



SPECTRA LOGIC BLACKPEARL NEARLINE GATEWAY

ADVANCED BUCKET MANAGEMENT GUIDE



SpectraLogic.com

COPYRIGHT

Copyright © 2024 Spectra Logic Corporation. All rights reserved. This item and the information contained herein are the property of Spectra Logic Corporation.

NOTICES

Except as expressly stated herein, Spectra Logic Corporation makes its products and associated documentation on an “AS IS” BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, BOTH OF WHICH ARE EXPRESSLY DISCLAIMED. In no event shall Spectra Logic be liable for any loss of profits, loss of business, loss of use or data, interruption of business, or for indirect, special, incidental or consequential damages of any kind, even if Spectra Logic has been advised of the possibility of such damages arising from any defect or error.

Information furnished in this manual is believed to be accurate and reliable. However, no responsibility is assumed by Spectra Logic for its use. Due to continuing research and development, Spectra Logic may revise this publication from time to time without notice, and reserves the right to change any product specification at any time without notice.

TRADEMARKS

Attack Hardened, BlackPearl, BlueScale, RioBroker, Spectra, SpectraGuard, Spectra Logic, Spectra Vail, StorCycle, TeraPack, TFinity, and TranScale are registered trademarks of Spectra Logic Corporation. All rights reserved worldwide. All other trademarks and registered trademarks are the property of their respective owners.

PART NUMBER

90990172 Revision A

REVISION HISTORY

Revision	Date	Description
A	February 2024	Initial Release.

- Notes:**
- To make sure you have the most current version of this guide, see the Spectra Logic Technical Support portal at support.spectralogic.com/documentations/user-guides/.
 - To make sure you have the release notes for the most current version of the BlackPearl OS software, see the Spectra Logic Technical Support portal at support.spectralogic.com/documentations/software-release-notes. The release notes may contain updates to the *User Guide* since the last time it was revised.

END USER LICENSE AGREEMENT

1. READ CAREFULLY

YOU SHOULD READ THE FOLLOWING TERMS AND CONDITIONS BEFORE ACCEPTING THIS END-USER LICENSE AGREEMENT ("EULA"). THIS EULA IS A LEGAL AGREEMENT BETWEEN YOUR ORGANIZATION, THE END USER, AND SPECTRA LOGIC CORPORATION ("SPECTRA") FOR THE SPECTRA SOFTWARE PRODUCT WHICH INCLUDES COMPUTER SOFTWARE AND MAY INCLUDE ASSOCIATED MEDIA, PRINTED MEDIA, AND "ONLINE" OR ELECTRONIC DOCUMENTATION (COLLECTIVELY, "SOFTWARE PRODUCT"). BY INSTALLING, COPYING, OR OTHERWISE USING THE SOFTWARE PRODUCT, YOU AGREE TO BE BOUND BY THE TERMS OF THIS EULA. IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MAY NOT INSTALL, COPY, DOWNLOAD OR USE THE SOFTWARE PRODUCT. YOU AGREE THAT YOUR USE OF THE SOFTWARE ACKNOWLEDGES THAT YOU HAVE READ THIS AGREEMENT, UNDERSTAND IT, AND AGREE TO BE BOUND BY ITS TERMS AND CONDITIONS.

2. OWNERSHIP

It is understood and agreed that Spectra Logic Corporation, a Delaware corporation with offices at 6285 Lookout Road, Boulder, CO 80301 ("Licensor") is the owner of all right, title and interest to the Software Product, regardless of the media or form of the original download, whether by the World Wide Web, disk or otherwise. You, as licensee ("Licensee") through your downloading, installing, copying or use of this product do not acquire any ownership rights to the Software Product.

3. GENERAL

The Software Product is licensed, not sold, to you by Spectra for use only under the terms of this EULA. The Software Product is protected by copyright laws and international copyright treaties, as well as other intellectual property laws and treaties. The rights granted herein are limited to Spectra's and its licensors' intellectual property rights in the Software Product and do not include any other patents or intellectual property rights. The terms of this EULA will govern any software upgrades provided by Spectra that replace and/or supplement the original Software Product, unless such upgrade is accompanied by a separate license in which case the terms of that license will govern.

4. SOFTWARE PRODUCT

The Software Product, as used in this EULA, means, collectively and/or as applicable:

- Related explanatory written materials and instructions, and any other possible documentation related thereto ("Documentation"); and
- Upgrades, modified versions, updates, additions and copies of the Software Product (the "Upgrades"), if any, licensed to by Spectra under this EULA.
- The Software Product package;
- Any and all contents, components, attachments, software, media, and code with which this Agreement is provided and delivered;
- Any and all images, photographs, art, art work, clip art, fonts or other artistic works (the "Art Work");

5. GRANT OF LICENSE AND RESTRICTIONS

- A. Spectra grants you a non-exclusive, non-transferable End-User license right to install the Software Product solely for the purpose for which it was created.

-
- B.** Unless provided otherwise in the Documentation or by prior express written consent of Spectra, you shall not display, modify, reproduce and distribute any Art Work, or portion(s) thereof, included with or relating to the Software Product, if any. Any such authorized display, modification, reproduction and distribution shall be in full accord with this EULA. Under no circumstances will your use, display, modification, reproduction and distribution of the Art Work give you any Intellectual Property or Proprietary Rights of the Art Work. All rights, title, and interest belong solely to Spectra.
 - C.** Except for the initial loading of the Software Product, you shall not, without Spectra's express written consent:
 - Copy or reproduce the Software Product; or
 - Modify, adapt, or create derivative works based on the Software Product or any accompanying materials.

6. DESCRIPTION OF OTHER RIGHTS AND LIMITATIONS

- A.** Spectra will provide you with support services related to the Software Product ("Support"). Such Support will be provided in accordance with the Spectra Master Support Agreement, available for download and viewing on the Spectra Corporate Web site. Use of Support is governed by this EULA and Spectra's Master Support Agreement.
- B.** Any supplemental software, code, content, or media provided to you in the course of Support shall be considered part of the Software Product and subject to the terms and conditions of this EULA.
- C.** Spectra retains all right, title, and interest in and to the Software Product, and any rights not granted to you herein are reserved by Spectra. You hereby expressly agree not to extract information, reverse engineer, disassemble, decompile, or translate the Software Product, or otherwise attempt to derive the source code of the Software, except to the extent allowed under any applicable law. In the event that such activities are permitted by applicable law, any information you, or your authorized agent, discover shall be promptly disclosed to Spectra and shall be deemed the confidential information of Spectra.
- D.** You shall not modify, sublicense, assign, or transfer the Software Product or any rights under this EULA, except as expressly provided in this EULA. Any attempt to sublicense, assign, or transfer any of the rights, duties, or obligations will be void.
- E.** You may permanently transfer all of your rights under this EULA, provided you retain no copies. The other party must agree to accept the terms and conditions of the EULA.

7. ALL RESERVED

All rights not expressly granted herein are reserved by Spectra.

8. TERM

- A.** This License is effective until terminated. Licensee may terminate it at any time by destroying the Software Product with all copies, full or partial, and removing all of its component parts.
- B.** Your rights under this EULA will terminate automatically without notice from Spectra if you fail to comply with any term(s) or condition(s) of this EULA. In such event, no notice shall be required by Spectra to effect such termination.
- C.** Upon termination of this EULA, you shall cease all use of the Software Product and destroy all copies, full or partial, together with all backup copies, modifications, printed or written materials, and merged portions in any form and remove all component parts of the Software Product.

9. INTELLECTUAL PROPERTY RIGHTS

- A. Spectra shall retain all right, title, and interest in the Software Product and to any modifications or improvements made thereto, and any upgrades, updates or Documentation provided to End User. End User will not obtain any rights in the Software Product, its updates, upgrades, and Documentation, as a result of its responsibilities hereunder.
- B. End User acknowledges Spectra's exclusive rights in the Software Product and that the Software Product is unique and original to Spectra and that Spectra is owner thereof. Unless otherwise permitted by law, End User shall not, at any time during or after the effective Term of the Agreement, dispute or contest, directly or indirectly, Spectra's exclusive right and title to the Software Product or the validity thereof.

10. U.S. GOVERNMENT END USERS

The Software Product and related documentation are "Commercial Items," as that term is defined at 48 C.F.R. §2.101, consisting of "Commercial Computer Software" and "Commercial Computer Software Documentation," as such terms are used in 48 C.F.R. §12.212 or 48 C.F.R. §§227.7202-1 through 227.7202-4, as applicable. The Commercial Computer Software and Commercial Computer Software Documentation are being licensed to U.S. Government end users (a) only as Commercial Items and (b) with only those rights as are granted to all other End Users pursuant to the terms and conditions herein. Unpublished rights reserved under the copyright laws of the United States.

11. EXPORT LAW ASSURANCES

You may not use or otherwise export or re-export the Software Product except as authorized by United States law and the laws of the jurisdiction in which the Software Product was obtained. In particular, but without limitation, the Software Product may not be exported or re-exported (a) into (or to a nation or resident of) any U.S. embargoed countries or (b) to anyone on the U.S. Treasury Department's list of Specially Designated Nationals or the U.S. Department of Commerce Denied Persons List or Entity List. By installing or using any component of the Software Product, you represent and warrant that you are not located in, under control of, or a national or resident of any such country or on any such list.

12. DISCLAIMER OF WARRANTIES

YOU EXPRESSLY ACKNOWLEDGE AND AGREE THAT USE OF THE SOFTWARE PRODUCT IS AT YOUR SOLE RISK AND THAT THE ENTIRE RISK AS TO SATISFACTORY QUALITY, PERFORMANCE, ACCURACY AND EFFORT IS WITH YOU. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, AND EXCEPT AS MAY BE STATED IN THE SPECTRA MASTER SERVICE AGREEMENT, THE SOFTWARE PRODUCT IS PROVIDED "AS IS," WITH ALL FAULTS AND WITHOUT WARRANTY OF ANY KIND, AND SPECTRA AND SPECTRA'S AFFILIATES (COLLECTIVELY REFERRED TO AS "SPECTRA" FOR THE PURPOSES OF SECTIONS 12 AND 13) HEREBY DISCLAIM ALL WARRANTIES AND CONDITIONS WITH RESPECT TO THE SOFTWARE PRODUCT, EITHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES AND/OR CONDITIONS OF MERCHANTABILITY, OF SATISFACTORY QUALITY, OF FITNESS FOR A PARTICULAR PURPOSE, OF ACCURACY, OF QUIET ENJOYMENT, AND NON-INFRINGEMENT OF THIRD-PARTY RIGHTS. SPECTRA DOES NOT WARRANT AGAINST INTERFERENCE WITH YOUR ENJOYMENT OF THE SOFTWARE PRODUCT THAT THE FUNCTIONS CONTAINED IN THE SOFTWARE PRODUCT WILL MEET YOUR REQUIREMENTS, THAT THE OPERATION OF THE SOFTWARE PRODUCT WILL BE UNINTERRUPTED OR ERROR-FREE, OR THAT DEFECTS IN THE SOFTWARE PRODUCT WILL BE CORRECTED. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY SPECTRA OR A SPECTRA AUTHORIZED REPRESENTATIVE SHALL CREATE A WARRANTY. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES OR LIMITATION ON APPLICABLE STATUTORY RIGHTS OF A CONSUMER, SO THE ABOVE EXCLUSION AND LIMITATIONS MAY NOT APPLY TO YOU.

13. LIMITATION OF LIABILITY

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT SHALL SPECTRA, ITS AFFILIATES OR LICENSEES, BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, OR ANY OTHER PECUNIARY LOSS) ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE PRODUCT OR THE PROVISION OF OR FAILURE TO PROVIDE SUPPORT SERVICES, EVEN IF SPECTRA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN ANY CASE, SPECTRA'S ENTIRE LIABILITY UNDER ANY PROVISION OF THIS EULA SHALL BE LIMITED TO THE AMOUNT ACTUALLY PAID BY YOU FOR THE SOFTWARE PRODUCT; PROVIDED HOWEVER, IF YOU HAVE ENTERED INTO A MASTER SUPPORT AGREEMENT, SPECTRA'S ENTIRE LIABILITY REGARDING SUPPORT SERVICES SHALL BE GOVERNED BY THE TERMS OF THAT AGREEMENT. BECAUSE SOME STATES AND JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF LIABILITY, THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

14. CONTROLLING LAW AND SEVERABILITY

This EULA will be governed by and construed in accordance with the laws of the State of Colorado, as applied to agreements entered into and to be performed entirely within Colorado between Colorado residents. This EULA shall not be governed by the United Nations Convention on Contracts for the International Sale of Goods, the application of which is expressly excluded. If for any reason a court of competent jurisdiction finds any provision, or portion thereof, to be unenforceable, the remainder of this EULA shall continue in full force and effect.

SYSTEM BIOS

Resetting the system BIOS when not authorized by Spectra Logic Technical Support invalidates the system configuration. Spectra Logic reserves the right to charge for time and materials to reconfigure and recertify the system.

CONTACTING SPECTRA LOGIC

To Obtain General Information	
Spectra Logic Website: spectralogic.com	
United States Headquarters	European Office
Spectra Logic Corporation 6285 Lookout Road Boulder, CO 80301 USA Phone: 1.800.833.1132 or 1.303.449.6400 International: 1.303.449.6400 Fax: 1.303.939.8844	Spectra Logic Europe Ltd. 329 Doncastle Road Bracknell Berks, RG12 8PE United Kingdom Phone: 44 (0) 870.112.2150 Fax: 44 (0) 870.112.2175
Spectra Logic Technical Support	
Technical Support Portal: support.spectralogic.com	
United States and Canada Phone: Toll free US and Canada: 1.800.227.4637 International: 1.303.449.0160	Europe, Middle East, Africa Phone: 44 (0) 870.112.2185 Deutsch Sprechende Kunden Phone: 49 (0) 6028.9796.507 Email: spectralogic@stortrec.de
Mexico, Central and South America, Asia, Australia, and New Zealand Phone: 1.303.449.0160	
Spectra Logic Sales	
Website: shop.spectralogic.com	
United States and Canada Phone: 1.800.833.1132 or 1.303.449.6400 Fax: 1.303.939.8844 Email: sales@spectralogic.com	Europe Phone: 44 (0) 870.112.2150 Fax: 44 (0) 870.112.2175 Email: eurosales@spectralogic.com
To Obtain Documentation	
Spectra Logic Website: support.spectralogic.com/documentations	

Table of Contents

About This Guide	11
Intended Audience	11
Related Information	12
Related Publications	12
Tape Library User Guides	14
Online Forum	14
Chapter 1 - Understanding Advanced Bucket Management	15
Goals of Advanced Bucket Management	17
DS3 Overview	17
DS3 Clients	18
BlackPearl Cache	19
Storage Domains	20
Data Policies	20
Data Persistence Rules	20
Data Replication Rules	21
Tape Export Strategy	21
Tape and Disk Partitions	23
Tape Partitions	23
Tape Drive Reservation	23
Tape Media Inspections	24
Object Storage Disk	25
Storage Disk Pools	25
Special Considerations for Exporting Tapes	26
Tape Export Best Practices	26
Special Considerations for Reading Tapes in a Non-BlackPearl Nearline Gateway Environment	27
Example Configurations	29
Single Copy on Tape	29
Dual Copy on Tape	33
Single Copy on Nearline Object Storage Disk	38
Single Copy on Nearline Object Storage Disk and Tape	42
Single Copy on Nearline Object Storage Disk and Dual Copy on Tape	47
Additional Configurations	53

Chapter 2 - Configuring Advanced Bucket Management	54
Create a Disk Pool	55
Create a Nearline Storage Disk Pool	56
Create an Online Disk Pool	59
Create a Disk Partition	62
Create a Tape Partition	63
Create a Replication Target	64
Create a BlackPearl Target	64
Create an Amazon S3 Target	68
Create a Microsoft Azure Target	72
Create a Storage Domain	75
Add a Storage Domain Member to a Storage Domain	81
Create a Data Policy	83
Add Data Persistence Rules and Replication Rules to a Data Policy	86
Create Data Policy ACLs	92
Create a Bucket	95
Create a New Bucket ACL for a Group	96
Create a New Bucket ACL for a User	97
Transfer Data	99
Chapter 3 - Managing Advanced Bucket Management Settings	100
Manage Buckets	102
Show Bucket Physical Placement	102
Edit a Bucket ACL	102
Delete a Bucket ACL	103
Edit a Bucket	103
Delete a Bucket	104
Manage a Storage Domain	106
Edit a Storage Domain Member	106
Exclude a Storage Domain Member	107
Cancel Storage Domain Member Exclusion	108
Delete a Storage Domain Member	108
Edit a Storage Domain	109
Delete a Storage Domain	110
Manage Data Replication Rules	112
Edit a BlackPearl Data Replication Rule	112

Edit an Amazon S3 Data Replication Rule	113
Edit a Microsoft Azure Data Replication Rule	114
Delete a Data Replication Rule	115
Manage a Data Policy	116
Edit a Data Policy	116
Delete a Data Policy	117
Delete a Data Policy ACL	118
Edit a Data Persistence Rule	118
Delete a Data Persistence Rule	119
Manage Replication Targets	121
Edit a BlackPearl Replication Target	121
Edit an Amazon Replication Target	122
Edit an Azure Replication Target	123
Verify a Replication Target	123
Put a Replication Target in Standby State	124
Activate a Replication Target	125
Delete a Replication Target	126
Manage a Disk Partition	127
Edit a Disk Partition	127
Delete a Disk Partition	128
Manage Online and Nearline Disk Pools	129
Edit a Nearline or Online Disk Pool	129
Import a Nearline or Online Disk Pool	130
Delete a Nearline or Online Disk Pool	131
Change the BlackPearl Gateway Cache Configuration	133

ABOUT THIS GUIDE

This guide describes the concepts and use of Advanced Bucket Management in the Spectra® BlackPearl® Nearline gateway master node.

