

Microsemi Smart Storage Controllers
User Guide
maxView Storage Manager

Released
October 2020



a  **MICROCHIP** company

Revision History

Revision	Revision Date	Details of Change
12	October 2020	maxView Storage Manager v3.06
11	October 2019	maxView Storage Manager v3.03.00
10	August 2019	maxView Storage Manager v3.02.00
9	March 2019	maxView Storage Manager v3.01.00
8	January 2019	maxView Storage Manager v3.0.0
7	May 2018	Removed "Proprietary and Confidential" markings in footer text.
6	April 2018	maxView Storage Manager v2.06
5	September 2017	maxView Storage Manager v2.04
4	May 2017	maxView Storage Manager v2.03
3	September 2016	V1.2.0 Maintenance Release. maxView Storage Manager v2.02
2	April 2016	V1.1.0 Maintenance Release. maxView Storage Manager v2.01
1	December 2015	First Production release. maxView Storage Manager v2.00

Microsemi Adaptec® Product Support

If you have questions about installing or using your Microsemi Adaptec® product, check this document first—you will find answers to most of your questions. If you need further assistance, use the support options listed below. To expedite your service, have your computer in front of you.

Note: Please visit our Support site at start.microsemi.com for the most up to date contact information.

Self Help and Support in English

- Search the Microsemi Support Knowledgebase (ASK) at ask.microsemi.com for articles, troubleshooting tips, and frequently asked questions for your product.
- For support through email, submit your question at ask.microsemi.com.
- To contact Technical Support, visit our product support site at start.microsemi.com.

Technische Informationen und Support in Deutsch

- Suchen Sie in der Adaptec Support Knowledgebase (ASK) unter ask-de.microsemi.com nach Artikeln, Tipps zur Fehlerbehebung und häufig gestellten Fragen zu Ihrem Produkt.
- Support per Email erhalten Sie unter ask-de.microsemi.com.
- Um den Technischen Support zu kontaktieren, besuchen Sie uns bitte unter start.microsemi.com und klicken Sie auf „Support kontaktieren“, für Auswahlmöglichkeiten.

Техническая поддержка и информация на русском языке

- База знаний Microsemi (ASK) на сайте ask-ru.microsemi.com ask-ru.adaptec.com – статьи, советы по устранению неисправностей и часто задаваемые вопросы о Вашем продукте.
- Для поддержки по электронной почте отправьте Ваш запрос на сайте ask-ru.microsemi.com
- Для обращения в службу Технической Поддержки, пожалуйста, посетите наш web сайт start.microsemi.com и используйте ссылку "Contact Support".

日本語での技術情報とサポート

- ask.microsemi.co.jp の Microsemi Support Knowledgebase (ASK)で、お使いの製品の情報トラブルシューティングのヒント、よくある質問を検索してください。
- Eメールでのサポートには ask.microsemi.co.jp から質問を送ってください。
- テクニカルサポートへコンタクトするには、弊社ウェブサイト start.microsemi.com をご覧になり、"Contact Support"をクリックして下さい。

Software License Agreement

PLEASE READ CAREFULLY: THE USE OF THIS SOFTWARE IS SUBJECT TO THE SOFTWARE LICENSE TERMS OF MICROSEMI CORPORATION, AND OTHER LICENSORS WHOSE SOFTWARE MAY BE BUNDLED WITH THIS PRODUCT.

BY YOUR USE OF THE SOFTWARE INCLUDED WITH THIS PRODUCT YOU AGREE TO THE LICENSE TERMS REQUIRED BY THE LICENSOR OF THAT SOFTWARE, AS SET FORTH DURING THE INSTALLATION PROCESS. IF YOU DO NOT AGREE TO THE LICENSE TERMS APPLICABLE TO THE SOFTWARE, YOU MAY RETURN THE ENTIRE UNUSED PRODUCT FOR A FULL REFUND.

In return for acquiring a license to use the Microsemi software, which may include software from third party licensors and patches made available by Microsemi (“Software”), and the related documentation, you agree to the following terms and conditions:

1. License. This Agreement grants you, the Licensee, a license to:
 - a. Use the Software on a single computer system, which is not intended for use by more than five (5) users; and:
 - b. Make one copy of the Software in machine readable form solely for back-up purposes, provided you reproduce Microsemi's copyright proprietary legends. Notwithstanding the foregoing, the Software may be used on the home, laptop or other secondary computer of the principal user of the Software, and an additional copy of the Software may be made to support such use. As used in this license, the Software is “in use” when it is either loaded into RAM or installed on a hard disk or other permanent memory device. The Software may be “in use” on only one computer at any given time. (Different license terms and fees are applicable for networked or multiple user applications.) As a specific condition of this license, you agree to use the Software in compliance with all applicable laws, including copyright laws, and that you will not copy, transmit, perform or distribute any audio or other content using the Software without obtaining all necessary licenses or permissions from the owner of the content.
2. Restrictions. You may not distribute copies of the Software to others or electronically transfer the Software from one computer to another over a network. You may not post or otherwise make available the Software, or any portion thereof, in any form, on the Internet. You may not use the Software in a computer service business, including in time sharing applications. The Software contains trade secrets and, in order to protect them, you may not decompile, reverse engineer, disassemble, or otherwise reduce the Software to a human-perceivable form. YOU MAY NOT MODIFY, ADAPT, TRANSLATE, RENT, LEASE, LOAN, RESELL FOR PROFIT, DISTRIBUTE, NETWORK OR CREATE DERIVATIVE WORKS BASED UPON THE SOFTWARE OR ANY PART THEREOF.
3. Ownership of Software. As Licensee, you own the media upon which the software is recorded or fixed, but Microsemi and its licensors retain title and ownership of the Software recorded on the original media and all subsequent copies of the Software, regardless of the form or media in which or on which the original and other copies may exist. This license is not a sale of the Software or any copy.
4. Confidentiality. You agree to maintain the Software in confidence and that you will not disclose the Software to any third party without the express written consent of Microsemi. You further agree to take all reasonable precautions to preclude access of unauthorized persons to the Software.
5. Term. This license is effective until January 1, 2045, unless terminated earlier. You may terminate the license at any time by destroying the Software (including the related documentation) together with all copies or modifications in any form. Microsemi will have the right to terminate our license immediately if you fail to comply with any term or condition of this Agreement. Upon any termination, including termination by you, you must destroy the Software (including the related documentation), together with all copies or modifications in any form.

6. **Special Terms Applicable to Databases.** Where a database is included with the Software, you acknowledge that it is licensed only in connection with the use of the Software to perform disc creation, and that the database and all data derived therefrom must be maintained in confidence in accordance with the provisions of Section 4. This license does not grant you any rights to distribute or disclose such database or data.
7. **Limited Warranty.** Microsemi and its Licensor warrant only that the media upon which the Software is furnished will be free from defects in material or workmanship under normal use and service for a period of thirty (30) days from the date of delivery to you. MICROSEMI AND ITS LICENSORS DO NOT AND CANNOT WARRANT THE PERFORMANCE OR RESULTS YOU MAY OBTAIN BY USING THE SOFTWARE OR DOCUMENTATION. THE FOREGOING STATES THE SOLE AND EXCLUSIVE REMEDIES MICROSEMI AND ITS LICENSORS WILL PROVIDE FOR BREACH OF WARRANTY. EXCEPT FOR THE FOREGOING LIMITED WARRANTY, MICROSEMI AND ITS LICENSORS MAKE NO WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED, AS TO NON-INFRINGEMENT OF THIRD PARTY RIGHTS, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow the exclusion of implied warranties or limitations on how long an implied warranty may last, so the above limitations may not apply to you. This warranty gives you specific legal rights and you may also have other rights which vary from state to state.
8. **The entire liability of Microsemi and its licensors, and your exclusive remedy for a breach of this warranty, shall be:**
 - a. The replacement of any media not meeting the above limited warranty which is returned to Microsemi; or
 - b. if Microsemi or its distributor is unable to deliver replacement media which is free from defects in materials or workmanship, you may terminate this Agreement by returning the Software and your money will be refunded.
9. **Limitation of Liability.** IN NO EVENT WILL MICROSEMI OR ITS LICENSORS BE LIABLE TO YOU FOR ANY INCIDENTAL, CONSEQUENTIAL OR INDIRECT DAMAGES, INCLUDING ANY LOST PROFITS, LOST SAVINGS, OR LOSS OF DATA, EVEN IF MICROSEMI OR A LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, OR FOR ANY CLAIM BY ANY OTHER PARTY. Some states do not allow the exclusion or limitation of special, incidental, or consequential damages, so the above limitation or exclusion may not apply to you.
10. **Export.** You acknowledge that the laws and regulations of the United States and other countries may restrict the export and re-export of the Software. You agree that you will not export or re-export the Software or documentation in any form in violation of applicable United States and foreign law.
11. **Government Restricted Rights.** The Software is subject to restricted rights as follows. If the Software is acquired under the terms of a GSA contract: use, reproduction or disclosure is subject to the restrictions set forth in the applicable ADP Schedule contract. If the Software is acquired under the terms of a DoD or civilian agency contract, use, duplication or disclosure by the Government is subject to the restrictions of this Agreement in accordance with 48 C.F.R. 12.212 of the Federal Acquisition Regulations and its successors and 49 C.F.R. 227.7202-1 of the DoD FAR Supplement and its successors.
12. **General.** You acknowledge that you have read this Agreement, understand it, and that by using the Software you agree to be bound by its terms and conditions. You further agree that it is the complete and exclusive statement of the agreement between Microsemi and you, and supersedes any proposal or prior agreement, oral or written, and any other communication between Microsemi and you relating to the subject matter of this Agreement. No additional or any different terms will be enforceable against Microsemi unless Microsemi gives its express consent, including an express waiver of the terms of this Agreement, in writing signed by an officer of Microsemi. You assume full responsibility for the use of the Software and agree to use the Software legally and responsibly. This Agreement shall be governed by California law, except as to copyright matters, which are covered by Federal law. This Agreement is deemed entered into at Sunnyvale, California by both parties. Should any provision of this Agreement be declared unenforceable in any jurisdiction, then such provision shall be deemed severable from this

Agreement and shall not affect the remainder hereof. All rights in the Software not specifically granted in this Agreement are reserved by Microsemi.

Should you have any questions concerning this license, contact:

Microsemi Corporation
Legal Department
One Enterprise
Aliso Viejo, CA 92656
(949) 380-6100

Contents

1 About this Guide.....	1
1.1 What You Need to Know Before You Begin.....	1
1.2 Terminology Used in this Guide.....	2
1.3 How to Find More Information.....	2
2 Introduction to maxView Storage Manager.....	3
2.1 Getting Started.....	3
2.2 About maxView Storage Manager.....	3
2.2.1 About maxView Storage Manager Redfish Server	3
2.2.2 About the maxView Storage Manager Web Server.....	3
2.3 System Requirements.....	4
2.4 Browser Support.....	4
2.5 Typical Storage Space Configurations.....	5
2.5.1 A Simple Storage Space.....	5
2.5.2 An Advanced Storage Space.....	5
2.5.3 Continuing to Grow Your Storage Space.....	6
3 Installing maxView Storage Manager.....	8
3.1 Before You Begin the Installation.....	8
3.1.1 Gather Installation Information.....	8
3.1.2 Download the Installation Package.....	8
3.2 Installing on Windows.....	10
3.3 Installing on Red Hat, Citrix XenServer, CentOS, or SuSE Linux.....	11
3.4 Installing on Debian or Ubuntu Linux.....	11
3.5 Installing on VMware.....	12
3.6 Running maxView Storage Manager from a Bootable USB Image	13
3.7 Uninstalling maxView Storage Manager.....	14
3.7.1 Uninstalling from Windows.....	14
3.7.2 Uninstalling from Red Hat, Citrix XenServer, CentOS, or SuSE Linux.....	14
3.7.3 Uninstalling from Ubuntu Linux.....	14
3.7.4 Uninstalling from VMware.....	14
4 Exploring maxView Storage Manager.....	15
4.1 Starting maxView Storage Manager and Logging In	15
4.2 Working in maxView Storage Manager.....	15
4.3 Overview of the Main Window.....	16
4.3.1 The Enterprise View	16
4.3.2 The Ribbon	19
4.3.3 The Storage Dashboard.....	19
4.4 Checking System Status from the Main Window.....	20
4.5 Revealing More Device Information	21
4.6 Getting Help.....	22
4.7 Logging Out of maxView Storage Manager.....	23
5 Building Your Storage Space.....	24
5.1 Overview.....	24
5.2 Choosing a Management System.....	24
5.2.1 ‘Local’ or ‘Remote’?.....	24
5.2.2 Logging in on the Local System	25
5.3 Logging into Remote Systems from the Local System.....	25

5.4 Creating Arrays and Logical Drives.....	26
5.4.1 Creating a Logical Drive on a New Array	27
5.4.2 Creating a Logical Drive on an Existing Array	30
5.4.3 Partitioning and Formatting Your Logical Drives.....	33
5.4.4 Creating Logical Drives on Other Systems in Your Storage Space.....	34
5.4.5 Controller Support for 4K Drives.....	34
6 Protecting Your Data.....	40
6.1 Dedicated Spare or Auto-Replace Spare?.....	40
6.2 Hot Spare Limitations.....	40
6.3 Assigning a Dedicated Hot Spare	41
6.4 Assigning an Auto-Replace Hot Spare.....	42
6.5 Removing a Hot Spare.....	44
6.6 Setting the Spare Activation Mode	46
6.7 Controller Sanitize Lock Freeze/Anti-Freeze	47
6.7.1 Sanitize Lock Property in Controller Node Properties Tab.....	48
6.7.2 Physical Device Sanitize Lock Freeze/Anti-Freeze.....	49
6.7.3 Secure Erase Pattern.....	50
7 Modifying Your Storage Space.....	52
7.1 Understanding Arrays and Logical Drives.....	52
7.2 Creating and Modifying Logical Drives.....	52
7.2.1 Including Different-sized Disk Drives in a Logical Drive.....	53
7.3 Enabling Background Consistency Check.....	54
7.4 Optimizing Logical Drive Performance.....	55
7.4.1 Enabling Cache Optimizations.....	55
7.4.2 Enabling SSD I/O Bypass.....	57
7.5 Moving a Logical Drive.....	57
7.6 Moving an Array.....	59
7.7 Modifying an Array.....	61
7.8 Working with Mirrored Arrays.....	62
7.8.1 Creating a Split Mirror Backup	63
7.8.2 Re-mirroring, Rolling Back, or Reactivating a Split Mirror Backup	63
7.9 Changing the RAID Level of a Logical Drive.....	64
7.10 Increasing the Capacity of a Logical Drive.....	66
7.11 Changing the Logical Drive Rebuild Priority.....	67
7.12 Renaming a Logical Drive.....	68
7.13 Deleting an Array or Logical Drive	68
7.14 Maintaining an Energy-Efficient Storage Space	68
8 Working with maxCache Devices.....	70
8.1 maxCache Limitations.....	70
8.2 Creating a maxCache Device.....	70
8.3 Changing the Write Cache Mode	72
8.4 Deleting the maxCache Device.....	72
8.5 Analyzing maxCache Performance	73
9 Working with maxCrypto Devices.....	74
9.1 maxCrypto Initial Setup	74
9.1.1 Managing maxCrypto Accounts	76
9.1.2 Logging In and Logging Out.....	77
9.1.3 Checking maxCrypto Status	79
9.2 Modifying the maxCrypto Configuration	80
9.3 Creating an Encrypted Logical Drive.....	81
9.4 Converting Plaintext Data to Encrypted Data.....	82
9.5 Re-Keying a Logical Drive	83

9.6 Clearing the maxCrypto Configuration	84
9.7 Erasing an Encrypted Logical Drive.....	85
9.8 Importing a Foreign Master Key	86
10 Maintaining Physical Devices.....	88
10.1 Viewing Device Properties.....	88
10.2 Locating Drives in Your Storage Space.....	88
10.2.1 Locating Disk Drives.....	89
10.2.2 Locating Physical Disks in an Array or Logical Drive.....	90
10.2.3 Locating SSDs in the maxCache Device.....	90
10.3 Working with Physical Device Error Counters.....	91
10.3.1 Viewing Physical Device Error Counters	92
10.3.2 Clearing Error Counters from a Physical Device.....	93
10.3.3 Clearing Error Counters from a Controller.....	94
10.4 Working with Failed or Failing Disk Drives.....	94
10.4.1 Replacing Disk Drives in a Logical Drive.....	94
10.4.2 Setting a Disk Drive to 'Failed'.....	95
10.5 Erasing a Disk Drive	95
10.5.1 Restricted/Unrestricted Secure Erase.....	96
10.6 Initializing and Uninitializing Disk Drives.....	96
10.6.1 Uninitializing a Disk Drive.....	96
10.6.2 Initializing/Uninitializing all Drives on a Controller	97
10.7 Setting the Physical Drive Boot Priority.....	98
10.8 Working with Controllers.....	99
10.8.1 Rescanning a Controller.....	99
10.8.2 Optimizing Controller Performance.....	100
10.8.3 Changing the Connector Operating Mode.....	101
10.8.4 Changing the Backplane Discovery Protocol Settings.....	102
10.8.5 Changing the Expander Minimum Scan Duration Settings.....	103
10.9 Updating Controller, Enclosure, and Disk Drive Firmware.....	104
10.9.1 Before You Begin.....	104
10.9.2 Updating the Controller Firmware	104
10.9.3 Updating the Disk Drive Firmware	108
10.9.4 Updating the Enclosure Firmware	109
11 Monitoring Status and Activity.....	111
11.1 Monitoring Options.....	111
11.2 Checking Status from the Main Window	111
11.2.1 Viewing Activity Status in the Event Log.....	112
11.2.2 Viewing Task Status in the Task Log.....	113
11.2.3 Viewing Component Status in the Storage Dashboard.....	113
11.2.4 Viewing Storage Space Usage in Chart View.....	116
11.3 Notifying Users by Email About Status and Activity	116
11.3.1 Setting Up Email Notifications.....	117
11.3.2 Sending a Test Message.....	119
11.3.3 Modifying or Removing an Email Recipient	120
11.3.4 Modifying Email Server Settings.....	121
11.3.5 Disabling Email Notifications.....	121
11.4 Monitoring Advanced Statistics about Activity in Your Storage Space.....	122
11.4.1 Setting up Statistics Logging.....	122
11.4.2 Viewing Advanced Statistics.....	124
11.4.3 Resetting the Statistics Counters.....	125
11.5 Changing an Operating System's Event Log Setting.....	125
12 Managing Your Storage Space.....	127
12.1 Deploying Servers	127

12.1.1 Creating a Server Template File.....	127
12.1.2 Duplicating the Server Template.....	128
12.2 Managing Remote Systems.....	130
12.2.1 Adding Remote Systems with the Wizard.....	131
12.2.2 Manually Adding a Remote System.....	131
12.2.3 Removing a Remote System.....	131
12.2.4 Changing the Auto-Discovery Settings.....	132
12.2.5 Importing and Exporting Remote Systems.....	133
12.3 Clearing the Controller Configuration.....	137
12.4 Changing the Web Server Port.....	138
12.5 Granting Standard Users Admin Privilege.....	139
13 Solving Problems.....	140
13.1 General Troubleshooting Tips.....	140
13.2 Identifying a Failed or Failing Component.....	140
13.3 Recovering from a Disk Drive Failure.....	141
13.3.1 Failed Disk Drive Protected by a Hot Spare.....	141
13.3.2 Failed Disk Drive <i>Not</i> Protected by a Hot Spare.....	142
13.3.3 Failure in Multiple Logical Drives Simultaneously.....	142
13.3.4 Disk Drive Failure in a RAID 0 Logical Drive.....	142
13.3.5 Forcing a Logical Drive with Multiple Drive Failures Back Online.....	142
13.3.6 Healing an Array.....	143
13.4 Rebuilding Logical Drives.....	144
13.5 Creating a Support Archive File.....	145
Appendix A Silent Installation on Windows and Linux.....	146
A.1 Completing a Silent Installation.....	146
A.1.1 Windows Silent Installation.....	146
A.1.2 Linux Silent Installation.....	147
A.2 Example Command Line Installations.....	148
Appendix B Configuring SNMP Notifications on Windows and Linux.....	149
B.1 Setting Up SNMP Notifications on Windows.....	149
B.2 Setting Up SNMP Notifications on Linux.....	150
Appendix C Using the maxView Plugin for VMware vSphere Web Client.....	151
C.1 Installing the maxView Plugin for vSphere Web Client.....	151
C.2 Starting the maxView Plugin for vSphere Web Client.....	152
C.3 Monitoring maxView Resources in vSphere Web Client.....	153
Appendix D Using the maxView Plugin for VMware vSphere HTML5.....	156
D.1 Installing the maxView Plugin for vSphere HTML5 Client.....	156
D.2 Starting the maxView Plugin for vSphere HTML Client.....	157
D.3 Monitoring maxView Resources in vSphere HTML Client.....	159
D.3.1 maxView vSphere Plugin Resources.....	160
Appendix E Using maxView Storage Manager with HBAs and Non-RAID Mode Controllers	162
Appendix F Selecting the Best RAID Level.....	164
F.1 Comparing RAID Levels.....	164
F.2 Non-redundant Logical Drives (RAID 0).....	165
F.3 RAID 1 Logical Drives	165
F.4 RAID 1 (ADM) Logical Drives	166
F.5 RAID 10 Logical Drives.....	167
F.6 RAID 10 (ADM) Logical Drives	167
F.7 RAID 5 Logical Drives.....	169
F.8 RAID 50 Logical Drive.....	170


F.9 RAID 6 Logical Drives.....	171
F.10 RAID 60 Logical Drives.....	171
Appendix G Icons At-a-Glance.....	172
Appendix H Advanced Statistics Descriptions.....	178
Appendix I Smart Controller Device Status.....	181
Appendix J Display Properties of a Controller, Array, Logical Device, and a Physical Device.....	185
Appendix K maxView Video Tutorials.....	193

1 About this Guide

maxView Storage Manager™ is a browser-based software application that helps you build a storage space using Microsemi Smart Storage Controllers, disk drives, and enclosures, and then manage your stored data, whether you have a single controller installed in a server or multiple controllers, servers, and enclosures.

This guide describes how to install and use maxView Storage Manager to build and manage *direct attached storage*; that is, storage where the controller and disk drives reside inside, or are directly attached to, the computer accessing them, similar to the basic configurations shown in the figures below.

Note: This guide focuses on using maxView Storage Manager with Microsemi Smart Storage Controllers (SmartRAID/SmartHBA/SmartIOC/SmartROC). For information about using maxView Storage Manager with Microsemi Adaptec Series 8 (legacy) RAID controllers, see [How to Find More Information](#) on page 2.



1.1 What You Need to Know Before You Begin

This guide is written for data storage and IT professionals who want to create a storage space for their online data. You should be familiar with computer hardware, operating system administration, and Redundant Array of Independent Disks (RAID) technology.

If you are using maxView Storage Manager as part of a complex storage system, with multiple servers, enclosures and Microsemi Smart Storage Controllers, you should be familiar with network administration, have knowledge of Local Area Networks (knowledge of storage area networks (SANs) is not required), and be familiar with the input/output (I/O) technology of the storage devices on your network, such as Serial ATA (SATA) or Serial Attached SCSI (SAS).

1.2 Terminology Used in this Guide

Because this guide provides information that can be used to manage multiple Smart Storage Controllers in a variety of configurations, the generic term “storage space” is used to refer to the controller(s), disk drives, and systems being managed with maxView Storage Manager.

For efficiency, the term “component” or “components” is used when referring generically to the physical and virtual parts of your storage space, such as systems, disk drives, controllers, and logical drives.

Many of the terms and concepts referred to in this guide are known to computer users by multiple names. In this guide, this terminology is used:

- Controller (also known as adapter, board, or I/O card)
- Disk drive (also known as hard disk, hard drive, or hard disk drive)
- Solid State Drive (also known as SSD or non-rotating storage media)
- Logical drive (also known as a logical device)
- Array (also known as a storage pool or container)
- System (also known as a server, workstation, or computer)
- Enclosure (also known as a storage enclosure or disk drive enclosure)

1.3 How to Find More Information

You can find more information about your Microsemi Smart Storage Controller, management software, and utilities by referring to these documents, available for download at start.microsemi.com and the Microsemi customer portal at www.pmcs.com/myPMC:

- *SmartIOC 2100/SmartROC 3100 Installation and User's Guide, SmartIOC 2000 Installation and User's Guide*—Describes how to install drivers and configure the SmartIOC/SmartROC controller for initial use
- *ARCCONF Command Line Utility User's Guide for Smart Storage Controllers, SmartIOC 2000 Command Line Utility User's Guide*—Describes how to use the ARCCONF utility to perform RAID configuration and storage management tasks from an interactive command line.
- *SmartIOC 2100/SmartROC 3100 Software/Firmware Release Notes, SmartIOC 2000 Software/Firmware Release Notes*—Provides driver, firmware, and release package information, and known issues.
- *README: maxView Storage Manager & ARCCONF Command Line Utility*—Provides product information, installation notes, and known issues for maxView Storage Manager and ARCCONF command line utility.
- *Microchip Adaptec® SmartRAID 3100 Series and SmartHBA 2100 Series Host Bus Adapters Installation and User's Guide*—Describes how to install drivers and configure the SmartRAID 3100 or SmartHBA 2100 Series Host Bus Adapter.
- *HBA 1100 Software/Firmware Release Notes*—Provides driver, firmware, and release package information, and known issues.
- *SmartHBA 2100 and SmartRAID 3100 Software/Firmware Release Notes*—Provides driver, firmware, and release package information, and known issues.

For information about using maxView Storage Manager with Microsemi Adaptec Series 8 (legacy) RAID controllers, see the *maxView Storage Manager User's Guide for Microsemi ARC Controllers* (CDP-00285-06-A).

2 Introduction to maxView Storage Manager

This section introduces the maxView Storage Manager software, explains the concept of a “storage space,” and provides a checklist of getting-started tasks.

2.1 Getting Started

The first part of this guide provides the information you need to install, start, and begin to use maxView Storage Manager. Follow these general steps:

Step 1: Familiarize yourself with the software components of maxView Storage Manager, review the system requirements, and study the configuration examples that illustrate how to build and grow your storage space (described in the remainder of this chapter).

Step 2: Install maxView Storage Manager on every system that will be part of your storage space (see [Installing maxView Storage Manager](#) on page 8).

Step 3: Start maxView Storage Manager and explore its graphical user interface (see [Exploring maxView Storage Manager](#) on page 15).

Step 4: Build your storage space (see [Building Your Storage Space](#) on page 24).

2.2 About maxView Storage Manager

maxView Storage Manager is a browser-based software application that helps you build a storage space for your data, using RAID controllers, disk drives, Solid State Drives (SSDs), and enclosures.

With maxView Storage Manager, you can group disk drives into arrays and logical drives and build in redundancy to protect your data and improve system performance. You can also use maxView Storage Manager to monitor and maintain all the controllers, enclosures, and disk drives in your storage space from a single location.

The maxView Storage Manager GUI, or *graphical user interface*, runs on most contemporary Web browsers (for a list of supported browsers, see [Browser Support](#) on page 4). A software stack comprising a Web server, and Redfish server allows maxView Storage Manager to communicate with the controller(s) in your storage space and coordinate activity in your system.

A flexible installation model allows you to install all software components on a single machine, or distribute components on different machines across your network, with the maxView Storage Manager GUI and Web server on one machine, and the and Redfish server on others.

2.2.1 About maxView Storage Manager Redfish Server

The maxView Storage Manager Redfish Server is an instance of Nodejs. On Windows and Linux systems, the Redfish Server manages the hardware, which monitors the controllers in your system and provide notifications to the maxView Storage Manager. The maxView Storage Manager Redfish Server is installed automatically with the maxView Storage Manager.

2.2.2 About the maxView Storage Manager Web Server

The maxView Storage Manager Web Server is an instance of the open-source Apache Tomcat servlet container. It runs the maxView Storage Manager Web application, and serves static and dynamic content to the maxView Storage Manager GUI. The maxView Web Server is installed automatically with the maxView Storage Manager GUI.

2.3 System Requirements

To install maxView Storage Manager, each system in your storage space must meet these requirements:

- PC-compatible computer with Intel Pentium processor, or equivalent
- At least 4 GB of RAM
- 350 MB of free disk drive space
- One of these operating systems:
 - Microsoft® Windows® Server, Windows SBS, Windows 10, Windows 8.1, Windows 7
 - Red Hat® Enterprise Linux
 - SuSE Linux Enterprise Server
 - Ubuntu Linux
 - CentOS
 - Hypervisors:
 - VMware vSphere, VMware ESXi
 - Citrix XenServer
 - Microsoft Hyper-V

See the *maxView Storage Manager and ARCCONF Command Line Utility Readme* for a complete list of supported operating system versions.

Note: maxView Storage Manager can also be used *before* an operating system is installed. See [Running maxView Storage Manager from a Bootable USB Image](#) on page 13 for more information.

2.4 Browser Support

To run the maxView Storage Manager GUI, each system in your storage space must be running one of these Web browsers:


- Microsoft® Edge browser for Windows 10
- Google® Chrome™ 32 or newer
- Mozilla Firefox® 31 or newer

2.5 Typical Storage Space Configurations

The following examples show typical storage spaces that you can build with maxView Storage Manager. You can grow your storage space as your requirements change by adding more systems, controllers, disk drives, and enclosures, and by adding redundant logical drives for protection against data loss.


2.5.1 A Simple Storage Space

This example shows a simple storage space that might be appropriate for a small business. This storage space includes one RAID controller and three disk drives installed in a server. For data protection, the disk drives have been used to build a RAID 5 logical drive.



2.5.2 An Advanced Storage Space


This example shows how you can grow your storage space as the requirements of your application change. On the first server, segments from each disk drive have been used to build two RAID 5 logical drives. A second server connected to two 12-disk enclosures has been added. The additional storage space has been used to create two RAID 50 logical drives. The Administrator of this storage space can create and modify logical drives and monitor both controllers, disk drives, and enclosures from a single system running the maxView Storage Manager GUI.



2.5.3 Continuing to Grow Your Storage Space

For more advanced applications, such as high-volume transaction processing in a "cloud" or data center environment, maxView Storage Manager helps you grow your storage space to include multiple controllers, storage enclosures, and disk drives in multiple locations.

In this example, multiple systems, servers, disk drives, and enclosures have been added to the storage space. The Administrator can create and modify logical drives and monitor all the controllers, enclosures, and disk drives in the storage space from any system running the maxView Storage Manager GUI.



3 Installing maxView Storage Manager

This section describes how to install and uninstall maxView Storage Manager on the supported operating systems. It also describes how to run maxView Storage Manager from a *bootable USB image*, before the application is installed on an operating system.

Note: To perform a silent installation without messages or user interaction, follow the instructions in [Silent Installation on Windows and Linux](#) on page 146.

