBT-00001DB2 12/14/01 10:13:44 Ready to read from saveset 2001121410125267.
on volume BEB026
SBT-00001DB2 12/14/01 10:13:47 Finished restoring saveset 2001121410125267.
from volume BEB026
SBT-00001DB2 12/14/01 10:13:49 Dismounting volume set member: BEB026 RVN 1
SBT-00001DB2 12/14/01 10:13:49 Deallocating drive TLZ88D
```

You are now ready to start using Oracle's Recovery Manager to backup your Oracle database to tape. The following section describes what Oracle's Recovery Manager commands allow you to control how you do backups using SBT.

9.5 Using System Backup to Tape with Oracle's Recovery Manager

This section describes how to use SBT to backup your Oracle database. If you have configured SBT correctly using the above steps, you are now ready to use SBT with Oracle's Recovery Manager. How to use Oracle's Recovery Manager is described in Oracle's documentation. This section describes what Oracle Recovery Manager command keywords and parameters affect SBT.

The following topics are covered in this section:

- Specifying an archive
- Specifying a catalog
- Specifying I/O block size
- Specifying archives for duplex backups

9.5.1 Specifying an Archive

In order to have SBT select the archive that you want, you must specify the archive in the Oracle Recovery Manager allocate command. The archive is specified in the `parms` keyword for the allocate command. Example 9–8 shows how to code an Oracle Recovery Manager script for the archive `OFFSITE_ARCH`. The Oracle server creates the process logical `MDMS$SBT_ARCHIVE`.

Example 9–8 Specifying the Archive in the Allocate Command

```
run
{
  allocate channel t1 type 'sbt_tape'
    parms="ENV=(MDMS$SBT_ARCHIVE=OFFSITE_ARCH)";
  backup tablespace system;
  release channel t1;
}
```

Note

Everything between the quotes in `parms` MUST be uppercase characters.

If you are doing a parallel backup, you need to specify an archive for each channel as shown in Example 9–9. Also, in this example, I show how to specify a catalog.

Example 9–9 Specifying the Archive for each Channel

```
run
{
  allocate channel t1 type 'sbt_tape'
    parms="ENV=(MDMS$SBT_ARCHIVE=ONSITE_ARCH)";
  allocate channel t2 type 'sbt_tape'
    parms="ENV=(MDMS$SBT_ARCHIVE=ONSITE_ARCH)";
  allocate channel t3 type 'sbt_tape'
    parms="ENV=(MDMS$SBT_ARCHIVE=ONSITE_ARCH)";
  backup
    ( tablespace tbs1,tbs2 channel t1 )
    ( tablespace tbs3,tbs4 channel t2 )
    ( tablespace tbs5,tbs6, system channel t3 );
  release channel t1;
  release channel t2;
  release channel t3;
}
```

If you have an archive you want to use as a default, you can define `MDMS$SBT_ARCHIVE` as a system wide logical. Then if you want something different for a particular backup, you can define it in the allocate command. Also, you can specify different archives for each channel you allocate.

9.5.2 Specifying a Catalog

In order to have SBT select the catalog to use, you must specify the catalog in the Oracle Recovery Manager allocate command. The catalog is specified in the `parms` keyword for the allocate command. Example 9–10 shows how to code an Oracle Recovery Manager script for the catalog `OFFSITE_CAT`.

System Backup to Tape for Oracle Databases

9.5 Using System Backup to Tape with Oracle's Recovery Manager

Example 9–10 Specifying the Catalog in the Allocate Command

```
run
{
  allocate channel t1 type 'sbt_tape'
    parms="ENV=(MDMS$SBT_CATALOG=OFFSITE_CAT)";
  restore tablespace tbs6;
  recover tablespace tbs6;
  release channel t1;
}
```

If you only have one catalog named ORACLE_DB, you never have to specify MDMS\$SBT_CATALOG. I only used different catalogs here for examples.

For a Recovery Manager backup operation, the catalog in the archive is always used. The MDMS\$SBT_CATALOG in the parms keyword of the allocate command is ignored.

For a Recovery Manager restore, crosscheck, or delete expired command the catalog that SBT uses is determined by the logical MDMS\$SBT_CATALOG, the catalog in the archive, or the default catalog. If the logical MDMS\$SBT_CATALOG is defined, SBT uses that catalog. If MDMS\$SBT_CATALOG is not defined, SBT checks to see if MDMS\$SBT_ARCHIVE is defined. If MDMS\$SBT_ARCHIVE is defined, SBT uses the catalog in the archive. If MDMS\$SBT_ARCHIVE is not defined, the default catalog ORACLE_DB is used.

9.5.3 Specifying an I/O Block Size

To help tune your output to different devices, SBT allows you to specify an I/O block size. In the case of a tape device, the block size is how much data is written to the tape device at one time. The default is the maximum of 65024 bytes. Example 9–11 shows how to code an Oracle Recovery Manager script for the I/O block size of 32768.

Example 9–11 Specifying the I/O Block Size in the Allocate Command

```
run
{
  allocate channel t1 type 'sbt_tape'
    parms="ENV=(MDMS$SBT_ARCHIVE=ONSITE_ARCH,
      MDMS$SBT_IO_BLOCK_SIZE=32768)";
  allocate channel t2 type 'sbt_tape'
    parms="ENV=(MDMS$SBT_ARCHIVE=ONSITE_ARCH,
      MDMS$SBT_IO_BLOCK_SIZE=32768)";
  allocate channel t3 type 'sbt_tape'
    parms="ENV=(MDMS$SBT_ARCHIVE=ONSITE_ARCH,
      MDMS$SBT_IO_BLOCK_SIZE=32768)";
  backup filesper set 1
    database;
  backup
    current controlfile;
  release channel t1;
  release channel t2;
  release channel t3;
}
```

9.5.4 Specifying Archives for Duplex Backups

SBT allows you specify an archive for each stream of a duplex backup. The duplex mode allows duplicate backups to separate tape volumes. This is accomplished in SBT by having a different archive for each stream. The different archives are passed into SBT using Oracle's Recovery Manager allocate command. The parms keyword uses the keyword of ENV. By using MDMS\$SBT_ARCHIVE_1, MDMS\$SBT_ARCHIVE_2, and so forth, you can specify which archive to use for which copy. Example 9–12 shows an example of doing a duplex backup with two archives: OFFSITE_ARCH and ONSITE_ARCH.

Example 9–12 Duplex Command using two Archives

```
run
{
  set duplex=2;
  allocate channel t1 type 'sbt_tape'
  parms="ENV=(MDMS$SBT_ARCHIVE_1=OFFSITE_ARCH,
            MDMS$SBT_ARCHIVE_2=ONSITE_ARCH)";
  backup tablespace system;
  release channel t1;
}
```

9.6 Using the Show Catalog Command

Using the show catalog command in MDMS allows you to look up information about a piece name. Because there are different types of catalogs not all of the qualifiers pertain to an ORACLE_DB type of catalog. Also, because the ORACLE_DB catalog type is a variation of another catalog, all of the qualifiers are not what you expect. Example 9–13 shows a simple command to retrieve information about piece name vedbk5ha_1_1.

Example 9–13 Simple Show Catalog Example

```
$ MDMS SHOW CATALOG ORACLE_DB -
_$/SAVE/FULL/PIECE_NAME="vedbk5ha_1_1" 1
  Catalog Name: ORACLE_DB
  Catalog Node: MOE
  Date Archived: 13-DEC-2001 20:23:07
  Source Node: MOE
  Database: EMPLOYEE
  Block Size: 262144 2
  Archive: RMAN_TAPE_TL893_ARCH
  Environment: 3E6EE027-F007-11D5-9421-5441524E2020 3
  Save: 3E6EE028-F007-11D5-941F-5441524E2020
  Save Type: () 4
  Owner: ABS
  Saveset Format: ORACLE_DB_SBT 5
  Archive Type: TAPE 6
  Saveset Location: AIF049,AIF050 7
  Saveset Name: 2001121320230791. 8
  Saveset Position: 288 9
  Status: Completed with success 10
  Severity: OP_SUCCESS
  Piece Name: vedbk5ha_1_1
```

The following describes the command and the different fields in the display that I think might not be obvious to you or I think needs an explanation:

1. The MDMS show catalog command—you must use the /SAVE qualifier for an ORACLE_DB catalog type. The /PIECE_NAME lets you look up a particular piece name. The piece name must be in quotes if the piece name is not all uppercase letters. If you do not use the /PIECE_NAME qualifier, you will get every entry in the catalog. You need to use the /FULL qualifier or you will not get much information in the modal default of /BRIEF. I hope to change this in a subsequent release.
2. Block Size—this is the block size that was sent down from Oracle. When a restore command is issued from the Oracle Recovery Manager, the block size must be the same.
3. Environment and Save—these UID's mean nothing for an ORACLE_DB type catalog. They are here because the ORACLE_DB type catalog is a variation of another ABS catalog.

System Backup to Tape for Oracle Databases

9.6 Using the Show Catalog Command

4. Save Type-this means nothing for an ORACLE_DB type catalog.
5. Saveset Format-this is the format type of the saveset on the I/O device. SBT has its own format and can only be read by SBT.
6. Archive Type-TAPE archive type is the only type supported in this version.
7. Saveset Location-this is the tape volume(s) that the saveset is on. In this example, there are two tape volumes: AIF049 and AIF050. This means that the saveset was started on tape volume AIF049 and finished on tape volume AIF050. If you need to do a restore of this piece name, you need to have both of these volumes onsite. You can use the MDMS show volume command to look up these tape volumes to see if they are onsite or offsite.
8. Saveset Name-this is the name of the file on the tape volume. We are limited to 17 characters so we change piece name into a saveset name on a particular volume.
9. Saveset Position-this is the number of tape marks down from the beginning of the tape where the saveset starts. I doubt if you will ever use this.
10. Status and Severity-these will always be success in the present version. The piece name is only put in the catalog if the backup completed successfully.

You may want to look up all of the piece names for a particular database. The /INCLUDE qualifier is the one to use. The /INCLUDE qualifier was there before we created the ORACLE_DB type catalog. It accesses the database field of an ORACLE_DB type catalog. Example 9-14 shows a lookup of all records for the EMPLOYEE database.

Example 9-14 Catalog Lookup of All Records for the EMPLOYEE Database

```
Oracle9i> MDMS SHOW CATALOG ORACLE_DB/SAVE/FULL/INCLUDE=EMPLOYEE
Catalog Name: ORACLE_DB
Catalog Node: MOE
Date Archived: 14-DEC-2001 10:04:53
Source Node: MOE
Database: EMPLOYEE
Block Size: 262144
Archive: RMAN_TAPE_TL875_ARCH
Environment: 08E036C7-F07A-11D5-BFEF-5441524E2020
Save: 08E036C8-F07A-11D5-BFED-5441524E2020
Save Type: ()
Owner: ABS
Saveset Format: ORACLE_DB_SBT
Archive Type: TAPE
Saveset Location: DEC031
Saveset Name: 2001121410045349.
Saveset Position: 9
Status: Completed with success
Severity: OP_SUCCESS
Piece Name: 75dbllm4_1_1
Catalog Name: ORACLE_DB
Catalog Node: MOE
Date Archived: 14-DEC-2001 10:00:46
.
.
.
Database: EMPLOYEE
Block Size: 262144
Archive: RMAN_TAPE_TL893_ARCH
Environment: 9868BCE3-E3EA-11D5-8CAB-5441524E2020
Save: 9868BCE2-E3EA-11D5-8CAA-5441524E2020
Save Type: ()
Owner: ABS
Saveset Format: ORACLE_DB_SBT
Archive Type: TAPE
Saveset Location: AFW892
Saveset Name: 2001112810163927.
Saveset Position: 0
```

System Backup to Tape for Oracle Databases 9.7 Using the MDMS Scheduler

