

Dell EMC PowerMax Family Product Guide

PowerMaxOS

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Contents

| | |
|--|-----------|
| Figures..... | 8 |
| Tables..... | 9 |
| Preface..... | 10 |
| Chapter 1: PowerMax with PowerMaxOS..... | 18 |
| Introduction to PowerMax with PowerMaxOS..... | 18 |
| PowerMax arrays..... | 18 |
| PowerMaxOS operating environment..... | 19 |
| Software packages..... | 19 |
| The Essentials software package..... | 19 |
| The Pro software package..... | 20 |
| The zEssentials software package | 21 |
| The zPro software package..... | 21 |
| Package availability..... | 22 |
| PowerMaxOS..... | 22 |
| PowerMaxOS emulations..... | 22 |
| Container applications | 24 |
| Data protection and integrity..... | 27 |
| Data efficiency..... | 32 |
| Chapter 2: Management Interfaces..... | 35 |
| Management interface versions..... | 35 |
| Unisphere for PowerMax..... | 35 |
| Workload Planner..... | 36 |
| FAST Array Advisor..... | 36 |
| Unisphere 360..... | 36 |
| CloudIQ..... | 36 |
| CyberSecIQ..... | 38 |
| Solutions Enabler..... | 39 |
| Mainframe Enablers..... | 39 |
| Geographically Dispersed Disaster Restart (GDDR)..... | 40 |
| SMI-S Provider..... | 40 |
| VASA Provider..... | 40 |
| eNAS management interface | 40 |
| Storage Resource Management (SRM)..... | 40 |
| vStorage APIs for Array Integration..... | 41 |
| SRDF Adapter for VMware vCenter Site Recovery Manager..... | 41 |
| SRDF/Cluster Enabler | 41 |
| Product Suite for z/TPF..... | 42 |
| SRDF/TimeFinder Manager for IBM i..... | 42 |
| AppSync..... | 43 |
| EMC Storage Analytics (ESA)..... | 43 |
| Chapter 3: Open Systems Features..... | 44 |

| | |
|---|-----------|
| PowerMaxOS support for open systems..... | 44 |
| PowerPath..... | 44 |
| Operational overview..... | 44 |
| Host registration..... | 45 |
| Device status..... | 45 |
| Automatic creation of Initiator Groups..... | 45 |
| Management..... | 45 |
| More information..... | 46 |
| Backup and restore using PowerProtect Storage Direct and Data Domain..... | 46 |
| Backup..... | 46 |
| Restore..... | 47 |
| Storage Direct agents..... | 47 |
| Features used for Storage Direct backup and restore..... | 48 |
| Storage Direct and traditional backup..... | 48 |
| More information..... | 48 |
| VMware Virtual Volumes..... | 48 |
| vVol components..... | 49 |
| vVol scalability..... | 49 |
| vVol workflow..... | 49 |
| Chapter 4: Mainframe Features..... | 51 |
| PowerMaxOS support for mainframe..... | 51 |
| IBM Z Systems functionality support..... | 51 |
| Global Mirror support..... | 52 |
| Transparent Cloud Tiering support..... | 53 |
| IBM 2107 support..... | 54 |
| Logical control unit capabilities..... | 54 |
| Disk drive emulations..... | 55 |
| Cascading configurations..... | 55 |
| Chapter 5: Provisioning..... | 56 |
| Thin provisioning..... | 56 |
| Pre-configuration for thin provisioning..... | 57 |
| Thin devices (TDEVs)..... | 57 |
| Thin device oversubscription..... | 58 |
| Internal memory usage..... | 58 |
| Open Systems-specific provisioning..... | 58 |
| Multi-array provisioning..... | 60 |
| Chapter 6: Service levels..... | 61 |
| Definition of service levels..... | 61 |
| Defined service levels..... | 61 |
| Service level priorities..... | 63 |
| Default service levels..... | 63 |
| Availability of service levels..... | 63 |
| Use of service levels to maintain system performance..... | 64 |
| Usage examples..... | 65 |
| Manage service levels..... | 65 |

| | |
|---|-----------|
| Chapter 7: Automated data placement..... | 66 |
| Environment..... | 66 |
| Operation..... | 66 |
| Service level biasing..... | 66 |
| Compression and deduplication..... | 66 |
| Availability..... | 66 |
| Chapter 8: Native local replication with TimeFinder..... | 67 |
| About TimeFinder..... | 67 |
| Interoperability with legacy TimeFinder products..... | 68 |
| Targetless snapshots..... | 68 |
| Secure snaps..... | 68 |
| Provision multiple environments from a linked target..... | 69 |
| Cascading snapshots..... | 69 |
| Accessing point-in-time copies..... | 70 |
| Mainframe SnapVX and zDP..... | 70 |
| Snapshot policy..... | 71 |
| Chapter 9: Remote replication..... | 73 |
| Native remote replication with SRDF..... | 73 |
| SRDF 2-site solutions..... | 74 |
| SRDF multi-site solutions..... | 76 |
| Interfamily compatibility..... | 77 |
| SRDF device pairs..... | 77 |
| Dynamic device personalities..... | 81 |
| SRDF modes of operation..... | 81 |
| SRDF groups..... | 82 |
| Director boards, links, and ports..... | 82 |
| SRDF consistency..... | 83 |
| Data migration..... | 83 |
| More information..... | 84 |
| SRDF/Metro..... | 85 |
| Deployment options..... | 85 |
| SRDF/Metro Resilience..... | 85 |
| Disaster recovery facilities..... | 87 |
| Mobility ID with ALUA..... | 89 |
| More information..... | 89 |
| RecoverPoint..... | 89 |
| Remote replication using eNAS..... | 90 |
| Chapter 10: Cloud Mobility..... | 91 |
| Cloud Mobility for Dell EMC PowerMax..... | 91 |
| Chapter 11: Blended local and remote replication..... | 93 |
| Integration of SRDF and TimeFinder..... | 93 |
| R1 and R2 devices in TimeFinder operations..... | 93 |
| SRDF/AR..... | 93 |
| SRDF/AR 2-site configurations..... | 94 |

| | |
|---|------------|
| SRDF/AR 3-site configurations..... | 95 |
| TimeFinder and SRDF/A..... | 95 |
| TimeFinder and SRDF/S..... | 96 |
| Chapter 12: Data migration..... | 97 |
| Overview..... | 97 |
| Data migration for open systems..... | 98 |
| Non-Disruptive Migration..... | 98 |
| Open Replicator..... | 105 |
| PowerPath Migration Enabler..... | 106 |
| Data migration using SRDF/Data Mobility..... | 107 |
| Space and zero-space reclamation..... | 107 |
| Data migration for IBM System i..... | 107 |
| Data migration for mainframe..... | 107 |
| Volume migration using z/OS Migrator..... | 108 |
| Dataset migration using z/OS Migrator..... | 108 |
| Chapter 13: Online Device Expansion..... | 110 |
| Introduction..... | 110 |
| General features..... | 110 |
| Standalone devices..... | 111 |
| SRDF devices..... | 111 |
| LREP devices..... | 112 |
| Management facilities..... | 113 |
| Solutions Enabler..... | 113 |
| Unisphere..... | 113 |
| Mainframe Enablers..... | 114 |
| Chapter 14: System security..... | 115 |
| User authentication and authorization..... | 115 |
| Roles and permissions..... | 115 |
| Roles and their hierarchy..... | 115 |
| Permissions for roles..... | 116 |
| Secure Reads policy..... | 119 |
| View permissions required for an operation..... | 119 |
| Lockbox..... | 119 |
| Stable System Values (SSVs)..... | 119 |
| Lockbox passwords..... | 120 |
| Default Lockbox password..... | 120 |
| Client/server communications..... | 120 |
| Secure device deletion..... | 120 |
| Appendix A: Mainframe Error Reporting..... | 121 |
| Error reporting to the mainframe host..... | 121 |
| SIM severity reporting..... | 121 |
| Environmental errors..... | 122 |
| Operator messages..... | 124 |
| Appendix B: Licensing..... | 126 |

| | |
|--|------------|
| eLicensing..... | 126 |
| Capacity measurements..... | 127 |
| Open systems licenses..... | 128 |
| License packages..... | 128 |
| Individual licenses..... | 130 |
| Ecosystem licenses..... | 130 |
| PowerMax Mainframe software packaging options..... | 130 |
| Index..... | 132 |


| | | |
|----|---|-----|
| 1 | D@RE architecture..... | 28 |
| 2 | EEEE components..... | 30 |
| 3 | Inline compression and over-subscription..... | 33 |
| 4 | Data flow during a backup operation to Data Domain..... | 47 |
| 5 | Two-site Global Mirror..... | 52 |
| 6 | TCT environment with PowerMax and DLm..... | 54 |
| 7 | Auto-provisioning groups..... | 59 |
| 8 | SnapVX targetless snapshots..... | 69 |
| 9 | SnapVX cascaded snapshots..... | 70 |
| 10 | zDP operation..... | 71 |
| 11 | R1 and R2 devices..... | 78 |
| 12 | R11 device in concurrent SRDF..... | 79 |
| 13 | R21 device in cascaded SRDF..... | 80 |
| 14 | R22 devices in cascaded and concurrent SRDF/Star..... | 80 |
| 15 | Migrating data and removing a secondary (R2) array..... | 84 |
| 16 | SRDF/Metro..... | 85 |
| 17 | SRDF/Metro Smart DR..... | 87 |
| 18 | Disaster recovery for SRDF/Metro..... | 88 |
| 19 | Cloud Mobility for PowerMax - high level architecture..... | 91 |
| 20 | SRDF/AR 2-site solution..... | 94 |
| 21 | SRDF/AR 3-site solution..... | 95 |
| 22 | Configuration of a VMAX3, VMAX All Flash or PowerMax migration..... | 99 |
| 23 | Configuration of a VMAX migration..... | 101 |
| 24 | Open Replicator hot (or live) pull..... | 106 |
| 25 | Open Replicator cold (or point-in-time) pull..... | 106 |
| 26 | z/OS volume migration..... | 108 |
| 27 | z/OS Migrator dataset migration..... | 108 |
| 28 | Expand Volume dialog in Unisphere..... | 113 |
| 29 | z/OS IEA480E acute alert error message format (call home failure)..... | 124 |
| 30 | z/OS IEA480E service alert error message format (Disk Adapter failure)..... | 125 |
| 31 | z/OS IEA480E service alert error message format (SRDF Group lost/SIM presented against unrelated resource)..... | 125 |
| 32 | z/OS IEA480E service alert error message format (mirror-2 resynchronization)..... | 125 |
| 33 | z/OS IEA480E service alert error message format (mirror-1 resynchronization)..... | 125 |
| 34 | eLicensing process..... | 126 |

| | | |
|----|---|-----|
| 1 | Typographical conventions used in this content..... | 16 |
| 2 | PowerMaxOS emulations..... | 22 |
| 3 | eManagement resource requirements..... | 24 |
| 4 | eNAS configurations by array..... | 26 |
| 5 | Unisphere tasks..... | 35 |
| 6 | Health score categories..... | 38 |
| 7 | vVol architecture component management capability..... | 49 |
| 8 | vVol-specific scalability..... | 49 |
| 9 | Logical control unit maximum values..... | 54 |
| 10 | Maximum LPARs per port..... | 55 |
| 11 | RAID options..... | 57 |
| 12 | SRDF 2-site solutions..... | 74 |
| 13 | SRDF multi-site solutions..... | 76 |
| 14 | SIM severity alerts..... | 121 |
| 15 | Environmental errors reported as SIM messages..... | 122 |
| 16 | PowerMax product title capacity types..... | 127 |
| 17 | PowerMax license packages..... | 128 |
| 18 | Individual licenses for open systems environment..... | 130 |
| 19 | Individual licenses for open systems environment..... | 130 |
| 20 | PowerMax Mainframe software packaging options (PowerMax 8000 only)..... | 131 |

Preface

As part of an effort to improve its product lines, Dell EMC periodically releases revisions of its software and hardware. Functions that are described in this document may not be supported by all versions of the software or hardware. The product release notes provide the most up-to-date information about product features.

Contact your Dell EMC representative if a product does not function properly or does not function as described in this document.

 **NOTE:** This document was accurate at publication time. New versions of this document might be released on Dell EMC Online Support (<https://www.dell.com/support/home>). Check to ensure that you are using the latest version of this document.

Purpose

This document introduces the features of the Dell EMC PowerMax arrays running PowerMaxOS 5978. The descriptions of the software capabilities also apply to VMAX All Flash arrays running PowerMaxOS 5978, except where noted.

Audience

This document is intended for use by customers and Dell EMC representatives.

Related documentation

The following documentation portfolios contain documents related to the hardware platform and manuals needed to manage your software and storage system configuration. Also listed are documents for external components that interact with the PowerMax array.

Hardware platform documents:

Dell EMC PowerMax Family Site Planning Guide Provides planning information regarding the purchase and installation of a PowerMax 2000, 8000 with PowerMaxOS.

Dell EMC Best Practices Guide for AC Power Connections for PowerMax 2000, 8000 with PowerMaxOS Describes the best practices to assure fault-tolerant power to a PowerMax 2000 or PowerMax 8000 array.

Dell EMC PowerMax Family Security Configuration Guide Shows how to securely deploy PowerMax arrays running PowerMaxOS.

Unisphere documents:

Dell EMC Unisphere for PowerMax Release Notes Describes new features and any known limitations for Unisphere for PowerMax.

Dell EMC Unisphere for Provides installation instructions for Unisphere for PowerMax.

**PowerMax
Installation Guide**

**Dell EMC
Unisphere for
PowerMax Online
Help** Describes the Unisphere for PowerMax concepts and functions.

**Dell EMC
Unisphere for
PowerMax REST
API Concepts and
Programmer's
Guide** Describes the Unisphere for PowerMax REST API concepts and functions.

**Dell EMC
Unisphere 360
Release Notes** Describes new features and any known limitations for Unisphere 360.

**Dell EMC
Unisphere 360
Installation Guide** Provides installation instructions for Unisphere 360.

**Dell EMC
Unisphere 360
Online Help** Describes the Unisphere 360 concepts and functions.

Solutions Enabler documents:

**Dell EMC
Solutions
Enabler, VSS
Provider, and
SMI-S Provider
Release Notes** Describes new features and any known limitations.

**Dell EMC
Solutions Enabler
and VSS Provider
Release Notes** Describes new features and any known limitations.

**Dell EMC
Solutions Enabler
Installation and
Configuration
Guide** Provides host-specific installation instructions.

**Dell EMC
Solutions Enabler
CLI Reference
Guide** Documents the SYMCLI commands, daemons, error codes and option file parameters provided with the Solutions Enabler man pages.

**Dell EMC
Solutions Enabler
Array Controls
and Management
CLI User Guide** Describes how to configure array control, management, and migration operations using SYMCLI commands for arrays running HYPERMAX OS and PowerMaxOS.

**Dell EMC
Solutions Enabler
Array Controls
and Management
CLI User Guide** Describes how to configure array control, management, and migration operations using SYMCLI commands for arrays running Enginuity.

**Dell EMC
Solutions Enabler
SRDF Family CLI
User Guide** Describes how to configure and manage SRDF environments using SYMCLI commands.

| | |
|--|---|
| Dell EMC Solutions Enabler SRDF Family State Tables Guide | Describes the applicable pair states for various SRDF operations. |
| SRDF Interfamily Connectivity Information | Defines the versions of PowerMaxOS, HYPERMAX OS and Enginuity that can make up valid SRDF replication and SRDF/Metro configurations, and can participate in Non-Disruptive Migration (NDM). |
| Dell EMC SRDF Introduction | Provides an overview of SRDF, its uses, configurations, and terminology. |
| Dell EMC Solutions Enabler TimeFinder SnapVX CLI User Guide | Describes how to configure and manage TimeFinder SnapVX environments using SYMCLI commands. |
| Dell EMC Solutions Enabler TimeFinder Family (Mirror, Clone, Snap, VP Snap) Version 8.2 and higher CLI User Guide | Describes how to configure and manage TimeFinder Mirror, Clone, Snap, VP Snap environments for Enginuity and HYPERMAX OS using SYMCLI commands. |
| Dell EMC Solutions Enabler SRM CLI User Guide | Provides Storage Resource Management (SRM) information that is related to various data objects and data handling facilities. |
| Dell EMC SRDF/Metro vWitness Configuration Guide | Describes how to install, configure, and manage SRDF/Metro using vWitness. |
| Dell EMC Events and Alerts for PowerMax and VMAX User Guide | Documents the SYMAPI daemon messages, asynchronous errors and message events, SYMCLI return codes, and how to configure event logging. |

PowerPath documents:

| | |
|--|--|
| PowerPath/VE for VMware vSphere Release Notes | Describes any new or modified features and any known limitations. |
| PowerPath/VE for VMware vSphere Installation and Administration Guide | Shows how to install, configure, and manage PowerPath/VE. |
| PowerPath Family CLI and System Messages Reference | Documents the PowerPath CLI commands and system messages. |
| PowerPath Family Product Guide | Provides a description of the products in the PowerPath family. |
| PowerPath Management | Shows how to install and configure the PowerPath Management Appliance. |

***Appliance
Installation and
Configuration
Guide***

***PowerPath
Management
Appliance
Release Notes*** Describes new features and any known limitations.

***PowerPath
Migration
Enabler User
Guide*** Shows how to carry out data migration using the PowerPath Migration Enabler.

Embedded NAS (eNAS) documents:

***Dell EMC
PowerMax eNAS
Release Notes*** Describes the new features and identify any known functionality restrictions and performance issues that may exist in the current version.

***Dell EMC
PowerMax eNAS
Quick Start
Guide*** Describes how to configure eNAS on a PowerMax storage system.

***Dell EMC
PowerMax eNAS
File Auto
Recovery with
SRDF/S*** How to install and use File Auto Recovery with SRDF/S.

***Dell EMC
PowerMax eNAS
CLI Reference
Guide*** A reference for command-line users and script programmers that provides the syntax, error codes, and parameters of all eNAS commands.


PowerProtect Storage Direct documents:

***Dell EMC
PowerProtect
Storage Direct
Solutions Guide*** Provides Storage Direct information that is related to various data objects and data handling facilities.

***Dell EMC File
System Agent
Installation and
Administration
Guide*** Shows how to install, configure, and manage the Storage Direct File System Agent.

***Dell EMC
Database
Application
Agent
Installation and
Administration
Guide*** Shows how to install, configure, and manage the Storage Direct Database Application Agent.

***Dell EMC
Microsoft
Application
Agent
Installation and
Administration
Guide*** Shows how to install, configure, and manage the Storage Direct Microsoft Application Agent.

 **NOTE:** ProtectPoint has been renamed to Storage Direct and it is included in PowerProtect, Data Protection Suite for Apps, or Data Protection Suite Enterprise Software Edition.

Mainframe Enablers documents:

| | |
|--|--|
| <i>Dell EMC Mainframe Enablers Installation and Customization Guide</i> | Describes how to install and configure Mainframe Enablers software. |
| <i>Dell EMC Mainframe Enablers Release Notes</i> | Describes new features and any known limitations. |
| <i>Dell EMC Mainframe Enablers Message Guide</i> | Describes the status, warning, and error messages generated by Mainframe Enablers software. |
| <i>Dell EMC Mainframe Enablers ResourcePak Base for z/OS Product Guide</i> | Describes how to configure VMAX system control and management using the EMC Symmetrix Control Facility (EMCSCF). |
| <i>Dell EMC Mainframe Enablers AutoSwap for z/OS Product Guide</i> | Describes how to use AutoSwap to perform automatic workload swaps between VMAX systems when the software detects a planned or unplanned outage. |
| <i>Dell EMC Mainframe Enablers Consistency Groups for z/OS Product Guide</i> | Describes how to use Consistency Groups for z/OS (ConGroup) to ensure the consistency of data remotely copied by SRDF in the event of a rolling disaster. |
| <i>Dell EMC Mainframe Enablers SRDF Host Component for z/OS Product Guide</i> | Describes how to use SRDF Host Component to control and monitor remote data replication processes. |
| <i>Dell EMC Mainframe Enablers TimeFinder SnapVX and zDP Product Guide</i> | Describes how to use TimeFinder SnapVX and zDP to create and manage space-efficient targetless snaps. |
| <i>Dell EMC Mainframe Enablers TimeFinder/Clone Mainframe Snap Facility Product Guide</i> | Describes how to use TimeFinder/Clone, TimeFinder/Snap, and TimeFinder/CG to control and monitor local data replication processes. |
| <i>Dell EMC Mainframe Enablers TimeFinder/Mirror for z/OS Product Guide</i> | Describes how to use TimeFinder/Mirror to create Business Continuance Volumes (BCVs) which can then be established, split, reestablished and restored from the source logical volumes for backup, restore, decision support, or application testing. |

**Dell EMC
Mainframe
Enablers
TimeFinder
Utility for z/OS
Product Guide**

Describes how to use the TimeFinder Utility to condition volumes and devices.

Geographically Dispersed Disaster Recovery (GDDR) documents:

**Dell EMC GDDR
for SRDF/S
with ConGroup
Product Guide**

Describes how to use Geographically Dispersed Disaster Restart (GDDR) to automate business recovery following both planned outages and disaster situations.

**Dell EMC GDDR
for SRDF/S
with AutoSwap
Product Guide**

Describes how to use Geographically Dispersed Disaster Restart (GDDR) to automate business recovery following both planned outages and disaster situations.

**Dell EMC GDDR
for SRDF/Star
Product Guide**

Describes how to use Geographically Dispersed Disaster Restart (GDDR) to automate business recovery following both planned outages and disaster situations.

**Dell EMC GDDR
for SRDF/Star
with AutoSwap
Product Guide**

Describes how to use Geographically Dispersed Disaster Restart (GDDR) to automate business recovery following both planned outages and disaster situations.

**Dell EMC GDDR
for SRDF/SQAR
with AutoSwap
Product Guide**

Describes how to use Geographically Dispersed Disaster Restart (GDDR) to automate business recovery following both planned outages and disaster situations.

**Dell EMC GDDR
for SRDF/A
Product Guide**

Describes how to use Geographically Dispersed Disaster Restart (GDDR) to automate business recovery following both planned outages and disaster situations.

**Dell EMC GDDR
Message Guide**

Describes the status, warning, and error messages generated by GDDR.

**Dell EMC GDDR
Release Notes**

Describes new features and any known limitations.

**Dell EMC GDDR
for Star-A
Product Guide**

Describes the basic concepts of Dell EMC Geographically Dispersed Disaster Restart (GDDR), how to install it, and how to implement its major features and facilities.

z/OS Migrator documents:

**Dell EMC z/OS
Migrator Product
Guide**

Describes how to use z/OS Migrator to perform volume mirror and migrator functions as well as logical migration functions.

**Dell EMC
z/OS Migrator
Message Guide**

Describes the status, warning, and error messages generated by z/OS Migrator.

**Dell EMC z/OS
Migrator Release
Notes**

Describes new features and any known limitations.

z/TPF documents:

**Dell EMC
ResourcePak for
z/TPF Product
Guide**

Describes how to configure VMAX system control and management in the z/TPF operating environment.

**Dell EMC SRDF
Controls for**

Describes how to perform remote replication operations in the z/TPF operating environment.

z/TPF Product Guide

Dell EMC TimeFinder Controls for z/TPF Product Guide Describes how to perform local replication operations in the z/TPF operating environment.

Dell EMC z/TPF Suite Release Notes Describes new features and any known limitations.

Typographical conventions

Dell EMC uses the following type style conventions in this document:

Table 1. Typographical conventions used in this content

| | |
|-------------------------|--|
| Bold | Used for names of interface elements Examples: Names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user selects or clicks) |
| <i>Italic</i> | Used for full titles of publications referenced in text |
| Monospace | Used for: <ul style="list-style-type: none">• System code• System output, such as an error message or script• Pathnames, filenames, prompts, and syntax• Commands and options |
| <i>Monospace italic</i> | Used for variables |
| Monospace bold | Used for user input |
| [] | Square brackets enclose optional values. |
| | A vertical bar indicates alternate selections. The bar means "or". |
| { } | Braces enclose content that the user must specify, such as x or y or z. |
| ... | Ellipses indicate nonessential information that is omitted from the example. |

Where to get help

Dell EMC support, product, and licensing information can be obtained as follows:

| | |
|-----------------------------------|---|
| Product information | Dell EMC technical support, documentation, release notes, software updates, or information about Dell EMC products can be obtained at https://www.dell.com/support/home (registration required) or https://www.dell.com/en-us/documentation/vmax-all-flash-family.htm . |
| Technical support | To open a service request through the Dell EMC Online Support (https://www.dell.com/support/home) site, you must have a valid support agreement. Contact your Dell EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account. |
| Additional support options | <ul style="list-style-type: none">• Support by Product — Dell EMC offers consolidated, product-specific information on the Web at: https://support.EMC.com/products The Support by Product web pages offer quick links to Documentation, White Papers, Advisories (such as frequently used Knowledgebase articles), and Downloads, as well as more dynamic content, such as presentations, discussion, relevant Customer Support Forum entries, and a link to Dell EMC Live Chat.• Dell EMC Live Chat — Open a Chat or instant message session with a Dell EMC Support Engineer. |

e-Licensing support

To activate your entitlements and obtain your license files, go to the Service Center on Dell EMC Online Support (<https://www.dell.com/support/home>). Follow the directions on your License Authorization Code (LAC) letter that is emailed to you.

- Expected functionality may be unavailable because it is not licensed. For help with missing or incorrect entitlements after activation, contact your Dell EMC Account Representative or Authorized Reseller.
- For help with any errors applying license files through Solutions Enabler, contact the Dell EMC Customer Support Center.
- Contact the Dell EMC worldwide Licensing team if you are missing a LAC letter or require further instructions on activating your licenses through the Online Support site.
 - licensing@emc.com
 - North America, Latin America, APJK, Australia, New Zealand: SVC4EMC (800-782-4362) and follow the voice prompts.
 - EMEA: +353 (0) 21 4879862 and follow the voice prompts.

Your comments

Your suggestions help improve the accuracy, organization, and overall quality of the documentation. Send your comments and feedback to: VMAXContentFeedback@emc.com

PowerMax with PowerMaxOS

This chapter introduces PowerMax systems and the PowerMaxOS operating environment.

Topics:

- [Introduction to PowerMax with PowerMaxOS](#)
- [Software packages](#)
- [PowerMaxOS](#)

Introduction to PowerMax with PowerMaxOS

PowerMax arrays

The PowerMax family of arrays has two models:

- PowerMax 2000 with a maximum capacity of 1 PBe (Petabytes effective) that can operate in open systems environments
- PowerMax 8000 with a maximum capacity of 4 PBe that can operate in open systems, mainframe, or mixed open systems and mainframe environments

PowerMax systems are modular enabling them to expand to meet the future needs of the customer.

System building blocks

Each PowerMax array is made up of one or more building blocks each known as a PowerMax Brick in an open systems array or a PowerMax zBrick in a mainframe array. A PowerMax Brick or PowerMax zBrick consists of:

- An engine with two directors (the redundant data storage processing unit)
- Flash storage in two Drive Array Enclosures (DAEs) each with 24 slots
- Minimum storage capacity:
 - PowerMax 2000: 13 TBu (Terabytes usable)
 - PowerMax 8000 in an open systems environment: 53 TBu
 - PowerMax 8000 in a mainframe environment: 13 TBu
 - PowerMax 8000 in a mixed open systems and mainframe environment: 66 TBu

Hardware expansion

Customers can increase the initial storage capacity in 13 TBu units each known as a Flash Capacity Pack (in an open systems environment) or a zFlash Capacity Pack (in a mainframe environment). The addition of Flash Capacity Packs or zFlash Capacity Packs to an array is known as *scaling up*.

Also, customers can add further PowerMax Bricks or PowerMax zBricks to increase the capacity and capability of the system. A PowerMax 2000 array can have a maximum of two PowerMax Bricks. A PowerMax 8000 can have a maximum of eight PowerMax Bricks or PowerMax zBricks. The addition of bricks to an array is known as *scaling out*.

Finally, customers can increase the internal memory of the system. A PowerMax 2000 system can have 512 GB, 1 TB, or 2 TB of memory on each engine. A PowerMax 8000 system can have 1 TB or 2 TB of memory on each engine.

Storage devices

Starting with PowerMaxOS 5978.444.444 there are two types of storage devices available for a PowerMax array:

- NVMe flash drive

- Storage Class Memory (SCM) drive

SCM drives are available with PowerMaxOS 5978.444.444 and later. Previous versions of PowerMaxOS 5978 work with NVMe flash drives only.

SCM drives are new, high-performance drives that have a significantly lower latency than the NVMe flash drives. An eligible array can have any mix of SCM drives and NVMe drives.

In SCM-based systems:

- Customers can increase the capacity of SCM drives in increments of 5.25 TBu.
- The minimum starting capacity of a SCM-based system is 21 TBu.

System specifications

Detailed specifications of the PowerMax arrays are available from the [Dell EMC](#) website.

PowerMaxOS operating environment

PowerMaxOS 5978 is the software operating environment for PowerMax arrays. It manages the storage and controls communications with the host systems. There are additional features for PowerMaxOS that provide specific capabilities such as remote replication. The software for a PowerMax is available in packages that consist of a core set of features and additional, optional features. There are two packages for open systems arrays and two for mainframe arrays.

Further information:

- [Software packages](#) on page 19 has information about the software packages, their contents, and their availability on the various PowerMax platforms.
- [PowerMaxOS](#) on page 22 has information about the capabilities and components of PowerMaxOS.

PowerMaxOS can also run on VMAX All Flash arrays.

Software packages

There are four software packages for PowerMax arrays. The Essentials and Pro software packages are for open system arrays while the zEssentials and zPro software packages are for mainframe arrays.

The Essentials software package

Standard features

The standard features in the Essentials software package are:

| Feature | For more information, see |
|--|---|
| PowerMaxOS | PowerMaxOS on page 22 |
| Embedded Management (eManagement) ^a | Management Interfaces on page 35 |
| Compression and deduplication | Inline compression on page 33 and Inline deduplication on page 34 |
| SnapVX | About TimeFinder on page 67 |
| AppSync Starter Pack | AppSync on page 43 |
| Migration | Data migration for open systems on page 98 |
| VMware Virtual Volumes support | VMware Virtual Volumes on page 48 |
| Cloud Mobility | Cloud Mobility for Dell EMC PowerMax on page 91 |

a. eManagement includes embedded Unisphere, Solutions Enabler, and SMI-S.

Optional features

The optional features in the Essentials software package are:

| Feature | For more information, see |
|--|--|
| SRDF | Remote replication on page 73 |
| SRDF/Metro | SRDF/Metro on page 85 |
| Embedded Network Attached Storage (eNAS) | Embedded Network Attached Storage (eNAS) on page 25 |
| Data at Rest Encryption (D@RE) | Data at Rest Encryption on page 27 |
| SRM | Storage Resource Management (SRM) on page 40 |
| Unisphere 360 | Unisphere 360 on page 36 |
| PowerProtect Storage Direct | Backup and restore using PowerProtect Storage Direct and Data Domain on page 46 |
| PowerPath | PowerPath on page 44 and PowerPath Migration Enabler on page 106 |
| RecoverPoint | RecoverPoint on page 89 |
| AppSync Full Suite | AppSync on page 43 |

The Pro software package

Standard features

The Pro software package contains all the standard features of the Essentials software package plus:

| Feature | For more information, see |
|------------------------|--|
| D@RE | Data at Rest Encryption on page 27 |
| SRDF | Remote replication on page 73 |
| SRDF/Metro | SRDF/Metro on page 85 |
| eNAS | Embedded Network Attached Storage (eNAS) on page 25 |
| Unisphere 360 | Unisphere 360 on page 36 |
| SRM | Storage Resource Management (SRM) on page 40 |
| PowerPath ^a | PowerPath on page 44 and PowerPath Migration Enabler on page 106 |
| AppSync Full Suite | AppSync on page 43 |

a. The Pro software package contains 75 PowerPath licenses. Extra licenses are available separately.