Note: VMware users can install the maxView plugin for vSphere Web Client for Windows and monitor storage resources without using the maxView Storage Manager GUI. For more information, see [Using the maxView Plugin for VMware vSphere Web Client](#) on page 151.

3.1 Before You Begin the Installation

Complete the following steps before you begin the installation.

3.1.1 Gather Installation Information

Prepare the following information:

- Redfish Server port number. The default port is recommended (8081). If the default port is not available, another port number will be automatically assigned. For more information on the Redfish Server, see [About maxView Storage Manager Redfish Server](#) on page 3.
- maxView Web Server port number. The default port is recommended (8443). If the default port is not available, another port number will be automatically assigned. For more information on the Web Server, see [About the maxView Storage Manager Web Server](#) on page 3.

Note: You can install maxView Storage Manager over an existing installation if it is no more than two versions older than the current release. Otherwise, you must remove the old version first, before beginning a new installation. See [Uninstalling maxView Storage Manager](#) on page 14 for details.

3.1.1.1 Check Network Configuration

Check your network configuration to ensure that it meets the prerequisites for a standard (non-Standalone Mode) installation:

- Ensure that the system is configured with an IP address.
- Ensure that the OS hostname is per standard.
- Ensure that the hostname-to-IP address mapping is updated in DNS. At minimum, ensure that the hostname-to-IP mapping is entered in the `/etc/hosts` file.

3.1.2 Download the Installation Package

Complete these steps to download the installation package for your operating system(s):

1. Open a browser window, then type `start.microsemi.com` in the address bar.
2. Select your controller family and controller model.
3. Select **Storage Manager Downloads**, then select the appropriate installer package from the list; for instance, maxView Storage Manager for Windows x64 or maxView Storage Manager for Linux.
4. Click **Download Now** and accept the license agreement.

5. When the download completes, extract the package contents to a temporary location on your machine.

Note: See the *Release Notes* for a complete list of installer packages for the supported operating systems.

3.2 Installing on Windows

This section describes how to install maxView Storage Manager on Windows systems.

Note: You need administrator privileges to install maxView Storage Manager. For details on verifying privileges, refer to your operating system documentation.


1. Open Windows Explorer or My Computer, then change to the directory where the Windows installer package is located (see [Download the Installation Package](#) on page 8 for details).

2. Double-click the setup program for your operating system version:

Option	Description
Windows 64-bit	setup_asm_x64.exe

The Installation wizard opens.


3. Click **Next** to begin the installation.
The License Agreement screen on the Installation wizard appears.
4. Select **I accept the terms in the license agreement** option, then click **Next**.
5. Accept or modify the default server ports in the maxView Storage Manager Configuration screen:
 - a. Web Server Port: 8443 (default)
 - b. Redfish Server Port: 8081 (default)



6. To *disable* remote system management from the GUI, click the **Standalone Mode** check box.

Note: In Standalone mode, maxView Storage Manager displays the system name as "localhost" and events as "127.0.0.1/localhost".

7. Click **Next**, then click **OK** to verify the Web Server port and the Redfish Server port numbers. The **Direct Attached Storage Setup** screen appears on the Installation wizard.
8. Ensure that **GUI and/or Redfish Server** is selected. Optionally, select **CLI Tools**. Click **Next**.



9. Click **Install** to begin the installation.

10. Repeat these steps to install maxView Storage Manager on every Windows system that will be part of your storage space.

When the installation is complete you receive a confirmation message and the maxView Storage Manager icon is placed on your desktop.

3.3 Installing on Red Hat, Citrix XenServer, CentOS, or SuSE Linux

This section describes how to install maxView Storage Manager on systems running Red Hat Linux, CentOS, XenServer, or SuSE Linux. For a list of supported Linux operating systems, see [System Requirements](#) on page 4.

1. Open a shell window, then change to the directory where the Linux installer package is located (see [Download the Installation Package](#) on page 8 for details).

2. Run the `.bin` file for your operating system version (x.xx-xxxx=version-build number):

Option	Description
Linux 64-bit	<code>./StorMan-X.XX-XXXXX.x86_64.bin</code>

3. When prompted for configuration details, enter the following:

Standalone Mode: [default: No]

Note: Standalone Mode *disables* remote system management from the GUI. maxView Storage Manager displays the system name as "localhost", and events as "127.0.0.1/localhost".

4. Repeat these steps to install maxView Storage Manager on every Linux system that will be part of your storage space.

When the installation completes a confirmation message is displayed and the maxView Storage Manager icon is placed on your desktop.

3.4 Installing on Debian or Ubuntu Linux

This section describes how to install maxView Storage Manager on systems running Debian or Ubuntu Linux.

1. Open a shell window, then change to the directory where the Linux installer package is located (see [Download the Installation Package](#) on page 8 for details).

2. Install the .deb package for your operating system version (x.xx-xxxxx=version-build number).

Option	Description
Linux 64-bit	dpkg -i StorMan-X.XX-XXXXX_amd64.deb

3. When prompted for configuration details, enter the following:

Standalone Mode: [default: No]

Note: Standalone Mode *disables* remote system management from the GUI. maxView Storage Manager displays the system name as "localhost", and events as "127.0.0.1/localhost".

4. Repeat these steps to install maxView Storage Manager on every Debian and Ubuntu Linux system that will be part of your storage space.
5. Before upgrading/re-installing maxView Storage Manager on an existing Ubuntu/Debian installation, enable the upgrade switch before installing the maxView .deb package:


```
export maxView_Upgrade=true
dpkg -i StorMan-*.deb
```

When the installation is complete you receive a confirmation message and the maxView Storage Manager icon is placed on your desktop.

3.5 Installing on VMware

Use the following procedure to install the .vib files for a VMware ESXi system. Perform the installation from a remote system running a Telnet/SSH client. Use a terminal emulator to access the ESXi server remotely.

1. Copy the following files from the installer download location to the /tmp directory on your local ESXi server (see [Download the Installation Package](#) on page 8 for details):
 - vmware-esx-provider-arconf.vib
 - vmware-esx-provider-arc-cim-provider.vib

The arconf.vib is for command line communication.

The arc-cim-provider.vib is for remote management communication.

2. Check for existing installation of arconf.


```
esxcli software vib list | grep arconf
```
3. Remove the existing arconf package.


```
esxcli software vib remove -n arconf
```

 When the package is removed, you receive the message "Reboot Required: false."
4. Check for an existing installation of arc-cim-provider.


```
esxcli software vib list | grep arc-cim-provider
```
5. Remove the existing arc-cim-provider package.


```
esxcli software vib remove -n arc-cim-provider
```

 When the package is removed, you receive the message "Reboot Required: true."
6. Set the installation acceptance level to either PartnerSupported or CommunitySupported:


```
esxcli software acceptance set --level=PartnerSupported
```
7. Install the arconf package.


```
esxcli software vib install -v /tmp/vmware-esx-provider-arconf.vib
```

 When the package is installed, you receive the message "Reboot Required: false."
8. Install the arc-cim-provider package.


```
esxcli software vib install -v /tmp/vmware-esx-provider-arc-cim-provider.vib
```

 When the package is installed, you receive the message "Reboot Required: true."
9. Reboot the system.

10. Continue with *Managing Remote Systems* on page 130 to add a remote system.

3.6 Running maxView Storage Manager from a Bootable USB Image

Running maxView Storage Manager from a *bootable USB image* allows you to configure your controller before you install the operating system. The procedure consists of three basic steps:

1. Download the bootable USB image from the Microsemi web site
2. Create a "live" image on a USB flash drive
 - Note:** We recommend using Rufus bootable USB create (<http://rufus.akeo.ie/>).
3. Boot from the USB flash drive, login to maxView Storage Manager and configure your controller

The bootable USB image is not a substitute for running maxView Storage Manager as an installed application. Many of the features and functions described in this guide are not available when you run maxView Storage Manager from a bootable USB image. Use the bootable USB image only to configure your controller before you install an operating system.

Note: Before you begin, ensure that your system is set up to boot from a USB drive. Check the system BIOS to see if the USB drive is included in the boot sequence. (For more information, see your system's documentation.) You will need a USB drive with at least 2 GB of storage to complete this task. To run the bootable USB image, the target machine must have at least **4 GB** of memory.

To run maxView Storage Manager from a bootable USB image:

1. Download the bootable USB image:
 - a. Open a browser window, then type `start.microsemi.com` in the address bar.
 - b. Select your controller family and controller model.
 - c. Select **Storage Manager Downloads**.
 - d. Download the bootable USB image (zip file archive).
 - e. Extract the contents of the bootable image archive file to a temporary location. The archive contains one file: the maxView Storage Manager bootable iso image.
2. Create a "live" image on the USB drive:
 - a. Run the USB Creator utility setup program at <http://rufus.akeo.ie/>.
 - b. Start USB Creator from the Windows All Programs menu.
 - c. In the Use Existing Live CD field, click **Browse**, then locate and select the maxView Storage Manager bootable ISO image.
 - d. In the Target Device field, select the USB flash drive (e:\, for instance).
 - e. Click **Create Live USB**.
3. Insert the USB drive on the machine you want to configure. The Boot menu opens in a shell window.
4. Select **Launch maxView** from the menu. After a minute or so, the maxView Storage Manager login screen opens in a browser window.
 - Note:** If you prefer to configure the controller from the command line, select **Launch arconf** from the Boot menu, then enter `root`, with no password, for the login credentials.
5. Enter `root/root` for the login credentials.
6. Continue with *Creating Arrays and Logical Drives*.

3.7 Uninstalling maxView Storage Manager

To uninstall maxView Storage Manager, follow the instructions for your operating system.

3.7.1 Uninstalling from Windows

To uninstall maxView Storage Manager from a Windows system, use the Add or Remove Programs tool in the Control Panel. All maxView Storage Manager components are uninstalled.

When the uninstall process is complete, you receive a confirmation message and the maxView icon is removed from your desktop.

3.7.2 Uninstalling from Red Hat, Citrix XenServer, CentOS, or SuSE Linux

This section describes how to uninstall maxView Storage Manager from systems running Red Hat, XenServer, CentOS, SuSE Linux.

1. Type the command `rpm -e StorMan`

When the uninstall process is complete, you receive a confirmation message and the maxView icon is removed from your desktop.

3.7.3 Uninstalling from Ubuntu Linux

This section describes how to uninstall maxView Storage Manager from systems running Ubuntu Linux.

1. Type the command `dpkg -r StorMan`
2. Type the command to uninstall maxView after the upgrade `export maxView_Upgrade=false dpkg -r storman`

When the uninstall process is complete, you receive a confirmation message and the maxView icon is removed from your desktop.

3.7.4 Uninstalling from VMware

Use the following procedure to remove maxView Storage Manager from a VMware ESXi 5.5 or 6.0 system.

1. Log in with the user name: root
2. List the installed packages:

```
esxcli software vib list | grep arconf
esxcli software vib list | grep arc-cim-provider
```
3. Remove the arconf package:

```
esxcli software vib remove -n arconf
```
4. Remove the arc-cim-provider package:

```
esxcli software vib remove -n arc-cim-provider
```
5. Reboot the system.

To verify that maxView Storage Manager is uninstalled, repeat Step 2. If no results are shown, the software was uninstalled successfully.

4 Exploring maxView Storage Manager


This section familiarizes you with the main features of the maxView Storage Manager graphical user interface. It describes how to start and log into maxView Storage Manager. It also explains how to get help and log out of maxView Storage Manager when you are finished working with the application.

4.1 Starting maxView Storage Manager and Logging In

The procedure for starting and logging in to maxView Storage Manager is the same for all operating systems with a graphical desktop. You can login as the Administrator, with full management-level access to your storage space, or as a Standard user, with restricted access to your storage space (see *Working in maxView Storage Manager* on page 15 for more information about access permissions).

1. On the desktop, double-click the maxView Storage Manager desktop icon.

The login window opens in the default browser.



Note: If you do not have an icon for maxView Storage Manager on your desktop, open a browser window, then type this URL in the address bar and press **Return**:
<https://127.0.0.1:8443/maxview/manager/login.xhtml>.

2. For full management-level access to your storage space, enter the Administrator account username and password for your operating system. For Standard-level access to your storage space, enter your regular network login credentials. Then click **Login**.

The maxView Storage Manager main window opens.

4.2 Working in maxView Storage Manager

You can perform most tasks in maxView Storage Manager by:

- Selecting storage components in the Enterprise View (controllers, hard drives, logical drives, and so on)
- Clicking icons on the *ribbon*, at the top of the maxView Storage Manager main window
- Working with information in the *Storage Dashboard* and *Chart View*
- Checking status in the Event Log and Task Log

If you are logged in as the Administrator, you have full access to manage and modify the components of your storage space, using all of the features of maxView Storage Manager. If you are logged in as a Standard user, you have restricted "view-only" access to your storage space, with limited ability to perform non-destructive operations, as described in the table below.

Note: maxView Storage Manager allows you to give Standard users Administrator privileges. For details, see *Granting Standard Users Admin Privilege* on page 139.


Standard users can:	Standard users can't:
Rescan controllers	Create arrays and logical drives
Save activity logs	Modify arrays and logical drives
Identify physical devices, logical devices, and enclosures	Delete arrays and logical drives
Silence alarms	Perform data migrations
View component properties on the Storage Dashboard	Clear the controller configuration

4.3 Overview of the Main Window

The main window of maxView Storage Manager has three main panels—left, right, and bottom—plus the ribbon, at the top of the window.

The left panel always shows the Enterprise View. The bottom panel shows the Event Log and Task Log. The right panel shows the Storage Dashboard and Chart View. Different information appears in the right panel depending on which component is selected in the Enterprise View.

In example below, a controller is selected in the Enterprise View, and the right panel displays the Storage Dashboard for the controller, with a chart view of its storage space.




You can resize the panels and scroll horizontally or vertically as needed, to view more or less information.

4.3.1 The Enterprise View


The Enterprise View is an expandable “tree” that shows the physical and logical components of your storage space. The Enterprise View lists the local system (the system you’re working on) and any remote systems that you have logged in to from the local system. (See *‘Local’ or ‘Remote’?* on page 24 for more information.) It also lists the *maxCache Devices* in your system.

Note: maxCache is not supported on all Microsemi Adaptec Smart Storage Controllers. See the Readme for more information. For more information about maxCache, see *Working with maxCache Devices* on page 70.



Expand a system in the Enterprise View to see its controllers, arrays, logical drives (“devices”), physical drives, enclosures, and maxCache Devices.

In the figure below, a controller is expanded in the Enterprise View, revealing the physical and logical devices associated with that controller.



You can perform most tasks in maxView Storage Manager by selecting a component in the Enterprise View, such as a controller or disk drive, then using the related commands on the *ribbon*, as described in the section below.

4.3.1.1 What do the Enterprise View Icons Mean?

Icon	Description
	System with controller and directly attached disk drives or enclosures
	Controller
	Enclosure

Icon	Description
	Logical drive (encrypted) ¹
	maxCache Device (healthy) ²
	Array (healthy)
	Hard disk drive
	Solid State Drive (SSD)
	SMR (Shingled Magnetic Recording) ³ drive
	Connector or other physical device

4.3.2 The Ribbon

Most tasks in maxView Storage Manager are available from the *ribbon*, at the top of the main window. The ribbon replaces toolbars and menus in maxView Storage Manager to help you quickly find the commands you need to complete a task.



The ribbon is organized into groups of related tasks for Systems, Controllers, Arrays, Logical Devices, Physical Devices, and maxCache Devices. The Home group (on the left) provides commands for working with remote systems (see [Managing Remote Systems](#) on page 130). Active options on the ribbon vary, depending on which type of component is selected in the Enterprise View.

For instance, if you select a controller in the Enterprise View, these options are activated:

- Create Logical Drive in the Logical Device group
- Spare Management in the Physical Device group
- Create maxCache Device in maxCache group (if the controller supports maxCache)
- All options in the Controller group

If you select an array in the Enterprise View, options in the Array group are highlighted; selecting a disk drive highlights options in the Physical Device group; and so on.

For a description of the icons on the ribbon, see [Icons At-a-Glance](#) on page 172.

4.3.3 The Storage Dashboard

When you select a component in the Enterprise View, maxView Storage Manager displays detailed information about that component on the *Storage Dashboard*. Occupying the largest portion of the main


¹ A lock in the Enterprise View means that the device is encrypted. For more information, see [Working with maxCrypto Devices](#) on page 74.

² A green check mark in the Enterprise View means that the device is healthy with no problems or issues. For more information, see [Identifying a Failed or Failing Component](#) on page 140.

³ Not supported on all controllers. See the Readme for more information.

window in maxView Storage Manager, the Storage Dashboard provides status information, physical and logical device properties, resources, usage statistics, and reliability indicators for hard drives and SSDs. It also provides a chart view of free and used space in your system.

Tabs provide quick access to component information



The screenshot shows the Storage Dashboard interface with several tabs: Summary, Properties, Resources, Connectors, maxCrypto, and Events. The 'Resources' tab is active, displaying a list of resources categorized by type and state. To the right, the 'Total Storage' section features a pie chart showing 78% free space (12.954 TB) and 22% used space (3.829 TB).

Controller Info

Model	MSCC Adapter SmartRAID 3162-01-w
Status	Optimal
Serial Number	5A12F3037DA
WWN	500001E0C210000
Firmware Version	2.09(2)
Hardware Revision	0
Manufacturing Part Number	Not Applicable
Manufacturing Spare Part Number	Not Applicable
Manufacturing Wellness Log	Not Applicable
Indirect Memory	2048 MB
Cache Memory	1024 MB
Mode	Wired
Pending Mode	Not Applicable
Temperature	31°C / 122°F (Normal)
Power Consumption	Not Available
NVRAM Checksum Status	Passed

Resources

Array(s)	3
Data Array(s)	3
Cache Array(s)	0
Logical Device(s)	3
Data Logical Device(s)	3
Cache Logical Device(s)	0
Physical Device(s)	20
By Type	
Hard Disk Drive(s)	18
Solid State Drive(s)	2
NVMe	0
Others	0
By State	
Data Drive(s)	4
Spare Drive(s)	0
Unassigned Drive(s)	16
Transient Data Drive(s)	0
Erase Completed Drive(s)	0

Total Storage

Chart view of free and used space in your system

Free Space (12.954 TB)
Used Space (3.829 TB)

For more information about the types of information provided on the Storage Dashboard for each component in your storage space, see [Viewing Component Information and Status in the Storage Dashboard](#); also see [Revealing More Device Information](#).

4.4 Checking System Status from the Main Window


maxView Storage Manager includes an Event Log and Task Log for at-a-glance status and activity information for all managed systems. The Event Log provides status information and messages about activity (or *events*) occurring in your storage space. The Task Log provides information about current or recurring processes in your storage space, such as rebuilding a logical device. Single-click any event or task to see more information in an easier-to-read format. For more information about the Event Log and Task Log, see [Viewing Activity Status in the Event Log](#) on page 112 and [Working with Tasks](#).

Click this tab to open the Task Log
Click this tab to open the Event Log



Severity	Date and Time	System	Source	Component	Description
Information	Wed Feb 06 2019 14:21:19 EST	WINAN/RD0LE	Controller 1	Controller 1	Controller Reduced Priority changed from High to Medium-high controller 1 (MSCC Adapter SmartRAID 3162-01-w #5A12F3037DA Physical Slot 3)
Information	Wed Feb 06 2019 14:21:19 EST	WINAN/RD0LE	Controller 1	Controller 1	Controller unassigned drive write cache policy setting changed from Default to Enabled controller 1 (MSCC Adapter SmartRAID 3162-01-w #5A12F3037DA Physical Slot 3)
Information	Wed Feb 06 2019 14:21:19 EST	WINAN/RD0LE	Controller 1	Controller 1	Controller unassigned drive write cache policy setting changed from Default to Enabled controller 1 (MSCC Adapter SmartRAID 3162-01-w #5A12F3037DA Physical Slot 3)
Information	Wed Feb 06 2019 14:21:19 EST	WINAN/RD0LE	Controller 1	Controller 1	Controller Expand Priority changed from Medium to High controller 1 (MSCC Adapter SmartRAID 3162-01-w #5A12F3037DA Physical Slot 3)
Information	Wed Feb 06 2019 14:21:19 EST	WINAN/RD0LE	Controller 1	Controller 1	Consistency check inconsistency notify is enabled controller 1 (MSCC Adapter SmartRAID 3162-01-w #5A12F3037DA Physical Slot 3)

Warning- and Error-level icons appear next to components in the Enterprise View affected by a failure or error, creating a trail, or *rapid fault isolation*, that helps you identify the source of a problem when it occurs. See [Identifying a Failed or Failing Component](#) on page 140 for more information.



If your storage space includes a drive enclosure with a temperature sensor, temperature, fan, and power module status is displayed on the Storage Dashboard (see [Monitoring Enclosure Status](#) on page 115).


For more information about checking status from the main window, see [Monitoring Status and Activity](#).

4.5 Revealing More Device Information

You can reveal more information about disk drive, array, and logical drive usage in your storage space (including maxCache Devices) with the Resources view on the Storage Dashboard.

To reveal disk drive usage by logical drive (and vice-versa), select a controller in the Enterprise View, then open the **Resources** tab on the Storage Dashboard. As shown in the figures below, you can click a logical drive to see its member disk drives and spares; similarly, you can click a physical disk to see which array (if any) it belongs to. In the following figure the disk in Slot 5 belongs to Array A.

Note: Click the Arrow icons, on the right side of the Resources table, to jump to that resource in the Enterprise View tree.




4.6 Getting Help

maxView Storage Manager provides online help that includes conceptual information and descriptions of on-screen items and dialog boxes, in addition to step-by-step instructions for completing tasks.


To open the online help, click the **Help** button at the upper-right corner of the main window.



For help with a dialog box or wizard, click the question-mark icon, in the lower corner of the dialog box, for help with that specific procedure.



For help with individual options in the Set Properties dialog box (for controllers, logical drives, and physical drives), or specific information fields on the Storage Dashboard, mouse over any field or option name for a brief description of that option.



4.7 Logging Out of maxView Storage Manager

To log out of maxView Storage Manager:

1. In the Enterprise View, click on the local system.
2. Click the **Logout** button at the upper-right corner of the main window:



Click here to log out

You are logged out of maxView Storage Manager and the main window is closed.

5 Building Your Storage Space

Follow the instructions in this section to choose a management system, log in to each system in your storage space, and create arrays and logical drives.

Note: Before beginning the tasks in this chapter, ensure that maxView Storage Manager is installed on every system that will be part of your storage space.

5.1 Overview

To build your storage space, complete these steps:

1. Choose at least one management system (see [Choosing a Management System](#)).
2. Start and log in to maxView Storage Manager on the management system (see [Starting maxView Storage Manager and Logging In](#) on page 15).
3. Log in to all other systems from the management system (see [Logging into Remote Systems from the Local System](#) on page 25).
4. Create arrays and logical drives for all systems in your storage space (see [Creating Arrays and Logical Drives](#)).

As your storage requirements change, you can add systems, controllers, and disk drives, then modify the arrays and logical drives in your storage space by following the instructions in [Modifying Your Storage Space](#) on page 52.

5.2 Choosing a Management System


You must designate at least one system as a *management system*, a system from which you will manage the storage on all systems in your storage space.

The management system can be any system on your network that has a video monitor and can run the maxView Storage Manager GUI and Web server.

5.2.1 ‘Local’ or ‘Remote’?

Whenever you're working in maxView Storage Manager, the system that you're working on is the *local* system. All other systems in your storage space are *remote* systems. ‘Local’ and ‘remote’ are relative terms, as shown in the following figure—when you are working on system A (local system), system B is a remote system; when you are working on system B (local system), system A is a remote system.

For the purposes of this guide, the ‘local system’ is the management system.



5.2.2 Logging in on the Local System

To log in on the local system, see [Starting maxView Storage Manager and Logging In](#) on page 15.

5.3 Logging into Remote Systems from the Local System

Once maxView Storage Manager is running on all systems in your storage space, you can log into the remote systems from the local system.

Once you have logged in to a remote system, it automatically appears in the Enterprise View each time you start maxView Storage Manager on the local system. You can work with a remote system's controllers, disk drives, and logical drives as if they were part of your local system.

To log in to a remote system:

1. On the ribbon, in the Home group, click **Add System**.




The Add System window opens, showing a list of "discovered" systems; that is, systems on your network that are running the Redfish.

Note:

The list of discovered systems appear only when Auto Discovery option is enabled in maxView. For more details on how to change the auto-discovery settings, see [Changing the Auto-Discovery Settings](#) on page 132.

2. Select the systems you want to add to the Enterprise View, then enter the systems' login credentials (username/password) in the space provided. The **Single Sign-On** option gets enabled if more than one system is selected. Also, ensure that the selected systems should have same login credentials.



<input type="checkbox"/>	System Name/IP Address	User Name	Password
<input type="checkbox"/>	10.187.74.157		
<input type="checkbox"/>	10.187.74.46		
<input type="checkbox"/>	10.187.74.155		
<input type="checkbox"/>	10.187.74.25		
<input type="checkbox"/>	10.187.74.69		
<input type="checkbox"/>	10.187.74.123		
<input type="checkbox"/>	10.187.74.156		
<input type="checkbox"/>	10.187.75.48		

Note: You can add a system manually if you don't see the system in the list. For more information, see [Manually Adding a Remote System](#).

3. Click **Add**.

maxView Storage Manager connects to the remote system(s) and adds them to the list of managed systems in the Enterprise View.

For more information about working with remote systems, see [Managing Remote Systems](#).

5.4 Creating Arrays and Logical Drives

maxView Storage Manager provides a wizard to help you create, or *configure*, the arrays and logical drives in your storage space. You can choose from two configuration methods:

- Create logical drive on new array—Helps you set the RAID level for the logical drive, group disk drives and SSDs, determine logical drive size and other advanced settings.

For instructions, see [Creating a Logical Drive on a New Array](#).

- Create logical drive on existing array—Helps you select an array on which to create the logical drive, set the RAID level, group disk drives and SSDs, determine logical drive size and configure advanced settings.

For instructions, see [Creating a Logical Drive on an Existing Array](#)

If maxCrypto is enabled, you can create encrypted or plaintext volumes. (For more information, see [Working with maxCrypto Devices](#) on page 74.)

Note:

1. Mixing SAS and SATA drives within the same logical drive is not supported. The wizard does not allow you to select a combination of SAS and SATA drive types.
2. maxView Storage Manager supports SMR HA⁴ and SMR DM drives for all RAID levels. However, mixing SMR and PMR⁵ drives within the same logical drive is not supported. maxView Storage Manager displays a warning message if you try to create a logical drive using a combination of SMR and PMR device types.

5.4.1 Creating a Logical Drive on a New Array

Before you can create a logical drive, you must create an array. Use the On New Array configuration method to step through the process of creating a logical drive on a new array, setting the RAID level, and configuring other settings.

To create a logical drive on an existing array, see [Creating a Logical Drive on an Existing Array](#).


By default, maxView Storage Manager uses all available disk space to maximize the capacity of a new logical drive.

To create a logical drive on a new array:

1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Logical Device group, click **Create Logical Device**.




3. When the wizard opens, select **On New Array**, then click **Next**.



⁴ SMR: Shingled Magnetic Recording. HA: Host Aware (backward compatible with standard HDD). DM: Device Managed (backward compatible with standard HDD).


⁵ PMR: Perpendicular Magnetic Recording; standard HDD recording technology.

4. Select a RAID level for the logical drive, then click **Next**.




Note: Not all RAID levels are supported by all controllers. (See the Release Notes for more information.) See [Selecting the Best RAID Level](#) for more information about RAID levels.

5. Select the disk drives you want to include in the logical drive, then click **Next**. Be sure the drive type is the same for all drives (SAS or SATA, not mixed), and that you select the right number of drives for the RAID level you selected.



6. (Optional) In the RAID Attributes panel, customize the logical drive settings.



The screenshot shows the 'Create Logical Device' window for Controller 1. The settings are as follows:


Property	Value
Name	LogicalDrv 4
Size	1862.964 GB (Max checked)
Stripe Size (KB)	256
Controller Caching	Enabled
Initialization Method	Default
Encrypt Data	Checked

You can:

- Enter a name for the logical drive. Names can include any combination of letters, numbers, and spaces.
- Set the size and unit of measure for the logical drive. (By default, a new logical drive uses all available disk space.)
- Change the stripe size—the amount of data, in bytes, written per disk in the logical drive. (The default stripe size usually provides the best performance.)
- Enable or disable controller caching.
- Set the initialization method to Default or Build. The initialization method determines how the logical drive is prepared for reading and writing, and how long initialization will take:
 - **Default**—Initializes parity blocks in the background while the logical drive is available for access by the operating system. A lower RAID level results in faster parity initialization.
 - **Build**—Overwrites both the data and parity blocks in the foreground. The logical drive remains invisible and unavailable to the operating system until the parity initialization process completes. All parity groups are initialized in parallel, but initialization is faster for single parity groups (RAID 5). RAID level does not affect performance during Build initialization.
- Create an encrypted or plaintext logical drive (for more information, see [Working with maxCrypto Devices](#) on page 74)

7. Click **Next**, then review the array and logical drive settings.

This example shows a RAID 5 logical drive ready to be created on Array A.



8. Click Finish.

maxView Storage Manager builds the array and logical drive. Use the Event Log and Task Log to track build progress.

- 9.** If you have other disk drives or available disk space and want to create additional arrays on the controller, repeat Steps [2]-[8].
- 10.** Repeat Steps [1]-[9] for each controller in your storage space.
- 11.** Partition and format your logical drives. See [Partitioning and Formatting Your Logical Drives](#) on page 33.

5.4.2 Creating a Logical Drive on an Existing Array

After you create an array, you can continue to build your storage space by creating more logical drives on that array. Use the On Existing Array configuration method to step through the process of creating a logical drive on an existing array, setting the RAID level, and configuring other settings.

To create a logical drive on a new array, see [Creating a Logical Drive on a New Array](#).

By default, maxView Storage Manager uses all available disk space to maximize the capacity of a new logical drive.

Note:


You can also add/create Logical drives by selecting the existing array from the Enterprise view.

To create a logical drive on an existing array:


1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Logical Device group, click **Create Logical Device**.




3. When the wizard opens, select **On Existing Array**, then click **Next**.



4. Select the array on which to create the logical drive, then click **Next**.



5. Select a RAID level for the logical drive, then click **Next**.




The screenshot shows the 'Create Logical Device' wizard for Controller 1. The 'RAID Level' step is active, displaying a list of RAID options with their descriptions. The 'RAID 0' option is selected.