Advanced bucket management is a collection of expert-level settings for the BlackPearl gateway. Spectra logic highly recommends working with Spectra Logic Professional Services before creating or modifying advanced bucket management settings. See [Contacting Spectra Logic](#) on page 7.

INTENDED AUDIENCE

This guide is intended for data center administrators and operators who maintain and operate file storage systems. The information in this guide assumes a familiarity with computing terminology and with network connectivity protocols such as SAS, Fibre Channel, and Ethernet. If your BlackPearl system installation includes a tape library, knowledge of tape-based backup systems and how to use the library is required. You also need to be familiar with installing, configuring, and using data file storage and data management software.

RELATED INFORMATION

This section contains information about this document and other documents related to the Spectra BlackPearl Nearline Gateway.

Typographical Conventions

This document uses the following conventions to highlight important information:

**WARNING**

Read text marked by the “Warning” icon for information you must know to avoid personal injury.

**CAUTION**

Read text marked by the “Caution” icon for information you must know to avoid damaging the hardware or losing data.

**IMPORTANT**

Read text marked by the “Important” icon for information that helps you complete a procedure or avoid extra steps.

Note: Read text marked with “Note” for additional information or suggestions about the current topic.

Related Publications

For additional information about the Spectra BlackPearl Nearline gateway and the DS3 interface, refer to the publications listed in this section.

Spectra BlackPearl Nearline Gateway

The following documents related to the Spectra BlackPearl Nearline gateway are available on the Support Portal website at support.spectralogic.com, and from the Documentation screen in the BlackPearl user interface.

- The [Spectra BlackPearl Nearline Gateway User Guide](#) provides detailed information about configuring, using, and maintaining your BlackPearl gateway.
- The [Spectra BlackPearl Site Preparation Guide](#) provides important information that you should know before installing a BlackPearl gateway in your storage environment.
- The [Spectra BlackPearl Rack Mounting Instructions Guide](#) provides detailed instructions for installing a Gen1 BlackPearl gateway in a standard rack.
- The [Spectra BlackPearl Network Setup Tips](#) document provides helpful instructions for troubleshooting common connectivity problems.

- The *Spectra BlackPearl DS3 API Reference* provides information on understanding and using the DS3 API.
- The *Spectra BlackPearl HotPair Installation & Configuration Guide* document provides detailed information on installing and using the BlackPearl gateway in a HotPair configuration.

The following documents are available after logging into your Support portal account at: support.spectralogic.com.

- The *Spectra BlackPearl Release Notes and Documentation Updates* provide the most up-to-date information about the BlackPearl gateway, including information about the latest software releases and documentation updates.
- The *BlackPearl Eon Browser User Guide* provides installation and usage information for the Spectra Eon browser.
- The *Spectra 12- & 36-Drive Chassis HBA Installation Guide* provides instructions for installing an HBA in a Gen1 master node.
- The *Spectra 12- & 36-Drive Chassis Boot Drive Replacement Guide* provides instructions for replacing a failed boot drive in a Gen1 master node.
- The *Spectra 12-, 36- & 45-Drive Chassis Drive Replacement Guide* provides instructions for replacing a failed data drive in a Gen1 master node or 44-bay expansion node.
- The *Spectra 12-, 36- & 45-Drive Chassis Fan Replacement Guide* provides instructions for replacing a failed fan in a Gen1 master node or 44-bay expansion node.
- The *Spectra 12-, 36- & 45-Drive Chassis Power Supply Replacement Guide* provides instructions for replacing a failed power supply in a Gen1 master node or 44-bay expansion node.
- The *Spectra 12-Drive Chassis HBA Replacement Guide* and *Spectra 36-Drive Chassis HBA Replacement Guide* provide instructions for replacing a failed HBA in a Gen1 master node.
- The *Spectra 96-Bay Chassis Drive Replacement Guide* provides instructions for replacing a failed data drive in the 96-bay expansion node.
- The *Spectra 96-Bay Chassis Fan Replacement Guide* provides instructions for replacing a failed fan in the 96-bay expansion node.
- The *Spectra 96-Bay Chassis Power Supply Replacement Guide* provides instructions for replacing a failed power supply in the 96-bay expansion node.
- The *Spectra 96-Bay Chassis I/O Module Replacement Guide* provides instructions for replacing a failed I/O module in the 96-bay expansion node.
- The *Spectra 107-Bay Expansion Node FRU Guide* provides instructions for replacing fans, power supplies, drives, and SAS expanders in the 77-bay and 107-bay expansion node.

Tape Library User Guides

Spectra Logic Tape Libraries

User Guides for Spectra Logic tape libraries are posted on the Support Portal website at: support.spectralogic.com/documentations/user-guides.

IBM Tape Libraries

User Guides for compatible IBM® tape libraries are posted on the IBM Knowledge Center website at: ibm.com/support/knowledgecenter/products/.

ONLINE FORUM

Need help with Spectra Logic's S3 software development kits or the DS3 API? Post your question at the Spectra Logic S3-SDK discussion forum located at: <https://developer.spectralogic.com/forums>

CHAPTER 1 - UNDERSTANDING ADVANCED BUCKET MANAGEMENT

This chapter explains the concepts of the BlackPearl Nearline gateway advanced bucket management. It is important to understand the information in this chapter before you begin designing the storage architecture of the BlackPearl gateway.



IMPORTANT

It is difficult and time consuming to change a data policy once the gateway writes data to a bucket using the data policy. Make sure that you understand the concepts in this chapter and have thoughtfully planned your data policies and consulted with Spectra Logic Professional Services before you start using the BlackPearl gateway to store data.

Goals of Advanced Bucket Management	17
DS3 Overview	17
DS3 Clients	18
BlackPearl Cache	19
Storage Domains	20
Data Policies	20
Data Persistence Rules	20
Data Replication Rules	21
Tape Export Strategy	21
Tape and Disk Partitions	23
Tape Partitions	23
Tape Drive Reservation	23
Tape Media Inspections	24
Object Storage Disk	25
Storage Disk Pools	25
Special Considerations for Exporting Tapes	26
Tape Export Best Practices	26
Special Considerations for Reading Tapes in a Non-BlackPearl Nearline Gateway Environment	27
Example Configurations	29

Single Copy on Tape	29
Dual Copy on Tape	33
Single Copy on Nearline Object Storage Disk	38
Single Copy on Nearline Object Storage Disk and Tape	42
Single Copy on Nearline Object Storage Disk and Dual Copy on Tape	47
Additional Configurations	53

GOALS OF ADVANCED BUCKET MANAGEMENT

The BlackPearl gateway provides a DS3 front end interface to disk, tape and cloud storage. The BlackPearl Advanced Bucket Management (ABM) feature automates many aspects of deep storage including policy based multiple copies on diverse media types without the need for expensive middleware to operate the libraries and stream data to tape drives. The BlackPearl gateway delivers seamless data management enabling infinite retention, seamless growth, and unlimited retrieval of data for as low as pennies per Gigabyte.

DS3 OVERVIEW

The Spectra BlackPearl Nearline Gateway allows data to move seamlessly into deep storage in a way not previously possible. DS3 is the first native REST-based interface to deep storage which enables easy archiving of large amounts of bulk data. It enables users to deploy tape, nearline disk, and online disk storage that is cost effective, easy to manage, and scalable to exabytes of data.

DS3 utilizes the standard Amazon S3 operations plus additional operations specifically designed to optimize the transport of data objects to and from deep storage. The additional operations define the job so that BlackPearl gateway interacts with the objects efficiently and define the data policy to customize where and for how long specific data is stored.

The first of these additional operations is called START BULK PUT. It is an HTTP PUT operation that provides BlackPearl gateway with information about the objects that the client wants to send as a single job for storing on tape. The Create Bulk Put command is sent with a payload that is made up of a list of object names and corresponding object sizes. This information allows the BlackPearl gateway to plan the initial storage of the objects in its cache, and how it will store the data on tape. The response to the Create Bulk Put command is a specifically ordered list of how the BlackPearl gateway wants those files (objects) sent.

The second command is called Create Bulk Get. The Create Bulk Get command is actually an HTTP PUT command because it too contains a payload for the BlackPearl gateway. This payload is a list of objects that the client wants to get from the BlackPearl gateway. It is not necessary for the request payload to contain the size of the files because the BlackPearl gateway already knows the sizes of the objects (files). The response to the request is again an ordered list of the objects and information about the objects, including if they are already in the cache and ready to be retrieved from cache by a GET command.

Knowing the files that the client wants to retrieve, the BlackPearl gateway can make the best use of its resources in retrieving the objects. For example, if the list of objects spans across four different tapes and there are four tape drives available, those four tapes can all be loaded into drives and the objects can be read back in parallel, greatly improving the speed at which the client can get all of the objects. Without the Create Bulk Get request, the client would be asking the BlackPearl gateway for those objects in a less efficient manner.

Storing large amounts of bulk data on tape has historically presented challenges. DS3 addresses these challenges:

- Tape drives are sequential block storage devices, with data laid out in a sequential manner along the full length of the tape. This makes it inefficient to retrieve data out of order. DS3 plans and queues a large amount of data to be efficiently written to tape; it logically groups data on tape in a way that reflects how the client is likely to read it back.
- Because of the mechanical nature of the tape media and drives, tape drives demand a large amount of data to be available, via a fast connection. When data is not efficiently streamed for the tape drives to write (due to slow data buffering or a slow connection to the drive), the result is poor write performance. This poor performance is due to a phenomenon referred to as “shoe-shining”. When a drive is sent a small amount of data, it writes the data and then is forced to stop. Because the tape cannot stop instantaneously, the drive overshoots a small amount and the tape is not in position for the next write operation. To compensate, the tape drive rewinds to get back to the correct position for the next write. If the next write also has a small amount of data, then the drive writes the next portion and again stops, overshoots, and rewinds, causing a back and forth “shoe-shining” like action. DS3 caches data on the BlackPearl Nearline Gateway before starting the transfer to tape, which prevents the shoe-shining behavior from occurring.
- Classically, different tape storage devices wrote data to tape in unique ways, locking you into a proprietary and single vendor solution to retrieve previously written data. DS3 writes data to tape using the open source Linear Tape File System (LTFS). With LTFS, data is always accessible with any LTFS enabled system.

For more information on the DS3 interface, see the [Spectra BlackPearl DS3 API Reference](#).

DS3 Clients

Users can leverage a library of existing DS3 clients available through the [Spectra Logic Developer Program](#), or develop their own client. The user moves data through the client to the BlackPearl gateway and then the gateway handles all interaction with the data storage hardware.

BLACKPEARL CACHE

The BlackPearl cache is allocated physical storage on either HDDs or SSDs installed in the gateway. The cache functions as a transient location for all data transferred to the BlackPearl gateway from a client, or transferred from tape storage to the BlackPearl gateway.

The capacity available for cache is managed by the BlackPearl data planner, where active jobs reserve various amounts of cache capacity known as 'chunks', and chunk size can vary.

When writing data to cache destined for tape storage, or restoring data from tape storage to the BlackPearl gateway, the chunk size is typically 2% of the capacity of a single tape cartridge. If the total job size is less than that amount, the chunk size reduces in size to match the job size.

STORAGE DOMAINS

A storage domain is a collection of data partitions and, when applicable, media type combinations. Storage domains define the possible places where data sent to the BlackPearl Nearline Gateway can be stored. Data persistence rules and data policies further define where and for how long specific data is stored.

Entire data partition/media type combinations are members of storage domains. When additional capacity is required, a single storage pool or tape is allocated out of the member data partitions to fulfill the capacity requirement.

DATA POLICIES

A data policy defines data integrity policies (checksum type and end-to-end CRC requirements), default job priorities, and data persistence rules, which define where data should be written and for how long it should be kept. A data policy may be used by multiple buckets, but a bucket uses precisely one data policy.

A data policy consists of one or more permanent persistence rules, zero or more temporary persistence rules, and zero or more retired persistence rules. A persistence rule can be permanent, meaning that data is kept in the specified storage domain at all times, or temporary, meaning that data is kept in the specified storage domain under certain circumstances, and then it can be deleted from that storage domain. Existing permanent and temporary persistence rules may be retired so that the rule is not applied for any new incoming data, but continues to retain data previously written.

Data Persistence Rules

Each data policy must have one or more permanent persistence rules. Each persistence rule targets a specified storage domain. There are three types of persistence rules:

- **Permanent** — A copy of the data is placed in the specified storage domain initially and maintained there permanently.
- **Temporary** — A copy of the data is placed in the specified storage domain initially and maintained there at least until the specified retention period expires.
- **Retired** — The rule is not applied for any new incoming data, but continues to retain data previously written.

Data is written to every storage domain for which there is a persistence rule with the type configured as permanent or temporary.

The same storage domain cannot be specified multiple times using different persistence rules in the same data policy. The same storage domain can be referred to across different data policies.

Data persistence rules must specify the level of physical isolation required for the data retention. There are two types of data retention:

- **Standard** — Data is isolated according to the standard storage domain isolation requirements. When more storage is needed, a tape or pool is assigned to the storage domain. Any buckets using that storage domain can have data on the pool or tape, which can make it difficult to export all of the tapes for a single bucket.
- **Bucket Isolated** — Data from different buckets cannot be mixed on the same physical storage media.

- Notes:**
- The **Standard** isolation level provides the best capacity utilization and overall performance.
 - **Bucket Isolated** allocates an entire tape or pool to a bucket when needed. Allocating an entire pool to a bucket may use up resources quickly and is not recommended.

Data Replication Rules

Data policies may also contain data replication rules. The BlackPearl gateway supports replicating data to the following targets:

- **BlackPearl target** — A BlackPearl gateway remote to the local gateway that stores replicated data.
- **Amazon S3 target** — An AWS S3 instance remote to the BlackPearl gateway that stores replicated data.
- **Microsoft Azure target** — A Microsoft Azure instance remote to the BlackPearl gateway that stores replicated data.

Tape Export Strategy

A tape export strategy must be considered as part of a data policy. Spectra recommends keeping at least one copy of all archived data in the tape library at all times. Libraries can be easily upgraded by purchasing more slot licenses, or, if the slots become completely full, upgrading the library itself to one with more slots using the exclusive Spectra TranScale technology.