```
Status: Completed with success
Severity: OP_SUCCESS
Piece Name: jpda8rm4_1_1
```

For more information about the MDMS show catalog command, refer to the reference manual.

9.7 Using the MDMS Scheduler

Because Oracle's Recovery Manager has no way of scheduling scripts to run periodically, you need to use an external scheduler. MDMS provides a scheduling capability. See the command reference guide for creating schedule objects. Example 9–15 shows how to create a schedule to run a command procedure which has Oracle Recovery Manager commands in it.

Example 9–15 Creating an MDMS schedule

```
$ MDMS CREATE SCHEDULE BACKUP_DB_TAPE -
  /COMMAND="@DISK$ORACLE5:[ORACLE9I]BACKUP_DB_TAPE.COM" -
  /TIMES=21:00
$ MDMS SHOW SCHED BACKUP_DB_TAPE
  Schedule: BACKUP_DB_TAPE
  Description:
  Access Control: NONE
    Owner: MOE::ORACLE9I
  After Schedule:
  After When: NONE
  Command: @DISK$ORACLE5:[ORACLE9I]BACKUP_DB_TAPE.COM
  Dates: 1-31
  Days: MON-SUN
  Exclude:
  Include:
  Months: JAN-DEC
  Times: 21:00
  Last Start Date: NONE
  Next Start Date: 07-DEC-2001 21:00:00
```

The following example shows my BACKUP_DB_TAPE.COM file:

```
$ set nover
$ @oracle_home:login
$ set verify
$ rman nocatalog target sys/admin@orcl
run
{
  allocate channel t1 type 'sbt_tape'
    parms='ENV=(MDMS$SBT_ARCHIVE=ONSITE_ARCH)';
  backup
    format 'Empt%t_s%s_p%p'
    database;
  backup
    format 'Empcct%t_s%s_p%p'
    current controlfile;
  release channel t1;
}
exit
$ exit
```

Note

When specifying a format in the backup command you must put in a % character that creates unique names being sent to SBT.

You can put Oracle's Recovery Manager command line in the scheduler command and execute a stored script through a cmdfile.

System Backup to Tape for Oracle Databases

9.8 System Backup to Tape Defaults

9.8 System Backup to Tape Defaults

The SBT feature of ABS/MDMS has defaults. This section describes these defaults.

9.8.1 Archive Name

If you do not pass the parms="ENV=(MDMS\$SBT_ARCHIVE=archive_name)" in an Oracle Recovery Manager allocate command, SBT will use the default ORACLE_DB_ARCHIVE archive.

9.8.2 Catalog Name

If you do not pass the parms="ENV=(MDMS\$SBT_CATALOG=catalog_name)" in an Oracle Recovery Manager allocate command, SBT will use the default ABS\$CATALOG:ORACLE_DB catalog.

9.8.3 I/O Block Size

If you do not pass the parms="ENV=(MDMS\$SBT_IO_BLOCK_SIZE=block_size)" in an Oracle Recovery Manager allocate command, SBT will use a block size of 65024 bytes when writing blocks on the I/O device. This has nothing to do with the block size that Oracle sends down to be written on the I/O device.

9.8.4 System Backup to Tape Logicals Names

Table 9–1 describes the logical names used to control your SBT sessions. These logical names can be a system wide logical (DEFINE/SYSTEM) or be defined in your Oracle Recovery Manager scripts. When using in Oracle's Recovery Manager scripts, the logical is declared when you allocate a channel with the keyword of ENV for the keyword parms. The following example shows how the MDMS\$SBT_ARCHIVE and MDMS\$SBT_IO_BLOCK_SIZE logicals are defined as a process logical with Oracle's Recovery Manager allocate command:

```
run
{
  allocate channel t1 for 'sbt_tape'
  parms="ENV=(MDMS$SBT_ARCHIVE=OFFSITE_ARCHIVE,
             MDMS$SBT_IO_BLOCK_SIZE=32768)";
  ...
}
```

Note

Everything between the quotes in parms MUST be uppercase characters.

Table 9–1 SBT Logical Names

Logical Name	Description
MDMS\$SBT_ARCHIVE	This logical name is the name of the archive used during a backup of the Oracle database.
MDMS\$SBT_ARCHIVE_n	These logical names are the names of the archives used during a backup of the Oracle database when using the duplex feature of Oracle's Recovery Manager. MDMS\$SBT_ARCHIVE_1 is for copy 1, MDMS\$SBT_ARCHIVE_2 is for copy 2, and so forth. Example 9–12 is an example of using these logical.
MDMS\$SBT_IO_BLOCK_SIZE	This logical name is the size of the block that is written on the tape volume. The default size is 65024 bytes (127 * 512 bytes). You cannot specify a value larger than 65024 bytes.

Table 9–1 SBT Logical Names

Logical Name	Description
MDMS\$SBT_CATALOG	This logical name is the name of the catalog used for the following Oracle Recovery Manager commands: <ul style="list-style-type: none"> • restore • validate • list backup
MDMS\$SBT_TRACE_LEVEL	This logical allows you to define how much tracing appears in the trace file. The logical is defined in SYS\$STAR-TUP:MDMS\$SYSTARTUP.COM. By default the values TR_GENINFO and TR_MEDINFO are defined.

9.9 System Backup to Tape Restrictions

This section describes restrictions in the use of this version of SBT.

9.9.1 Doing Parallel Backups

When doing a parallel backup and you do not have all of the SBT resources needed could cause a backup not to complete. When doing a parallel backup, the controlling Oracle server does not let any of the Oracle servers doing the backup end until all servers say they have completed their task. I believe this is to keep the database consistent. However, this can cause the backup to not complete if SBT resources are not available.

Note

This is not really a SBT restriction, but it is placed here to make you aware of the possibility of parallel backups not completing.

The following three scenarios can cause a parallel backup not to complete:

- Setting the Drives List in the archive to a number of drives less than the number of parallel streams that Oracle's Recovery Manager starts. If you do not need to restrict the Oracle server backups to certain drives, I suggest you do not put anything in this attribute. If you do use this attribute, be sure you have enough drives for the number of streams in the parallel backup.
- Setting the Maximum Saves in the Archive to a number less than the number of parallel streams that Oracle's Recovery Manager starts. I suggest you set it at 36 and control the number of backups from the Oracle Recovery Manager.
- Starting more parallel backups than you have tape drives. Do not start two or more parallel backups at the same time unless you have drives available for all streams in the parallel backups.

To illustrate the problems stated above, I will give an example of setting the Maximum Saves in the archive to 2 and then using a Oracle Recovery Manager script that starts 3 parallel streams. Each stream starts the backup, the first two are allowed to get a volume set and start their backup. The third server starts up and finds that there is no volume set available at this time. It keeps trying to get a volume set. In the mean time, the first two servers finish doing the backup. However, the controlling server will not let them end until the third server is finished. Therefore, they hold on to the resources: drive and volume set. The third server can not finish for lack of resources. Any of the above problems stated above can cause this deadlock.

System Backup to Tape for Oracle Databases

9.10 Troubleshooting Tips

9.9.2 Piece Name Length Greater than 254 Characters

A piece name length greater than 254 characters is not supported. This restriction may be lifted in subsequent versions to 511 characters for a piece name length.

9.9.3 Using RDF Drives with SBT

You can NOT use Remote Device Facility (RDF) drives when using Oracle's tape I/O slaves. RDF drives may be used with SBT if you are not using Oracle's tape I/O slaves.

9.10 Troubleshooting Tips

This section gives you some tips that may help you troubleshoot SBT if you have problems.

9.10.1 Using the logical MDMS\$SBT_TRACE_LEVEL

The logical MDMS\$SBT_TRACE_LEVEL allows you define how much tracing appears in the trace file, ORA_DUMP:SBTIO.LOG. By default any error that is detected by SBT, is traced in the trace file. You can define the logical MDMS\$SBT_TRACE_LEVEL to display more information. The definition of MDMS\$SBT_TRACE_LEVEL is defined in SYS\$STAR-TUP:MDMS\$SYSTARTUP.COM. If the logical is not defined on your system, see Section 9.2 to setup the logical.

There are only three values that you might want to use. The rest of the values are for support use. By default the following values are enabled:

- TR_GENINFO - general information about the backup and restore. With this value defined, you will see message about starting the backup or restore.
- TR_MEDINFO - information about drives and volumes. With this value define, you will see messages about allocating, loading, mounting, unloading, and deallocating tape volumes and drives.

As a very minimum, we suggest that you enable the above two values so that you can check the progress of backups and restores. If the backup or restore can not get a volume set, it will loop, once a minute, waiting for the volume set to become available. If a tape drive is unavailable, SBT loops waiting for a drive to become available. SBT tries to allocate the tape drive every minute. SBT reports that it can not allocate the tape drive after five minutes and then after 10 minutes and so forth.

Example 9–16 shows an example of ORA_DUMP:SBTIO.LOG with TR_GENINFO and TR_MEDINFO values enabled. Note that the SBT-00008140 is the pid, in hexadecimal, of the process doing the backup with SBT- prepended.

Example 9–16 ORA_DUMP:SBTIO.LOG with TR_GENINFO and TR_MEDINFO Values Enabled

```
EDINFO Values Enabled
SBT-00008140 12/17/01 09:27:11 Using archive RMAN_TAPE_TL875_ARCH
SBT-00008140 12/17/01 09:27:12 Starting backup of rldbtagjd_1_1 for DB:
EMPLOYEE
SBT-00008140 12/17/01 09:27:12 Using catalog ORACLE_DB
SBT-00008140 12/17/01 09:27:12 Allocated drive: MKC200 Device: MOE$MKC200:
SBT-00008140 12/17/01 09:27:13 Drive is in jukebox TLZ875
SBT-00008140 12/17/01 09:27:13 Allocated volume AIF078
SBT-00008140 12/17/01 09:27:17 Unloading drive MKC200
SBT-00008140 12/17/01 09:27:46 Loading volume AIF078 on drive MKC200
SBT-00008140 12/17/01 09:28:39 Deallocating drive MKC200
SBT-00008140 12/17/01 09:28:39 Initializing volume AIF078
SBT-00008140 12/17/01 09:29:00 Allocated drive: MKC200 Device: MOE$MKC200:
SBT-00008140 12/17/01 09:29:00 Drive is in jukebox TLZ875
SBT-00008140 12/17/01 09:29:00 Loading/mounting volume AIF078 on drive MKC200
SBT-00008140 12/17/01 09:29:00 Loading volume AIF078 on drive MKC200
SBT-00008140 12/17/01 09:29:09 Mounting volume AIF078 on device MOE$MKC200:
```

System Backup to Tape for Oracle Databases 9.10 Troubleshooting Tips