RAID Level	Description
<input checked="" type="radio"/> RAID 0	Striping; 1 or more drives required; not redundant
<input type="radio"/> Arrays with RAID 0	Striping; Multiple array with one drive each; not redundant
<input type="radio"/> RAID 1	Mirroring; 2 drives required; redundant
<input type="radio"/> RAID 1 (ADM)	Advanced data mirroring; 3 drives required; redundant
<input type="radio"/> RAID 5	Striping with parity; 3 or more drives required; redundant
<input type="radio"/> RAID 6	Striping with dual parity; 4 or more drives required; redundant
<input type="radio"/> RAID 10	Stripped mirrors; even number of drives required; redundant
<input type="radio"/> RAID 10 (ADM)	Stripped mirrors in trios; Multiples of three drives required; redundant
<input type="radio"/> RAID 50	RAID 5 groups stripped together; 6 or more drives required; redundant
<input type="radio"/> RAID 60	RAID 6 groups stripped together; 8 or more drives required; redundant

Note: Not all RAID levels are supported by all controllers. (See the Release Notes for more information.) See [Selecting the Best RAID Level](#) for more information about RAID levels.

6. (Optional) In the RAID Attributes panel, customize the logical drive settings.



The screenshot shows the 'Create Logical Device' wizard for Controller 1, now in the 'RAID Attributes' step. The 'RAID Level' is confirmed as RAID 0. The following settings are visible:

- Name:** LogicalDrv 4
- Size:** 1862.964 GB (Max checked)
- Stripe Size (KB):** 256
- Controller Caching:** Enabled
- Initialization Method:** Default
- Encrypt Data:** Checked

You can:

- Enter a name for the logical drive. Names can include any combination of letters, numbers, and spaces.
- Set the size and unit of measure for the logical drive. (By default, a new logical drive uses all available disk space.)
- Change the stripe size—the amount of data, in bytes, written per disk in the logical drive. (The default stripe size usually provides the best performance.)
- Enable or disable controller caching.
- Set the initialization method to Default or Build. The initialization method determines how the logical drive is prepared for reading and writing, and how long initialization will take:


- **Default**—Initializes parity blocks in the background while the logical drive is available for access by the operating system. A lower RAID level results in faster parity initialization.
- **Build**—Overwrites both the data and parity blocks in the foreground. The logical drive remains invisible and unavailable to the operating system until the parity initialization process completes. All parity groups are initialized in parallel, but initialization is faster for single parity groups (RAID 5). RAID level does not affect performance during Build initialization.

Note: Not all initialization methods are available for all RAID levels.

- Create an encrypted or plaintext logical drive (for more information, see [Working with maxCrypto Devices](#) on page 74)

7. Click **Next**, then review the array and logical drive settings.

This example shows a RAID 5 logical drive ready to be created on Array A.



8. Click **Finish**.

maxView Storage Manager builds the logical drive on the array. Use the Event Log and Task Log to track build progress.

9. If you have other disk drives or available disk space and want to create more logical drives on an existing array, repeat Steps [2]-[8] .
10. Repeat Steps [1]-[9] for each controller in your storage space.
11. Partition and format your logical drives. See [Partitioning and Formatting Your Logical Drives](#) on page 33.

5.4.3 Partitioning and Formatting Your Logical Drives

The logical drives you create appear as physical disk drives on your operating system. You *must* partition and format these logical drives before you can use them to store data.

Note: Logical drives that have not been partitioned and formatted cannot be used to store data.

Refer to your operating system documentation for more information.

5.4.4 Creating Logical Drives on Other Systems in Your Storage Space

If you installed maxView Storage Manager and Smart Storage controllers on more than one system, continue building your storage space as follows:


- From each individual system, log in to maxView Storage Manager and repeat the steps to create logical drives on new or existing arrays, *or*
- From your *local* system (the system you're working on), log in to all other systems in your storage space as *remote* systems (see [Logging into Remote Systems from the Local System](#)), then repeat the steps to create logical drives on new or existing arrays, *or*
- From your local system, create a *server template file* and deploy the configuration to the remote systems in your storage space (see [Deploying Servers Server Template File](#)).

5.4.5 Controller Support for 4K Drives

This section describes how to use the maxView GUI with 4K drives to create and modify logical drives and spares.


5.4.5.1 Creating a Logical Drive

You can create a logical device using 4K drives. 512-byte drives cannot be mixed with 4K drives. This can be done by selecting the **Device Type** as HDD SATA 4K or HDD SAS 4K. This will ensure that only HDD SATA 4K or HDD SAS 4K devices are displayed.




5.4.5.2 Moving a Logical Drive


A 4K SAS or 4K SATA logical device can be moved to another array of 4K SAS or 4K SATA drives, but cannot be moved to an array with 512-byte drives.



- Moving to a new array: all SATA and SAS 4K drives that are available to move to a new array are listed.




- Moving to an existing array: if the logical device has already been created in a different array using 4K drives, then the user can move a logical device to the existing array of the same block size SAS/SATA 4K drives. Only arrays created using 4K drives will be listed (512-byte arrays will not be listed).




5.4.5.3 Modifying a Logical Drive


Arrays created using 4K drives can be modified.



- Moving drive(s): you can move a drive from one array to another array that uses the same interface type. For example, if an array is created using 4K SATA drives, then you can move a drive(s) from that array to a separate array that also uses 4K SATA drives.




- Changing drive types: you can change the drive interface type from SAS to SATA or from SATA to SAS. For example, if an array is created using 4K SAS drives, you can change the drive type to 4K SATA drives only.




5.4.5.4 Assigning Spares at the Array Level

Spares for 4K logical drives can be assigned at the array level.




1. Dedicated Hot Spare: If the array/logical device is created using 4K SATA drives, then only the 4K SATA devices can be assigned as spares.
2. Auto Replace Hot Spare: The process is the same as the Dedicated Hot Spare.




5.4.5.5 Assigning Spares at the Physical Device Level

Spares for 4K logical drives can be assigned at the physical device level.



- If array/logical device is created with 4K SAS drives, then only logical devices that were created with 4K SAS drives are listed.



Note:

- maxCache cannot be created using 4K SATA drives.
- 512-byte maxCache cannot be assigned to 4K logical devices.
- Drive interface types and drive block sizes cannot be mixed.

For example, SATA drives and SAS drives of the same block size cannot be mixed; 512-byte drives and 4K drives of the same interface type cannot be mixed.

6 Protecting Your Data

In addition to standard RAID (RAID 0, RAID 1, RAID 5, RAID 10), controllers provide additional methods of protecting your data, including dedicated and auto-replace hot spare drives.


A *hot spare* is a disk drive or SSD (Solid State Drive) that automatically replaces any failed drive in a logical drive, and can subsequently be used to rebuild that logical drive. (For more information, see [Recovering from a Disk Drive Failure](#) on page 141.)

6.1 Dedicated Spare or Auto-Replace Spare?

A *dedicated* hot spare is assigned to one or more arrays. It will protect any redundant logical drive on those arrays.

After using a dedicated hot spare to rebuild a failed logical drive, data is moved back to its original location, using a process called *copyback*, once the controller detects that the failed drive has been replaced. Once the data is copied back, the hot spare becomes available again. You must create an array before you can assign a dedicated hot spare to protect it. To assign a dedicated hot spare, see [Assigning a Dedicated Hot Spare](#) on page 41.

An *auto-replace* hot spare is assigned to a specific array. It will protect any redundant logical drive on that array. After using an auto-replace spare to rebuild a failed logical drive, it becomes a permanent part of the array. You must create an array before you can assign an auto-replace hot spare to protect it. To assign an auto-replace hot spare, see [Assigning an Auto-Replace Hot Spare](#) on page 42.



6.2 Hot Spare Limitations

- Hot spares protect redundant logical drives only. To protect non-redundant logical drives, set the spare activation mode of the controller to predictive activation.
- You cannot create a hot spare from a disk drive that is already part of an array.
- You should select a disk drive that is at least as big as the smallest disk drive in the array that it might replace.
- You must designate a SAS hot spare drive for an array comprised of SAS disk drives, and a SATA hot spare drive for an array comprised of SATA disk drives.
- You can designate a SMR HA⁶ or SMR DM drive for all hot spare types. A SMR drive cannot protect a PMR drive,⁷ or vice-versa.

⁶ SMR: Shingled Magnetic Recording. HA: Host Aware (backward compatible with standard HDD). DM: Device Managed (backward compatible with standard HDD).

⁷ PMR: Perpendicular Magnetic Recording; standard HDD recording technology.

6.3 Assigning a Dedicated Hot Spare

A dedicated hot spare is assigned to one or more arrays. It will protect any redundant logical drive on those arrays.

Note: You must create the array before you can assign a dedicated hot spare to protect it.


To assign a dedicated spare:

1. In the Enterprise View, select a controller, an array on that controller, or a Ready physical drive.
2. On the ribbon, in the Physical Device group, click **Spare Management**.




The Spare Management wizard opens.


3. Select the **Dedicated** spare type, then click **Next**.



4. If you selected a physical drive in the Enterprise view, select the arrays you want to protect with a dedicated spare, then click **Next**.



- If you selected an array in the Enterprise view, select the physical drive(s) you want to dedicate as hot spares, then click **Next**. (See [Hot Spare Limitations](#) on page 40 for help selecting drives.)



- Review the summary of dedicated spares and protected arrays, then click **Finish**.

6.4 Assigning an Auto-Replace Hot Spare

An auto-replace hot spare is assigned to a specific array. After using an auto-replace spare to rebuild a failed logical drive, it becomes a permanent part of the array.

To assign an auto-replace hot spare to an array:

- In the Enterprise View, select an array on that controller.

Note: The auto-replace option is not available, if you select an array with a non-redundant logical device when the controller's "spare activation mode" is set to "failure activation". However, when you select a physical device itself, the option is


available only if one or more auto-replace spares already exist. Otherwise, you can just assign Dedicated spares in the wizard.

2. On the ribbon, in the Physical Device group, click **Spare Management**.




The Spare Management wizard opens.


3. Select the **Auto-Replace** spare type, then click **Next**.



4. If you selected a controller in the Enterprise view, select the array you want to protect with an auto-replace spare, then click **Next**.



5. Select the physical drive(s) you want to assign as auto-replace hot spares, then click **Next**. (See [Hot Spare Limitations](#) on page 40 for help selecting drives.)



6. Review the summary of auto-replace spares and protected arrays, then click **Finish**.

6.5 Removing a Hot Spare

You can remove a dedicated or auto-replace hot spare from an array. Removing the last hot spare from an array returns the drive to the Ready state.

You may want to remove a hot spare to:

- Make disk drive space available for another array or logical drive.
- Convert an auto-replace hot spare into a dedicated hot spare.
- Remove the 'hot spare' designation from a drive that you no longer want to use as a spare.


To remove a hot spare:

1. In the Enterprise View, select an array or an existing hot spare drive.
2. On the ribbon, in the Physical Device group, click **Spare Management**.




The Spare Management wizard opens.


3. Select **Un-Assign**, then click **Next**. (Un-Assign is preselected for an existing hot spare.)



4. If you selected a hot spare in the Enterprise view, select the array(s) from which to remove the spare, then click **Next**.



- If you selected an array in the Enterprise view, select the hot spare(s) to remove from the array, then click **Next**.



- Review the summary of affected hot spares and arrays, then click **Finish**.
If the spare protects only one array, it is deleted and the drive becomes available for other uses in your storage space. If the spare protects more than one array, it is removed from the selected array(s) but continues to protect the other arrays to which it is assigned.

6.6 Setting the Spare Activation Mode

The spare activation mode determines when a hot spare is used to rebuild a failed logical drive. You can choose to activate a spare when:

- A data drive fails; this is the default mode.
- A data drive reports a predictive failure (SMART) status.

In normal operations, the firmware starts rebuilding a failed logical drive with a spare only when a data drive fails. With the predictive failure activation mode, rebuilding can begin before the drive fails, reducing the likelihood of data loss.

The spare activation mode applies to all arrays on a controller.


To set the spare activation mode:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



The Set Properties window opens.

3. Click the **Data Protection** tab.
4. From the Spare Activation Mode drop-down list, select **Failure** (default) or **Predictive**, then click **OK**.



6.7 Controller Sanitize Lock Freeze/Anti-Freeze

The Sanitize Lock Freeze/Anti-Freeze feature provides the controller level of sanitize lock, which helps prevent accidental erasing of user data on the disk after initiating a sanitize command. To accomplish this, you have the option of applying a controller-wide Sanitize Lock Freeze/Anti-Freeze policy. The freeze and anti-freeze commands will be used to block and unblock the sanitize commands that would erase user data on the disk.


The sanitize lock feature has three options:

- Freeze: prevents the user from performing any sanitize erase operation.
- Anti-Freeze: locks the freeze command, enabling the user to perform any sanitize erase operation.
- None: enables the user to perform any sanitize erase operation

This is applicable only to SATA drives which support Sanitize Erase, Freeze, and Anti-Freeze.


To set the Sanitize Lock:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



The Set Properties window opens.

3. Click the **Data Protection** tab.
4. From the Sanitize Lock drop-down list, select one of the three following options: **None** (default), **Freeze**, or **Anti-Freeze**.



Note:


If the Sanitize Lock is set to any value other than **None**, the following warning message will be displayed in the menu header:

Changing the Sanitize Lock will require a reboot to apply the new state to the controller, and require all physical devices to be power cycled or hot-plugged for the lock state to be applied to the physical devices.

5. Click **OK**.

6.7.1 Sanitize Lock Property in Controller Node Properties Tab

The properties of the Sanitize Lock feature are displayed in the controller node properties tab as shown in the following screen capture.




The Sanitize Lock property will display the current setting in which the controller is operating.

When the Sanitize Lock property is changed in the Set Properties dialog, the pending Sanitize Lock property will show the changed value.

When the machine is rebooted, the pending Sanitize Lock value will be "Not Applicable", and the Sanitize Lock value will be set to the previous pending Sanitize Lock value.

6.7.2 Physical Device Sanitize Lock Freeze/Anti-Freeze

This feature is supported only on SATA drives that are connected to the controller. If the drive supports the Sanitize Lock Freeze feature, it may or may not support the Sanitize Lock Anti-freeze. Based on the support bit on the drive, the Sanitize Lock policy can be set from the controller and it will be applied on the drives that support Sanitize Freeze/Anti-Freeze.



The Sanitize Lock property is dependent upon the following conditions:

- If the drive does not support Sanitize Erase, the Sanitize Lock property is not displayed.
- If the drive supports Sanitize Erase but does not support Freeze/Anti-Freeze, then the Sanitize Lock property will be listed as "Not Applicable".
- If the controller Sanitize Lock is in the Freeze state, then no Sanitize Erase can be performed.
- If the controller Sanitize Lock is in the Anti-Freeze or None state, then all Sanitize Erase commands can be performed.

Once the controller Sanitize Lock is in the freeze state, then no Sanitize Erase operations will be listed during the secure erase operation.


6.7.3 Secure Erase Pattern

If the drive or controller Sanitize Lock is in the freeze state, then all the Sanitize Erase patterns will not be listed when you click on the Secure Erase ribbon icon in the physical device ribbon group.




Secure Erase

Secure erase completely destroys the data on the drive; the data is completely and irretrievably eradicated.



Only three secure erases can be performed.

If the drive and controller Sanitize Lock is in Anti-Freeze or None states, then the Sanitize Erase pattern will be listed.



Note: When you perform the Sanitize Erase operation, it sets the controller Sanitize Lock to freeze, and reboots the system, the drive will remember the percentage completion for the Sanitize Secure Erase after the reboot. The freeze state will be applied only after the Sanitize Erase is completed and the sanitize erase operation cannot be stopped.

7 Modifying Your Storage Space


This section provides additional scenarios for creating and modifying arrays and logical drives. It explains how to check your logical drives for bad or inconsistent data; optimize controller and logical drive performance; move arrays and logical drives; and perform advanced operations, such as creating a split mirror backup array.

7.1 Understanding Arrays and Logical Drives

A *logical drive* is a group of physical disk drives that appears to your operating system as a single drive that can be used to store data.

The group of physical drives containing the logical drive is called a *drive array*, or just *array*. An array can contain several logical drives, each of a different size.

You can include the same disk drive in two different logical drives by using just a portion of the space on the disk drive in each, as shown in the following figure.



Disk drive space that has been assigned to a logical drive is called a *segment*. A segment can include all or just a portion of a disk drive's space. A disk drive with one segment is part of one logical drive, a disk drive with two segments is part of two logical drives, and so on. When a logical drive is deleted, the segments that comprised it revert to available space (or *free segments*).

A logical drive can include redundancy, depending on its RAID level. (See [Selecting the Best RAID Level](#) for more information.)

You can also protect your logical drives by assigning one or more hot spares to them. (See [Protecting Your Data](#) on page 40 for more information.)

7.2 Creating and Modifying Logical Drives

For basic instructions for creating logical drives, see [Building Your Storage Space](#) on page 24.

This section describes additional scenarios for creating logical drives:


- To create a logical drive from different-sized disk drives, see [Including Different-sized Disk Drives in a Logical Drive](#) on page 53.
- To create a logical drive using available *segments* of disk drives, see [Creating a Logical Drive Using Available Segments](#).

7.2.1 Including Different-sized Disk Drives in a Logical Drive

You can combine disk drives of different sizes in the same logical drive. If the logical drive includes redundancy, however, the size of each segment can be no larger than the size of the smallest disk drive. (See [Selecting the Best RAID Level](#) for more information about redundancy.)

Note: You cannot combine SAS and SATA disk drives and also different block size like 512 bytes or 4K within the same array or logical drive.

To create a logical drive with disk drives of different sizes, follow the instructions in [Creating a Logical Drive on a New Array](#). When the wizard displays the RAID Members panel, select different size drives, as shown in the figure below, then complete the wizard.



When the logical drive is created, check its resources on the Storage Dashboard: it should appear similar to the next figure, where a RAID 5 logical drive includes two disk drives of one size and one of another.

Physical Device(s)					
Device	State	Interface	Block Size	Total Size	
CN0: Enclosure 1: Slot 0	Optimal	SAS	512 Bytes	931.5 GB	
CN0: Enclosure 1: Slot 1	Failed	SAS	512 Bytes	931.5 GB	
CN0: Enclosure 1: Slot 2	Ready	SATA	512 Bytes	7.2 TB	
CN0: Enclosure 1: Slot 3	Ready	SAS	4 KB	1.8 TB	
CN0: Enclosure 1: Slot 4	Optimal	SATA	512 Bytes	3.6 TB	
CN0: Enclosure 1: Slot 5	Optimal	SATA	512 Bytes	7.2 TB	
CN0: Enclosure 1: Slot 6	Optimal	SATA	512 Bytes	931.5 GB	
CN0: Enclosure 1: Slot 7	Ready	SAS	512 Bytes	931.5 GB	
CN0: Enclosure 1: Slot 8	Optimal	SATA	512 Bytes	931.5 GB	
CN0: Enclosure 1: Slot 9	Ready	SATA	512 Bytes	111.7 GB	
CN0: Enclosure 1: Slot 10	Ready	SATA	512 Bytes	111.7 GB	
CN0: Enclosure 1: Slot 11	Ready	SATA	512 Bytes	931.5 GB	
CN1: Enclosure 1: Slot 1	Ready	SATA	512 Bytes	931.5 GB	
CN1: Enclosure 1: Slot 2	Ready	SATA	512 Bytes	931.5 GB	
CN1: Enclosure 1: Slot 3	Ready	SATA	512 Bytes	931.5 GB	
CN1: Enclosure 1: Slot 4	Ready	SATA	512 Bytes	931.5 GB	
CN1: Enclosure 1: Slot 5	Ready	SATA	512 Bytes	931.5 GB	
CN1: Enclosure 1: Slot 6	Ready	SATA	512 Bytes	931.5 GB	
CN1: Enclosure 1: Slot 7	Ready	SATA	512 Bytes	931.5 GB	
CN1: Enclosure 1: Slot 8	Ready	SATA	512 Bytes	931.5 GB	

RAID5 with different size disks. No segment is larger than smallest disk drive (931.5 GB, in this example).

7.3 Enabling Background Consistency Check

When background consistency check is enabled, maxView Storage Manager continually and automatically checks your logical drives for bad or inconsistent data, and then fixes any problems. Enabling consistency check ensures that you can recover data if a logical drive fails. The scanning process checks physical drives in fault-tolerant logical drives for bad sectors. It also verifies the consistency of parity data, if applicable. The available modes are High, Disable, Idle. If you select the Idle mode, you must also specify a delay value and parallel scan count.


To enable or disable background consistency check:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



The Set Properties window opens.

3. Click the **Data Protection** tab.



4. In Consistency Check Mode drop-down list, select High, Disabled, or Idle.
5. If you selected the Idle mode, enter the consistency check delay (in seconds) and parallel consistency check count:
 - **Consistency Check Delay**—Amount of time the controller must be inactive before the consistency check is started. Enter a value from 0-30. A value of 0 disables the scan.
 - **Parallel Consistency Check Count**—Number of logical drives on which the controller will perform the consistency check in parallel.
6. Click **OK**.

7.4 Optimizing Logical Drive Performance

This section describes how to enable controller cache optimizations and SSD I/O bypass acceleration to improve I/O throughput on the logical drives in your storage space. You can apply cache optimizations independently on a per controller or per logical drive basis. You can apply I/O bypass acceleration on arrays comprised of SSDs only.

7.4.1 Enabling Cache Optimizations

Use this option to enable the following cache optimizations on the controllers in your storage space. You can apply cache optimizations independently on a per controller or per logical drive basis.

Note: You cannot use controller caching and maxCache caching concurrently. Controller caching is available only if maxCache is not enabled on the controller. For more information about maxCache, see [Working with maxCache Devices](#) on page 70.

Option	Description
Cache Ratio	Sets the global Read:Write cache ratio.
Write Cache Bypass Threshold	Sets the write cache block size threshold, above which data is written directly to the drive.
No Battery Write Cache	Enables write caching on controllers without a backup module.
Wait for Cache Room	Waits for cache space (if none is available) before completing the request.

Option	Description
Global Physical Devices Write Cache Policy	Sets the write cache policy for the physical drives on the controller. Caution: Enabling drive write caching can improve performance. However, a power, device, system failure, or dirty shut down may result in data loss or file-system corruption.


To enable cache optimizations on a controller:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



When the Set Properties window opens, click the **Cache** tab.

3. Adjust cache settings, as needed.



4. Click **OK**.

7.4.1.1 Enabling Cache Optimization for a Logical Drive

You can enable/disable cache optimization for each logical drive in your storage space:


1. In the Enterprise View, select a controller, then select a logical drive.
2. On the ribbon, in the Logical Device group, click **Set Properties**.
3. In the Controller Caching drop down-list, select **Disabled** or **Enabled**.
4. Click **OK**.

7.4.2 Enabling SSD I/O Bypass

Use this option to enable I/O Bypass acceleration for logical drives comprised of SSDs only. This option enables I/O requests to bypass the controller firmware and access SSDs directly. This process accelerates reads for all RAID levels and writes for RAID 0.


To enable I/O Bypass acceleration:

1. In the Enterprise View, select a controller, then select an array on the controller.
2. On the ribbon, in the Array group, click **Set Properties**.



The Set Properties window opens; the **General** tab is selected, by default.

3. From the SSD I/O Bypass drop-down, select **Enabled** or **Disabled**.



4. Click **OK**.

7.5 Moving a Logical Drive

maxView Storage Manager allows you to move a single logical drive from one array to another array. You can choose the following destinations:

- Move Logical Drive To a New Array
- Move Logical Drive To an Existing Array

If you move the logical drive to a new array, the array is created automatically. If you move the logical drive to an existing array, it must have sufficient space and member disk drives to store the logical drive data and accommodate the RAID level; for example, three drives, minimum, for a RAID 5.


Note: Moving a logical drive can be a time-consuming process. All data in the logical drive is moved onto the new or existing array, and the controller continues to service I/O requests to other logical drives.

To move a logical drive:


1. In the Enterprise View, select a logical drive.
2. On the ribbon, in the Logical Device group, click **Move Logical Device**.



3. When the wizard opens, select **To New Array** or **To Existing Array**, then click **Next**.




4. If you are moving the logical drive to a new array, select the physical drives for the array. Be sure the drive type is the same for all drives (SAS or SATA, not mixed).



Note: The drives must have sufficient capacity to store the logical drive data.

- If you are moving the logical drive to an existing array, expand the Arrays and Logical Devices list, then select the destination array.



- Click **Next**, review the summary information, then click **Finish**. maxView Storage Manager moves the logical drive onto the new or existing array. If you moved the last logical drive on an array, maxView Storage Manager deletes the array and removes it from the Enterprise View.

7.6 Moving an Array

You can move an array by replacing its physical drives with drives of the same type or different type. For example, you can replace SAS drives in the array with other SAS drives, or replace SAS drives with SATA drives. You cannot combine drive types in the same array, however. If you choose to replace SAS drives

with SATA drives, for example, *all* drives in the array must be replaced with SATA drives. The replacement drives must be in the Ready state; that is, not part of any array or assigned as a spare.

Moving an array automatically removes any previously assigned spare drives. Replaced drives in the array are freed and become Ready drives that can be used in other arrays, logical drives, or as spares.


Note: Moving an array can be a time-consuming process. All data in each logical drive is copied to the replacement drives, and the controller continues to service I/O requests to other logical drives.

To move an array:


1. In the Enterprise View, select an array.
2. On the ribbon, in the Array group, click **Modify Array**.



3. When the wizard opens, select an action, then click **Next**:
 - Select **Move Drives** to replace array drives with drives of the same type.
 - Select **Change Drive Type** to replace array drives with drives of a different type.



4. Select one or more drives. For Move Drives, the wizard displays only physical devices of the same type. For Change Drive Type, the wizard displays only physical devices of a different type. The RAID level determines the number of drives you need to select.



Note: The drives must have sufficient capacity to hold all of the logical drives in the source array.

5. Click **Next**, review the summary information, then click **Finish**.

7.7 Modifying an Array

maxView Storage Manager allows you to perform different actions to reconfigure an array. You can choose the following destinations:


- Add Drives To an Array
- Move Drives To an Array

If you add the logical drives, you are expanding the array by adding the data drives. If you move the logical drive to an existing array, you need to replace the array drives with the drives of the same type.


Note: Moving a logical drive can be a time-consuming process. All data in the logical drive is moved onto the new or existing array, and the controller continues to service I/O requests to other logical drives.

To add or remove drives in an array:


1. In the Enterprise View, select an array.
2. On the ribbon, in the Array group, click **Modify Array**.



3. When the wizard opens, select **Add Drive(s)** or **Move Drive(s)**, then click **Next**.



4. If you are adding the new drives to an array, select the physical drives for the array. Be sure the drive type is the same for all drives (SAS or SATA, not mixed).



Note: The drives must have sufficient capacity to store the logical drive data.

5. Click **Next**, review the summary information, then click **Finish**.
maxView Storage Manager moves the logical drive onto the new or existing array. If you moved the last logical drive on an array, maxView Storage Manager deletes the array and removes it from the Enterprise View.

7.8 Working with Mirrored Arrays

maxView Storage Manager allows you to split a mirrored array and then recombine it. This process entails splitting a RAID 1, RAID 1(ADM), RAID 10, or RAID 10(ADM) array into two identical new arrays consisting of RAID 0 logical drives. Arrays with other RAID configurations cannot be split.

7.8.1 Creating a Split Mirror Backup

Use this option to split a mirrored array, consisting of one or more RAID 1, RAID 1(ADM), RAID 10, or RAID 10(ADM) logical drives, into two arrays: a primary array and a backup array, with these characteristics:

- The primary array and backup array will contain identical RAID 0 logical drives.
- The primary array continues to be fully accessible to the operating system.
- The backup array is hidden from the operating system and data on the drive is frozen.


Note: You can use the backup array to restore the primary array with its original contents. See [Re-mirroring, Rolling Back, or Reactivating a Split Mirror Backup](#) on page 63.

- The primary array includes the designation "Split Mirror Set Primary" as the device type.
- The backup array includes the designation "Split Mirror Set Backup" as the device type.

If the array is protected by a spare drive, the drive is unassigned after the split.

To create a split mirror backup:

1. In the Enterprise View, select a mirrored array.
2. On the ribbon, in the Array group, click **Split Mirror Backup**.



3. When prompted to create the backup array, click **OK**.



7.8.2 Re-mirroring, Rolling Back, or Reactivating a Split Mirror Backup

When you re-mirror a split mirrored array, you recombine the primary array and backup array into a single array. You can:


- Re-mirror the array and preserve the existing data; the backup array is discarded. This option re-creates the original mirrored array with the current contents of the primary array.
- Re-mirror the array and roll back to the contents of the backup array; existing data is discarded. This option re-creates the mirrored array but restores its original contents from the backup array.

You can also reactivate the split mirror backup. This option makes the backup array fully accessible to the operating system. maxView Storage Manager removes the "Split Mirror Set Backup" designation and re-designates it as a Data Array.

To re-mirror, roll back, or reactivate a split mirror backup:


1. In the Enterprise View, select the Split Mirror Set Primary array; that is, an array with an existing split mirror backup.

Note: Use the Summary tab on the Storage Dashboard to verify the array type.
2. On the ribbon, in the Array group, click **Remirror/Activate Backup**.



3. When prompted to select a re-mirroring task, choose: `Re-mirror array`, `Re-mirror with roll-back`, or `Activate Backup`.

Note: Microsemi recommends that you do not perform a re-mirror with roll back if the logical drive to be rolled back is mounted or in use by the operating system.



4. Click **OK**.

7.9 Changing the RAID Level of a Logical Drive


If your storage needs or application requirements change, you can change, or *migrate*, the RAID level of your logical drives to another, more suitable, RAID level. You might want to change the RAID level to add redundancy, further protect your data, or to improve data availability for speedier access. See [Selecting the Best RAID Level](#) for more information.

The following RAID level migrations are supported:

- RAID 0 to RAID 5
- RAID 0 to RAID 10
- RAID 5 to RAID 6
- RAID 6 to RAID 5
- RAID 10 to RAID 5
- RAID 5 to RAID 10
- RAID 1 to RAID 5


To change the RAID level of a logical drive:

1. In the Enterprise View, select a controller, then select the logical drive that you want to migrate.
2. On the ribbon, in the Logical Device group, click **Expand/Migrate**.




The Expand/Migrate Logical Device wizard opens.

3. Click **Migrate**, then click **Next**.



4. Select a new RAID level, then click **Next**. Only valid RAID level options are offered.
5. Select the sub array count for RAID 50 and RAID 60.



6. Select the logical drive stripe size from the drop-down list.

Note: The default stripe size usually provides the best performance.

7. Click **Next**.
8. Review the summary of logical drive settings. To make changes, click **Back**.
9. Click **Finish**.
The logical drive is reconfigured and migrates to the new RAID level.

7.10 Increasing the Capacity of a Logical Drive


You can add more disk drive space, or *expand*, a logical drive, to increase its capacity.

The expanded logical drive must have a capacity that is greater than or equal to the original logical drive.

Note: You can expand a logical drive only into the free space of the host array. To add physical drives in an array, see [Modifying an Array](#) on page 61


To increase the capacity of a logical drive:

1. In the Enterprise View, select a controller, then select the logical drive you want to expand.
2. On the ribbon, in the Logical Device group, click **Expand/Migrate**.