A tape library user or administrator may decide to export media cartridges from a tape library for any of the reasons described below:

- **Exporting a copy for off-site disaster recovery:** The BlackPearl gateway allows a user to make multiple copies of data automatically. A typical use case is to create a “tape first copy” that is intended to be left in the library for easy retrieval as well as an “export copy” intended to be removed from the library once full for archival at an alternate site for safety. See [Configuring Advanced Bucket Management on page 54](#) for information on setting up multiple copies and exporting a copy, and , or your [Tape Library User Guide](#), for details on the physical process of exporting and importing tapes into the library.
- **Exporting a copy of data for transfer to another location:** In some work flows, a user exports a tape to transfer the data to another facility. Individual tapes can be exported manually using the BlackPearl user interface (see).
- **Exporting tapes to free up space in the library:** Some work flows and budgets require older or unused media to be exported, making it not readily available to the BlackPearl gateway, in order to free up space in the tape library. After tapes are exported, new tapes are imported to provide the BlackPearl gateway with new media for storage operations.

TAPE AND DISK PARTITIONS

Tape and disk partitions are external data storage targets cabled to the BlackPearl gateway through SAS or Fibre Channel connections. Once created, partitions are assigned to storage domains.

Tape Partitions

Tape partitions refer to data partitions configured on Spectra Logic or other supported tape libraries. When you create a partition on a tape library attached to a BlackPearl gateway, the gateway automatically detects the tape library partition and adds it to the list of available partitions in the BlackPearl user interface. The gateway also automatically creates two commonly used storage domains: tape first copy, and tape second copy. For more information on storage domains, see [Storage Domains on page 20](#).

- Notes:**
- Cleaning partitions are not added to the BlackPearl user interface.
 - Tape drive cleaning is typically handled by the Spectra tape library. For more information on cleaning the library tape drives using the tape library, see your *Tape Library User Guides*.
 - Tape drive cleaning is handled automatically by the IBM TS4500 tape library.
 - Tape drive cleaning for Spectra Logic libraries can be initiated through the DS3 API. See the [Spectra BlackPearl DS3 API Reference](#) for more information.
 - Tape drive cleaning cannot be initiated through the BlackPearl user interface.
 - If the BlackPearl gateway is configured to transfer data to tape, make sure the partitions configured for use by the BlackPearl gateway do not contain WORM (Write Once-Read Many) media. The BlackPearl gateway is not compatible with WORM media.

Tape Drive Reservation

Tape drive reservation allows you to control how the tape drives are used to transfer data, by dedicating drives to accept only read commands or write commands, and to accept only jobs of a specified priority level or higher. With a large number of tape drives, using drive reservation can increase efficiency and reduce latency when either reading or writing data. Reserving tape drives for either reading or writing, or for a specified job priority level, is not required and is typically only used when read or write throughput and drive availability are important enough to dedicate tape drives to that function.

- Note:** Tape drive reservation is not recommended for BlackPearl gateways connected to two or fewer tape drives.

Tape drive reservation is configured on both the drive, and library partition level.

- When reserving an individual tape drive, you can exclude the drive from performing reads, writes, or jobs lower than a specified level.
- You can also configure the library partition to reserve a specified number of drives for either reads or writes. This can prevent individual tape drive failures, or unavailable drives, from impacting the desired number of drives available for either read or write commands.

Note: Tape drives always allow inspection and verify tasks.

**IMPORTANT**

Spectra Logic does not recommend setting both a minimum reservation priority and reserved task type for the same drive.

Tape Media Inspections

When new tape media is added to the BlackPearl gateway, the gateway inspects the tape cartridge as configured in the DS3 service. However, all tapes that are new to the gateway require inspection before they are usable in a managed state. Additionally, under certain circumstances, the gateway may override the configured behavior for inspections on already managed tapes.

Inspection Behavior when the Tape Library is in Standby

With BlackPearl OS 5.2 or earlier, if there was an inventory change in the tape library while in standby, the BlackPearl gateway could react by re-inspecting all tapes in the library. Starting with BlackPearl OS 5.3, the gateway no longer re-inspects tapes when there is an inventory change while the partition is in standby if the S3 service is set to "Never Inspect".

Inspection Behavior when the Tape Library is in Active Use

Prior to BlackPearl OS 5.3, if a tape library associated with the BlackPearl gateway goes offline, or is otherwise made unavailable to the BlackPearl gateway, the tapes are marked as "Lost" by the BlackPearl gateway. When this occurs, the tape is considered to be a new tape, and is re-inspected when the tape library is made available. Starting with BlackPearl OS 5.3, the BlackPearl gateway automatically quiesces the tape partition and prevents the tape cartridge being marked "Lost".

Object Storage Disk

Object Storage Disk uses external SAS expansion nodes to provide both Online Storage Disk and Nearline Storage Disk partitions. The disk partitions are behind the BlackPearl DS3 interface, but can also be accessed using the Spectra StorCycle application, or the Spectra Vail application. Nearline Object Storage disk can also be used as a cache for data stored on tape.

Storage Disk Pools

Disk partitions are comprised of disk storage pools on external SAS expansion nodes. There are two types of disk partitions. **Online Storage Disk** partitions are created on Spectra 44-bay expansion nodes, while **Nearline Storage Disk** partitions are created on Spectra 96-bay expansion nodes. Both the 77-bay and 107-bay expansion nodes provide either **Online Storage Disk** or **Nearline Storage Disk** partitions, depending on the type of drives installed.

- **Online Storage Disk** — Is used as a temporary, high-performance storage target for highly transactional data.
- **Nearline Storage Disk** — Is a cost effective storage target for deep storage. Nearline storage is not recommended for frequent reads or writes.

Disk storage pools are manually created using the BlackPearl user interface after you attach an expansion node. See [Create a Disk Pool on page 55](#).

Once storage pools are created, you must manually add them to disk partitions. For information on creating a disk partition, see [Create a Disk Partition on page 62](#).

SPECIAL CONSIDERATIONS FOR EXPORTING TAPES

If you plan to export tapes, Spectra Logic recommends the following storage domain and data policy settings:

- Set the Write Optimization setting to **Capacity** when configuring a storage domain so that data is written to as few tapes as possible.
- Enable **Bucket Isolation** when configuring a data policy. This setting configures the bucket to have its own unique set of tapes. This ensures that tape media containing one bucket of information is not mixed with another bucket, making it easier to export a bucket.

Tape Export Best Practices

The system Administrator **must** be configured to receive emails with both Informational and Warning message severity to be notified when tape media is exported. This allows the user to retrieve the tape media when it is exported. Do not leave tape media in the library Entry/Exit port for long periods of time. Tape media left in the Entry/Exit port may interfere with other automatic tape export operations, or import of new or requested tape media.

By configuring email alerts, the user is also notified when a GET job is requesting an object from exported tape media, so it can be imported into the tape library to complete the GET job.

Note: All **request to import tape media** messages and emails list the required tape cartridge barcode(s), which help you quickly identify the tape cartridges to import.

It is important to not export tape media from the library directly. The BlackPearl Nearline gateway controls the movement of media in the library.

If you are using a Spectra Logic T120 or T50e tape library in with your BlackPearl Nearline gateway, it must be configured with only a single storage partitions. Multiple partitions, including a cleaning partition, are not supported.

SPECIAL CONSIDERATIONS FOR READING TAPES IN A NON-BLACKPEARL NEARLINE GATEWAY ENVIRONMENT

The BlackPearl gateway stores data on LTO-5 or later generation Ultrium, or TS11xx technology tape media using the LTFS format. If you plan to export tapes and read them in a non-BlackPearl gateway, you must follow the guidelines for [Special Considerations for Exporting Tapes](#) on the previous page as well as the guidelines below when configuring your storage domains and data policies.

- The LTFS file name option should be set to **Object Name** when configuring a storage domain. This setting configures LTFS file names to use the format *{bucket name}/{object name}*, for example *bucket1/video1.mov*. If the tapes are exported from the library attached to the BlackPearl gateway and loaded into a non-BlackPearl gateway, the file names match the object names. If you do not configure this option, object names are assigned a UUID string, which is not human readable, but can be translated back to the actual file name using an external conversion tool.
- Object names must comply with LTFS file naming rules:
 - The colon character (:) is not allowed in LTFS file names and therefore not allowed in BlackPearl object names. The slash character (/) is also technically not allowed in LTFS file names; however, the BlackPearl software can accommodate a slash in the object name and translates it as a directory in the LTFS file system (for example, *directory1/directory2/video1.mov*).
 - File names with multiple consecutive slash characters (//) are not allowed.
 - Directory names have a limit of 255 characters.
 - File names have a variable character limit. If you are using English ASCII characters, the limit is 1024 characters. If you are using a graphical language, such as Japanese, the limit is 512 characters.

- Spectra Logic does not recommend the following characters in LTFS file names or BlackPearl object names for reasons of cross-platform compatibility:
 - Asterisk (*)
 - Question mark (?)
 - Question mark (?)
 - Forward slash (/)
 - Backslash (\)
 - Vertical bar / pipe (|)
 - Left curly brace ({)
 - Right curly brace (})
 - Caret (^)
 - Percent character (%)
 - Grave accent / back tick (`)
 - Right square bracket (])
 - Left square bracket ([)
 - Double quotation marks (")
 - Greater Than symbol (>)
 - Less Than symbol (<)
 - Tilde (~)
 - Pound character (#)
 - Control characters such as carriage return (CR) and line feed (LF),
 - Non-printable ASCII characters (128–255 decimal characters)

Note: Spectra Logic does not recommend accented characters in LTFS file names or BlackPearl object names because LTFS normalizes them before objects are written to tape and there could be conflicts with two objects having the same normalized name.

- **Blobbing Enabled** should be cleared when configuring a data policy. Blobbing allows an object larger than 1 TB to be broken into multiple blobs and then stored on multiple tapes. Tapes created with blobbing disabled are always readable by a non-BlackPearl gateway; tapes created with blobbing enabled may not be readable by a non-BlackPearl gateway when very large objects span across multiple tapes. With blobbing disabled, all files must have a size of 1 TB or less.
- **Minimize Spanning** should be cleared when configuring a data policy. This helps further insure that blobs do not span across tapes and provides great LTFS compatibility.
- The **Keep Latest** setting cannot be used for a data policy which uses a storage domain configured with the **LTFS File Name** option set to **Object Name**.

EXAMPLE CONFIGURATIONS

Below are explanations of the preconfigured data policies on the BlackPearl gateway, which include data persistence rules and storage domain targets.



IMPORTANT

It is difficult and time consuming to change a data policy once the gateway writes data to a bucket using the data policy. Make sure that you understand the concepts in this chapter and have thoughtfully planned your data policies and consulted with Spectra Logic Professional Services before you start using the BlackPearl gateway to store data.

Note: For information on additional data policy settings that are not available through the BlackPearl user interface, see the [Spectra BlackPearl DS3 API Reference](#).

Single Copy on Tape

This data policy is the most basic of the preconfigured data policies on the gateway. This policy creates a single copy of each object sent to the gateway on tape media. Once data is written on tape, it is removed from the BlackPearl cache if the gateway detects that more cache space is needed for incoming data.

Note: This data policy is automatically created when the gateway detects a tape partition.

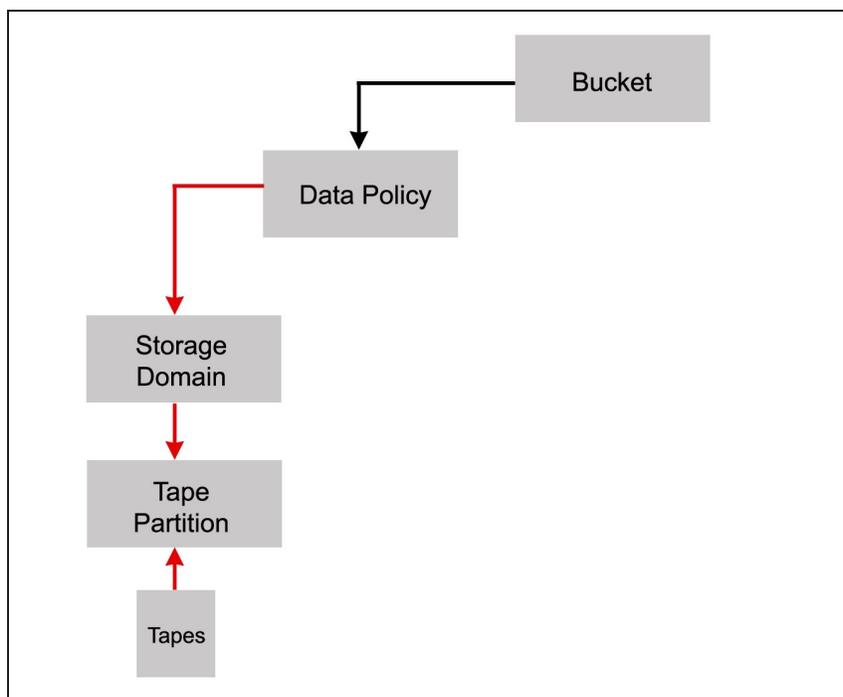


Figure 1 The Single Copy on Tape workflow.

The single copy on tape data policy is configured with the following attributes:

Parameter	Value	Description
Storage Domain - Tape First Copy - Data is written to the primary tape storage domain. This domain is created automatically when the gateway detects the first partition created on the tape library.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.
Write Optimization	Capacity	Job chunks are written across as few pieces of media as possible. Fewer tape drives are used, so performance is lower.
LTFS File Naming	Object ID	File names use the format <code>{bucket name}/{object id}</code> , for example, <code>bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82</code> .
Media Export Allowed	cleared	Media export is not allowed.
Storage Domain Member for Tape First Copy - The first tape partition created on the tape library and detected by the gateway.		
Tape Type	varies	The tape type matching the latest generation of tape drive in the partition.
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Auto Compaction Threshold	20	The percentage of a tape with deleted objects at which auto compaction is triggered. The default is 95. The minimum is 10.
Data Policy - Single Copy on Tape		
Blobbing Enabled	selected	Allows an object to be broken into multiple blobs.
Minimize Spanning	cleared	Jobs larger than 1 TB are allowed to span across multiple tapes or pools as needed to maximize capacity utilization and performance.

Parameter	Value	Description
Default GET Job Priority	High	<p>When the storage domain receives a GET job, it is processed with high priority, which is before low and normal priority jobs. Jobs of different types are put in order based on priority.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a GET job by specifying a different priority in the Get Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default PUT Job Priority	Normal	<p>When the storage domain receives a PUT job, it is processed with normal priority, which is after high priority jobs but before low and normal priority jobs. Jobs of different types are put in order based on priority.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a PUT job by specifying a different priority in the Put Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default VERIFY Job Priority	Low	<p>When the storage domain receives a VERIFY job, it is processed with low priority, which is the lowest setting.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a VERIFY job by specifying a different priority in the Create Verify Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default Verify After Write	cleared	Data is not verified after a write.
Rebuild Priority	Low	If data is lost from tape media, the data is rebuilt using low priority, which is the lowest setting.
Checksum Type	MD5	Data using this storage domain is CRC checked using the MD5 checksum type. Data is CRC checked when it is written to cache with a PUT job, or read back from tape media with a GET job.
End-to-end CRC	No	This data policy does not use end-to-end CRC checking.

Parameter	Value	Description
Versioning	None	In order to upload a new version of an object already PUT to the gateway, the first version must be deleted.
Always Accept Replicated PUT Jobs	cleared	PUT jobs created for this data policy fail if one or more replication targets the gateway must PUT to are unavailable.
Data Persistence Rule for Tape First Copy		
Type	Permanent	Data is moved to tape and maintained on tape media until data is deleted from a bucket. Note: When data is deleted from a bucket, it is removed from tape media, but individual tape media cartridges are not reclaimed for use until all data is deleted from a tape cartridge.
Bucket Isolation Level	Standard	Data from different buckets can be mixed on to the same piece of media.

Dual Copy on Tape

This data policy persists two copies of data for each PUT job the policy receives. Both copies of data are moved to tape media.

This data policy is useful if you want to export one copy of the bucket on tape media for storage off site, which provides enhanced data security in the case of ransomware attacks or disaster recovery.