```
SBT-00008140 12/17/01 09:29:18 Ready to write to saveset 2001121709291582.  
on volume AIF078  
SBT-00008140 12/17/01 09:30:47 Using catalog ORACLE_DB  
SBT-00008140 12/17/01 09:30:48 Finished writing saveset 2001121709291582.  
on volume AIF078
```

We hope the only thing that might not be obvious in Example 9–16 is the saveset name. We are limited to 17 characters so I change piece name, r1dbtgid_1_1, into a saveset name, 2001121709291582., on a particular volume.

Another value that you may want to use is TR_TAPSTAT. This gives you statistics about the reading/writing to a tape volume. Example 9–17 shows an example of tape statistics.

Example 9–17 Trace of Tape Statistics

```
SBT-00008140 12/17/01 09:50:49 I/O Statistics:  
SBT-00008140 12/17/01 09:50:49      DB block size: 262144 bytes 1  
SBT-00008140 12/17/01 09:50:50      I/O block size: 65024 bytes 2  
SBT-00008140 12/17/01 09:50:50      Total I/Os: 620 3  
SBT-00008140 12/17/01 09:50:50      Total I/O wait: 23279 milliseconds 4  
SBT-00008140 12/17/01 09:50:50      Maximum I/O wait: 936 milliseconds 5  
SBT-00008140 12/17/01 09:50:50      Average I/O wait: 37 milliseconds 6  
SBT-00008140 12/17/01 09:50:50      Total Kbytes: 40300 7  
SBT-00008140 12/17/01 09:50:50      Total bytes: 40314880.000 8  
SBT-00008140 12/17/01 09:50:50      Total seconds: 23 9  
SBT-00008140 12/17/01 09:50:50      MBytes/sec: 1.752 10  
SBT-00008140 12/17/01 09:50:50      Bytes/sec : 1752820.870 11
```

The following describes the entries in Example 9–17:

1. DB block size:-this is the size of blocks that Oracle sends to SBT.
2. I/O block size:-this the size of blocks that SBT writes to the I/O device. You can change this with the logical MDMS\$SBT_IO_BLOCK_SIZE which can be specified in an Oracle Recovery Manager script or as a system wide logical.
3. Total I/Os:-total I/Os to write the I/O block size blocks to the I/O device. In this example: 620 total I/Os writing 65024 byte blocks to the I/O device.
4. Total I/O wait:-this is the number of milliseconds that SBT waited while the blocks are being written to the I/O device.
5. Maximum I/O wait:-this is the longest wait that SBT made to write a block to the I/O device.
6. Average I/O wait:-this the average wait of the total I/Os for the total I/O wait time.
7. Total Kbytes:-total Kbytes transferred to the I/O device. It is a truncated value. See total bytes below.
8. Total bytes:-total bytes transferred to the I/O device.
9. Total seconds:-this is the number of seconds that SBT waited while the blocks are being written to the I/O device. This does not have the time that Oracle took to provide the information to write to the I/O device. The value is truncated from the total I/O wait in milliseconds.
10. Mbytes/sec:-Mbytes per second transferred based on total Kbytes and total seconds.

System Backup to Tape for Oracle Databases

9.10 Troubleshooting Tips

11. Bytes/sec:-bytes per second transferred based on total bytes and total seconds.

Note

These values are the time that SBT waited for the I/O device to give back control. The values have nothing to do with how long it took Oracle to give information to SBT.

9.10.2 Fatal Internal Error

If you should ever get an Internal error, you should report it to the customer support center. Example 9–18 shows an example of a fatal internal error. Anyplace that SBT could have received a value that it did not expect, it is captured and reported as a fatal internal error.

Example 9–18 Fatal Internal Error Example

```
SBT-00001DB2 12/17/01 13:53:45 Internal error
SBT-00001DB2 12/17/01 13:53:45 Extended Status:
The invalid archive type, 23
SBT-00001DB2 12/17/01 13:53:57 ABEND (abnormal end ) was passed
in sbtend flag
```

9.10.3 Check ORA_DUMP:SBTIO.LOG for Errors

Any time you receive an error in Oracle's Recovery Manager that is related to SBT, you should check ORA_DUMP:SBTIO.LOG for errors. We are limited to 250 characters returned to the Oracle Recovery Manager. Therefore, you may not receive all of the information about the error. However, I am not limited to what I put in the ORA_DUMP:SBTIO.LOG file. So be sure to look in ORA_DUMP:SBTIO.LOG when getting an error message that is not all there.

Example 9–19 shows an example of a fatal error reported in Oracle's Recovery Manager. Example 9–20 shows what is reported in ORA_DUMP:SBTIO.LOG for the same error. The File:, Function: and line are information for me to troubleshoot the problem if it is a software error. In this case, there were no volumes available. Look for it in the ORA_DUMP:SBTIO.LOG file. In this case, you can look in the archive to see what the media type, pool, and location were.

Example 9–19 Fatal Error in Oracle's Recovery Manager

```
ORA-27028: skgfgcre: sbtbackup returned error
ORA-19511: Error received from media manager layer, error text:
Fatal media movement error
Failed to allocate tape volume
MDMS object: RMAN_TAPE_TL875_ARCH
System Error: %MDMS-E-NOVOLUMES, no free volumes match selection criteria
%MDMS-E-NOVOLUMES, no free volumes match selection criteria
%MDMS-I-NOVOLSPPOOL, no
```

Example 9–20 Fatal Error in ORA_DUMP:SBTIO.LOG

```
SBT-0000819F 12/17/01 14:10:20 Fatal media movement error
SBT-0000819F 12/17/01 14:10:20 Extended Status:
Failed to allocate tape volume
MDMS object: RMAN_TAPE_TL875_ARCH
System Error: %MDMS-E-NOVOLUMES, no free volumes match selection criteria
%MDMS-E-NOVOLUMES, no free volumes match selection criteria
%MDMS-I-NOVOLSPPOOL, no free volumes in the specified pool were found
File: WRK$ROOT:[SRC]MDMS_SBT_API_MEDIA.C;1, Function:
sbt_media_allocate_volume,
Line 416
Failed to allocate volume with attributes:
Pool:
Media Type: TLZ88M
Location: 110281
```

9.10.4 Using Tape I/O Slaves

When using tape I/O slaves, you may not receive the error message from SBT in Oracle's Recovery Manager. Example 9–21 shows the type of error you may get reported for a tape volume being offsite. Example 9–22 shows the tape volume offsite error reported when not using tape I/O slaves. So when troubleshooting, also look in the trace log file ORA_DUMP:SBTIO.LOG. In both cases the error was reported in the trace log file.

Example 9–21 Volume Offsite Error using Tape I/O Slave

```
RMAN-08031: released channel: t1
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03006: non-retryable error occurred during execution of command: vali
ate
RMAN-07004: unhandled exception during command execution on channel t1
RMAN-10035: exception raised in RPC: ORA-00447: fatal error in background
process
RMAN-10031: ORA-19583 occurred during call to DBMS_BACKUP_RESTORE.RESTORE
ACKUPP
IECE
RMAN>
```

Example 9–22 Volume Offsite Error Not using Tape I/O Slave

```
RMAN-07004: unhandled exception during command execution on channel t1
RMAN-10035: exception raised in RPC: ORA-19507: failed to
retrieve sequential fi LE, handle="6sd8vntq_1_1", parms=""
ORA-27029: skgfrtrv: sbtrestore returned error
ORA-19511: Fatal catalog access error
Piece 6sd8vntq_1_1 cannot be restored because volume AHI164 is offsite
RMAN-10031: ORA-19624 occurred during call to
DBMS_BACKUP_RESTORE.RESTOREBACKUPPIECE
RMAN>
```

Example 9–23 Volume Offsite Error in Trace File