Optional features

The optional features of the Pro software package are:

| Feature | For more information, see |
|-----------------------------|---|
| PowerProtect Storage Direct | Backup and restore using PowerProtect Storage Direct and Data Domain on page 46 |
| RecoverPoint | RecoverPoint on page 89 |

The zEssentials software package

Standard features

The standard features in the zEssentials software package are:

| Feature | For more information, see |
|--|--|
| PowerMaxOS | PowerMaxOS on page 22 |
| Embedded Management (eManagement) ^a | Management Interfaces on page 35 |
| Snap VX | About TimeFinder on page 67 |
| Mainframe Essentials | Mainframe Features on page 51 |

a. eManagement includes embedded Unisphere, Solutions Enabler, and SMI-S.

Optional Features

The optional features in the zEssentials software package are:

| Feature | For more information, see |
|---------------------------|---|
| SRDF | Remote replication on page 73 |
| D@RE | Data at Rest Encryption on page 27 |
| Unisphere 360 | Unisphere 360 on page 36 |
| zDP | Mainframe SnapVX and zDP on page 70 |
| AutoSwap | Mainframe SnapVX and zDP on page 70 |
| GDDR | Geographically Dispersed Disaster Restart (GDDR) on page 40 |
| Mainframe Essentials Plus | Mainframe Features on page 51 |
| Product Suite for z/TPF | Product Suite for z/TPF on page 42 |

The zPro software package

Standard features

The zPro software package contains all the standard features of the zEssentials software package plus:

| Feature | For more information, see |
|---------------|---|
| SRDF | Remote replication on page 73 |
| D@RE | Data at Rest Encryption on page 27 |
| Unisphere 360 | Unisphere 360 on page 36 |
| zDP | Mainframe SnapVX and zDP on page 70 |
| AutoSwap | Mainframe SnapVX and zDP on page 70 |

Optional features

The optional features in the zPro software package are:

| Feature | For more information, see |
|---------------------------|---|
| GDDR | Geographically Dispersed Disaster Restart (GDDR) on page 40 |
| Mainframe Essentials Plus | Mainframe Features on page 51 |
| Product Suite for z/TPF | Product Suite for z/TPF on page 42 |

Package availability

The availability of the PowerMaxOS software packages on the PowerMax platforms is:

| Software package | Platforms |
|------------------------------|---------------|
| Essentials software package | PowerMax 8000 |
| Pro software package | PowerMax 2000 |
| zEssentials software package | PowerMax 8000 |
| zPro software package | |

PowerMaxOS

This section summarizes the main features of PowerMaxOS.

PowerMaxOS emulations

PowerMaxOS provides emulations (executables) that perform specific data service and control functions in the PowerMaxOS environment. The available emulations are:

Table 2. PowerMaxOS emulations

| Area | Emulation | Description | Protocol Speed ^a |
|------------|-----------|---|-----------------------------|
| Back-end | DN | Back-end connection in the array that communicates with the drives, DN is also known as an internal drive controller. | NVMe - 8 Gb/s |
| | DX | Back-end connections that are not used to connect to hosts. Used by Storage Direct. | FC - 16 Gb/s |
| | | Storage Direct links Data Domain to the array. DX ports must be configured for the FC protocol. | |
| Management | IM | Separates infrastructure tasks and emulations. By separating these tasks, emulations can focus on I/O-specific work only, while IM manages and runs common infrastructure tasks, such as environmental monitoring, Field Replacement Unit (FRU) monitoring, and vaulting. | N/A |

Table 2. PowerMaxOS emulations (continued)

| Area | Emulation | Description | Protocol Speed ^a |
|--------------------|---|--|--|
| | ED | Middle layer used to separate front end and back-end I/O processing. It acts as a translation layer between the front-end, which is what the host knows about, and the back-end, which is the layer that reads, writes, and communicates with physical storage in the array. | N/A |
| Host connectivity | FA - Fibre Channel SE - iSCSI EF - FICON ^b FN - FC-NVMe | Front-end emulation that: <ul style="list-style-type: none"> • Receives data from the host or network and commits it to the array • Sends data from the array to the host or network | FC - 16 Gb/s and 32 Gb/s SE - 10 Gb/s and 25 Gb/s EF - 16 Gb/s FN - 32 Gb/s ^{cd} |
| Remote replication | RF - Fibre Channel RE - GbE | Interconnects arrays for SRDF. | RF - 16 Gb/s and 32 Gb/s SRDF RE - 10 GbE and 25 GbE SRDF |

a. The 16 Gb/s module autonegotiates to 16/8/4 Gb/s using optical SFP and OM2/OM3/OM4 cabling.

b. Only on PowerMax 8000 arrays.

c. Available on PowerMax arrays only.

d. The 32 Gb/s module autonegotiates to 32/16/8 Gb/s.

Container applications

PowerMaxOS provides an open application platform for running data services. It includes a lightweight hypervisor that enables multiple operating environments to run as virtual machines on the storage array.

Application containers are virtual machines that provide embedded applications on the storage array. Each container virtualizes the hardware resources that are required by the embedded application, including:

- Hardware needed to run the software and embedded application (processor, memory, PCI devices, power management)
- VM ports, to which LUNs are provisioned
- Access to necessary drives (boot, root, swap, persist, shared)

Embedded Management

The eManagement container application embeds management software (Solutions Enabler, SMI-S, Unisphere for PowerMax) on the storage array, enabling you to manage the array without requiring a dedicated management host.

With eManagement, you can manage a single storage array and any SRDF attached arrays. To manage multiple storage arrays with a single control pane, use the traditional host-based management interfaces: Unisphere and Solutions Enabler. To this end, eManagement allows you to link-and-launch a host-based instance of Unisphere.

eManagement is typically preconfigured and enabled at the factory. However, eManagement can be added to arrays in the field. Contact your support representative for more information.

Embedded applications require system memory. The following table lists the amount of memory unavailable to other data services.

Table 3. eManagement resource requirements

| PowerMax model | CPUs | Memory | Devices supported |
|----------------|------|--------|-------------------|
| PowerMax 2000 | 4 | 16 GB | 200K |
| PowerMax 8000 | 4 | 20 GB | 400K |

Virtual machine ports

Virtual machine (VM) ports are associated with virtual machines to avoid contention with physical connectivity. VM ports are addressed as ports 32-63 on each director FA emulation.

LUNs are provisioned on VM ports using the same methods as provisioning physical ports.

A VM port can be mapped to one VM only. However, a VM can be mapped to multiple ports.

Embedded Network Attached Storage (eNAS)

eNAS is fully integrated into the PowerMax array. eNAS provides flexible and secure multi-protocol file sharing (NFS 2.0, 3.0, 4.0/4.1, and CIFS/SMB 3.0) and multiple file server identities (CIFS and NFS servers). eNAS enables:

- File server consolidation/multi-tenancy
- Built-in asynchronous file level remote replication (File Replicator)
- Built-in Network Data Management Protocol (NDMP)
- VDM Synchronous replication with SRDF/S and optional automatic failover manager File Auto Recovery (FAR)
- Built-in creation of point-in-time logical images of a production file system using SnapSure
- Anti-virus

eNAS provides file data services for:

- Consolidating block and file storage in one infrastructure
- Eliminating the gateway hardware, reducing complexity and costs
- Simplifying management

Consolidated block and file storage reduces costs and complexity while increasing business agility. Customers can leverage data services across block and file storage including storage provisioning, dynamic Host I/O Limits, and Data at Rest Encryption.

eNAS solutions and implementation

eNAS runs on standard array hardware and is typically pre-configured at the factory. There is a one-off setup of the Control Station and Data Movers, containers, control devices, and required masking views as part of the factory pre-configuration. Additional front-end I/O modules are required to implement eNAS. eNAS can be added to arrays in the field. Contact your support representative for more information.

eNAS uses the PowerMaxOS hypervisor to create virtual instances of NAS Data Movers and Control Stations on PowerMax controllers. Control Stations and Data Movers are distributed within the PowerMax array based upon the number of engines and their associated mirrored pair.

By default, PowerMax arrays have:

- Two Control Station virtual machines
- Two or more Data Mover virtual machines. The number of Data Movers differs for each model of the array. All configurations include one standby Data Mover.

eNAS configurations

The storage capacity required for arrays supporting eNAS is at least 680 GB. This table lists eNAS configurations and front-end I/O modules.

Table 4. eNAS configurations by array


| Component | Description | PowerMax 2000 | PowerMax 8000 |
|--|--------------------------------|-----------------|----------------------------|
| Data Movers ^a virtual machine | Maximum number | 4 | 8 ^b |
| | Max capacity/DM | 512 TB | 512 TB |
| | Logical cores ^c | 12/24 | 16/32/48/64 ^b |
| | Memory (GB) ^c | 48/96 | 48/96/144/192 ^b |
| | I/O modules (Max) ^c | 12 ^d | 24 ^d |
| Control Station virtual machines (2) | Logical cores | 4 | 4 |
| | Memory (GB) | 8 | 8 |
| NAS Capacity/Array | Maximum | 1.15 TB | 3.5 TB |

- a. Data Movers are added in pairs and must have the same configuration.
- b. The PowerMax 8000 can be configured through Sizer with a maximum of four Data Movers. However, six and eight Data Movers can be ordered by RPQ. As the number of data movers increases, the maximum number of I/O cards , logical cores, memory, and maximum capacity also increases.
- c. For 2, 4, 6, and 8 Data Movers, respectively.
- d. A single 2-port 10GbE Optical I/O module is required by each Data Mover for initial PowerMax configurations. However, that I/O module can be replaced with a different I/O module (such as a 4-port 1GbE or 2-port 10GbE copper) using the normal replacement capability that exists with any eNAS Data Mover I/O module. Also, additional I/O modules can be configured through a I/O module upgrade/add as long as standard rules are followed (no more than three I/O modules per Data Mover, all I/O modules must occupy the same slot on each director on which a Data Mover resides).

Replication using eNAS

The replication methods available for eNAS file systems are:

- Asynchronous file system level replication using VNX Replicator for File.
- Synchronous replication with SRDF/S using File Auto Recovery (FAR).
- Checkpoint (point-in-time, logical images of a production file system) creation and management using VNX SnapSure.

 **NOTE:** SRDF/A, SRDF/Metro, and TimeFinder are not available with eNAS.

Data protection and integrity

PowerMaxOS provides facilities to ensure data integrity and to protect data in the event of a system failure or power outage:

- RAID levels
- Data at Rest Encryption
- End-to-end efficient encryption
- Data erasure
- Block CRC error checks
- Data integrity checks
- Drive monitoring and correction
- Physical memory error correction and error verification
- Drive sparing and direct member sparing
- Vault to flash

RAID levels

PowerMax arrays can use the following RAID levels:

- PowerMax 2000: RAID 5 (7+1) (Default), RAID 5 (3+1), RAID 6 (6+2), and RAID 1
- PowerMax 8000: RAID 5 (7+1), RAID 6 (6+2), and RAID 1

Data at Rest Encryption

Securing sensitive data is an important IT issue, that has regulatory and legislative implications. Several of the most important data security threats relate to protection of the storage environment. Drive loss and theft are primary risk factors. Data at Rest Encryption (D@RE) protects data by adding back-end encryption to an entire array.

D@RE provides hardware-based encryption for PowerMax arrays using I/O modules that incorporate AES-XTS inline data encryption. These modules encrypt and decrypt data as it is being written to or read from a drive. This protects your information from unauthorized access even when drives are removed from the array.

D@RE can use either an internal embedded key manager, or one of these external, enterprise-grade key managers:

- Gemalto SafeNet KeySecure
- IBM Security Key Lifecycle Manager

D@RE accesses an external key manager using the Key Management Interoperability Protocol (KMIP). The Dell EMC E-Lab Interoperability Matrix (<https://www.emc.com/products/interoperability/elab.htm>) lists the external key managers for each version of PowerMaxOS.

When D@RE is active, all configured drives are encrypted, including data drives, spares, and drives with no provisioned volumes. Vault data is encrypted on Flash I/O modules.

D@RE provides:

- Secure replacement for failed drives that cannot be erased—For some types of drive failures, data erasure is not possible. Without D@RE, if the failed drive is repaired, data on the drive may be at risk. With D@RE, deletion of the applicable keys makes the data on the failed drive unreadable.
- Protection against stolen drives—When a drive is removed from the array, the key stays behind, making data on the drive unreadable.
- Faster drive sparing—The drive replacement script destroys the keys associated with the removed drive, quickly making all data on that drive unreadable.
- Secure array retirement—Delete all copies of keys on the array, and all remaining data is unreadable.

In addition, D@RE:

- Is compatible with all array features and all supported drive types or volume emulations
- Provides encryption without degrading performance or disrupting existing applications and infrastructure

Enabling D@RE

D@RE is a licensed feature that is installed and configured at the factory. Upgrading an existing array to use D@RE is possible, but is disruptive. The upgrade requires re-installing the array, and may involve a full data back up and restore. Before upgrading, plan how to manage any data already on the array. Dell EMC Professional Services offers services to help you implement D@RE.

D@RE components

Embedded D@RE uses the following components, all of which reside on the primary Management Module Control Station (MMCS):

- Dell EMC Key Trust Platform (KTP) (embedded)—This component adds embedded key management functionality to the KMIP Client.
- Lockbox—Hardware- and software-specific encrypted repository that securely stores passwords and other sensitive key manager configuration information. The lockbox binds to a specific MMCS.

External D@RE uses the same components as embedded D@RE, and adds the following:

- Dell EMC Key Trust Platform (KTP)—Also known as the KMIP Client, this component resides on the MMCS and communicates with external key managers using the OASIS Key Management Interoperability Protocol (KMIP) to manage encryption keys.
- External Key Manager—Provides centralized encryption key management capabilities such as secure key generation, storage, distribution, audit, and enabling Federal Information Processing Standard (FIPS) 140-2.
- Cluster/Replication Group—Multiple external key managers sharing configuration settings and encryption keys. Configuration and key lifecycle changes made to one node are replicated to all members within the same cluster or replication group.

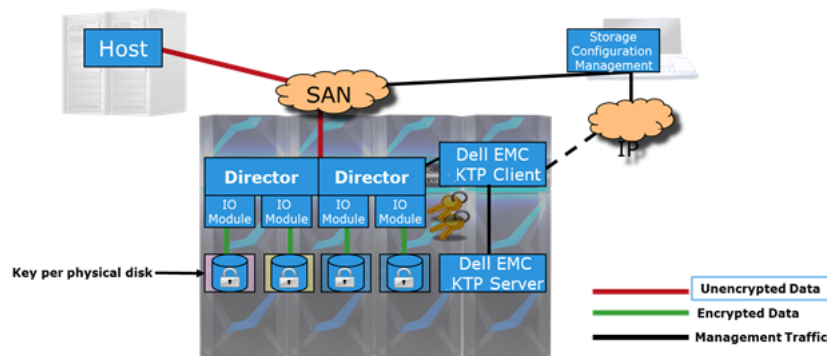


Figure 1. D@RE architecture

External Key Managers

D@RE's external key management is provided by Gemalto SafeNet KeySecure and IBM Security Key Lifecycle Manager. Keys are generated and distributed using industry standards (NIST 800-57 and ISO 11770). With D@RE, there is no need to replicate keys across volume snapshots or remote sites. D@RE external key managers can be used with FIPS 140-2 level 1 validated software.

Encryption keys must be highly available when they are needed, and tightly secured. Keys, and the information required to use keys (during decryption), must be preserved for the lifetime of the data. This is critical for encrypted data that is kept for many years.

Key accessibility is vital in high-availability environments. D@RE caches the keys locally. So connection to the Key Manager is necessary only for operations such as the initial installation of the array, replacement of a drive, or drive upgrades.

Lifecycle events involving keys (generation and destruction) are recorded in the array's Audit Log.

Key protection

The local keystore file is encrypted with a 256-bit AES key derived from a randomly generated password file. This password file is secured in the Lockbox. The Lockbox is protected using MMCS-specific stable system values (SSVs) of the primary MMCS. These are the same SSVs that protect Secure Service Credentials (SSC).

Compromising the MMCS drive or copying Lockbox and keystore files off the array causes the SSV tests to fail. Compromising the entire MMCS only gives an attacker access if they also successfully compromise SSC.

There are no backdoor keys or passwords to bypass D@RE security.

Key operations

D@RE provides a separate, unique Data Encryption Key (DEK) for each physical drive in the array, including spare drives. To ensure that D@RE uses the correct key for a given drive:

- DEKs stored in the array include a unique key tag and key metadata when they are wrapped (encrypted) for use by the array. This information is included with the key material when the DEK is wrapped (encrypted) for use in the array.
- During encryption I/O, the expected key tag associated with the drive is supplied separately from the wrapped key.
- During key unwrap, the encryption hardware checks that the key unwrapped correctly and that it matches the supplied key tag.
- Information in a reserved system LBA (Physical Information Block, or PHIB) verifies the key used to encrypt the drive and ensures the drive is in the correct location.
- During initialization, the hardware performs self-tests to ensure that the encryption/decryption logic is intact. The self-test prevents silent data corruption due to encryption hardware failures.

Audit logs

The audit log records major activities on an array, including:


- Host-initiated actions
- Physical component changes
- Actions on the MMCS
- D@RE key management events
- Attempts blocked by security controls (Access Controls)


The Audit Log is secure and tamper-proof so event contents cannot be altered. Users with the Auditor access can view, but not modify, the log.

End-to-end efficient encryption integration

End-to-end efficient encryption (EEEE) integration with array data services enables data reduction on host-encrypted data.

User data encrypted at the host or application for devices can be compressed and deduplicated using PowerMax data reduction services, when deployed with Thales' 'Efficient Storage' solution.

 **NOTE:** To ensure a complete end-to-end encryption of user data, this feature is only available on systems which have embedded D@RE enabled.

 **NOTE:** Use of end-to-end efficient encryption requires an RPQ.

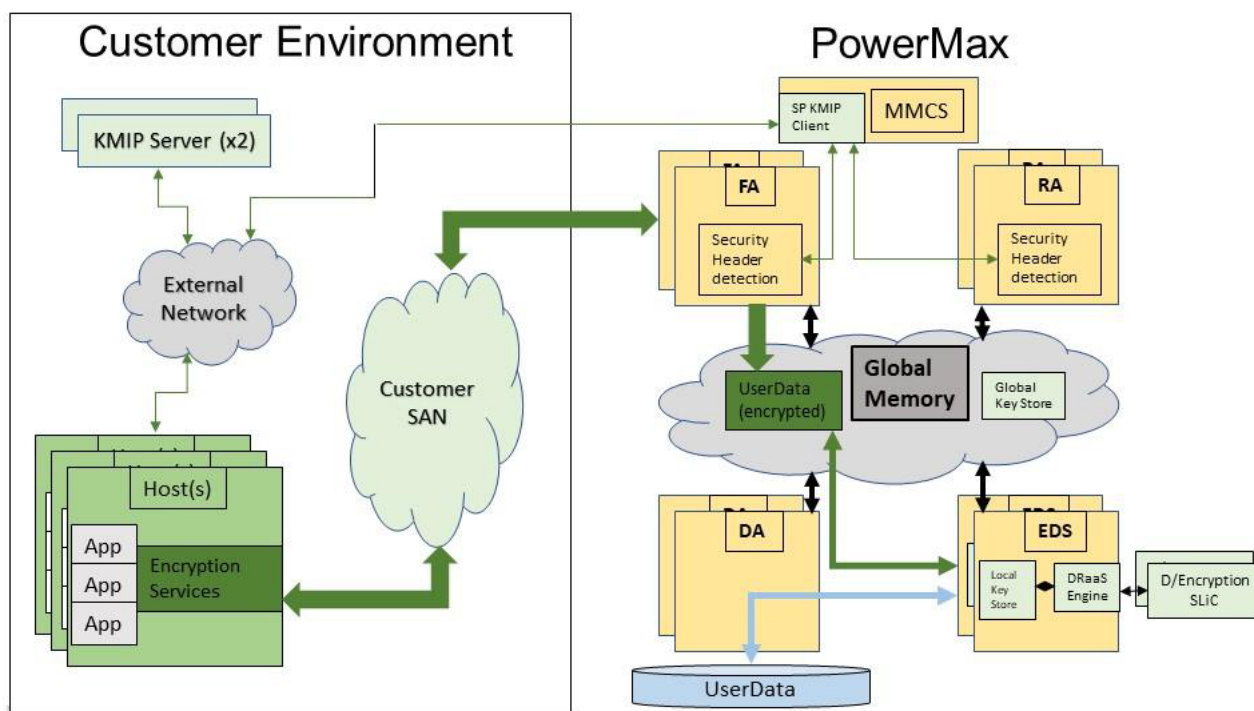


Figure 2. EEEE components

Data erasure

Dell EMC Data Erasure uses specialized software to erase information on arrays. It mitigates the risk of information dissemination, and helps secure information at the end of the information lifecycle. Data erasure:

- Protects data from unauthorized access
- Ensures secure data migration by making data on the source array unreadable
- Supports compliance with internal policies and regulatory requirements

Data Erasure overwrites data at the lowest application-addressable level to drives. The number of overwrites is configurable from three (the default) to seven with a combination of random patterns on the selected arrays.

An optional certification service is available to provide a certificate of erasure. Drives that fail erasure are delivered to customers for final disposal.

For individual flash drives, Secure Erase operations erase all physical flash areas on the drive which may contain user data.

The available data erasure services are:

- Dell EMC Data Erasure for Full Arrays—Overwrites data on all drives in the system when replacing, retiring or re-purposing an array.
- Dell EMC Data Erasure/Single Drives—Overwrites data on individual drives.
- Dell EMC Disk Retention—Enables organizations that must retain all media to retain failed drives.
- Dell EMC Assessment Service for Storage Security—Assesses your information protection policies and suggests a comprehensive security strategy.

All erasure services are performed on-site in the security of the customer's data center and include a Data Erasure Certificate and report of erasure results.

Block CRC error checks

PowerMaxOS provides:

- Industry-standard, T10 Data Integrity Field (DIF) block cyclic redundancy check (CRC) for track formats.

For open systems, this enables host-generated DIF CRCs to be stored with user data by the arrays and used for end-to-end data integrity validation.

- Additional protections for address/control fault modes for increased levels of protection against faults. These protections are defined in user-definable blocks supported by the T10 standard.
- Address and write status information in the extra bytes in the application tag and reference tag portion of the block CRC.

Data integrity checks

PowerMaxOS validates the integrity of data at every possible point during the lifetime of that data. From the time data enters an array, it is continuously protected by error detection metadata. This metadata is checked by hardware and software mechanisms any time data is moved within the array. This allows the array to provide true end-to-end integrity checking and protection against hardware or software faults.

The protection metadata is appended to the data stream, and contains information describing the expected data location as well as the CRC representation of the actual data contents. The expected values to be found in protection metadata are stored persistently in an area separate from the data stream. The protection metadata is used to validate the logical correctness of data being moved within the array any time the data transitions between protocol chips, internal buffers, internal data fabric endpoints, system cache, and system drives.

Drive monitoring and correction

PowerMaxOS monitors medium defects by both examining the result of each disk data transfer and proactively scanning the entire disk during idle time. If a block on the disk is determined to be bad, the director:

1. Rebuilds the data in the physical storage, if necessary.
2. Rewrites the data in physical storage, if necessary.

The director keeps track of each bad block detected on a drive. If the number of bad blocks exceeds a predefined threshold, the array proactively invokes a sparing operation to replace the defective drive, and then alerts Customer Support to arrange for corrective action, if necessary. With the deferred service model, immediate action is not always required.

Physical memory error correction and error verification

PowerMaxOS corrects single-bit errors and reports an error code once the single-bit errors reach a predefined threshold. In the unlikely event that physical memory replacement is required, the array notifies Customer Support, and a replacement is ordered.

Drive sparing and direct member sparing

When PowerMaxOS 5978 detects a drive is about to fail or has failed, it starts a direct member sparing (DMS) process. Direct member sparing looks for available spares within the same engine that are of the same or larger capacity and performance, with the best available spare always used.

With direct member sparing, the invoked spare is added as another member of the RAID group. During a drive rebuild, the option to directly copy the data from the failing drive to the invoked spare drive is available. The failing drive is removed only when the copy process is complete. Direct member sparing is automatically initiated upon detection of drive-error conditions.

Direct member sparing provides the following benefits:

- The array can copy the data from the failing RAID member (if available), removing the need to read the data from all of the members and doing the rebuild. Copying to the new RAID member is less CPU intensive.
- If a failure occurs in another member, the array can still recover the data automatically from the failing member (if available).
- More than one spare for a RAID group is supported at the same time.

Vault to flash

PowerMax arrays initiate a vault operation when the system is powered down, goes offline, or if environmental conditions occur, such as the loss of a data center due to an air conditioning failure.

Each array comes with Standby Power Supply (SPS) modules. On a power loss, the array uses the SPS power to write the system mirrored cache to flash storage. Vaulted images are fully redundant; the contents of the system mirrored cache are saved twice to independent flash storage.

The vault operation

When a vault operation starts:

- During the save part of the vault operation, the PowerMax array stops all I/O. When the system mirrored cache reaches a consistent state, directors write the contents to the vault devices, saving two copies of the data. The array then completes the power down, or, if power down is not required, remains in the offline state.
- During the restore part of the operation, the array's startup program initializes the hardware and the environmental system, and restores the system mirrored cache contents from the saved data (while checking data integrity).

The system resumes normal operation when the SPS modules have sufficient charge to complete another vault operation, if required. If any condition is not safe, the system does not resume operation and notifies Customer Support for diagnosis and repair. This allows Customer Support to communicate with the array and restore normal system operations.

Vault configuration considerations

- To support vault to flash, the PowerMax arrays require the following number of flash I/O modules:
 - PowerMax 2000 two to four per engine/PowerMax Brick
 - PowerMax 8000 four to eight per engine/PowerMax Brick/PowerMax zBrick
- The size of the flash module is determined by the amount of system cache and metadata required for the configuration.
- The vault space is for internal use only and cannot be used for any other purpose when the system is online.
- The total capacity of all vault flash partitions is sufficient to keep two logical copies of the persistent portion of the system mirrored cache.

Data efficiency

Data efficiency is a feature of PowerMax systems that is designed to make the best available use of the storage space on a storage system. Data efficiency has two elements:

- Inline compression
- Deduplication

They work together to reduce the amount of storage that an individual storage group requires. The space savings achieved through data efficiency is measured as the Data Reduction Ratio (DRR). Data efficiency operates on individual storage groups so that a system can have a mix of storage groups that use data efficiency and those that do not.

Inline compression

Inline compression is a feature of storage groups. When enabled (this is the default setting), new I/O to a storage group is compressed when written to disk, while existing data on the storage group starts to compress in the background. After turning off compression, new I/O is no longer compressed, and existing data remains compressed until it is written again, at which time it decompresses.

Inline compression, deduplication, and over-subscription complement each other. Over-subscription allows presenting larger than needed devices to hosts without having the physical drives to fully allocate the space represented by the thin devices ([Thin device oversubscription](#) on page 58 has more information on over-subscription). Inline compression further reduces the data footprint by increasing the effective capacity of the array.

The example in [Inline compression and over-subscription](#) on page 33 shows this. Here, 1.3 PB of host attached devices (TDEVs) is over-provisioned to 1.0 PB of back-end (TDATs), that reside on 1.0 PB of Flash drives. Following data compression, the data blocks are compressed, by a ratio of 2:1, reducing the number of Flash drives by half. Basically, with compression enabled, the array requires half as many drives to support a given front-end capacity.

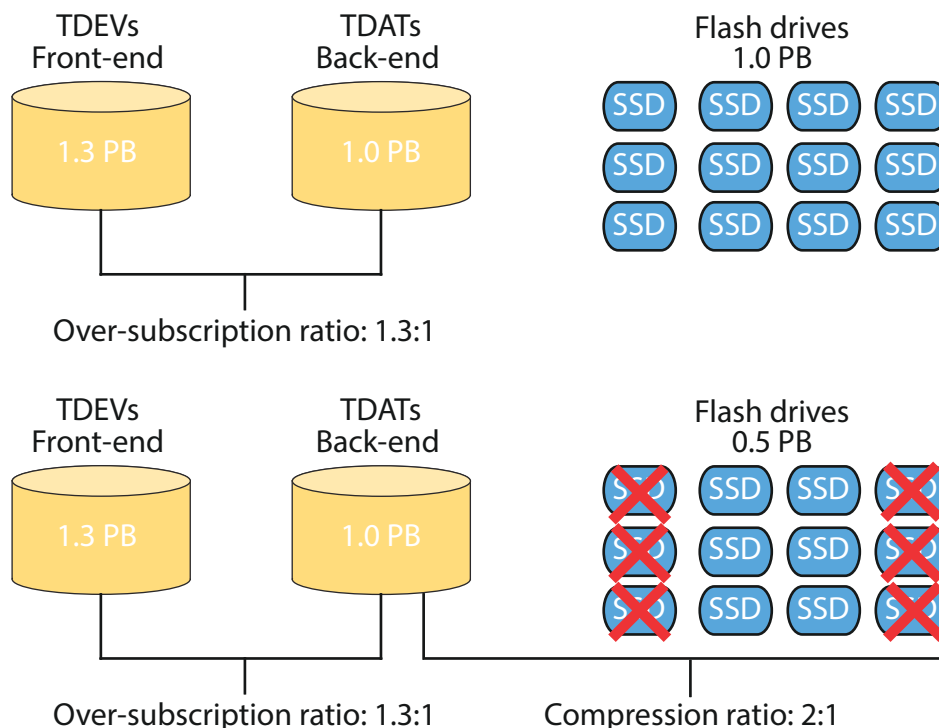


Figure 3. Inline compression and over-subscription

Further characteristics of compression are:

- All supported data services, such as SnapVX, SRDF, and encryption are supported with compression.
- Compression is available on open systems (FBA) only (including eNAS). It is not available for CKD arrays, including those with a mix of FBA and CKD devices. Any open systems array with compression enabled, cannot have CKD devices added to it.
- Compression is switched on and off through Solutions Enabler and Unisphere.
- Compression efficiency can be monitored for SRPs, storage groups, and volumes.
- Activity Based Compression: the most active tracks are held in cache and not compressed until they move from cache to disk. This feature helps improve the overall performance of the array while reducing wear on the flash drives.

Software compression

PowerMaxOS 5978 introduces software compression for PowerMax arrays. Software compression is an extension of regular, inline compression and is available on PowerMax systems only. It operates on data that was previously compressed but has not been accessed for 35 days or more. Software compression recompresses this data using an algorithm that may produce a much greater DRR. The amount of extra compression that can be achieved depends on the nature of the data.

The criteria that software compression uses to select a data extent for recompression are:

- The extent is in a storage group that is enabled for compression
- The extent has not already been recompressed by software compression
- The extent has not been accessed in the previous 35 days

Software compression runs in the background, using CPU cycles that would otherwise be free. Therefore, it does not impact the performance of the storage system. Also, software compression does not require any user intervention as it automatically selects and recompresses idle data.

Inline deduplication

Deduplication works in conjunction with inline compression to further improve efficiency in the use of storage space. It reduces the number of copies of identical tracks that are stored on back-end devices. Depending on the nature of the data, deduplication can provide additional data reduction over and above the reduction that compression provides.

The storage group is the unit that deduplication works on. When it detects a duplicated track in a group, deduplication replaces it with a pointer to the track that already resides on back-end storage.

Availability

Deduplication is available only on PowerMax arrays that run PowerMaxOS. In addition, deduplication works on FBA data only. A system with a mix of FBA and CKD devices can use deduplication, even when the FBA and CKD devices occupy separate SRPs.

Relationship with inline compression

The combination of compression and deduplication is called DRR (Data Reduction Ratio). DRR is enabled by default for all storage groups but it can be disabled if required.

In addition, deduplication operates across an entire system. It is not possible to use compression only on some storage groups and compression with deduplication on others.

Compatibility

Deduplication is compatible with the Dell EMC Live Optics performance analyzer. An array with deduplication can participate in a performance study of an IT environment.

User management

Solutions Enabler or Unisphere for PowerMax have facilities to manage deduplication, including:

- Selecting the storage groups to use deduplication
- Monitoring the performance of the system

[Management Interfaces](#) on page 35 contains an overview of Solutions Enabler and Unisphere for PowerMax.