Note: This data policy is created automatically when the gateway detects the first and second tape partitions created on the tape library.

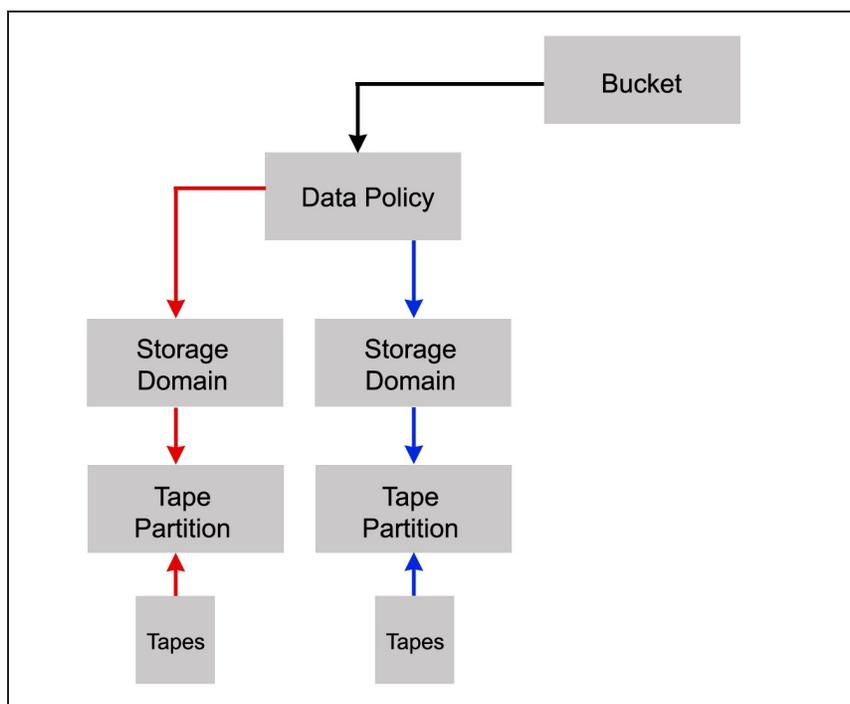


Figure 2 The Dual Copy on Tape workflow.

The dual copy on tape data policy is configured with the following attributes:

Parameter	Value	Description
Storage Domain - Tape First Copy - Data is written to the primary tape storage domain. This domain is created automatically when the gateway detects the first tape partition created on the tape library.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.
Write Optimization	Capacity	Job chunks are written across as few pieces of media as possible. Fewer tape drives are used, so performance is lower.
LTFS File Naming	Object ID	File names use the format <i>{bucket name}/{object id}</i> , for example, bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82.
Media Export Allowed	cleared	Media export is not allowed.
Storage Domain Member for Tape First Copy		
Tape Type	varies	The tape type matching the latest generation of tape drive in the partition.
Auto Compaction Threshold	20	The percentage of a tape with deleted objects at which auto compaction is triggered. The default is 95. The minimum is 10.
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Storage Domain - Tape Second Copy - A second copy of the data is written to a tape storage domain optimized for tape export. This domain is created automatically when the gateway detects a tape partition created on the tape library.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.

Parameter	Value	Description
Write Optimization	Capacity	Job chunks are written across as few pieces of media as possible. Fewer tape drives are used, so performance is lower.
LTFS File Naming	Object ID	File names use the format <i>{bucket name}/{object id}</i> , for example, bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82.
Media Exort Allowed	selected	Media export is allowed.
Auto Export on Job Completion	cleared	Media is not auto exported upon job completion.
Auto Export on Job Cancel	cleared	Media is not auto exported upon job cancellation.
Auto Export on Media Full	cleared	Media is not auto exported upon media full.
Scheduled Auto Export	cleared	Media is not auto exported on a schedule.
Tape Type	varies	The tape type matching the latest generation of tape drive in the partition.
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Data Policy - Dual Copy on Tape		
Blobbing Enabled	selected	Allows an object to be broken into multiple blobs.
Minimize Spanning	cleared	Jobs larger than 1 TB are allowed to span across multiple tapes or pools as needed to maximize capacity utilization and performance.

Parameter	Value	Description
Default GET Job Priority	High	<p>When the storage domain receives a GET job, it is processed with high priority, which is before low and normal priority jobs. Jobs of different types are put in order based on priority.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a GET job by specifying a different priority in the Get Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default PUT Job Priority	Normal	<p>When the storage domain receives a PUT job, it is processed with normal priority, which is after high priority jobs but before low and normal priority jobs. Jobs of different types are put in order based on priority.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a PUT job by specifying a different priority in the Put Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default VERIFY Job Priority	Low	<p>When the storage domain receives a VERIFY job, it is processed with low priority, which is the lowest setting.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a VERIFY job by specifying a different priority in the Create Verify Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default Verify After Write	cleared	Data is not verified after a write.
Rebuild Priority	Low	If data is lost from tape media, the data is rebuilt using low priority, which is the lowest setting.
Checksum Type	MD5	Data using this storage domain is CRC checked using the MD5 checksum type. Data is CRC checked when it is written to cache with a PUT job, or read back from media with a GET job.
End-to-end CRC	No	This data policy does not use end-to-end CRC checking.

Parameter	Value	Description
Versioning	None	In order to upload a new version of an object already PUT to the gateway, the first version must be deleted.
Always Accept Replicated PUT Jobs	cleared	PUT jobs created for this data policy fail if one or more replication targets the gateway must PUT to are unavailable.
Data Persistence Rule for Tape First Copy		
Type	Permanent	Data is moved to tape and maintained on tape media until data is deleted from a bucket. Note: When data is deleted from a bucket, it is removed from tape media, but individual tape media cartridges are not reclaimed for use until all data is deleted from a tape cartridge.
Bucket Isolation Level	Standard	Data from different buckets can be mixed on to the same piece of media.
Data Persistence Rule for Tape Second Copy		
Type	Permanent	Data is moved to tape and maintained on tape media until data is deleted from a bucket. Note: When data is deleted from a bucket, it is removed from tape media, but individual tape media cartridges are not reclaimed for use until all data is deleted from a tape cartridge.
Bucket Isolation Level	Standard	Data from different buckets can be mixed on to the same piece of media.

Single Copy on Nearline Object Storage Disk

This data policy persists a single copy of each job sent to the gateway on to nearline object storage disk, which is provided by 77-bay, 96-bay, and 107-bay expansion nodes, or in the Gen2 S Series or Gen3 H Series master node. Once data is written on nearline storage disk, it is removed from the BlackPearl cache if the gateway detects that more cache space is needed for incoming data.

Note: This data policy is automatically created when the gateway detects a nearline storage disk partition.

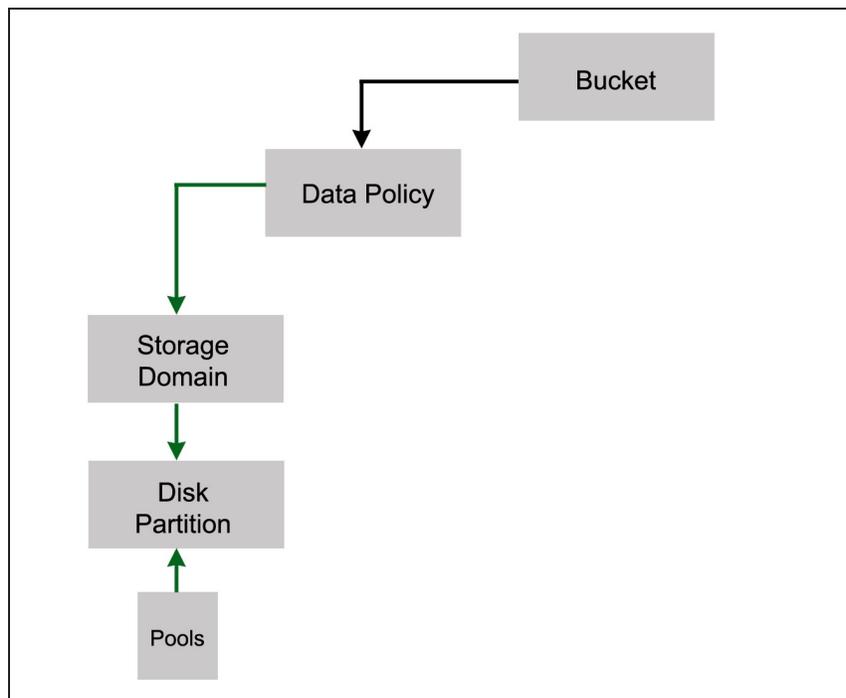


Figure 3 The Single Copy on Nearline Object Storage Disk workflow.

The single copy on nearline storage disk data policy is configured with the following attributes:

Parameter	Value	Description
Storage Domain - Pool First Copy - Data is written to the primary nearline storage domain. This domain is created automatically after you create a nearline object storage disk partition.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.
Write Optimization	Capacity	Job chunks are written across as few pools as possible.
LTFS File Naming	Object ID	File names use the format <i>{bucket name}/{object id}</i> , for example, bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82.
Media Export Allowed	selected	Media export is allowed.
Auto Export on Job Completion	cleared	Media is not auto exported upon job completion.
Auto Export on Job Cancel	cleared	Media is not auto exported upon job cancellation.
Auto Export on Media Full	cleared	Media is not auto exported upon media full.
Scheduled Auto Export	cleared	Media is not auto exported on a schedule.
Storage Domain Member for Pool First Copy - The first nearline disk partition created.		
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Data Policy - Single Copy on Nearline Disk		
Blobbing Enabled	selected	Allows an object to be broken into multiple blobs.

Parameter	Value	Description
Minimize Spanning	cleared	Jobs larger than 1 TB are allowed to span across multiple tapes or pools as needed to maximize capacity utilization and performance.
Default GET Job Priority	High	When the storage domain receives a GET job, it is processed with high priority, which is before low and normal priority jobs. Jobs of different types are put in order based on priority. Note: When using a DS3 client, this setting can be overridden when sending a GET job by specifying a different priority in the Get Job command. See the Spectra BlackPearl DS3 API Reference for more information.
Default PUT Job Priority	Normal	When the storage domain receives a PUT job, it is processed with normal priority, which is after high priority jobs but before low and normal priority jobs. Jobs of different types are put in order based on priority. Note: When using a DS3 client, this setting can be overridden when sending a PUT job by specifying a different priority in the Put Job command. See the Spectra BlackPearl DS3 API Reference for more information.
Default VERIFY Job Priority	Low	When the storage domain receives a VERIFY job, it is processed with low priority, which is the lowest setting. Note: When using a DS3 client, this setting can be overridden when sending a VERIFY job by specifying a different priority in the Create Verify Job command. See the Spectra BlackPearl DS3 API Reference for more information.
Default Verify After Write	cleared	Data is not verified after a write.
Rebuild Priority	Low	If data is lost from disk, the data is rebuilt using low priority, which is the lowest setting.
Checksum Type	MD5	Data using this storage domain is CRC checked using the MD5 checksum type. Data is CRC checked when it is written to cache with a PUT job, or read back from media with a GET job.

Parameter	Value	Description
End-to-end CRC	No	This data policy does not use end-to-end CRC checking.
Versioning	None	In order to upload a new version of an object already PUT to the gateway, the first version must be deleted.
Always Accept Replicated PUT Jobs	cleared	PUT jobs created for this data policy fail if one or more replication targets the gateway must PUT to are unavailable.
Type	Permanent	Data is moved to nearline disk and maintained on nearline disk until data is deleted from a bucket.
Bucket Isolation Level	Standard	Data from different buckets can be mixed on to the same piece of media.
Data Persistence Rule for Pool First Copy		
Type	Permanent	Data is moved to disk and maintained on disk media until data is deleted from a bucket.
Isolation Level	Standard	Data from different buckets can be mixed on to the same piece of media.

Single Copy on Nearline Object Storage Disk and Tape

This data policy persists a single copy of each job sent to the gateway on to both nearline storage disk and tape media. Nearline object storage disk is provided by 77-bay, 96-bay, and 107-bay expansion nodes, or in the Gen2 S Series or Gen3 H Series master node. Tape storage is provided by a Spectra Logic tape library.

Once data is written on both nearline object storage disk and tape, it is removed from the BlackPearl cache if the gateway detects that more cache space is needed for incoming data.

This configuration offers secure glacier storage on tape with the benefit of faster restores of data from nearline object storage disk.

Note: This data policy is automatically created when the gateway detects a nearline disk partition and a tape partition.

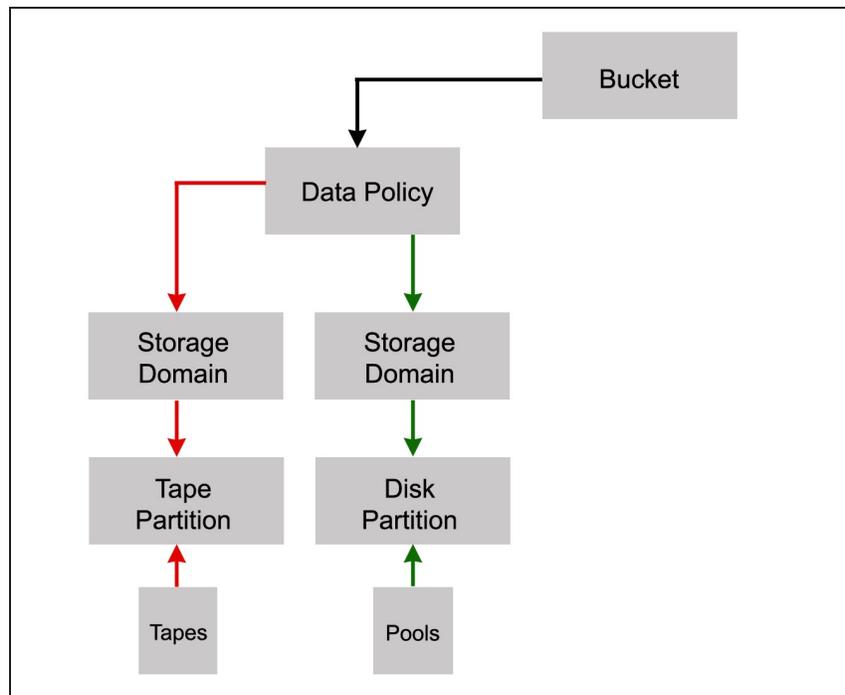


Figure 4 The Single Copy on Nearline Object Storage Disk and Tape workflow.

The single copy on nearline storage disk and tape data policy is configured with the following attributes:

Parameter	Value	Description
Storage Domain - Tape First Copy - Data is written to the primary tape storage domain. This domain is created automatically when the gateway detects the first tape partition created on the tape library.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.
Write Optimization	Capacity	Job chunks are written across as few pieces of media as possible. Fewer tape drives are used, so performance is lower.
LTFS File Naming	Object ID	File names use the format <code>{bucket name}/{object id}</code> , for example, <code>bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82</code> .
Media Export Allowed	cleared	Media export is not allowed.
Storage Domain Member for Tape First Copy		
Tape Type	varies	The tape type matching the latest generation of tape drive in the partition.
Auto Compaction Threshold	20	The percentage of a tape with deleted objects at which auto compaction is triggered. The default is 95. The minimum is 10.
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Storage Domain - Pool First Copy - Data is written to the primary nearline storage domain. This domain is created automatically after you create a nearline disk partition.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.

Parameter	Value	Description
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.
Write Optimization	Capacity	Job chunks are written across as few pools as possible.
LTFS File Naming	Object ID	File names use the format <i>{bucket name}/{object id}</i> , for example, bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82.
Media Export Allowed	selected	Media export is allowed.
Auto Export on Job Completion	cleared	Media is not auto exported upon job completion.
Auto Export on Job Cancel	cleared	Media is not auto exported upon job cancellation.
Auto Export on Media Full	cleared	Media is not auto exported upon media full.
Scheduled Auto Export	cleared	Media is not auto exported on a schedule.
Storage Domain Member for Pool First Copy - The first nearline disk partition created.		
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Data Policy - Single Copy on Nearline Disk and Tape		
Blobbing Enabled	selected	Allows an object to be broken into multiple blobs.
Minimize Spanning	cleared	Jobs larger than 1 TB are allowed to span across multiple tapes or pools as needed to maximize capacity utilization and performance.