The Expand/Migrate Logical Device wizard opens.

3. Click **Expand**, then click **Next**.



4. Enter the new logical drive size in the space provided. It must be greater than or equal to the current size.

5. Click **Next**.
6. Review the summary of logical drive settings. To make changes, click **Back**.
7. Click **Finish**.
The logical drive is expanded and its capacity is increased to the new size.

7.11 Changing the Logical Drive Rebuild Priority

The Rebuild Priority setting determines the urgency with which the controller treats an internal command to rebuild a failed logical drive:

- At the low setting, normal system operations take priority over a rebuild.
- At the medium setting, normal system operations and rebuilds get equal priority.
- At the medium high setting, rebuilds get higher priority than normal system operations.
- At the high setting, rebuilds take precedence over all other system operations.

If the logical drive is part of an array with an online spare, rebuilding begins automatically when drive failure occurs. If the array does not have an online spare, rebuilding begins when the failed physical drive is replaced. For more information, see [Rebuilding Logical Drives](#) on page 144.


To change the rebuild priority:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



The Set Properties window opens.

3. In Rebuild Priority Mode drop-down list, select **Low**, **Medium**, **Medium High**, or **High**. The Rebuild Priority field is circled in the next figure.



4. Click **OK**.

7.12 Renaming a Logical Drive

To change the name of a logical drive:

1. In the Enterprise View, select a controller, then select the logical drive you want to rename.
2. On the ribbon, in the Logical Device group, click **Set Properties**.



The Set Properties window opens.

3. In the Logical Device Name field, type the new name, then click **OK**. Names can include any combination of letters, numbers, and spaces.

maxView Storage Manager updates the logical drive name and displays the new name in the Enterprise View.

7.13 Deleting an Array or Logical Drive

When you delete an array or logical drive, it is removed from the Enterprise View and the disk drives or segments in the logical drive(s) become available to use in a new array or logical drive (see [Creating a Logical Drive Using Available Segments](#)).

Caution: When you delete an array you lose all data on the logical drive(s) within the array, in addition to the array itself. When you delete a logical drive, you lose all data stored on that logical drive. Be sure you no longer need the data on the array or logical drive before you delete it.

To delete an array or logical drive:

1. In the Enterprise View, select the array or logical drive you want to delete.
2. On the ribbon, in the Array group or Logical Device group (shown below), click **Delete**.



3. When prompted to continue, click **Delete** to delete the array or logical drive.

Note: If a deleted logical drive is the only logical in the array, the array itself is also deleted.

7.14 Maintaining an Energy-Efficient Storage Space

The power management options in maxView Storage Manager control the power profile of the physical drives on a controller. They offer a balance between maximum performance and minimum power usage. To ensure continued operations when temperature thresholds are exceeded, you can enable Survival mode to throttle dynamic power settings to their minimum values.


To set the power management options for a controller:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



The Set Properties window opens.

3. Click the **Power Management** tab.



4. In the Power Mode drop-down list, select:
 - Minimum Power—Set power settings to lowest possible values and reduce power dynamically, based on workload.
 - Maximum Performance—Set power settings to highest possible values and do not reduce power dynamically.
5. In the Survival Mode drop-down list, select:
 - Enabled—Allows the controller to throttle back dynamic power settings to their minimum values when temperatures exceed the warning threshold.

Note: Enabling Survival mode allows the server to continue running in more situations, but may affect performance.
 - Disabled—Disables Survival mode.
6. Click **OK**.

8 Working with maxCache Devices

Microsemi Adaptec Smart Storage Controllers support an advanced SSD caching technology called maxCache™. maxCache uses a reserved logical drive, called the *maxCache Device*, to support read and redundant write caching for storage connected directly to your controller. The maxCache Device is comprised of SSDs only.

With maxCache read caching enabled, the system copies frequently read "hot" data to the maxCache Device for faster retrieval. With maxCache write caching enabled, the maxCache Device is populated with certain "hot" blocks from the logical drives on the controller. All writes to these hot blocks go directly to the maxCache Device. The data remain on the maxCache Device until it is full or some other "hotter" data replaces it.

8.1 maxCache Limitations

- maxCache is not supported on all Microsemi Adaptec Smart Storage Controllers. For more information, see *PMC-2153191 maxView Storage Manager and ARCCONF Command Line Utility Readme*.
- If the maxCache controller has a green backup module, the super capacitor must be fully charged.
- Following are the limitations on maxCache device:
 - It must be created with SSDs
 - It must have logical block size of 512 bytes
 - Minimum maxCache device capacity is 16 GB
 - Maximum aggregate maxCache device sizes can be ~1.7TB for 64KB cache line size, ~6.8TB for 256KB cache line size.
- Following are the limitations on the data logical device for which the maxCache device to be assigned:
 - It must have the capacity at least as large as the maxCache device
 - It must have logical block size of 512 bytes
 - Maximum data logical device size can be 256TB for the maxCache created with 64KB cache line size, 1024TB for the maxCache created with 256KB cache line size
- The following operations are not available when maxCache is enabled:
 - Expand Array/Logical Device
 - Move Logical Device
 - Replace Array Drives
 - Split Mirror
 - Heal Array
 - Migrate Array


8.2 Creating a maxCache Device

To create a maxCache Device:


1. In the Enterprise View, select a system, then select a controller on that system. You can also create a maxCache device by selecting a logical device node.
2. On the ribbon, in the maxCache group, click **Create maxCache**.



3. When the wizard opens, select a RAID level for the maxCache Device, then click **Next**.
See [Selecting the Best RAID Level](#) for more information about RAID levels.
4. Select the SSDs that you want to include in the maxCache Device, then click **Next**. Be sure to select the right number of SSDs for the RAID level you selected.



5. Select the data logical drive (16 GB minimum), then click **Next**.



6. (Optional) In the Cache Attributes panel, customize the maxCache Device settings. You can:
 - Set a smaller logical drive size. (By default, the maxCache Device uses up to the maxSize of the data logical drive.)
 - Set the write cache mode to Write-Back (default) or Write-Through.
 - Select the Cache Line Size. The Cache Line Size impacts the cache performance and the maximum size selected. To create a large size maxCache you need to select the larger cache line size supported by that controller.
7. Click **Next**, then review the logical drive settings.
8. Click **Finish**, then click **OK**.
 maxView Storage Manager updates the configuration, then adds an array and logical device to the maxCache Device tree in the Enterprise View.

8.3 Changing the Write Cache Mode

The write cache mode determines when data is stored on the maxCache Device and when the controller communicates with the operating system. You can set the Write Cache mode to:

- Write-Through—The controller sends (or *writes*) the data to the maxCache Device, then sends confirmation to the operating system that the data was received. Use this setting when performance is less important than data protection.
- Write-Back—The controller sends confirmation to the operating system that the data was received, then writes the data to the maxCache Device. Use this setting when performance is more important than data protection.

To change the maxCache write cache mode:

1. In the Enterprise View, select a controller, then select a maxCache Device on that controller.
2. On the ribbon, in the maxCache group, click **Set Properties**.



The Set Properties window opens.

3. In the Write Cache drop-down list, select `Write-Back` or `Write-Through`.
4. Click **OK**.

8.4 Deleting the maxCache Device

When you delete the maxCache Device, the component SSDs become available and can be used to create a new logical drive, hot spare, or new maxCache Device.

Caution: Be sure that the maxCache controller is quiescent before deleting the maxCache Device; otherwise you may lose data.

Note:

Also, the delete maxCache option is available only when the write-cache policy is set to "write-through" and is the last maxCache logical device in the maxCache array.

To delete the maxCache Device:

1. In the Enterprise View, select a controller, then select the maxCache Device.

2. On the ribbon, in the maxCache group, click **Delete**.



3. When prompted, click **Delete**, then click **OK**. Click **Cancel** to cancel the action.

8.5 Analyzing maxCache Performance

maxView Storage Manager provides advanced usage statistics about the maxCache Devices on your Microsemi Smart Storage controllers. You can use these statistics to gain a better understanding of how maxCache is performing in your storage space. Use the Statistics Viewer to view the maxCache statistics; see [Viewing Advanced Statistics](#) on page 124 for more information.

9 Working with maxCrypto Devices

Microsemi Smart Storage Controllers support an advanced controller-based encryption (CBE) technology called maxCrypto™. It provides an enterprise-class encryption solution that protects sensitive data on storage connected directly to your controller.

maxCrypto supports two roles for managing encryption services:

- A Crypto Officer (Admin) role that can perform all encryption operations
- A User role with reduced privileges

maxCrypto allows you to selectively encrypt arrays and logical drives, regardless of RAID level; create storage spaces with mixed encrypted and plaintext volumes; and convert plaintext volumes to encrypted volumes.

Note: Encryption of non-RAID volumes (such as physical, raw, or pass-through devices) is not supported. Consequently, HBAs and controllers operating in HBA Mode do not support maxCrypto.

9.1 maxCrypto Initial Setup


Before you can begin using maxCrypto in your storage space, you must complete the initial setup task to create the Crypto Officer account and configure the initial maxCrypto settings, including the Crypto Officer login credentials, maxCrypto master encryption key, and other basic information. You must also accept the maxCrypto Certificate of Use.

To setup maxCrypto:

1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Controller group, click **maxCrypto Settings**.




The Set maxCrypto Configuration window opens; the Initial Setup tab is selected, by default.



3. In the maxCrypto Mode drop-down, select `Enabled` to activate maxCrypto. Select `Disabled` to deactivate maxCrypto.
4. In the Allow New Plaintext Logical Device(s) drop-down, select `Enabled` to allow plaintext logical devices to be created, in addition to encrypted logical devices. Select `Disabled` to allow only encrypted logical devices to be created.

Note: To create plaintext logical devices, both maxCrypto Mode and Allow New Plaintext Logical Device(s) must be enabled.
5. In the Master Key field, enter the maxCrypto master encryption key. The Master Key is a 10-32 character string, comprising all printable ASCII characters.

Caution: Be sure to record the master key and store in a safe place. Once set, the Master Key cannot be displayed or recovered, only reset. Failure to provide the Master Key may result in encrypted data being irretrievable.
6. In the Enter Crypto Password field, enter the Crypto Officer password. The password is a 8-16 character string, comprising all printable ASCII characters. It must include at least one uppercase character, one lowercase character, one numeric, and one special character (`#,!,@,...`).
7. In the Re-Enter Crypto Password field, re-enter the Crypto Officer password.
8. Click **OK**.
9. When the maxCrypto Certificate of Use window opens, click **Agree** to complete the maxCrypto activation.



9.1.1 Managing maxCrypto Accounts

After initial setup is complete, the Crypto Officer account is logged in, by default. As the Crypto Officer, you can perform all account management operations, including configuring the standard User account, changing passwords, and enabling password recovery options. If you are logged into the standard User account, you have the ability to perform a limited number of maxCrypto operations, as described in the table below.


Standard users can:	Standard users can't:
Log into maxCrypto	Perform initial setup
Lock/Unlock firmware update	Recover password
Enter/Re-Enter new password	Allow new plaintext logical device(s)
Create encrypted and plaintext logical drives	Set/Change master key
Convert plaintext array/logical drive to encrypted data	Import master key
Crypto erase array/logical drive	Set/Change password recovery question
Clear maxCrypto configuration	

To manage maxCrypto accounts:

1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Controller group, click **maxCrypto Settings**.



When the Set maxCrypto Configuration window opens, click the **Account** tab.



3. In the User Role drop-down, select the `Crypto Officer` or `User` account.
4. In the Enter New Password field, enter the password for the account.

The password is a 8-16 character string, comprising all printable ASCII characters. It must include at least one uppercase character, one lowercase character, one numeric, and one special character (`#,!,@,...`).

Note: The first time you enable the User account, this entry defines the initial login credentials for the account. For an existing account, this entry changes the login credentials.

5. In the Re-Enter New Password field, re-enter the password for the user account.
6. For the `Crypto Officer` account, set/change the password recovery question and answer:
 - a. In the Set/Change Password Recovery Question field, enter a recovery question for a forgotten password.
 - b. In the Set/Change Password Recovery Answer field, enter the answer to the recovery question.

Note: Password recovery is available only for the `Crypto Officer`. The recovery question/answer fields are deactivated for the `User` account.


7. Click **OK**.

9.1.2 Logging In and Logging Out

You must be logged into maxCrypto to use the encryption features in maxView Storage Manager, from encrypting a logical drive, to enabling plaintext volumes, to simply changing a maxCrypto password.

After initial setup is complete, the `Crypto Officer` account is logged in, by default.

To logout of the account (ending the maxCrypto session), simply click the **Logout** button on the **Account** tab or **General** tab in the Set maxCrypto Configuration window.



Set maxCrypto Configuration

Controller 1 Setup controller maxCrypto account parameters

Initial Setup Login **General** Account

User Role Crypto User(Admin) v

Enter New Password ***** i

Re-Enter New Password ***** i

Set/Change Password Recovery Question

Set/Change Password Recovery Answer


Logout OK Cancel

To login back in, or to login to a different account:

1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Controller group, click **maxCrypto Settings**.



When the Set maxCrypto Configuration window opens, click the **Login** tab.



Set maxCrypto Configuration

Controller 1 New Crypto password will be applied only when the recovery answer is correct

Initial Setup **Login** General Account

User Role Crypto User(Admin) v

Password *****

Forgot Password

Recovery Question Which is your first company

Recovery Answer

Enter New Crypto Password ***** i

Re-Enter New Crypto Password ***** i

OK Cancel

3. In the User Role drop-down, select the **Crypto Officer** or **User** account.

4. In the Password field, enter the password for the user role.

Note: If you're logging in as the Crypto Officer and forgot the password, you can reset the password by following the steps below.
5. Click **OK**.

To reset the Crypto Officer password:


1. In the Login tab, select the `Crypto Officer` user role, as described above.
2. Click **Forgot Password**.
3. In the Recovery Answer field, enter the answer to the recovery question.
4. Enter a new password for the Crypto Officer account.

The password is a 8-16 character string, comprising all printable ASCII characters. It must include at least one uppercase character, one lowercase character, one numeric, and one special character (#,!,@,...).

5. Re-enter the new Crypto Officer password.
6. Click **OK**.

9.1.3 Checking maxCrypto Status

To check the maxCrypto status, select a controller in the Enterprise view, then click the **maxCrypto** tab on the storage dashboard. Before completing the initial setup steps, most maxCrypto properties will be listed as "Not Configured". After completing the initial setup steps, the main maxCrypto properties will be Configured, as described in the table below.



The screenshot shows the maxCrypto configuration interface with three main panels:

- maxCrypto Info:**
 - maxCrypto/SSD Smart Path... Supported
 - maxCrypto/maxCache Mixing Not Supported
- Settings and Status:**
 - maxCrypto Status Enabled
 - Allow New Plaintext Logica... Enabled
 - Master Key Configured
 - Firmware Locked for Update Unlocked
 - Encrypted Logical Device C... 0
 - Encrypted Physical Device ... 0
- Account Info:**
 - Login Status Logged-in as Crypto
 - Crypto Officer Password Configured
 - User Password Not Configured
 - Crypto Password Unlock A... 10
 - User Password Unlock Att... 10
 - Crypto Officer Password R... Configured

Property	Description
maxCrypto Status	When Enabled, both Encrypted and Plaintext logical devices can be created, based on the Allow New Plaintext Volumes property. When Disabled, only Plaintext logical devices can be created.
Allow New Plaintext Volumes	When Enabled, both Plaintext and Encrypted logical devices can be created. When Disabled, only Encrypted logical devices can be created, when maxCrypto Status is Enabled.

Property	Description
Master Key	After Initial setup, Master Key value is displayed as Configured.
Firmware Locked for Update	If Unlocked, firmware upgrade is enabled. If Locked, firmware upgrade is disabled.
Number of maxCrypto Logical Devices	Total count of the encrypted logical devices.
Number of maxCrypto Physical Devices	Total count of encrypted physical devices.
Crypto Officer Password	After Initial setup, Crypto Officer Password will be displayed as Configured.
Login Status	If logged-in as the Crypto Officer (Admin), displays "Logged-in as Crypto". If logged in as User, displays "Logged-in as User". If maxCrypto session timed out, displays "Timeout". If maxCrypto is logged-out, displays as "Not Logged In".
User Password	If the User account is configured, displays as "Configured"; otherwise, displays as "Not Configured".
Crypto Password Unlock Attempts Remaining	Countdown of Crypto Officer attempts remaining after a failed login. Note: After 10 failed login attempts, the Crypto Officer account is locked for 15 minutes.
User Password Unlock Attempts Remaining	Countdown of User role attempts remaining after a failed login. Note: After 10 failed login attempts, the User account is locked for 15 minutes.
Crypto Officer Password Recovery Parameters	Displays as "Configured" when Crypto Officer sets the password recovery question and answer. Displays as "Not Configured" when Crypto Officer has not set the password recovery question and answer.

9.2 Modifying the maxCrypto Configuration

To modify the maxCrypto configuration, use the General tab on the Set maxCrypto Configuration window.

If you are logged in as the Crypto Officer, you can:

- Enable/Disable maxCrypto
- Enable/Disable new plaintext volumes
- Lock/Unlock firmware upgrade
- Change the maxCrypto master encryption key


If you are logged in as the maxCrypto User, you can only Lock/Unlock firmware upgrade.

To modify the maxCrypto configuration:

1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Controller group, click **maxCrypto Settings**.



When the Set maxCrypto Configuration window opens, click the **General** tab.



3. Adjust the maxCrypto settings, as needed:

- In the maxCrypto field, select **Enabled** or **Disabled** to enable/disable the maxCrypto system.
- In the Allow New Plaintext Volumes field, select **Enabled** to allow plaintext volumes and encrypted volumes in your storage space. Select **Disabled** to allow encrypted volumes only in your storage space.
- In the Firmware Locked for Update field, select **Unlocked** to allow firmware upgrades. Select **Locked** to block firmware upgrades.
- To change the master encryption key, click **Change Master Key**, then enter the new key in the Set/Change Master Key field.

The Master Key is a 10-32 character string, comprising all printable ASCII characters.


Caution: Be sure to record the master key and store in a safe place. Once set, the Master Key cannot be displayed or recovered, only reset. Failure to provide the Master Key may result in encrypted data being irretrievable.

4. Click **OK**.

9.3 Creating an Encrypted Logical Drive

Use the Create Logical Device wizard to create encrypted and plaintext logical drives on an existing array or a new array. See [Creating Arrays and Logical Drives](#) on page 26.

When the wizard reaches the RAID Attributes panel, click **Create Encrypted Logical(s)** to encrypt the logical device. Un-check the check box to create a plaintext logical device.

**Note:**

1. If maxCrypto status is Disabled, only plaintext logical drives can be created.
2. If maxCrypto status is Enabled and Allow New Plaintext Volumes property is Enabled, both encrypted and plaintext logical drives can be created; the default is encrypted.

Note: You must be logged in to maxCrypto to create plaintext volumes, even if maxCrypto status and Allow New Plaintext Volumes are both Enabled; see [Logging In and Logging Out](#) on page 77.

3. If maxCrypto status is Enabled and Allow New Plaintext Volumes property is Disabled, only encrypted logical drives can be created.

9.4 Converting Plaintext Data to Encrypted Data


Use this operation to convert plaintext data to encrypted data. You can convert an unencrypted logical drive or all unencrypted logical drives on an array. Additionally, you can choose to preserve or discard the existing data during conversion.

To convert plaintext data:

1. In the Enterprise View, select a controller, then select a plaintext (unencrypted) logical drive or an array.
2. On the ribbon, in the Logical Device group or Array group, click **Set Properties**.



When Set Properties window opens, click the **maxCrypto tab**.



Note: The User Role and Password fields are enabled only when you are not logged in to maxCrypto.

3. Click the **Convert Plaintext Data to Encrypted Data** check box.
4. To preserve the data during conversion, click the **Preserve Existing Data** check box. Leave it un-checked to discard the data.

Note: When you convert a plaintext volume to an encrypted volume *without* preserving existing data, the old plaintext data remains on disk and is still available if the logical device is deleted.

5. Click **OK**.

9.5 Re-Keying a Logical Drive

maxCrypto allows you to re-key a logical drive for added security. The logical drive key is used with the master key to encrypt the device.


Optionally, you can re-key an array to generate a new key for all encrypted logical drives on the array.

To re-key a logical drive:

1. In the Enterprise View, select a controller, then select an encrypted logical drive or array.
2. On the ribbon, in the Logical Device group or Array group, click **Set Properties**.



When Set Properties window opens, click the **maxCrypto** tab.



3. Click the **Logical Device(s) Key ReKey** check box.
4. Click **OK**.

9.6 Clearing the maxCrypto Configuration

You can clear the maxCrypto configuration to restore the default maxCrypto settings. Clearing the maxCrypto configuration resets all keys, passwords, and users, including the maxCrypto Officer account and User account.


Note: To clear the maxCrypto configuration, ensure that there is no encrypted logical drive available in your storage space.

To clear the maxCrypto configuration:

1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Controller group, click **Manage Configuration**.



The Manage Configuration wizard opens.



3. Select **Clear maxCrypto**, then click **Next**.
4. Review the Summary information, then click **Finish**.

9.7 Erasing an Encrypted Logical Drive

You can cryptographically erase the data on any encrypted logical drive in the Optimal state. Crypto erase performs an instant/quick erase. The logical drive remains in the Enterprise View and ready to store new data.

Optionally, you can cryptographically erase an array, to erase the data on all encrypted logical drives on the array.


Caution: When you cryptographically erase a logical drive, you lose all data on that drive.

To cryptographically erase a logical drive:

1. In the Enterprise View, select a controller, then select an encrypted logical drive or array.
2. On the ribbon, in the Logical Device group or Array group, click **Crypto Erase**.



The Crypto Erase Logical Device window opens.



3. Click **Erase**.

9.8 Importing a Foreign Master Key

When an encrypted logical drive is moved from another controller, the master key used to encrypt the logical drive is needed to decrypt it. Use the Import Foreign Master Key option to import the master key so that the logical drive data can be accessed and managed on the new controller.


To import a foreign master key:

1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Controller group, click **maxCrypto Settings**.



When the Set maxCrypto Configuration window opens, click the **General** tab.

3. In the Import Foreign Local Master Key field, enter the master key originally used to encrypt the logical drive.
The Master Key is a 10-32 character string, comprising all printable ASCII characters.



4. Click **OK**.


10 Maintaining Physical Devices

This section describes how to manage the controllers, disk drives, solid state drives, and enclosures in your storage space.

10.1 Viewing Device Properties

Click on any physical device in the Enterprise View then, on the Storage Dashboard, click the **Properties** tab to view version numbers, status, model numbers, features, and other information about the device.

The properties listed vary, depending on which type of device you select. The figure below shows the properties for a controller. For more information about using the Storage Dashboard to monitor the components in your storage space, see [Viewing Component Status in the Storage Dashboard](#).



10.2 Locating Drives in Your Storage Space

You can blink the LEDs on disk drives and SSDs to identify where they are physically located in your storage space. This table describes how to locate specific devices.

Note:

Once any of the device is located, the timeout value will be overwritten with the latest timeout value for all the located devices.

To Locate...	Select...
A disk drive	Disk Drive icon: 
All disk drives on a controller	Controller icon: 
All disk drives included in an array	Array icon:

To Locate...	Select...
	
All disk drives included in a logical drive	Logical Drive icon: 
All SSDs in the maxCache Device	maxCache Device icon: 

10.2.1 Locating Disk Drives

To locate an individual disk drive or all physical drives on the controller:

1. In the Enterprise View, select a controller or an individual drive on the controller.
2. On the ribbon, in the Physical Device group, click **Locate**.



The Locate Physical Device window opens.

3. From the drop-down list, select the timeout period (1 hour, 4 hours, 24 hours).
4. Click the **Locate** button.
The LED on the disk drive(s) begin to blink.
5. To stop blinking the drive(s) immediately, click **Stop**.



6. Click **Cancel** to close the Locate Physical Device window.
The LED(s) continue to blink for the duration of the timeout period.

10.2.2 Locating Physical Disks in an Array or Logical Drive


To locate all physical disks in an array or logical drive:

1. In the Enterprise View, open the Arrays and Logical Devices tree for a controller, then select an array or logical drive.
2. On the ribbon, in the Array group or Logical Device group (shown below), click **Locate**.



The Locate Logical Device window opens and displays a list of the physical disks associated with the array or logical drive.

3. Select the timeout period (1 hour, 4 hours, 24 hours), then click **Locate**.



The LEDs on the disk drives begin to blink.

4. Click **Cancel** to close the Locate window.
The LEDs continue to blink for the duration of the timeout period.
5. Click **Stop** to stop blinking the drives immediately.

10.2.3 Locating SSDs in the maxCache Device


To locate the Solid State Drives (SSDs) in the maxCache Device:

1. In the Enterprise View, select a controller; then select the maxCache Device on that controller. You can select the maxCache array or logical device.
2. On the ribbon, in the maxCache group, click **Locate**.



The Locate maxCache window opens, displaying a list SSDs comprising the maxCache Device.

3. Select the time-out period from the drop-down list: 1 hour, 4 hours, 24 hours.
4. Click the **Locate** button.



The LEDs on the SSDs begin to flash.

5. Click **Stop** to stop blinking the SSDs.
6. Click **Cancel** to close the Locate maxCache window.


10.3 Working with Physical Device Error Counters

This section explains how to view the physical device error counters and how to clear the error counters from a physical device and a controller.

The clear device error counters feature provides an option to clear the device error counters on the physical devices. This option is available at the controller level to clear the device error counters on all the physical devices connected to it and at the physical device level to clear the error counters on the specific device.

10.3.1 Viewing Physical Device Error Counters

To view the physical device error counters for a hard drive or SSD, select the drive in the Enterprise View then, on the Storage Dashboard, click the **Error Counters** tab. The table below describes the error counters.



Error	Counter
Aborted Command	0
Bad Target Error	0
ECC Recovered Read Error	0
Failed Read Recover	0
Failed Write Recover	0
Format Error	0
Hardware Error	0
Hard Read Error	0
Hard Write Error	0
Hot-Plug Count	0
Media Failure	0
Not Ready Error	0
Time-Out Error	0
Predictive Failure	0
Retry Recovered Read Error	0

Error Counter	Description
Aborted Command	Number of times a drive was failed due to aborted commands that could not be retried successfully.
Bad Target Error	Number of times that this drive did something that did not conform to the SCSI Bus Protocol. It will cause a reset of the SCSI bus that this drive is attached to.
ECC Recovered Read Errors	Number of ECC-corrected read errors.
Failed Read Recover	Number of times a recover of another physical drive in the logical volume failed due to a hard read error from this drive.
Failed Write Recover	Number of times a recover of this physical drive failed due to an error occurring on this drive during a write operation.
Format Error	Number of times a Format command (used when remapping defects) failed. A failed remap operation may cause the controller to fail a drive.
Hardware Error	Number of times a drive returned a bad hardware status. The drive may be failed if retries do not work.
Hard Read Error	Number of unrecoverable read errors.
Hard Write Error	Number of unrecoverable write errors.
Hot-Plug Count	Number of times this drive was hot-plugged (removed) from a box.
Media Failure	Number of times a drive was failed due to unrecoverable media errors.
Not Ready Error	Number of times the drive was failed because it never became ready after the “spin up” command was issued. If retries or drive spin-ups fail, the drive will be failed.
Other Timeouts	Timeouts other than Data ReQuest Timeouts (DRQ).
Predictive Failure	Number of times that the drive returned a predictive failure error.

Error Counter	Description
Retry Recovered Read Error	Number of retry-recovered read errors.
Retry Recovered Write Error	Number of retry-recovered write errors.
SCSI Bus Fault	Number of “bus faults”, which we define as SCSI bus parity errors, overrun/underrun conditions, etc.
Service Hours	Number of service hours since the last power cycle.
Sectors Written	Number of sectors written to media.
Sectors Read	Number of sectors read from the media. This value will include sectors read into the on-drive cache buffer only if the drive keeps track of this value. Otherwise, only sectors requested through the drive interface are counted.

10.3.2 Clearing Error Counters from a Physical Device


Use this option to clear the device error counters of a specific physical device.

To clear the error counters from a physical device:

1. In the Enterprise View, select a physical drive node.
2. On the ribbon, in the Physical Device group, click **Set Properties**.



The Set Properties window opens.



3. To clear the device errors, select the Clear Device Error Counter(s) check box.

4. Click **OK**.

10.3.3 Clearing Error Counters from a Controller


Use this option to clear the the device error counters of all the physical devices from controller.

To clear the the device errors from a controller:

1. In the Enterprise View, select a controller, then select a physical drive on the controller.
2. On the ribbon, in the Physical Device group, click **Set Properties**.



The Set Properties window opens.



3. To clear the device errors, select the Clear Device Error Counter(s) check box.
4. Click **OK**.

10.4 Working with Failed or Failing Disk Drives

This section describes how to use maxView Storage Manager to manage failed or failing disk drives in your storage space.

10.4.1 Replacing Disk Drives in a Logical Drive

You can replace one or more disk drives in a logical drive. You may want to replace a drive to upgrade to larger disk drives, or to make disk drive size uniform across the logical drive.

Caution: If another disk drive in the logical drive fails during rebuild (see [Rebuilding Logical Drives](#) on page 144), you may lose data. For help solving disk drive problems, see [Recovering from a Disk Drive Failure](#).

To replace a disk drive in a logical drive:

1. In the Physical Devices tree in the Enterprise View, find the disk drive you want to replace; note its size and location (for instance Slot 1 in Enclosure 0).
2. Set the drive state to failed. (See [Setting a Disk Drive to 'Failed'](#) on page 95.)
3. Remove and replace the disk drive with one of equal or greater size.
4. Wait for the logical drive to rebuild. (See [Rebuilding Logical Drives](#) on page 144.)
5. Repeat these steps for each disk drive you want to replace.

10.4.2 Setting a Disk Drive to 'Failed'

Before you can remove a disk drive, you should set it to the Failed state to protect your data. To fail a disk drive (or SSD), use the Force Offline option for physical devices.

You can set a disk drive to the Failed state if:

- The disk drive is not part of a logical drive, *or*
- The disk drive is part of a redundant, healthy logical drive

Once you force a drive offline, it can be brought online again only after power-cycling the controller.

Caution: You may lose data or damage your disk drive if you remove a disk drive without first setting it to a failed state.

To set a disk drive to Failed:

1. In the Enterprise View, select a controller then, in the Physical Devices tree, select the drive you want to set to Failed.
2. On the ribbon, in the Physical Devices group, click **Force Offline**.



The Force Offline window opens.

3. Click **Force**.

The drive is taken offline and set to the Failed state.

Note: If the drive is part of a healthy logical drive, the drive is degraded and a warning message is displayed in the Event Log.