Management Interfaces

This chapter introduces the tools for managing arrays.

Topics:

- [Management interface versions](#)
- [Unisphere for PowerMax](#)
- [Unisphere 360](#)
- [CloudIQ](#)
- [Solutions Enabler](#)
- [Mainframe Enablers](#)
- [Geographically Dispersed Disaster Restart \(GDDR\)](#)
- [SMI-S Provider](#)
- [VASA Provider](#)
- [eNAS management interface](#)
- [Storage Resource Management \(SRM\)](#)
- [vStorage APIs for Array Integration](#)
- [SRDF Adapter for VMware vCenter Site Recovery Manager](#)
- [SRDF/Cluster Enabler](#)
- [Product Suite for z/TPF](#)
- [SRDF/TimeFinder Manager for IBM i](#)
- [AppSync](#)
- [EMC Storage Analytics \(ESA\)](#)

Management interface versions

The following components provide management capabilities for PowerMaxOS:

- Unisphere for PowerMax V9.2
- CloudIQ
- Solutions Enabler V9.2
- Mainframe Enablers V8.5
- GDDR V5.3
- Migrator V9.2
- SMI-S V9.2
- SRDF/CE V4.2.1
- SRA V6.3
- VASA Provider V9.2
- EMC Storage Analytics V4.5

Unisphere for PowerMax

Unisphere for PowerMax is a web-based application that provides provisioning, management, and monitoring of arrays.

With Unisphere you can perform the following tasks:

Table 5. Unisphere tasks

| Section | Allows you to: |
|---------|---|
| Home | View and manage functions such as array usage, alert settings, authentication options, system preferences, user authorizations, and link and launch client registrations. |

Table 5. Unisphere tasks (continued)

| Section | Allows you to: |
|-----------------|--|
| Storage | View and manage storage groups and storage tiers. |
| Hosts | View and manage initiators, masking views, initiator groups, array host aliases, and port groups. |
| Data Protection | View and manage local replication, monitor and manage replication pools, create and view device groups, and monitor and manage migration sessions. |
| Performance | Monitor and manage array dashboards, perform trend analysis for future capacity planning, and analyze data. Set preferences, such as, general, dashboards, charts, reports, data imports, and alerts for performance management tasks. |
| Databases | Troubleshoot database and storage issues, and launch Database Storage Analyzer. |
| System | View and display dashboards, active jobs, alerts, array attributes, and licenses. |
| Events | View alerts, the job list, and the audit log. |
| Support | View online help for Unisphere tasks. |

Unisphere also has a Representational State Transfer (REST) API. With this API you can access performance and configuration information, and provision storage arrays. You can use the API in any programming environment that supports standard REST clients, such as web browsers and programming platforms that can issue HTTP requests.

Workload Planner

Workload Planner displays performance metrics for applications. Use Workload Planner to:

- Model the impact of migrating a workload from one storage system to another.
- Model proposed new workloads.
- Assess the impact of moving one or more workloads off of a given array running PowerMaxOS.
- Determine current and future resource shortfalls that require action to maintain the requested workloads.

FAST Array Advisor

The FAST Array Advisor wizard guides you through the steps to determine the impact on performance of migrating a workload from one array to another.

If the wizard determines that the target array can absorb the added workload, it automatically creates all the auto-provisioning groups required to duplicate the source workload on the target array.

Unisphere 360

Unisphere 360 is an on-premises management solution that provides a single window across arrays running PowerMaxOS at a single site. Use Unisphere 360 to:

- Add a Unisphere server to Unisphere 360 to allow for data collection and reporting of Unisphere management storage system data.
- View the system health, capacity, alerts and capacity trends for your Data Center.
- View all storage systems from all enrolled Unisphere instances in one place.
- View details on performance and capacity.
- Link and launch to Unisphere instances running V8.2 or higher.
- Manage Unisphere 360 users and configure authentication and authorization rules.
- View details of visible storage arrays, including current and target storage.

CloudIQ

Cloud IQ is a web-based application for monitoring multiple PowerMax arrays simultaneously. However, CloudIQ is more than a passive monitor. It uses predictive analytics to help with:

- Visualizing trends in capacity usage
- Predicting potential shortcomings in capacity and performance so that early action can be taken to avoid them
- Troubleshooting performance issues

CloudIQ is available with PowerMaxOS 5978.221.221 and later, and with Unisphere for PowerMax V9.0.1 and later. It is free for customers to use.

Periodically, a data collector runs that gathers and packages data about the arrays that Unisphere manages and their performance. The collector then sends the packaged data to CloudIQ. On receiving the data, CloudIQ unpacks it, processes it, and makes it available to view in a UI.

CloudIQ is hosted on Dell EMC infrastructure that is secure, highly available, and fault tolerant. In addition, the infrastructure provides a guaranteed, 4-hour disaster recovery window.

The rest of this section contains more information on CloudIQ and how it interacts with a PowerMax array.

Connectivity

The data collector communicates with CloudIQ through a Secure Remote Services (SRS) gateway. SRS uses an encrypted connection running over HTTPS to exchange data with CloudIQ. The connection to the Secure Remote Services gateway is either through the secondary Management Modules Control Station (MMCS) within a PowerMax array, or through a direct connection from the management host that runs Unisphere. Connection through the MMCS requires that the array runs PowerMaxOS 5978.444.444.

The data collector is a component of Unisphere for PowerMax. So, it is installed along with Unisphere and you manage it with Unisphere.

Registration

Before you can monitor an array you register it with SRS using the Settings dialog in Unisphere for PowerMax. To be able to register an array you need a current support contract with Dell EMC. Once an array is registered, data collection can begin. If you wish you can exclude any array from data collection and hence being monitored by CloudIQ.

Data collection

The data collector gathers four categories of data and uses a different collection frequency for each category:

| Type of data | Collection frequency |
|---------------|----------------------|
| Alerts | 5 minutes |
| Performance | 5 minutes |
| Health | 5 minutes |
| Configuration | 1 hour |

In the Performance category, CloudIQ displays bandwidth, latency and IOPS (I/O operations). The values are calculated from these data items, collected from the array:

- Throughput read
- Throughput write
- Latency read
- Latency write
- IOPS read
- IOPS write

The Configuration category contains information on configuration, capacity, and efficiency for the overall array, each SRP (Storage Resource Pool), and each storage group.

CloudIQ provides the collector with configuration data that defines the data items to collect and their collection frequency. CloudIQ sends this configuration data once a day (at most). As CloudIQ gets new features, or enhancements to existing features, the data it requires changes accordingly. It communicates this to the data collector in each registered array in the form of revised configuration data.

Monitor facilities

CloudIQ has a comprehensive set of facilities for monitoring a storage array:

- A summary page gives an overview of the health of all the arrays.
- The systems page gives a summary of the state of each individual array.
- The details gives information about an individual array, its configuration, storage capacity, performance, and health.
- The health center provides details of the alerts that individual arrays have raised.
- The hosts page lists host systems connected to the monitored arrays.

The health score can help you see where the most severe health issues are, based on five core factors, shown in the following table.

Table 6. Health score categories

| Category | Description |
|-----------------|--|
| Components | Overall health of the physical components of a system. |
| Configuration | System configuration, such as the high availability status of the hosts attached to this system. |
| Capacity | System capacity, for example, whether there is imminent risk of running out of capacity. |
| Performance | System performance, for example, whether the system could better balance resource usage. |
| Data Protection | System data protection, for example, whether protection policies are in compliance. |

The differentiator for CloudIQ, however, is the use of predictive analytics. CloudIQ analyzes the data it has received from each array to determine the normal range of values for various metrics. Using this it can highlight when the metric goes outside of this normal range.

Support services

SRS provides more facilities than simply sending data from an array to CloudIQ:

- An array can automatically open service requests for critical issues that arise.
- Dell EMC support staff can access the array to troubleshoot critical issues and to obtain diagnostic information such as log and dump files.

Security

Each customer with access to CloudIQ has a dedicated access portal through which they can view their own arrays only. A customer does not have access to any other customer's arrays or data. In addition, SRS uses point-to-point encryption over a dedicated VPN, multi-factor authentication, customer-controlled access policies, and RSA digital certificates to ensure that all customer data is securely transported to Dell EMC.

The infrastructure that CloudIQ uses is regularly scanned for vulnerabilities with remediation taking place as a result of these scans. This helps to maintain the security and privacy of all customer data.

CyberSecIQ

CyberSecIQ is an as a service cloud-based storage security analytics application that provides security assessment and measures the overall cyber security risk level of storage systems using intelligent, comprehensive, and predictive analytics.

CyberSecIQ uses Secure Remote Services to collect system logs, system configurations, security configurations and settings, alerts, and performance metrics from the Unisphere system.

Prerequisites for the application include:

- The Secure Remote Services gateway has already been registered in Unisphere.
- The Secure Remote Services gateway must be directly connected to Unisphere.

- Sending data to CloudIQ setting must be enabled.
- There must be at least one local array in Unisphere.

Solutions Enabler

Solutions Enabler provides a comprehensive command line interface (SYMCLI) to manage your storage environment.

SYMCLI commands are invoked from a management host, either interactively on the command line, or using scripts.

SYMCLI is built on functions that use system calls to generate low-level I/O SCSI commands. Configuration and status information is maintained in a host database file, reducing the number of enquiries from the host to the arrays.

Use SYMCLI to:

- Configure array software (for example, TimeFinder, SRDF, Open Replicator)
- Monitor device configuration and status
- Perform control operations on devices and data objects

Solutions Enabler also has a Representational State Transfer (REST) API. Use this API to access performance and configuration information, and provision storage arrays. It can be used in any programming environment that supports standard REST clients, such as web browsers and programming platforms that can issue HTTP requests.

Mainframe Enablers

The Dell EMC Mainframe Enablers are software components that allow you to monitor and manage arrays running PowerMaxOS in a mainframe environment:

- ResourcePak Base for z/OS
Enables communication between mainframe-based applications (provided by Dell EMC or independent software vendors) and PowerMax/VMAX arrays.
- SRDF Host Component for z/OS
Monitors and controls SRDF processes through commands executed from a host. SRDF maintains a real-time copy of data at the logical volume level in multiple arrays located in physically separate sites.
- Dell EMC Consistency Groups for z/OS
Ensures the consistency of data remotely copied by SRDF feature in the event of a rolling disaster.
- AutoSwap for z/OS
Handles automatic workload swaps between arrays when an unplanned outage or problem is detected.
- TimeFinder SnapVX
With Mainframe Enablers V8.0 and higher, SnapVX creates point-in-time copies directly in the Storage Resource Pool (SRP) of the source device, eliminating the concepts of target devices and source/target pairing. SnapVX point-in-time copies are accessible to the host through a link mechanism that presents the copy on another device. TimeFinder SnapVX and PowerMaxOS support backward compatibility to traditional TimeFinder products, including TimeFinder/Clone, TimeFinder VP Snap, and TimeFinder/Mirror.
- Data Protector for z Systems (zDP™)
With Mainframe Enablers V8.0 and higher, zDP is deployed on top of SnapVX. zDP provides a granular level of application recovery from unintended changes to data. zDP achieves this by providing automated, consistent point-in-time copies of data from which an application-level recovery can be conducted.
- TimeFinder/Clone Mainframe Snap Facility
Produces point-in-time copies of full volumes or of individual datasets. TimeFinder/Clone operations involve full volumes or datasets where the amount of data at the source is the same as the amount of data at the target. TimeFinder VP Snap leverages clone technology to create space-efficient snaps for thin devices.
- TimeFinder/Mirror for z/OS
Allows the creation of Business Continuance Volumes (BCVs) and provides the ability to ESTABLISH, SPLIT, RE-ESTABLISH and RESTORE from the source logical volumes.
- TimeFinder Utility

Conditions SPLIT BCVs by relabeling volumes and (optionally) renaming and recataloging datasets. This allows BCVs to be mounted and used.

Geographically Dispersed Disaster Restart (GDDR)

GDDR automates business recovery following both planned outages and disaster situations, including the total loss of a data center. Using the PowerMax architecture and the foundation of SRDF and TimeFinder replication families, GDDR eliminates any single point of failure for disaster restart plans in mainframe environments. GDDR intelligence automatically adjusts disaster restart plans based on triggered events.

GDDR does not provide replication and recovery services itself. Rather GDDR monitors and automates the services that other Dell EMC products and third-party products provide that are required for continuous operations or business restart. GDDR facilitates business continuity by generating scripts that can be run on demand. For example, scripts to restart business applications following a major data center incident, or resume replication following unplanned link outages.

Scripts are customized when invoked by an expert system that tailors the steps based on the configuration and the event that GDDR is managing. Through automatic event detection and end-to-end automation of managed technologies, GDDR removes human error from the recovery process and allows it to complete in the shortest time possible.

The GDDR expert system is also invoked to automatically generate planned procedures, such as moving compute operations from one data center to another. This is the gold standard for high availability compute operations, to be able to move from scheduled DR test weekend activities to regularly scheduled data center swaps without disrupting application workloads.

SMI-S Provider

Dell EMC SMI-S Provider supports the SNIA Storage Management Initiative (SMI), an ANSI standard for storage management. This initiative has developed a standard management interface that resulted in a comprehensive specification (SMI-Specification or SMI-S).

SMI-S defines the open storage management interface, to enable the interoperability of storage management technologies from multiple vendors. These technologies are used to monitor and control storage resources in multivendor or SAN topologies.

Solutions Enabler components required for SMI-S Provider operations are included as part of the SMI-S Provider installation.

VASA Provider

The VASA Provider enables PowerMax management software to inform vCenter of how VMDK storage, including vVols, is configured and protected. These capabilities are defined by Dell EMC and include characteristics such as disk type, type of provisioning, storage tiering and remote replication status. This allows vSphere administrators to make quick and informed decisions about virtual machine placement. VASA offers the ability for vSphere administrators to complement their use of plugins and other tools to track how devices hosting vVols are configured to meet performance and availability needs. Details about VASA Provider replication groups can be viewed on the Unisphere vVols dashboard.

eNAS management interface

You manage eNAS block and file storage using the Unisphere File Dashboard. Link and launch enables you to run the block and file management GUI within the same session.

The configuration wizard helps you create storage groups (automatically provisioned to the Data Movers) quickly and easily. Creating a storage group creates a storage pool in Unisphere that can be used for file level provisioning tasks.

Storage Resource Management (SRM)

SRM provides comprehensive monitoring, reporting, and analysis for heterogeneous block, file, and virtualized storage environments.

Use SRM to:

- Visualize applications to storage dependencies

- Monitor and analyze configurations and capacity growth
- Optimize your environment to improve return on investment

Virtualization enables businesses to simplify management, control costs, and guarantee uptime. However, virtualized environments also add layers of complexity to the IT infrastructure that reduce visibility and can complicate the management of storage resources. SRM addresses these layers by providing visibility into the physical and virtual relationships to ensure consistent service levels.

As you build out a cloud infrastructure, SRM helps you ensure storage service levels while optimizing IT resources — both key attributes of a successful cloud deployment.

SRM is designed for use in heterogeneous environments containing multi-vendor networks, hosts, and storage devices. The information it collects and the functionality it manages can reside on technologically disparate devices in geographically diverse locations. SRM moves a step beyond storage management and provides a platform for cross-domain correlation of device information and resource topology, and enables a broader view of your storage environment and enterprise data center.

SRM provides a dashboard view of the storage capacity at an enterprise level through Watch4net. The Watch4net dashboard view displays information to support decisions regarding storage capacity.

The Watch4net dashboard consolidates data from multiple ProSphere instances spread across multiple locations. It gives a quick overview of the overall capacity status in the environment, raw capacity usage, usable capacity, used capacity by purpose, usable capacity by pools, and service levels.

vStorage APIs for Array Integration

VMware vStorage APIs for Array Integration (VAAI) optimize server performance by offloading virtual machine operations to arrays running PowerMaxOS.

The storage array performs the select storage tasks, freeing host resources for application processing and other tasks.

In VMware environments, storage arrays supports the following VAAI components:

- Full Copy — (Hardware Accelerated Copy) Faster virtual machine deployments, clones, snapshots, and VMware Storage vMotion® operations by offloading replication to the storage array.
- Block Zero — (Hardware Accelerated Zeroing) Initializes file system block and virtual drive space more rapidly.
- Hardware-Assisted Locking — (Atomic Test and Set) Enables more efficient meta data updates and assists virtual desktop deployments.
- UNMAP — Enables more efficient space usage for virtual machines by reclaiming space on datastores that is unused and returns it to the thin provisioning pool from which it was originally drawn.
- VMware vSphere Storage APIs for Storage Awareness (VASA).

VAAI is native in PowerMaxOS and does not require additional software, unless eNAS is also implemented. If eNAS is implemented on the array, support for VAAI requires the VAAI plug-in for NAS. The plug-in is available from the Dell EMC support website.

SRDF Adapter for VMware vCenter Site Recovery Manager

Dell EMC SRDF Adapter is a Storage Replication Adapter (SRA) that extends the disaster restart management functionality of VMware vCenter Site Recovery Manager 5.x to arrays running PowerMaxOS.

SRA allows Site Recovery Manager to automate storage-based disaster restart operations on storage arrays in an SRDF configuration.

SRDF/Cluster Enabler

Cluster Enabler (CE) for Microsoft Failover Clusters is a software extension of failover clusters functionality. Cluster Enabler enables Windows Server 2012 (including R2) Standard and Datacenter editions running Microsoft Failover Clusters to operate across multiple connected storage arrays in geographically distributed clusters.

SRDF/Cluster Enabler (SRDF/CE) is a software plug-in module to Dell EMC Cluster Enabler for Microsoft Failover Clusters software. The Cluster Enabler plug-in architecture consists of a CE base module component and separately available plug-in modules, which provide your chosen storage replication technology.

SRDF/CE supports:

- Synchronous and asynchronous mode ([SRDF modes of operation](#) on page 81 summarizes these modes)
- Concurrent and cascaded SRDF configurations ([SRDF multi-site solutions](#) on page 76 summarizes these configurations)

Product Suite for z/TPF

The Dell EMC Product Suite for z/TPF is a suite of components that monitor and manage arrays running PowerMaxOS from a z/TPF host. z/TPF is an IBM mainframe operating system characterized by high-volume transaction rates with significant communications content. The following software components are distributed separately and can be installed individually or in any combination:

- SRDF Controls for z/TPF

Monitors and controls SRDF processes with functional entries entered at the z/TPF Prime CRAS (computer room agent set).

- TimeFinder Controls for z/TPF

Provides a business continuance solution consisting of TimeFinder SnapVX, TimeFinder/Clone, and TimeFinder/Mirror.

- ResourcePak for z/TPF

Provides PowerMax and VMAX configuration and statistical reporting and extended features for SRDF Controls for z/TPF and TimeFinder Controls for z/TPF.

SRDF/TimeFinder Manager for IBM i

Dell EMC SRDF/TimeFinder Manager for IBM i is a set of host-based utilities that provides an IBM i interface to SRDF and TimeFinder.

This feature allows you to configure and control SRDF or TimeFinder operations on arrays attached to IBM i hosts, including:

- SRDF: Configure, establish, and split SRDF devices, including:
 - SRDF/A
 - SRDF/S
 - Concurrent SRDF/A
 - Concurrent SRDF/S
- TimeFinder:
 - Create point-in-time copies of full volumes or individual data sets.
 - Create point-in-time snapshots of images.

Extended features

SRDF/TimeFinder Manager for IBM i extended features provide support for the IBM independent ASP (IASP) functionality.

IASPs are sets of switchable or private auxiliary disk pools (up to 223) that can be brought online/offline on an IBM i host without affecting the rest of the system.

When combined with SRDF/TimeFinder Manager for IBM i, IASPs let you control SRDF or TimeFinder operations on arrays attached to IBM i hosts, including:

- Display and assign TimeFinder SnapVX devices.
- Execute SRDF or TimeFinder commands to establish and split SRDF or TimeFinder devices.
- Present one or more target devices containing an IASP image to another host for business continuance (BC) processes.

Access to extended features control operations include:

- From the SRDF/TimeFinder Manager menu-driven interface.
- From the command line using SRDF/TimeFinder Manager commands and associated IBM i commands.

AppSync

Dell EMC AppSync offers a simple, SLA-driven, self-service approach for protecting, restoring, and cloning critical Microsoft and Oracle applications and VMware environments. After defining service plans, application owners can protect, restore, and clone production data quickly with item-level granularity by using the underlying Dell EMC replication technologies. AppSync also provides an application protection monitoring service that generates alerts when the SLAs are not met.

AppSync supports the following applications and storage arrays:

- Applications—Oracle, Microsoft SQL Server, Microsoft Exchange, and VMware vStorage VMFS and NFS datastores and File systems.
- Replication Technologies—SRDF, SnapVX, RecoverPoint, XtremIO Snapshot, VNX Advanced Snapshots, VNXe Unified Snapshot, and ViPR Snapshot.

On PowerMax arrays:

- The Essentials software package contains AppSync in a starter bundle. The AppSync Starter Bundle provides the license for a scale-limited, yet fully functional version of AppSync. For more information, see the *AppSync Starter Bundle with PowerMax Product Brief* available on the Dell EMC Online Support Website.
- The Pro software package contains the AppSync Full Suite.

EMC Storage Analytics (ESA)

ESA links VMware vRealize Operations manager for storage with EMC Adapter. The vRealize Operations Manager shows performance and capacity metrics from storage systems with data that the adapter provides by:

- Connecting to and collecting data from storage system resources
- Converting the data into a format that vRealize Operations Manager can process
- Passing the data to the vRealize Operations Manager collector

vRealize Operations Manager presents the aggregated data through alerts and dashboards, and in predefined reports that end users can interpret easily. EMC Adapter is installed with the vRealize Operations Manager administrative user interface.

ESA complies with VMware management pack certification requirements and has received VMware Ready certification. It is a free download for PowerMax customers.

Open Systems Features

This chapter introduces the open systems features of PowerMax arrays.

Topics:

- [PowerMaxOS support for open systems](#)
- [PowerPath](#)
- [Backup and restore using PowerProtect Storage Direct and Data Domain](#)
- [VMware Virtual Volumes](#)

PowerMaxOS support for open systems

PowerMaxOS provides FBA device emulations for open systems and D910 for IBM i.

Any logical device manager software installed on a host can be used with the storage devices.

PowerMaxOS increases scalability limits from previous generations of arrays, including:

- Maximum device size is 64 TB
- Maximum host addressable devices is 64,000 for each array
- Maximum storage groups, port groups, initiator groups, and masking views is 16,000 for each object type for each array
- Maximum devices addressable through each port is 4,000

[Open Systems-specific provisioning](#) on page 58 has more information on provisioning storage in an open systems environment.

The Dell EMC Support Matrix in the E-Lab Interoperability Navigator at <http://elabnavigator.emc.com> has the most recent information on PowerMaxOS open systems capabilities.

PowerPath

PowerPath runs on an application host and manages data paths between the host and LUNs on a storage array. PowerPath is available for various operating systems including AIX, Microsoft Windows, Linux, and VMware.

This section is high-level summary of the PowerPath capabilities for PowerMax arrays. It also shows where to get detailed information including instructions on how to install, configure, and manage PowerPath.

Operational overview

A data path is a physical connection between an application host and a LUN on a PowerMax array. The path has several components including:

- Host-based adapter (HBA) port
- Cables
- Switches, PowerMax port
- The LUN

PowerPath manages the use of the paths between a host and a LUN to optimize their use and to take corrective action should an error occurs.

There can be multiple paths to a LUN enabling PowerPath to:

- Balance the I/O load across the available paths. In turn, this:
 - Optimizes the use of the paths
 - Improves overall I/O performance
 - Reduces management intervention
 - Eliminates the need to configure paths manually

- Automatic failover should a path become unavailable due to the failure of one or more of its components. That is, if a path becomes unavailable, PowerPath reroutes I/O traffic to alternative paths without manual intervention.

Host registration

Each host that uses PowerPath to access an array registers itself with the array. The information that PowerPath sends to the array is:

- Host name
- Operating system and version
- Hardware
- PowerPath version
- Name of the cluster the host is part of and the host's cluster name (if applicable)
- WWN of the host
- Name of each VM on the host and the operating system that each runs

The array stores this information in memory.

PowerPath repeats the registration process every 24 hours. In addition, it checks the host information at hourly intervals. If the name or IP address of the host have changed, PowerPath repeats the registration process with the array immediately.

Rather than wait for the next registration check to occur, a system administrator can register a change immediately using the PowerPath CLI. If necessary a site can control whether automatic registration occurs both for an individual host and for an entire array.

In addition, the array deletes information on any host that has not registered over the last 72 hours. This prevents a build up of out-of-date host data.

Device status

In addition to host information, PowerPath sends device information to the array. The device information includes:

- Date of last usage
- Mount status
- Name of the process that owns the device
- PowerPath I/O statistics (these are in addition to the I/O statistics that the array itself gathers)

The array stores this information in memory.

Benefits of the device information include:

- Early identification of potential I/O problems
- Better long-term planning of array and host usage
- Recover and redeploy unused storage assets

Automatic creation of Initiator Groups

Solutions Enabler includes the ability to automatically create an Initiator Group (IG) from a host name. This feature is introduced in PowerMaxOS 5978.144.144. When this feature is switched on, the storage administrator can use a modified form of the `symaccess` command that simplifies the creation of an IG.

From PowerMaxOS 5978.144.144 `symaccess` has an additional `-host` qualifier that takes a host name as its value. On issuing this command, Solutions Enabler searches the host information received from PowerPath to find all the WWNs associated with that host. From the results of this search Solutions Enabler creates a IG with no further intervention by the storage administrator.

Management

Solutions Enabler and Unisphere have facilities to:

- View host information
- View device information

- View PowerPath performance data
- Register PowerPath hosts with an array
- Control automatic registration of host systems

More information

There is more information about PowerPath, how to configure it, and manage it in:

- *PowerPath Installation and Administration Guide*
- *PowerPath Release Notes*
- *PowerPath Family Product Guide*
- *PowerPath Family CLI and System Messages Reference*
- *PowerPath Management Appliance Installation and Configuration Guide*
- *PowerPath Management Appliance Release Notes*

There are *Installation and Administration Guide* and *Release Notes* documents for each supported operating system.

Backup and restore using PowerProtect Storage Direct and Data Domain

Dell EMC Storage Direct provides data backup and restore facilities for a PowerMax array. A remote Data Domain array stores the backup copies of the data.

Storage Direct uses existing features of the PowerMax and Data Domain arrays to create backup copies and to restore backed up data if necessary. There is no need for any specialized or additional hardware and software.

This section is a high-level summary of Storage Direct backup and restore facilities. It also shows where to get detailed information about the product, including instructions on how to configure and manage it.

Backup

A LUN is the basic unit of backup in Storage Direct. For each LUN, Storage Direct creates a *backup image* on the Data Domain array. You can group backup images to create a *backup set*. One use of the backup set is to capture all the data for an application as a point-in-time image.

Backup process

To create a backup of a LUN, Storage Direct:

1. Uses SnapVX to create a local snapshot of the LUN on the PowerMax array (the primary storage array).
After the snapshot is created, Storage Direct and the application can proceed independently each other and the backup process has no further impact on the application.
2. Copies the snapshot to a vdisk on the Data Domain array where it is deduplicated and cataloged.
On the primary storage array, the vdisk appears as a FAST.X encapsulated LUN. The copy of the snapshot to the vdisk uses existing SnapVX link copy and PowerMax destaging technologies.

When the vdisk contains all the data for the LUN, Data Domain converts the data into a *static image*. This image then has metadata added to it and Data Domain catalogs the resultant backup image.

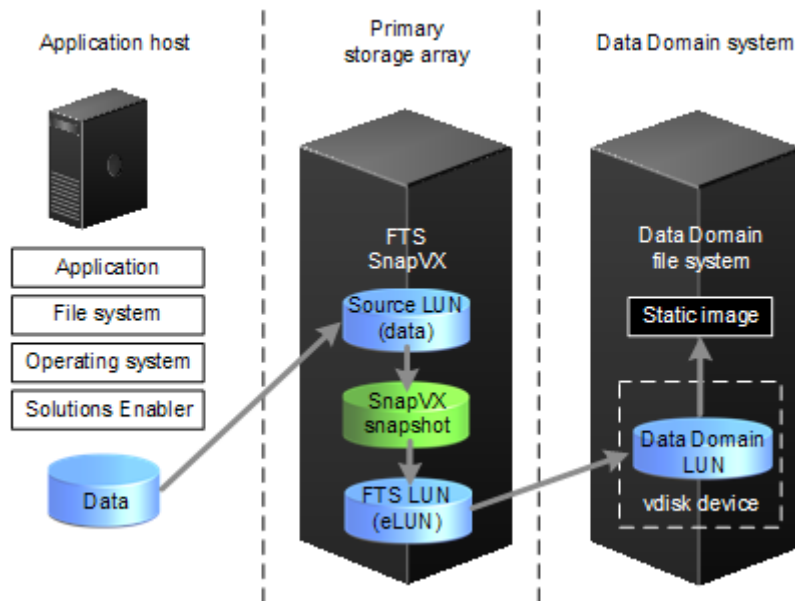


Figure 4. Data flow during a backup operation to Data Domain

Incremental data copy

The first time that Storage Direct backs up a LUN, it takes a complete copy of its contents using a SnapVX snapshot. While taking this snapshot, the application assigned to the LUN is paused for a short time. This ensures that Storage Direct has a copy of the LUN that is application consistent. To create the first backup image of the LUN, Storage Direct copies the entire snapshot to the Data Domain array.

For each subsequent backup of the LUN, Storage Direct copies only those parts of the LUN that have changed. This makes best use of the communication links and minimizes the time that is required to create the backup.

Restore

Storage Direct provides two forms of data restore:

- Object level restore from a selected backup image
- Full application rollback restore

Object level restore

For an object level restore, Data Domain puts the static image from the selected backup image on a vdisk. As with the backup process, this vdisk on the Data Domain array appears as a FAST.X encapsulated LUN on the PowerMax array. The administrator can now mount the file system of the encapsulated LUN, and restore one or more objects to their final destination.

Full application rollback restore

In a full application rollback restore, all the static images in a selected backup set are made available as vdisks on the Data Domain array and available as FAST.X encapsulated LUNs on the PowerMax array. From there, the administrator can restore data from the encapsulated LUNs to their original devices.

Storage Direct agents

Storage Direct has three agents, each responsible for backing up and restoring a specific type of data:

File system agent Provides facilities to back up, manage, and restore application LUNs.

| | |
|------------------------------------|--|
| Database application agent | Provides facilities to back up, manage, and restore DB2 databases, Oracle databases, or SAP with Oracle database data. |
| Microsoft application agent | Provides facilities to back up, manage, and restore Microsoft Exchange and Microsoft SQL Server databases. |

Features used for Storage Direct backup and restore

Storage Direct uses existing features of PowerMaxOS and Data Domain to provide backup and restore services:

- PowerMaxOS:
 - SnapVX
 - FAST.X encapsulated devices
- Data Domain:
 - Block services for Storage Direct
 - vdisk services
 - FastCopy

Storage Direct and traditional backup

The Storage Direct workflow can provide data protection in situations where more traditional approaches cannot successfully meet the business requirements. This is often due to small or nonexistent backup windows, demanding recovery time objective (RTO) or recovery point objective (RPO) requirements, or a combination of both.

Unlike traditional backup and recovery, Storage Direct does not rely on a separate process to identify the data that must be backed up and additional actions to move that data to backup storage. Instead of using dedicated hardware, host, and network resources, Storage Direct uses existing application and storage capabilities to create point-in-time copies of large datasets. The copies are transported across a storage area network (SAN) to Data Domain systems to protect the copies while providing deduplication to maximize storage efficiency.

Storage Direct minimizes the time that is required to protect large datasets, and allows backups to fit into the smallest of backup windows to meet demanding RTO or RPO requirements.