Parameter	Value	Description
Default GET Job Priority	High	<p>When the storage domain receives a GET job, it is processed with high priority, which is before low and normal priority jobs. Jobs of different types are put in order based on priority.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a GET job by specifying a different priority in the Get Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default PUT Job Priority	Normal	<p>When the storage domain receives a PUT job, it is processed with normal priority, which is after high priority jobs but before low and normal priority jobs. Jobs of different types are put in order based on priority.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a PUT job by specifying a different priority in the Put Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default VERIFY Job Priority	Low	<p>When the storage domain receives a VERIFY job, it is processed with low priority, which is the lowest setting.</p> <p>Note: When using a DS3 client, this setting can be overridden when sending a VERIFY job by specifying a different priority in the Create Verify Job command. See the Spectra BlackPearl DS3 API Reference for more information.</p>
Default Verify After Write	cleared	Data is not verified after a write.
Rebuild Priority	Low	If data is lost from tape media, the data is rebuilt using low priority, which is the lowest setting.
Checksum Type	MD5	Data using this storage domain is CRC checked using the MD5 checksum type. Data is CRC checked when it is written to cache with a PUT job, or read back from media with a GET job.
End-to-end CRC	No	This data policy does not use end-to-end CRC checking.

Parameter	Value	Description
Versioning	None	In order to upload a new version of an object already PUT to the gateway, the first version must be deleted.
Always Accept Replicated PUT Jobs	cleared	PUT jobs created for this data policy fail if one or more replication targets the gateway must PUT to are unavailable.
Data Persistence Rule for Tape First Copy		
Type	Permanent	Data is moved to tape and maintained on tape media until data is deleted from a bucket. Note: When data is deleted from a bucket, it is removed from tape media, but individual tape media cartridges are not reclaimed for use until all data is deleted from a tape cartridge.
Bucket Isolation Level	Standard	Data from different buckets can be mixed on to the same piece of media.
Data Persistence Rule for Pool First Copy		
Type	Permanent	Data is moved to disk and maintained on disk media until data is deleted from a bucket.
Isolation Level	Standard	Data from different buckets can be mixed on to the same media.

Single Copy on Nearline Object Storage Disk and Dual Copy on Tape

This data policy persists a single copy of each job sent to the gateway on to nearline object storage disk, and two copies of the job on to tape media. Nearline storage is provided by 77-bay, 96-bay, and 107-bay expansion nodes, or in the Gen2 S Series or Gen3 H Series master node, while tape storage is provided by a Spectra Logic tape library.

Once data is written on both nearline object storage disk and tape, it is removed from the BlackPearl cache if the gateway detects that more cache space is needed for incoming data.

This configuration offers secure glacier storage on tape with the benefit of faster restores of data from nearline object storage disk.

This data policy is helpful if you want to export one copy of the bucket on tape media for storage off site, which provides enhanced data security in the case of ransomware attacks or disaster recovery.

Note: This data policy is automatically created when the gateway detects a nearline storage disk partition and two tape partitions.

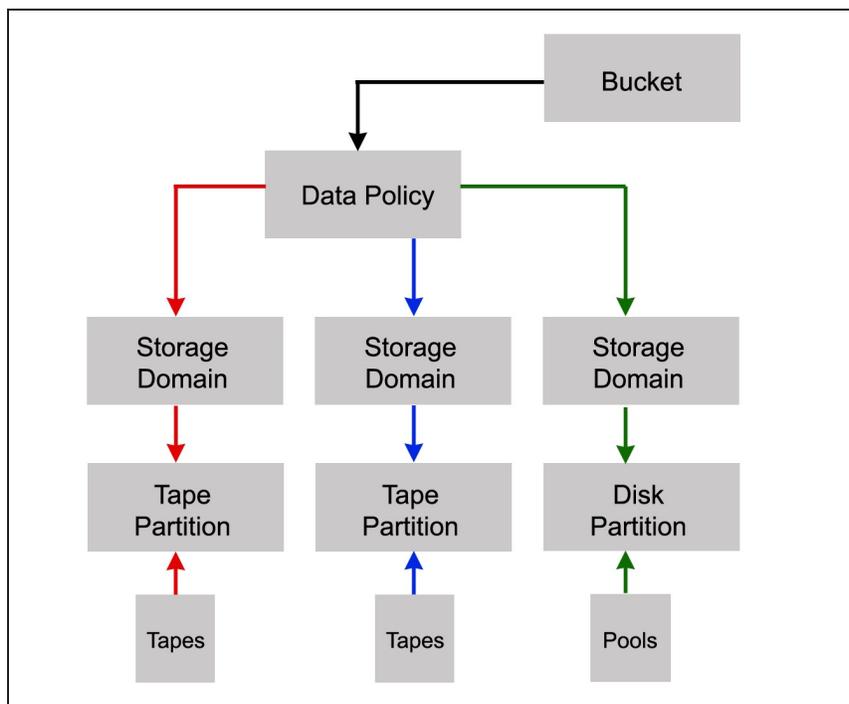


Figure 5 The Single Copy on Nearline Object Storage Disk and Dual Copy on Tape workflow.

The single copy on nearline storage disk and dual copy on tape data policy is configured with the following attributes:

Parameter	Value	Description
Storage Domain - Tape First Copy - Data is written to the primary tape storage domain. This domain is created automatically when the gateway detects the first tape partition created on the tape library.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.
Write Optimization	Capacity	Job chunks are written across as few pieces of media as possible. Fewer tape drives are used, so performance is lower.
LTFS File Naming	Object ID	File names use the format <code>{bucket name}/{object id}</code> , for example, <code>bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82</code> .
Media Export Allowed	cleared	Media export is not allowed.
Storage Domain Member for Tape First Copy		
Tape Type	varies	The tape type matching the latest generation of tape drive in the partition.
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Storage Domain - Tape Second Copy - A second copy of the data is written to a tape storage domain optimized for tape exports. This domain is created automatically when the gateway detects a tape partition created on the tape library.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.

Parameter	Value	Description
Write Optimization	Capacity	Job chunks are written across as few pieces of media as possible. Fewer tape drives are used, so performance is lower.
LTFS File Naming	Object ID	File names use the format <i>{bucket name}/{object id}</i> , for example, bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82.
Media Export Allowed	selected	Media export is allowed.
Auto Export on Job Completion	cleared	Media is not auto exported upon job completion.
Auto Export on Job Cancel	cleared	Media is not auto exported upon job cancellation.
Auto Export on Media Full	cleared	Media is not auto exported upon media full.
Scheduled Auto Export	cleared	Media is not auto exported on a schedule.
Storage Domain Member for Tape Second Copy		
Tape Type	<i>varies</i>	The tape type matching the latest generation of tape drive in the partition.
Auto Compaction Threshold	20	The percentage of a tape with deleted objects at which auto compaction is triggered. The default is 95. The minimum is 10.
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Storage Domain - Pool First Copy- Data is written to the primary nearline storage domain. This domain is created automatically after you create a nearline storage disk partition.		
Days to wait before verifying data	null	Data integrity verification is not performed automatically.
Secure Media Allocation	cleared	Media allocated to the storage domain may be reused by another storage domain if all data is deleted.

Parameter	Value	Description
Write Optimization	Capacity	Job chunks are written across as few pools as possible.
LTFS File Naming	Object ID	File names use the format <i>{bucket name}/{object id}</i> , for example, bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82.
Media Export Allowed	selected	Media export is allowed.
Auto Export on Job Completion	cleared	Media is not auto exported upon job completion.
Auto Export on Job Cancel	cleared	Media is not auto exported upon job cancellation.
Auto Export on Media Full	cleared	Media is not auto exported upon media full.
Scheduled Auto Export	cleared	Media is not auto exported on a schedule.
Storage Domain Member for Pool First Copy - The first nearline disk partition created.		
Write Preference	Normal	The gateway uses the partition after partitions with High write preference and before a partition with Low or Never Select write preference.
Data Policy - Single Copy on Nearline Disk and Dual Copy on Tape		
Blobbing Enabled	selected	Allows an object to be broken into multiple blobs.
Minimize Spanning	cleared	Jobs larger than 1 TB are allowed to span across multiple tapes or pools as needed to maximize capacity utilization and performance.
Default GET Job Priority	High	When the storage domain receives a GET job, it is processed with high priority, which is before low and normal priority jobs. Jobs of different types are put in order based on priority. Note: When using a DS3 client, this setting can be overridden when sending a GET job by specifying a different priority in the Get Job command. See the Spectra BlackPearl DS3 API Reference for more information.

Parameter	Value	Description
Default PUT Job Priority	Normal	When the storage domain receives a PUT job, it is processed with normal priority, which is after high priority jobs but before low and normal priority jobs. Jobs of different types are put in order based on priority. Note: When using a DS3 client, this setting can be overridden when sending a PUT job by specifying a different priority in the Put Job command. See the Spectra BlackPearl DS3 API Reference for more information.
Default VERIFY Job Priority	Low	When the storage domain receives a VERIFY job, it is processed with low priority, which is the lowest setting. Note: When using a DS3 client, this setting can be overridden when sending a VERIFY job by specifying a different priority in the Create Verify Job command. See the Spectra BlackPearl DS3 API Reference for more information.
Default Verify After Write	cleared	Data is not verified after a write.
Rebuild Priority	Low	If data is lost from tape media, the data is rebuilt using low priority, which is the lowest setting.
Checksum Type	MD5	Data using this storage domain is CRC checked using the MD5 checksum type. Data is CRC checked when it is written to cache with a PUT job, or read back from media with a GET job.
End-to-end CRC	No	This data policy does not use end-to-end CRC checking.
Versioning	None	In order to upload a new version of an object already PUT to the gateway, the first version must be deleted.
Always Accept Replicated PUT Jobs	cleared	PUT jobs created for this data policy fail if one or more replication targets the gateway must PUT to are unavailable.
Continue to the next page		

Parameter	Value	Description
Data Persistence Rule for Tape First Copy		
Type	Permanent	Data is moved to tape and maintained on tape media until data is deleted from a bucket. Note: When data is deleted from a bucket, it is removed from tape media, but individual tape media cartridges are not reclaimed for use until all data is deleted from a tape cartridge.
Bucket Isolation Level	Standard	Data from different buckets can be mixed on to the same piece of media.
Data Persistence Rule for Tape Second Copy		
Type	Permanent	Data is moved to tape and maintained on tape media until data is deleted from a bucket. Note: When data is deleted from a bucket, it is removed from tape media, but individual tape media cartridges are not reclaimed for use until all data is deleted from a tape cartridge.
Bucket Isolation Level	Standard	Data from different buckets can be mixed on to the same piece of media.
Data Persistence Rule for Pool First Copy		
Type	Permanent	Data is moved to disk and maintained on disk media until data is deleted from a bucket.
Isolation Level	Standard	Data from different buckets can be mixed on to the same media.

ADDITIONAL CONFIGURATIONS

In addition to the above standard configurations, many custom configurations can be created for BlackPearl Nearline gateway flash disks, spinning disk, and tape storage.

Contact Spectra Logic Professional Services for assistance. See [Contacting Spectra Logic](#) on page 7.

CHAPTER 2 - CONFIGURING ADVANCED BUCKET MANAGEMENT

This chapter provides instructions on how to configure Advanced Bucket Management features.



IMPORTANT

It is difficult and time consuming to change a data policy once the gateway writes data to a bucket using the data policy. Make sure that you understand the concepts in [Understanding Advanced Bucket Management on page 15](#) and have thoughtfully planned your data policies before you start using the BlackPearl gateway to store data.

Create a Disk Pool	55
Create a Nearline Storage Disk Pool	56
Create an Online Disk Pool	59
Create a Disk Partition	62
Create a Tape Partition	63
Create a Replication Target	64
Create a BlackPearl Target	64
Create an Amazon S3 Target	68
Create a Microsoft Azure Target	72
Create a Storage Domain	75
Add a Storage Domain Member to a Storage Domain	81
Create a Data Policy	83
Add Data Persistence Rules and Replication Rules to a Data Policy	86
Create Data Policy ACLs	92
Create a Bucket	95
Create a New Bucket ACL for a Group	96
Create a New Bucket ACL for a User	97
Transfer Data	99

CREATE A DISK POOL

A disk pool groups a set of physical drives together to create a virtual drive that the operating system treats as a single physical drive. There are two types of disk pools:

- **Nearline Storage Disk Pool** - If all drives in the pool are capable of setting an idle timer, Nearline storage disk pools can be configured to spin down after 60 minutes without I/O, for power savings.
- **Online Storage Disk Pool** - Online storage disk pools remain powered on at all times for fast access to data.

Online and Nearline Storage disk pools use compression. This allows the BlackPearl gateway to store more data.

If desired, select the check box to enable data compression with ZFS to allow the BlackPearl gateway to store more data. If the data being written is compressible there is typically an increase in store and restore operations, because less data is transferred to and from the disk drives. The size reduction of transferred files depends on how much the system can compress the data, and may fluctuate.

The data compression process uses CPU cycles to perform the compression. If compression is enabled for non-compressible data, for example JPEG images or movie files that use the H. 264 codec, the compression process may use an excessive number of CPU cycles, slowing the overall performance of the gateway. This impact is less evident with Gen2 and Gen 3 master nodes.

Note: When viewing the details of an online or nearline storage disk pool, the user interface displays the physically used space on the pool, not the logically used space.

Once a disk pool is created, it can be added to a disk partition.

If your BlackPearl gateway does not include disk storage, continue with [Create a Tape Partition](#) on page 63.

Create a Nearline Storage Disk Pool

Use the instructions in this section to create a nearline storage disk pool.

If you do not want to configure a nearline disk pool, continue with [Create an Online Disk Pool](#) on page 59.

Note: Nearline pools created on a 96-bay expansion node have a hard coded capacity utilization limit percentage. On gateways running BlackPearl OS 5.2 or later, this percentage is 95%. On gateways running BlackPearl OS 5.1.x or older, the capacity limit percentage is 87%.

- From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Policy Management** to display the Advanced Bucket Management screen.

Understanding Advanced Bucket Management

```

    graph TD
      Bucket[Bucket] --> DP[Data Policy]
      Bucket --> SD1[Storage Domain]
      Bucket --> SD2[Storage Domain]
      SD1 --> TP1[Tape Partition]
      SD1 --> DP1[Disk Partition]
      SD2 --> TP2[Tape Partition]
      SD2 --> DP2[Disk Partition]
      DP1 --> SP1[Storage Pool]
      DP1 --> SP2[Storage Pool]
      DP2 --> SP3[Storage Pool]
      DP2 --> SP4[Storage Pool]
    
```

System Pools

Name	Health	Raw Size	Available	Used	Overhead	Fault Tolerance
BlackPearl_Cac...	Online	72.8 TB	55.2 TB	759.2 GB	16.8 TB	Can sustain two drive failures per array.
BlackPearl_Dat...	Online	1.5 TB	720.7 GB	81.5 MB	769.7 GB	Can sustain one drive failure per mirror.

Double click on a row to see more detailed information

Storage Pools

Name	Health	S3 Pool Type	Partition
There are currently no items to display			

Tape Partitions

Name	State	Standby
912F005F29	Online	No
A02F002D0A	Online	No

Double click on a row to see more detailed information

Disk Partitions

Name	Type	Members
There are currently no items to display		

Storage Domains

Name	Write Optimization	Members
Db Backup First Copy	Performance	912F005F29 (LTO-6), 912F005F29 (LTO-5), A02F002D0A (LTO-7)
Db Backup Second Copy	Capacity	912F005F29 (LTO-6), A02F002D0A (LTO-7), 912F005F29 (LTO-5)
secure-capacity	Capacity	A02F002D0A (LTO-7)
Tape First Copy	Performance	A02F002D0A (LTO-7), 912F005F29 (LTO-5), 912F005F29 (LTO-6)
Tape Second Copy	Capacity	A02F002D0A (LTO-7), 912F005F29 (LTO-6), 912F005F29 (LTO-5)
test-performance-normal	Performance	A02F002D0A (LTO-7)
test-performance-secure	Performance	A02F002D0A (LTO-7)

Double click on a row to see more detailed information

Data Policies

Name	Default Get Job Priority	Default Put Job Priority	Default Verify Job Priority	Rebuild Priority	Checksum Type
cloud-policy	High	Normal	Low	Low	MD5
Database Backup	High	Normal	Low	Low	MD5
Dual Copy on Tape	High	Normal	Low	Low	MD5
Single Copy on Tape	High	Normal	Low	Low	MD5
test-capacity	High	Normal	Low	Low	MD5
test-capacity-secure-media	High	Normal	Low	Low	MD5
test-performance-normal	High	Normal	Low	Low	MD5
test-performance-secure	High	Normal	Low	Low	MD5

Double click on a row to see more detailed information

Hardware

Figure 6 The Advanced Bucket Management screen.