```
SBT-00000889 11/13/01 16:28:02 Fatal catalog access error
SBT-00000889 11/13/01 16:28:02 Extended Status:
Piece 6sd8vntq_1_1 cannot be restored because volume AHI164 is offsite
```


This chapter describes in more technical details the ABS and MDMS infrastructure and implementation.

10.1 The Server Process

Each OpenVMS node participating in an MDMS Domain runs a generic process called MDMS\$SERVER.

Each MDMS server process can implement 3 functions:

- Current access to the database, the **database server**
- **Forwarding a user request** to the current database server
- **Executing remote requests** on behalf of the database server

Domain

All nodes communicating with the same database server belong to the same MDMS Domain. Each MDMS Domain has its own database. Typically you have only one MDMS Domain in your network. But the architecture allows to setup more than one domain. However, one has to make sure that none of the nodes and none of the MDMS objects (i.e jukeboxes) are used in more than one domain.

10.1.1 The Database (DB) Server

10.1.1.1 Database

MDMS keeps all its permanent settings in files in a location defined by logical MDMS\$DATABASE. The summary of these files are called the MDMS Database.

Each MDMS server needs access to the MDMS database before it is fully functional. The server translates logical name MDMS\$DATABASE_SERVERS which contains a list of potential database server nodes. This logical is defined in MDMS\$SYSTARTUP.COM and contains the **network names** of other servers. Because the server has not yet accessed the database it cannot use an MDMS node name.

While scanning through the database servers list the server tries to contact the remote server using the appropriate network for a given network name:

- DECnet, if only alphanumeric characters, e.g. “STAR”
- DECnet-Plus, if network name contains “:.”, e.g. “VMS:.STAR”
- TCP/IP, if network name contains just dots “.” and a possible colon “:” followed by a number range, e.g. “star.vms.com” or “star.vms.com:2501-2510”

Architecture

10.1 The Server Process

Because the database server list is processed from left to right one can control the order by which server nodes are tried and which network to use. Choosing a network at this point is unrelated to how the node's transport is defined in the MDMS database. The requesting node and the contacted node must have the network for this server entry enabled otherwise the contact fails and the server continues on with the next entry in the list. The failed attempt is logged in the MDMS server logfile ("MDMS\$LOGFILE_LOCATION:MDMS\$LOGFILE_<node>.LOG" or "MDMS\$LOGFILE_LOCATION:MDMS\$LOGFILE_DBSERVER.LOG").

10.1.1.2 Becoming a DB Server

The MDMS server tries to match an entry in the database server list with one of its own network name definitions. The network name definitions are obtained by retrieving the following values or translating logicals:

- SCSNODE sysgen parameter
- SYS\$NODE for DECnet, stripping off the trailing "::"
- SYS\$NODE_FULLNAME for DECnet-Plus, stripping off the trailing "::"
- {UCX|TCPIP}\$INET_HOST and {UCX|TCPIP}\$INET_DOMAIN for TCP/IP, concatenating the two strings using a dot "." in between
- MDMS\$SERVER, if none of the above are available

If the server finds a match it tries to open the database files. If it successfully opens all the database files it declares itself the database server. Because the files are opened for exclusive write and shared read, no other MDMS server can open the database files after that.

A server remains to be a database server until it exits. At this point the database files are closed and the domain is without a database server until the next server has successfully opened the database files.

If the server finds the files already open it continues on with the search for a DB server.

10.1.1.3 Finding another DB Server

When contacting another server, the server passes all its network names on to the other node. If the other node happens to be a DB server it verifies that the requesting node is defined in the MDMS database. Only when **all** the node's network names are defined in the node's object the DB server grants access to the requesting node. Otherwise the DB server returns a MDMS_NODENOTENA ("node not in database or not fully enabled").

Once the node is granted access to the DB server the node updates its setting from the database. At this point the TRANSPORT setting of the node is in use. For example it is possible that a server contacted the DB server via DECnet but when it updates its TRANSPORT setting it is only allowed to use TCPIP. So from that point on this server only uses TCPIP to "talk" to the DB server.

Typically all nodes in a domain have the same definition of MDMS\$DATABASE_SERVER in their MDMS\$SYSTARTUP.COM. But the definitions do not have to match. For example each node could list itself first in the list to give a more round-robin behavior.

10.1.1.4 Failover of the DB Server

Once a MDMS server loses contact to the DB server it starts to search for a new DB server using its own search list in MDMS\$DATABASE_SERVER. The server tries the whole search list three times. The search for the DB server finally ends with either:

- a. the node became the DB server itself
- b. the node found another DB server

- c. the request failed with MDMS_NODBACC (“no access to database server”)

Once a new DB server has been established, all nodes start to forward requests to this server.

10.1.1.5 Role of the DB server

The DB server receives all user requests in an MDMS Domain. It coordinates all activities and accesses the MDMS database files. The DB server uses a writethrough cache to access the database. All database files are RMS index-sequential files and their key layout is defined by “.FDL” (File Definition Language) files in MDMS\$SYSTEM.

Most user requests can be executed entirely on the DB server. In some cases the DB server has to send remote requests to other servers in the domain. For example remote load volume requests or remote scheduling requests.

10.1.2 Server Communications

An MDMS server can establish three types of listeners:

- The Mailbox Listener
- The DECnet Listener
- The TCP/IP Listener

The Mailbox Listener is always enabled. The server receives user request through its mailbox described by logical MDMS\$MAILBOX. Each user process has its own mailbox to receive the response from the server.

The DECnet Listener is enabled during server startup if DECnet is available on this node indicated by the existence of logical name SYS\$NODE or logical SYS\$NODE_FULL_NAME. Once the server had access to the database and DECNET is not defined in its TRANSPORT setting the server shuts down the DECnet Listener.

The TCPIP Listener is enabled during server startup if TCP/IP is available on this node indicated by the existence of logical names {UCX|TCPI}\$INET_HOST and {UCX|TCPI}\$INET_DOMAIN. Once the server had access to the database and TCPIP is not defined in its TRANSPORT setting the server shuts down the TCPIP Listener.

Startup and shutdown of the listeners is logged in the MDMS server logfile. Also the “MDMS SHOW SERVER” display shows the current servers network names at the top and its current TRANSPORT setting which reflects the active network listeners.

Even though a DB server has received a request via DECnet it could use TCPIP to request a remote operation (e.g. load volume) at a third node. It all depends on the TRANSPORT setting of the individual nodes.

10.2 Scheduler Interface

MDMS calls the scheduler interface from the MDMS DB server process.

10.2.1 Option INT_QUEUE_MANAGER

MDMS uses the programming interface to the OpenVMS Queue Manager. A thread in the MDMS DB server submits due requests to the OpenVMS Queue Manager. The request will be submitted to batch queue ABS\$<execution_node>. If the batch queue is available on the local node or is within the OpenVMS cluster the MDMS DB server calls the local Queue Manager. For remote nodes the MDMS DB server forwards the request to the remote MDMS server on the execution node of the request. The remote MDMS server then submits the request to the local batch queue ABS\$<execution_node>.

Failures to call the OpenVMS Queue Manager will be logged in the servers’ logfiles.

Architecture

10.3 Catalogs

10.2.2 Option EXT_QUEUE_MANAGER

This option uses the same method as INT_QUEUE_MANAGER to schedule jobs locally or remote. But instead of calling the programming interface of the OpenVMS Queue Manager, a subprocess is created from the MDMS server process to run the command procedure MDMS\$SYSTEM:MDMS\$EXT_QUEUE_MANAGER.COM. The command procedure issues the DCL commands to create, delete, modify and show batch jobs. Also the command procedure has to return status about the commands and in some cases additional information. See the command procedure template file, MDMS\$SYSTEM:ABS\$EXT_QUEUE_MANAGER.TEMPLATE for more details.

Failures to execute the command procedure will be logged in the servers' logfiles. Each activation of the command procedure creates a logfile of MDMS\$LOG:MDMS\$EXT_QUEUE_MANAGER_<request_name>.LOG. The request name portion of the logfile name maybe truncated to a valid OpenVMS file specification.

10.2.3 Option EXT_SCHEDULER

This option uses the same method as EXT_QUEUE_MANAGER to interface with the scheduler. A subprocess is created to run the command procedure ABS\$SYSTEM:ABS\$EXT_SCHEDULER.COM. The command procedure issues the DCL commands to create, delete, modify and show jobs for third party scheduler product. Also the command procedure has to return status about the commands and in some cases additional information. See the command procedure template file MDMS\$SYSTEM:MDMS\$EXT_SCHEDULER.TEMPLATE for more details. In contrast to option EXT_QUEUE_MANAGER, ABS assumes that the third party scheduler product reschedules all requests locally and remote. So MDDBS will not call the scheduler if a request is due to run.

Failures to execute the command procedure will be logged in the servers' logfiles. Each activation of the command procedure creates a logfile of MDDBS\$LOG:MDMS\$EXT_SCHEDULER_<request_name>.LOG. The request name portion of the logfile name maybe truncated to a valid OpenVMS file specification

10.3 Catalogs

ABS can have multiple catalogs. Each catalog is comprised of three RMS Indexed Sequential Files:

- <catalog_name>_%TLE.DAT - Transaction Log Entry
- <catalog_name>_%AOE.DAT - Archive Object Entry – not used for FULL_RESTORE catalog type
- <catalog_name>_%*AOE_INSNC.DAT - Archive Entry Object Instance - not used for FULL_RESTORE catalog type, one file per volume set if VOLUME_SET catalog type

These files must reside in the same directory. Different catalogs can be in different directories or different disk volumes.

The **Transaction Log Entry** file contains two entries per save request executed. It contains among other data the save set name, the tape's volume ID and the expiration date of the save set. Depending on record compression the average record size on disk is about 300 bytes. Information in a transaction log entry can be displayed by showing catalog save entries.

The **Archive Object Entry** file contains one entry for each file backed up. It contains among other data the device and file name. Depending on record compression and depending on actual filename sizes the average record size on disk is about 300 bytes.

The **Archive Object Entry Instance** file contains an entry for every time a file is backed up. It does not contain the filename but a back pointer to the record in the AOE. Depending on record compression the average record size on disk is about 200 bytes. For a VOLUME_SET catalog type there is one file per volume set in use. The volume set name is part of the instance file name.

Information in the archive object entry and the archive object entry instance can be displayed by showing catalog file entries which contains information from both files.

10.3.1 Catalog Sizes

TLE: This grows to the average size of how many save requests are active.

- This file does not have size problems
- Low volatility to deletes
- 300 bytes times number of active save requests times retention period in days + some record overhead.

AOE: This grows to the number of files that are actively being backed up

- Medium volatility to deletes
- 300 bytes times number of active files + some record overhead

AOE_INSNC or AOEI: This can grow very large.

- Sized is based on how many files are being backup up and how long the retention time on the file is.
- High volatility to deletes.
- 200 bytes times average number of files backed up per day times the retention period in days.

Example 10–1

- 1 disk volume with 40,000 files
- full saves every week (40,000 files)
- incrementals 6 times a week (estimate 2,000 files/day)
- retention is 30 days for all backups
- TLE $300 \times 7 \times 30 = 63\text{K}$ bytes
- AOE: $300 \times 40,000 = 12 \text{ MB}$
- AOE_INSNC: $200 \times 7428 \times 30 = 44 \text{ MB}$

Example 10–2

- 1 disk volume, 40,000 files
- full saves every night (40,000 files)
- retention is 30 days for all backups
- TLE: small
- AOE: $300 \times 40,000 = 12 \text{ MB}$
- AOE_INSNC: $200 \times 40,000 \times 30 = 240 \text{ MB}$

Architecture

10.4 Coordinator

Example 10–3

- 10 disk volumes, total of 400,000 files
- full saves every week (400,000 files)
- incrementals 6 times a week (20,000 files)
- retention is 30 days for all backups
- TLE: small
- AOE: $300 \times 400,000 = 120 \text{ MB}$
- AOE_INSNC: $200 \times 74285 \times 30 = 445 \text{ MB}$

Example 10–4

- 10 disk volumes, 400,000 files,
- full saves every night (400,000 files)
- retention is 365 days for all backups
- TLE: small compared to rest
- AOE: $300 \times 400,000 = 120 \text{ MB}$
- AOE_INSNC: $200 \times 400,000 \times 365 = 29 \text{ GB}$
- ...and if you had 100 volumes: AOE_INSNC is 292 GB!!!

As you can see from example 10-4, catalogs can become quite large. Changing the backup schedule so that less files are saved and using shorter retention periods helps to maintain smaller catalogs. If this cannot be achieved extra disk space should be reserved for the ABS catalogs with space for future expansion.

10.4 Coordinator

The coordinator process is created when a SAVE or RESTORE request is scheduled to run. It starts out as a single process in a batch or scheduler job executing `ABS$$SYSTEM:ABS$COORDINATOR.COM`. This process prepares the drive and media for the individual backup agent to move the data. Once the media is ready to be used the coordinator spawns a subprocess using a Pseudo Terminal device to communicate with the subprocess.

The coordinator then “feeds” DCL commands to the subprocess which finally contains the command to execute the backup agent (e.g. `OpenVMS BACKUP`).

All output by the subprocess is received by the coordinator and checked against entries in the template files in `ABS$TEMPLATES`. Each backup agent has its own set of template files for the different type of save or restore operations. Even though these files can be changed it is not recommended. The original files have been checksummed for each release and any modification will be noted in the ABS save or restore logfile.

The coordinator starts a separate subprocess for each selection. If the `SEQUENCE OPTION` of the save or restore is set to `SEQUENTIAL` the coordinator will not start the next subprocess before the current one has completed. With `SEQUENCE_OPTION OVERLAPPED` the next subprocess will be started as soon as the backup agent in the current subprocess has reached a point where the archive (i.e. drive) is no longer needed. This is defined internally for each backup agent. For example `OpenVMS BACKUP` releases the tape drive being used while it executes the recording pass when `/RECORD` was specified.

10.4.1 Coordinator Cleanup

The coordinator cleanup process (“ABS\$COORD_CLEAN”) is responsible to cleanup after a failed save or restore request. It needs to run all the time to perform this task.

Each save or restore request enters a cleanup record into file ABS\$SYS-TEM:COORD_CLEANUP.DAT. The record contains:

- the PID of the process executing the save or restore
- the archive being used

The cleanup process reads this file every minute. If it finds an entry for which the PID field refers to a non-existent process it releases the volume set used in the archive so it can be used again.

10.4.2 Volume Sets

To synchronize access to volumes in a volume set ABS keeps pseudo volume records in the volume database. The pseudo volume starts with “&+” and the volume ID of the first volume in the set. To show the pseudo volumes you have to use the /ABS_VOLSET qualifier. The fields in the volume record are used as follows:

- **Brand:** PID of process which has the volume locked or locked the last time. Do not change!
- **Description:** A reservation bitmap displayed as a 32 hex-digit value. The low-order bit is the general locking bit which means the volume set is in use while the other bits represent which relative volume in the set is used for a write operation. For troubleshooting purposes this can be set to an all zero value by specifying exactly 32 zeroes.
- **Length:** Currently last volume in set by number. Do not change!
- **Mount Count:** Number of savesets on volume set. Do not change!
- **Pool:** The EOT tapemark position expressed in number of tapemarks and a version number. Do not change!

11.1 Save and Restore Requests

11.1.1 Notification of Save/Restore Completion

The first step to checking the status of save and restore requests is by using the notification options in the environment object. You may set several levels of notification which include start, complete, warning, error and fatal. The notification may be sent by OPCOM or by mail. If you have notification options set, you will receive notification when problems occur with your save and restore requests (or a message about start or completion).

In the MDMS GUI, doing a show of the save or restore request will display the last status of the request. A green (success) or red (error) box will be displayed in the upper right corner of the show output.

11.1.2 Log Files

Each save and restore request creates a log file in the ABS\$LOG directory when it is run. The log file is named by the request name. This log contains information about the request, the media management activities, the backup command and any output from the backup process. If errors occur it also contains trace information about the error. The last error message generally contains the actual cause of the error.

11.1.3 Logical Names

There are some logical names which may be defined at a system level which will cause ABS to log more information in the request log files. You should not set these logical names unless advised to by a COMPAQ customer support representative because the log files can grow quite large if you use them.

11.1.4 Alpha Stack Size Logical

If you are running your save/restore request on an OpenVMS Alpha system and you see either ACCVIO or CMA-F-EXCCOP errors in the logs, there is a stack size variable which may eliminate the problem. ABS\$COORD_ALPHA_STACKSIZE may be used to increase the stack size beyond the 65536 default. To use the logical, define it at system level to a value which is a multiple of 8192