More information

More information about Storage Direct, its components, how to configure them, and how to use them is available in:

- *PowerProtect Storage Direct Solutions Guide*
- *File System Agent Installation and Administration Guide*
- *Database Application Agent Installation and Administration Guide*
- *Microsoft Application Agent Installation and Administration Guide*

VMware Virtual Volumes

VMware Virtual Volumes (vVols) are a storage object developed by VMware to simplify management and provisioning in virtualized environments.

With vVols, the management process moves from the LUN (data store) to the virtual machine (VM). This level of detail allows VMware and cloud administrators to assign specific storage attributes to each VM, according to its performance and storage requirements. Storage arrays running PowerMaxOS implement vVols.

PowerMax 5978.669.669, and later, can remotely replicate vVols for disaster recovery purposes using SRDF/Asynchronous (SRDF/A).

For more information about vVols, see *Dell EMC SRDF Introduction* and *Dell EMC VASA Provider and Embedded VASA Provider for PowerMax Product Guide*.

vVol components

To support management capabilities of vVols, the storage/vCenter environment requires the following:

- Dell EMC PowerMax VASA Provider – The VASA Provider (VP) is a software plug-in that uses a set of out-of-band management APIs (VASA version 2.0). The VASA Provider exports storage array capabilities and presents them to vSphere through the VASA APIs. vVols are managed by way of vSphere through the VASA Provider APIs (create/delete) and not with the Unisphere for PowerMax user interface or Solutions Enabler CLI. After vVols are setup on the array, Unisphere and Solutions Enabler only support vVol monitoring and reporting.
- Storage Containers (SC)—Storage containers are chunks of physical storage used to logically group vVols. SCs are based on the grouping of Virtual Machine Disks (VMDKs) into specific Service Levels. SC capacity is limited only by hardware capacity. At least one SC per storage system is required, but multiple SCs per array are allowed. SCs are created and managed on the array by the Storage Administrator. Unisphere and Solutions Enabler CLI support management of SCs.
- Protocol Endpoints (PE)—Protocol endpoints are the access points from the hosts to the array by the Storage Administrator. PEs are compliant with FC and replace the use of LUNs and mount points. vVols are "bound" to a PE, and the bind and unbind operations are managed through the VP APIs, not with the Solutions Enabler CLI. Existing multi-path policies and NFS topology requirements can be applied to the PE. PEs are created and managed on the array by the Storage Administrator. Unisphere and Solutions Enabler CLI support management of PEs.

Table 7. vVol architecture component management capability

| Functionality | Component |
|---|---|
| vVol device management (create, delete) | VASA Provider APIs / Solutions Enabler APIs |
| vVol bind management (bind, unbind) | |
| Protocol Endpoint device management (create, delete) | Unisphere/Solutions Enabler CLI |
| Protocol Endpoint-vVol reporting (list, show) | |
| Storage Container management (create, delete, modify) | |
| Storage container reporting (list, show) | |

vVol scalability

The vVol scalability limits are:

Table 8. vVol-specific scalability

| Requirement | Value |
|--|-------------|
| Number of vVols/Array | 64,000 |
| Number of Snapshots/Virtual Machine ^a | 12 |
| Number of Storage Containers/Array | 16 |
| Number of Protocol Endpoints/Array | 1/ESXi Host |
| Maximum number of Protocol Endpoints/Array | 1,024 |
| Number of arrays supported /VP | 1 |
| Number of vCenters/VP | 2 |
| Maximum device size | 16 TB |

a. vVol Snapshots are managed through vSphere only. You cannot use Unisphere or Solutions Enabler to create them.

vVol workflow

Requirements

Install and configure these applications:

- Unisphere for PowerMax V9.0 or later
- Solutions Enabler CLI V9.0 or later
- VASA Provider V9.0 or later

Instructions for installing Unisphere and Solutions Enabler are in their respective installation guides. Instructions on installing the VASA Provider are in the *Dell EMC PowerMax VASA Provider Release Notes* .

Procedure

The creation of a vVol-based virtual machine involves both the storage administrator and the VMware administrator:

Storage administrator

The storage administrator uses Unisphere or Solutions Enabler to create the storage and present it to the VMware environment:

1. Create one or more storage containers on the storage array.
This step defines how much storage and from which service level the VMware user can provision.
2. Create Protocol Endpoints and provision them to the ESXi hosts.

VMware administrator

The VMware administrator uses the vSphere Web Client to deploy the VM on the storage array:

1. Add the VASA Provider to the vCenter.
This allows vCenter to communicate with the storage array,
2. Create a vVol datastore from the storage container.
3. Create the VM storage policies.
4. Create the VM in the vVol datastore, selecting one of the VM storage policies.

Mainframe Features

This chapter introduces the mainframe-specific features of PowerMax arrays.

Topics:

- [PowerMaxOS support for mainframe](#)
- [IBM Z Systems functionality support](#)
- [IBM 2107 support](#)
- [Logical control unit capabilities](#)
- [Disk drive emulations](#)
- [Cascading configurations](#)

PowerMaxOS support for mainframe

PowerMax 8000 arrays can be ordered with the zEssentials and zPro software packages to provide mainframe capabilities.

PowerMax arrays provide the following mainframe features:

- Mixed FBA and CKD drive configurations.
- Support for 64, 128, 256 FICON single and multi mode ports, respectively.
- Support for CKD 3380/3390 and FBA devices.
- Mainframe (FICON) and OS FC/iSCSI connectivity.
- High capacity flash drives.
- Up to 16 Gb/s FICON host connectivity.
- Support for Forward Error Correction, Query Host Access, and FICON Dynamic Routing.
- T10 DIF protection for CKD data along the data path (in cache and on disk) to improve performance for multi-record operations.
- D@RE external key managers. [Data at Rest Encryption](#) on page 27 provides more information on D@RE and external key managers.

IBM Z Systems functionality support

PowerMax arrays support the latest IBM Z Systems enhancements, ensuring that the array can handle the most demanding mainframe environments:

- zHPF, including support for single track, multi track, List Prefetch, bi-directional transfers, QSAM/BSAM access, and Format Writes
- zHyperWrite
- Non-Disruptive State Save (NDSS)
- Compatible Native Flash (Flash Copy)
- Multiple Incremental Flash Copy (up to 12 incremental flash copy target relationships to one source device)
- Remote Pair Flash Copy (RPFC)
- Concurrent Copy
- Multi-subsystem Imaging
- Parallel Access Volumes (PAV)
- Dynamic Channel Management (DCM)
- Dynamic Parallel Access Volumes/Multiple Allegiance (PAV/MA)
- Peer-to-Peer Remote Copy (PPRC) SoftFence
- PPRC event aggregation (PPRCSUM)
- Support for z/OS Storage Controller Health Check system messages
- Extended Address Volumes (EAV)

- Dynamic volume expansion for 3390 TDEVs, including devices that are part of an SRDF (except SRDF/Metro), SnapVX, Concurrent Copy, or SDDF configuration.
- Persistent I/O Pacing (Extended Distance FICON)
- HyperPAV
- SuperPAV
- PDS Search Assist
- Modified Indirect Data Address Word (MIDAW)
- Multiple Allegiance (MA)
- Sequential Data Striping
- Multi-Path Lock Facility
- Product Suite for z/TPF
- HyperSwap
- Secure Snapsets in SnapVX for zDP
- Global Mirror
- Transparent Cloud Tiering

Global Mirror support

Global Mirror is an IBM solution for long-distance replication that can be run through Geographically Dispersed Parallel Sysplex (GDPS) automation.

PowerMax (and VMAX) supports a two-site Global Mirror environment. A two-site Global Mirror configuration has Peer-to-Peer Remote Copy (PPRC) primary devices at Site A and PPRC secondary devices at Site B. There is also a relationship from the PPRC secondary devices to a third set of devices at Site B that is not usually active.

NOTE: Use of Global Mirror requires an RPQ.

Global Mirror requirements and restrictions are as follows:

- Global Mirror is supported on PowerMax 8000 and VMAX 950F arrays.
- Both PPRC primary and secondary sites run PowerMaxOS 5978.669.669, or later.
- One Global Mirror session is associated with one SRDF group.
- Up to 31 Global Mirror sessions can be run simultaneously.
- When multiple Global Mirror sessions are running on an array, each session maps to a unique SRDF group.
- All devices in the SRDF group are defined to the Global Mirror session.

NOTE: To add or remove devices from the Global Mirror session, stop or pause Global Mirror.

NOTE: A PowerMax array can participate in a z/OS Global Mirror (XRC) configuration only as a secondary.

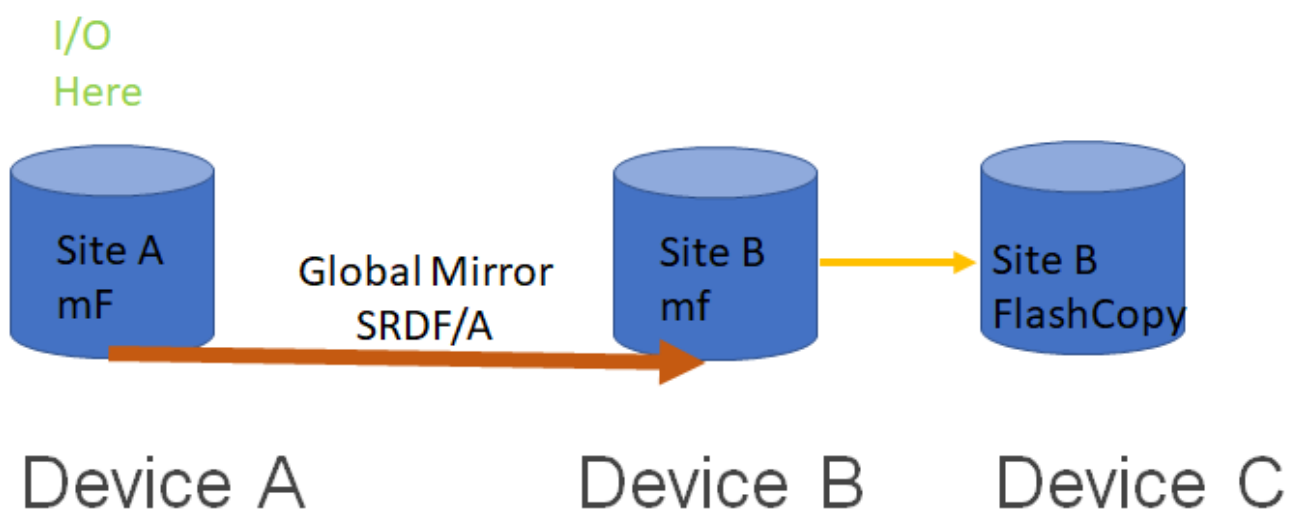


Figure 5. Two-site Global Mirror

Transparent Cloud Tiering support


PowerMax (or VMAX) can be used in an IBM Transparent Cloud Tiering (TCT) environment to store DASD files in, or retrieve them from, the cloud with minimal CPU requirements on the z/OS host.

PowerMax TCT support uses the Dell EMC Disk Library for mainframe (DLm) as its 'cloud'. Count Key Data (CKD) extents are stored as files on standard 3590 tape volumes. The DLm 3590 tape volumes and the DLm tape drives for TCT are separate from any z/OS-defined tape volumes and tape drives. TCT 3590 tapes are not accessible to, or managed, by z/OS. Tape volumes are written to the DLm from the PowerMax over a FICON connection. DLm then stores the data on any backend storage that DLm supports. Optionally, the DLm Long-Term Retention feature can then be used, independent of TCT, to move the data to a Dell EMC Elastic Cloud Storage (ECS) solution.

A cloud object store is required for TCT support to operate. This cloud must support the OpenStack SWIFT protocol and is used to store cloud metadata. If ECS is used as the cloud, the ECS can be the same or a different ECS from any ECS deployed with the DLm.

A REST API proxy server is required in each z/OS image accessing a TCT-enabled PowerMax. This proxy server runs as a separate address space in z/OS.

The *ResourcePak Base for z/OS Product Guide* provides additional information about TCT support, including TCT support requirements and restrictions. It also discusses how to set up and run the Dell EMC REST API proxy and the Dell EMC REST API utility.

 **NOTE:** Use of TCT in PowerMaxOS 5978 requires an RPQ.

PowerMax TCT environment

The PowerMax TCT environment includes the following components:

- z/OS host with DFSMS and Dell EMC REST API proxy:
 - The Data Facility Storage Management Subsystem (DFSMS) provides space management and data movement tools that are primary users of TCT services.
 - Dell EMC REST API proxy provides REST API services for DFSMS.

The REST API proxy requires a connection to a cloud service. The cloud service can be the same ECS cloud that DLm uses or another supported service (SWIFT protocol is required).

- PowerMax with TCT support—A PowerMax mainframe host adapter runs the Cloud Services Stack. The Cloud Services Stack writes data to and reads data from DLm virtual tape volumes.
- Disk Library for mainframe (DLm)—The TCT-enabled PowerMax stores data in a DLm, which acts as a cloud for PowerMax. DLm can be connected to a cloud provider, such as ECS.

The following figure illustrates a TCT environment with PowerMax and DLm.

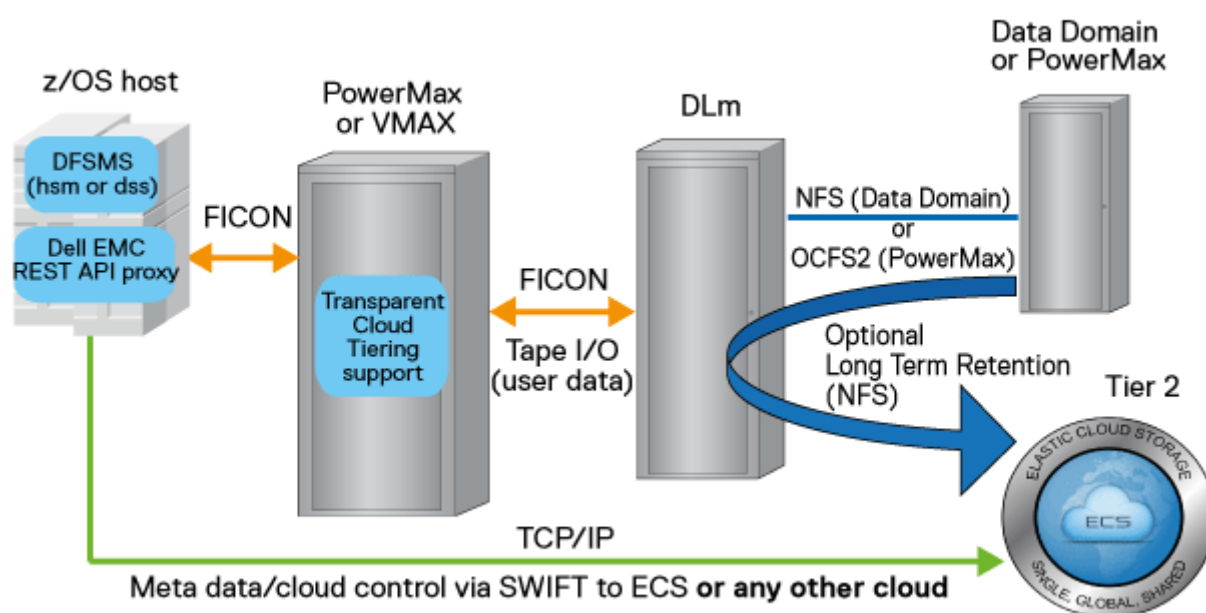


Figure 6. TCT environment with PowerMax and DLM

IBM 2107 support

When PowerMax arrays emulate an IBM 2107, they externally represent the array serial number as an alphanumeric number in order to be compatible with IBM command output. Internally, the arrays retain a numeric serial number for IBM 2107 emulations. PowerMaxOS handles correlation between the alphanumeric and numeric serial numbers.

Logical control unit capabilities

The following table lists logical control unit (LCU) maximum values:

Table 9. Logical control unit maximum values

| Capability | Maximum value |
|--|------------------------------------|
| LCUs per director slice (or port) | 255 (within the range of 00 to FE) |
| LCUs per split ^a | 255 |
| Splits per array | 16 (0 to 15) |
| Devices per split | 65,280 |
| LCUs per array | 512 |
| Devices per LCU | 256 |
| Logical paths per port | 2,048 |
| Logical paths per LCU per port (see Maximum LPARs per port on page 55) | 128 |
| Array system host address per array (base and alias) | 64K |
| I/O host connections per array engine | 32 |

a. A split is a logical partition of the storage array, identified by unique devices, SSIDs, and host serial number. The maximum storage array host address per array is inclusive of all splits.

The following table lists the maximum LPARs per port based on the number of LCUs with active paths:

Table 10. Maximum LPARs per port

| LCUs with active paths per port | Maximum volumes supported per port | Array maximum LPARs per port |
|---------------------------------|------------------------------------|------------------------------|
| 16 | 4K | 128 |
| 32 | 8K | 64 |
| 64 | 16K | 32 |
| 128 | 32K | 16 |
| 255 | 64K | 8 |

Disk drive emulations

When PowerMax arrays are configured to mainframe hosts, the data recording format is Extended CKD (ECKD). The supported CKD emulations are 3380 and 3390.

Cascading configurations

Cascading configurations greatly enhance FICON connectivity between local and remote sites by using switch-to-switch extensions of the CPU to the FICON network. These cascaded switches communicate over long distances using a small number of high-speed lines called interswitch links (ISLs). A maximum of two switches may be connected together within a path between the CPU and the storage array.

Use of the same switch vendors is required for a cascaded configuration. To support cascading, each switch vendor requires specific models, hardware features, software features, configuration settings, and restrictions. Specific IBM CPU models, operating system release levels, host hardware, and PowerMaxOS levels are also required.

The Dell EMC Support Matrix, available through E-Lab Interoperability Navigator (ELN) at <http://elabnavigator.emc.com> has the most up-to-date information on switch support.

Provisioning


This chapter introduces storage provisioning.

Topics:

- [Thin provisioning](#)
- [Multi-array provisioning](#)

Thin provisioning

PowerMax arrays are configured in the factory with thin provisioning pools ready for use. Thin provisioning improves capacity utilization and simplifies storage management. It also enables storage to be allocated and accessed on demand from a pool of storage that services one or many applications. LUNs can be “grown” over time as space is added to the data pool with no impact to the host or application. Data is widely striped across physical storage (drives) to deliver better performance than standard provisioning.

 **NOTE:** Data devices (TDATs) are provisioned/pre-configured/created while the host addressable storage devices TDEVs are created by either the customer or customer support, depending on the environment.

Thin provisioning increases capacity utilization and simplifies storage management by:

- Enabling more storage to be presented to a host than is physically consumed
- Allocating storage only as needed from a shared thin provisioning pool
- Making data layout easier through automated wide striping
- Reducing the steps required to accommodate growth

Thin provisioning allows you to:

- Create host-addressable thin devices (TDEVs) using Unisphere or Solutions Enabler
- Add the TDEVs to a storage group
- Run application workloads on the storage groups

When hosts write to TDEVs, the physical storage is automatically allocated from the default Storage Resource Pool.

Pre-configuration for thin provisioning

PowerMax arrays are custom-built and pre-configured with array-based software applications, including a factory pre-configuration for thin provisioning that includes:


- *Data devices (TDAT)* — an internal device that provides physical storage used by thin devices.
- *Virtual provisioning pool* — a collection of data devices of identical emulation and protection type, all of which reside on drives of the same technology type and speed. The drives in a data pool are from the same disk group.
- *Disk group*— a collection of physical drives within the array that share the same drive technology and capacity. RAID protection options are configured at the disk group level. Dell Technologies strongly recommends that you use one or more of the RAID data protection schemes for all data devices.

Table 11. RAID options

| RAID | Provides the following | Configuration considerations |
|--------|---|---|
| RAID 1 | The highest level of performance for all mission-critical and business-critical applications. Maintains a duplicate copy of a device on two drives. If a drive in the mirrored pair fails, the array automatically uses the mirrored partner without interruption of data availability. | <ul style="list-style-type: none">• Withstands failure of a single drive within the mirrored pair.• A drive rebuild is a simple copy from the remaining drive to the replaced drive.• The number of required drives is twice the amount required to store data (usable storage capacity of a mirrored system is 50%). |
| RAID 5 | Distributed parity and striped data across all drives in the RAID group. Options include: <ul style="list-style-type: none">• PowerMax 2000 only: RAID 5 (3 + 1) — Consists of four drives with parity and data striped across each device.• RAID-5 (7 + 1) — Consists of eight drives with parity and data striped across each device. | <ul style="list-style-type: none">• RAID-5 (3 + 1) provides 75% data storage capacity. Only available with PowerMax 2000 arrays.• RAID-5 (7 + 1) provides 87.5% data storage capacity.• Withstands failure of a single drive within the RAID-5 group. |
| RAID 6 | Striped drives with double distributed parity (horizontal and diagonal). The highest level of availability options include: <ul style="list-style-type: none">• RAID-6 (6 + 2) — Consists of eight drives with dual parity and data striped across each device. | <ul style="list-style-type: none">• RAID-6 (6 + 2) provides 75% data storage capacity. Only available with VMAX 250F arrays.• Withstands failure of two drives within the RAID-6 group. |

- *Storage Resource Pools* — one (default) Storage Resource Pool is pre-configured on the array. This process is automatic and requires no setup. You cannot modify Storage Resource Pools, but you can list and display their configuration. You can also generate reports detailing the demand storage groups are placing on the Storage Resource Pools.

Thin devices (TDEVs)

 **NOTE:** On PowerMax arrays the thin device is the only device type for front end devices.

Thin devices (TDEVs) have no storage allocated until the first write is issued to the device. Instead, the array allocates only a minimum allotment of physical storage from the pool, and maps that storage to a region of the thin device including the area targeted by the write.

These initial minimum allocations are performed in units called thin device extents. Each extent for a thin device is 1 track (128 KB).

When a read is performed on a device, the data being read is retrieved from the appropriate data device to which the thin device extent is allocated. Reading an area of a thin device that has not been mapped does not trigger allocation operations. Reading an unmapped block returns a block in which each byte is equal to zero.

When more storage is required to service existing or future thin devices, data devices can be added to existing thin storage groups.

Thin device oversubscription

A thin device can be presented for host use *before* mapping all of the reported capacity of the device.

The sum of the reported capacities of the thin devices using a given pool can exceed the available storage capacity of the pool. Thin devices whose capacity exceeds that of their associated pool are "oversubscribed".

Over-subscription allows presenting larger than needed devices to hosts and applications without having the physical drives to fully allocate the space represented by the thin devices.

Internal memory usage

Each TDEV uses an amount of the array's internal memory. In systems prior to PowerMaxOS 5978, the system allocated all of the internal memory required for an entire TDEV when the device was created. In extreme circumstances this could result in the system running out of memory even though there is still plenty of capacity on the back-end devices. This behavior changes from PowerMaxOS 5978 onwards. Now the system allocates only the amount of internal memory necessary for the amount of back-end storage actually consumed. Additional internal memory is allocated to the TDEV as it fills up. This results in more efficient use of the system memory and reduces the chances of a failure due to memory exhaustion.

Open Systems-specific provisioning

PowerMaxOS host I/O limits for open systems

On open systems, you can define host I/O limits and associate a limit with a storage group. The I/O limit definitions contain the operating parameters of the I/O per second or bandwidth limitations.

When an I/O limit is associated with a storage group, the limit is divided equally among all the directors in the masking view that is associated with the storage group. All devices in that storage group share that limit.

When applications are configured, you can associate the limits with storage groups that contain a list of devices. A single storage group can only be associated with one limit, and a device can only be in one storage group that has limits associated.

There can be up to 4096 host I/O limits.

Consider the following when using host I/O limits:

- Cascaded host I/O limits controlling parent and child storage groups limits in a cascaded storage group configuration.
- Offline and failed director redistribution of quota that supports all available quota to be available instead of losing quota allocations from offline and failed directors.
- Dynamic host I/O limits support for dynamic redistribution of steady state unused director quota.

Initiator bandwidth limits

Slow SAN drain is a congestion issue that is caused by slow drain devices (host bus adapters (HBAs) that have a lower link speed (BW) compared to storage array SLIC speeds (BW)). This issue can lead to severe fabric-wide performance degradation, but can be mitigated by configuring per initiator bandwidth limits.

At the host group level you can configure initiator-based host I/O, or bandwidth, limits. To set the limits, there must be fibre channel connectivity to the host, and at least one initiator configured on the host.

Auto-provisioning groups on open systems

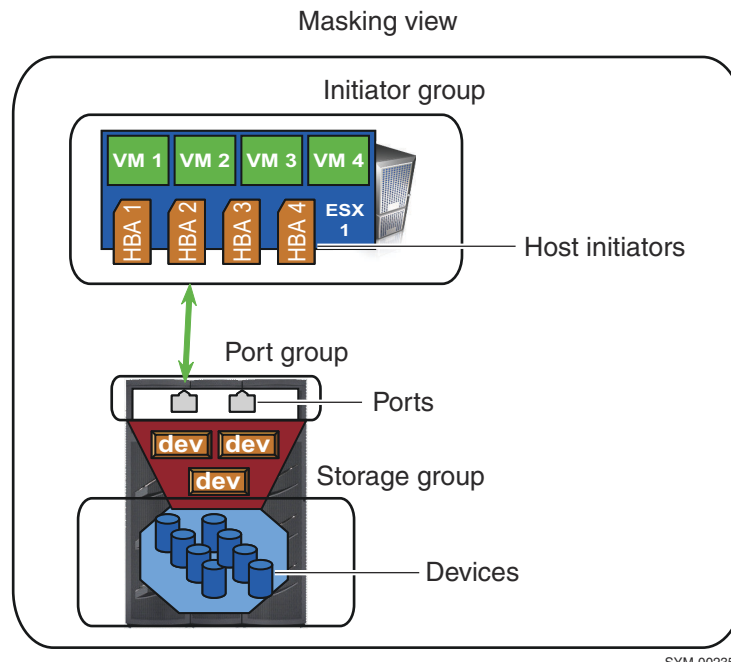
You can auto-provision groups on open systems to reduce complexity, execution time, labor cost, and the risk of error.

Auto-provisioning groups enables users to group initiators, front-end ports, and devices together, and to build masking views that associate the devices with the ports and initiators.

When a masking view is created, the necessary mapping and masking operations are performed automatically to provision storage.

After a masking view exists, any changes to its grouping of initiators, ports, or storage devices automatically propagate throughout the view, automatically updating the mapping and masking as required.

Components of an auto-provisioning group



SYM-002353

Figure 7. Auto-provisioning groups

Initiator group

A logical grouping of Fibre Channel initiators. An initiator group is limited to either a parent, which can contain other groups, or a child, which contains one initiator role. Mixing of initiators and child name in a group is not supported.

Port group

A logical grouping of Fibre Channel front-end director ports. A port group can contain up to 32 ports.

Storage group

A logical grouping of thin devices. LUN addresses are assigned to the devices within the storage group when the view is created if the group is either cascaded or stand alone. Often there is a correlation between a storage group and a host application. One or more storage groups may be assigned to an application to simplify management of the system. Storage groups can also be shared among applications.

Cascaded storage group

A parent storage group comprised of multiple storage groups (parent storage group members) that contain child storage groups comprised of devices. By assigning child storage groups to the parent storage group members and applying the masking view to the parent storage group, the masking view inherits all devices in the corresponding child storage groups.

Masking view

An association between one initiator group, one port group, and one storage group. When a masking view is created, the group within the view is a parent, the contents of the children are used. For example, the initiators from the children initiator groups and the devices from the children storage groups. Depending on the server and application requirements, each server or group of servers may have one or more masking views that associate a set of thin devices to an application, server, or cluster of servers.

Multi-array provisioning

The multi-array Provisioning Storage wizard simplifies the task of identifying the optimal target array and provisioning storage on that array.

Unisphere for PowerMax 9.2 provides a system-level provisioning launch point that takes array-independent inputs (storage group name, device count and size, and (optionally) response time target or initiator filter), selects ports that are based on current utilization and port group best practices, and returns component impact scores for all locally connected arrays running HYPERMAX OS 5977 or PowerMaxOS 5978.

You can also select a provisioning template and provision new storage using the wizard. Storage group capacity information and response time targets that are already part of the provisioning template are populated when the wizard opens. The most suitable ports (based on specified options) are selected and a list of all locally connected arrays (V3 and higher) are returned. The list is sorted by the impact of the new workload on the target arrays.

Host I/O limits (quotas) can be used to limit the amount of Front End (FE) Bandwidth and I/O operations per second (IOPS) that can be consumed by a set of storage volumes over a set of director ports. Host I/O limits are defined as storage group attributes – the maximum bandwidth (in MB per second) and the maximum IOPS. The Host I/O limit for a storage group can be either active or inactive.

Service levels

Service levels in PowerMaxOS enable a storage administrator to define quality-of-service criteria for individual storage groups.

Topics:

- [Definition of service levels](#)
- [Use of service levels to maintain system performance](#)
- [Usage examples](#)
- [Manage service levels](#)

Definition of service levels

A service level is a property of a storage group (SG) that is managed using Unisphere, Solutions Enabler, and Mainframe Enablers. The service level defines the target response times for I/O operations that involve the SG, a priority, and a compliance range:

- Target response time defines the average response time expected for an SG associated with a particular service level.
- Priority defines the importance of a service level relative to all other service levels. I/O associated with an SG that has a high priority service level is serviced more quickly than I/O to an SG with a lower-priority service level.
- Compliance range is a range of response times that Unisphere uses to determine whether to issue an alert. Should the response time for an SG fall outside this compliance range, Unisphere issues an alert.

The ranges that Unisphere uses are available on the **SG Compliance** tab.

Some service levels have a minimum response time (known as the *floor*). The floor defines the shortest time that each I/O operation on an SG with that service level takes to complete.

The storage administrator can use the service levels to help ensure that the performance of high-priority applications is not disrupted by lower-priority ones.

Defined service levels

PowerMaxOS defines six service levels:

- Diamond
- Platinum
- Gold
- Silver
- Bronze
- Optimized

These are the names as supplied, but the storage administrator can change any of them.


This table shows the target response time for each service level along with an indication of whether the service level has a floor value.

| Service Level | Target | Floor |
|---------------|---|----------------|
| Diamond | 0.4 ms (SCM storage) 0.6 ms (NVMe storage) | No |
| Platinum | 0.8 ms | No |
| Gold | 1 ms | No |
| Silver | 3.6 ms | approx. 3.6 ms |

| Service Level | Target | Floor |
|---------------|----------------|----------------|
| Bronze | 7.2 ms | approx. 7.2 ms |
| Optimized | Not applicable | Not applicable |

Service level priorities

Together, the five service levels create a priority list of service levels, each providing different quality-of-service criteria:

| Service level | Priority |
|---------------|---|
| Diamond | Highest |
| Platinum |  |
| Gold | |
| Silver | |
| Bronze | |
| | Lowest |

Optimized is not part of the priority scheme for service levels.

Default service levels

The default service level differs, depending on the tool you use to create a storage group:

| Management tool | Default service level |
|-----------------------|-----------------------|
| Unisphere | Diamond |
| Unisphere REST API | Optimized |
| Solutions Enabler CLI | Optimized |
| Mainframe Enablers | Optimized |

Availability of service levels

This table shows that availability of service levels in FBA and CKD environments.

| Service level | Availability | |
|---------------|--------------|-----|
| | FBA | CKD |
| Diamond | ✓ | ✓ |
| Platinum | ✓ | X |
| Gold | ✓ | X |
| Silver | ✓ | X |
| Bronze | ✓ | X |
| Optimized | ✓ | ✓ |

Use of service levels to maintain system performance

PowerMaxOS uses the service level property of each SG to maintain system performance.

Application of the floor value

The minimum response time for I/O operations to SGs that have the Silver or Bronze service level is the floor value, even if the storage array can provide a better response time. This behavior provides capacity for I/O requests for SGs with higher-priority service levels to complete promptly. Diamond, Platinum, and Gold service levels do not have an in-built delay and can complete immediately. However Diamond takes precedence over both Platinum and Gold, while Platinum takes precedence over Gold.