10.5 Erasing a Disk Drive

You can use maxView Storage Manager to *securely* erase existing data on any disk drive or SSD in the Ready state. Secure erase completely destroys the data on the drive; the data is completely and irretrievably eradicated.

Optionally, you can choose the erase pattern.

To securely erase a disk drive or SSD:

1. In the Enterprise View, select a controller then, in the Physical Devices tree, select the drive you want to erase.
2. On the ribbon, in the Physical Device group, click **Secure Erase**.



The Secure Erase Physical Device window opens.

3. From the drop-down list, select the erase pattern:
 - **Zero** (default)—Initializes all blocks to zero.
 - **Random Zero**—Initializes block to random value then zero.
 - **Random Random Zero**—Initializes block to random value, next block to random value, then zero.
 - **Sanitize Overwrite**—(HDD only) Fills every physical sector of the drive with a pattern.
 - **Sanitize Block Erase**—(SSD only) Sets the blocks on the drive to a vendor-specific value, removing all user data. It provides a very fast, complete, and robust erasure of the solid state device.
 - **Sanitize Crypto Scramble**— (Self-Encrypting Drives (SED) only) Changes the internal encryption keys that are used for user data, making the data irretrievable. This option is implemented for both HDD and SSD SED devices.
4. Click **Erase** to erase the drive.

10.5.1 Restricted/Unrestricted Secure Erase

For the Sanitize erase patterns (Overwrite, Block Erase, Crypto Scramble), the following erase methods are applicable if your drive supports the method:

- **Restricted:** the drive will be unusable until the sanitize operation is completed successfully. If a restricted sanitize operation fails, you are only allowed to start another sanitize operation.
- **Unrestricted:** the drive will be recoverable in the case that the sanitize erase operation fails. User data may still be present on the drive. Not all drives support this sanitize method.

For more information about Sanitize erase patterns, see [Erasing a Disk Drive](#) on page 95.

10.6 Initializing and Uninitializing Disk Drives

This section describes how to initialize or erase data and meta-data (including logical drive and partition information) from the disk drives (and SSDs) in your storage space. You can initialize or uninitialize individual disks, or use the wizard to initialize/uninitialize all disks on a controller.

10.6.1 Uninitializing a Disk Drive


You can uninitialize any disk drive (or SSD) containing Smart Controller configuration metadata. Uninitializing a disk drive clears the meta-data and reserved space from the drive and removes any OS partitions; existing data on the drive is destroyed.

Note: Uninitialized drives change from their current state to the Raw state. Raw drives are compatible with any Host Bus Adapter (HBA), including RAID controllers operating in HBA mode, and can be exchanged with drives on the motherboard's SATA interface.

To uninitialize a disk drive:


1. In the Enterprise View, select a controller; then, in Physical Devices tree, select the disk drive you want to uninitialize.

- On the ribbon, in the Physical Device group, click **Initialize**.



The Initialize/Uninitialize Physical Device window opens.

- Click the **Uninitialize** button, then click **OK**.



- When prompted, click **OK** to close the Initialize/Uninitialize Device window.

10.6.2 Initializing/Uninitializing all Drives on a Controller

To initialize or uninitialize all disk drives (or SSDs) on a controller, use the Initialize/Uninitialize Physical Devices wizard to clear the meta-data on all drives at once.


To initialize or uninitialize drives with the wizard:

- In the Enterprise View, select a controller.
- On the ribbon, in the Physical Device group, click **Initialize**.



The Initialize/Uninitialize Physical Devices wizard opens.

- Select Initialize or Uninitialize, then click **Next**.



4. Select drives on the controller to initialize or uninitialize, then click **Next**.
5. Review the Summary, then click **Finish**.

10.7 Setting the Physical Drive Boot Priority


Use this option to set the boot priority of a physical device to Primary, Secondary, Primary and Secondary, or None (default). A controller can have only one physical (or logical) device as the primary or secondary boot device. When you select a new physical device as the primary/secondary boot drive, the boot priority of the existing primary/secondary boot drive is overwritten and set to None.

To set the boot priority of a physical device:

1. In the Enterprise View, select a controller, then select a physical drive on the controller.
2. On the ribbon, in the Physical Device group, click **Set Properties**.



The Set Properties window opens.



3. From the Boot Type drop-down list, select Primary, Secondary, Primary and Secondary, or None.
4. Click **OK**.

10.8 Working with Controllers

This section describes how to use maxView Storage Manager to manage the controllers in your storage space:

- To re-scan a controller, see [Rescanning a Controller](#) on page 99.
- To optimize controller performance, see [Optimizing Controller Performance](#) on page 100.
- To change the operating mode of connectors on the controller, see [Changing the Connector Operating Mode](#) on page 101.

10.8.1 Rescanning a Controller

After you connect a disk drive or remove a Ready (non-failed) disk drive from a controller, maxView Storage Manager may not recognize the change until it rescans the controller.

To rescan a controller:

1. In the Enterprise View, select the controller.
2. On the ribbon, in the Controller group, click **Rescan**.



The Rescan window opens.

3. Click the **Rescan** button (on the Rescan window). maxView Storage Manager scans all the channels or ports on the controller you selected.
4. When the rescan is finished, click **Cancel** to close the Rescan window.

10.8.2 Optimizing Controller Performance

You can enable the following performance optimizations on a controller to improve I/O throughput and ensure optimal performance of the arrays and logical drives in your storage space.

Option	Description
Queue Depth	Sets the max drive request queue depth for the controller. Valid values are Automatic, 2, 4, 8, 16, and 32.
Monitor and Performance Delay	Sets the MNP Analysis delay for the controller, in seconds. Set the value to zero to disable Monitor and Performance Analysis. Default is 60 minutes (3600 seconds).
Elevator Sort	Sets the behavior of the drive's write Elevator sort algorithm, a scheduling optimization that prioritizes I/O requests such that disk arm and head motion continues in the same direction. Enabling the elevator sort improves seek times and disabling the elevator sort improves throughput.
Degraded Performance Optimization	For degraded RAID 5 logical drives, enabling this setting directs the controller to attempt to improve performance of large read requests by buffering physical drive requests. Disabling this setting forces the controller to read from the same drives multiple times.
Latency	Enables Flexible Latency Optimization for HDDs. When latency optimization is enabled, the controller detects high-latency I/O requests and applies a cutoff, or threshold, value, after which it suspends elevator sorting and services the request right away. You can set the latency optimization to low, medium, high, aggressive level 1, or aggressive level 2.


To enable/disable performance optimizations on a controller:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



When the Set Properties window opens, click the **Performance** tab.

3. Enable/disable performance optimizations, as needed.



4. Click **OK**.

10.8.3 Changing the Connector Operating Mode

Use this option to change the behavior of the connectors on your Microsemi Adaptec Smart Storage Controller. The connectors on the controller can operate in three modes:

- HBA Mode: exposes physical drives to the operating system
- RAID Mode: exposes only RAID volumes to the operating system
- Mixed Mode: exposes RAID volumes and physical drives to the operating system

By default, products with RAID support are configured to operate in Mixed Mode. Mode options vary, depending on the configuration of logical and physical devices on the connector. For example, you cannot switch the connector to HBA mode if the connector is already configured with a RAID volume.

A reboot is required for connector mode changes to take effect.

Note: Changing from Mixed Mode or HBA Mode to RAID Mode removes access to the physical drives from the operating system.


To change the connector mode on a controller:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



When the Set Properties window opens, click the **Connector** tab.

3. From the drop-down list, change the connector mode for each port to RAID, HBA, or Mixed.



4. Click **OK**.
5. Reboot the server.

10.8.4 Changing the Backplane Discovery Protocol Settings

A discovery protocol is the connector mode protocol to discover the attached backplane on the connectors of your Microsemi Adaptec Smart Storage Controller. Following are the supported backplane discovery protocols in maxView:

- AutoDetect (With SAS/SATA fallback)
- SGPIO
- UBM

While using an unsupported PCIe cable/backplane/expander that causes you to set the PHY mode policy to PCIe, then change the discovery protocol setting to PCIe.

Note:

In maxView, the Backplane Discovery Protocol related properties are displayed both at connector and controller level.


To change the backplane discovery mode protocol settings:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



When the Set Properties window opens, click the **Connector** tab.

3. From the drop-down list, change the backplane setting of each of the connector to Independent, SGPIO, or UBM. You can also change the backplane setting of all the connectors of that controller by selecting the mode from "All connectors" drop-down list.



Note: The Connector tab displays only those connectors that supports the backplane discovery mode.

4. Click **OK**.
5. Reboot the server.

10.8.5 Changing the Expander Minimum Scan Duration Settings

Some expanders do not scan the drives on time for the firmware to discover them. Once the controller is up, the logical device may be marked as failed. To overcome this situation, the firmware has added a new configurable value of minimum scan duration (in seconds) that will force the firmware to stop and wait during discovery.

Whenever the Expander Scan Duration value is changed, it will take affect on the next power cycle. The controller waits for the specified duration and then scans/discovers the drive attached to the expander.


To change the expander minimum scan duration settings:

1. In the Enterprise View, select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



The **General** tab on the Set Properties window opens.

3. To specify the waiting time for the controller, slide through the bar next to the Expander Minimum Scan Duration field.



Note: When the Expander Minimum Scan Duration setting is changed, the power cycle warning is displayed.

4. Click **OK**.

The value of Expander Scan Duration gets displayed in the Properties tab of the controller.

10.9 Updating Controller, Enclosure, and Disk Drive Firmware

Note: This task is recommended for advanced users only.

maxView Storage Manager includes a wizard to help you update the firmware on the controllers, enclosures, and disk drives in your storage space. The wizard updates the firmware for devices of the same type on the local or a remote system.

For example, if your storage space includes disk drives from two different manufactures, you must update the firmware for each manufacturer's drives separately, by running the wizard twice. Additionally, if you have more than one system in your storage space, you must run the wizard for each system separately.

To update the firmware on the controllers, enclosures, or disk drives in your storage space, review the prerequisites in *Before You Begin* on page 104, then follow one of these sets of instructions:

- *Updating the Controller Firmware* on page 104
- *Updating the Disk Drive Firmware* on page 108
- *Updating the Enclosure Firmware* on page 109

10.9.1 Before You Begin

Before you begin, download the latest firmware images from www.adaptec.com, or from your vendor's support site on the World Wide Web. Controller images come in sets of one or more files and have a .bin file extension. Disk drive and enclosure image file names vary by manufacturer.

10.9.2 Updating the Controller Firmware

Use the Firmware Update wizard to update the firmware for one or more controllers of the same type on the local or a remote system.


To update the controller firmware:

1. In the Enterprise View, select a system.


- On the ribbon, in the System group, click **Firmware Update**.



- When the wizard opens, select **Controller**, then click **Next**.




- Select the controllers you want to update, then click **Next**.




- Select flash options for the update, then click **Next**. Choose **Toggle Image** to replace the active image with the backup image.


Note: If you choose Toggle Image, Step [6], Select Files, is skipped.



6. Click **Choose**, browse the file system for the firmware update file (typically, a .bin file), click **Open** to select the file (the button label may be different on your browser).



7. When the file name appears in the Uploaded Firmware File(s) list, click **Next**.
8. Review the update summary, then click **Finish**.



Caution: Do *not* power down the controller(s) while the update is in progress!

9. When the update is complete, click **OK**. Restart the server to activate the new firmware image

10.9.3 Updating the Disk Drive Firmware


Use the Firmware Update wizard to update the firmware for one or more disk drives of the same type on the local or a remote system. The procedure is similar to updating the controller firmware (see [Updating the Controller Firmware](#) on page 104).

To update the disk drive firmware:

1. In the Enterprise View, select a system.
2. On the ribbon, in the System group, click **Firmware Update**.



3. When the wizard opens, select **Disk Drive**, then click **Next**.



4. In the Select Files panel, click **Choose**, browse the file system for the firmware update file, click **Open** to select the file (the button label may be different on your browser).
5. When the file name appears in the Uploaded Firmware File(s) list, click **Next**.
6. In the Attributes panel, select the Chunk Size, from 1-n, in kilobytes (KB), then click **Next**.
7. In the Select Devices panel, select the disk drives you want to update, then click **Next**.
8. Review the update summary, then click **Finish**.

Caution: Do *not* power down the controller(s) while the update is in progress!
9. When the update is complete, click **OK**. Restart the server to activate the new firmware image.

10.9.4 Updating the Enclosure Firmware


Use the Firmware Update wizard to update the firmware for one or more enclosures or expanders of the same type on the local or a remote system.

To update the enclosure/expander firmware:

1. In the Enterprise View, select a system.
2. On the ribbon, in the System group, click **Firmware Update**.



3. When the wizard opens, select **Enclosure/Expander**, then click **Next**.



4. In the Select Files panel, click **Choose**, browse the file system for the firmware update file, click **Open** to select the file (the button label may be different on your browser).

Note: If the upgrade requires multiple firmware update files, update one file at a time or use a combined firmware image to complete the upgrade.

5. When the file name appears in the Uploaded Firmware File(s) list, click **Next**.
6. Select the **Chunk Size**, from 1-n, in kilobytes (KB).
7. Select the firmware **Upgrade Type**:
 - **Firmware**—update the firmware image on the expander or enclosure
 - **Manufacturer**—update the manufacturing image (BOOT SEEPROM) on the expander or enclosure
 - **CPLD**—update the CPLD image on the expander or enclosure
8. Select the firmware upgrade **Mode**:
 - **Download Microcode Data Only**—transfer microcode to the device using one or more write buffer commands; requires system reset or power cycle to activate.

- **Download Microcode with Offsets and Activate**—transfer microcode to the device using one or more write buffer commands and activate immediately.
- **Download Microcode with Offsets, Save and Activate**—transfer microcode to the device using one or more write buffer commands, save to non-volatile storage, then activate.

Note: In this release, maxView Storage Manager supports option 3 only for expander firmware upgrade: Download Microcode with Offsets, Save and Activate.

9. When you are ready to continue, click **Next**.
10. In the Select Devices panel, select the enclosure(s) you want to update, then click **Next**.
11. Review the summary information, then click **Finish**.
 - Caution:** Do *not* power down the controller or enclosure(s) while the update is in progress!
12. When the update is complete, click **OK**. Restart the server to activate the new firmware image, as needed.

11 Monitoring Status and Activity

This section describes how maxView Storage Manager helps you monitor status and activity in your storage space.

11.1 Monitoring Options

maxView Storage Manager provides many ways to monitor the status of your storage space:

- **Event Log**—The main window of maxView Storage Manager features an event log that provides at-a-glance status information about activity (or *events*) occurring in your storage space. All Warning- and Error-level events are also recorded in your *operating system's* event log.

See [Viewing Activity Status in the Event Log](#) on page 112 and [Changing an Operating System's Event Log Setting](#) on page 125.

- **Task Log**—The main window also features a task log that provides status information about the progress of tasks in your storage space, such as the creation of a logical drive.

See [Viewing Task Status in the Task Log](#) on page 113.

- **Storage Dashboard**—Occupying the largest portion of the main window in maxView Storage Manager, the Storage Dashboard provides complete, at-a-glance, information about the components of your storage space, including status information, physical and logical device properties, resources, and reliability indicators for hard drives and SSDs.

See [Viewing Component Status in the Storage Dashboard](#).

- **Chart View**—Provides a visual representation of free and used space for a system, controller, or your entire storage space.

See [Viewing Storage Space Usage in Chart View](#) on page 116.

- **Notifications**—You can set maxView Storage Manager to email status notifications in your choice of format to help you monitor activities in your storage space, such as:
 - Changes in the status of physical devices, such as disk drive failures.
 - Changes on local or remote systems, such as the creation of a hot spare.
 - Changes in temperature in storage enclosures, or that fans or power supplies within an enclosure have failed.

See [Notifying Users by Email About Status and Activity](#) on page 116.

- **Advanced Statistics**—You can set maxView Storage Manager to log advanced I/O and usage statistics for the RAID controllers in your system.

See [Monitoring Advanced Statistics about Activity in Your Storage Space](#) on page 122.

- **Audible Alarm**—A series of beeps sounds whenever a serious event occurs on your storage space.

See [Working with System Alarms](#).

11.2 Checking Status from the Main Window


You can view status information and messages about the activity occurring in your storage space by looking at the *event log*, status icons, and *task log* in the main window of maxView Storage Manager. (You can also view all events for a system in its operating system event log; see [Changing an Operating System's Event Log Setting](#) on page 125.) Using the Storage Dashboard and Chart View, you can also monitor the physical and logical components of your storage space from the main window, including summary information and status, physical and logical device properties and resources, and usage and I/O statistics.

11.2.1 Viewing Activity Status in the Event Log

The Event Log lists activity occurring in your storage space, with the most recent event listed at the top. Status is indicated by icons (see [What Do the Event Status Icons Mean?](#) on page 112) in the left-hand column, as shown in the figure below.

You can view events as they occur in the bottom panel of the maxView Storage Manager main window. The main window displays the last 100 events in your storage space. To view more events, filtered by device (a controller, for example), open the **Event** tab on the Storage Dashboard (see [Viewing Component Status in the Storage Dashboard](#)).

Single-click any event to open the Event Log Detail window to see more information in an easier-to-read format. Use the up and down arrows to view previous or following events.



The screenshot shows the Event Log window with a list of events. The events are sorted by date, with the most recent at the top. The events are:

Severity	Source	Description	Date
Information	WIN-SCFOAT8PRBV.adaptec.com	Bus rescan complete: controller 1.	Tue Jul 10 09:12:00 PDT 2012
Information	WIN-SCFOAT8PRBV.adaptec.com	Started the bus rescan: controller 1.	Tue Jul 10 09:11:40 PDT 2012
Information	WIN-SCFOAT8PRBV.adaptec.com	Initialized drive: controller 1, enclosure 0, slot 4, S/N 4ND4JRH1.	Tue Jul 10 08:25:12 PDT 2012
Warning	WIN-SCFOAT8PRBV		Tue Jul 10 03:23:58 PDT 2014
Warning	WIN-SCFOAT8PRBV		
Information	WIN-SCFOAT8PRBV.adaptec.com		

An arrow points from the warning event in the list to the Event Log Detail window. The detail window shows the following information:

System	SCFOAT8PRBV.adaptec.com	Date and Time	Tue Jul 10 03:23:58 PST 2014
Source	Controller 1	Component	Controller 1
Severity	Warning		
Description			
Prototype controller firmware will expire: controller 1 (Adaptec ASR81605ZQ-EVAL #3131133B8DC Physical Slot: 5). Time remaining: 0 minutes.			
Recommended Action			
Not Supported			

Single-click to view event details.

To make it easier to find a specific event, click on the column heads to sort the events. For example, sorting the events by Severity can help you find specific Error- or Warning-level events quickly.

11.2.1.1 What Do the Event Status Icons Mean?

maxView Storage Manager indicates event status with icons. This table lists the three categories, or types, of events based on severity.

Icon	Status	Examples
	Information	The local system successfully connected to a remote system. A logical drive was created. A hot spare was deleted.


Icon	Status	Examples
	Warning	A logical drive is in a degraded state. A disk drive is being rebuilt. A controller is not responding to an enclosure.
	Error	A controller has failed. A logical drive has failed. A disk drive or hot spare has failed. An enclosure is overheating. Multiple fans or power supplies within an enclosure have failed. An enclosure is not responding.

Note: All Warning- and Error-level events also cause the audible alarm to sound. See [Working with Controllers](#) on page 99 for more information.

11.2.2 Viewing Task Status in the Task Log

The Task Log shows the status and progress of tasks in your storage space, with the most recent task listed at the top.

Single-click any task to open the Task Log Detail window to see more information in an easier-to-read format.




11.2.3 Viewing Component Status in the Storage Dashboard

The Storage Dashboard provides detailed information about the components of your storage space, including local and remote systems, controllers, arrays, logical drives, enclosures, disk drives and SSDs, and maxCache Devices. Occupying the largest portion of the main window in maxView Storage Manager, the Storage Dashboard organizes component information by category, with a tabs providing one-click access to summary information and status, properties, resources, and usage statistics.

The information on the Storage Dashboard varies, depending on which component is selected in the Enterprise View. The figure below shows the Storage Dashboard for a controller. Tabs provide access to summary information, controller properties, and resources. The Events tab shows filtered events for the selected device (see [Viewing Activity Status in the Event Log](#) on page 112).

Note: For information about Chart View, on the right side of the Storage Dashboard, see [Viewing Storage Space Usage in Chart View](#) on page 116.



The following table lists the categories and types of information provided on the Storage Dashboard for each component in your storage space. All top-level nodes in the Enterprise View (System, Controller, Arrays, Logical Drives, Physical Devices, and so on) include a Summary tab and Events tab.

Component	Categories/Tabs	Examples
System	Summary Properties	System name and IP address Operating system Number and type of controllers Alarm status (see Working with System Alarms) Web Server settings SMTP settings
Controller	Summary Properties Resources Connectors maxCrypto	Model, key features, manufacturing data, driver and firmware version, controller mode, and status Number of physical drives, arrays, logical drives, and status Power management features I2C address for PBSI interface (hex), I2C clock speed and clock stretching status maxCache status maxCrypto status (see Checking maxCrypto Status on page 79) Health and activity of flash backup module, if present ("Green backup" status) Connector functional mode Performance optimizations and other settings Physical drive assignments by logical device (see Revealing More Device Information on page 21)
Arrays	Summary Resources	Total size and unused size Spare rebuild mode Logical drive RAID level, size, status
Logical drives and maxCache Device	Summary Resources	Raid level, segment and group (RAID 10 only), size, mount point, status Member drives and sizes
Enclosure	Summary Resources Slots	Enclosure type, vendor, model and status Fan, power supply, and temperature status (see Monitoring Enclosure Status on page 115) Speaker status Slot allocation and usage

Component	Categories/Tabs	Examples
Hard drives and SSDs	Summary Resources SMART Statistics	Drive type (hard drive, SSD, SMR), interface type (SAS/SATA), vendor, and model Drive state (Ready, Optimal, Hot Spare), mount point Channel number and device ID Transfer speed Drive segment allocation SMART statistics (see Viewing SMART Statistics on page 115)

11.2.3.1 Monitoring Enclosure Status

If your storage space includes an enclosure with an enclosure management device, such as a SCSI Accessed Fault-Tolerant Enclosure (SAF-TE) processor, maxView Storage Manager displays temperature, fan, and power module status on the Storage Dashboard, as shown in the figure below.

Resources	
Fan(s)	6
Optimal	6
Malfunctioning	0
Not Installed	0
Power Supplies	2
Optimal	2
Malfunctioning	0
Not Installed	0
Temperature Sensor(s)	2
Normal	2
Abnormal	0
Not Installed	0
Speaker(s)	1
On	1
Off	0
Not Installed	0

11.2.3.2 Viewing SMART Statistics

You can use the Storage Dashboard to view various indicators of reliability for the SATA hard drives and SSDs in your storage space. maxView Storage Manager displays SMART statistics for the drives using *Self-Monitoring, Analysis and Reporting Technology* available on most contemporary hard drives and non-spinning storage devices. You can use this information to verify the health of your hard drives and SSDs and to predict drive failures.

To view the SMART statistics for a hard drive or SSD, select the drive in the Enterprise View, then click the **SMART** tab on the Storage Dashboard. For SSDs, the statistics include wear-level and longevity indicators, as shown in next figure. Refer to your drive vendor's data sheet for a description of individual report items.

Summary Resources SMART Events

Please refer to drive vendors data sheet for description

(2 of 3) 1 2 3

ID	Name	Norma
0xBB	Temperature	100
0xC2	Reported I/O Error Detection Code Errors	30
0xC3	Unknown Attribute	120
0xC4	Unknown Attribute	100
0xC9	Unknown Attribute	120
0xCC	Unknown Attribute	120
0xE6	Life Curve Status	100
0xE7	SSD Life Left	100
0xE9	Unknown Attribute	0
0xEA	Unknown Attribute	0


(2 of 3) 1 2 3

SSD wear-level and longevity indicators

11.2.4 Viewing Storage Space Usage in Chart View

Chart View provides a visual representation of the free and used space for a system, controller, array, or your entire storage space (all systems and controllers). Located on the right side of the Storage Dashboard in the maxView main window, Chart View displays a pie chart of storage space usage.

To view storage space usage in Chart View, simply select a component in the Enterprise View (a system, for instance); the chart view is updated immediately.



11.3 Notifying Users by Email About Status and Activity

You can set up maxView Storage Manager to send email messages (or *notifications*) to one or more email addresses when an event occurs on a system, such as the creation of a logical drive or the failure of a disk

drive. Email notifications can help you monitor activity on your entire storage space from any location, and are especially useful in storage spaces that include multiple systems running the maxView Storage Manager only.

Only the users you specify receive email notifications. You can specify which types of events generate email messages (Error, Informational, Warning). You can also specify if you want to be notified instantly when an event occurs to ensure that urgent issues receive immediate attention from the right people. Alternatively, you can specify that you want events “coalesced” and receive only one email message for each event type.

Follow the instructions in this section to:

- Set up email notifications (see [Setting Up Email Notifications](#) on page 117).
- Send a test email (see [Sending a Test Message](#) on page 119).
- Modify or remove an email recipient (see [Modifying or Removing an Email Recipient](#) on page 120).
- Modify email server settings (see [Modifying Email Server Settings](#) on page 121).
- Disable email notifications (see [Disabling Email Notifications](#) on page 121)

11.3.1 Setting Up Email Notifications


This section describes how to set up email notifications for one system. If you want to monitor multiple systems by email, you must complete the tasks in this section for each one separately.

Before you begin, note this information:


- The address of your Simple Mail Transfer Protocol (SMTP) server (host name and domain, or TCP/IP address)
- The email address of each person who will receive email notifications

To set up email notifications:


1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **System Settings**.




3. When the System settings window opens, click the **SMTP** tab.
4. Select **Enable Email Notifications**.
5. Enter the IP address of your SMTP server and the server's port number (or use the default port).



6. If authentication is enabled on your SMTP server (that is, the server requires authentication details before it will send messages to users), select **Use Secure Mail Server**, then enter the SMTP server's login credentials (username/password) in the space provided.
7. On the System settings window, click the **Email** tab. The Email Notifications Manager opens.




8. Click **Add Email**. When the Add Email window opens, enter the recipient's email address, select the level of events that will trigger an email notification for that recipient (Error, Error/Warning, Error/Warning/Informational), then select the notification type—Instant or Coalesced. To include a support archive file with the email, click **Attach Support.zip**, then click **OK**. (For more information about event levels, see [What Do the Event Status Icons Mean?](#) on page 112; for more information about the support archive file, see [Creating a Support Archive File](#) on page 145.)



Repeat this step to add more email recipients.

Each recipient appears in the Email Notifications Manager, as shown below:



	Email ID	Event Type	Notification Type
<input type="checkbox"/>	ann_thomas@lantec.com	Error, Warning, Information	Coalesced
<input type="checkbox"/>	richard_stanton@lantec.com	Error	Instant


9. When you're done adding email recipients, click **OK**.
The email recipients and your SMTP server settings are saved.
10. Repeat the steps in this section *for each system* you want to monitor with email notifications, then continue by sending test messages to all recipients (see [Sending a Test Message](#) on page 119).

11.3.2 Sending a Test Message


To ensure that an email recipient is receiving event notifications, you can send them a test message.

To send a test message:

1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **System Settings**.



- When the System settings window opens, click the **Email** tab. The Email Notifications Manager opens.
- Select one or more email addresses to send a test message to. To select all addresses, click the check box at the top of the list, as shown in the figure below.



- Click **Send Test Email**.

If the test is successful, the email recipient(s) receive the test message. If the test fails:


- Ensure that the recipient's email address is correct. (See [Modifying or Removing an Email Recipient](#) on page 120.)
- Ensure that your SMTP server address is correct. (See [Modifying Email Server Settings](#) on page 121.)
- Try sending the test message again.

11.3.3 Modifying or Removing an Email Recipient

This section describes how to modify a recipient's email address, change the types of event notifications the recipient receives, or stop sending email notifications to a recipient from a selected system.

To modify recipient information or to stop sending email notifications to a recipient:

- Select the Enterprise View node.
- On the ribbon, in the System group, click **System Settings**.



3. When the System settings window opens, click the **Email** tab.

The Email Notifications Manager opens.

4. Select the email recipient you want to modify or remove, then:
 - Click **Modify Email**, change the recipient information, as needed, then click **Modify** to save your changes.
 - Or,*
 - Click **Delete Email** to remove the recipient from the notification list.

The changes become effective immediately.

5. Click **OK** to close the Email Notifications Manager.


11.3.4 Modifying Email Server Settings

You can modify these email server settings, if required:

- Address and port of your SMTP server
- 'From' address that will appear in email notifications
- Secure server login credentials

To modify email server settings:

1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **System Settings**.



3. When the System settings window opens, click the **SMTP** tab.
4. Edit the SMTP server settings as required, then click **OK** to save your changes.


11.3.5 Disabling Email Notifications

This section describes how to disable email notifications on a selected system.

Note: If you disable email notifications, events continue to be generated but email messages won't be sent.

To disable email notifications:

1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **System Settings**.



3. When the System settings window opens, click the **SMTP** tab.
4. Clear the **Enable Email Notifications** check box.
5. Click **OK** to save your changes.

11.4 Monitoring Advanced Statistics about Activity in Your Storage Space

You can set maxView Storage Manager to log advanced I/O and usage statistics about the RAID controllers, hard drives, SSDs, and logical drives in your storage space, including maxCache caching statistics.

Follow the instructions in this section to:

- Enable statistics logging on a controller (see [Setting up Statistics Logging](#) on page 122).
- View the advanced statistics for a controller, hard drive, SSD, logical drive, or maxCache Device (see [Viewing Advanced Statistics](#) on page 124).
- Reset the statistics counters (see [Resetting the Statistics Counters](#) on page 125).

11.4.1 Setting up Statistics Logging


To set up statistics logging for the controllers in your storage space:

1. In the Enterprise View, select a system, then select a controller.
2. On the ribbon, in the Controller group, click **Set Properties**.



The Set Properties window opens.


3. In the Advanced Statistics drop-down list, select **Enabled**.



4. Click **OK**.


11.4.2 Viewing Advanced Statistics

Use the Statistics Viewer to view the advanced statistics for the controllers, hard drives, SSDs, logical drives, and maxCache Devices in your storage space. To view the statistics for a component, select it in the Enterprise View then, on the Storage Dashboard, click the **Advanced Statistics** tab. See [Advanced Statistics Descriptions](#) on page 178 for a complete list of Advanced Statistics descriptions.