```
$ DEFINE/SYSTEM ABS$COORD_ALPHA_STACKSIZE 8192 * x
```

11.1.5 Fast Skip Errors

If you receive an ABS_SKIPMARKS_FAILED error there is a logical name which may be defined at system level which turns off the ABS fast skip methods. To disable fast skip do the following command on the affected system:

```
$ DEFINE/SYSTEM ABS_NO_FAST_SKIP TRUE
```


A list of supported classes is available in MDMS HELP or in the MDMS Reference Manual.

To enable OPCOM on a terminal so that you may see and reply to the messages, type

```
$ SET REPLY/ENABLE=opcom_class
```

To disable OPCOM, type

```
$ REPLY/DISABLE
```

Operator privilege is required in order to enable OPCOM.

These message are particularly useful when an ABS save or restore request is hung waiting for volume. If MDMS is having difficulty obtaining or loading a volume the OPCOM message may be helpful in determining the problem.

11.2.3 MDMS Requests

Whenever an MDMS request is issued, you may view them using the command

```
$ MDMS SHOW REQUESTS
```

Or, you may view the requests by selecting the request tab in MDMSView GUI.

If a request is stalled for some reason you may be able to determine the problem by viewing the request. It is also useful to look in the MDMS\$LOGFILE_<node>_.LOG or MDMS\$LOGFILE_DBSERVER.LOG files.

11.2.4 Scheduling Problems

The MDMS database server acts as the scheduler for all ABS and MDMS schedules. The schedules are viewable by using the command

```
$ MDMS SHOW SCHEDULES
```

The MDMS domain contains the type of scheduling that you are using (Internal, External or Scheduler). In the MDMS\$LOG:MDMS\$LOGFILE_DBSERVER.LOG file, there will be a RUN SCHEDULE command for each schedule executed. If save/restore requests, or MDMS scheduled activities fail to run, there are several ways to track down the problem.

11.2.4.1 Internal Scheduling

If you are using the INTERNAL scheduler type there are log files generated in the MDMS\$LOG directory called MDMS\$RUN_<#>.LOG. These files contain information about every schedule that is run. You can search these files for the name of the request that you were expecting to be run

11.2.4.2 External Scheduling

If you are using the EXTERNAL scheduler type, MDMS invokes a command procedure to scheduler the job into a batch queue. This command procedure is called MDMS\$SYSTEM:MDMS\$EXT_QUEUE_MANAGER.COM. There are log files generated in the MDMS\$LOG directory called MDMS\$EXQ_<requestname>.LOG which contain the output from a set verify on the command procedure. These logs may give you information about errors generated when the job is being inserted into the batch queue. If you have modified the command procedure they may be especially useful for debugging your procedure.

11.2.4.3 Scheduler Scheduling

If you are using the SCHEDULER scheduler type, MDMS invokes a command procedure to schedule the job in the DECscheduler (or another scheduler product, if you have modified the command procedure).

Troubleshooting

11.3 MDMSView GUI

This command procedure is called MDMS\$SYSTEM:MDMS\$EXT_SCHEDULER.COM. There are log files generated in the MDMS\$LOG directory called MDMS\$EXS_<request-name>.LOG. They contain the output from a set verify on the command procedure. These logs may give you information about an error generated when the job is being inserted into the scheduler and may be especially useful if you have modified the command procedure and are debugging.

11.2.5 MDMS Scheduled Activities

There are six MDMS scheduled activities scheduled daily.

MDMS\$DEALLOCATE_VOLUMES

MDMS\$DELETE_RESTORES

MDMS\$DELETE_SAVES

MDMS\$MOVE_MAGAZINES

MDMS\$MOVE_VOLUMES

MDMS\$PURGE_LOGS

Each one of these generate a log depending on which scheduler type you are using (see above). If errors occur there may be information in these log files or in the MDMS\$LOG:MDMS\$LOGFILE_<node>_.LOG file.

11.3 MDMSView GUI

11.3.1 Running MDMSView GUI After ABS/MDMS Installation

After installing ABS/MDMS, you must logout and back in before running the MDMSView GUI on OpenVMS Alpha. Some of the required symbols for Java will be missing if you do not log back in and you may receive errors from the MDMS\$SYSTEM:MDMS\$START_GUI.COM procedure.

11.3.2 Windows Java Path

If you have Java installed in a location different from the normal default location, the GUI will not find Java. You must edit the MDMSView.bat file and include the correct path. The default in this file is

```
C:\Program Files\JavaSoft\JRE\1.2\bin\java.exe
```

11.3.3 MDMSView Log Screen

If you receive errors while running the GUI, there is a log screen that may be displayed. This window may show more information about the errors. This window comes up with the GUI by default and you can bring it up to the foreground by selecting MDMSView Log Screen from the View pulldown. The information displayed are the actual calls the GUI is sending to the MDMS server.

11.3.4 MDMSView Command Window

The window that initially brings up the GUI has additional information in it. This displays the Java error messages and operations.

11.3.5 MDMS\$LOGFILE_*.LOG

The MDMSView GUI is generating requests to the MDMS server, so any problems may be logging errors into the MDMS\$LOG:MDMS\$LOGFILE_<node>.LOG or MDMS\$LOG:MDMS\$LOGFILE_DBSERVER.LOG files. If you receive an MDMS error window when executing an action, check these files for errors.

If you receive an error MDMS\$ERROR, this means that the MDMS server did not respond correctly to the request. This error may need to be reported to COMPAQ.

11.4 ABS Catalogs

11.4.1 Staging Unpack

If you are using ABS catalogs which are set to use staging, the save/restore request logs will contain information about the staging files, the command procedure used to unpack the file, and the log file generated by the unpack process. The log file is generated in the ABS\$LOG directory and contains information about the unpack process. If there were errors during the unpack mail will be sent to the person(s) named as the MAIL recipient(s) in the MDMS domain.

If errors occur, the ABS\$CATALOG:*.STG and ABS\$CATALOG:*.COM files are not deleted. You may run the *.com file as a batch job with ABS as the user. This allows you to unpack the files once you have determined the reason that they failed. Some reasons may be that the catalog disk is full, the system went down, etc.

If there are errors in the unpack logs which indicate an error with the ABS\$CATALOG_UNPACK_STG program, you should report this problem to COMPAQ.

11.4.2 Catalog Cleanup

To clean expired entries from the catalog there is a process which runs in the ABS\$node batch queue called ABS_CLEAN_CATLG_node. This process is scheduled to run once a day at 12:30 pm. The scheduled time is set in the ABS\$SYS-TEM:ABS\$START_CATALOG_CLEANUP.COM procedure. You may modify the start time for the job or change the frequency of the job. If you do not have lot of expired entries daily, you may want to run the job less frequently.

The log file generated by this cleanup process is called ABS\$LOG:ABS\$CATALOG_CLEANUP.LOG. A lot of information about how many records were read from the catalog, how many were deleted, and any errors are kept in this log. Most errors seen should be reported to COMPAQ.

11.5 SWindows and Unix Clients

11.5.1 Windows Log File

Should you encounter problems when saving or restoring data using ABS for an NT client system, ABS provides way to help you troubleshoot the problem. Assign a system variable on the NT client system that, in turn, creates log files about the NT client system during ABS backup operations. These log files will assist you during the troubleshooting process.

Note

Assign this system variable only when you need troubleshooting assistance. Deassign the system variable when it is no longer needed. Do not leave the system variable assigned during normal, day-to-day operations. Because the log files can become extremely large, leaving the system variable assigned could cause performance problems.

Troubleshooting

11.5 SWindows and Unix Clients

To assign the system variable, use the procedure in Table 11–1.

Table 11–1 Assigning a System Variable for NT Troubleshooting

Step	Action
1.	Log into the administrator account on the NT client node.
2.	Bring up the registry editor (for example, regedt32 from command line)
3.	Go to the window for the following location: <code>HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ - ABSClient\Parameters</code>
4.	From the EDIT menu select “Add Value...”, enter ABSGtarLog as the value name.
5.	Select the data type as REG_DWORD.
6.	Enter a one (1) as the data in the DWORD window; select decimal.
7.	Click OK, exit from the registry.
8.	Run your save and restore requests as usual. Result: The log files generated during the save or restore operation will be located in the system directory. For example, on a NT Version 4.0 server system, the directory system name would be: <code>c:\Winnt40S\system32</code> The log files are named as follows: <ul style="list-style-type: none">• <code>abs_log_file.txt</code> - This log file contains information about the execution of the file <code>absgtar.exe</code>.• <code>absclient_log_file.txt</code> - This log file contains information about the execution of the file <code>absclient.exe</code>.
9.	When the log files are no longer needed, go to the same registry window and delete the entry. Do this by highlighting the entry; select Delete from the EDIT menu.

11.5.2 Windows Quotas

If you expect to have continuation volumes during your Windows backups, you should set the following parameter in the registry.

Modify the registry path

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters
```

with the following Windows parameter (set it to 20 or higher)

```
TcpMaxDataRetransmissions REG_DWORD 20
```

This change to the default built into Windows ensures that the TCPIP connection is not prematurely terminated with send failures.

Note

After making the changes to the parameter you need to reboot the system to allow the changes to take effect.