2. Select **Action > New Nearline Disk Pool**. The New Nearline Disk Pool dialog box displays.

Note: The **Storage Pool Preview** pane does not display until you have selected the disks you want to use in the disk pool.

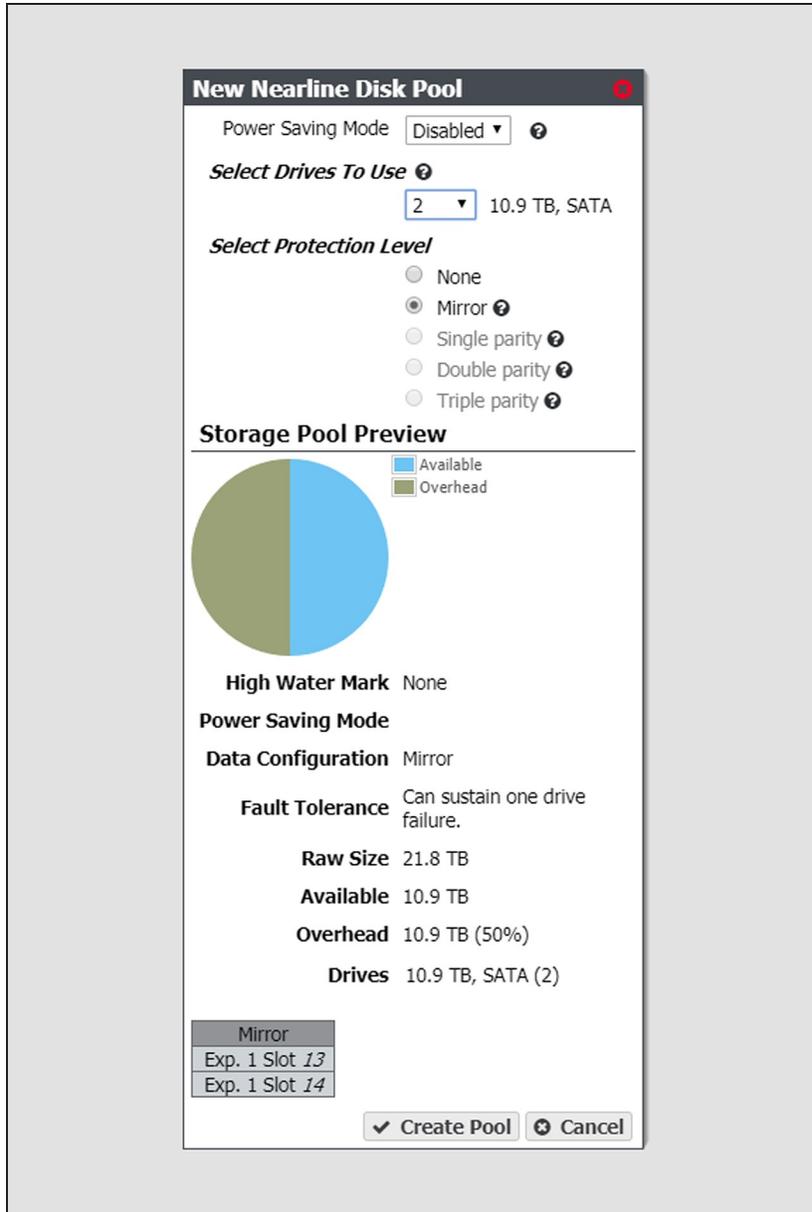


Figure 7 The New Nearline Disk Pool dialog box.

- Configure the disk pool as required for your environment. As you make changes, the screen updates to show the characteristics of the new pool.

For this option....	Do the following...
<p>Power Saving Mode</p>	<p>Using the drop-down menu, select the desired Power Saving Mode. Enabling the power saving mode sets the standby timer to 60 minutes for all drives in the pool, but only if all drives in the pool are capable of using a standby timer. When the disk pool is idle for 60 minutes, the drives spin-down to conserve power.</p> <p>Note: To use this feature, all drives in the storage pool must be power-saving compatible.</p>
<p>Select Drives To Use</p>	<p>Use the drop-down menu to select the number of drives to include in the pool. If your gateway contains more than one type of disk drive, multiple drop-down menus are present, but only one type can be assigned to a pool.</p> <p>Any drive not in a disk pool acts as a global spare. A global spare drive is activated as soon as a drive configured in a disk pool fails.</p>
<p>Select Protection Level</p>	<p>Use the radio buttons to select the protection level for the pool. Only one option can be selected. Use the Storage Pool Preview information to compare the fault tolerance and required overhead for each configuration.</p> <p>None—The pool is not configured to provide data protection. Any drive failure results in data loss.</p> <p>Note: Spectra Logic does not recommend setting protection to None.</p> <p>Mirror—Data is striped across two mirrors. Any detected data corruption is corrected using checksums. This type of RAID offers the best performance for small random reads and writes.</p> <p>Single parity—Data is striped across multiple single-parity arrays, which can tolerate one drive failure without data loss. This type of RAID has faster performance than double- and triple-parity based RAIDs.</p> <p>Double parity—Data is striped across multiple double-parity arrays, which can tolerate two drive failures without data loss. In most cases, double-parity provides the best balance between data protection, performance, and storage capacity.</p> <p>Triple parity—Data is striped across multiple triple-parity arrays, which can tolerate three drive failures without data loss. This type of RAID provides the most data protection.</p>

- Click **Create Pool**. The new nearline disk pool is listed on the Advanced Bucket Management screen.

Create an Online Disk Pool

Use the instructions in this section to create an online disk pool.

If you do not want to create an online disk pool, continue with [Create a Disk Partition](#) on page 62.

Use the instructions in this section to create an online disk pool.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Policy Management** to display the Advanced Bucket Management screen (see [Figure 6](#) on page 57).
2. Select **Action > New Online Disk Pool**. The New Online Disk Pool dialog box displays.

Note: The **Storage Pool Preview** pane does not display until you have selected the disks you want to use in the disk pool.

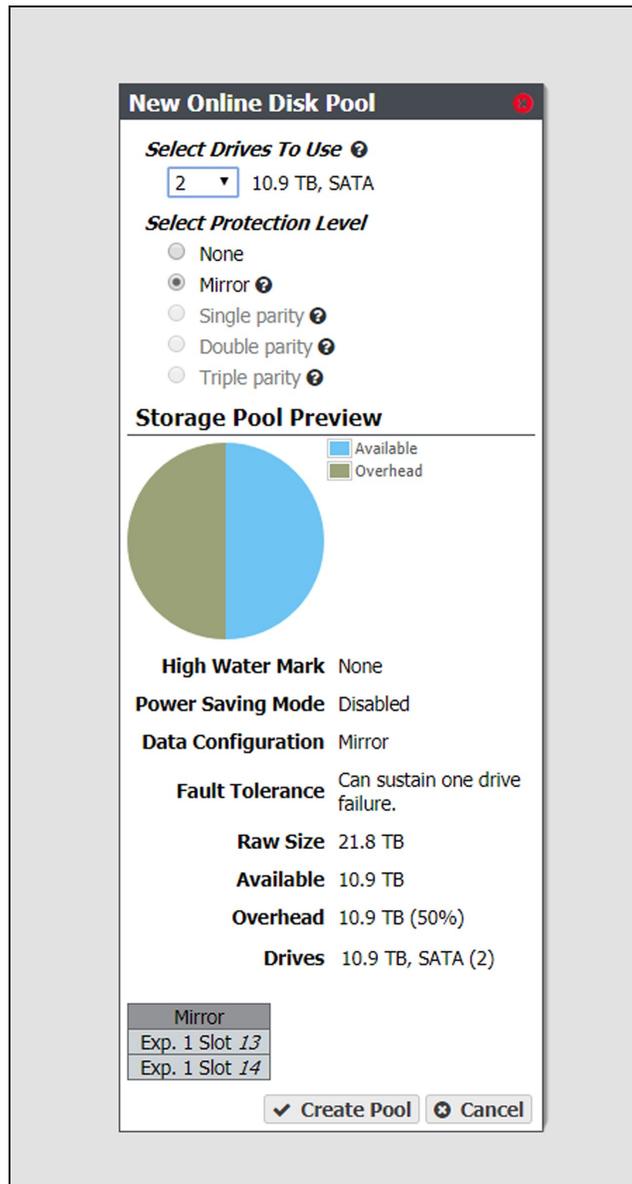


Figure 8 The New Online Disk Pool dialog box.

3. Configure the disk pool as required for your environment. As you make changes, the screen updates to show the characteristics of the new pool.

For this option....	Do the following...
<p>Select Drives To Use</p>	<p>Use the drop-down menu to select the number of drives to include in the pool. If your gateway contains more than one type of disk drive, multiple drop-down menus are present, but only one type can be assigned to a pool.</p> <p>Any drive not in a disk pool acts as a global spare. A global spare drive is activated as soon as a drive configured in a disk pool fails.</p>
<p>Select Protection Level</p>	<p>Use the radio buttons to select the protection level for the pool. Only one option can be selected. Use the Storage Pool Preview information to compare the fault tolerance and required overhead for each configuration.</p> <p>None—The pool is not configured to provide data protection. Any drive failure results in data loss.</p> <p>Mirror—Data is striped across two mirrors. Any detected data corruption is corrected using checksums. This type of RAID offers the best performance for small random reads and writes.</p> <p>Single parity—Data is striped across multiple single-parity arrays, which can tolerate one drive failure without data loss. This type of RAID has faster performance than double- and triple-parity based RAIDs.</p> <p>Double parity—Data is striped across multiple double-parity arrays, which can tolerate two drive failures without data loss. In most cases, double-parity provides the best balance between data protection, performance, and storage capacity.</p> <p>Triple parity—Data is striped across multiple triple-parity arrays, which can tolerate three drive failures without data loss. This type of RAID provides the most data protection.</p>

4. Click **Create Pool**. The new online disk pool is listed on the Advanced Bucket Management screen.

CREATE A DISK PARTITION

Disk partitions are collections of one or more disk pools. Disk partitions are specified in storage domains as storage targets.

Use the instructions in this section to create a new disk partition.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Policy Management**. The Advanced Bucket Management screen displays (see Figure 6 on page 57).
2. Select **Action > New Disk Partition**. The New Disk Partition dialog box displays.

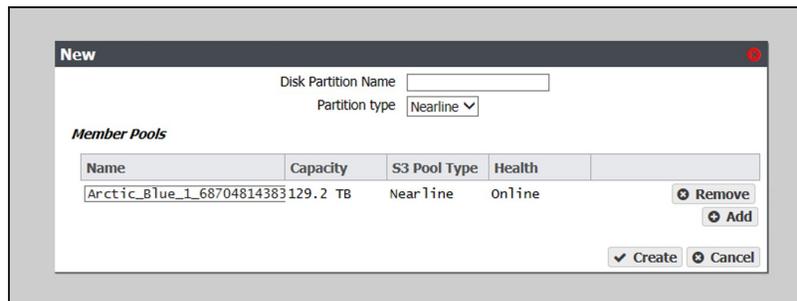


Figure 9 The New Disk Partition dialog box.

3. Enter a name for the disk partition in the **Disk Partition Name** field.
4. Use the drop-down menu to select the **Partition type**. You cannot mix different types of disk pools in a disk partition.
 - Select **Online** to use a disk pool that is always powered on and available for access.
 - Select **Nearline** to use a disk pool that can be configured to spin down after 60 minutes without I/O, for power savings.
5. Add a disk pool to the disk partition.
 - a. In the Member Pools pane, click **Add**. A new row appears in the pane.
 - b. Use the **Name** drop down menu to select a disk pool from the list of previously configured disk pools. The Capacity, Type, and Health of the disk pool display.

Note: It may take up to 1 minute after creating an online or nearline disk pool before it displays in the Member Pools list.

 - c. If desired, repeat **Step a** and **Step b** to add additional disk pools to the disk partition.
6. Click **Create**. The new disk partition displays on the Advanced Bucket Management screen.

CREATE A TAPE PARTITION

Use the *Tape Library User Guides on page 14* for your Spectra Logic or other supported tape library to create a partition. Once the BlackPearl gateway detects a partition on a tape library connected to it, the tape partition is automatically listed on the Advanced Bucket Management screen.

If your BlackPearl gateway does not have a tape library, continue with [Create a Replication Target](#) below.

Note: If the BlackPearl gateway is configured to transfer data to tape, make sure the partitions configured for use by the BlackPearl gateway do not contain WORM (Write Once-Read Many) media. The BlackPearl gateway is not compatible with WORM media.

CREATE A REPLICATION TARGET

Replication targets allow you to configure the BlackPearl gateway to automatically replicate data to another BlackPearl gateway, or to the Azure or Amazon S3 clouds.

If your BlackPearl gateway does not include the feature to replicate data to cloud targets, continue with [Create a Storage Domain on page 75](#).

Note: The instructions below describe configuring a target that is later associated with a data policy. For instructions on creating NAS replication, see .

Create a BlackPearl Target

Configuring a BlackPearl target allows a data policy on one BlackPearl gateway to replicate data to a second gateway. If data is sent to a data policy that is not configured for replication, the data is not replicated to the target gateway.

With replication enabled, as soon as data is PUT to the cache of the source gateway it begins replicating to the target gateway. Storing multiple copies of the same data on different BlackPearl gateways provides enhanced data security and disaster recovery if the source gateway fails. When you delete data from the source gateway, you can optionally specify to have the data deleted from the target gateway as well.

**IMPORTANT**

Spectra Logic recommends using the same versioning settings on both the source and target BlackPearl gateways.

Note: If the source BlackPearl gateway uses object versioning but the target BlackPearl gateway does not, when an object is deleted on the source gateway, the delete is replicated to the target gateway. However, when IOM validates the data on the two gateways, it detects that the object still exists on the source gateway, and self-heals the object on the target gateway again.

Use the instructions in this section to configure a BlackPearl target.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen.