Statistics Parameter	Value
Average Dirty Cache Lines	0
Free Largest Transfer Buffer Sectors	639486511
Average Free Total Transfer Buffer Sectors	639486511
Average Free Logical Requests	2659402
Average Free Physical Requests	2657900
Average Free Processor RAM in KB	66269535
Average Locked Stripes	0
Average Locked Stripes Waiting	0
Average Write Cache Sectors	981871980
Command List Count	129506
Command List Latency	636879670 (100ths of milliseconds)
Logical Request Count	105818
Maximum DMA Transfer Queue Depth	836879870
Maximum Outstanding Command List	1
Percent Active	2455065 (100ths of percent)
Percent Busy	Not Available
Percent Time Waiting For DMA	0 (100ths of percent)
Sample Interval	26314 (100ths of seconds)
Structure version	110
Volume Mask	13

The statistics listed vary, depending on which type of component you select. The following example shows the advanced statistics for a controller (on the left) and the maxCache statistics for a logical drive, including cache hits and misses and a histogram of cache hit distribution.



Controller 1

General

Adapter General

Statistics Parameter	Value
Statistics Version	3.0
SSD Cache Size	380721
Pages on Fetch List	0
Fetch List Candidates	0
Candidate Replacements	0
Pages on dirty List	0
Pages on failed List	0

Adapter Outstanding Commands

	0	1	2	4	8
	138	0	0	0	0

LogicalDrv4 - Optimal (99.989 GB)

MaxCache Statistics

Statistics Parameter	Value
Read hits	19461
Write hits	62523
Valid Pages	1
Updates on Writes	0
Invalidations by RW Balance	0
Invalidations by Updates	0
Invalidations by Others	0
Page Fetches	0

SSD Cache Page Hit Distribution

	0	1	2	4	8	16	32	64	128	256	512	1024	2048
	40514	8355	4314	330	1	0	0	0	1	0	0	0	0

11.4.3 Resetting the Statistics Counters

In maxView Storage Manager, advanced statistics are cumulative. That is, the statistics continue to accumulate until you reset the counters for a controller to zero.

To reset the statistics counters for a controller:

1. In the Enterprise View, select a system, then select the controller you want to reset.
2. On the ribbon, in the Controller group, click **Set Properties**.



The Set Properties window opens.

3. Click the **Reset Now** button, then click **Apply**.




The statistics counters are reset.

11.5 Changing an Operating System's Event Log Setting

In addition to the maxView Storage Manager event log, all Warning- and Error-level events on a system are recorded in its *operating system* event log. You can select the type of events that are recorded, or you can disable operating system event logging.


To change or disable operating system event logging on a system:

1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **System Settings**.



The System Settings window opens.

3. In the Save Events to OS Log drop-down list, select the type of events that you want to log, then click **OK**.



4. Restart maxView Storage Manager to apply the new setting.

12 Managing Your Storage Space

This section describes the advanced features in maxView Storage Manager that help you manage your storage space. You can:

- Deploy servers with a *server template file*
- Manage remote systems and auto-discovery tasks with the Remote System wizard
- Clear a controller configuration
- Change the Web Server port
- Grant Standard users Admin Privilege

12.1 Deploying Servers

maxView Storage Manager helps you deploy servers in your storage space without configuring each server manually. You can select an optimally configured server in your storage space, save its configuration to a *server template file*, then duplicate the configuration on servers throughout your network.

The basic procedure works like this:

1. Choose the system you want to use as the model for other servers in your storage space.
2. Save the configuration to a server template file.
3. Log in to each remote system in your storage space and restore the configuration from the server template file.


The following sections provide details on each of these steps.

12.1.1 Creating a Server Template File

This procedure saves the configuration of a system that you want to use as a model for other servers in your storage space. It creates a server template file in XML format, which defines the controller type(s), operational settings, physical drive size, logical drive size, RAID level, and more. The default name of the server template file is `ControllerConf.xml`.


To create a server template file:

1. In the Enterprise View, select a system.
2. On the ribbon, in the System group, click **Manage Configuration**.



The Manage Configuration wizard opens.

3. Select **Save Configuration**, then click **Next**.



4. Review the Summary information, then click **Finish**.
5. When the File Download window opens, click **Save File**, then click **OK**.
 - Note:** The procedure for downloading and saving the template file may vary, depending on the Web browser.
6. Continue with [Duplicating the Server Template](#) to deploy the same configuration on multiple systems in your storage space.

12.1.2 Duplicating the Server Template


When you are ready to duplicate the server template on other systems in your storage space, you can restore the configuration from the server template file.

Keep in mind that:

- The server template file (default, `ControllerConf.xml`) is editable. For example, you may need to change the disk drive capacity or logical drive size to accommodate the differences on each machine.
- Drives from the same vendor with slightly different capacities (147 GB vs. 150 GB, for instance) are considered interchangeable. If the logical drive capacity changes as a result of the size difference, it is scaled accordingly. For example, if the new drives have 4% more capacity due to vendor or model changes, then all logical drives are increased in size by 4%.


To duplicate the server template on another system:

1. In the Enterprise View, select a system.
2. On the ribbon, in the System group, click **Manage Configuration**.




The Manage Configuration wizard opens.

3. Select **Restore Configuration**, then click **Next**.




4. In the Configuration File panel, click **Choose**, navigate the file system to your server template file, then click **Open**. When the file name appears in the "selected file" area (circled in blue in figure below), click **Upload**, wait for the upload to complete, then click **Next**.



5. In the Restore Options panel, choose a Force option if a controller does not support all of the features of the template controller, or if the drive capacity on the new system does not match the configuration in the server template file. The default is Force None. You can choose to:

Option	Description
Force All	To force deployment of all features
Force Size	To force deployment of just the logical drives




6. To apply the configuration based on SlotID rather than DeviceID, click the **Slot ID** check box.
7. Click **Next**, review the summary information, then click **Finish**.
maxView Storage Manager duplicates the system configuration on the new controller.

12.2 Managing Remote Systems

maxView Storage Manager has a wizard to help you manage the remote systems in your storage space. The wizard simplifies the process of connecting to remote systems from the local system and adding them to the Enterprise View.

When you start maxView Storage Manager, an “auto-discovery” task runs in the background, continuously searching your network for systems running the maxView Redfish server. The wizard presents a list of discovered systems (see figure below). You can select systems to add to the Enterprise View when you start maxView Storage Manager; add systems manually if they are not discovered automatically; and remove systems that you no longer want to manage.



12.2.1 Adding Remote Systems with the Wizard

For basic instructions for adding remote systems with the wizard, see [Logging into Remote Systems from the Local System](#). Once you add a system in the wizard, it automatically appears in the Enterprise View each time you start maxView Storage Manager. You can work with a remote system's controllers, disk drives, and logical drives as if they were part of your local system.

Note: The wizard adds all selected systems to the Enterprise view even if login fails on some systems. For those systems, try running the wizard again with different login credentials.

12.2.2 Manually Adding a Remote System


You can add a remote system manually if auto-discovery fails to find the system on your network.

To manually add a remote system:

1. On the ribbon, in the Home group, click **Add System**.



2. When the Add System window opens, click **System Not Present**. The Add System Manually window opens.
3. Enter the system name and login credentials in the space provided. Select the Management Protocol from the drop-down list and specify the Port number, then click **Add**.



The window closes and the system appears in the discovered systems list in the Add System window.

4. Select the system, then click **Add**. maxView Storage Manager connects to the remote system and adds it to the Enterprise View.


12.2.3 Removing a Remote System

If you no longer want to manage a remote system, you can remove it from the Enterprise View.

Note: Removing a remote system from the Enterprise View does not take it off-line.


To remove a remote system:

1. On the ribbon, in the Home group, click **Delete System**.



The Delete System window opens.

2. Select the system(s) you want to remove. To select all systems in the list, click the checkbox at the top of the window.



3. Click **Delete**.


maxView Storage Manager removes the remote system(s) from the Enterprise View.

12.2.4 Changing the Auto-Discovery Settings


Auto-discovery, in maxView Storage Manager, is disabled by default. If enabled, the auto-discovery task runs in the background each time maxView Storage Manager is started. You can enable the auto-discovery if desired, and configure the auto-discovery settings described below.

To change the auto-discovery settings on a system:

1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **System Settings**.



The System Settings window opens for that system. The auto-discovery settings appear at the bottom.



3. To enable/disable auto-discovery, select `Enable Auto Discovery`. (This option toggles between enabled and disabled.)
4. Update the auto-discovery settings. In the `Auto Discovery Interval` field, enter the number of seconds between each auto-discovery check. This number determines how often maxView Storage Manager checks for changes in remote system resources.
5. Click **OK** to save the changes.

12.2.5 Importing and Exporting Remote Systems

maxView provides the 'Import and Export systems' feature to add multiple systems and export the added systems in "SystemConf.json" file, which can be used later to import the added systems in maxView running on another system.

The Import and Export feature gets enabled at the "Enterprise View" level in the Manage Configuration ribbon icon.

Note:

Export feature is applicable only when maxView GUI manages at least one remote system.

To export a system:

1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **Manage Configuration**.




Manage Configuration


Save and Restore configuration of controller(s) on system node. Import and Export system(s) on the enterprise view node.

The Manage Configuration wizard opens for that system.

3. Select **Export System(s)** option, then click **Next**.



4. Select the systems that need to be exported. Click **Next**.



Note:

To get the details of the respective systems, hover the cursor on the system name. It shows details like system name, IP address, operating system, and communication protocol.

The **Manage Configuration Summary** page appears.

5. Click Finish.


The exported systems are downloaded in a "SystemConf.json" file.

To import a system:


1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **Manage Configuration**.

The **Manage Configuration** wizard opens for that system.

3. Select Import System(s) option, then click Next.



4. Click Choose to specify the path of the "SystemConf.json" file.



The file gets uploaded under "Uploaded Configuration File" field.


5. Click **Next**.

The **Select Systems** screen appears.


6. Select the system name(s) and specify the login credentials. Select **Single Sign-On** option to specify the **User Name** and **Password** for all the selected systems that have same credentials. Otherwise, specify each system's credentials manually.

Note:

Single sign-on option is enabled only when more than one system is selected for import.



Hover the cursor on the system name to get the details of the respective systems. It shows details like system name, IP address, operating system, and communication protocol.



7. Click **Next**.

The **Manage Configuration Summary** page appears that shows the list of imported systems.

8. Click **Finish**.

The imported systems will appear in the Enterprise View.

12.3 Clearing the Controller Configuration

You can clear the configuration of a controller to accommodate changes in your storage space. For example, you may want to clear a controller if you upgraded your hardware or if you plan to move the controller to another machine. Clearing the configuration destroys the controller meta-data, including array and logical device information, partition information, maxCache information, and so on. Once you clear the controller configuration, your online data is no longer accessible.

Caution: When you clear a controller configuration, you lose all data stored on that controller. Be sure you no longer need the data on the controller before proceeding.


To clear the controller configuration:

1. In the Enterprise View, select a system, then select a controller on that system.
2. On the ribbon, in the Controller group, click **Manage Configuration**.



The Manage Configuration wizard opens.

3. Select **Clear Configuration**, then click **Next**.




4. Review the Summary information, then click **Finish**.

12.4 Changing the Web Server Port


You can change the port used by the maxView Storage Manager Web Server, if needed, to accommodate changes in your network or IT requirements. The Web Server can use any open port for communication. The default port is 8443. If you change the port, you must restart maxView Storage Manager for the change to take effect.

To change the Web Server port:

1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **System Settings**.



When the System Settings window opens, click the **Web Server** tab.




3. Enter the new Web Server port. Optionally, click **Secured Protocol** to enable/disable secure communication over https.
4. Click **Apply**.
5. Restart maxView Storage Manager.

12.5 Granting Standard Users Admin Privilege

Normally, Standard users have restricted access to your storage space, with limited ability to perform non-destructive operations in maxView Storage Manager (see [Working in maxView Storage Manager](#) on page 15). You can grant Standard users Admin Privilege to accommodate changes in your system policies or IT requirements.

To grant Admin privilege to Standard users:

1. Select the Enterprise View node.
2. On the ribbon, in the System group, click **System Settings**.



The System Settings window opens.

3. Click the **Allow Standard User as Admin** check box, then click **Apply**.
4. Restart the webserver.

13 Solving Problems

This section describes how to troubleshoot the components in your storage space.

13.1 General Troubleshooting Tips


If you experience problems installing or using maxView Storage Manager, try these troubleshooting tips first:

- Ensure that all managed systems are powered on and that you are logged in to any remote systems that you want to manage. (See [Logging into Remote Systems from the Local System](#) on page 25 for more information.)
- Check all cable connections.
- Try uninstalling and reinstalling maxView Storage Manager.
- Check the Release Notes for compatibility issues and known problems.

13.2 Identifying a Failed or Failing Component

When you receive notice of a Warning- or Error-level event, use maxView Storage Manager's *rapid fault isolation* feature to quickly identify the source of the problem.

For instance, in this example, a disk drive has failed. To find the failed disk drive, expand the tree in the Enterprise View, look for the orange and red warning and error icons, then continue tracing the problem to its source.



Enterprise View

WIN-V07FSIQGHQE Error is on local system...

Controller 1

Controller 2 ...on Controller 2

Arrays and Logical Devices

Array4

Array3

LogicalDrive3

LogicalDrive2 ...affecting two Logical Drives

Array2

LogicalDrive1

Array1

Physical Devices

Enclosure 0

Slot 1

Slot 2 ...and a physical device in Slot 2 in Enclosure 0. Click Slot 2 to check the device status on the Storage Dashboard and continue tracing the fault to its source...

Summary Properties Resources Events

Device Status

Type : Disk drive State : Failed

Vendor : Unknown Negotiated transfer speed : Failed

Model : ST3250620NS Write-cache mode : Write back

Serial No : 5QE4MAS1 S.M.A.R.T warnings : 0

Firmware level : 3.AEK SSD(non-spinning) : false

Reported channel : 0 MaxCache capable : Not supported

...a disk drive failure.

13.3 Recovering from a Disk Drive Failure

This section describes how to recover when a disk drive or SSD fails:

- If the logical drive is protected by a hot spare, see [Failed Disk Drive Protected by a Hot Spare](#) on page 141.
- If the logical drive is *not* protected by a hot spare, see [Failed Disk Drive Not Protected by a Hot Spare](#).
- If there is a disk drive failure in more than one logical drive simultaneously, see [Failure in Multiple Logical Drives Simultaneously](#) on page 142.
- If it is a RAID 0 logical drive, see [Disk Drive Failure in a RAID 0 Logical Drive](#) on page 142.
- If multiple disk drives fail within the same logical drive, see [Forcing a Logical Drive with Multiple Drive Failures Back Online](#) on page 142.

13.3.1 Failed Disk Drive Protected by a Hot Spare

If a disk drive in a logical drive fails and that logical drive is protected by a hot spare, the hot spare is automatically incorporated into the logical drive and takes over for the failed drive.

For example, if a disk drive fails in a RAID 5 logical drive, the logical drive is automatically *rebuilt*, with its data reconstructed using the hot spare in place of the failed drive. You can access the logical drive while it's rebuilding.

To recover from the failure:

1. Remove and replace the failed disk drive, following the manufacturer's instructions.

2. If the logical drive is protected with a *dedicated* hot spare, data is moved back to its original location once the controller detects that the failed drive has been replaced. Once the data is copied back, the hot spare becomes a permanent part of the array on failure of array. No further action is necessary. If the logical drive is protected with an *auto-replace* hot spare, the spare becomes a permanent part of the array. You must designate a new hot spare to protect the logical drive(s) on that array.

See [Protecting Your Data](#) on page 40 for more information about managing spares.

13.3.2 Failed Disk Drive *Not* Protected by a Hot Spare

If a disk drive in a logical drive fails when the logical drive is not protected by a hot spare, remove and replace the failed disk drive. The controller detects the new disk drive and begins to rebuild it. You can access the logical drive while it's rebuilding.

For example, when one of the disk drives fails in a RAID 1 logical drive, the logical drive is *not* automatically rebuilt. The failed disk drive must be removed and replaced before the logical drive can be rebuilt.

If the controller fails to rebuild the logical drive, check that the cables, disk drives, and controllers are properly installed and connected. Then, if necessary, follow the instructions in [Rebuilding Logical Drives](#).

13.3.3 Failure in Multiple Logical Drives Simultaneously

If a disk drive fails in more than one logical drive at the same time (one failure per logical drive), and the logical drives have hot spares protecting them, the controller rebuilds the logical drives with these limitations:

- A hot spare must be at least as big as the smallest disk drive in the array that it might replace.
- Failed disk drives are replaced with hot spares in the order in which they failed. (The logical drive that includes the disk drive that failed first is rebuilt first, assuming an appropriate hot spare is available—see the previous bullet.)

Note: If the number of disk drive failures exceeds the number of hot spares, see [Failed Disk Drive Not Protected by a Hot Spare](#).

13.3.4 Disk Drive Failure in a RAID 0 Logical Drive

Because RAID 0 volumes do not include redundancy, if a disk drive fails in a RAID 0 logical drive, the data cannot be recovered.

Correct the cause of the failure or replace the failed disk drives. Then, restore your data from backup, if available. To protect the RAID 0 logical drive, set the spare activation mode to "predictive".

13.3.5 Forcing a Logical Drive with Multiple Drive Failures Back Online

If multiple disk drives fail in the same logical drive, you may be able to recover the data by forcing the logical drive back online. For instance, if two drives fail in a RAID 5, forcing it online may allow you to access the data, depending on which disk drives failed.

Caution: This procedure is not guaranteed to successfully recover your logical drive. The surest way to recover your data is to restore the failed logical drive from backup.

To force a logical drive online:

1. In the Enterprise view, select the failed logical drive (see [Identifying a Failed or Failing Component](#) on page 140).
2. On the ribbon, in the Logical Device group, click **Force Online**.



3. Click **Force**, then click **OK**.

13.3.6 Healing an Array

You can use the Heal Array operation to replace failed physical drives in the array with healthy physical drives. After replacement, the original array and logical drive numbering is unaffected.

The Heal Array operation is part of the Modify Array wizard (see [Moving an Array](#) on page 59). It is available in the wizard only if:

- The array has at least one failed drive.
- The array is not rebuilding to a spare.
- A sufficient number of Ready physical drives of the same type and correct size are available to replace each failed physical drive in the array.

Note: The correct size is defined as a drive as large as the smallest drive on the array, but no larger than the smallest spare.


For a RAID 0 volume, the heal operation recreates the volume. For other RAID volume types, the heal operation rebuilds the volume.

To heal an array:


1. In the Enterprise View, select an array.
2. On the ribbon, in the Array group, click **Modify Array**.



3. When the wizard opens, select **Heal Array**, then click **Next**.



4. Select one or more drives to replace the failed drives in the array.



Note: The drives must have sufficient capacity to hold all of the logical drives in the array.

5. Click **Next**, review the summary information, then click **Finish**.

13.4 Rebuilding Logical Drives

A *hot-swap rebuild* occurs when a controller detects that a failed disk drive in a logical drive has been removed and then reinserted.

Note: You can use the Heal Array operation as an alternative to a hot-swap rebuild if you have a sufficient number of Ready physical drives of the same type in your storage space. See [Healing an Array](#) on page 143.

To start a hot-swap rebuild:

1. Following manufacturer's instructions, gently pull the failed disk drive from the server without fully removing it, then wait for it to spin down fully before continuing.
2. If there is nothing wrong with the disk drive, reinstall it, following manufacturer's instructions.
If necessary, replace the failed disk drive with a new disk drive of equal or larger size.
3. The controller detects the reinserted (or new) disk drive and begins to rebuild the logical drive.

13.5 Creating a Support Archive File


Your support representative might ask you to create a support archive file to help diagnose a problem with your system. Saved information includes device logs, drive logs, event logs, error logs, controller logs, history logs, and SMART statistics.

To create the support archive file:

1. In the Enterprise View, select the system on which the problem is occurring. (Look for the orange or red error icons in the Enterprise View.)
2. On the ribbon, in the System group, click **Archive Management**.



The Archive Management window opens.



3. Select the logs you want to save.
4. Click **Save**.
5. When the File Download window opens, click **OK**.
6. In the Archive Management window, click **Clear All Logs** to clear, or **Cancel** to exit.

A Silent Installation on Windows and Linux

This appendix describes how to complete a silent installation of maxView Storage Manager on Windows and Linux systems. A silent installation uses command line parameters to complete an installation without messages or user interaction.

Note: Silent installation on Linux is supported on Red Hat, CentOS, and SLES only.

A.1 Completing a Silent Installation

This section describes the silent installation process for Windows and Linux.

A.1.1 Windows Silent Installation

To complete a silent installation on Windows:

1. Open a command prompt window, then change to the directory where you downloaded the Windows setup program (see [Downloading the Installer Package](#) for details).
2. Run the silent installation from the command line:

Option	Description
Windows 64-bit	<code>setup_asm_x64.exe /s /v"/qn <properties>"</code>

where *<properties>* is one or more of the options listed in [Switches, Properties, and Values](#) on page 147.

Separate properties with spaces and enclose all properties after /v in quotes, with NO leading space. Separate feature names for the ADDLOCAL property with commas.

A.1.1.1 Switches, Properties, and Values

This section describes the command line options for Windows silent installation. These options are not supported on Linux.

Switch or Property	Description
/s (required)	Suppress dialog boxes.
/v (required)	Pass command line parameters to the setup program.
/qn	Suppress progress bar during installation.
/qb	Show progress bar during installation.
/lv* <path> (optional)	Generate verbose installation log at <path>. Example: /lv* c:\pmc.log
INSTALLDIR (optional)	Specifies the installation path. If specified, the installation path must be enclosed in escaped quotation marks. Example: INSTALLDIR="C:\Program Files\Adaptec\maxView Storage Manager\ Note: The default installation path is "C:\Program Files\Adaptec\maxView Storage Manager".
ADDLOCAL (optional)	<ul style="list-style-type: none"> ALL (default)—Installs the maxView Storage Manager Console (GUI and Redfish server), and ARCCONF (CLI). If you specify ALL, do not specify any of the following values. ARCCONF—Installs the Command Line Interface tool (ARCCONF) RestfulServer—Installs the maxView Storage Manager Redfish Server Console—Installs the maxView Storage Manager GUI Note: Use commas to separate multiple values.

A.1.2 Linux Silent Installation

To complete a silent installation on Red Hat Linux, CentOS, or SLES:

1. Open a shell window, then change to the directory where you downloaded the Linux installer package (see [Downloading the Installer Package](#) for details).
2. Run the silent installation from the command line using one of these commands (x.xx-xxxx=version-build number):

Option	Description
Linux 64-bit	./StorMan-X.XX-XXXX.x86_64.bin --silent

Note: Linux systems also support silent upgrade and silent removal. See [Example Command Line Installations](#) on page 148.

A.2 Example Command Line Installations

This section shows typical command line installations for Windows and Linux. In the Linux examples, <x.xx>-<xxxxx>=version-build number.

- Normal Windows Installation:

```
setup_asm_x64.exe /s /v"/qb /lv* c:\pmc.log"
```

- Install to Specific Location on Windows:

```
setup_asm_x64.exe /s /v"/qb INSTALLDIR="\C:\Program Files\Adaptec\maxView Storage Manager\""
```

- Install Specific Feature on Windows:

```
setup_asm_x64.exe /s /v"/qb ADDLOCAL=ARCCONF /lv* c:\pmc.log"
```

- Normal Linux Installation:

```
./StorMan-<x.xx>-<xxxxx>.x86_64.bin --silent
```

- Linux Software Upgrade:

```
./StorMan-<x.xx>-<xxxxx>.x86_64.bin --upgrade
```

- Linux uninstallation (removal):

```
./StorMan-<x.xx>-<xxxxx>.x86_64.bin --remove
```

Alternative: `rpm -e StorMan`

B Configuring SNMP Notifications on Windows and Linux

This appendix describes how to enable SNMP trap notifications on Windows and Linux.

After installing and configuring the SNMP service, you can monitor activity in your storage space with the maxView Storage Manager GUI or any OS monitoring tool, such as a Mib Browser.

B.1 Setting Up SNMP Notifications on Windows

1. Install and enable the SNMP service on your Windows system. Define the SNMP community to which to send trap messages ("public", for instance). Then designate that name as an Accepted Community in the SNMP Service Properties.

For details on installing and configuring SNMP on Windows, refer to your operating system documentation.

2. On Windows Server 2012 and Windows 8.x systems, the SNMP sub-agent does not have permission to open a socket over TCP/IP or UDP, preventing it from communicating with the maxView Storage Manager. Use the following PowerShell scripts to allow the SNMP sub-agent to communicate with the maxView Storage Manager and send trap notifications:

- a. Outbound Rule for Port 34572:

```
$OutBound = @{
    DisplayName = "Maxview Outbound Rule on TCP port 34572 for SNMP Service"
    Direction = "Outbound"
    InterfaceType = "Any"
    Action = "Allow"
    Protocol = "TCP"
    Service = "snmp"
    Program = "${Env:systemdrive}\WINDOWS\system32\snmp.exe"
    Enabled = "TRUE"
    RemotePort = "34572"
    PolicyStore = "ConfigurableServiceStore"
}
New-NetFirewallRule @OutBound
```

- b. Inbound Rule for Port 34572:

```
$InBound = @{
    DisplayName = "Maxview Inbound Rule on TCP port 34572 for SNMP Service"
    Direction = "Inbound"
    InterfaceType = "Any"
    Action = "Allow"
    Protocol = "TCP"
    Service = "snmp"
    Program = "${Env:systemdrive}\WINDOWS\system32\snmp.exe"
    Enabled = "TRUE"
    RemotePort = "34572"
    PolicyStore = "ConfigurableServiceStore"
}
New-NetFirewallRule @InBound
```

3. To remove the NetFirewall rules (as needed):

- a. Outbound Rule for Port 34572:

```
Remove-NetFirewallRule -DisplayName "Maxview Outbound Rule on TCP port 34572 for
SNMP Service" -PolicyStore "ConfigurableServiceStore"
```

b. Inbound Rule for Port 34572:

```
Remove-NetFirewallRule -DisplayName "Maxview Inbound Rule on TCP port 34572 for
SNMP Service" -PolicyStore "ConfigurableServiceStore"
```

B.2 Setting Up SNMP Notifications on Linux

1. Install the Net-SNMP RPM packages:

- net-snmp
- libsnmp15
- snmp-mibs

2. In /etc/snmp/snmpd.conf configuration file:

a. Comment out the com2sec entry:

```
# com2sec notConfigUser default public
```

b. Add the following lines at the end of the file:

```
rocommunity public
trapsink localhost
master agentx
```

c. (SLES 10 only) Register the agentx socket:

```
agentxsocket /var/agentx/master
```

3. Copy aus.mib from /usr/StorMan to /usr/share/snmp/mibs/:

```
#cp /usr/StorMan/aus.mib /usr/share/snmp/mibs
```

4. Restart the SNMP agent:

```
#service snmpd restart
```

5. Start aus-snmpd from /usr/StorMan:

```
#./aus-snmpd
```

C Using the maxView Plugin for VMware vSphere Web Client

The maxView plugin for VMware vSphere Web Client is a monitoring tool that lets you explore your storage resources directly from the vSphere Web client, without using maxView Storage Manager as a separate Web GUI. It replicates most of the information on the maxView Storage Manager Dashboard (see [The Storage Dashboard](#)) for the controllers, logical devices, physical devices, enclosures, and systems (hosts) in your storage space. It offers a vSphere "look and feel" and native navigation.

Note: The maxView plugin is a monitoring tool only. To modify your storage space, such as creating or deleting a logical drive or designating a hot spare, you must use the maxView Storage Manager GUI or CLI (ARCCONF).

C.1 Installing the maxView Plugin for vSphere Web Client

Follow the instructions in this section to install the maxView plugin for vSphere Web Client on a Windows x64 system.

Note: Before you begin, ensure that VMware vSphere Web Client is already installed on the Windows host. You need administrator privileges to install the maxView plugin.


1. Download the setup program for the maxView vSphere plugin from the Web site (see [Download the Installation Package](#)).

2. Double-click the setup program:

Option	Description
VMware ESXi 6.0 client	setup_maxView_VWC6_Plugin_x64.exe
VMware ESXi 6.5 client	setup_maxView_VWC65_Plugin_x64.exe

The Installation wizard opens.

3. Click **Next** to begin the installation, accept the terms in the license agreement, then click **Next**.
4. Enter the following configuration settings for your VMware ESXi server, then click **Next**:
 - a. Hypervisor IP address
 - b. Hypervisor user name
 - c. Hypervisor password



- In the Features screen, ensure that **maxView VWC Plugin** is selected. Optionally, select **CLI Tools**. Then click **Next**.




- Click **Install** to begin the installation.

When the installation is complete, you receive a confirmation message that the plugin is installed.

C.2 Starting the maxView Plugin for vSphere Web Client

- Launch the VMware vSphere Web Client: double-click the vSphere desktop icon, then enter your login credentials.
- In the Monitoring section on the vSphere Home screen (on the right), click the **maxView** icon; the Controller information screen opens.



- Alternatively, in the Navigation pane (on the left), click **vCenter**, scroll down to the maxView tree, then select a resource category, such as Controller, Logical Device, Physical Device, or Enclosure, to open the information screen for that item.



C.3 Monitoring maxView Resources in vSphere Web Client

For each maxView resource in your storage space—controller, logical device, physical device, and so on—you can view summary information about the resource (or "object") and view its related resources, such as the physical devices in a logical drive, the logical drives on a controller, or the controllers on a host.

For example, the figure below shows the summary information screen for a logical drive, with expandable information fields containing general information about the logical drive, its settings, and resources. (Each expandable field has an arrow next to it.)

The screenshot shows the summary page for a logical device named 'con2-log-1'. The page includes a summary section with a cylinder icon, a table of properties, and three expandable sections: Logical Device Information, Settings, and Resources.

con2-log-1	
ID:	0
RAID level:	1
Status:	Optimal
Date created:	04/07/2014
Interface type:	Serial Attached SCSI


Logical Device Information	
Data space	930.99 GB
Parity space	931 GB
Stripe size	Not Applicable
Partitioned	No
Bootable	Yes
Block size	512
Protected by hot spare	No
Failed stripes	No

Settings	
Power management	Disabled
Controller cache	On
maxCache	Not Supported


Resources	
Name	Status
Slot 1	Optimal
Slot 0	Optimal

Click on the **Related Objects** tab (next to the Summary tab) to show the physical devices comprising the logical drive, the logical drive's controller, or a filtered list of events for that logical drive.


Button bar provides quick access to related object information



To drill down further, click on any item in the related objects table, below the button bar. For example, click on a physical device in the table shown below to view summary information for that device, its related objects, and so on.



When you're ready to monitor a different resource, click the resource category in the Navigation pane (on the left). Or click on the History window one or more times (in the upper-left corner of the vSphere client) to return to the vCenter top-level screen; then select a resource in the maxView tree.



Once you return to the top-level screen or choose a resource in the Navigation pane, the procedure for viewing the resource summary and drilling down for more detail is identical for all maxView resources in your storage space.

For more information about the types of information provided on the summary screen and related objects screens for each maxView resource, see [Viewing Component Status in the Storage Dashboard](#).