11.5.3 Permission Denied Errors

If you receive these errors during a Windows save, you will need to set access to files on the Windows system:

```
ABSgtar: can't add file C:\AFILE.EXT: Permission Denied
```

```
ABSgtar: can't open directory C:\ADIR: Invalid Command
```

The cause of these errors is either the directory or file is open for write access by a user or application, or the system has been denied read access to the file or directory. ABS runs under the system account.

To get around this error, close all open files or set the access on the files for the SYSTEM account.

To set the file access, select the file from a fileview window. Select Properties from the File pull-down. Click the Security tab and then select Permissions. Select Add and highlight SYSTEM. Add the types of access (full control is best so you can restore the files). Click the Add button. This gives the SYSTEM account access to the files.

11.5.4 UBS FAILURE

If you received %ABS --UBS FAILURE -- errors, you may have your TCP/IP parameters set too low on the OpenVMS system executing the backup. To view the parameters issue the command

```
$ TCPIP SHOW PROTOCOL/PARAMETERS
```

The receive and send parameters should be set to 50,000 or higher. If they are not, change them by using the command

```
$ TCPIP SET PROTOCOL TCP/QUOTA=(SEND=50000,RECEIVE=50000)
```

Note

If you have to reboot the machine, make sure that you reset these values after rebooting.

11.5.5 Considerations for Saving Large Disks on UNIX and NT Clients

ABS stores data on tape based on ANSI Standard X3.27-1987, File Structure and Labeling of Magnetic Tapes for Information Exchange. This standard requires that the block length (number of bytes per block for a file) be stored in the header section and the block count (number of blocks in a file) be stored in the end of file section. Together these fields determine the maximum number of bytes that the file contains on tape. So, in theory the following formula is implemented:

$$\text{block length} * \text{block count} = \text{number of bytes}$$

These fields on tape are stored in an ASCII format with the block length being five digits, and the block count being six digits. This allows for a maximum save request disk size of 99999 * 999999 = 99,998,900,001 bytes (approximately 99 gigabytes (GB)).

ABS uses a default block length of 10240 bytes/block when it stores data to tape. As a result, the maximum disk size by default is 10240* 999999 = 10,239,989,760 (approximately 10 GB). If the actual number of bytes exceeds this amount, then ABS\$UBS will raise the following assertion and the save request will fail:

```
assert error: expression = section_block_count <= 999999
```

Troubleshooting

11.5 SWindows and Unix Clients

The value of the block length is specified to the underlying gtar backup engine as a blocking factor. The blocking factor is defined as a multiple of 512 bytes. The default block length passed to gtar is “-b20”. To determine an appropriate blocking factor or block length for a specific situation, follow these steps:

Step 1. Divide the size of the disk (in bytes) by 99999

Step 2. Divide the resulting number by 512

Step 3. Round up to the next whole number

For example, if the disk size is approximately 30,000,000,000 bytes (30 GB), use the following formula:

$$30,000,000,000 / 99999 / 512 = 58.59 \text{ or } 59$$

This results in a blocking factor of “-b59”.

You can modify the default block length from the GUI for an NT or UNIX save or restore request on the Agent Qualifiers window (see Section or Section). Specify this value in the Agent Qualifiers window.

Restriction:

ABS will not produce the correct results if the value exceeds “-b127”. If the disk is large enough to exceed this amount, create more than one save request for that particular disk.

To modify the blocking factor, use the procedure described in Table 11–2.

Table 11–2 Modifying the Blocking Factor

Step	Action
1.	Invoke ABS GUI as described in Chapter 6
2.	Click Modify or Delete Requests & Policies
3.	Click Save Request
4.	Double-Click the save request name
5.	Double-Click on the file name in the Data To Save area
6.	Click Agent Qualifiers
7.	Click Backup Agent-Specific Qualifiers and enter the value of <i>-bnn</i> Where <i>nn</i> is a numeric value which when multiplied by 512 bytes gives the block length

Restore requirement:

When restoring data from a save request where the blocking factor has been modified, you must specify the same blocking factor that was specified on the save request. Otherwise, the restore request will fail due to an invalid block size on the tape. As a default, ABS uses 10240.

11.5.6 Files Larger than 2gb

If you are attempting to backup files larger than 2gb on a Windows or Unix system, you may receive errors indicating that the file was not saved. The Windows gtar image we provide is not able to backup these files. Also, some Unix operating systems do not, by default, compile the gtar image with the 64 bit variables required to backup the large files.

See the documentation provided with your Unix operating system for information about compiling with 64 bit variables.

11.6 RDF (Remote Device Facility)

When errors occur with RDF (RDEV) devices, you should check your RDF setup and log files in the directories pointed to by the logical names TTI_RDEV and TTI_RDF. There are log files called RDCLIENT_<node>.LOG and RDSERVER_<node>.LOG. Also see the RDF documentation and the chapter about Remote Devices in the ABS Guide to Operations for more information.

11.7 Information Required When Reporting Problems

If you report a problem to your COMPAQ support organization, the following information should be included.

- If the problem is related to a save request:
 - MDMS SHOW SAVE/SELECTIONS save
 - MDMS SHOW ARCHIVE archive
 - MDMS SHOW ENVIRONMENT environment
 - MDMS SHOW SCHEDULE schedule
 - The log file of the save request
- If the problem is related to a restore request:
 - MDMS SHOW RESTORE/SELECTIONS restore
 - MDMS SHOW ARCHIVE archive
 - MDMS SHOW ENVIRONMENT environment
 - MDMS SHOW SCHEDULE schedule
 - The log file of the restore request
 - The log file for a corresponding save request which saved the data
 - MDMS SHOW CATALOG/FILES of the data being requested in the restore request
- If the problem is related to MDMS:
 - MDMS SHOW output of the related volumes, drives, etc
 - Output from OPCOM messages issued by MDMS
 - Pertinent information from the MDMS server log
- If the problem is related to the MDMSView GUI
 - GUI version
 - Steps taken to reproduce the error
 - Error message(s)
 - MDMSView Log Screen information
 - Java version
 - MDMS Command Window information
- Other information may be required, but will be addressed as needed.

Configuration Example

Getting ABS/MDMS up and running is very easy with the MDMS objects configuration command procedure and then create a save.

First you need to setup your MDMS configuration. Using the MDMS\$ROOT:[SYSTEM]MDMS\$CONFIGURE.COM procedure you can configure your MDMS domain. However, you need the following information to start:

- Media type - TLZ06 (a media type you make up)
- Onsite location - COMP_ROOM_1 (name you make up)
- Offsite location - IRON_MOUNTAIN (name you make up)
- IP domain name for node - 78.12.53.81 (if using IP)
- Name of your jukebox - TLZ06J (a name you make up)
- Robot name - GKB601: (OpenVMS device name controlling the robot)
- Drive name - TLZ06D (name you make up)
- OpenVMS device name - MOE\$MKB600:
- Volumes - TLZ000-TLZ012 (made up name or bar code labels)

After configuring MDMS objects, then you can create a save.

The following is a sample run of MDMS\$ROOT:[SYSTEM]MDMS\$CONFIGURE.COM using the information above.

```
$ @MDMS$ROOT:[SYSTEM]MDMS$CONFIGURE.COM
```

```
MDMS Domain Configuration Procedure
```

```
Copyright (c) 2000 Compaq Computer Corporation
```

```
Use this procedure to configure MDMS for the first time or to add objects to the configuration.
```

```
Do not use this procedure to convert from MDMS V2.9x - use  
MDMS$CONVERT_V2_TO_V3.COM instead
```

```
Type "?" to any question for help
```

```
Type "??" to any question for help and list of values
```

```
Type "<return>" to any question for [default] value
```

```
Media, Device, and Management Services for ABS and HSM
```

```
Command Line Version: V4.0(439)
```

```
Shareable Image Version: V4.0(439)
```

Configuration Example

Server Version: V4.0(439)

* Have you used this procedure before [NO]: no

This command procedure prompts you to enter information that is used to configure the media and device management (MDMS) portion of your ABS and HSM environment. If you are running the procedure for the first time you should say "Yes" to "...configure all objects". If you are refining your configuration, you should say "No" to "...configure all objects", and you will be prompted for the types of objects you want to configure.

With the exception of volumes, all object names are strings consisting of the letters A-Z, the numbers 0-9 and the "_" underscore character. White space in object names is not supported. The object names must be unique in the domain and may be from 1 to 31 characters in length. Volume names have a maximum of 6 characters. You can type the answer to any question in upper or lower case and conversions will automatically be performed as needed.

There are a total of 10 types of objects in MDMS, and these are summarized as follows:

* Domain - The entire scope of MDMS operations, which can span geographic locations. There is one predefined domain which you can configure using this procedure.

* Location - A physical location, configurable as a hierarchy, that may contain volumes, nodes and jukeboxes, and is used as one selection criteria for allocating volumes and drives.

* Node - An OpenVMS computer system capable of running MDMS and accessing drives and jukeboxes.

* Jukebox - A robotic device capable of automatically loading and unloading volumes into drives. Jukeboxes contains drives and volumes, and optionally slots, ports, CAPS depending on the type of jukebox.

* Drive - A tape drive capable of supporting read and/or write operations for ABS and HSM applications.

* Pool - A logical object containing a set of volumes that can be allocated by authorized users.

* Media Type - A logical object describing a type of media associated with volumes.

* Volume - A physical piece of tape media used for storing and retrieving data.

* Group - A group of nodes with something in common (e.g. cluster members) that can be specified instead of a list of nodes.

Configuration Example

* Magazine - A logical set of volumes which are moved as a whole and are contained in a physical magazine cartridge. Magazines are not configured using this procedure. If you want to configure magazines, you should do that manually later.

You will be guided through the following configuration steps during this procedure:

1. Configure the domain - define default values applicable across the domain
2. Configure locations - define physical locations that may contain nodes, jukeboxes, magazines and volumes
3. Configure nodes - define OpenVMS nodes that run MDMS in your domain and optionally assign them to groups
4. Configure jukeboxes and drives - define jukebox devices in the domain and their associated drives and optionally inventory the jukeboxes
5. Configure standalone drives and stackers - define drives that are not contained in jukeboxes
6. Configure volumes - configure tape volumes, together with media types and pools, and optionally inventory jukeboxes and initialize volumes

*You may execute or skip any step. If you say "configure all objects" you will automatically execute all steps. However, you can always exit a step by entering <return> when asked to configure an object.

Type "?" to any question for help

Type "??" to any question for help and a list of values

Type <return> to accept the [default]

* Do you want to configure all objects [YES]: yes

Configuring domain...

* Enter domain default media type: TLZ06

* Apply to default ABS archives? [YES]: YES

* Enter domain default onsite location: COMP_ROOM_1

* Enter domain default offsite location: IRON_MOUNTAIN

* Enter domain default scratch time [365]:

* Enter domain default maximum scratch time [365]:

* Enter domain default transition time [14]:

* Enter domain default deallocation state [TRANSITION]:

* Enter domain default mail notification [SYSTEM]:

Configuration Example

- * Enter domain default OPCOM classes [TAPES]:
- * Enter domain default volume protection [SY:RW, OW:RW, GR:R]:

Configuring locations...

- * Enter a location to be configured [NONE]:

Configuring nodes...