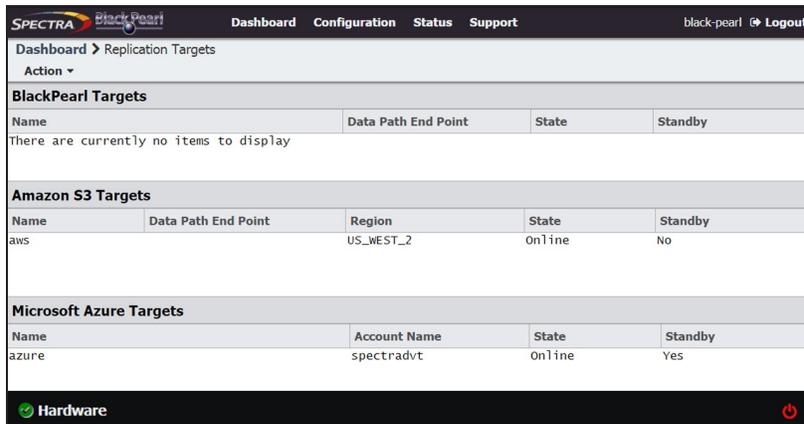


Figure 10 The Replication Targets screen

2. Select **Action > New BlackPearl Target**. The New BlackPearl Target dialog box displays.

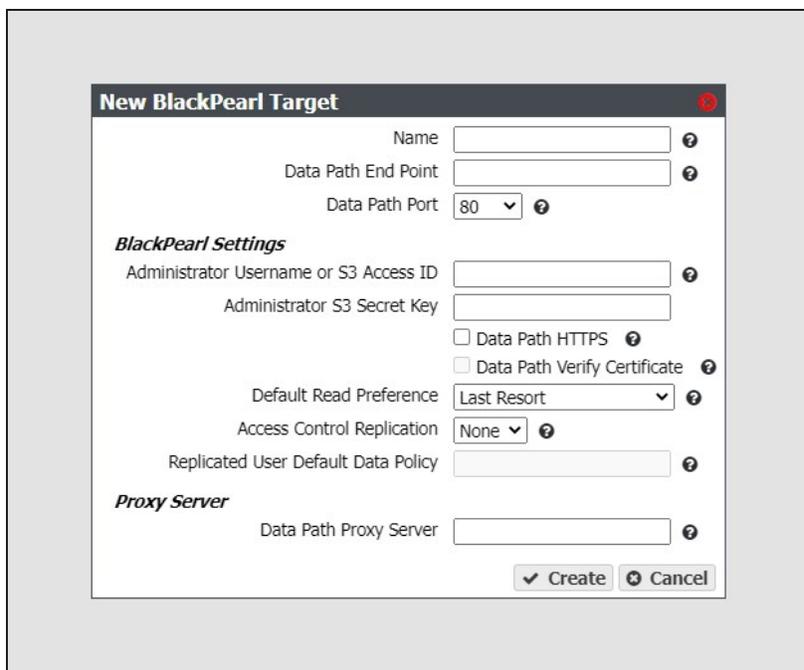


Figure 11 The New BlackPearl Target dialog box.

3. Enter a name for the BlackPearl target in the **Name** field.
4. Enter the system name of the target gateway, or the IP address of the target gateway’s data port, as the **Data Path End Point**.

Note: Do not use the IP address of the target gateway’s management port.

5. Using the drop-down menu, select a value for the **Data Path Port**. Set this to the value of the port on which the target gateway's S3 service is running.

Note: Port selections for secure transfer (443/8443) are grayed-out and not able to be selected until you select Data Path HTTPS in [Step 8](#) below.

6. Enter the username or S3 Access ID of a user with administrator privileges on the target gateway in the **Administrator Username or S3 Access ID** field.

Note: Administrator credentials are used to configure and maintain the source/target relationship. They are not used for user driven replication operations.

7. In the **Administrator S3 Secret Key** field, enter the S3 Secret Key of the user you entered in [Step 6](#).

8. Select **Data Path HTTPS** to enable secure data transfer with the BlackPearl target. When this option is selected, the Data Path Port setting automatically changes to 443. Repeat [Step 5](#) if you want to change the data path port to 8443.



IMPORTANT

Using HTTPS for data transfer greatly impacts data transfer speed. Spectra Logic recommends leaving this disabled if it is not required for your data storage environment.

9. If you enabled Data Path HTTPS, you can optionally select **Data Path Verify Certificate** to verify the SSL certificate of the BlackPearl target. This option is not available if you did not enable Data Path HTTPS.

Note: Do not enable this option if the BlackPearl target uses the default self-signed SSL certificate.

10. Using the drop-down menu, select a value for the **Default Read Preference**. Data is normally read from the source gateway whenever possible. This setting determines from what location data is read back from the target gateway, if needed.

Name	Description
Last Resort	The source gateway only reads data from the target gateway if the source gateway cannot read from any of its own data partitions.
Minimum Latency	The source gateway reads the data from the data partition with the least latency no matter whether it is connected to the source gateway or the target gateway. For example, if the source gateway only has the data on tape and the target gateway has the data on pool, the data is read from the target pool. Note: Only use MINIMUM LATENCY when the network between the source and target is very inexpensive.
After Online Pool	The source gateway only reads data from the target gateway if the source gateway cannot read from an online pool.
After Nearline Pool	The source gateway only reads data from the target gateway if the source gateway cannot read from a nearline pool.
After Non-Exportable Tape	The source gateway only reads data from the target gateway if the source gateway cannot read from secure media.
Never	Data is never read from the target gateway.

11. Using the drop-down menu, select a value for **Access Control Replication**.

Name	Description
None	No access control information is replicated to the BlackPearl target. Note: The Administrator secret key on both the source and target BlackPearl gateways must be identical when setting Access Control Replication to None .
Users	User creation, modification, and deletion is replicated to the BlackPearl target.

12. If you selected Users in [Step 11](#), you can optionally enter the name of a data policy previously configured on the target gateway to use as the **Replicated User Default Data Policy**. If configured, the gateway uses this target data policy as the default data policy for any users replicated to the target.
13. Optionally, enter the IP address of the **Data Path Proxy Server**. If configured, the source gateway uses the specified proxy to connect to the target gateway.
14. Click **Create**. The new BlackPearl target appears on the Advanced Bucket Management screen.

Create an Amazon S3 Target

Configuring an Amazon S3 target allows a data policy on the BlackPearl gateway to replicate data to the Amazon S3 cloud. With replication enabled, as soon as data is PUT to the cache of the source gateway it begins replication to the Amazon S3 cloud.

Note: Only Amazon Web Services (AWS) S3 is qualified as an Amazon S3 target. Other S3 services have not been tested.

Restrictions

The following restrictions apply to creating an Amazon S3 target:

- You cannot create two Amazon S3 targets using the same Data Path End Point and Access Key.
- You cannot create two Amazon S3 targets using the same Region and Access Key when the Data Path End Point has no value.
- You cannot link multiple Amazon S3 targets to the same Data Policy when both targets have no value for the Data Path End Point, and the prefix and suffix are the same for both targets.

Use the instructions in this section to configure an Amazon S3 target.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen (see [Figure 10 on page 65](#)).

2. Select **Action > New Amazon S3 Target**. The New Amazon S3 Target dialog box displays.

New Amazon S3 Target

S3 File Naming Object ID ?
 Object Name ?

Name ?

Data Path End Point ?

Region

HTTPS
 Restricted Access ?

Amazon S3 Settings

Access Key

Secret Key

Cloud Bucket Prefix

Cloud Bucket Suffix

Staged Data Expiration Days ?

BlackPearl Settings

Default Read Preference ?

Proxy Server

Proxy Domain

Proxy Host

Proxy Port

Proxy Username

Proxy Password

Figure 12 The New Amazon S3 Target dialog box.

3. Select the type of **S3 File Naming** to use for the target.

- **Object ID** - objects display a UUID when viewed on the Amazon target.
- **Object Name** - objects display their name when viewed on the Amazon target.

4. Enter a name for the Amazon S3 target in the **Name** field.

- Enter a **Data Path End Point** (system name or IP address) or a **Region** to identify the remote Amazon S3 target.

Acceptable regions are: Note: Dashes (-) in the standard AWS S3 region code must be replaced by underscores (_) in the text entered in the Region field.		
<ul style="list-style-type: none"> • us_east_1 • us_east_2 • us_west_1 • us_west_2 • eu_west_1 • eu_west_2 • eu_central_1 	<ul style="list-style-type: none"> • ap_south_1 • ap_southeast_1 • ap_southeast_2 • ap_northeast_1 • ap_northeast_2 	<ul style="list-style-type: none"> • sa_east_1 • cn_north_1 • ca_central_1 • gov_cloud

- Notes:**
- If you enter both a **Data Path End Point** and a **Region**, the gateway uses the **Data Path End Point** and ignores the **Region**.
 - You cannot use the same **Data Path End Point** or **Region** for multiple Amazon S3 targets.

- By default, **HTTPS** is selected so that the replication uses a secure connection. If desired, clear **HTTPS** to use HTTP.

- If desired, select **Restricted Access** to limit access to a specific set of credentials and buckets set in your Amazon S3 account. This setting removes the verification the BlackPearl gateway uses to confirm valid credentials when a data path endpoint and region are entered in Step 5.

Note: Spectra Logic recommends against using **Restricted Access**.

- Enter the S3 Access Key of a user with administrator privileges for the Amazon S3 account in the **Access Key** field.

Note: Administrator credentials are used to configure and maintain the source/target relationship. They are not used for user driven replication operations.

- In the **Secret Key** field, enter the S3 Secret Key of the user you entered in Step 8.

- Optionally, enter a **Cloud Bucket Prefix** and/or **Cloud Bucket Suffix**. Bucket names on the BlackPearl gateway must be unique within the gateway, but bucket names in AWS S3 must be unique across the world. To permit friendlier, shorter local bucket names on the BlackPearl gateway while avoiding naming conflicts with AWS S3, the gateway adds the defined **Cloud Bucket Prefix** and **Cloud Bucket Suffix** to the BlackPearl bucket name when it replicates the bucket. For example, if **Cloud Bucket Prefix**=`prefix`, **Cloud Bucket Suffix**=`suffix`, and the bucket name=`name`, the resulting name of the bucket on the Amazon S3 target is `prefix-name-suffix`.

Note: The prefix and/or suffix must adhere to the replication target naming requirements.

11. Enter a **Staged Data Expiration** time in days using any value between 1 and 365. The default is 30. When data is pre-staged by the S3 service so that the BlackPearl gateway can retrieve the data in an S3-standard manner, you must specify an expiration period in days. This is the minimum number of days before the pre-staged copy expires. If the gateway does not retrieve all of the data before the copy expires, it has to pre-stage the data again, incurring additional delays and costs.

Note: Spectra Logic strongly discourages configuring a **Staged Data Expiration** of less than 7 days as any potential cost savings are offset by the possibility of multiple stagings.

12. Using the drop-down menu, select a value for the **Default Read Preference**. Data is normally read from the source gateway whenever possible. This setting determines when data is read back from the Amazon S3 target, if needed.

Note: Spectra Logic recommends that **Default Read Preference** be kept at the default of **Last Resort**.

Name	Description
Last Resort	The source gateway only reads data from the target if the source gateway cannot read from any of its own data partitions.
Minimum Latency	The source gateway reads the data from the data partition with the least latency no matter whether it is connected to the source gateway or the target.
After Online Pool	The source gateway only reads data from the target if the source gateway cannot read from an online pool.
After Nearline Pool	The source gateway only reads data from the target if the source gateway cannot read from a nearline pool.
After Non-Exportable Tape	The source gateway only reads data from the target gateway if the source gateway cannot read from secure media.
Never	Data is never read from the target.

13. Optionally, enter the information for a proxy server:

Field	Description
Proxy Domain	Domain name for the proxy server.
Proxy Host	The host name or IP address for the proxy server through which the gateway connects.
Proxy Port	The proxy server port through which the gateway connects.
Proxy Username	The username used when connecting through the proxy server.
Proxy Password	The password used when connecting through the proxy server.

14. Click **Create**. The new Amazon S3 target appears on the Replication Targets screen.

Create a Microsoft Azure Target

Configuring a Microsoft Azure target allows a data policy on the BlackPearl gateway to replicate data to the Microsoft Azure cloud. With replication enabled, as soon as data is PUT to the cache of the source gateway it begins replication to the Microsoft Azure cloud.

Use the instructions in this section to configure a Microsoft Azure target.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen(see Figure 10 on page 65).

2. Select **Action > New Microsoft Azure Target**. The New Microsoft Azure Target dialog box displays.

Figure 13 The New Microsoft Azure Target dialog box.

3. Enter a name for the Microsoft Azure target in the **Name** field.

Note: Each Azure target name must be unique. You cannot create two Azure targets with the same name.
4. By default, **HTTPS** is selected so that the replication uses a secure connection. If desired, clear the **HTTPS** check box to use HTTP.
5. Enter the account name for the Microsoft Azure account in the in the **Account Name** field.

Note: You can not use the same **Account Name** for multiple Microsoft Azure targets.
6. In the **Account Key** field, enter the account key associated with the account entered in Step 5.
7. Optionally, enter a **Cloud Bucket Prefix** and/or **Cloud Bucket Suffix**. Bucket names on the BlackPearl gateway must be unique within the gateway, but bucket names in Microsoft Azure must be unique across the world. To permit friendlier, shorter local bucket names on the BlackPearl gateway while avoiding naming conflicts with Microsoft Azure, the gateway adds the defined **Cloud Bucket Prefix** and **Cloud Bucket Suffix** to the BlackPearl bucket name when it replicates the bucket. For example, if **Cloud Bucket Prefix**=`prefix`, **Cloud Bucket Suffix**=`suffix`, and the bucket name=`name`, the resulting name of the bucket on the Azure target is `prefix-name-suffix`.

Note: The prefix and/or suffix must adhere to the replication target naming requirements.

8. Using the drop-down menu, select a value for the **Default Read Preference**. Data is normally read from the source gateway whenever possible. This setting determines when data is read back from the Microsoft Azure target, if needed.

Note: Spectra Logic recommends that **Default Read Preference** be kept at the default of **Last Resort**.

Name	Description
Last Resort	The source gateway only reads data from the target if the source gateway cannot read from any of its own data partitions.
Minimum Latency	The source gateway reads the data from the data partition with the least latency no matter whether it is connected to the source gateway or the target.
After Online Pool	The source gateway only reads data from the target if the source gateway cannot read from an online pool.
After Nearline Pool	The source gateway only reads data from the target if the source gateway cannot read from a nearline pool.
After Non-Exportable Tape	The source gateway only reads data from the target gateway if the source gateway cannot read from secure media.
Never	Data is never read from the target.

9. Click **Create**. The new Microsoft Azure target appears on the Replication Targets screen.

CREATE A STORAGE DOMAIN

A storage domain is a named collection of member data partitions and, when applicable, media type combinations. Storage domains define the possible places where the BlackPearl Nearline Gateway stores data that is sent to it. Data persistence rules and data policies further define where and for how long to store specific data.

Entire data partition/media type combinations are members of storage domains. When a bucket requires additional capacity, a single disk partition or tape is allocated out of the members to fulfill the capacity requirement.

Use the instructions in this section to create a new storage domain.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Select **Action > New Storage Domain**. The New Storage Domain dialog box displays.

The dialog box titled "New" contains the following fields and options:

- Storage Domain Name:
- Days to wait before verifying data: ?
- Secure Media Allocation ?
- Write Optimization**
 - Capacity ?
 - Performance ?
- LTFS File Naming**
 - Object ID ?
 - Object Name ?
- Media Export Allowed**
- Auto Export on Job Completion
- Auto Export on Job Cancel
- Auto Export on Media Full ?
- Export Media with an Available Capacity of Only: GIB ▾
- Auto Export Verify Task Priority: ▾
- Scheduled Auto Export
 - Hourly
 - Daily
 - Weekly
 - Start Time: e.g. 3:00 AM
 - Every days ?

Buttons:

Figure 14 The New Storage Domain dialog box.

3. Enter a name for the storage domain in the **Storage Domain Name** field.

4. Enter a value for **Days to wait before verifying data**. The gateway automatically performs a data integrity verification for all tape media in the storage domain that are unchanged after the specified number of days pass, to ensure the data written to the tape cartridge is still viable. If null, data integrity verification is not performed automatically.

- Notes:**
- By default, all data on the tape is verified. You can customize the amount of data to be verified in .
 - When this verification completes, the **Last Verified** field on the tape details screen is updated.
 - While the verification is in progress, client access has priority over the data integrity verification.
 - You can also initiate data integrity verification for tape media manually. See for more information.
 - Disk pools are not subject to automatic data integrity verification. However, you can initiate data integrity verification for disk pools manually. See for more information.

5. Select or clear **Secure Media Allocation**. If enabled, Secure Media Allocation ensures that media allocated to the storage domain always remains in the storage domain. Even if all data on the media is deleted, the media will not be reallocated to another storage domain.

Note: Secure Media Allocation should only be enabled when, for compliance purposes, the user must be certain which media ever contained any data for the storage domain (usually, to physically destroy the media once the data is no longer needed), or to force rotating through media when new backups are created and old backups are deleted.

6. Select the **Write Optimization** for the storage domain. This setting specifies whether job chunks are written as quickly as possible or across as few pieces of media as possible.