Imposed delays on response times

The load on the system could increase such that I/O on SGs with high-priority service levels cannot meet their target response times. In this case, PowerMaxOS delays I/O to other SGs with lower-priority service levels. This policy enables I/O for higher-priority SGs to bypass that for the lower-priority SGs and so maintain the necessary service. The process of slowing down I/O for SGs with lower-priority service levels is known as *throttling*. The maximum delay that PowerMaxOS can apply to the response time of an SG is approximately 10 times the service level's target response time.

The delay applied to a response time gradually increases over time as the system load increases. Similarly, the delay applied to a response time gradually decreases as the system load decreases. The strategy of gradually increasing and decreasing delays to response times avoids spikes (the sudden increase in a response time from a low to a high value) and thrashing.

The cascade of delays to response times

Exceeding the target response time for an SG with a specific service level causes throttling to occur on groups that have lower-priority service levels. Throttling cascades down from the service level that cannot meet its response time criterion:

- Diamond throttles Platinum, Gold, Silver, and Bronze SGs
- Platinum throttles Gold, Silver, and Bronze SGs
- Gold throttles Silver and Bronze SGs
- Silver throttles Bronze SGs

SGs with the Optimized service level

SGs that have the Optimized service level are exempt from this performance maintenance regime. I/O requests for such groups are not throttled. However, the response time for I/O to Optimized SGs may degrade as the system load increases even though PowerMaxOS does not intentionally throttle I/O to those SGs. By the same token, Optimized SGs that experience increased response times do not cause delays to be applied to SGs that have any other service level.

Usage examples

Here are three examples of using service levels:

- Protected application
- Service provider
- Relative application priority

Protected application

A storage administrator wants to ensure that a set of SGs is protected from the performance impact of other, noncritical applications that use the storage array. In this case, the administrator assigns the Diamond service level to the critical SGs and sets lower-priority service levels on all other SGs.

For instance:

- An enterprise-critical OLTP application requires almost immediate response to each I/O operation. The storage administrator may assign the Diamond level to its SGs.
- A batch program that runs overnight has less stringent requirements. So, the storage administrator may assign the Bronze level to its SGs.

Service provider

A provider of storage for external customers has a range of prices. Storage with lower response times is more costly than that with longer response times. In this case, the provider uses service levels to establish SGs that provide the required range of performance. An important part of this strategy is the use of the Silver and Bronze service levels to introduce delays even though the storage array could provide a shorter response time.

Relative application priority

A site wants to have the best possible performance for all applications. However, there is a relative priority among the protected applications. To achieve this, the storage administrator can assign Diamond, Platinum, and Gold to the SGs that the applications use. The SGs for the higher priority applications have the Diamond service level. The Platinum and Gold service levels are assigned to the remaining SGs depending on the relative priority of the applications.

In normal conditions, there is no delay to any SG because the array has the capacity to handle all I/O requests. However, should the workload increase and it is not possible to service all I/O requests immediately, the SGs with Platinum and Gold service levels begin to experience delay. This delay, however, is in proportion to the service level allocated to each SG.

Manage service levels

Solutions Enabler, Mainframe Enablers, and Unisphere for PowerMax have facilities to manage service levels:

- Create an SG and assign it a service level
- Set the service level for an SG
- Change the service level of an SG
- Remove the service level from an SG
- View the service level of an SG
- View the available service levels
- Rename a service level
- View the service level compliance of all SGs or a specific SG (Unisphere only)
- Set up, modify, remove, and view alert policies that monitor the performance of an SG relative to its service level (Unisphere only)
- Create service level compliance reports (Unisphere only)

Automated data placement

Automated data placement is a feature of PowerMaxOS 5978.444.444 and later on PowerMax arrays. It takes advantage of the superior performance of SCM drives to optimize access to frequently accessed data and data with high-priority service levels.

Topics:

- [Environment](#)
- [Operation](#)
- [Service level biasing](#)
- [Compression and deduplication](#)
- [Availability](#)

Environment

The performance of SCM drives is an order of magnitude better than NVMe drives. So an array that contains both types of drive effectively has two storage tiers: the higher performance SCM drives and the NVMe drives.

Automated data placement takes advantage of the performance difference to optimize access to data that is frequently accessed. The feature can also help to optimize access to storage groups that have higher priority service levels.

An array that contains only SCM drives or NVMe drives has only one tier of storage. So that type of array cannot use automated data placement.

Operation

Automated data placement monitors the frequency that the application host accesses data in the array. As a piece of data becomes accessed more frequently, automated data placement promotes that data to the SCM drives. Similarly, when a piece of data is accessed less frequently, automated data placement relegates it to the NVMe devices. Should more data need to be promoted but there is no available space in the SCM drives, automated data placement relegates data that has been accessed least frequently. This algorithm ensures that the SCM drives contain the most frequently accessed data.

Service level biasing

Automated data placement takes into account service levels when deciding whether to promote or relegate FBA data. That is, data associated with the Diamond service level has priority over that associated with other service levels. Data associated with the Silver and Bronze service level is never promoted to the SCM drives.

Compression and deduplication

When automated data placement is in operation, FBA data on the SCM devices is not compressed, but is deduplicated. Data on the NVMe drives has the same compression and deduplication characteristics as in previous releases of PowerMaxOS 5978 (see [Data efficiency](#) on page 32).

An array that contains SCM devices only has a slightly different compression strategy. Here, the most frequently accessed data is not compressed, but all other data is compressed.

Availability

Automated data placement is available for arrays that contain any combination of FBA and CKD devices.

Native local replication with TimeFinder

This chapter introduces the local replication features.

Topics:

- [About TimeFinder](#)
- [Mainframe SnapVX and zDP](#)
- [Snapshot policy](#)

About TimeFinder

Dell EMC TimeFinder delivers point-in-time copies of volumes that can be used for backups, decision support, data warehouse refreshes, or any other process that requires parallel access to production data.

Previous VMAX families offered multiple TimeFinder products, each with their own characteristics and use cases. These traditional products required a target volume to retain snapshot or clone data.

PowerMaxOS and HYPERMAX OS introduce TimeFinder SnapVX which provides the best aspects of the traditional TimeFinder offerings combined with increased scalability and ease-of-use.

TimeFinder SnapVX emulates the following legacy replication products:

- FBA devices:
 - TimeFinder/Clone
 - TimeFinder/Mirror
 - TimeFinder VP Snap
- Mainframe (CKD) devices:
 - TimeFinder/Clone
 - TimeFinder/Mirror
 - TimeFinder/Snap
 - Dell EMC Dataset Snap
 - IBM FlashCopy (Full Volume and Extent Level)

TimeFinder SnapVX dramatically decreases the impact of snapshots and clones:

- For snapshots, this is done by using redirect on write technology (ROW).
- For clones, this is done by storing changed tracks (deltas) directly in the Storage Resource Pool of the source device - sharing tracks between snapshot versions and also with the source device, where possible.

There is no need to specify a target device and source/target pairs. SnapVX supports up to 256 snapshots per volume. Each snapshot can have a name and an automatic expiration date.

Access to snapshots

With SnapVX, a snapshot can be accessed by *linking* it to a host accessible volume (known as a target volume). Target volumes are standard PowerMax TDEVs. Up to 1024 target volumes can be linked to the snapshots of the source volumes. The 1024 links can all be to the same snapshot of the source volume, or they can be multiple target volumes linked to multiple snapshots from the same source volume. However, a target volume may be linked only to one snapshot at a time.

Snapshots can be cascaded from linked targets, and targets can be linked to snapshots of linked targets. There is no limit to the number of levels of cascading, and the cascade can be broken.

SnapVX links to targets in the following modes:

- **Nocopy Mode (Default):** SnapVX does not copy data to the linked target volume but still makes the point-in-time image accessible through pointers to the snapshot. The target device is modifiable and retains the full image in a space-efficient manner even after unlinking from the point-in-time.
- **Copy Mode:** SnapVX copies all relevant tracks from the snapshot's point-in-time image to the linked target volume. This creates a complete copy of the point-in-time image that remains available after the target is unlinked.

If an application needs to find a particular point-in-time copy among a large set of snapshots, SnapVX enables you to link and relink until the correct snapshot is located.

Online device expansion

PowerMaxOS provides facilities for the online expansion of devices in a TimeFinder.

Interoperability with legacy TimeFinder products

TimeFinder SnapVX and PowerMaxOS emulate legacy TimeFinder and IBM FlashCopy replication products to provide backwards compatibility. You can run your legacy replication scripts and jobs on PowerMax arrays running TimeFinder SnapVX and PowerMaxOS without altering them.


Arrays that run PowerMaxOS 5978.444.444 and later enable coexistence and interoperability of SnapVX with legacy TimeFinder products. On such an array, a device can simultaneously be the source of a SnapVX operation and the source of one of these legacy TimeFinder products:

- TimeFinder/Clone
- TimeFinder/Mirror
- TimeFinder VP Snap

The target device of a legacy TimeFinder product cannot be the source device for SnapVX. Similarly, the target device of SnapVX cannot be the source device for a legacy TimeFinder product.

Uses for the coexistence of SnapVX with legacy TimeFinder products include:

- A site wants to keep its current, legacy configuration in place while trying out SnapVX.
- Moving to SnapVX may require the deletion of existing legacy sessions and that violates local business policies.

 **NOTE:** Coexistence of SnapVX and legacy TimeFinder products is not available when the source of a SnapVX session is undergoing a restore operation.

Targetless snapshots

With the TimeFinder SnapVX management interfaces you can take a snapshot of an entire PowerMax Storage Group using a single command. With this in mind, PowerMax supports up to 64K storage groups. The number of groups is enough even in the most demanding environment to provide one for each application. The storage group construct already exists in most cases as they are created for masking views. TimeFinder SnapVX uses this existing structure, so reducing the administration required to maintain the application and its replication environment.

Creation of SnapVX snapshots does not require preconfiguration of extra volumes. In turn, this reduces the amount of cache that SnapVX snapshots use and simplifies implementation. Snapshot creation and automatic termination can easily be scripted.

The following Solutions Enabler example creates a snapshot with a 2-day retention period. The command can be scheduled to run as part of a script to create multiple versions of the snapshot. Each snapshot shares tracks where possible with the other snapshots and the source devices. Use a cron job or scheduler to run the snapshot script on a schedule to create up to 256 snapshots of the source volumes; enough for a snapshot every 15 minutes with 2 days of retention:

```
symsnapvx -sid 001 -sg StorageGroup1 -name sg1_snap establish -ttl -delta 2
```

If a restore operation is required, any of the snapshots created by this example can be specified.

When the storage group transitions to a restored state, the restore session can be terminated. The snapshot data is preserved during the restore process and can be used again should the snapshot data be required for a future restore.

Secure snaps

Secure snaps prevent administrators or other high-level users from deleting snapshot data, intentionally or not. Also, Secure snaps are also immune to automatic failure resulting from running out of Storage Resource Pool (SRP) or Replication Data Pointer (RDP) space on the array.

When the administrator creates a secure snapshot, they assign it an expiration date and time. The administrator can express the expiration either as a delta from the current date or as an absolute date. Once the expiration date passes, and if the snapshot has no links, PowerMaxOS automatically deletes the snapshot. Before its expiration, administrators can only extend

the expiration date; they cannot shorten the date or delete the snapshot. If a secure snapshot expires, and it has a volume linked to it, or an active restore session, the snapshot is not deleted. However, it is no longer considered secure.

NOTE: Secure snapshots may only be terminated after they expire or by customer-authorized Dell EMC support. Refer to Knowledgebase article 498316 for more information.

Provision multiple environments from a linked target

Use SnapVX to create multiple test and development environments using linked snapshots. To access a point-in-time copy, create a link from the snapshot data to a host mapped target device.

Each linked storage group can access the same snapshot, or each can access a different snapshot version in either no copy or copy mode. Changes to the linked volumes do not affect the snapshot data. To roll back a test or development environment to the original snapshot image, perform a relink operation.

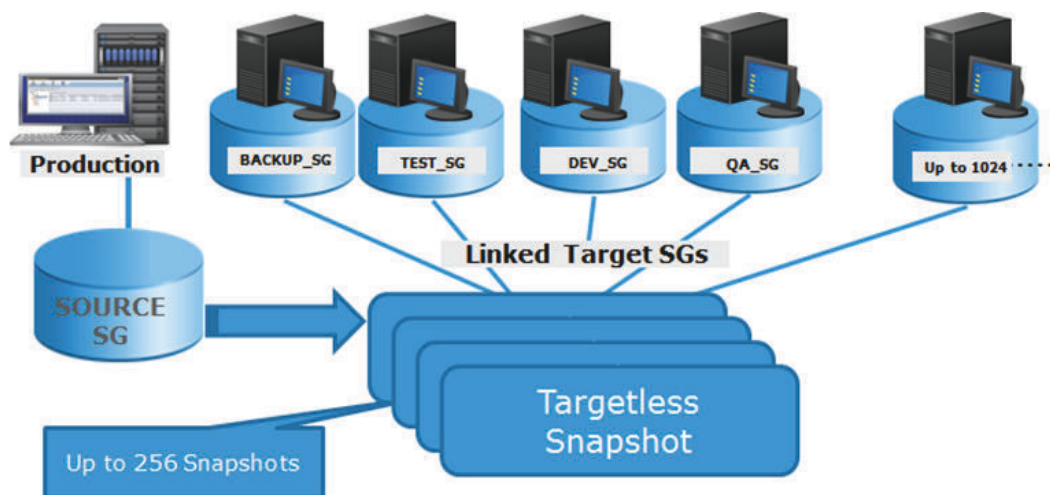


Figure 8. SnapVX targetless snapshots

NOTE: Unmount target volumes before issuing the relink command to ensure that the host operating system does not cache any filesystem data. If accessing through VPLEX, ensure that you follow the procedure outlined in the technical note *VPLEX: Leveraging Array Based and Native Copy Technologies*, available on the Dell EMC support website.

Once the relink is complete, volumes can be remounted.

Snapshot data is unchanged by the linked targets, so the snapshots can also be used to restore production data.

Cascading snapshots

Presenting sensitive data to test or development environments often requires that the source of the data be disguised beforehand. Cascaded snapshots provides this separation and disguise, as shown in the following image.

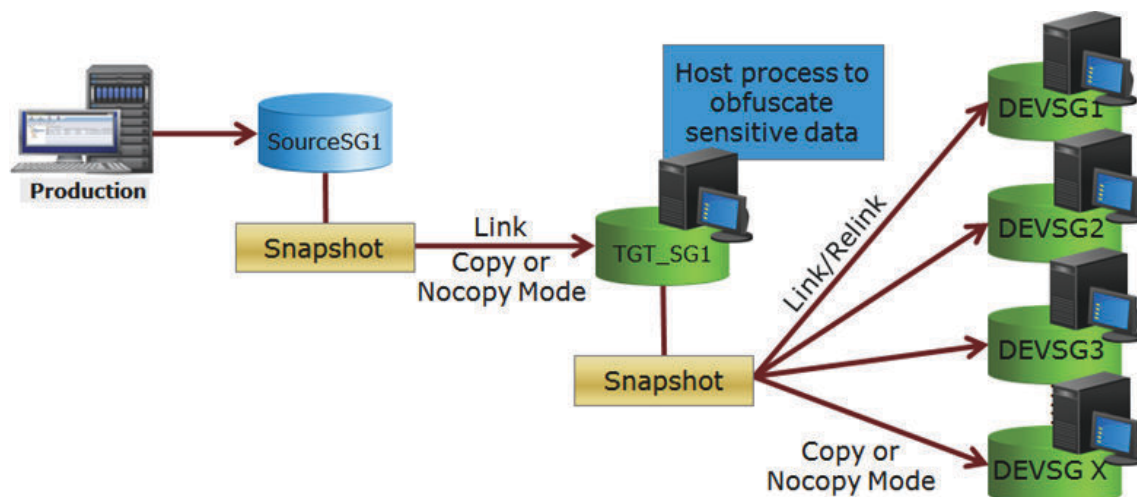


Figure 9. SnapVX cascaded snapshots

If no change to the data is required before presenting it to the test or development environments, there is no need to create a cascaded relationship.

Accessing point-in-time copies

To access a point-in-time-copy, create a link from the snapshot data to a host mapped target device. The links may be created in Copy mode for a permanent copy on the target device, or in NoCopy mode for temporary use. Copy mode links create full-volume, full-copy clones of the data by copying it to the target device's Storage Resource Pool. NoCopy mode links are space-saving snapshots that only consume space for the changed data that is stored in the source device's Storage Resource Pool.

PowerMaxOS supports up to 1,024 linked targets per source device.

NOTE: When a target is first linked, all of the tracks are undefined. This means that the target does not know where in the Storage Resource Pool the track is located, and host access to the target must be derived from the SnapVX metadata. A background process eventually defines the tracks and updates the thin device to point directly to the track location in the source device's Storage Resource Pool.

Mainframe SnapVX and zDP

Data Protector for z Systems (zDP) is a mainframe software solution that is layered on SnapVX on PowerMax arrays. Using zDP you can recover from logical data corruption with minimal data loss. zDP achieves this by providing multiple, frequent, consistent point-in-time copies of data automatically. You can then use these copies to recover an application or the environment to a point prior to the logical corruption.

By providing easy access to multiple different point-in-time copies of data (with a granularity of minutes), precise recovery from logical data corruption can be performed using application-based recovery procedure. zDP results in minimal data loss compared to other methods such as restoring data from daily or weekly backups.

As shown in [zDP operation](#) on page 71, you can use zDP to create and manage multiple point-in-time snapshots of volumes. Each snapshot is a pointer-based, point-in-time image of a single volume. These images are created using the SnapVX feature of PowerMaxOS. SnapVX is a space-efficient method for making snapshots of thin devices and consuming additional storage capacity only when changes are made to the source volume.

There is no need to copy each snapshot to a target volume as SnapVX separates the capturing of a point-in-time copy from its usage. Capturing a point-in-time copy does not require a target volume. Using a point-in-time copy from a host requires linking the snapshot to a target volume.

From PowerMaxOS 5978.444.444 onwards, there can be up to 1024 snapshots of each source volume. On earlier versions of PowerMaxOS, HYPERMAX OS, and Enginuity there can be up to 256 snapshots for each source volume. PowerMaxOS 5978.444.444 also provides facilities for creating a snapshot on demand.




Figure 10. zDP operation

These snapshots share allocations to the same track image whenever possible while ensuring they each continue to represent a unique point-in-time image of the source volume. Despite the space efficiency achieved through shared allocation to unchanged data, additional capacity is required to preserve the pre-update images of changed tracks captured by each point-in-time snapshot.

zDP includes the secure snap facility (see [Secure snaps](#) on page 68).

The process of implementing zDP has two phases — the planning phase and the implementation phase.

- The planning phase is done in conjunction with your Dell EMC representative who has access to tools that can help size the capacity needed for zDP if you are currently a PowerMax or VMAX All Flash user.
- The implementation phase uses the following methods for z/OS:
 - A batch interface that allows you to submit jobs to define and manage zDP.
 - A zDP run-time environment that executes under SCF to create snapshots.

For details on zDP usage, refer to the *TimeFinder SnapVX and zDP Product Guide*. For details on zDP usage in z/TPF, refer to the *TimeFinder Controls for z/TPF Product Guide*.

Snapshot policy

The Snapshot policy feature provides snapshot orchestration at scale (1024 snaps per storage group). The feature simplifies snapshot management for standard and cloud snapshots.

Snapshots can be used to recover from data corruption, accidental deletion, or other damage, offering continuous data protection. A large number of snapshots can be difficult to manage. The Snapshot policy feature provides an end to end solution to create, schedule and manage standard (local) and cloud snapshots.

The snapshot policy (Recovery Point Objective (RPO)) specifies how often the snapshot should be taken and how many of the snapshots should be retained. The snapshot may also be specified to be secure (these snapshots cannot be terminated by users before their time to live (TTL), derived from the snapshot policy's interval and maximum count, has expired.) Up to four policies can be associated with a storage group, and a snapshot policy can be associated with many storage groups.

The following rules apply to snapshot policies:

- The maximum number of snapshot policies that can be created on a storage system is 20. Multiple storage groups can be associated with a snapshot policy.
- A maximum of four snapshot policies can be associated with an individual storage group.
- A storage group or device can have a maximum of 256 manual snapshots.
- A storage group or device can have a maximum of 1024 snapshots.
- The oldest unused snapshots are removed or recycled in accordance with the specified policy max_count value.
- When devices are added to a snapshot policy storage group, snapshot policies that apply to the storage group are applied to the added devices.

- When devices are removed from a snapshot policy storage group, snapshot policies that apply to the storage group are no longer applied to the removed devices.
- If overlapping snapshot policies are applied to storage groups, they run and take snapshots independently.

Compliance information is provided for each snapshot policy that is directly associated with (not inherited to) a storage group.

Snapshot compliance for a storage group is taken as the lowest compliance value for any of the snapshot policies that are directly associated with the storage group.

Compliance for a snapshot policy that is associated with a storage group is based on the number of good snapshots within the retention count. The retention count is translated to a retention period for compliance calculation. The retention period is the snapshot interval multiplied by the snapshot maximum count. For example, a 1 hr interval with a 30 snapshot count means a 30-hour retention period.

The compliance threshold for green to yellow change is the maximum count, that is, all snapshots must be good and in place for the compliance to be green. If there is one snapshot short (missing or failed), then the compliance turns yellow.

The compliance threshold value for yellow to red is stored in the snapshot policy definition. Once the number of good snapshots falls below this value, compliance turns red.

Snapshot compliance is calculated by polling the storage system once an hour for SnapVX related information for storage groups which have snapshot policies that are associated with them. The returned snapshot information is summarized into the required information for the database compliance entries.

When the maximum count of snapshots for a snapshot policy is changed, this changes the compliance for the storage group or service level combination. Compliance values are updated accordingly simultaneously.

If compliance calculation is performed during the creation of a snapshot, then an establish-in-progress state may be detected. This is acceptable for the most recent snapshot but is considered failed for any older snapshot.

When a storage group and service level have only recently been associated and the full maximum count of snapshots has not yet been reached, the calculation is scaled to the number of snapshots that are available and represents compliance accordingly until the full maximum count of snapshots has been reached. If a snapshot failed to be taken for a reason, such as the storage group or service level was suspended, or a snapshot was manually terminated before the maximum snapshot count was reached, the compliance is reported as out of compliance appropriately.

When the service level interval is changed, the compliance window changes and the number of snapshots may not exist for correct compliance.

If a service level is suspended or a storage group or service level combination is suspended, snapshots are not created and older snapshots fall outside the compliance window and the maximum count of required snapshot is not found.

Manual termination of snapshots inside the compliance window results in the storage group or service level combination falling out of compliance.

Remote replication

This chapter introduces the remote replication facilities.

Topics:

- [Native remote replication with SRDF](#)
- [SRDF/Metro](#)
- [RecoverPoint](#)
- [Remote replication using eNAS](#)

Native remote replication with SRDF

The Dell EMC Symmetrix Remote Data Facility (SRDF) family of products offers a range of array-based disaster recovery, parallel processing, and data migration solutions for Dell EMC storage systems, including:

- PowerMaxOS for PowerMax 2000 and 8000 arrays and for VMAX All Flash 450F and 950F arrays
- HYPERMAX OS for VMAX All Flash 250F, 450F, 850F, and 950F arrays
- HYPERMAX OS for VMAX 100K, 200K, and 400K arrays
- Enginuity for VMAX 10K, 20K, and 40K arrays

SRDF disaster recovery solutions use “active, remote” mirroring and dependent-write logic to create consistent copies of data. Dependent-write consistency ensures transactional consistency when the applications are restarted at the remote location. You can tailor your SRDF solution to meet various Recovery Point Objectives and Recovery Time Objectives.

Using SRDF, you can create complete solutions to:

- Create real-time or dependent-write-consistent copies at 1, 2, or 3 remote arrays.
- Move data quickly over extended distances.
- Provide 3-site disaster recovery with zero data loss recovery, business continuity protection and disaster-restart.

You can integrate SRDF with other Dell EMC products to create complete solutions to:

- Restart operations after a disaster with zero data loss and business continuity protection.
- Restart operations in cluster environments. For example, Microsoft Cluster Server with Microsoft Failover Clusters.
- Monitor and automate restart operations on an alternate local or remote server.
- Automate restart operations in VMware environments.

PowerMaxOS provides facilities for the online expansion of devices in a SRDF configuration.

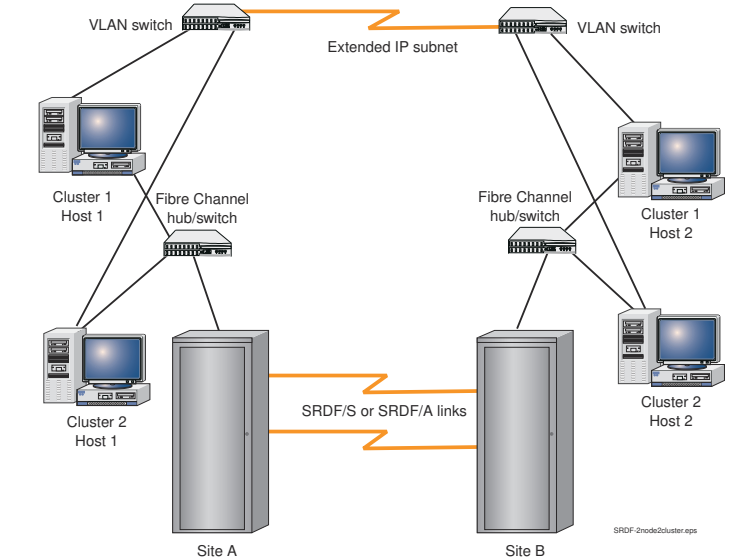
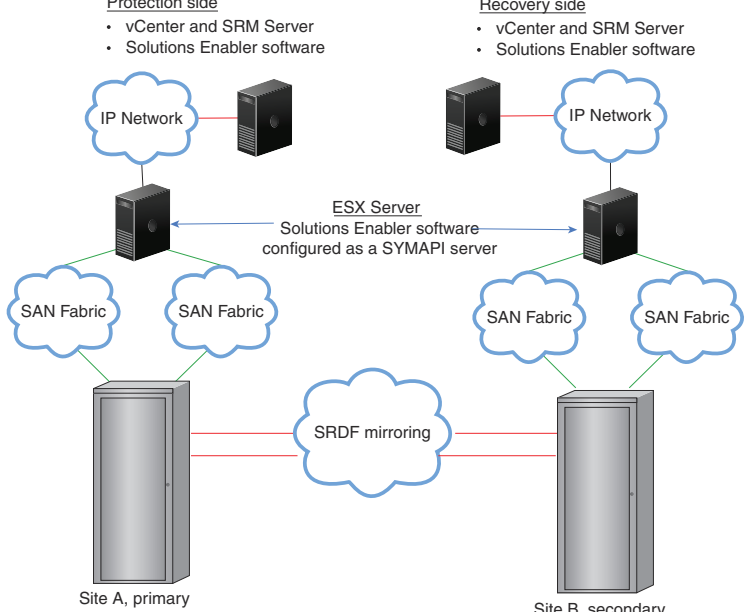
SRDF 2-site solutions

The following table describes SRDF 2-site solutions.

Table 12. SRDF 2-site solutions

| Solution highlights | Site topology |
|--|---------------|
| SRDF/Synchronous (SRDF/S) Maintains a real-time copy of production data at a physically separated array. <ul style="list-style-type: none">No data exposure.Ensured consistency protection with SRDF/Consistency Group.Recommended maximum distance of 200 km (125 miles) between arrays as application latency may rise to unacceptable levels at longer distances.^a | |
| SRDF/Asynchronous (SRDF/A) Maintains a dependent-write consistent copy of the data on a remote secondary site. The sites can be an unlimited distance apart. The copy of the data at the secondary site is seconds behind the primary site. | |
| SRDF/Data Mobility (SRDF/DM) Enables the fast transfer of data from R1 to R2 devices over extended distances. <ul style="list-style-type: none">Uses adaptive copy mode to transfer data.Designed for migration or data replication purposes, not for disaster restart solutions. | |
| SRDF/Automated Replication (SRDF/AR) <ul style="list-style-type: none">Combines SRDF and TimeFinder to optimize bandwidth requirements and provide a long-distance disaster restart solution.Operates in 2-site solutions that use SRDF/DM in combination with TimeFinder. | |

Table 12. SRDF 2-site solutions (continued)

| Solution highlights | Site topology |
|--|---|
| <p>SRDF/Cluster Enabler (CE)</p> <ul style="list-style-type: none"> Integrates SRDF/S or SRDF/A with Microsoft Failover Clusters (MSCS) to automate or semi-automate site failover. Complete solution for restarting operations in cluster environments (MSCS with Microsoft Failover Clusters). Expands the range of cluster storage and management capabilities while ensuring full protection of the SRDF remote replication. |  |
| <p>SRDF and VMware Site Recovery Manager</p> <p>Completely automates storage-based disaster restart operations for VMware environments in SRDF topologies.</p> <ul style="list-style-type: none"> The Dell EMC SRDF Adapter enables VMware Site Recovery Manager to automate storage-based disaster restart operations in SRDF solutions. Can address configurations in which data are spread across multiple storage arrays or SRDF groups. Requires that the adapter is installed on each array to facilitate the discovery of arrays and to initiate failover operations. Implemented with: <ul style="list-style-type: none"> SRDF/S SRDF/A SRDF/Star TimeFinder |  |

- a. In some circumstances, using SRDF/S over distances greater than 200 km may be feasible. Contact your Dell EMC representative for more information.

SRDF multi-site solutions

The following table describes SRDF multi-site solutions.

Table 13. SRDF multi-site solutions

| Solution highlights | Site topology |
|---|---------------|
| <p>SRDF/Automated Replication (SRDF/AR)</p> <ul style="list-style-type: none"> Combines SRDF and TimeFinder to optimize bandwidth requirements and provide a long-distance disaster restart solution. Operates in a 3-site environment that uses a combination of SRDF/S, SRDF/DM, and TimeFinder. | |
| <p>Concurrent SRDF</p> <p>3-site disaster recovery and advanced multi-site business continuity protection.</p> <ul style="list-style-type: none"> Data on the primary site is concurrently replicated to 2 secondary sites. Replication to remote site can use SRDF/S, SRDF/A, or adaptive copy. | |
| <p>Cascaded SRDF</p> <p>3-site disaster recovery and advanced multi-site business continuity protection.</p> <p>Data on the primary site (Site A) is synchronously mirrored to a secondary site (Site B), and then asynchronously mirrored from the secondary site to a tertiary site (Site C).</p> | |
| <p>SRDF/Star</p> <p>3-site data protection and disaster recovery configuration with zero data loss recovery, business continuity protection and disaster restart.</p> <ul style="list-style-type: none"> Available in 2 configurations: <ul style="list-style-type: none"> Cascaded SRDF/Star Concurrent SRDF/Star Differential synchronization allows rapid reestablishment of mirroring among surviving sites in a multi-site disaster recovery implementation. | |

Table 13. SRDF multi-site solutions (continued)

| Solution highlights | Site topology |
|--|---------------|
| <ul style="list-style-type: none">Implemented using SRDF consistency groups (CG) with SRDF/S and SRDF/A. | |

Interfamily compatibility

SRDF supports connectivity between different operating environments and arrays. Arrays running PowerMaxOS can connect to legacy arrays running older operating environments. In mixed configurations where arrays are running different versions, SRDF features of the lowest version are supported.

PowerMax arrays can connect to:

- PowerMax arrays running PowerMaxOS
- VMAX 250F, 450F, 850F, and 950F arrays running HYPERMAX OS
- VMAX 100K, 200K, and 400K arrays running HYPERMAX OS
- VMAX 10K, 20K, and 40K arrays running Enginuity 5876 with an Enginuity ePack

NOTE: When you connect between arrays running different operating environments, limitations may apply. Information about which SRDF features are supported, and applicable limitations for 2-site and 3-site solutions is in the *SRDF Interfamily Connectivity Information*.

This interfamily connectivity allows you to add the latest hardware platform/operating environment to an existing SRDF solution, enabling technology refreshes.

SRDF device pairs

An SRDF device pair is a logical device that is paired with another logical device that resides in a second array. The arrays are connected by SRDF links.