- * Enter a node to be configured [NONE]: MOE
- * Does the node support TCPIP communications [YES]: YES
- * Does the node support DECnet communications [YES]: YES
- * Enter IP domain name for node []: 78.12.53.81
- * Enter DECnet-plus domain name []:
- * Enter the location of the node [COMP_ROOM_1]:
- * Is this node eligible to be a database server [YES]:
- * Enter group names for the node []:
- *** Proceed (YES, NO/REENTER, QUIT) [YES]:
- * Enter a node to be configured [NONE]:

Configuring jukeboxes...

- * Enter a jukebox to be configured [NONE]: TLZ06J
- * Enter jukebox control type (MRD or DCSC) [MRD]:
- * Enter robot name controlling jukebox: GKB601:
- * Enter nodes that directly access jukebox: MOE
- * Enter location of jukebox [COMP_ROOM_1]:
- *** Proceed (YES, NO/REENTER, QUIT) [YES]:
- * Enter media types for jukebox drives [TLZ06]:
- * Enter jukebox drive 0 to be configured [NONE]: TLZ06D
- * Enter OpenVMS device name of drive [TLZ06D]: MOE\$MKB600:
- *** Proceed (YES, NO/REENTER, QUIT) [YES]:
- * Enter jukebox drive 1 to be configured [NONE]:
- * Do you want to perform an inventory of the jukebox [NO]: NO
- * Enter a jukebox to be configured [NONE]:

Configuring standalone drives and stackers...

* Enter a drive to be configured [NONE]:

Configuring volumes...

* Enter volume range [NONE]: TLZ000-TLZ012

* Enter media type for volumes [TLZ06]:

* Enter pool for volumes:

* Enter placement (Onsite, Offsite or Jukebox) [ONSITE]:

* Enter the onsite location of volumes [COMP_ROOM_1]:

* Enter the offsite location of volumes [IRON_MOUNTAIN]:

* Are volumes initialized [NO]:

* Do you want to initialize the volumes [NO]: NO

*** Proceed (YES, NO/REENTER, QUIT) [YES]:

MDMS configuration is complete.

The following objects now exist in the database:

Domain definition...

```
      Description: Default MDMS Domain
      Access Control: NONE
      Last Updated By: MOE::SMITH
      Mail: SYSTEM
      Offsite Location: IRON_MOUNTAIN
      Onsite Location: COMP_ROOM_1
      Check Access: NO
      Deallocate State: TRANSITION
      Default Access: YES
      Default Media Type: TLZ06
      Opcom Class: TAPES
      Request ID: 35
      Protection: S:RW,O:RW,G:R,W
      DB Server Node: MOE
      DB Server Date: 20-DEC-2001 14:17:00
      Scheduler Type: INTERNAL
      Max Scratch Time: NONE
      Scratch Time: 0365 00:00:00
      Transition Time: 0014 00:00:00
```

Locations...

Location Name	In Location
---------------	-------------

Configuration Example

COMP_ROOM_1
IRON_MOUNTAIN

Groups...

%MDMS-E-NOOBJECTS, no such objects currently exist

Nodes...

Node Name	Database	Transports
MOE	YES	TCPIP,DECNET

Drives...

Drive Name	Allocated	State	Number	Jukebox
TLZ06D	NO	EMPTY	0	TLZ06J

Jukeboxes...

Jukebox Name	State
TLZ06J	AVAILABLE

Media types...

Media Type
TLZ06

Pools...

MDMS-E-NOOBJECTS, no such objects currently exist

Volumes...

Volume ID	State	Scratch Date	Placement
TLZ000	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ001	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ002	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ003	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ004	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ005	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ006	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ007	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ008	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ009	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ010	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ011	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1
TLZ012	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1

Configuration Example

If you completed the procedure successfully and completely, your system should now be ready for most operations using ABS and/or HSM. If you require further custom configuration, refer to the Guide to Operations.

Now that you have configured MDMS, you need to move the volumes into the jukebox. In this example, the volumes were already in the jukebox. I had to move them into the jukebox in the database. This is why I used /NOASSIST/NOPHYSICAL. The following command moved the volumes into the jukebox in the database. If you have a vision jukebox the volumes will have been configured in the jukebox in the MDMS\$CONFIGURE.COM procedure.

```
$ MDMS MOVE VOL TLZ000-TLZ011 TLZ06J/SLOT=0-11/NOASSIST/NOPHYSICAL
```

```
$ MDMS SHOW VOL
```

Volume ID	State	Scratch Date	Placement	
TLZ000	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 0
TLZ001	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 1
TLZ002	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 2
TLZ003	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 3
TLZ004	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 4
TLZ005	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 5
TLZ006	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 6
TLZ007	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 7
TLZ008	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 8
TLZ009	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 9
TLZ010	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 10
TLZ011	UNINITIALIZED	NONE	JUKEBOX	TLZ06J, SLOT 11
TLZ012	UNINITIALIZED	NONE	ONSITE	COMP_ROOM_1

You can show the contents of your jukebox using the following command:

```
$ MDMS SHOW JUKE TLZ06J/CONTENTS
```

```
Jukebox: TLZ06J
Description:
Access Control: NONE
  Owner: MOE::SMITH
  Nodes: MOE
Groups:
Location: COMP_ROOM_1
Disabled: NO
Auto Reply: YES
  Access: ALL
  State: AVAILABLE
Control: MRD
Threshold: 0
Free Volumes: 0
  Robot: GKB601
Slot Count: 12
Usage: NOMAGAZINE
```

Configuration Example

Jukebox TLZ06J contents:

Number	Drive Name	Allocated	State	Volume
0	TLZ06D	NO	EMPTY	

Slot	Volume ID	State	Scratch date	Magazine Slot	
0	TLZ000	UNINITIALIZED	NONE	---	-
1	TLZ001	UNINITIALIZED	NONE	---	-
2	TLZ002	UNINITIALIZED	NONE	---	-
3	TLZ003	UNINITIALIZED	NONE	---	-
4	TLZ004	UNINITIALIZED	NONE	---	-
5	TLZ005	UNINITIALIZED	NONE	---	-
6	TLZ006	UNINITIALIZED	NONE	---	-
7	TLZ007	UNINITIALIZED	NONE	---	-
8	TLZ008	UNINITIALIZED	NONE	---	-
9	TLZ009	UNINITIALIZED	NONE	---	-
10	TLZ010	UNINITIALIZED	NONE	---	-
11	TLZ011	UNINITIALIZED	NONE	---	-

Before you can use the volumes, you have to initialize the volumes. If this jukebox would have been a vision jukebox, you could initialize them in the MDMS\$CONFIGURE.COM procedure.

```
$ MDMS INIT VOL TLZ000-TLZ0011/OVER
```

```
$ MDMS SHOW VOL
```

Volume ID	State	Scratch Date	Placement
TLZ000	FREE	NONE	JUKEBOX TLZ06J, SLOT 0
TLZ001	FREE	NONE	JUKEBOX TLZ06J, SLOT 1
TLZ002	FREE	NONE	DRIVE TLZ06D
TLZ003	FREE	NONE	JUKEBOX TLZ06J, SLOT 3
TLZ004	FREE	NONE	JUKEBOX TLZ06J, SLOT 4
TLZ005	FREE	NONE	JUKEBOX TLZ06J, SLOT 5
TLZ006	FREE	NONE	JUKEBOX TLZ06J, SLOT 6
TLZ007	FREE	NONE	JUKEBOX TLZ06J, SLOT 7
TLZ008	FREE	NONE	JUKEBOX TLZ06J, SLOT 8
TLZ009	FREE	NONE	JUKEBOX TLZ06J, SLOT 9
TLZ010	FREE	NONE	JUKEBOX TLZ06J, SLOT 10
TLZ011	FREE	NONE	JUKEBOX TLZ06J, SLOT 11
TLZ012	UNINITIALIZED	NONE	ONSITE COMP_ROOM_1

Check the SYSTEM_BACKUPS_ENV. This environment was created when you installed ABS.

```
$ MDMS SHOW ENV SYSTEM_BACKUPS_ENV
```

```
Environment: SYSTEM_BACKUPS_ENV
Description:
Access Control: MOE::ABS (READ, WRITE, EXECUTE, DELETE, SET,
SHOW,
CONTROL)
Owner: MOE::ABS
Action: RECORD_DATE
```


Configuration Example

```
        Assist: YES
Compression: NONE
Data Safety: CRC,FULL_VERIFY,XOR
Drive Count: 1
    Epilogue:
        Interval: NONE
        Links Only: YES
Listing Option: NONE
        Lock: YES
Notification -
    - Opcom: TAPES
    - Type: BRIEF
    - When: FATAL
Notification -
    - Mail: <REQUESTER>
    - Type: BRIEF
    - When: FATAL
Profile -
    - Cluster: *
      - Node: *
- Privileges:
    - Rights:
      - User: ABS
Prologue:
Retry Limit: 0
Span Filesystems: YES
```

Check the SYSTEM_BACKUPS archive. This archive is created when you installed ABS. Make sure that it has the media type of your volumes.

\$ MDMS SHOW ARCHIVE SYSTEM_BACKUPS

```
        Archive: SYSTEM_BACKUPS
Description:
Access Control: MOE::ABS (READ, WRITE, EXECUTE, DELETE, SET,
SHOW,
                CONTROL)
        Owner: MOE::ABS
Archive Type: TAPE
Catalog -
    - Name: ABS_CATALOG
    - Nodes:
Consolidation -
    - Interval: 0007 00:00:00
    - Savesets: 0
    - Volumes: 0
Destination:
    Drives:
Expiration Date: NONE
    Location:
Maximum Saves: 1
    Media Type: TLZ06
    Pool:
```

Configuration Example

Retention Days: 365
Volume Sets:

Now create a save with the following attributes:

Name - SYSTEM_WFD_SR
Frequency - DAILY_FULL_WEEKLY
Include - \$1\$DKA0:
Environment - system_backups_env
Archive - system_backups
Start = 21:00

```
$ MDMS CREATE SAVE SYSTEM_WFD_SR -
_$ /FREQUENCY=DAILY_FULL_WEEKLY -
_$ /INCLUDE=$1$DKA0: -
_$ /ENVIRONMENT=system_backups_env
_$ /ARCHIVE=SYSTEM_BACKUPS -
_$ /START=21:00
$ MDMS SHOW SAVE SYSTEM_WFD_SR

          Save: SYSTEM_WFD_SR
Description:
Access Control: NONE
          Owner: MOE::SMITH
          Archive: SYSTEM_BACKUPS
          Base Date: 20-DEC-2001 21:00:00
Delete Interval: NONE
          Environment: SYSTEM_BACKUPS_ENV
          Epilogue:
Execution Nodes: MOE
Explicit Interval:
          Frequency: DAILY_FULL_WEEKLY
          Groups:
Incremental: NO
          Job Number: 0
          Prologue:
          Schedule: SYSTEM_WFD_SR_SAVE_SCHED
Sequence Option: SEQUENTIAL
          Skip Time: NONE
          Start Date: 20-DEC-2001 21:00:00
Transaction Status:
          Selections: SYSTEM_WFD_SR_SAVE_SEL_DEF

Default Selection -
- Data Select Type: VMS_FILES
  - Include: $1$DKA0:
  - Exclude:
- Source Node:
```

All done! You can check the results of the daily backups in
ABS\$LOG:SYSTEM_WFD_SR.LOG.