D Using the maxView Plugin for VMware vSphere HTML5

The maxView plugin for VMware vSphere Web Client is a monitoring tool that lets you explore your storage resources directly from the vSphere HTML client, without using maxView Storage Manager as a separate Web GUI.

D.1 Installing the maxView Plugin for vSphere HTML5 Client

Follow the instructions in this section to install the maxView vSphere Plugin HTML5 on a vCenter Server Appliance.

Note: You need to have VMware ESXi 7 and above and vCenter Server Appliance (VCSA) 7 and above versions running on the machine before proceeding with the installation steps.

1. Create an ESXi setup and mount the VCSA .iso file on a machine.
2. Navigate to the following directory on the mounted file:
`\vcsa-ui-installer\win32\installer.exe`
3. Double-click on the installer and follow the instructions to create a virtual machine in the ESXi machine with the VMWare Photon OS.
4. Perform either of the following two steps to install the maxView vSphere plugin:
 - a. Download and install the maxView vSphere plugin from the Web site.
 - b. Install the `maxView-plugin.zip` locally by copying the downloaded `maxView-plugin.zip` in `/etc/vmware-vpx/locale/ maxView-plugin.zip`.
5. Execute the following steps to install the maxView vSphere plugin.
 - a. Extract the `vsphere-client-sdk-7.zip` provided by the VMWare.
 - b. Navigate to `\html-client-sdk\tools\vCenter plugin registration\prebuilt\location` that contains the plugin registration script.
 - c. Run the following command to register the plugin from the Windows machine:
 - Using Microsemi Web URL:

```
extension-registration.bat -action
    registerPlugin --key com.pmc.maxview.maxviewvsphere -url
    https://IP_ADDRESS/sdk --username Username --version 1.0.0

    --vcenterServerThumbprint vcenterServerThumbprint
    --pluginUrl microsemiwebsite -serverThumbprint
    serverThumbprint
```

- Installing the `maxView-plugin.zip` locally:

```
extension-registration.bat -action
    registerPlugin --key com.pmc.maxview.maxviewvsphere -url
    https://IP_ADDRESS/sdk --username Username --version 1.0.0

    --vcenterServerThumbprint vcenterServerThumbprint
    --pluginUrl https:// IP_ADDRESS
    /catalog/maxView-plugin.zip -serverThumbprint
    serverThumbprint
```

Note: For more detail on the commands, check with VMWare.

where:

- `IP_Address` is the IP address of the vCenter Server Appliance
- `Username` is the username of the vCenter Server Appliance
- `vcenterServerThumbprint` is the vCenter Server Appliance SHA1 Thumbprint. To retrieve SHA1 Thumbprint, perform the following steps:
 - i. Launch the vCenter Server appliance in the Web browser.
 - ii. Click the **Lock** icon on the address bar to get the certificate.
 - iii. View the SHA1 Thumbprint on the the certificate.
- `microsemiwebsite` is Microsemi Web URL, where the maxView plugin is placed. Or, the IP address of the machine where `maxView_plugin.zip` is placed.
- `serverThumbprint` is required only if Microsemi website is not fully secured. Otherwise, ignore this option.

A password is prompted. Enter the password of the vCenter Server Appliance.

- d. Run the following command to de-register the plugin from Windows machine.

```
extension-registration.bat -action unregisterPlugin
--key com.pmc.maxview.maxviewvsphere -url https://IP_ADDRESS /sdk
--username Username --vcenterServerThumbprint
vcenterServerThumbprint
```

Note: For more detail on the commands, check with VMWare.

where:

- `IP_Address` is the IP address of the vCenter Server Appliance
- `Username` is the username of the vCenter Server Appliance
- `vcenterServerThumbprint` is the vCenter Server Appliance SHA1 Thumbprint. Retrieve the SHA1 Thumbprint certificate as described in the preceding step.

A password is prompted. Enter the password of vCenter Server Appliance.

Once the maxView vSphere plugin is registered, the maxView icon is displayed in shortcut of vCenter Server Appliance (see https://VCSA_IP/ui/app/shortcuts).


6. Click on the maxView icon to list the ESXi systems added.

Note: Initially, none of the system are added.
7. Click on **Add system** to add the ESXi systems.
8. Add the firewall entry by using URL https://VCSA_IP: 5480. Click on the firewall and add the entry.


Note: The configuration or events are not delivered, if the firewall is not added.

D.2 Starting the maxView Plugin for vSphere HTML Client


1. Launch the VMware vSphere Web Client and enter your login credentials.
2. From the menu, select the **Shortcuts** option.
3. In the Monitoring section on the vSphere client's Shortcut page screen, click the **maxView** icon; the Host information screen opens.



4. The maxView plugin loads and displays the Host list page.



5. Click on the **Add Host** icon to add and manage the ESXi host. The **Add Host** dialog box opens.


 The screenshot shows the 'Add Host' dialog box. It has a title bar with 'Add Host' and a close button. The dialog contains three input fields: 'Hostname or IP Address' with the placeholder text 'Hostname/IP Address', 'Host Username' with the placeholder text 'Username', and 'Host Password' with the placeholder text 'Password' and a toggle for password visibility. At the bottom, there are two buttons: 'CANCEL' and 'ADD'.

6. Enter the following credentials:
- Hostname or IP address of the ESXi Server
 - Username and password of the ESXi Server
7. Click the **Add** button.
The status dialog is displayed with success or failure status.
8. Click **OK**.
9. Click on **Global Refresh** icon to display the new added system in the list.




Note: Before adding the system, make sure to install the `arcconf` and `arc-cim-provider` vib on the ESXi server.

D.3 Monitoring maxView Resources in vSphere HTML Client

For each maxView resource in your storage space—controller, logical device, physical device, and so on—you can view summary information about the resource (or "object") and view its related resources, such as the physical devices in a logical drive, the logical drives on a controller, or the controllers on a host.

For example, the following figure shows the added ESXi server.



Click on the system IP address to navigate to the two tabs listed in the following table.




Table 1 • Host Tab Details

Tab	Details
Getting Started	Provide details on the ESXi server
More Objects	List the controller details

D.3.1 maxView vSphere Plugin Resources


The following section describes the controller details of the maxView vSphere plugin.

The controller has the following four tabs:

- Getting Started
- Summary
- Device Alert
- More Objects (VMware specific)


Getting Started Tab

The following figure shows the **Getting Started** tab that displays the basic information of a controller.




Summary Tab

The **Summary** tab displays the status and the properties of the controller.




Device Alert Tab

The **Device Alert** tab display the information(s), warning(s), and errors(s) of the controller.



More Objects Tab

The **More Objects** tab is a VMware specific tab, which displays the related resources of the controller.



Note:

You can also obtain the similar details of an Array, Enclosure, Event, Logical Device, and Physical device by clicking on the respective tabs.

E Using maxView Storage Manager with HBAs and Non-RAID Mode Controllers

maxView Storage Manager allows you to manage storage resources on Microsemi Adaptec Host Bus Adapters (HBAs) and RAID controllers operating in HBA mode (see [Changing the Connector Operating Mode](#) on page 101). In the Enterprise View, maxView Storage Manager displays HBAs and non-RAID mode controllers in the controller list. Drives on the HBA are surfaced to the OS as Raw devices; that is, storage devices without Microsemi Adaptec RAID meta-data.

Note: maxView Storage Manager identifies the drive as a Raw device even if it has an OS partition.



With HBAs and non-RAID mode controllers, maxView Storage Manager limits access to features that are not used to configure and maintain RAID volumes (see table below). For example, on the Ribbon, you can use the options in the Controller group to manage your controller, but not options in the Array group or Logical Device group (because HBAs don't support logical volumes); similarly, you can use options in the System group to upgrade the controller firmware, but not the Spare Management option in the Physical Devices group (because HBAs don't support spares); and so on.

Ribbon	Options for HBAs / Non-RAID Mode Controllers
System Group	Firmware Update
Controller Group	Rescan, Properties (non-RAID mode controllers only)
Array Group	None
Logical Device Group	None
Physical Devices Group	Force Offline, Uninitialize, Locate
maxCache Group	None

The Storage Dashboard provides detailed information about the HBAs and non-RAID mode controllers in your storage space (similar to its function for RAID controllers), including the enclosures, disk drives, and SSDs connected to them (for more information about the dashboard, see [Viewing Component Status in the Storage Dashboard](#)).

Tabs on the dashboard provide quick access to summary information, controller properties, resources, and the connector configuration. The Events tab shows filtered events for the device (see [Viewing Activity Status in the Event Log](#) on page 112).

The following table lists the categories and types of information provided on the Storage Dashboard for HBAs and connected devices.

Component	Categories	Examples
Controller	Summary Properties Resources Connectors	Model, WWN, key features, firmware version, controller mode, status, number and type of physical devices. Slot, driver version, bus type and speed, number of ports, settings (mostly disabled) Physical drive assignments by connector, including protocol, state, free and used space Connector name, number of devices, functional mode
Physical Devices (node)	Summary	Physical drive assignments by connector, including protocol, state, free and used space
Connector	Summary	Functional mode, number of devices
Enclosure	Summary	Enclosure type, vendor, model, ID, channel, firmware version, status Fan, power supply, and temperature status (see Monitoring Enclosure Status on page 115) Slot allocation and usage
Hard drives and SS-Ds	Summary Resources SMART	Drive type (hard drive, SSD), vendor, interface (SAS/SATA), and model Block size, total size, rotational speed Boot type Firmware version, WWN, transfer speed Free space, used space, reserved space SMART statistics (see Viewing SMART Statistics on page 115)

F Selecting the Best RAID Level

When you create logical drives in maxView Storage Manager, you can assign a RAID level to protect your data.

Each RAID level offers a unique combination of performance and redundancy. RAID levels also vary by the number of disk drives they support.

This section provides a comparison of all the RAID levels supported by maxView Storage Manager, and provides a basic overview of each to help you select the best level of protection for your storage system.

Note: Not all RAID levels are supported by all controllers. See the Release Notes for supported RAID levels on specific controller models.

F.1 Comparing RAID Levels

Use this table to select the RAID levels that are most appropriate for the logical drives on your storage space, based on the number of available disk drives and your requirements for performance and reliability.

RAID Level	Redundancy	Disk Drive Usage	Read Performance	Write Performance	Built-in Hot Spare	Minimum Disk Drives
RAID 0	No	100%	***	***	No	2
RAID 1	Yes	50%	**	**	No	2
RAID 1(ADM)	Yes	33%	**	**	No	3
RAID 1E	Yes	50%	**	**	No	3
RAID 10	Yes	50%	**	**	No	4
RAID 10(ADM)	Yes	33%	**	**	No	6
RAID 5	Yes	67 – 94%	***	*	No	3
RAID 50	Yes	67 – 94%	***	*	No	6
RAID 6	Yes	50 – 88%	**	*	No	4
RAID 60	Yes	50 – 88%	**	*	No	8


Disk drive usage, read performance, and write performance depend on the number of drives in the logical drive. In general, the more drives, the better the performance.

F.2 Non-redundant Logical Drives (RAID 0)

A logical drive with RAID 0 includes one or more disk drives and provides data *striping*, where data is distributed evenly across the disk drives in equal-sized sections. However, RAID 0 logical drives do not maintain redundant data, so they offer *no data protection*.

Compared to an equal-sized group of independent disks, a RAID 0 logical drives provides improved I/O performance.


Drive segment size is limited to the size of the smallest disk drive in the logical drive. For instance, an array with two 250 GB disk drives and two 400 GB disk drives can create a RAID 0 drive segment of 250 GB, for a total of 1000 GB for the volume, as shown in this figure.



F.3 RAID 1 Logical Drives

A RAID 1 logical drive is built from two disk drives, where one disk drive is a *mirror* of the other (the same data is stored on each disk drive). Compared to independent disk drives, RAID 1 logical drives provide improved performance, with up to twice the read rate and an equal write rate of single disks. However, capacity is only 50 percent of independent disk drives.


If the RAID 1 logical drive is built from different- sized disk drives, the free space, drive segment size is the size of the smaller disk drive, as shown in this figure.



F.4 RAID 1 (ADM) Logical Drives

RAID 1 Advanced Data Mirroring, or ADM, is similar to RAID 1, but creates fault tolerance by maintaining redundant copies of data using three disk drives, rather than two. All three drives contain mirrored duplicated user data.


If a drive fails, the remaining drives provide backup copies of the files and normal system operations are not interrupted.



F.5 RAID 10 Logical Drives

A RAID 10 logical drive is built from two or more equal-sized RAID 1 logical drives. Data in a RAID 10 logical drive is both striped and mirrored. Mirroring provides data protection, and striping improves performance.


Drive segment size is limited to the size of the smallest disk drive in the logical drive. For instance, an array with two 250 GB disk drives and two 400 GB disk drives can create two mirrored drive segments of 250 GB, for a total of 500 GB for the logical drive, as shown in this figure.



F.6 RAID 10 (ADM) Logical Drives

RAID 10 Advanced Data Mirroring, or ADM, is similar to RAID 10, but creates fault tolerance by maintaining redundant copies of data using at least six disk drives. Data is striped across two or more sets of RAID 1 (ADM) drives for rapid access.

If a drive fails, the remaining drives provide backup copies of the files and normal system operations are not interrupted.




F.7 RAID 5 Logical Drives

A RAID 5 logical drive is built from a minimum of three disk drives, and uses data striping and *parity* data to provide redundancy. Parity data provides data protection, and striping improves performance.

Parity data is an error-correcting redundancy that's used to re-create data if a disk drive fails. In RAID 5 logical drives, parity data (represented by Ps in the next figure) is striped evenly across the disk drives with the stored data.

Drive segment size is limited to the size of the smallest disk drive in the logical drive. For instance, an array with two 250 GB disk drives and two 400 GB disk drives can contain 750 GB of stored data and 250 GB of parity data, as shown in this figure.




F.8 RAID 50 Logical Drive

A RAID 50 logical drive is built from six to forty-eight disk drives configured as two or more RAID 5 arrays, and stripes stored data and parity data across all disk drives in both RAID 5 logical drives. (For more information, see [RAID 5 Logical Drives](#).)

The parity data provides data protection, and striping improves performance. RAID 50 logical drives also provide high data transfer speeds.

Drive segment size is limited to the size of the smallest disk drive in the logical drive. For example, three 250 GB disk drives and three 400 GB disk drives comprise two equal-sized RAID 5 logical drives with 500 GB of stored data and 250 GB of parity data. The RAID 50 logical drive can therefore contain 1000 GB (2 x 500 GB) of stored data and 500 GB of parity data.

In this figure, P represents the distributed parity data.




F.9 RAID 6 Logical Drives

A RAID 6 logical drive—also known as dual drive failure protection—is similar to a RAID 5 logical drive because it uses data striping and parity data to provide redundancy. However, RAID 6 logical drives include *two* independent sets of parity data instead of one. Both sets of parity data are striped separately across all disk drives in the logical drive.

RAID 6 logical drives provide extra protection for your data because they can recover from two simultaneous disk drive failures. However, the extra parity calculation slows performance (compared to RAID 5 logical drives).

RAID 6 logical drives must be built from at least four disk drives. Maximum stripe size depends on the number of disk drives in the logical drive.



F.10 RAID 60 Logical Drives

Similar to a RAID 50 logical drive (see [RAID 50 Logical Drive](#)), a RAID 60 logical drive—also known as dual drive failure protection—is built from eight disk drives configured as two or more RAID 6 logical drives, and stripes stored data and two sets of parity data across all disk drives in both RAID 6 logical drives.

Two sets of parity data provide enhanced data protection, and striping improves performance. RAID 60 logical drives also provide high data transfer speeds.

G Icons At-a-Glance

The following is a complete list of icons used in maxView Storage Manager. It contains the icons on the ribbon, in the Enterprise View, and on tabs and dialog boxes.

See [Overview of the Main Window](#) for more information.

Ribbon Home Icons

Icon	Description
	Remote system add
	Remote system delete

Ribbon System Icons

Icon	Description
	System settings
	Manage configuration (save/restore)
	Firmware update
	Save archive file

Ribbon Controller Icons

Icon	Description
	Controller settings
	Manage configuration (clear)
	Controller rescan
	maxCrypto settings

Ribbon Array Icons

Icon	Description
	Array settings
	Array modify

Icon	Description
	Array split/mirror
	Array locate
	Array delete

Ribbon Logical Device Icons

Icon	Description
	Logical drive settings
	Logical drive create
	Logical drive expand/migrate
	Logical drive locate
	Logical drive delete
	Logical drive force online

Ribbon Physical Device Icons

Icon	Description
	Physical disk properties
	Assign/unassign physical disk as spare
	Force physical disk offline
	Physical disk secure erase
	Physical disk locate

Ribbon maxCache Icons

Icon	Description
	maxCache Device create
	maxCache Device set properties
	maxCache Device locate

Icon	Description
	maxCache Device delete
	maxCache Device disable

Enterprise View Icons

Icon	Description
	Enterprise View
	Local or remote system
	Controller
	Enclosure
	Logical disk
	Logical disks
	Physical disk
	Hard disk drive
	Solid State drive
	SMR drive
	Physical disks
	Enclosure
	Connector or other physical device

Enterprise View Status Icons

Icon	Description
	Enterprise OK
	Enterprise error
	Enterprise no access

Icon	Description
	Enterprise warning

Enterprise View System Icons

Icon	Description
	System OK
	System error
	System missing
	System no access
	System warning

Enterprise View Connector Icons

Icon	Description
	Connector OK
	Connector failed
	Connector missing
	Connector warning

Enterprise View Controller Icons

Icon	Description
	Controller OK
	Controller failed
	Controller missing
	Controller warning
	Controller encrypted

Enterprise View Enclosure Icons

Icon	Description
	Enclosure Management OK
	Enclosure Management failed
	Enclosure missing
	Enclosure Management warning

Enterprise View Physical Disk Icons

Icon	Description
	Physical disk OK
	Physical disks OK
	Physical disks failure
	Physical disks missing
	Physical disks warning
	Physical disk encrypted

Enterprise View MaxCache Device Icons

Icon	Description
	maxCache Device error
	maxCache Device missing
	maxCache Device OK
	maxCache Device warning

Tab Icons

Icon	Description
	Summary

Icon	Description
	Properties
	Resources
	Events
	Task
	maxCrypto

Dialog Box Icons

Icon	Description
	E-mail notification
	Chart

H Advanced Statistics Descriptions

The following table describes the advanced statistics properties by device type.

See [Viewing Advanced Statistics](#) on page 124 for more information.

Table 2 • Advanced Statistics Descriptions

Type	Property	Description	Units
Controller	Average Dirty Cache Lines	The average number of dirty cache lines that are waiting to be flushed. Dirty cache lines that are in the process of being flushed are not included in this count. Note that lines in the read-ahead cache may also be "dirty," so the number of dirty cache lines may exceed the number of write cache lines.	
Controller	Free Largest Transfer Buffer Sectors	The average number of consecutive transfer buffer sectors available in the controller's internal pool.	
Controller	Average Free Total Transfer Buffer Sectors	The average number of transfer buffer sectors available in the controller's internal pool.	
Controller	Average Free Logical Requests	The average number of logical request structures available in the controller's internal pool.	
Controller	Average Free Processor RAM in kB	The average amount of available processor RAM in kB available in the controller's internal pool.	
Controller	Average Locked Stripes	The average number of stripes that are locked. In this context, a stripe is a group of physical sectors across a RAID group that is busy performing a write operation (other operations to the same sectors are locked out during this process).	
Controller	Average Locked Stripes Waiting	The average number of requests that are stalled waiting to access stripes that are locked. A large value indicates random writes over a small area of a RAID 5 array.	
Controller	Average Write Cache Sectors	The average number of 512-byte blocks that there is room for in the write cache. Controller cache is currently organized in 16 kB "lines," so divide this value by 32 to determine the number of available write cache lines.	
Controller	Command List Count	Number of command lists submitted by the host.	
Controller	Command List Latency	Average command latency in hundredths of milliseconds.	Hundredths of milliseconds
Controller	Logical Request Count	Number of logical requests submitted by the host.	
Controller	Maximum DMA Transfer Queue Depth	The maximum number of DMA transfer (WCXC bus master) requests that have been queued up at any time since performance monitoring was started.	
Controller	Maximum Outstanding Command List	Maximum number of outstanding command lists.	
Controller	Percent Active	Percentage of time (in tenths of a percent) that at least one logical request is outstanding, including event-notification commands.	Tenths of a percent

Type	Property	Description	Units
Controller	Percent Busy	Percentage of time (in percent) that the controller's CPU is not just sitting in the "idle" task. Note that background activity such as surface analysis is not considered "idle" time.	Percent
Controller	Percent Time Waiting For DMA	The percentage of time (in tenths of a percent) that the local CPU has been stalled waiting for DMA (WCXC bus master) transfers.	Tenths of a percent
Controller	Sample Interval	Time since performance monitoring was last cleared, in tenths of seconds.	Tenths of seconds
Logical Drive	Average Queue Depth	Average number of logical read and write requests queued.	
Logical Drive	Average Read Latency	Average latency for logical read requests in hundredths of milliseconds.	Hundredths of milliseconds
Logical Drive	Average Write Latency	Average latency for logical write requests in hundredths of milliseconds.	Hundredths of milliseconds
Logical Drive	Cache Hits	Number of cache hits (host-generated logical read requests which could be serviced entirely by the cache without any disk access).	
Logical Drive	Cache Misses	Number of cache misses.	
Logical Drive	Coalesced Requests	Number of host requests to this volume that were coalesced by firmware.	
Logical Drive	Flush Read Requests	Number of read-fill requests issued by the cache flush task. The flush task generates read requests to make a dirty cache line completely valid so that it can be flushed in a single write request.	
Logical Drive	Flush Write Requests	Number of flush requests issued by the cache flush task.	
Logical Drive	Logical Reads	Number of logical read requests submitted by the host, excluding requests coalesced by firmware.	
Logical Drive	Logical Writes	Number of logical write requests submitted by the host, excluding requests coalesced by firmware.	
Logical Drive	Maximum Read Request Latency	Maximum latency for logical read requests to this volume in hundredths of milliseconds.	
Logical Drive	Maximum Write Request Latency	Maximum latency for logical write requests to this volume in hundredths of milliseconds.	Hundredths of milliseconds
Logical Drive	Non-Sequential Reads	Number of non-consecutive or non-sequential read requests from the host. A non-sequential request is defined as a logical request with a starting logical block that is not equal to 1 plus the ending logical block of the previously received logical request.	
Logical Drive	Non-Sequential Writes	Number of non-consecutive or non-sequential write requests from the host. A non-sequential request is defined as a logical request with a starting logical block that is not equal to 1 plus the ending logical block of the previously received logical request.	
Logical Drive	Posted Writes	Number of host write operations that were absorbed by the posted write cache.	

Type	Property	Description	Units
Logical Drive	Read Ahead Sector Count	Number of sectors pre-fetched.	
Logical Drive	Sectors Flushed	Total number of sectors in all flush requests issued by the cache flush task.	
Logical Drive	Sectors Read	Number of logical blocks read from this volume by the host. Does not include any internally generated read requests (such as read-ahead, capacity expansion, etc.).	
Logical Drive	Sectors Written	Number of logical blocks written to this volume by the host. Does not include any internally generated write requests (such as read-ahead, capacity expansion, etc.).	
Logical Drive	Unaligned Reads	Number of unaligned read requests from the host. An unaligned logical request is defined as one that spans multiple physical drives but has a block count smaller than or equal to the distribution factor.	
Logical Drive	Unaligned Writes	Number of unaligned write requests from the host. An unaligned logical request is defined as one that spans multiple physical drives but has a block count smaller than or equal to the distribution factor.	
Physical Drive	Average Request Latency	Average latency for physical read/write requests in hundredths of milliseconds.	Hundredths of milliseconds
Physical Drive	Maximum Queue Depth	Maximum number of physical (SCSI) requests in the queue for this drive, including both those queued up on the drive and those on the array controller's elevator queue.	
Physical Drive	Maximum Request Latency	Maximum latency for physical read/write requests in hundredths of milliseconds.	Hundredths of milliseconds
Physical Drive	Maximum Wait Time Between Completions	Maximum amount of time between completion of outstanding requests (when multiple requests are outstanding), in hundredths of milliseconds. A high value may indicate that the drive is experiencing errors.	Hundredths of milliseconds
Physical Drive	Read Requests	Number of physical (SCSI) read requests generated from host and cache read-ahead logical requests (not counting internal requests such as surface analysis, rebuild, expand, snapshot, M&P, etc.).	
Physical Drive	Write Requests	Number of physical (SCSI) write requests generated from host and cache-flush logical requests (not counting internal requests such as surface analysis, rebuild, expand, snapshot, M&P, RIS saves, etc.).	
Physical Drive	Aborted Commands	The number of times a drive was failed due to aborted commands that could not be retried successfully.	
Physical Drive	Medium Errors	The number of times a drive was failed due to unrecoverable media errors.	
Physical Drive	Parity Errors	The total number of times that write memory errors were detected on transfers to/from cache memory from other components. This includes parity errors, ECC corrected errors, and ECC un-corrected errors, depending upon the controller.	
Physical Drive	Hardware Errors	The number of times a drive returned a bad hardware status.	

I Smart Controller Device Status

The following is a complete list of the Smart Controller devices, their status, and their descriptions used in maxView Storage Manager.

Status Details of a Controller

Status	Description
Optimal	Controller is healthy.
Controller Has Incompatible Driver	The controller has the incompatible driver.
Failed	Controller is not in working condition.
Inaccessible	Controller communication failure error.
Down/Offline	Controller offline error.
Controller Lockup Error	Controller failed with a lockup error.
Missing SG Module	Controller missing module error. This controller requires that the scsi_generic (sg) module be loaded in order to be configured.
Controller Powered OFF	Controller poweroff error.
Sys PQI Driver Conflict	This controller has an incompatible driver.
Unknown	Controller unknown error.

Status Details of an Array

Status	Description
Ok	The array has all the logical device(s) in optimal state
Has Initializing Logical Device	One or more member logical device(s) has RPI in progress
Logical Devices Not Contiguous	The logical devices in this array are not in contiguous order. Perform consolidate space operation to consolidate all the free space to the end of the array.
Has Failed Physical Device	Array has a bad or missing physical device
Has Failed Logical Device	One or more logical device(s) in the array has failed
Failed	A physical device or logical device in the array has failed.
Has Erasing Drive	The array currently has a drive erase operation queued, running, stopped or completed on a logical or physical device.
Spare Drive Size Too Small	The array has a spare drive assigned which is smaller than the smallest data drive in the array.

Status Details of a Logical Device

Status	Description
Optimal	The logical device is healthy and is readily accessible by the host
Queued for Expansion	The logical device is queued for expansion
Expanding	The logical device is undergoing online capacity expansion
Ready for Recovery	The logical device is queued to be recovered from a failed physical device
Recovering	The logical device is rebuilding a physical device from fault tolerant data
Wrong Drive Replaced	A wrong physical device was replaced
RPI In Progress	Rapid parity initialization is currently in progress on this logical device
RPI Queued	Rapid parity initialization is currently queued on this logical device. It will start once other progress tasks are completed.
Unsupported on The Controller	Logical device is unsupported on this controller. Host access to this volume is denied. Logical device can still be deleted/reconfigured with data loss.
Encrypted Logical Device Without Key	The encrypted logical device is exported from a foreign controller with different master key. Please import the foreign master key to access the logical device.
Encryption Migration	The logical device is being migrated between plaintext and ciphertext
Encrypted Logical Device Rekeying	The logical device is encrypted and all data is being re-keyed using the background 'online capacity expansion' transformation task. The cache memory is being used to keep track of progress.
Encrypted Logical Device With maxCrypto Off	The logical device is encrypted, exported from a foreign controller and cannot be accessed as the controller does not have encryption enabled/not configured.
Encryption Migration Requested	The logical device has received a request to migrate from plaintext to ciphertext. But this process has not yet started. The plaintext volume is currently online.
Encrypted Logical Device Rekey Requested	The logical device is encrypted and has received a request to re-key all data with a new encryption key
Unknown	The status of logical device is unknown
Erase In Progress	The logical device is offline and has erase in progress
Ejected	The logical device is offline from being ejected. Reinstall the removed physical devices.
Not Yet Available	An expand, shrink, or move operation on the array is in progress. This logical device will remain in this state until all expand, shrink, or move operations on this array are completed. All I/O requests sent to the logical drive in this state will be rejected.
Not Configured	The logical device is not yet configured
Interim Recovery	The logical device has a bad or missing drive. Logical device is operating with reduced performance and a further physical drive failure may result in data loss depending on the fault tolerance. To correct this problem, check the data and power connections to the physical drives or replace the failed drive.
Failed	The logical device has bad or missing physical device(s).

Status	Description
Disabled From SCSI ID Conflict	A conflict with an existing SCSI ID exists. Check all SCSI components to make sure they all have a unique SCSI ID.
Drive Improperly Connected	A physical device is not properly connected.
Hardware Has Overheated	A physical device temperature has crossed the threshold value.
Hardware Is Overheating	A physical device temperature is about to reach the threshold value.
Optimal(Background Parity Initialization)	Logical device is undergoing the Parity initialization in background.
Rapid Parity Initialization	Logical device is undergoing the Rapid Parity initialization and may not available until it is completed.
Offline Parity Initialization	Logical device is undergoing Offline parity initialization.
Logical Device Reconfiguring	Logical device is reconfiguring.
Plaintext Volume Rejected In Encrypting Mode	Logical device is plaintext and cannot be accessed as the controller is in encryption-only mode.

Status Details of a Physical Device

Status	Details
Ready	The physical device is readily available for RAID configuration
Optimal	The physical device is part of an array/logical device
Waiting For Rebuild	The physical device is waiting to be rebuilt
Rebuilding	The data on the physical device is being rebuilt. The physical device will be accessible. But performance will be less than optimal during the rebuilding process.
Queued For Erase	The physical device is currently queued for erase and the will not be available for use until the erase operation is completed
Erase In Progress	The physical device is currently being erased and the will not be available for use until the erase operation is completed
Erase Completed	Erase process has been completed on the physical device and the physical device is offline. The physical device may now be brought online through the initialize operation.
Erase Failed	The physical device erase process is failed and the is offline. The physical device may now be brought online through the initialize operation.
Erase Aborted	The physical device is offline due to a aborted erase process
Predictive Failure	This physical device is predicted to fail soon. Backup all the data on the drive and replace the drive.
Transient Data Drive	The physical device is in transition from being a member of an array to being an unassigned physical device as a result of shrink array/move array operation
Failed	The physical device is bad or missing

Status	Details
Failed Due To Predictive Spare Activation	The physical device has been failed by the controller after completing a predictive spare activation
Unsupported	The physical device is not supported by the controller
Not Supported	The controller firmware version does not support this physical device. Replace the physical device with the one supported by the controller.
Dedicated Hot Spare	A dedicated hot spare is assigned to one or more arrays.
Auto Replace Hot Spare	An auto-replace hot spare is assigned to a specific array. After using an auto-replace spare to rebuild a failed logical drive, it becomes a permanent part of the array.
Raw	A physical device is in RAW state which has no or unknown file system.
Size Not Valid	Physical device size is not valid.