Encapsulated Data Domain devices that are used for Storage Direct cannot be part of an SRDF device pair.

R1 and R2 devices

An R1 device is the member of the device pair at the source (production) site. R1 devices are generally Read/Write accessible to the application host.

An R2 device is the member of the device pair at the target (remote) site. During normal operations, host I/O writes to the R1 device are mirrored over the SRDF links to the R2 device. In general, data on R2 devices is not available to the application host while the SRDF relationship is active. In SRDF synchronous mode, however, an R2 device can be in Read Only mode that allows a host to read from the R2.

In a typical environment:

- The application production host has Read/Write access to the R1 device.
- An application host connected to the R2 device has Read Only (Write Disabled) access to the R2 device.




Figure 11. R1 and R2 devices

R11 devices

R11 devices operate as the R1 device for two R2 devices. Links to both R2 devices are active.

R11 devices typically occur in 3-site concurrent configurations where data on the R11 site is mirrored to two secondary (R2) arrays:




Figure 12. R11 device in concurrent SRDF

R21 devices

R21 devices have a dual role and are used in cascaded 3-site configurations where:

- Data on the R1 site is synchronously mirrored to a secondary (R21) site, and then
- Asynchronously mirrored from the secondary (R21) site to a tertiary (R2) site:




Figure 13. R21 device in cascaded SRDF

The R21 device acts as a R2 device that receives updates from the R1 device, and as a R1 device that sends updates to the R2 device.

When the R1->R21->R2 SRDF relationship is established, no host has write access to the R21 device.

In arrays that run Enginuity, the R21 device can be diskless. That is, it consists solely of cache memory and does not have any associated storage device. It acts purely to relay changes in the R1 device to the R2 device. This capability requires the use of thick devices. Systems that run PowerMaxOS or HYPERMAX OS contain thin devices only, so setting up a diskless R21 device is not possible on arrays running those environments.

R22 devices

R22 devices:

- Have two R1 devices, only one of which is active at a time.
- Are typically used in cascaded SRDF/Star and concurrent SRDF/Star configurations to decrease the complexity and time required to complete failover and failback operations.
- Enables recovery to occur without removing old SRDF pairs and creating new ones.




Figure 14. R22 devices in cascaded and concurrent SRDF/Star

Dynamic device personalities

SRDF devices can dynamically swap “personality” between R1 and R2. After a personality swap:

- The R1 in the device pair becomes the R2 device, and
- The R2 becomes the R1 device.

Swapping R1/R2 personalities allows the application to be restarted at the remote site without interrupting replication if an application fails at the production site. After a swap, the R2 side (now R1) can control operations while being remotely mirrored at the primary (now R2) site.

An R1/R2 personality swap is not supported:

- If the R2 device is larger than the R1 device.
- If the device to be swapped is participating in an active SRDF/A session.
- In SRDF/EDP topologies diskless R11 or R22 devices are not valid end states.
- If the device to be swapped is the target device of any TimeFinder or Dell EMC Compatible flash operations.

SRDF modes of operation

The SRDF mode of operation determines:

- How R1 devices are remotely mirrored to R2 devices across the SRDF links
- How I/O operations are processed
- When the acknowledgment is returned to the application host that issued an I/O write command

In SRDF there are three principal modes:

- Synchronous
- Asynchronous
- Adaptive copy

Synchronous mode

Synchronous mode maintains a real-time mirror image of data between the R1 and R2 devices over distances up to 200 km (125 miles). Host data is written to both arrays in real time. The application host does not receive the acknowledgment until the data has been stored in the cache of both arrays.

Asynchronous mode

Asynchronous mode maintains a dependent-write consistent copy between the R1 and R2 device over unlimited distances. On receiving data from the application host, SRDF on the R1 side of the link writes that data to its cache. Also it batches the data received into *delta sets*. Delta sets are transferred to the R2 device in timed cycles. The application host receives the acknowledgment once data is successfully written to the cache on the R1 side.

Adaptive copy modes

Adaptive copy modes:

- Accumulate write requests that are destined for the R2 device on the R1 side, but not in cache memory.
- A background copy process sends the outstanding write requests to the R2 device.
- Allow the R1 and R2 devices to be out of synchronization by user-defined *maximum skew* value. Once the skew value is exceeded, SRDF transfers the batched data to the R2 device.
- Send the acknowledgment to the application host once the data is successfully written to cache on the R1 side.

Unlike asynchronous mode, the adaptive copy modes do not guarantee a dependent-write copy of data on the R2 devices.

SRDF groups

An SRDF group defines the logical relationship between SRDF devices and directors on both sides of an SRDF link.

Group properties

The properties of an SRDF group are:

- Label (name)
- Set of ports on the local array used to communicate over the SRDF links
- Set of ports on the remote array used to communicate over the SRDF links
- Local group number
- Remote group number
- One or more pairs of devices

The devices in the group share the ports and associated CPU resources of the port's directors.

Types of group

There are two types of SRDF group:

- Static: which are defined in the local array's configuration file.
- Dynamic: which are defined using SRDF management tools and their properties that are stored in the array's cache memory.

On arrays running PowerMaxOS or HYPERMAX OS all SRDF groups are dynamic.

Director boards, links, and ports

SRDF links are the logical connections between SRDF groups and their ports. The ports are physically connected by cables, routers, extenders, switches and other network devices.

 **NOTE:** Two or more SRDF links per SRDF group are required for redundancy and fault tolerance.

The relationship between the resources on a director (CPU cores and ports) varies depending on the operating environment.

PowerMaxOS and HYPERMAX OS

On arrays running PowerMaxOS or HYPERMAX OS :

- The relationship between the SRDF emulation and resources on a director is configurable:
 - One director/multiple CPU cores/multiple ports
 - Connectivity (ports in the SRDF group) is independent of compute power (number of CPU cores). You can change the amount of connectivity without changing compute power.
- Each director has up to 16 front end ports, any or all of which can be used by SRDF. Both the SRDF Gigabit Ethernet and SRDF Fibre Channel emulations can use any port.
- The data path for devices in an SRDF group is not fixed to a single port. Instead, the path for data is shared across all ports in the group.

Mixed configurations: PowerMaxOS or HYPERMAX OS and Enginuity 5876

For configurations where one array is running Enginuity 5876, and the other array is running PowerMaxOS or HYPERMAX OS, these rules apply:

- On the 5876 side, an SRDF group can have the full complement of directors, but no more than 16 ports on the PowerMaxOS or HYPERMAX OS side.
- You can connect to 16 directors using one port each, 2 directors using 8 ports each or any other combination that does not exceed 16 per SRDF group.

SRDF consistency

Many applications, especially database systems, use dependent write logic to ensure data integrity. That is, each write operation must complete successfully before the next can begin. Without write dependency, write operations could get out of sequence resulting in irrecoverable data loss.

SRDF implements write dependency using the *consistency group* (also known as SRDF/CG). A consistency group consists of a set of SRDF devices that use write dependency. For each device in the group, SRDF ensures that write operations propagate to the corresponding R2 devices in the correct order.

However, if the propagation of any write operation to any R2 device in the group cannot complete, SRDF suspends propagation to all group's R2 devices. This suspension maintains the integrity of the data on the R2 devices. While the R2 devices are unavailable, SRDF continues to store write operations on the R1 devices. It also maintains a list of those write operations in their time order. When all R2 devices in the group become available, SRDF propagates the outstanding write operations, in the correct order, for each device in the group.

SRDF/CG is available for both SRDF/S and SRDF/A.

Data migration

Data migration is the one-time movement of data from one array to another. Once the movement is complete, the data is accessed from the secondary array. A common use of migration is to replace an older array with a new one.

Dell EMC support personnel can assist with the planning and implementation of migration projects.

SRDF multisite configurations enable migration to occur in any of these ways:

- Replace R2 devices.
- Replace R1 devices.
- Replace both R1 and R2 devices simultaneously.

For example, this diagram shows the use of concurrent SRDF to replace the secondary (R2) array in a 2-site configuration:




Figure 15. Migrating data and removing a secondary (R2) array

Here:

- The top section of the diagram shows the original, 2-site configuration.
- The lower left section of the diagram shows the interim, 3-site configuration with data being copied to two secondary arrays.
- The lower right section of the diagram shows the final, 2-site configuration where the new secondary array has replaced the original one.

The *Dell EMC SRDF Introduction* contains more information about using SRDF to migrate data.

More information

Here are other Dell EMC documents that contain more information about the use of SRDF in replication and migration:

SRDF Introduction

SRDF and NDM Interfamily Connectivity Information

SRDF/Cluster Enabler Plug-in Product Guide

Using the Dell EMC Adapter for VMware Site Recovery Manager Technical Book

Dell EMC SRDF Adapter for VMware Site Recovery Manager Release Notes

SRDF/Metro

In traditional SRDF configurations, only the R1 devices are Read/Write accessible to the application hosts. The R2 devices are Read Only and Write Disabled.

In SRDF/Metro configurations, however:

- Both the R1 and R2 devices are Read/Write accessible to the application hosts.
- Application hosts can write to both the R1 and R2 side of the device pair.
- R2 devices assume the same external device identity as the R1 devices. The identity includes the device geometry and device WWN.

This shared identity means that R1 and R2 devices appear to application hosts as a single, virtual device across two arrays.

Deployment options

SRDF/Metro can be deployed in either a single, multipathed host environment or in a clustered host environment:




Figure 16. SRDF/Metro

Hosts can read and write to both the R1 and R2 devices:

- In a single host configuration, a single host issues I/O operations. Multipathing software directs parallel reads and writes to each array.
- In a clustered host configuration, multiple hosts issue I/O operations. Those hosts access both sides of the SRDF device pair. Each cluster node has dedicated access to one of the storage arrays.
- In both configurations, writes to the R1 and R2 devices are synchronously copied to the paired device in the other array. SRDF/Metro software resolves any write conflicts to maintain consistent images on the SRDF device pairs.

SRDF/Metro Resilience

If either of the devices in a SRDF/Metro configuration become Not Ready, or connectivity between the devices is lost, SRDF/Metro must decide which side remains available to the application host. There are two mechanisms that SRDF/Metro can use : Device Bias and Witness.

Device Bias

Device pairs for SRDF/Metro are created with a *bias* attribute. By default, the create pair operation sets the bias to the R1 side of the pair. That is, if a device pair becomes Not Ready (NR) on the SRDF link, the R1 (bias side) remains accessible to the hosts, and the R2 (nonbias side) becomes inaccessible. However, if there is a failure on the R1 side, the host loses all connectivity to the device pair. The Device Bias method cannot make the R2 device available to the host.

Witness

A witness is a third party that mediates between the two sides of a SRDF/Metro pair to help:

- Decide which side remains available to the host

- Avoid a "split brain" scenario when both sides attempt to remain accessible to the host despite the failure

The witness method allows for intelligently choosing on which side to continue operations when the bias-only method may not result in continued host availability to a surviving, nonbiased array.

There are two forms of the Witness mechanism:

- **Array Witness:** The operating environment of a third array is the mediator.
- **Virtual Witness (vWitness):** A daemon running on a separate, virtual machine is the mediator.

When both sides run PowerMaxOS 5978 SRDF/Metro takes these criteria into account when selecting the side to remain available to the hosts (in priority order):

1. The side that has connectivity to the application host (requires PowerMaxOS 5978.444.444or later)
2. The side that has a SRDF/A DR leg
3. Whether the SRDF/A DR leg is synchronized
4. The side that has more than 50% of the RA or FA directors that are available
5. The side that is currently the bias side

The first of these criteria that one array has, and the other does not, stops the selection process. The side with the matched criteria is the preferred winner.

Disaster recovery facilities

Devices in SRDF/Metro groups can simultaneously be in other groups that replicate data to a third, disaster recovery site. There are two replication solutions. The solutions available in any SRDF/Metro configuration depends on the version of the operating environment that the participating arrays run:

- Highly-available disaster recovery – in configurations that consist of arrays that run PowerMaxOS 5978.669.669 and later
- Independent disaster recovery – in configurations that run all supported versions of PowerMaxOS 5978 and HYPERMAX OS 5977

Highly available disaster recovery (SRDF/Metro Smart DR)

SRDF/Metro Smart DR maintains a single, disaster recovery (DR) copy of the data in a SRDF/Metro pair on a third, remote array. This diagram shows the SRDF/Metro Smart DR configuration:

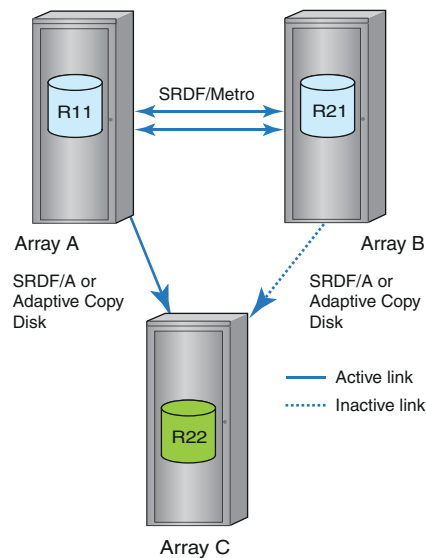


Figure 17. SRDF/Metro Smart DR

Notice that the device names differ from a standard SRDF/Metro configuration. This difference reflects the change in the device functions when SRDF/Metro Smart DR is in operation. For instance, the R1 side of the SRDF/Metro on Array A now has the name R11, because it is the R1 device to both the:

- R21 device on Array B in the SRDF/Metro configuration
- R22 device on Array C in the SRDF/Metro Smart DR configuration

Arrays A and B both have SRDF/Asynchronous or Adaptive Copy Disk connections to the DR array (Array C). However, only one of those connections is active at a time (in this example the connection between Array A and Array C). The two SRDF/A connections are known as the active and standby connections.

If a problem prevents Array A replicating data to Array C, the standby link between Array B and Array C becomes active and replication continues. Array A and Array B keep track of the data replicated to Array C to enable replication and avoid data loss.

Independent disaster recovery

Devices in SRDF/Metro groups can simultaneously be part of device groups that replicate data to a third, disaster-recovery site.

Either or both sides of the Metro region can be replicated. An organization can choose which ever configuration that suits its business needs. The following diagram shows the possible configurations:

NOTE: When the SRDF/Metro session is using a witness, the R1 side of the Metro pair can change based on the witness determination of the preferred side.




Figure 18. Disaster recovery for SRDF/Metro

The device names differ from a stand-alone SRDF/Metro configuration. This difference reflects the change in the devices' function when disaster recovery facilities are in place. For instance, when the R2 side is replicated to a disaster recovery site, its name changes to R21 because it is both the:

- R2 device in the SRDF/Metro configuration
- R1 device in the disaster-recovery configuration

When an SRDF/Metro uses a witness for resilience protection, the two sides periodically renegotiate the winning and losing sides. If the winning and losing sides switch as a result of renegotiation:

- An R11 device becomes an R21 device. That device was the R1 device for both the SRDF/Metro and disaster recovery configurations. Now the device is the R2 device of the SRDF/Metro configuration but it remains the R1 device of the disaster recovery configuration.
- An R21 device becomes an R11 device. That device was the R2 device in the SRDF/Metro configuration and the R1 device of the disaster recovery configuration. Now the device is the R1 device of both the SRDF/Metro and disaster recovery configurations.

Mobility ID with ALUA

Mobility ID with Asymmetric Logical Unit Access (ALUA) assigns a unique identifier to a device in a system. This identifier enables the device to be moved between arrays without the need for any reconfiguration on the host. PowerMaxOS brings Mobility ID with ALUA capabilities to SRDF/Metro. So, when both sides run PowerMaxOS you can specify the Mobility ID in the create pair operation in place of the regular device identifier.

More information

Here are other Dell EMC documents that contain more information on SRDF/Metro:

SRDF Introduction

SRDF/Metro vWitness Configuration Guide

SRDF Interfamily Connectivity Information

RecoverPoint

RecoverPoint is a comprehensive data protection solution designed to provide production data integrity at local and remote sites. RecoverPoint also provides the ability to recover data from a point in time using journaling technology.

The primary reasons for using RecoverPoint are:

- Remote replication to heterogeneous arrays
- Protection against Local and remote data corruption
- Disaster recovery
- Secondary device repurposing
- Data migrations

RecoverPoint systems support local and remote replication of data that applications are writing to SAN-attached storage. The systems use existing Fibre Channel infrastructure to integrate seamlessly with existing host applications and data storage subsystems. For remote replication, the systems use existing Fibre Channel connections to send the replicated data over a WAN, or use Fibre Channel infrastructure to replicate data asynchronously. The systems provide failover of operations to a secondary site in the event of a disaster at the primary site.

Previous implementations of RecoverPoint relied on a splitter to track changes made to protected volumes. The current implementation relies on a cluster of RecoverPoint nodes, provisioned with one or more RecoverPoint storage groups, leveraging SnapVX technology, on the storage array. Volumes in the RecoverPoint storage groups are visible to all the nodes in the cluster, and available for replication to other storage arrays.

RecoverPoint allows data replication of up to 8,000 LUNs for each RecoverPoint cluster and up to eight different RecoverPoint clusters attached to one array. Supported array types include PowerMax, VMAX All Flash, VMAX3, VMAX, VNX, VPLEX, and XtremIO.

RecoverPoint is licensed and sold separately. For more information about RecoverPoint and its capabilities see the *Dell EMC RecoverPoint Product Guide*.

Remote replication using eNAS

File Auto Recovery (FAR) allows you to manually failover or move a virtual Data Mover (VDM) from a source eNAS system to a destination eNAS system. The failover or move leverages block-level SRDF synchronous replication, so it incurs zero data loss in the event of an unplanned operation. This feature consolidates VDMs, file systems, file system checkpoint schedules, CIFS servers, networking, and VDM configurations into their own separate pools. This feature works for a recovery where the source is unavailable. For recovery support in the event of an unplanned failover, there is an option to recover and clean up the source system and make it ready as a future destination.

Cloud Mobility

This chapter introduces cloud mobility.

Topics:

- [Cloud Mobility for Dell EMC PowerMax](#)

Cloud Mobility for Dell EMC PowerMax

Cloud Mobility is configured within an embedded guest running on the PowerMaxOS hypervisor.

Management of Cloud Mobility is performed using the Embedded Management (eManagement) Unisphere for PowerMax. Communication between the embedded Unisphere and Cloud Mobility is through REST API over a PowerMax internal private network connection.

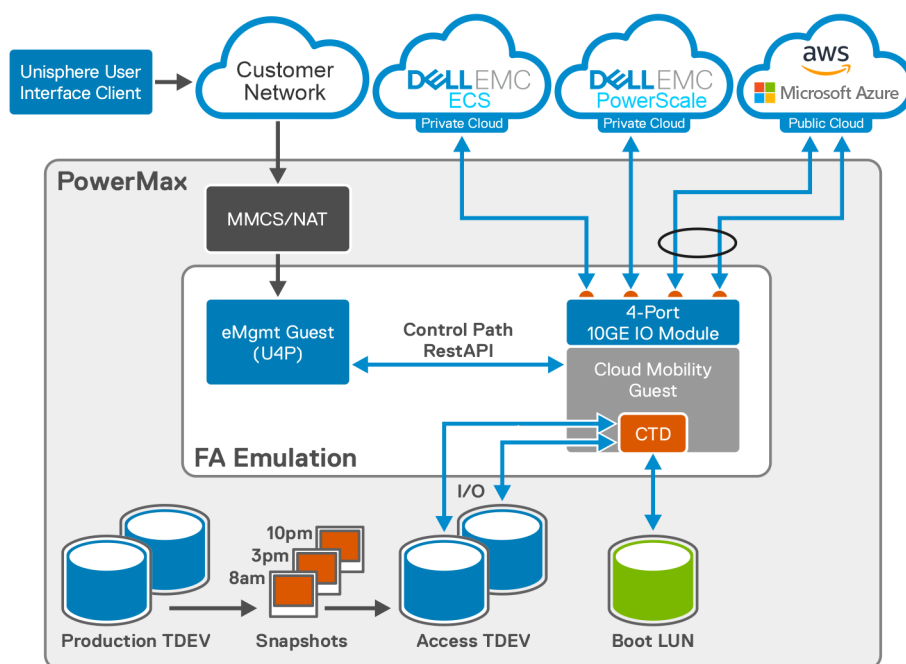


Figure 19. Cloud Mobility for PowerMax - high level architecture

Management of Cloud Mobility is performed on the Cloud Mobility Dashboard in Unisphere for PowerMax. Through the dashboard you can view cloud system alerts, configure and manage cloud snapshot policies, and view performance metrics for the cloud provider, as well as other operations. For details see the *Unisphere Online Help*.

Cloud snapshot policies

Cloud Mobility snapshot policies can be managed using Unisphere for PowerMax and REST API, and can be created and modified with Embedded Management Unisphere. External Unisphere instances can add and remove storage groups from existing cloud snapshot policies.

You can assign up to four policies to each storage group (SG). There is no limit to the number of SGs that can be assigned to a single policy.

PowerMax Cloud Mobility functionality allows you to move snapshots off the storage system and on to the cloud. Snapshots can also be restored back to the original storage system.

For more information see the *Cloud Mobility for Dell EMC PowerMax White Paper*.

Blended local and remote replication

This chapter introduces TimeFinder integration with SRDF.

Topics:

- Integration of SRDF and TimeFinder
- R1 and R2 devices in TimeFinder operations
- SRDF/AR
- TimeFinder and SRDF/A
- TimeFinder and SRDF/S

Integration of SRDF and TimeFinder

You can use TimeFinder and SRDF products to complement each other when you require both local and remote replication. For example, you can use TimeFinder to create local gold copies of SRDF devices for recovery operations and for testing disaster recovery solutions.

The key benefits of TimeFinder integration with SRDF include:

- Remote controls simplify automation—Use Dell EMC host-based control software to transfer commands across the SRDF links. A single command from the host to the primary array can initiate TimeFinder operations on both the primary and secondary arrays.
- Consistent data images across multiple devices and arrays—SRDF/CG guarantees that a dependent-write consistent image of production data on the R1 devices is replicated across the SRDF links.

You can use TimeFinder/CG in an SRDF configuration to create dependent-write consistent local and remote images of production data across multiple devices and arrays.

NOTE: Using a SRDF/A single session guarantees dependent-write consistency across the SRDF links and does not require SRDF/CG. SRDF/A MSC mode requires host software to manage consistency among multiple sessions.

NOTE: Some TimeFinder operations are not supported on devices that SRDF protects. The *Dell EMC Solutions Enabler TimeFinder SnapVX CLI User Guide* has further information.

The rest of this chapter summarizes the ways of integrating SRDF and TimeFinder.

R1 and R2 devices in TimeFinder operations

You can use TimeFinder to create local replicas of R1 and R2 devices. The following rules apply:

- You can use R1 devices and R2 devices as TimeFinder source devices.
- R1 devices can be the target of TimeFinder operations as long as there is no host accessing the R1 during the operation.
- R2 devices can be used as TimeFinder target devices if SRDF replication is not active (writing to the R2 device). To use R2 devices as TimeFinder target devices, first suspend the SRDF replication session.

SRDF/AR

SRDF/AR combines SRDF and TimeFinder to provide a long-distance disaster restart solution. SRDF/AR can be deployed over 2 or 3 sites:

- In 2-site configurations, SRDF/DM is deployed with TimeFinder.
- In 3-site configurations, SRDF/DM is deployed with a combination of SRDF/S and TimeFinder.

The time to create the new replicated consistent image is determined by the time that it takes to replicate the deltas.

SRDF/AR 2-site configurations

The following image shows a 2-site configuration where the production device (R1) on the primary array (Site A) is also a TimeFinder target device:

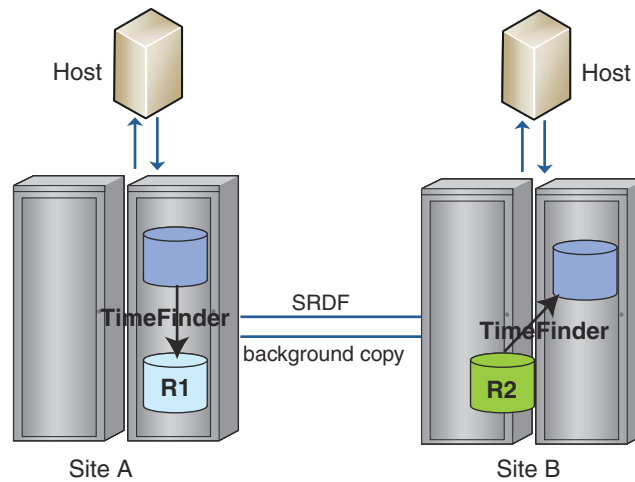


Figure 20. SRDF/AR 2-site solution

In this configuration, data on the SRDF R1/TimeFinder target device is replicated across the SRDF links to the SRDF R2 device.

The SRDF R2 device is also a TimeFinder source device. TimeFinder replicates this device to a TimeFinder target device. You can map the TimeFinder target device to the host connected to the secondary array at Site B.

In a 2-site configuration, SRDF operations are independent of production processing on both the primary and secondary arrays. You can utilize resources at the secondary site without interrupting SRDF operations.

Use SRDF/AR 2-site configurations to:

- Reduce required network bandwidth using incremental resynchronization between the SRDF target sites.
- Reduce network cost and improve resynchronization time for long-distance SRDF implementations.

SRDF/AR 3-site configurations

SRDF/AR 3-site configurations provide a zero data loss solution at long distances in the event that the primary site is lost.

The following image shows a 3-site configuration where:

- Site A and Site B are connected using SRDF in synchronous mode.
- Site B and Site C are connected using SRDF in adaptive copy mode.

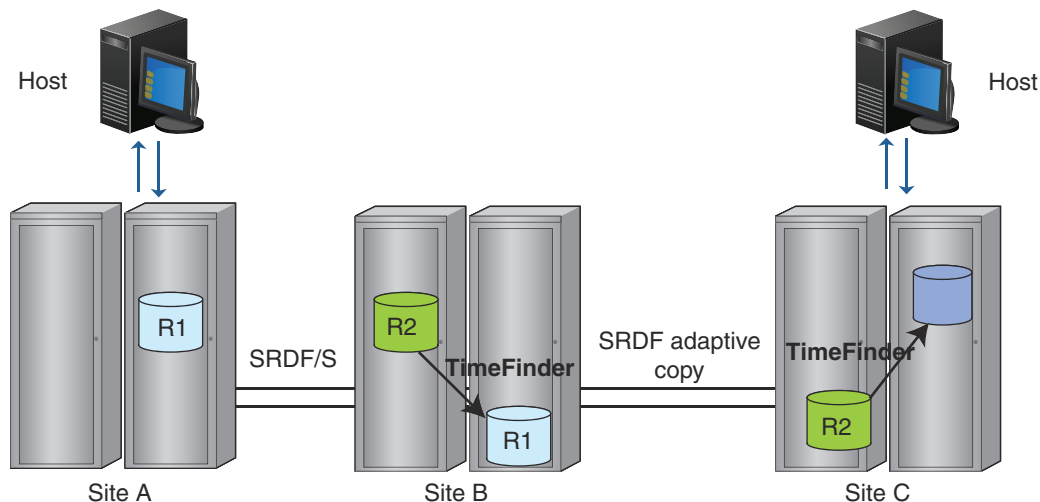


Figure 21. SRDF/AR 3-site solution

If Site A (primary site) fails, the R2 device at Site B provides a restartable copy with zero data loss. Site C provides an asynchronous restartable copy.

If both Site A and Site B fail, the device at Site C provides a restartable copy with controlled data loss. The amount of data loss is a function of the replication cycle time between Site B and Site C.

SRDF and TimeFinder control commands to R1 and R2 devices for all sites can be issued from Site A. No controlling host is required at Site B.

Use SRDF/AR 3-site configurations to:

- Reduce required network bandwidth using incremental resynchronization between the secondary SRDF target site and the tertiary SRDF target site.
- Reduce network cost and improve resynchronization time for long-distance SRDF implementations.
- Provide disaster recovery testing, point-in-time backups, decision support operations, third-party software testing, and application upgrade testing or the testing of new applications.

Requirements/restrictions

In a 3-site SRDF/AR multi-hop configuration, SRDF/S host I/O to Site A is not acknowledged until Site B has acknowledged it. This can cause a delay in host response time.

TimeFinder and SRDF/A

In SRDF/A solutions, device pacing:

- Prevents cache utilization bottlenecks when the SRDF/A R2 devices are also TimeFinder source devices.
- Allows R2 or R22 devices at the middle hop to be used as TimeFinder source devices.

NOTE: Device write pacing is not required in configurations that include PowerMaxOS 5978 and Engenuity 5876.

TimeFinder and SRDF/S

SRDF/S solutions support any type of TimeFinder copy sessions running on R1 and R2 devices as long as the conditions described in [R1 and R2 devices in TimeFinder operations](#) on page 93 are met.

Data migration

This chapter introduces data migration solutions.

Topics:

- [Overview](#)
- [Data migration for open systems](#)
- [Data migration for IBM System i](#)
- [Data migration for mainframe](#)

Overview

Data migration is a one-time movement of data from one array (the source) to another array (the target). Typical examples are data center refreshes where data is moved from an old array after which that array is retired or re-purposed. Data migration is *not* data movement due to replication (where the source data is accessible after the target is created) or data mobility (where the target is continually updated).

After a data migration operation, applications that access the data reference it at the new location.

To plan a data migration, consider the potential impact on your business, including the:

- Type of data to be migrated
- Site location(s)
- Number of systems and applications
- Amount of data to be moved
- Business needs and schedules

PowerMaxOS provides migration facilities for:

- Open systems
- IBM System i
- Mainframe

Data migration for open systems

The data migration features available for open system environments are:

- Non-disruptive migration
- Open Replicator
- PowerPath Migration Enabler
- Data migration using SRDF/Data Mobility
- Space and zero-space reclamation

Non-Disruptive Migration

Non-Disruptive Migration (NDM) is a method for migrating data without application downtime. The migration takes place over a metro distance, typically within a data center.

NDM Updates is a variant of NDM introduced in PowerMaxOS 5978.444.444. NDM Updates requires that the application associated with the migrated data is shut down for part of the migration process. This is due to the fact that the NDM is heavily dependent on the behavior of multipathing software to detect, enable, and disable paths none of which is under the control of Dell EMC (except for supported products such as PowerPath). NDM is the term that covers both non-disruptive and disruptive migration.

Starting with PowerMaxOS 5978 there are two implementations of NDM each for a different type of source array:

- Either:
 - PowerMax array running PowerMaxOS 5978
 - VMAX3 or VMAX All Flash array running HYPERMAX OS 5977.1125.1125 or later with an ePack
- VMAX array running Enginuity 5876 with an ePack

When migrating to a PowerMax array, these are the only configurations for the source array.

The *SRDF Interfamily Connectivity Information* lists the Service Packs and ePacks required for HYPERMAX OS 5977 and Enginuity 5876. In addition, the NDM support matrix has information on array operating systems support, host support, and multipathing support for NDM operations. The support matrix is available on the eLab Navigator.

Regulatory or business requirements for disaster recovery may require the use of replication to other arrays attached to source array, the target array, or both using SRDF/S, during the migration. In this case, refer to the *SRDF Interfamily Connectivity Information* for information on the Service Packs and the ePacks required for the SRDF/S configuration.

Migration from a VMAX3, VMAX All Flash or PowerMax array

Migrating from a VMAX3, VMAX All Flash or PowerMax array uses a modified form of SRDF/Metro. This means that in the normal workflow, both the source and target arrays are visible to the application host while the migration takes place. Indeed, both arrays are read/write accessible to the host. The following picture shows the logical structure of a migration from VMAX3, VMAX All Flash or PowerMax including the connections required.