Index

Numerics

100 times per day 3-33

A

ABS 1-1
 catalogs 2-1
 policy objects 2-1
ABS Operational Environment 2-1
ABS_CATALOG 3-2
ABS_EXECUTION_ENVIRONMENT 3-13
ABS_NODE_NAME 3-13
ABS_OS_DMT_n 3-25
ABS_OS_INCREMENTAL_LEVEL_n 3-25
ABS_OS_LAST_RVN_n 3-25
ABS_OS_OBJECT_SET_n 3-25
ABS_OS_OBJECT_TYPE_n 3-25
ABS_OS_SAVESET_FORMAT_n 3-26
ABS_OS_SAVESET_NAME_n 3-26
ABS_OS_START_FILE_POSITION_n 3-26
ABS_OS_START_RVN_n 3-25
ABS_OS_STATUS_n 3-26
ABS_OS_VOLUME_SET_n 3-25
ABS_OUTPUT_DEVICE 3-13
ABS_RESTORE_REQUEST_NAME 3-13
ABS_SAVE_REQUEST_NAME 3-13
ABS_STORAGE_CLASS 3-13
Action 3-11
After Schedule 3-30
after schedule when 3-30
Agent Qualifiers 3-27
ALL 3-30
ANNUALLY 3-22
ANSI-imposed maximum 3-2
Archive 3-16
archive 3-15
Archive Backup System 1-1
Archive Type 3-2
archive type 3-3, 3-4, 3-26
Archives 2-3
authorized users 3-4

B

Backup agent 2-1, 2-6
backup agent 3-27
Base Date 3-16

Before Date 3-17, 3-27
BIWEEKLY 3-22
BRIEF 3-14

C

Catalog 2-1, 2-5, 3-2, 3-4, 3-17
 creating
 SLS type 3-5
 staging type 3-5
 improving performance 3-8
Catalog Commands 9-13
 for Oracle database lookup 9-14
 for System Backup to Tape 9-6
Catalogs 3-4
Combining Dates, Days and Months 3-32
Command 3-30
COMPLETE 3-14
Complex Backup Schedules 3-23
Compression 3-11
Conflict Options 3-28
Consolidation 3-2
consolidation criteria 3-2, 3-3
consolidation interval 3-2
CRC 3-12
CUSTOM 3-23, 3-29
Custom 3-20, 3-29

D

DAILY 3-20
Data Safety 3-12
DATA TYPE 3-18, 3-28
Data Type 3-18, 3-28
Database
 catalog 2-5
Database Management Services 1-1
database names 3-18
Date Archived 3-17
Date Specifications 3-31
Date Type 3-27
Dates 3-31
Day Specifications 3-31
Days 3-31
DCL 1-2
default archives 3-1
default consolidation criteria 3-3
default domain media type. 3-4
default environments 3-11
default onsite location 3-4
default retry interval 3-13
DEFAULT_ENV 3-11

- Defaults 9-16
 - Archive Name 9-16
 - Catalog Name 9-16
 - I/O Block Size 9-16
- Delete Interval 3-19
- DELETE_FILE 3-11
- Destination 3-3, 3-20
- destination 3-2
- DISASTER_RECOVERY_ENV 3-11
- DISK 3-2
- disk names 3-18
- Disk, File, Path and Database Specification Formats 3-19, 3-29
- domain 3-4
- drive 3-4
- Drive Count 3-12
- Drives 3-3
- drives 3-4

E

- Environment 3-20
- environment 3-11, 3-15
- Environment Name 3-11
- Environment policy 2-4, 2-5
- Environments 2-4, 3-11
- Epilogue 3-12, 3-25
- ERROR 3-14
- EXCLUDE 3-18, 3-28, 3-32
- Exclude 3-18, 3-28, 3-32
- execution node 3-24
- Expiration Date 3-3
- EXPLICIT 3-23
- Explicit 3-20, 3-29
- EXPLICIT INTERVAL 3-20, 3-29
- Explicit Interval 3-20, 3-23
- EXTERNAL 2-7

F

- FATAL 3-14, 3-30
- file names 3-18
- Frequency 3-20
- frequency. 3-16
- FULL 3-14
- full backups 3-3
- FULL_BACKUP 3-12

G

- get 2-1
- Groups 3-24

- GZIP Compression 3-11

I

- INCLUDE 3-18, 3-26, 3-27, 3-28, 3-32
- Include 3-18, 3-28, 3-32
- Incremental 3-24
- incremental backups 3-3
- inherit 3-1
- inherit restore requests 3-1
- inheriting attributes 3-1
- Interfaces
 - CLI 2-7
 - GUI 2-7
- INTERNAL 2-7
- INTERVAL 3-2
- Interval 3-13

J

- jukeboxes 3-4

K

- Keep 3-19

L

- Links Only 3-13
- Listing Option 3-14
- Location 3-4
- location 3-2, 3-3, 3-4
- Lock 3-14
- LOG-2 3-22
- LOG-3 3-22
- Logical 9-16
 - MDMS\$SBT_ARCHIVE 9-16
 - MDMS\$SBT_ARCHIVE_n 9-16
 - MDMS\$SBT_CATALOG 9-17
 - MDMS\$SBT_IO_BLOCK_SIZE 9-16
 - MDMS\$SBT_TRACE_LEVEL 9-6, 9-17, 9-18
- Logical Names 3-13
- Logical Names in Save/Restore Prologues and Epilogues 3-25

M

- MAIL 3-14
- Maximum Saves 3-4
- MDMS 1-1, 2-7
- MDMSview 1-2
- Media Management Services 1-1

- Media Type 3-4
- media type 3-2, 3-3, 3-4
- Media, Device and Management Services 1-1, 2-7
- Month Specifications 3-31
- MONTHLY 3-22
- Months 3-31

N

- NEVER 3-23
- NEW VERSION 3-28
- next start date 3-16
- No compression 3-11
- NO_CHANGE 3-11
- Nodes 3-24
- nolock 3-14
- NONE 3-30
- NORMAL 3-14
- Notification 3-14
- NTclient
 - large disk considerations 11-7

O

- ON DEMAND 3-20
- ONE TIME ONLY 3-19, 3-20
- OPCOM 3-14
- Oracle Databases 9-1
 - Backing up to Tape 9-2
 - See System Backup to Tape 9-1
- Oracle Rdb Database 3-28
- Oracle Rdb Database Options 3-18
- Oracle Rdb Databases 3-19, 3-29
- Oracle Rdb Storage Area 3-18, 3-28
- Oracle Rdb Storage Areas 3-19, 3-29
- Oracle's Recovery Manager 9-10
- Oracle's Recovery Manager
 - Specifying a Catalog 9-11
 - Specifying an I/O Block Size 9-12
 - Specifying Archives for Duplex Backups 9-12
 - Using with System Backup to Tape 9-10
- OVERLAPPED 3-26
- OVERLAY VERSION 3-28

P

- path names 3-18
- Policy
 - objects 2-1
- Policy objects 2-1
 - environment policy 2-4, 2-5
 - restore request 2-2

- save request 2-2
- storage policy 2-3
- Pool 3-4
- pool 3-2, 3-4
- Privileges 9-3
 - authorizing for SBT 9-3
- Profile 3-15
- Prologue 3-12, 3-25

Q

- QUARTERLY 3-22

R

- Rdb Databases 3-27
- Rdb Storage Areas 3-27
- RECORD_DATE 3-11
- Relationships Between ABS Objects 3-16
- REPLACE VERSION 3-28
- Request
 - restore 2-2
 - save 2-2
- Reschedule 3-26
- Restore Name 3-16
- Restore request 2-2
- restore request 3-15
- Restores 2-2, 3-15
- Restrictions 9-17
 - Doing Parallel Backups 9-17
 - granting for SBT 9-3
 - Piece Name Length Greater than 254 Characters 9-18
- RETAIN VERSION 3-28
- Retention Days 3-3
- retry interval 3-13
- Retry Limit 3-13
- Rights 9-3
 - Granting for SBT 9-3

S

- Save Name 3-16
- Save request 2-2
- save request 3-15
- Saves 2-2, 3-15
- SAVESETS 3-2
- schedule 3-15
- SCHEDULER 2-7
- Schedules 2-5, 3-29
- Scheduling 9-15
 - System Backup to Tape Operations 9-13

- Scheduling Services 1-1
- scratch pool 3-4
- Security Services 1-1
- SELECTIONS 3-18, 3-26, 3-27
- Selections 2-4, 3-26, 3-27
- selections 3-15
- SEMI_ANNUALLY 3-22
- Sequence Option 3-26
- SEQUENTIAL 3-26
- Since Date 3-17, 3-27
- Skip Time 3-16
- skip time 3-17
- source location 3-20
- SOURCE NODE 3-19, 3-24, 3-28
- Source Node 3-18, 3-28
- Span Filesystems 3-13
- START 3-14
- Start Date 3-16
- start date 3-16, 3-17
- Storage policy 2-3
- SUCCESS 3-30
- System 9-1
- System Backup to Tape 9-1
 - Archive Name 9-16
 - Catalog Name 9-16
 - Configuring 9-6
 - Creating an Archive 9-7
 - Creating an ORACLE_DB Catalog 9-7
 - Defaults 9-16
 - Defining the logical
 - MDMS\$SBT_TRACE_LEVEL 9-6
 - Editing Oracle's Command Procedures 9-3
 - Editing Oracle's Link Option File 9-3
 - I/O Block Size 9-16
 - Introduction 9-1
 - Linking with the Oracle Server 9-2
 - Logicals 9-16
 - Logicals Names 9-16
 - Privileges 9-3
 - Relinking the ORA_RDBMS
 - executables 9-5
 - Restrictions 9-17
 - Rights 9-3
 - Specifying a Catalog 9-11
 - Specifying an I/O Block Size 9-12
 - Specifying Archives for Duplex Backups 9-12
 - Testing the Configuration of 9-9
 - Testing with Oracle's Recovery Manager 9-2
 - Troubleshooting 9-18
 - Using the logical MDMS\$SBT_TRACE_LEVEL
 - 9-18
 - Using the MDMS Scheduler with 9-15

- Using the Show Catalog Command 9-13
- Using with Oracle's Recovery Manager 9-10
- What versions of Oracle supported 9-1
- System backups
 - NTclient 11-7
 - UNIX client 11-7
- SYSTEM_BACKUPS_ENV 3-11

T

- TAPE 3-2
- Times 3-33
- two archives 3-3, 3-15, 3-16
- TYPE 3-14

U

- UNIX client
 - large disk considerations 11-7
- UNIX Compression 3-11
- UNIX Files 3-18, 3-27, 3-28
- UNIX files 3-19, 3-24, 3-29
- UNIX_BACKUPS_ENV 3-11, 3-12
- Use of Base Date, Start Date and Skip Time 3-17
- USER_BACKUPS_ENV 3-11

V

- VMS Files 3-18, 3-19, 3-27, 3-28, 3-29
- volume set 3-2, 3-4
- Volume Sets 3-4
- VOLUMES 3-2
- volumes 3-4

W

- WARNING 3-14, 3-30
- WEEKLY 3-22
- WEEKLY FULL, DAILY INCREMENTAL 3-21
- WHEN 3-14
- Windows Files 3-18, 3-27, 3-28
- Windows files 3-19, 3-24, 3-29

X

- XOR 3-12