J Display Properties of a Controller, Array, Logical Device, and a Physical Device

This section lists the display properties of a controller, array, logical device, and physical device.

Table 3 • Controller Display Properties

Property	Tooltip Details
Model	Model of the Controller
Status	Overall status of the controller based on its resources.
Serial Number	A unique number assigned to the controller, used for identification and inventory purposes.
WWN	A World Wide Name(WWN) is a unique identifier of the controller.
Firmware Version	Active firmware version of the controller
Hardware Revision	Describes the hardware revision information about the controller.
Manufacturing Part Number	Describes the hardware part number information about the controller.
Manufacturing Spare Part Number	Describes the hardware spare part number information about the controller.
Manufacturing Wellness Log	Describes the hardware wellness log information about the controller.
Installed Memory	Size of Dynamic Random Access Memory(DRAM) installed on the controller
Cache Memory	Cache memory size on controller.
Mode	Mode of the controller on which it is operating
Pending Mode	Pending mode of the controller which will reflect on reboot.
Temperature	Current temperature of the controller
Power Consumption	Power Consumption
NVRAM Checksum Status	NVRAM Checksum Status
maxCache	The maxCache software uses a reserved logical drive comprised of SSDs only, called the maxCache device, for fast read and write caching.
maxCrypto	maxCrypto feature ensures the sensitive data is encrypted and protected by secure 256 bit AES, in-line encryption.
Physical Slot	PCI slot number to which the controller is connected.
Driver Version	Current version of driver installed on the system.
Driver Name	Driver name describes the name of the driver.
Host Bus Type	Type of host expansion bus standard.

Property	Tooltip Details
Host Bus Speed	Speed of host expansion bus standard in MHz.
Host Bus Link Width	Speed of host expansion bus standard in MHz.
PCI Address (Bus:Device:Function)	PCI address describes the PCI address for the controller.
I2C Address	I2C address describes the Inter-Integrated Circuit(I2C) slave address.
I2C Clock Speed	I2C clock speed describes the Inter-Integrated Circuit(I2C) clock speed.
I2C Clock Stretching Status	I2C stretching status describes the Inter-Integrated Circuit(I2C) clock status.
NCQ	Native Command Queuing, or NCQ, lets SATA disk drives arrange commands into the most efficient order for optimum performance.
Number of Ports	Number of ports describes number of internal and external ports of the controller.
Backup Power Status	Status of battery backup unit.
Battery / Capacitor Pack Count	Number of battery pack connected to controller
Hardware Error	Hardware error type occurred on battery backup unit
Post Prompt Timeout	Post prompt timeout describes the F1/F2 POST prompt timeout for the controller during system boot.
Rebuild Priority	Rebuild priority determines the urgency with which the controller treats an internal command to rebuild a failed logical drive. At the low setting, normal system operations take priority over a rebuild. At the medium setting, rebuilding occurs for half of the time, and normal system operations occur for the rest of the time. At the medium high setting, rebuilding is given a higher priority over normal system operations. At the high setting, the rebuild takes precedence over all other system operations.
Expand Priority	Expand Priority setting determines the urgency with which the controller treats an internal command to expand an array. At the low setting level, normal system operations take priority over an array expansion. At the medium setting, expansion occurs for half of the time, and normal system operations occur for the rest of the time. At the high setting, the expansion takes precedence over all other system operations.
Advanced Statistics	Use the Statistics Viewer to view the advanced statistics for the controllers, hard drives, SSDs, logical drives, and maxCache devices in your storage space, including virtual volumes and the maxCache Container.
Consistency Check Priority	Consistency Check Mode is an automatic background process that ensures that you can recover data if a drive failure occurs. The scanning process checks physical drives in fault-tolerant logical drives for bad sectors and it also verifies the consistency of parity data if applicable. The available modes are disable, high, or idle. The idle mode must also specify a delay value.
Consistency Check Delay	Consistency Check Delay determines the time interval for which a controller must be inactive before a consistency check is started on the physical drives that are connected to it. The value can be between 0 and 30 to specify the duration of the delay in seconds. A value of 0 disables the scan.
Parallel Consistency Check Count	Parallel consistency check count describes the number of logical devices on which the controller will perform consistency check in parallel.

Property	Tooltip Details
Raid 6/60 Alternate Inconsistency Repair Policy	RAID 6/60 alternate inconsistency repair policy searches for a single inconsistent strip and repairs the strip on that one drive only.
Consistency Check Inconsistency Notify	Consistency Check Inconsistency Notify property enables the event notification messages and serial debug log messages for mirrored volumes.
Spare Activation Mode	Spare activation mode feature enables the controller firmware to activate a spare drive. The firmware starts rebuilding a spare drive only when a data drive fails when the mode is Failure. With the predictive failure activation mode, rebuilding can begin before the drive fails when a data drive reports a predictive failure (SMART) status which will reduce the likelihood of data loss that could occur if an additional drive fails.
Maximum Drive Request Queue Depth	Queue Depth controls the behavior of the cache write queue. This option is used to tune controller performance for video applications. The valid values are 2, 4, 8, 16, 32, or Automatic.
Monitor and Performance Delay	MNP Delay option controls the behavior of the controller Monitor and Performance Analysis Delay and is expressed in values ranging from 0 to 1440. This option is primarily used to tune controller performance for video applications.
Physical Drive Request Elevator Sort	Elevator Sort option controls the behavior of the controller cache write Elevator sort algorithm. This option is used to tune controller performance for video applications. The possible options are Enable or Disable.
Degraded Mode Performance Optimization	Degraded Mode Performance Optimization setting applies to RAID 5 logical devices in degraded mode only. Enabling this setting directs the controller to attempt to improve performance of large read requests by buffering physical drive requests. Disabling this feature forces the controller to read from the same drives multiple times. This option is used to tune controller performance for video applications. The possible options are Enable or Disable.
HDD Flexible Latency Optimization	Latency describes Flexible Latency Schedule (FLS) setting. Flexible Latency Scheduler (FLS) is a controller option where the controller can re-prioritize I/O requests to prevent some requests to HDDs from timing out. Under normal operation (when FLS is disabled, or in controllers that don't support FLS), the controller will sort incoming requests in order to minimize the amount of travel for the HDD's read heads (Elevator Sort). This strategy works well for workloads that access sequential data, or workloads that require multiple requests from localized sectors in the drive. For highly random workloads, such as transaction processing, some requests will end up on the wrong side of the disk platter and, due to their high latency, will be marked as timed out. When FLS is enabled, it will detect these high-latency requests and apply a cutoff value, after which it will suspend elevator sorting and service the request right away.
Primary Boot Volume	Primary Boot Volume describes which logical device or physical device is the primary boot volume on the current controller.
Secondary Boot Volume	Secondary Boot Volume describes which logical device or physical device is the secondary boot volume on the current controller.
Sanitize Lock	Set the sanitize lock policy of the controller. This policy will be applied to all SATA physical devices that support the feature. 1. None : No freeze lock or anti-freeze lock commands are sent to any physical device. 2. Freeze : Supported physical devices are freeze locked and sanitize is not allowed. 3. Anti-Freeze : Supported physical devices are anti-freeze locked and freezing the physical devices is not allowed.

Property	Tooltip Details
Pending Sanitize Lock	Sanitize lock is in pending state, reboot the system and require all physical devices to be power cycled or hot-plugged for the lock state to be applied to the physical devices.
Expander Minimum Scan Duration	Controller waits for the specified seconds to scans/discover the drives attached to the expander on the next power cycle. Set this to a non-zero value if some devices do not appear in the topology after controller boot or rescan requests.
Intelligent Power Management	
Current Power Mode	Power mode setting determines controller static settings based on work load. MINIMUM POWER: Set static settings to lowest possible. Reduce dynamically based on workload. MAXIMUM PERFORMANCE : Set static settings to highest possible. Do not reduce dynamically.
Pending Power Mode	Power mode setting determines controller static settings based on work load before system reboot. MINIMUM POWER: Set static settings to lowest possible. Reduce dynamically based on workload. MAXIMUM PERFORMANCE : Set static settings to highest possible. Do not reduce dynamically.
Survival Mode	
Controller Cache	
Cache Status	Determines the preservation status of the cache module.
Cache Ratio	The controller cache ratio setting determines the controller ability to adjust the amount of memory for read-ahead cache versus write cache.
Write Cache Bypass Threshold (KB)	All writes larger than the specified value will bypass the write cache and be written directly to the disk for non-parity RAID volumes.
No-Battery Write Cache	No-Battery Write Cache setting decides the controller to enable write cache when no battery is present or when the battery fails. Values are Enable or Disable.
Wait for Cache Room	Wait for Cache Room setting causes the controller to always wait for room in the read/write cache when full instead of automatically bypassing it in favor of higher performance.
maxCache	
Status	maxCache support status on controller.
Version	maxCache version
Drive Cache	
Write Cache Policy for Configured Drives	This option allows to configure the write cache policy on a controller. Setting to default allows the controller to optimize the drive write cache policy of those drives. Enabling drive write cache can increase write performance but risks losing the data in the cache on sudden power loss.
Write Cache Policy for Unconfigured Drives	This option allows to configure the write cache policy on a controller. Setting to default for unconfigured drives uses the drive's existing write cache policy. Enabling drive write cache can increase write performance but risks losing the data in the cache on sudden power loss.
Write Cache Policy for HBA Drives	This option allows to configure the write cache policy on a controller. Setting to default uses the drive's existing write cache policy. Enabling drive write cache

Property	Tooltip Details
	can increase write performance but risks losing the data in the cache on sudden power loss.

Table 4 • Array Display Properties

Property	Tooltip Details
ID	ID describes unique array identifier within the controller.
Name	Name describes unique name of array
Status	Status of array is based on health of member disk drives.
Device Type	Type describes the type of the array such as data array, backup array etc.
Interface Type	Disk drives which are the member of array can have interface type such SAS, SATA, SAS SSD and SATA SSD. The interface type of array is based on the member disk drives interface type.
Total Size	Total usable size is the total space available in the array for creating logical device.
Used Size	The total disk space used by the logical device(s) on the given array.
Unused Size	Unused size is the free space available to create new logical device to store the data.
Member Device(s) Block Size	Block size indicates the maximum size of data block on disk drives which are member of array (can be 512 Bytes or 4K).
Transformation Status	Transformation status indicates whether the array is transforming or not.
Protected by Hot Spare	Protected by Hot Spare indicates whether the array is protected by Hot Spare.
Spare Rebuild Mode	Spare rebuild mode describes the spare type for the array. It can be "dedicated" or "auto replace" if the array is valid.
SSD I/O Bypass	SSD I/O Bypass enables an optimized data path to high performance solid state drives. The optimized path bypasses the controllers RAID processing components and sends I/O directly to the drives.
Member Logical Device(s)	Number of logical device(s) present in the array.
Member Physical Device(s)	Number of physical device(s) used to create the array.
Spare Drive(s)	Number of spare drives associated to this array. If a drive fails in the array, the controller automatically rebuilds the data onto the spare drive.

Table 5 • Logical Device Display Properties

Property	Tooltip Details
ID	Describes unique ID of logical device listed.
RAID Level	RAID level on which the logical device has been created.
Device Type	Drive type indicates the type of logical device like data and etc.

Property	Tooltip Details
Interface Type	Disk drive which are RAID member of logical device can have interface type such SAS or SATA will also reflect as interface type of logical device. A logical drive can be combination of SAS and SATA interface.
Data Space	Data space is where actual data is striped across the disk drives.
Stripe Size	Stripe size is the amount of data (in KB) written to one disk drive, before moving to the next disk drive in the logical device. Stripe size options vary, depending on your controller and RAID level.
Full Stripe Size	Full stripe size refers to the combined size of all the strips across all physical drives, excluding parity-only drives.
Member Device(s) Block Size	Maximum size of data block on disk drives which are RAID member of logical device (can be 512 Bytes or 4K).
Volume Unique Identifier	The logical device unique identifier.
Heads	Heads indicates the pre-defined space set aside for RAID redundant information on a logical device.
Sectors Per Track	Sectors Per Track specifies the number of sectors that are to comprise each track.
Cylinders	Cylinders indicates the set of all of tracks of equal diameter in a logical device.
Status	Status of logical device based on health of RAID members of logical device.
Name	Logical device name can be of maximum 64 characters and it should contain only ASCII characters
Disk Name	Name of the logical disk drive
OS Location	Operating system location of the logical disk drive
Mounted	Mount points describes the Operating system device names of the logical device.
Controller Caching	This option toggles the controller cache preservation state. When enabled, the system preserves the controller's cache to prevent data loss in the event of a system failures like power loss or shutdown
Acceleration Method	Logical Device Acceleration Method indicates whether caching for logical device enabled through controller cache. Defaults to enable.
Boot Type	A bootable logical device is a logical device that the system can attempt to boot from after a system power-on. A controller can have up to two bootable logical device, where one is a primary boot logical device and the other a secondary boot logical device. When the system looks at a controller for a boot logical device, it will first attempt to boot from a primary boot logical device, and if that fails, then it will attempt to boot from a secondary boot logical device.
Protected by Hot Spare	Protected by Hot Spare indicates whether the logical device is protected by Hot Spare.
maxCrypto	

Property	Tooltip Details
Encrypted	Indicate user whether the logical device is encrypted or not.
maxCache	
State	State of the associated maxCache logical device.
Write Cache Policy	The current write cache policy used by the associated maxCache for data write operations. This will indicate whether the posted write operations to this logical device, are transferred via the write cache memory.
Write Cache Policy Preferred	The write cache policy preferred by the user for data write operations to this logical device.
Write Cache Policy Status	The status of current write cache policy used by the associated maxCache.

Table 6 • Physical Device Display Properties

Property	Tooltip Details
Vendor	Physical device manufacturer name.
Model	Product model name of the physical device.
Serial Number	Serial number of physical device.
Interface Type	Interface type supported by the physical device.
Total Size	Total data storage capacity of the physical device.
Block Size	Maximum size of data block on disk drives which are RAID member of logical device (can be 512 Bytes or 4K).
Physical Block Size	Physical block size is the unit of data that can be physically read or write to the disk.
Rotational Speed	Indicates the rotational speed of the physical device.
Device Type	Type of physical device such as "hard disk drive", "solid state drive" or "shingled magnetic recording hard disk drive".
Firmware Level	Firmware version of the physical device.
WWN	Reported world wide name provided by manufacturer.
Unique ID	ID to uniquely identify the physical device.
Reported Channel	The channel to which the physical device is connected.
Reported SCSI Device ID	The SCSI ID for a physical device reported by controller.
NCQ Supported	Specifies whether this physical device supports native command queuing.
NCQ Status	Indicates whether the native command queuing is enabled/disabled on this physical device.
Sanitize Erase	Specifies whether the sanitize erase is supported by this physical device.

Property	Tooltip Details
Sanitize Lock Freeze	Specifies whether the sanitize lock freeze is supported by this physical device.
Sanitize Lock Anti-Freeze	Specifies whether the sanitize lock anti-freeze is supported by this physical device.
Encryption Capability	A SED (or Self-Encrypting Drive) is a type of hard drive that automatically and continuously encrypts the data on the drive without any user interaction.
State	Current state of physical device based on the operations done on it.
Negotiated Transfer Speed	Negotiated data transfer rate of selected physical device.
Configuration Type	Determines the presence/type of logical devices of which this physical device is a part of.
Encryption Status	Current status of SED (or Self-Encrypting Drive).
Boot Type	A bootable physical device is a physical device that the system can attempt to boot from after a system power-on. A controller can have up to two bootable physical device, where one is a primary boot physical device and the other a secondary boot physical device. When the system looks at a controller for a boot physical device, it will first attempt to boot from a primary boot physical device, and if that fails, then it will attempt to boot from a secondary boot physical device.
Exposed to OS	Indicates whether the physical device is exposed to the operating system.
Disk Name	Name of the physical disk drive
OS Location	Operating system location of the physical disk drive
Partitioned	Partition describes that the physical device is exposed to operating system or not. The drive must be partitioned and formatted for storing data.
Mounted	Mount point(s) describes the operating system device names of the physical device.
Has Stale RIS Data	Specifies whether the physical device has stale RIS data.
S.M.A.R.T. Error	Any SMART error reported on physical device.
Current Temperature	Current temperature of the physical device.
Maximum Temperature	The maximum temperature reported by the physical device.
Threshold Temperature	The threshold temperature value of the physical device.
Encrypted Drive	Indicates whether this physical device is a part of any encrypted logical device.
Negotiated Physical Link Rate	Negotiated Physical Link Rate
Negotiated Logical Link Rate	Negotiated Logical Link Rate
Maximum Link Rate	Maximum Link Rate

K maxView Video Tutorials

The following table provides the list of maxView video tutorials, their descriptions, and the web links to access them online.

Table 7 • maxView Video Tutorials

S.No	Topic	Details	Links
1	maxView Remote System Management for Adaptec® RAID Adapters and Host Bus Adapters	This video shows how to remotely view, monitor, and configure all Adaptec® SAS/SATA RAID adapters and Host Bus Adapters (HBAs) in your storage infrastructure with Adaptec® maxView Storage Manager.	https://youtu.be/Wzcz9KuL0U
2	How to flash firmware using Adaptec® maxView Storage Manager	This video provides step-by-step instructions to flash firmware on Adaptec® SAS/SATA RAID adapters and Host Bus Adapters (HBAs) using maxView Storage Manager.	https://www.youtube.com/watch?v=9xPTgZASiSc
3	Export and Import Remote Systems Using maxView Storage Manager	This video provides instructions on how to import and / or export remote systems from your network using the Adaptec® maxView Storage Manager user interface	https://www.youtube.com/watch?v=Nu-1I4eidIM&feature=youtu.be&ab_channel=MicrochipTechnology
4	How to Diagnose Errors using Adaptec® maxView Storage Manager	This video allows users to quickly identify, isolate, review and fix errors using an intuitive, browser-based software application. maxView Storage Manager can be used locally or remotely across storage platforms featuring Adaptec® storage adapters and popular third-party vendor solutions that have integrated remote management.	https://www.youtube.com/watch?v=2ff13Lm5LEE

Index

A

activity in storage space [20](#)
 event log [20](#)
 task log [20](#)
 activity on storage space [111](#)
 adapters [2](#)
 Add System wizard [131](#)
 archive file [145](#)
 arrays [68](#)
 deleting [68](#)
 auto-discovery [130](#)
 auto-replace hot spares [40](#)

B

background consistency check [54](#)
 backplane discovery protocol [102](#)
 boards [2](#)
 bootable USB image [13](#)
 build method [29, 33](#)
 Build [29, 33](#)
 Default [29, 33](#)
 building storage space [27, 30, 70](#)
 adding maxCache SSD caching [70](#)
 on existing array [30](#)
 on new array [27](#)

C

cache optimizations [55](#)
 cards [2](#)
 components [2, 88](#)
 defined [2](#)
 viewing properties [88](#)
 Configuration Event Detail window [112](#)
 configuring [26–27, 30](#)
 See also building storage space
 connector modes [101](#)
 controller [21, 137](#)
 clearing configuration [137](#)
 controller connector modes [101](#)
 controller performance optimization [100](#)
 controllers [79, 88, 99, 104](#)
 maxCrypto properties [79](#)
 properties [88](#)
 rescanning [99](#)
 updating firmware [104](#)

D

dedicated hot spares [40–41, 44](#)
 creating [41](#)
 removing [44](#)
 deleting logical drives [68](#)

direct attached storage [1](#)
 defined [1](#)
 discovery protocol [102](#)
 disk drive [108](#)
 updating firmware [108](#)
 disk drives [28, 52–53, 88, 94–97, 104, 141–142](#)
 conceptual graphic [52](#)
 different sizes in same logical drive [53](#)
 failed state [95](#)
 failure recovery [141–142](#)
 multiple disk drives [142](#)
 multiple logical drives [142](#)
 RAID 0 logical drives [142](#)
 with hot spare [141](#)
 without hot spare [142](#)
 identifying [88](#)
 initializing [97](#)
 locating [88](#)
 properties [88](#)
 replacing in logical drives [94](#)
 secure erasing [95](#)
 uninitializing [96](#)
 updating firmware [104](#)
 within logical drives [28](#)
 drives [88](#)
 identifying [88](#)
 locating [88](#)

E

email notifications [117, 119–121](#)
 changing "from" address [121](#)
 changing SMTP server [121](#)
 disabling [121](#)
 failed test emails [120](#)
 modifying recipient information [120](#)
 recipient addresses [117](#)
 removing recipients [120](#)
 setup [117](#)
 SMTP server settings [117](#)
 test emails [119](#)
 Email Notifications Manager [117, 119](#)
 opening [117](#)
 test emails [119](#)
 enclosure [109](#)
 updating firmware [109](#)
 enclosures [21, 88, 115](#)
 fan status [115](#)
 identifying [88](#)
 locating [88](#)
 monitoring [115](#)
 power status [115](#)
 status [21](#)
 temperature status [115](#)
 Enterprise View [16, 18](#)
 icons [18](#)
 event log [20, 111](#)

event log (operating system) [125](#)
 events [20](#), [111](#)
 expander [109](#)
 updating firmware [109](#)
 expander minimum scan duration [103](#)
 expander scan duration settings [103](#)
 expanding logical drives [66](#)

F

failed disk drives [142](#)
 multiple logical drives [142](#)
 RAID 0 logical drives [142](#)
 without hot spare [142](#)
 failed logical drive [142](#)
 force online [142](#)
 firmware [104](#), [108–109](#)
 Firmware Update wizard [104](#)
 force online [142](#)
 formatting logical drives [33](#)

G

Green Backup status [114](#)

H

hard disk [2](#)
 hard disk drive [2](#)
 hard drive [2](#)
 help [22](#)
 hot spare activation [46](#)
 hot spares [40–42](#), [44](#)
 auto-replace [42](#)
 creating dedicated [41](#)
 limitations [40](#)
 removing [44](#)

I

icons [18](#), [112](#)
 Enterprise View [18](#)
 event status icons [112](#)
 identifying components [88](#)
 initializing disk drives [97](#)

L

local systems [24](#)
 locating drives [88](#)
 logging in [15](#), [25](#)
 logging into remote systems [25](#)
 logging out [23](#)
 logical drives [26–31](#), [33](#), [52–54](#), [64](#), [66–68](#), [81](#), [94](#), [141](#),
 [144](#), [165–167](#), [169–171](#)
 build method [29](#), [33](#)
 building on existing array [30](#)
 building on new array [27](#)
 changing RAID levels [64](#)
 defined [52](#)

logical drives (*continued*)
 deleting [68](#)
 different-sized disk drives [53](#)
 disabling background consistency check [54](#)
 enabling background consistency check [54](#)
 encrypting [81](#)
 expanding [66](#)
 formatting [33](#)
 hot-swap rebuilding [144](#)
 increasing capacity [66](#)
 non-redundant [165](#)
 options for creating [26](#)
 partitioning [33](#)
 RAID 1 [165](#)
 RAID 1 (ADM) [166](#)
 RAID 10 [167](#)
 RAID 10 (ADM) [167](#)
 RAID 5 [169](#)
 RAID 50 [170](#)
 RAID 6 [171](#)
 RAID 60 [171](#)
 rebuild priority [67](#)
 rebuilding [141](#)
 renaming [68](#)
 replacing disk drives [94](#)
 selecting disk drives [28](#)
 selecting RAID levels [28](#), [31](#)
 Shingled Magnetic Recording device [26](#)
 SMR drive [26](#)
 Device Managed (SMR DM) [26](#)
 Host Aware (SMR HA) [26](#)
 Host Managed (SMR HM) [26](#)
 Perpendicular Magnetic Recording device (PMR) [26](#)

M

main window [16](#), [23](#), [111–113](#)
 event details [112](#)
 event log [111](#)
 Logout button [23](#)
 monitoring activity and status [111](#)
 notification types [111–112](#)
 event log [111](#)
 event status icons [112](#)
 overview [16](#)
 status icons [112](#)
 Storage Dashboard [113](#)
 task details [113](#)
 maxCache [70](#), [72](#), [124](#)
 cache hits and misses [124](#)
 limitations [70](#)
 write cache setting [72](#)
 write caching
 maxCache Device [70–72](#)
 creating [70](#)
 deleting [72](#)
 selecting a RAID level [71](#)
 selecting SSDs [71](#)
 maxCache statistics [124](#)
 maxCrypto [74](#), [76](#), [79–80](#), [83–84](#), [86](#)
 accounts [76](#)

maxCrypto (*continued*)
 checking status [79](#)
 configuration [80](#)
 importing a master key [86](#)
 re-keying a logical drive [83](#)
 resetting [84](#)
 setup [74](#)
 maxView Storage Manager [3–4](#), [16](#), [18–19](#), [25](#), [111](#)
 beginning steps [3](#)
 Enterprise View [16](#)
 event log [111](#)
 icons [18](#)
 introduction [3](#)
 logging into remote systems [25](#)
 main window overview [16](#)
 maxCache devices [16](#)
 Redfish Server [3](#)
 ribbon [19](#)
 Ribbon [16](#)
 system requirements [4](#)
 monitoring [20](#), [88](#), [111–113](#), [117](#), [119–121](#)
 device properties [88](#)
 email notifications [117](#), [119–121](#)
 changing "from" address [121](#)
 changing SMTP server [121](#)
 failed test emails [120](#)
 modifying recipient information [120](#)
 recipient addresses [117](#)
 removing recipients [120](#)
 setup [117](#)
 SMTP server settings [117](#)
 test emails [119](#)
 Email Notifications Manager [117](#)
 event details [112](#)
 event log [20](#), [111](#)
 options [111](#)
 status icons [112](#)
 Storage Dashboard [113](#)
 task details [113](#)
 task log [20](#)

N

network requirements [8](#)
 non-RAID Mode controllers [162](#)
 non-redundant logical drives [165](#)
 notification types [117](#)
 email notifications [117](#)

O

online help [22](#)
 operating system event log [125](#)
 operating systems [4](#)

P

partitioning logical drives [33](#)
 physical device error counters [92](#)
 power management [68](#)

Properties tab [88](#)
 pure HBAs [162](#)

R

RAID [165–167](#), [169–171](#)
 non-redundant logical drives [165](#)
 RAID 0 [165](#)
 RAID 1 [165](#)
 RAID 1 (ADM) [166](#)
 RAID 10 [167](#)
 RAID 10 (ADM) [167](#)
 RAID 5 [169](#)
 RAID 50 [170](#)
 RAID 6 [171](#)
 RAID 60 [171](#)
 RAID levels [28](#), [31](#), [64](#)
 changing [64](#)
 selecting [28](#), [31](#)
 RAW drive [96](#)
 RAW state [96](#)
 reactivating split mirror backup [63](#)
 rebuild priority [67](#)
 rebuilding (defined) [141–142](#)
 rebuilding logical drives [144](#)
 Redfish [3](#)
 introduction [3](#)
 remirroring split mirror backup [63](#)
 remote systems [24](#), [131](#)
 adding manually [131](#)
 removing [131](#)
 renaming logical drives [68](#)
 rescanning controllers [99](#)
 Resources View [21](#)
 ribbon [19](#), [25](#), [64](#), [66](#), [68](#), [99](#), [145](#)
 Add system [25](#)
 Change logical device name [68](#)
 Expand or change logical device [64](#), [66](#)
 overview [19](#)
 Rescan [99](#)
 Save support archive [145](#)
 Ribbon [16](#)

S

secure erase [95](#)
 segments [52](#)
 Self-Monitoring Analysis and Reporting Technology (SMART) [116](#)
 Server Template File [ix](#), [127](#)
 Shingled Magnetic Recording device [40](#)
 silent installation [147](#)
 switches, properties, and values [147](#)
 SMART statistics [116](#)
 SMR drive [40](#)
 Device Managed (SMR DM) [40](#)
 Host Aware (SMR HA) [40](#)
 Host Managed (SMR HM) [40](#)
 Perpendicular Magnetic Recording device (PMR) [40](#)
 SMTP server [121](#)

split mirror backup array [63](#)
SSD IO Bypass [57](#)
statistics counters [125](#)
resetting [125](#)
statistics logging [122](#), [124–125](#)
counters [125](#)
enabling [122](#)
viewing [124](#)
Statistics Viewer [124](#)
status [20–21](#), [111](#)
enclosures [21](#)
event log [20](#)
task log [20](#)
status icons [112](#)
main window [112](#)
Storage Dashboard [88](#), [113](#)
storage space [2](#), [5](#)
examples [5](#)
support archive file [145](#)
Survival mode [68](#)
system requirements [4](#)

systems [24](#), [88](#)
local systems [24](#)
properties [88](#)
remote systems [24](#)

T

task log [20](#)
temperature status [115](#)
terminology [2](#)

U

uninitializing disk drives [96–97](#)
updating firmware [104](#), [108–109](#)

W

write cache [72](#)
write-back [72](#)
write-through [72](#)

**Microsemi**

2355 W. Chandler Blvd.
 Chandler, AZ 85224 USA

Within the USA: +1 (480) 792-7200
 Fax: +1 (480) 792-7277

www.microsemi.com © 2020 Microsemi and its corporate affiliates. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation and its corporate affiliates. All other trademarks and service marks are the property of their respective owners.

Microsemi's product warranty is set forth in Microsemi's Sales Order Terms and Conditions. Information contained in this publication is provided for the sole purpose of designing with and using Microsemi products. Information regarding device applications and the like is provided only for your convenience and may be superseded by updates. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is your responsibility to ensure that your application meets with your specifications. THIS INFORMATION IS PROVIDED "AS IS." MICROSEMI MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION, INCLUDING BUT NOT LIMITED TO ITS CONDITION, QUALITY, PERFORMANCE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL MICROSEMI BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL OR CONSEQUENTIAL LOSS, DAMAGE, COST OR EXPENSE WHATSOEVER RELATED TO THIS INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROSEMI HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROSEMI'S TOTAL LIABILITY ON ALL CLAIMS IN RELATED TO THIS INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, YOU PAID DIRECTLY TO MICROSEMI FOR THIS INFORMATION. Use of Microsemi devices in life support, mission-critical equipment or applications, and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend and indemnify Microsemi from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microsemi intellectual property rights unless otherwise stated.

Microsemi Corporation, a subsidiary of Microchip Technology Inc. (Nasdaq: MCHP), and its corporate affiliates are leading providers of smart, connected and secure embedded control solutions. Their easy-to-use development tools and comprehensive product portfolio enable customers to create optimal designs which reduce risk while lowering total system cost and time to market. These solutions serve more than 120,000 customers across the industrial, automotive, consumer, aerospace and defense, communications and computing markets. Headquartered in Chandler, Arizona, the company offers outstanding technical support along with dependable delivery and quality. Learn more at www.microsemi.com.

PMC-2153109