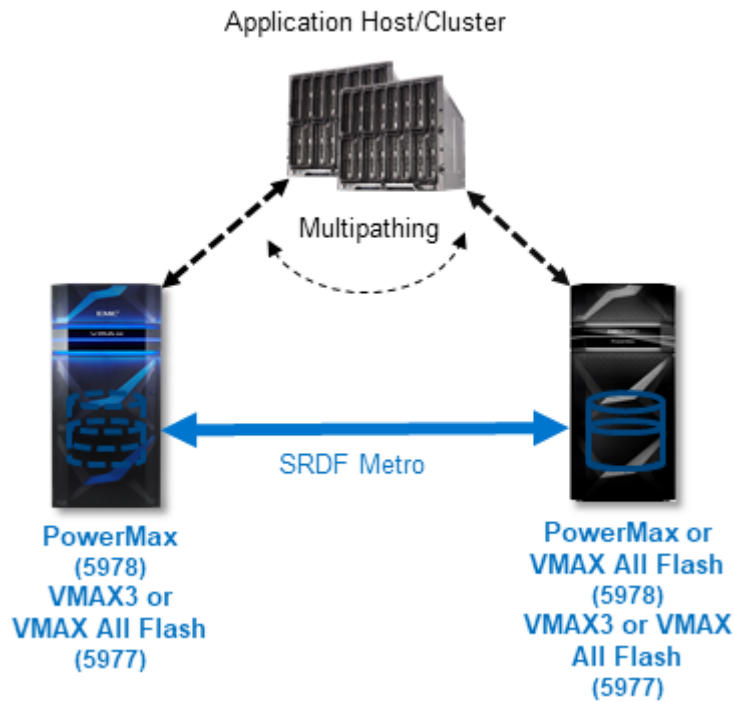


Figure 22. Configuration of a VMAX3, VMAX All Flash or PowerMax migration

Process

Normal flow

The steps in the migration process that is normally followed are:

1. Set up the migration environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.
3. If using NDM Updates, shut down the application associated with the storage group.
4. Create the migration session optionally specifying whether to move the identity of the LUNs in the storage group to the target array – copy the content of the storage group to the target array using SRDF/Metro.
During this time the source and target arrays are both accessible to the application host.
5. When the data copy is complete:
 - a. If the migration session did not move the identity of the LUNs, reconfigure the application to access the new LUNs on the target array.
 - b. Commit the migration session – remove resources from the source array and those used in the migration itself.
6. If using NDM Updates, restart the application.
7. To migrate further storage groups, repeat steps 2 on page 99 to 6 on page 99.
8. After migrating all the required storage groups, remove the migration environment.

Alternate flow

There is an alternative process that pre-copies the data to the target array before making it available to the application host. The steps in this process are:

1. Set up the migration environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.
3. Use the precopy facility of NDM to copy the selected data to the target array.
Optionally, specify whether to move the identity of the LUNs in the storage group to the target array.
While the data copy takes place, the source array is available to the application host, but the target array is unavailable.
4. When the copying of the data is complete: use the Ready Target facility in NDM to make the target array available to the application host also.
 - a. If the migration session did not move the identity of the LUNs, reconfigure the application to access the new LUNs on the target array.
 - b. If using NDM Updates, restart the application.
 - c. Commit the migration session – remove resources from the source array and those used in the migration itself. The application now uses the target array only.
5. To migrate further storage groups, repeat steps 2 on page 100 to 4 on page 100.
6. After migrating all the required storage groups, remove the migration environment.

Other functions

Other NDM facilities that are available for exceptional circumstances are:

- Cancel – to cancel a migration that has not yet been committed.
- Sync – to stop or start the synchronization of writes to the target array back to source array. When stopped, the application runs on the target array only. Used for testing.
- Recover – to recover a migration process following an error.

Other features

Other features of migrating from VMAX3, VMAX All Flash, or PowerMax to PowerMax are:

- Data can be compressed during migration to the PowerMax array
- Allows for nondisruptive revert to the source array
- There can be up to 50 migration sessions in progress simultaneously
- Does not require an additional license as NDM is part of PowerMaxOS
- The connections between the application host and the arrays use FC; the SRDF connection between the arrays uses FC or GigE

Devices and components that cannot be part of an NDM process are:

- CKD devices
- eNAS data
- Storage Direct and FAST.X relationships along with their associated data

Migration from a VMAX array

Migrating from a VMAX array uses SRDF technology. For NDM purposes, the source is a VMAX array running Enginuity 5876, with an ePack. The target is a PowerMax array running PowerMaxOS 5978. The following picture shows the logical structure of a migration from VMAX including the connections required:

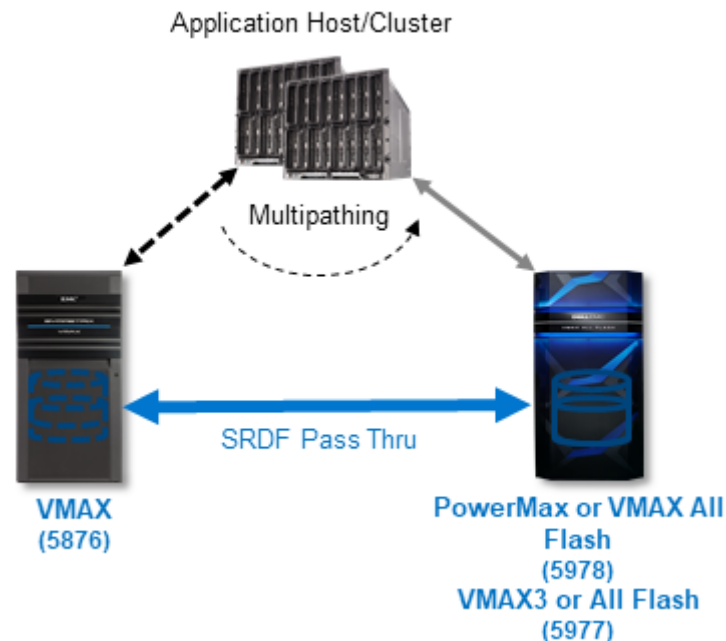


Figure 23. Configuration of a VMAX migration

Process

The steps in the migration process are:

1. Set up the environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.
3. If using NDM Updates, shut down the application associated with the storage group.
4. Create the migration session – copy the content of the storage group to the target array using SRDF.

When creating the session, optionally specify whether to move the identity of the LUNs in the storage group to the target array.

5. When the data copy is complete:
 - a. If the migration session did not move the identity of the LUNs, reconfigure the application to access the new LUNs on the target array.
 - b. Cutover the storage group to the PowerMax array.
 - c. Commit the migration session – remove resources from the source array and those used in the migration itself. The application now uses the target array only.
6. If using NDM Updates, restart the application.
7. To migrate further storage groups, repeat steps 2 on page 101 to 6 on page 101.
8. After migrating all the required storage groups, remove the migration environment.

Other features

Other features of migrating from VMAX to PowerMax are:

- Data can be compressed during migration to the PowerMax array
- Allows for nondisruptive revert to the source array
- There can be up to 50 migration sessions in progress simultaneously
- NDM does not require an additional license as it is part of PowerMaxOS

- The connections between the application host and the arrays use FC; the SRDF connection between the arrays uses FC or GigE

Devices and components that cannot be part of an NDM process are:

- CKD devices
- eNAS data
- Storage Direct and FAST.X relationships along with their associated data

Environmental requirements for NDM

There are requirements associated with both arrays in a migration and the host system.

Storage arrays

- The eligible combinations of operating environments running on the source and target arrays are:

| Source | Targets |
|----------------------------|---|
| PowerMaxOS 5978.444.444 | PowerMaxOS 5978.444.444 |
| PowerMaxOS 5978.221.221 | PowerMaxOS 5978.444.444 |
| PowerMaxOS 5978.221.221 | PowerMaxOS 5978.221.221 |
| HYPERMAX OS 5977.1125.1125 | PowerMaxOS 5978.444.444 HYPERMAX OS 5977.1125.1125 |
| Enginuity 5876 | PowerMaxOS 5978 HYPERMAX OS 5977.1125.1125 |

- The source array is one of:
 - A PowerMax array running PowerMaxOS 5978.221.221 or later
 - A VMAX3 or VMAX All Flash array running HYPERMAX OS 5977.1125.1125
 - A VMAX array running Enginuity 5876

The source array may require a Service Pack or an ePack. The *SRDF Interfamily Connectivity Information* lists the required packs (if any).

- SRDF is used for data migration, so zoning of SRDF ports between the source and target arrays is required. An SRDF license is not required, as there is no charge for NDM.
- The NDM SRDF group requires a minimum of two paths on different directors for redundancy and fault tolerance. If more paths are found, up to eight paths are configured.
- If SRDF is not normally used in the migration environment, it may be necessary to install and configure RDF directors and ports on both the source and target arrays and physically configure SAN connectivity.

Management host

- Wherever possible, use a host system separate from the application host to initiate and control the migration (the control host).
- The control host requires visibility of and access to both the source and target arrays.

Pre-migration rules and restrictions for NDM

In addition to general configuration requirements of the migration environment, the following rules and restrictions apply before starting a migration:

- A storage group is the data container that is migrated, and the requirements that apply to the group and its devices are:
 - Storage groups must have masking views. All devices in the group on the source array must be visible only through a masking view. Each device must be mapped only to a port that is part of the masking view.
 - Multiple masking views on a storage group using the same initiator group are valid only when:

- Port groups on the target array exist for each masking view, and
 - Ports in the port group are selected
- A storage group must be a parent or stand-alone group. A child storage group with a masking view on the child group is not supported.
- If the selected storage group is a parent, its child groups are also migrated.
- The names of storage groups and their children (if any) must not exist on the target array.
- Gatekeeper devices in a storage group are not migrated.
- Devices cannot:
 - Have a mobility ID
 - Have a nonbirth identity, when the source array runs Enginuity 5876
 - Have the BCV attribute
 - Be encapsulated
 - Be RP devices
 - Be Data Domain devices
 - Be vVOL devices
 - Be R2 or Concurrent SRDF devices
 - Be masked to FCoE (in the case of source arrays), iSCSI, non-ACLX, or NVMe over FC ports
 - Be part of another data migration operation
 - Be part of an ORS relationship
 - Be in other masked storage groups
 - Have a device status of Not Ready
- Devices can be part of TimeFinder sessions.
- Devices can act as R1 devices but cannot be part of a SRDF/Star or SRDF/SQAR configuration.
- The names of masking groups to migrate must not exist on the target array.
- The names of initiator groups to migrate may exist on the target array. However, the aggregate set of host initiators in the initiator groups that the masking groups use must be the same. Also, the effective ports flags on the host initiators must have the same setting on both arrays.
- The names of port groups to migrate may exist on the target array, as long as the groups on the target array are in the logging history table for at least one port.
- The status of the target array must be as follows:
 - If a target-side Storage Resource Pool (SRP) is specified for the migration, that SRP must exist on the target array.
 - The SRP to be used for target-side storage must have enough free capacity to support the migration.
 - The target side must be able to support the additional devices required to receive the source-side data.
 - All initiators provisioned to an application on the source array must also be logged into ports on the target array.

Migration infrastructure - RDF device pairing

RDF device pairing is done during the create operation, with the following actions occurring on the device pairs.

- NDM creates RDF device pairs, in a DM RDF group, between devices on the source array and the devices on the target array.
- Once device pairing is complete NDM controls the data flow between both sides of the migration process.
- Once the migration is complete, the RDF pairs are deleted when the migration is committed.
- Other RDF pairs may exist in the DM RDF group if another migration is still in progress.

Due to differences in device attributes between the source and target array, the following rules apply during migration:

- Any source array device that has an odd number of cylinders is migrated to a device on the target array that has Geometry Compatibility Mode (GCM).
- Any source array meta device is migrated to a non-meta device on the target array.

Once the copying of data to the target array has begun, the target devices can have SRDF mirrors (R2 devices) added to them for remote replication. However, the mirror devices cannot be:

- Enabled for MSC or Synchronous SRDF Consistency
- Part of a SRDF/Star, SRDF/SQAR, or SRDF/Metro configuration

Rules and restrictions while the migration is in progress

There are rules and restrictions that apply from the time the migration starts until the Commit operation completes:

- The source and target masking of the application that is migrating cannot be changed, except for:
 - Changing the service levels or SRPs on the storage groups
 - Changing the compression attribute on the storage groups
 - Changing Host I/O limits on the storage groups
 - Adding ports to port groups
- Once the Cutover stage in a migration is complete, the devices on the target array can have TimeFinder sessions added to them.
- Source or target devices with TimeFinder sessions cannot be the target of a data copy operation during the migration session. For example, TimeFinder sessions cannot copy data to the source devices being migrated.
- Source or target devices in a Storage Direct session cannot be the target of a rollback or restore operation during the migration session.
- Once the copying of data to the target array has begun, the target devices can have SRDF mirrors (R2 devices) added to them for remote replication. However, the mirror devices cannot be:
 - Enabled for MSC or Synchronous SRDF Consistency
 - Part of a SRDF/Star, SRDF/SQAR, or SRDF/Metro configuration
- The source and target devices cannot be part of an ORS relationship.


Open Replicator

Open Replicator enables copying data (full or incremental copies) from qualified arrays within a storage area network (SAN) infrastructure to or from arrays running PowerMaxOS. Open Replicator uses the Solutions Enabler SYMCLI `symrcopy` command.

Use Open Replicator to migrate and back up/archive existing data between arrays running PowerMaxOS. Open Replicator uses the Solutions Enabler SYMCLI `symrcopy` and third-party storage arrays within the SAN infrastructure without interfering with host applications and ongoing business operations.

Use Open Replicator to:

- Pull from source volumes on qualified remote arrays to a volume on an array running PowerMaxOS. Open Replicator uses the Solutions Enabler SYMCLI `symrcopy`.
- Perform online data migrations from qualified storage to an array running PowerMaxOS. Open Replicator uses the Solutions Enabler SYMCLI `symrcopy` with minimal disruption to host applications.

 **NOTE:** Open Replicator cannot copy a volume that is in use by TimeFinder.

Open Replicator operations

Open Replicator uses the following terminology:

Control


The recipient array and its devices are referred to as the control side of the copy operation.

Remote

The donor Dell EMC arrays or third-party arrays on the SAN are referred to as the remote array/devices.

Hot

The Control device is Read/Write online to the host while the copy operation is in progress.

 **NOTE:** Hot push operations are not supported on arrays running PowerMaxOS. Open Replicator uses the Solutions Enabler SYMCLI `symrcopy`.

Cold

The Control device is Not Ready (offline) to the host while the copy operation is in progress.

Pull

A pull operation copies data to the control device from the remote device(s).

Push

A push operation copies data from the control device to the remote device(s).

Pull operations

On arrays running PowerMaxOS, Open Replicator uses the Solutions Enabler SYMCLI `symrcopy` support for up to 4096 pull sessions.

For pull operations, the volume can be in a live state during the copy process. The local hosts and applications can begin to access the data as soon as the session begins, even before the data copy process has completed.

These features enable rapid and efficient restoration of remotely vaulted volumes and migration from other storage platforms.

Copy on First Access ensures the appropriate data is available to a host operation when it is needed. The following image shows an Open Replicator hot pull.




Figure 24. Open Replicator hot (or live) pull

The pull can also be performed in cold mode to a static volume. The following image shows an Open Replicator cold pull.




Figure 25. Open Replicator cold (or point-in-time) pull

Disaster Recovery

When the control array runs PowerMaxOS it can also be the R1 side of a SRDF configuration. That configuration can use SRDF/A, SRDF/S, or Adaptive Copy Mode to provide data protection during and after the data migration.

PowerPath Migration Enabler

Dell EMC PowerPath is host-based software that provides automated data path management and load-balancing capabilities for heterogeneous server, network, and storage deployed in physical and virtual environments. PowerPath includes a migration tool called PowerPath Migration Enabler (PPME). PPME enables non-disruptive or minimally disruptive data migration between storage systems or within a single storage system.

PPME allows applications continued data access throughout the migration process. PPME integrates with other technologies to minimize or eliminate application downtime during data migration.

PPME works in conjunction with underlying technologies, such as Open Replicator, SnapVX, and Host Copy.

NOTE: PowerPath Multipathing must be installed on the host machine.

The following documentation provides additional information:

- *Dell EMC Support Matrix PowerPath Family Protocol Support*
- *Dell EMC PowerPath Migration Enabler User Guide*

Data migration using SRDF/Data Mobility

SRDF/Data Mobility (DM) uses SRDF's adaptive copy mode to transfer large amounts of data without impact to the host.

SRDF/DM supports data replication or migration between two or more arrays running PowerMaxOS. Adaptive copy mode enables applications using the primary volume to avoid propagation delays while data is transferred to the remote site. SRDF/DM can be used for local or remote transfers.

[Data migration](#) on page 83 has a more information about using SRDF to migrate data.

Space and zero-space reclamation

Space reclamation reclaims unused space following a replication or migration activity from a regular device to a thin device in which software tools, such as Open Replicator and Open Migrator, copied-all-zero, unused space to a target thin volume.

Space reclamation deallocates data chunks that contain all zeros. Space reclamation is most effective for migrations from standard, fully provisioned devices to thin devices. Space reclamation is non-disruptive and can be executed while the targeted thin device is fully available to operating systems and applications.

Zero-space reclamations provides instant zero detection during Open Replicator and SRDF migration operations by reclaiming all-zero space, including both host-unwritten extents (or chunks) and chunks that contain all zeros due to file system or database formatting.

Solutions Enabler and Unisphere can be used to initiate and monitor the space reclamation process.

Data migration for IBM System i

NDM is also available for IBM i systems. The process of migrating is the same as [Non-Disruptive Migration](#) on page 98 explains, but with a few differences:

- A separate, open-systems host is necessary to manage and control the migration process, using Solutions Enabler or Unisphere. It is not possible to run the migration directly from an IBM System i.
- Migration is available for D910 devices only.

The migration can take place between these source and target systems:

- PowerMaxOS 5978 to PowerMaxOS 5978
- HYPERMAX OS 5977 to PowerMaxOS 5978
- Enginuity 5876 to PowerMaxOS 5978
- Enginuity 5876 to HYPERMAX OS 5977

Data migration for mainframe

For mainframe environments, z/OS Migrator provides non-disruptive migration from any vendor storage to PowerMax, VMAX All Flash, or VMAX arrays. z/OS Migrator can also migrate data from one PowerMax, VMAX All Flash, or VMAX array to another PowerMax, VMAX All Flash, or VMAX array. With z/OS Migrator, you can:

- Introduce new storage subsystem technologies with minimal disruption of service.
- Reclaim z/OS UCBs by simplifying the migration of datasets to larger volumes (combining volumes).
- Facilitate data migration while applications continue to run and fully access data being migrated, eliminating application downtime usually required when migrating data.
- Eliminate the need to coordinate application downtime across the business, and eliminate the costly impact of such downtime on the business.
- Improve application performance by facilitating the relocation of poor performing datasets to lesser used volumes/storage arrays.
- Ensure all metadata always accurately reflects the location and status of datasets being migrated.

Refer to the *Dell EMC z/OS Migrator Product Guide* for detailed product information.

Volume migration using z/OS Migrator

Dell EMC z/OS Migrator is a host-based data migration facility that performs traditional volume migrations as well as host-based volume mirroring. Together, these capabilities are referred to as the volume mirror and migrator functions of z/OS Migrator.




Figure 26. z/OS volume migration

Volume level data migration facilities move logical volumes in their entirety. z/OS Migrator volume migration is performed on a track for track basis without regard to the logical contents of the volumes involved. Volume migrations end in a volume swap which is entirely non-disruptive to any applications using the data on the volumes.

Volume migrator

Volume migration provides host-based services for data migration at the volume level on mainframe systems. It provides migration from third-party devices to devices on Dell EMC arrays as well as migration between devices on Dell EMC arrays.

Volume mirror

Volume mirroring provides mainframe installations with volume-level mirroring from one device on a Dell EMC array to another. It uses host resources (UCBs, CPU, and channels) to monitor channel programs scheduled to write to a specified primary volume and clones them to also write to a specified target volume (called a mirror volume).

After achieving a state of synchronization between the primary and mirror volumes, Volume Mirror maintains the volumes in a fully synchronized state indefinitely, unless interrupted by an operator command or by an I/O failure to a Volume Mirror device. Mirroring is controlled by the volume group. Mirroring may be suspended consistently for all volumes in the group.

Dataset migration using z/OS Migrator

In addition to volume migration, z/OS Migrator provides for logical migration, that is, the migration of individual datasets. In contrast to volume migration functions, z/OS Migrator performs dataset migrations with full awareness of the contents of the volume, and the metadata in the z/OS system that describe the datasets on the logical volume.




Figure 27. z/OS Migrator dataset migration

Thousands of datasets can either be selected individually or wild-carded. z/OS Migrator automatically manages all metadata during the migration process while applications continue to run.

Online Device Expansion

Online device expansion (ODE) is a mechanism to increase the capacity of a device without taking it offline. This is an overview of the ODE capabilities:

Topics:

- [Introduction](#)
- [General features](#)
- [Standalone devices](#)
- [SRDF devices](#)
- [LREP devices](#)
- [Management facilities](#)

Introduction

ODE enables a storage administrator to provide more capacity on a storage device while it remains online to its application. This particularly benefits organizations where applications need to remain available permanently. If a device associated with an application runs low on space, the administrator can increase its capacity without affecting the availability and performance of the application.

Standalone devices, devices in a SRDF configuration and those in an LREP configuration can all be expanded using ODE.

General features

Features of ODE that are applicable to stand-alone, SRDF, and LREP devices are:

- ODE is available for both FBA and CKD devices.
- ODE operates on thin devices (TDEVs).
- A device can be expanded to a maximum capacity of 64 TB (1,182,006 cylinders for a CKD device).
- A device can expand only.

There are no facilities for reducing the capacity of a device.

- During expansion, a device is locked.

This prevents operations such as adding a device to an SRDF configuration until the expansion is complete.

- An administrator can expand the capacity of multiple devices using one management operation.

A thin device presents a given capacity to the host, but consumes only the physical storage necessary to hold the data that the host has written to the device ([Thin devices \(TDEVs\)](#) on page 57 has more information). Increasing the capacity of a device using ODE does not allocate any additional physical storage. Only the configured capacity of the device as seen by the host increases.

Failure of an expansion operation for a stand-alone, SRDF, or LREP device may occur because:

- The device does not exist.
- The device is not a TDEV.
- The requested capacity is less than the current capacity.
- The requested capacity is greater than 64 TB.
- There is insufficient space in the storage system for expansion.
- There are insufficient PowerMax internal resources to accommodate the expanded device.
- Expanding the device to the requested capacity would exceed the oversubscription ratio of the physical storage.
- A reclaim, deallocation, or free-all operation is in progress on the device.

There are other reasons specific to each type of device. These are listed in the description of device expansion for that type of device.

Standalone devices

The most basic form of device expansion is of a device that is associated with a host application and is not part of a SRDF or LREP configuration. Additional features of ODE in this environment are:

- ODE can expand vVols in addition to TDEVs.
vVols are treated as a special type of TDEV.
- ODE for a standalone device is available in PowerMaxOS 5978, HYPERMAX OS 5977.691.684 or later (for FBA devices), and HYPERMAX OS 5977.1125.1125 or later (for CKD devices).

Each expansion operation returns a status that indicates whether the operation succeeded or not. The status of an operation to expand multiple devices can indicate a partial success. In this case at least one of the devices was successfully expanded but one or more others failed.

Another reason why an expansion operation might fail is if the device is not a vVol.

SRDF devices


PowerMaxOS 5978 introduces online device expansion for SRDF configurations. The administrator can expand the capacity of thin devices in an SRDF relationship without any service disruption in a similar way to expanding stand-alone devices.

Devices in an asynchronous, synchronous, adaptive copy mode, SRDF/Metro, SRDF/Star (mainframe only), or SRDF/SQAR (mainframe only) configuration are all eligible for expansion. However, this feature is not available in RecoverPoint, Storage Direct, NDM, or NDM Updates configurations.

Also, device expansion is available only on storage arrays in an SRDF configuration that run PowerMaxOS (PowerMaxOS 5978.444.444 or later for SRDF/Metro) on both sides. Any attempt to expand an SRDF device in a system that runs an older operating environment fails.

Other features of ODE in an SRDF environment are for expanding:

- An individual device on either the R1 or R2 side
- An R1 device and its corresponding device on the R2 side in one operation
- A range of devices on either the R1 or R2 side
- A range of devices on the R1 side and their corresponding devices on the R2 side in one operation
- A storage group on either the R1 or R2 side
- A storage group on the R1 side and its corresponding group on the R2 side in one operation

 **NOTE:** An SRDF/Metro configuration does not allow the expansion of devices on one side only. Both sides, whether it is a device, a range of devices, or a storage group, must be expanded in one operation.

Basic rules of device expansion are:

- The R1 side of an SRDF pair cannot be larger than the R2 side.
- In an SRDF/Metro configuration, both sides must be the same size.

When both sides are available on the SRDF link, Solutions Enabler, Mainframe Enablers, and Unisphere (the tools for managing ODE) enforce these rules. When either device is not available on the SRDF link, the management tools allow you to make the R1 larger than the R2. However, before the devices can be made available on the link, the capacity of the R2 must increase to at least the capacity of the R1 device.

Similar considerations apply to multiple site configurations:

- **Cascaded SRDF:** The size of R1 must be less than or equal to the size of R21. The size of R21 must be less than or equal to the size of R2.
- **Concurrent SRDF:** The size of R11 must be less than or equal to the size of both R2 devices.

Other reasons why an expansion operation may fail in an SRDF environment are:

- One or more of the devices is on a storage system that does not run PowerMaxOS 5978 (or PowerMaxOS 5978.444.444 or later for SRDF/Metro).
- One or more of the devices is a vVol.
- One or more devices are part of a Storage Direct, RecoverPoint, NDM, or MDM configuration.

- The operation would result in an R1 device being larger than its R2 device.

LREP devices

PowerMaxOS 5978 also introduces online device expansion for LREP (local replication) configurations. As with stand-alone and SRDF devices, this means that an administrator can increase the capacity of thin devices that are part of an LREP relationship without any service disruption. Devices eligible for expansion are those that are part of:

- SnapVX sessions
- Legacy sessions that use CCOPY, SDDF, or Extent

ODE is not available for:

- SnapVX emulations such as Clone, TimeFinder Clone, TimeFinder/Mirror, TimeFinder/Snap, and VP Snap
- RecoverPoint and Storage Direct devices
- vVols
- PPRC

This is to maintain compatibility with the limitations that IBM place on expanding PPRC devices.

By extension, ODE is not available for a product that uses any of these technologies. For example, it is not available for Remote Pair FlashCopy since that uses PPRC.

Other ODE features in an LREP environment are:

- Expand Snap VX source or target devices.
- Snapshot data remains the same size.
- The ability to restore a smaller snapshot to an expanded source device.
- Target link and relink operations are dependent on the size of the source device when the snapshot was taken not its size after expansion.

There are additional reasons for an ODE operation to fail in an LREP environment. For instance, when the LREP configuration uses one of the excluded technologies.

Management facilities

Solutions Enabler, Unisphere, and Mainframe Enablers all provide facilities for managing ODE. With any of these tools you can:

- Expand a single device
- Expand multiple devices in one operation
- Expand both sides of an SRDF pair in one operation

Solutions Enabler

Use the `symdev modify` command in Solutions Enabler to expand one or more devices. Some features of this command are:

- Use the `-cap` option to specify the new capacity for the devices.

Use the `-captype` option with `-cap` to specify the units of the new capacity. The available units are cylinders, MB, GB, and TB.

- Use the `-devs` option to define the devices to expand. The argument for this option consists of a single device identifier, an range of device identifiers, or a list of identifiers. Each element in the list can be a single device identifier or a range of device identifiers.
- Use the `-rdfg` option to specify the SRDF group of the devices to be expanded. Inclusion of this option indicates that both sides of the SRDF pair associated with the group are to be expanded in a single operation.

The *Dell EMC Solutions Enabler Array Controls and Management CLI User Guide* has details of the `symdev modify` command, its syntax and its options.

Examples:

- Expand a single device on array 005 to a capacity of 4TB:

```
symdev modify lfe0 -sid 005 -cap 4 -captype tb -nop -tdev
```

- Expand four devices on SRDF group 33 and their corresponding R2 devices:

```
symdev modify -sid 85 -tdev -cap 1000 -captype mb -dev 007D2:007D5 -v -rdfg 33 -nop
```

Unisphere

Unisphere provides facilities to increase the capacity of a device, a range of devices (FBA only), and SRDF pairs. The available units for specifying the new device capacity are cylinders, GB, and TB. The Unisphere Online Help has details on how to select and expand devices.

For example, this is the dialog for expanding a standalone device:

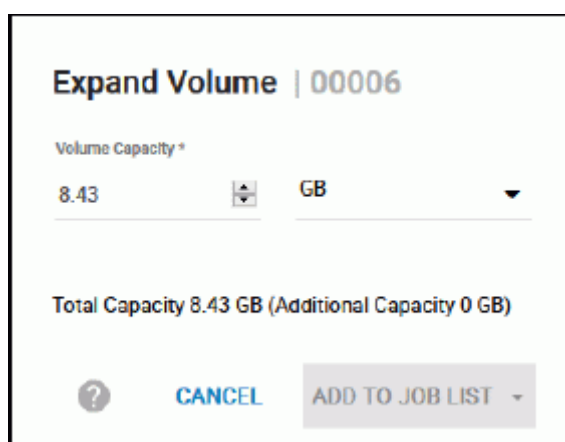
The image shows a screenshot of the 'Expand Volume' dialog box in Unisphere. The title bar reads 'Expand Volume | 00006'. Below the title, there is a section for 'Volume Capacity *' with a numeric input field containing '8.43', a unit dropdown menu currently set to 'GB', and a small expand/collapse icon. Below this, the text 'Total Capacity 8.43 GB (Additional Capacity 0 GB)' is displayed. At the bottom of the dialog, there are three buttons: a help button with a question mark icon, a 'CANCEL' button, and an 'ADD TO JOB LIST' button with a dropdown arrow.

Figure 28. Expand Volume dialog in Unisphere

Audit log

Storage system audit records come from the SYMAPI database and include all actions that are taken on that storage system.

The audit log resides on the storage system and has a maximum size of 40 MB. When the 40 MB limit is reached, the log begins to overwrite itself. In Unisphere release 9.2, operations that are run through Unisphere UI have been added to the audit log.

The audit log:

- Has advanced search and filter capabilities, such as:
 - Specific time period (for example, last 24 hours, last week, specific date)
 - Specific user (username, user role)
 - Operation type and category
 - Application type (Unisphere, CLI)
 - Hostname
 - Search phrase
- Records and the filtered audit log list can be exported to a .log file. This is facilitated with a REST endpoint.

Mainframe Enablers

Mainframe Enablers provides the `DEV,EXPAND` command in the Symmetrix Control Facility (SCF) to increase the capacity of a device. Some features of this command are:

- Use the `DEVIce` parameter to specify a single device or a range of devices to expand.
- Use the `CYLinders` parameter to specify the new capacity of the devices, in cylinders.
- Use the `RDFG` parameter to specify the SRDF group associated with the devices and so expand the R1 and R2 devices in a single operation.

The *Dell EMC Mainframe Enablers ResourcePak Base for z/OS Product Guide* has details of the `DEV,EXPAND` command and its parameters.

Example:

Expand device 8013 to 1150 cylinders:

```
DEV,EXPAND,DEV(8013),CYL(1150)
```

System security

This is an overview of some of the security features of PowerMaxOS. For more detailed information, see the *Dell EMC PowerMax Family Security Configuration Guide*.

Topics:

- [User authentication and authorization](#)
- [Roles and permissions](#)
- [Lockbox](#)
- [Client/server communications](#)
- [Secure device deletion](#)

User authentication and authorization

Access to an array's management functions must be available only to those users who have responsibility for administering the array in one way or another.

All users must identify themselves when they want to access an array. That is, they go through an authentication process.

Once authenticated, the management function grants one or more permissions to the user that define what the user can do on the system. That is, what the user is authorized to do.

Each management function has its own way of user authentication. For instance:

- Solutions Enabler maintains a list of host usernames associated with users who have access to the SYMAPI.
- Unisphere requires a user to log in using a separate username and password.

In all cases, roles and permissions define what the user is authorized to do.

Roles and permissions

Role-based access controls (RBAC) are central to security in a PowerMaxOS system. Each host in the storage system has an access ID attached to it. Host access information can be modified by attaching or detaching the host access ID, or by renaming the host. In RBAC, a user can be assigned one or more roles. In turn, a role consists of a number of permissions that define what the user can do.

For added refinement, some roles can be restricted to one or more storage groups. In this case, the user can carry out the operations that the permissions grant on these storage groups only. For example, a user has responsibility for managing storage for a particular application. The roles associated with the user are restricted to the storage groups for that application. The user has no access to any other storage group in the array.

A user can have up to four roles. The scope of the roles are independent of each other. Each of a user's roles can be associated with any set of storage groups or have array-wide effect.

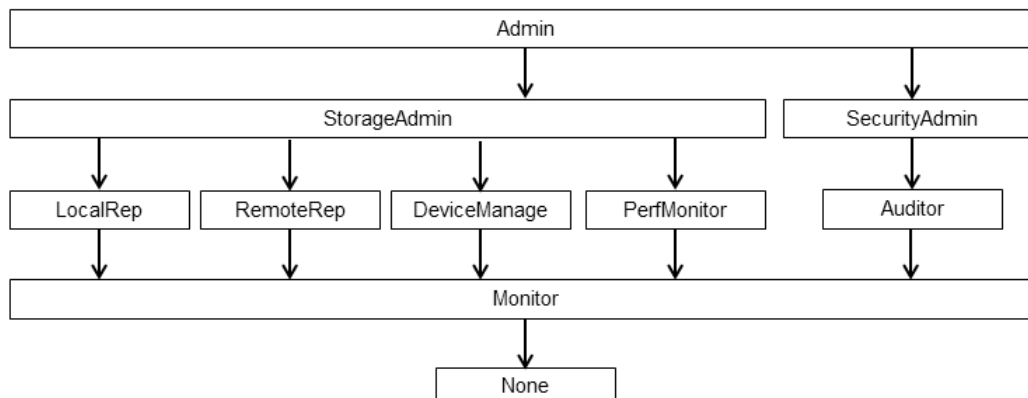
Roles and their hierarchy

There are nine user roles:

- Monitor
- PerfMonitor (performance monitor)
- Auditor
- DeviceManage (device manager)
- RemoteRep (remote replication)
- LocalRep (local replication)
- StorageAdmin (storage administrator)

- SecurityAdmin (security administrator)
- Admin (system administrator)

These roles form a hierarchy:



The higher a role is in the hierarchy the more permissions, and hence capabilities, it has.

Permissions for roles

The permissions associated with each role define what a user with that role can and cannot do on a storage system.

Monitor

The Monitor role allows a user to use show, list, and view operations to monitor a system.

Allowed operations

Examples of the operations that the Monitor role allows are:

- View array information
- View masking objects (storage groups, initiator groups, port groups, and masking views)
- View device information
- View the RBAC rules defined on this array.

This is available only when the Secure Reads policy is not in effect. [Secure Reads policy](#) on page 119 has more information the Secure Reads policy and its management.

Prevented operations

The Monitor role does not allow the user to view:

- Security-related data such as array ACLs and the array's Audit Log file
- The RBAC roles defined on this system, when the Secure Reads policy is in effect

PerfMonitor

The PerfMonitor role allows a user to configure performance alerts and thresholds in Unisphere. The PerfMonitor role also has the permissions of the Monitor role.

Auditor

The Auditor role allows a user to view the security settings on a system.

Allowed operations

Examples of operations that the Auditor role allows are:

- View the array's ACL settings
- View RBAC rules and settings
- View the array's Audit Log file

The Auditor role also has the permissions of the Monitor role.

Prevented operations

The Auditor role does not allow the user to modify any security setting.

DeviceManage

The DeviceManage role allows a user to configure and manage devices.

Allowed operations

Examples of operations that the DeviceManage role allows are:

- Control operations on devices, such as Ready, Not-Ready, Free
- Configuration operations on devices, such as setting names, or setting flags
- Link, Unlink, Relink, Set-Copy, and Set-NoCopy operations on SnapVX link devices
- Restore operations to SnapVX source devices

This is available only when the user also has the LocalRep role.

When the role is restricted to one or more storage groups, it allows these operations on the devices in those groups only.

The DeviceManage role also has the permissions of the Monitor role.

Prevented operations

The DeviceManage role does not allow the user to create, expand, or delete devices. However, if the role is associated with a storage group, those operations are allowed on the devices within the group.

LocalRep

The LocalRep role allows the user to carry out local replication using SnapVX, or the legacy operations of Snapshot, Clone, and BCV.

Allowed operations

Examples of operations that the LocalRep role allows are:

- Create, manage, and delete SnapVX snapshots

For operations that result in changes to the contents of any device, the user may also need the DeviceManage role:

- SnapVX restore operations require both the LocalRep and DeviceManage roles.
- SnapVX Link, Unlink, Relink, Set-Copy, and Set-No_Copy operations require the DeviceManage role on the link devices and the LocalRep role on the source devices.

When the role is restricted to one or more storage groups, it allows all these operation on the devices within those groups only.

The LocalRep role also has the permissions of the Monitor role.

Prevented operations

The LocalRep role does not allow the user to create Secure SnapVX snapshots.

RemoteRep

The RemoteRep role allows a user to carry out remote replication using SRDF.

Allowed operations

Examples of operations that the RemoteRep role allows are:

- Create, manage, and delete SRDF device pairs
When the role is restricted to storage groups, it allows these operations on devices within those groups only.
- Set attributes that are not associated with SRDF/A on a SRDF group
This is available only if the role is applied to the entire array.

When the role is restricted to one or more storage groups, it allows these operations on the devices in those groups only.

The RemoteRep role also has the permissions of the Monitor role.

Prevented operations

The RemoteRep role does not allow the user to:

- Create and delete SRDF groups
- Set attributes that are not associated with SRDF/A on a SRDF group when the role is restricted to a set of storage groups

StorageAdmin

The StorageAdmin role allows a user to perform any storage operation, except those related to security.

Allowed operations

Examples of operations that the StorageAdmin role allows are:

- Perform array configuration operations
- Provision storage
- Delete storage
- Create, modify, and delete masking objects (storage groups, initiator groups, port groups, and masking views)
- Create and delete Secure SnapVX Snapshots
- Any operation allowed for the LocalRep, RemoteRep, and DeviceManage roles

This role also has the permissions of the LocalRep, RemoteRep, DeviceManage, and Monitor roles.

SecurityAdmin

The SecurityAdmin role allows a user to view and modify the system security settings.

Allowed operations

Operations that the SecurityAdmin role allows are:

- Modify the array's ACL settings
- Modify the RBAC rules and settings

The SecurityAdmin role also has the permissions of the Auditor and Monitor roles.

Admin

The Admin role allows a user to carry out any operation on the array. It has the permissions of the StorageAdmin and SecurityAdmin roles.

Secure Reads policy

The Secure Reads policy determines whether all users can view all the RBAC roles defined on the array. The policy can be in force or not in force.

In force

Users with the Admin, SecurityAdmin, or Auditor roles can view all RBAC rules on the array. All other users can only see the rules that either apply to them, or that assign a role of Admin or SecurityAdmin to someone.

Not in force

All users, no matter what role they have, can view all RBAC rules in the array. This is the default setting for the policy.

Policy management

Both the Solutions Enabler SYMCLI and Unisphere provide facilities for controlling whether the policy is in force.

View permissions required for an operation

It can be difficult to know which roles, permissions, or ACL access types are required for any particular operation. PowerMaxOS can write the information to the Solutions Enabler log file about the facilities an operation required when it executed. You control this using an environment variable: SYMAPI_LOG_ACCESS_CHECKS.

Lockbox

Solutions Enabler uses a Lockbox to store and protect sensitive information. The Lockbox is associated with a particular host. This association prevents the Lockbox from being copied to a second host and used to obtain access.

The Lockbox is created at installation. During installation, the installer prompts the user to provide a password for the Lockbox, or if no password is provided at installation, a default password is generated and used with the Stable System values (SSVs, a fingerprint that uniquely identifies the host system). For more information about the default password, see [Default Lockbox password](#) on page 120.

Stable System Values (SSVs)

When Solutions Enabler is upgraded, values stored in the existing Lockbox are automatically copied to the new Lockbox.

Lockbox passwords

If you create the Lockbox using the default password during installation, change the password immediately after installation to best protect the contents in the Lockbox.

For maximum security, select a password that is hard to guess. It is very important to remember the password.

 **WARNING:** Loss of this password can lead to situations where the data stored in the Lockbox is unrecoverable. Dell EMC cannot recover a lost lockbox password.

Passwords must meet the following requirements:

- 8 - 256 characters in length
- Include at least one numeric character
- Include at least one uppercase and one lowercase character
- Include at least one of these non-alphanumeric characters: ! @ # % &

Lockbox passwords may include any character that can be typed in from US standard keyboard.

- The new password must not be the same as the previous password.

Default Lockbox password

When you install Solutions Enabler, you are asked whether you want to use the default password for the Lockbox. If you choose to use the default, the installation process establishes the default Lockbox password in the following format:


`nodename@SELockbox1`

where: *nodename* is the hostname of the computer on which you are installing.

Operating systems have different methods of determining the node name:

- UNIX: The installation program uses the `hostname` command to determine the node name. Normally, the node name is set in the `/etc/hosts` file.
- Windows: The value of the `COMPUTERNAME` system environment variable, converted to lower case.
- z/OS: The `gethostname()` function is used to get the node name of the machine.

If the value of *nodename* is stored in upper case letters, it is converted to lower case for the default password.

 **NOTE:** Dell Technologies strongly recommends that you change the default password. If you allow the installation program to use the default password, note it for future use. You need the password to reset the Lockbox Stable System values or generate or replace SSL certificates for client/server operations.


Client/server communications

All communications between client and hosts uses SSL to help ensure data security.

Secure device deletion

PowerMax provides secure TDEV (device) deletion.

The tracks (data) associated with each TDEV are stored in backend (TDAT) pools. After a TDEV is deleted, the tracks that were associated with that device cannot be read or accessed in any way regardless of where they are stored. For example, when old track space in the pools is reused (re-written), the new write data is padded with zeros to fill the track, preventing any leakage of old track data.

 **NOTE:** You cannot delete a TDEV that is part of a replication session (TF or SRDF).

Mainframe Error Reporting

This appendix lists the mainframe environmental errors.

Topics:

- [Error reporting to the mainframe host](#)
- [SIM severity reporting](#)

Error reporting to the mainframe host

PowerMaxOS can detect and report the following error types to the mainframe host in the storage systems:

- **Data Check** — PowerMaxOS detected an error in the bit pattern read from the disk. Data checks are due to hardware problems when writing or reading data, media defects, or random events.
- **System or Program Check** — PowerMaxOS rejected the command. This type of error is indicated to the processor and is always returned to the requesting program.
- **Overrun** — PowerMaxOS cannot receive data at the rate it is transmitted from the host. This error indicates a timing problem. Resubmitting the I/O operation usually corrects this error.
- **Equipment Check** — PowerMaxOS detected an error in hardware operation.
- **Environmental** — PowerMaxOS internal test detected an environmental error. Internal environmental tests monitor, check, and report failures of the critical hardware components. They run at the initial system power-up, upon every software reset event, and at least once every 24 hours during regular operations.

If an environmental test detects an error condition, it sets a flag to indicate a pending error and presents a unit check status to the host on the next I/O operation. The test that detected the error condition is then scheduled to run more frequently. If a device-level problem is detected, it is reported across all logical paths to the device experiencing the error. Subsequent failures of that device are not reported until the failure is fixed.

If a second failure is detected for a device while there is a pending error-reporting condition in effect, PowerMaxOS reports the pending error on the next I/O and then the second error.

PowerMaxOS reports error conditions to the host and to the Dell EMC Customer Support Center. When reporting to the host, PowerMaxOS presents a unit check status in the status byte to the channel whenever it detects an error condition such as a data check, a command reject, an overrun, an equipment check, or an environmental error.

When presented with a unit check status, the host retrieves the sense data from the storage array and, if logging action has been requested, places it in the Error Recording Data Set (ERDS). The EREP (Environment Recording, Editing, and Printing) program prints the error information. The sense data identifies the condition that caused the interruption and indicates the type of error and its origin. The sense data format depends on the mainframe operating system. For 2105, 2107, or 3990 controller emulations, the sense data is returned in the SIM format.

SIM severity reporting

PowerMaxOS supports SIM severity reporting that enables filtering of SIM severity alerts reported to the multiple virtual storage (MVS) console.

- All SIM severity alerts are reported by default to the EREP (Environmental Record Editing and Printing program).
- ACUTE, SERIOUS, and MODERATE alerts are reported by default to the MVS console.

The following table lists the default settings for SIM severity reporting.

Table 14. SIM severity alerts

| Severity | Description |
|----------|---|
| SERVICE | No system or application performance degradation is expected. No system or application outage has occurred. |

Table 14. SIM severity alerts (continued)

| Severity | Description |
|----------------|---|
| MODERATE | Performance degradation is possible in a heavily loaded environment. No system or application outage has occurred. |
| SERIOUS | A primary I/O subsystem resource is disabled. Significant performance degradation is possible. System or application outage may have occurred. |
| ACUTE | A major I/O subsystem resource is disabled, or damage to the product is possible. Performance may be severely degraded. System or application outage may have occurred. |
| REMOTE SERVICE | Dell EMC Customer Support Center is performing service/maintenance operations on the system. |
| REMOTE FAILED | The Service Processor cannot communicate with the Dell EMC Customer Support Center. |

Environmental errors

The following table lists the environmental errors in SIM format for PowerMaxOS 5978 or later.

Table 15. Environmental errors reported as SIM messages

| Hex code | Severity level | Description | SIM reference code |
|----------|----------------|--|--------------------|
| 04DD | MODERATE | MMCS health check error | 24DD |
| 043E | MODERATE | An SRDF Consistency Group was suspended. | E43E |
| 044D | MODERATE | An SRDF path was lost. | E44D |
| 044E | SERVICE | An SRDF path is operational after a previous failure. | E44E |
| 0461 | NONE | The M2 is resynchronized with the M1 device. This event occurs once the M2 device is brought back to a Ready state. ^a | E461 |
| 0462 | NONE | The M1 is resynchronized with the M2 device. This event occurs once the M1 device is brought back to a Ready state. ^a | E462 |
| 0463 | SERIOUS | One of the back-end directors failed into the IMPL Monitor state. | 2463 |
| 0465 | NONE | Device resynchronization process has started. ^a | E465 |
| 0467 | MODERATE | The remote storage system reported an SRDF error across the SRDF links. | E467 |
| 046D | MODERATE | An SRDF group is lost. This event happens, for example, when all SRDF links fail. | E46D |
| 046E | SERVICE | An SRDF group is up and operational. | E46E |

Table 15. Environmental errors reported as SIM messages (continued)

| Hex code | Severity level | Description | SIM reference code |
|------------------------------|----------------|--|--------------------|
| 0470 | ACUTE | OverTemp condition based on memory module temperature. | 2470 |
| 0471 | ACUTE | The Storage Resource Pool has exceeded its upper threshold value. | 2471 |
| 0473 | SERIOUS | A periodic environmental test (env_test9) detected the mirrored device in a Not Ready state. | E473 |
| 0474 | SERIOUS | A periodic environmental test (env_test9) detected the mirrored device in a Write Disabled (WD) state. | E474 |
| 0475 | SERIOUS | An SRDF R1 remote mirror is in a Not Ready state. | E475 |
| 0476 | SERVICE | Service Processor has been reset. | 2476 |
| 0477 | REMOTE FAILED | The Service Processor could not call the Dell EMC Customer Support Center (failed to call home) due to communication problems. | 1477 |
| 047A | MODERATE | AC power lost to Power Zone A or B. | 247A |
| 047B | MODERATE | Drop devices after RDF Adapter dropped. | E47B |
| 01BA 02BA 03BA 04BA | ACUTE | Power supply or enclosure SPS problem. | 24BA |
| 047C | ACUTE | The Storage Resource Pool has Not Ready or Inactive TDATs. | 247C |
| 047D | MODERATE | Either the SRDF group lost an SRDF link or the SRDF group is lost locally. | E47D |
| 047E | SERVICE | An SRDF link recovered from failure. The SRDF link is operational. | E47E |
| 047F | REMOTE SERVICE | The Service Processor successfully called the Dell EMC Customer Support Center (called home) to report an error. | 147F |
| 0488 | SERIOUS | Replication Data Pointer Meta Data Usage reached 90-99%. | E488 |
| 0489 | ACUTE | Replication Data Pointer Meta Data Usage reached 100%. | E489 |

Table 15. Environmental errors reported as SIM messages (continued)

| Hex code | Severity level | Description | SIM reference code |
|----------|----------------|--|--------------------|
| 0492 | MODERATE | Flash monitor or MMCS drive error. | 2492 |
| 04BE | MODERATE | Meta Data Paging file system mirror not ready. | 24BE |
| 04CA | MODERATE | An SRDF/A session dropped due to a non-user request. Possible reasons include fatal errors, SRDF link loss, or reaching the maximum SRDF/A host-response delay time. | E4CA |
| 04D1 | REMOTE SERVICE | Remote connection established. Remote control connected. | 14D1 |
| 04D2 | REMOTE SERVICE | Remote connection closed. Remote control rejected. | 14D2 |
| 04D3 | MODERATE | Flex filter problems. | 24D3 |
| 04D4 | REMOTE SERVICE | Remote connection closed. Remote control disconnected. | 14D4 |
| 04DA | MODERATE | Problems with task/threads. | 24DA |
| 04DB | SERIOUS | SYMPL script generated error. | 24DB |
| 04DC | MODERATE | PC related problems. | 24DC |
| 04E0 | REMOTE FAILED | Communications problems. | 14E0 |
| 04E1 | SERIOUS | Problems in error polling. | 24E1 |
| 052F | None | A sync SRDF write failure occurred. | E42F |
| 3D10 | SERIOUS | A SnapVX snapshot failed. | E410 |

a. Dell EMC recommendation: NONE.

Operator messages

Error messages

On z/OS, SIM messages are displayed as IEA480E Service Alert Error messages. They are formatted as shown below:

```
*IEA480E 1900,SCU,ACUTE ALERT,MT=2107,SER=0509-ANTPC, 266
REFCODE=1477-0000-0000,SENSE=00101000 003C8F00 40C00000 00000014
```

PC failed to call home due to communication problems.

Figure 29. z/OS IEA480E acute alert error message format (call home failure)




Figure 30. z/OS IEA480E service alert error message format (Disk Adapter failure)




Figure 31. z/OS IEA480E service alert error message format (SRDF Group lost/SIM presented against unrelated resource)

Event messages

The storage array also reports events to the host and to the service processor. These events are:

- The mirror-2 volume has synchronized with the source volume.
- The mirror-1 volume has synchronized with the target volume.
- Device resynchronization process has begun.

On z/OS, these events are displayed as IEA480E Service Alert Error messages. They are formatted as shown below:

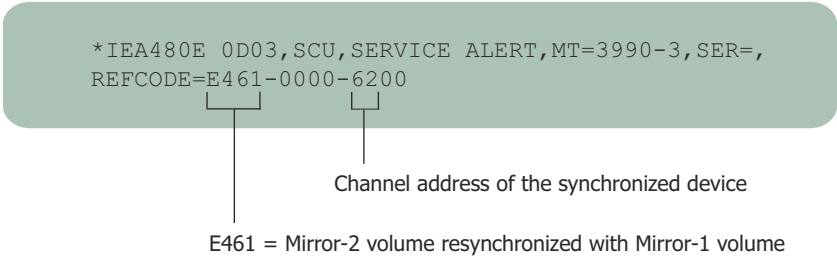


Figure 32. z/OS IEA480E service alert error message format (mirror-2 resynchronization)

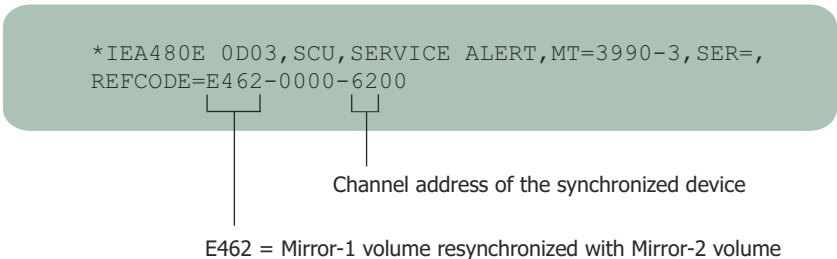


Figure 33. z/OS IEA480E service alert error message format (mirror-1 resynchronization)

Licensing

This appendix is an overview of licensing on arrays running PowerMaxOS.

Topics:

- [eLicensing](#)
- [Open systems licenses](#)
- [PowerMax Mainframe software packaging options](#)

eLicensing

Arrays running PowerMaxOS use *Electronic Licenses* (eLicenses).

NOTE: For more information on eLicensing, refer to Dell EMC Knowledgebase article 335235 on the Dell EMC Online Support website.

You obtain license files from Dell EMC Online Support, copy them to a Solutions Enabler or a Unisphere host, and push them out to your arrays. The following figure illustrates the process of requesting and obtaining your eLicense.

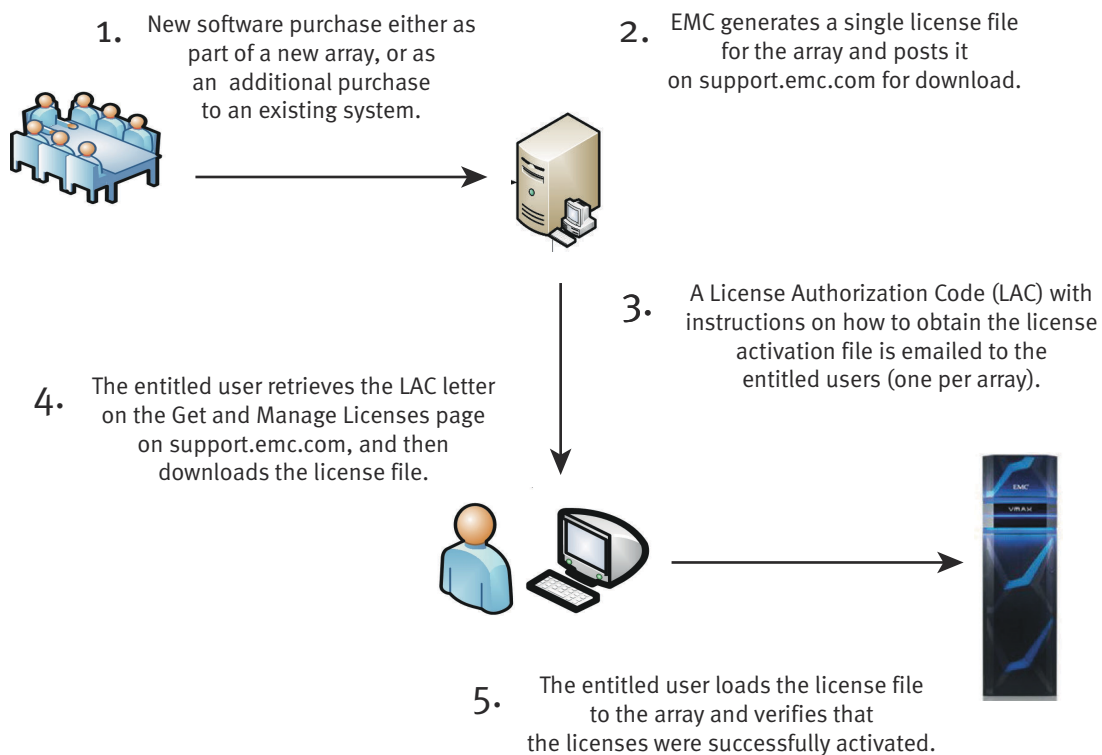


Figure 34. eLicensing process

NOTE: To install array licenses, follow the procedure described in the *Solutions Enabler Installation Guide* and *Unisphere Online Help*.

Each license file fully defines all of the entitlements for a specific system, including the license type and the licensed capacity. To add a feature or increase the licensed capacity, obtain and install a new license file.

Most array licenses are array-based, meaning that they are stored internally in the system feature registration database on the array. However, there are a number of licenses that are host-based.

Array-based eLicenses are available in the following forms:

- An *individual license* enables a single feature.
- A *license suite* is a single license that enables multiple features. License suites are available only if all features are enabled.
- A *license pack* is a collection of license suites that fit a particular purpose.

To view effective licenses and detailed usage reports, use Solutions Enabler, Unisphere, Mainframe Enablers, Transaction Processing Facility (TPF), or IBM i platform console.

Capacity measurements

Array-based licenses include a *capacity licensed* value that defines the scope of the license. The method for measuring this value depends on the license's *capacity type* (Usable or Registered).

Not all product titles are available in all capacity types, as shown below.

Table 16. PowerMax product title capacity types

| Usable | Registered | Other |
|---|----------------|-------------------------------------|
| All Essential software package titles | Storage Direct | PowerPath (if purchased separately) |
| All Pro software package titles | | Events and Retention Suite |
| All zEssentials software package titles | | |
| All zPro software package titles | | |
| RecoverPoint | | |

Usable capacity

Usable Capacity is defined as the amount of storage available for use on an array. The usable capacity is calculated as the sum of all Storage Resource Pool (SRP) capacities available for use. This capacity does not include any external storage capacity.

Registered capacity

Registered capacity is the amount of user data managed or protected by each particular product title. It is independent of the type or size of the disks in the array.

The methods for measuring registered capacity depends on whether the licenses are part of a bundle or individual.

Registered capacity licenses

Registered capacity is measured according to the following:

- Storage Direct
 - The registered capacity of this license is the sum of all Data Domain encapsulated devices that are link targets. When there are TimeFinder sessions present on an array with only a Storage Direct license and no TimeFinder license, the capacity is calculated as the sum of all Data Domain encapsulated devices with link targets and the sum of all TimeFinder allocated source devices and delta RDPs.

Open systems licenses

This section details the licenses available in an open system environment.

License packages

This table lists the license packages available in an open systems environment.

Table 17. PowerMax license packages

| License suite | Includes | Allows you to | With the command |
|-----------------------------|--|--|----------------------------------|
| Essentials software package | <ul style="list-style-type: none"> PowerMaxOS Priority Controls OR-DM Unisphere for PowerMax SL Provisioning Workload Planner Database Storage Analyzer | Create time windows | symoptmz symtw |
| | | Perform SL-based provisioning | symconfigure symmsg symcfg |
| | AppSync Starter Pack | Manage protection and replication for critical applications and databases for Microsoft, Oracle and VMware environments. | |
| | <ul style="list-style-type: none"> TimeFinder/Snap TimeFinder/SnapVX SnapSure | Create new native clone sessions | symclone |
| | | Create new TimeFinder/Clone emulations | symmir |
| | | <ul style="list-style-type: none"> Create new sessions Duplicate existing sessions | symsnap |
| | | <ul style="list-style-type: none"> Create snap pools Create SAVE devices | symconfigure |
| | | <ul style="list-style-type: none"> Perform SnapVX Establish operations Perform SnapVX snapshot Link operations | symsnapvx |
| Pro software package | Essentials software package | Perform tasks available in the Essentials software package. | |
| | <ul style="list-style-type: none"> SRDF SRDF/Asynchronous SRDF/Synchronous SRDF/Star Replication for File | <ul style="list-style-type: none"> Create new SRDF groups Create dynamic SRDF pairs in Adaptive Copy mode | symrdf |
| | | <ul style="list-style-type: none"> Create SRDF devices Convert non-SRDF devices to SRDF Add SRDF mirrors to devices in Adaptive Copy mode | symconfigure |

Table 17. PowerMax license packages (continued)

| License suite | Includes | Allows you to | With the command |
|---------------|----------|---|------------------|
| | | Set the dynamic-SRDF capable attribute on devices Create SAVE devices | |
| | | <ul style="list-style-type: none"> • Create dynamic SRDF pairs in Asynchronous mode • Set SRDF pairs into Asynchronous mode | symrdf |
| | | <ul style="list-style-type: none"> • Add SRDF mirrors to devices in Asynchronous mode <p>Create RDFA_DSE pools</p> <p>Set any of the following SRDF/A attributes on an SRDF group:</p> <ul style="list-style-type: none"> ○ Minimum Cycle Time ○ Transmit Idle ○ DSE attributes, including: <ul style="list-style-type: none"> ▪ Associating an RDFA-DSE pool with an SRDF group DSE Threshold DSE Autostart ○ Write Pacing attributes, including: <ul style="list-style-type: none"> ▪ Write Pacing Threshold ▪ Write Pacing Autostart ▪ Device Write Pacing exemption ▪ TimeFinder Write Pacing Autostart | symconfigure |
| | | <ul style="list-style-type: none"> • Create dynamic SRDF pairs in Synchronous mode • Set SRDF pairs into Synchronous mode | symrdf |
| | | Add an SRDF mirror to a device in Synchronous mode | symconfigure |
| | | | |
| | D@RE | Encrypt data and protect it against unauthorized access unless valid keys are provided. This prevents data from being accessed and provides a mechanism to quickly shred data. | |

Table 17. PowerMax license packages (continued)

| License suite | Includes | Allows you to | With the command |
|---------------|------------|--|------------------|
| | SRDF/Metro | <ul style="list-style-type: none"> Place new SRDF device pairs into an SRDF/Metro configuration. Synchronize device pairs. | |
| | SRM | Automate storage provisioning and reclamation tasks to improve operational efficiency. | |

Individual licenses

These items are available for arrays running PowerMaxOS and are not in any of the license suites:

Table 18. Individual licenses for open systems environment

| License | Allows you to | With the command |
|----------------|--|------------------|
| Storage Direct | Store and retrieve backup data within an integrated environment containing arrays running PowerMaxOS and Data Domain arrays. | |
| RecoverPoint | Protect data integrity at local and remote sites, and recover data from a point in time using journaling technology. | |

Ecosystem licenses

These licenses do not apply to arrays:

Table 19. Individual licenses for open systems environment

| License | Allows you to |
|----------------------------|---|
| PowerPath | Automate data path failover and recovery to ensure applications are always available and remain operational. |
| Events and Retention Suite | <ul style="list-style-type: none"> Protect data from unwanted changes, deletions and malicious activity. Encrypt data where it is created for protection anywhere outside the server. Maintain data confidentiality for selected data at rest and enforce retention at the file-level to meet compliance requirements. Integrate with third-party anti-virus checking, quota management, and auditing applications. |

PowerMax Mainframe software packaging options

This table lists the mainframe software packaging options for PowerMax.

Table 20. PowerMax Mainframe software packaging options (PowerMax 8000 only)

| Feature | zEssentials package include | zEssentials package options | zPro package included | zPro package options | Notes |
|---------------------------------------|-----------------------------|-----------------------------|-----------------------|----------------------|---|
| PowerMaxOS | Yes | | Yes | | |
| Embedded Management | Yes | | Yes | | Includes Unisphere for PowerMax REST APIs, SMI-S |
| Local Replication | Yes | | Yes | | Includes TimeFinder SnapVX, Compatible Flash (FlashCopy support) |
| Mainframe Essentials | Yes | | Yes | | Includes Compatible High Performance FICON (zHPF) and Compatible PAV (Dynamic, Hyper, and SuperPAV) support |
| Remote Replication Suite ^a | | Yes | Yes | | Includes SRDF/S/A/STAR, Mirror Optimizer, Compatible Peer (PPRC) |
| Unisphere 360 | | Yes | Yes | | |
| AutoSwap | | Yes | Yes | | |
| D@RE ^b | | Yes | Yes | | |
| zDP | | Yes | Yes | | |
| Mainframe Essentials Plus | | Yes | | Yes | zBoost PAV Optimizer |
| GDDR ^c | | Yes | | Yes | |

a. Software packages include software licensing. Order any additional required hardware separately.

b. Factory configured. Must be enabled during the ordering process.

c. Use of SRDF/STAR for mainframe requires GDDR.

Index

C

cloud mobility [91](#)

G

Global mirror [52](#)

S

Snapshot policy [71](#)

T

TCT environment [53](#)

Transparent Cloud Tiering [53](#)