The BlackPearl gateway writes to tape drives based on chunks, with default chunk size of approximately 128 GB, or 2% of the tape media capacity. When there is a queue of jobs, the BlackPearl gateway aggregates smaller jobs or smaller chunks into a size of approximately 128 GB for each tape drive read or write task.

When running in **Capacity** mode, the BlackPearl gateway uses as few tape cartridges or disk pools as possible. The gateway only allocates a new tape cartridge or disk pool when capacity is needed.

When running in **Performance** mode, the BlackPearl gateway spreads the chunks or aggregations across all available tape drives, or disk pools. The number of tape drives used can be limited by using tape drive reservations.

- The consequence of using performance mode with tape media is that during a restore or GET job, more tape drives and tapes cartridges are required to restore a data set that was initially spread across many tapes. This can drastically reduce overall performance during restores, as the gateway takes longer to get access to the full data set.

**IMPORTANT**

Spectra Logic highly recommends working with Spectra Logic Technical Support, or Spectra Professional Services before using Performance mode.

For more information on capacity and performance modes, see .

Note: If the storage domain is assigned to a data policy and "Minimize Spanning" is enabled for the data policy, it overrides the capacity mode and performance mode logic for a given job, up to 1 TB in size. Minimize spanning increases the chunk size to 1 TB, and always keeps one chunk on a single tape regardless of write mode. If the job is larger than 1 TB, then multiple chunks are used and the gateway uses the logic for capacity mode vs performance mode, where different chunks may transfer to different tapes.

7. Select the **LTFS File Name** option for the storage domain.

Note: This setting only applies to tape media. If the storage domain includes tape partition(s), you must specify the **LTFS File Name** option for the storage domain. This option specifies how the gateway names the file when it writes them to tape.

There are two options for the **LTFS File Name**:

- **Object Name** — LTFS file names use the format *{bucket name}/{object name}*, for example *bucket1/video1.mov*. Object names must comply with LTFS file naming rules. If the tapes are exported from the BlackPearl gateway and loaded into a non-BlackPearl tape partition, the file names match the object names.

**IMPORTANT**

If you select **Object Name**, you cannot assign this storage domain to a data policy that uses versioning.

- Notes:**
- The colon character (:) is not allowed in LTFS file names and therefore not allowed in BlackPearl object names.
 - The slash character (/) is not allowed in LTFS file names; however, the BlackPearl software can accommodate a slash in the object name and translates it as a directory in the LTFS file system (e.g. *directory1/directory2/video1.mov*).
 - File names with multiple consecutive slash characters (//) are not allowed.
 - Directory names have a limit of 255 characters.
 - File names have a variable character limit. If you are using English ASCII characters, the limit is 1024 characters. If you are using a graphical language, such as Japanese, the limit is 512 characters.

- Spectra Logic does not recommend the following characters in LTFS file names or BlackPearl object names for reasons of cross-platform compatibility:
 - Asterisk (*)
 - Question mark (?)
 - Question mark (?)
 - Forward slash (/)
 - Backslash (\)
 - Vertical bar / pipe (|)
 - Left curly brace ({)
 - Right curly brace (})
 - Caret (^)
 - Percent character (%)
 - Grave accent / back tick (`)
 - Right square bracket (])
 - Left square bracket ([)
 - Double quotation marks (")
 - Greater Than symbol (>)
 - Less Than symbol (<)
 - Tilde (~)
 - Pound character (#)
 - Control characters such as carriage return (CR) and line feed (LF),
 - Non-printable ASCII characters (128–255 decimal characters)
- Spectra Logic does not recommend accented characters in LTFS file names or BlackPearl object names because LTFS normalizes them before objects are written to tape and there could be conflicts with two objects having the same normalized name.
- **Object ID** — LTFS file names use the format *{bucket name}/{object id}*, for example `bucket1/1fc6f09c-dd72-41ea-8043-0491ab8a6d82`. Object names do not need to comply with LTFS file naming rules. The gateway saves object names as LTFS extended attributes allowing any third party application to reconstruct all the data including the object names.

**IMPORTANT**

If this storage domain is assigned to a data policy that uses versioning, after data is persisted, you cannot change this setting from Object ID to Object Name.

8. Optionally, select the **Media Export Allowed** check box to enable tape media export options for the storage domain. Clear the check box to disallow tape media export for the storage domain. This setting only applies to tape media. See [Tape Export Best Practices on page 26](#) for more information.

When a tape cartridge export occurs, a message displays in the BlackPearl user interface, and is also emailed to the system administrator. The system administrator **must** be configured to receive emails with both Informational and Warning message severity to be notified when tape media is exported. This allows the administrator to retrieve the tape media when it is exported. Do not leave tape media in the library Entry/Exit port for long periods of time.

Note: By configuring email alerts, the user is also notified when a GET job is requesting an object from exported tape media, so it can be imported into the library to complete the GET job.

Note: It is important to not export tape media from the library directly. The BlackPearl Nearline gateway controls the movement of media in the library.

If you select to allow media export, configure the following options:

- a. Select or clear options for auto media export.
 - **Auto Export on Job Completion** — select this option to have the gateway automatically export tape(s) when a job completes. This option is helpful if you plan to write a single job to tape and want to retrieve the media shortly after the job completes for secure archival or transfer of data to another tape library or BlackPearl Nearline gateway.
 - **Auto Export on Job Cancel** — select this option to have the gateway automatically export tape(s) when a user cancels a job. This option is helpful if you do not want append the next job to a partially filled tape.
 - **Auto Export on Media Full** — select this option to have the gateway automatically export tape(s) when a tape is full. This option is helpful to maximize the amount of data stored on tape media.

Note: If you select for media to automatically export when the media is full, you can optionally configure the **Export Media with an Available Capacity of Only** setting, which determines when the gateway marks a piece of media as full, and queues the piece of media for export. Select the desired unit size from the drop-down menu and enter a numerical value for the media full threshold in the text box to the left of the unit size drop-down menu.

- b. Using the **Auto Export Verify Task Priority** drop-down menu, select a task priority for tapes to be verified when they are automatically exported. Selecting **None** means that the gateway does not verify tapes before exporting them.

- c. Select or clear the **Scheduled Auto Export** check box. If enabled, this option automatically exports all tape media on a set schedule. This option is helpful if you need to move all tape media to off-site archival physical storage on a set schedule, regardless of the capacity of storage remaining on the tape cartridges. Use the instructions below to configure either hourly, daily, or weekly, automatic tape export.

Note: **Scheduled Auto Export** operates independently from the condition-based auto export options discussed in [Step a on page 79](#).

For example, if you select to have tape media auto export when full, the gateway exports a tape cartridge when it meets the media full threshold. Additionally, when the scheduled auto export time is met, the gateway exports **all** tape cartridges, regardless of whether they have reached the media full threshold.

Create an Hourly Schedule

- i. Select **Hourly** as the interval for the tape export schedule (see [Figure 14 on page 75](#)).
- ii. Enter numbers for **Every _ hours on minute _**. These values specify the interval in hours between tape exports and the number of minutes after the top of the hour when the export starts. For example, if the values are set to 4 and 15, tapes are exported every four hours at 15 minutes after the hour. The maximum setting for the **hours** field is 48, where tapes are exported every two days.

Note: Spectra Logic recommends offsetting the minutes after the hour for starting tape exports so that there are not a large number of jobs starting at exactly the same time.

Create a Daily Schedule

- i. Select **Daily** as the interval for the tape export schedule (see [Figure 14 on page 75](#)).
- ii. Enter a time value for **Start Time**, and include AM or PM after the value. This field is not case sensitive.
- iii. Enter a number for **Every _ days**. Allowed values are between 1 and 31. This value specifies the interval, in days, between scheduled exports. The value entered is enumerated from the first day of the month. The schedule resets at the beginning of each month. For example, if this value is set to 2, scheduled tape exports occur every two days, starting with the 1st of the month, at the time specified in [Step ii](#). A value of 30 runs on the 1st of the month, and then again on the 31st of the month (for months that have 31 days). To schedule exports on the first of every month, set the interval to 31 days.

Create a Weekly Schedule

- i. Select **Weekly** as the interval for the tape export schedule (see Figure 14 on page 75).
 - ii. Enter a time value for **Start Time**, and include AM or PM after the value. This field is not case sensitive.
 - iii. Select one or more days for **Every week on:**. This determines the day(s) of each week the gateway exports tapes.
9. Click **Create**. The new storage domain displays on the Advanced Bucket Management screen.

Add a Storage Domain Member to a Storage Domain

Once a storage domain is created, you must add storage domain members. Entire data partition/media type combinations are members of storage domains. When a bucket requires additional capacity, a single disk partition or tape cartridge is allocated out of the members to fulfill the capacity requirement.

Use the instructions in this section to add a storage domain member to a storage domain.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the storage domain for which you want to add a new storage domain member in the Storage Domains pane, or select the storage domain and select **Action > Show Details** from the menu bar. The Storage Domain details screen displays.

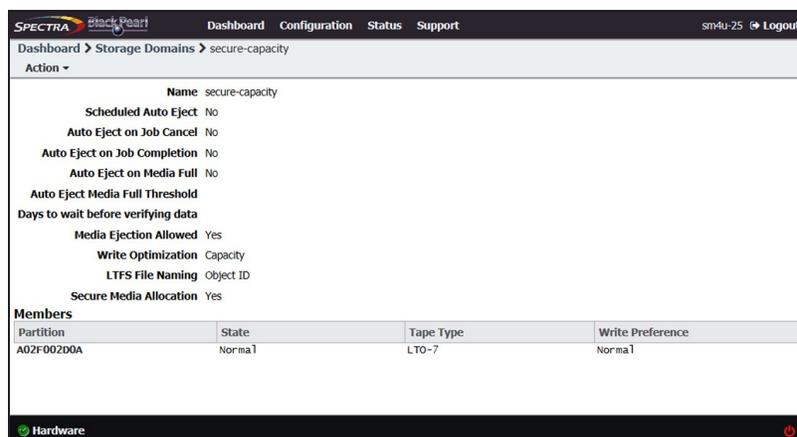


Figure 15 The Storage Domain details screen.

3. Select **Action > New Storage Domain Member**. The New Storage Domain Member dialog box displays.

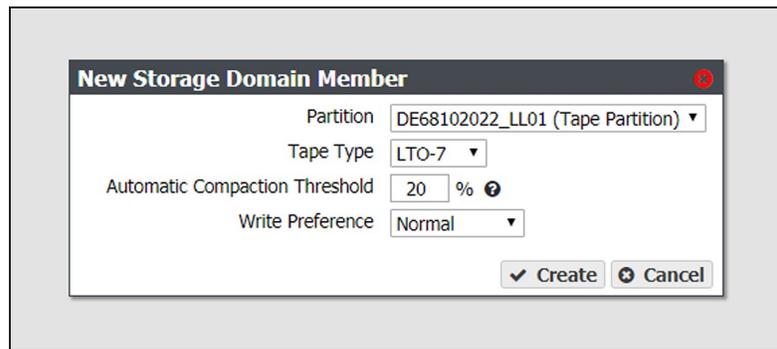


Figure 16 The New Storage Domain Member dialog box.

4. Use the **Partition Name** drop-down menu to select a tape or disk partition from the list of previously created partitions.
Note: You cannot add a disk partition to a storage domain that already uses a tape partition, and you cannot add a tape partition to a storage domain that already uses a disk partition.
5. Use the **Tape Type** drop-down menu to select the media type for a tape partition.
Notes:
 - You must select the media type that matches the media present in the tape library partition. If the partition contains multiple generations of media, select the highest version.
 - This option does not display if you selected a disk partition in [Step 4 on page 82](#).
6. Enter a percentage for the **Automatic Compaction Threshold**. Automatic compaction occurs when the percentage of deleted objects on a tape cartridge exceeds this value. The default percentage is 95.
Note: If you selected a disk partition in [Step 4 on page 82](#), this setting is unavailable.
7. Use the **Write Preference** drop-down menu to select the write preference for this member of the storage domain. This setting determines the preferred usage of the partition when additional capacity is needed. The gateway uses a partition with **High** write preference before a partition with **Normal** write preference, and so on. Use **Never Select** to indicate that a partition is read-only.
8. Click **Create**. The new storage domain member displays on the Storage Domain details screen.

CREATE A DATA POLICY

A data policy defines data integrity policies, default job attributes, and persistence and replication rules, which define where data is written and for how long it is kept. A data policy may be used by multiple buckets, but a bucket uses precisely one data policy.

Use the instructions in this section to create a new data policy.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Select **Action > New Data Policy**. The New Data Policy dialog box displays.

Figure 17 The New Data Policy dialog box.

3. Enter a name for the data policy in the **Data Policy Name** field.

4. Select or clear the **Blobbing Enabled** check box. When enabled this setting allows an object to be broken into multiple blobs. If disabled, an object must be exactly one blob. Blobbing must be enabled to handle objects larger than 1 TB, to use multi-part upload, or to break up an object into multiple blobs.

Note: Disabling blobbing guarantees that an object will never span multiple tapes or disk pools, since a blob cannot span multiple media.

5. Select or clear the **Minimize Spanning** check box. When enabled, this setting minimizes the spanning of data across multiple tapes or pools. Jobs less than 1 TB never span media.

Notes:

- When "Minimize Spanning" is enabled, it overrides a storage domain's capacity mode and performance mode logic for a given job, up to 1 TB in size. Minimize spanning increases the chunk size to 1 TB, and always keeps one chunk on a single tape regardless of write mode. If the job is larger than 1 TB, then multiple chunks are used and the gateway uses the logic for capacity mode vs performance mode, where different chunks may transfer to different tapes.

- Enabling this option can adversely affect capacity utilization and performance.

6. Select the **Performance** characteristics for the data policy. Each priority determines the resources assigned and the processing order. Jobs with priority **Urgent** can use up all of the resources and prevent other jobs from making progress. Use this priority sparingly.

- a. Use the drop-down menu to select the **Default GET Priority**.

- b. Use the drop-down menu to select the **Default PUT Priority**.

- c. Use the drop-down menu to select the **Default VERIFY Priority**.

- d. Select or clear **Default Verify After Write**. Clients may specify whether or not to verify data immediately after writes when creating a PUT job. If the client does not specify a policy for verification of data after writes, this selection determines whether a verify is done. If done, the verification uses the checksum type specified in [Step 7 on page 85](#).

- Notes:**
- After the PUT job completes, the tape remains in the drive during data verification.
 - Only the data just written by the PUT job is verified.
 - This verification does not update the **Last Verified** field on the tape details screen.
 - Selecting **Default Verify After Write** reduces gateway write throughput by up to 50%.
 - This setting does not apply to replication targets.
- e. Use the drop-down menu to select the **Rebuild Priority**.

7. Select the **Data Security** options for the data policy.
 - a. Use the drop-down menu to select the **Checksum type**. This setting specifies the type of checksum used to verify data integrity for data in any bucket using this data policy, and the type of checksum required for end-to-end CRC, if specified.
- Notes:**
 - CRC, MD5, and SHA-512 perform the best for their corresponding cryptographic strengths on the BlackPearl gateway.
 - Using SHA-256 and SHA-512 reduces single stream performance and may reduce throughput capabilities of the gateway.
- b. If you want to enable end-to-end security for each GET or PUT job, select the **Require end-to-end CRC** check box.
8. Select the type of **Versioning** you want to use for the data policy. You must select either **None**, **Keep Latest**, or **Keep Multiple Versions**.

**IMPORTANT**

If you select **Keep Multiple Versions**, you are not able to change this setting after the data policy is created.

- **None**—Only one version of an object may exist at any time and the version number of the object is always 1.
- **Keep Latest**—Only one version of the data is available at a time. When a new version of an object is written, the old version is retained until the new version is fully written in compliance with the data policy, and then the old version is deleted.
- **Keep Multiple Versions** (*default*)—When a new version of an object is written, it is added as the latest version of the object. Any previous versions of the object, up to the value specified, are retained and accessible. The default value of 1000 is pre-entered.

- Notes:**
- You cannot assign a Storage Domain configured with the LTFS option set to **Object Name** when using the **Keep Latest** or **Keep Multiple Versions** setting. See [Create a Storage Domain on page 75](#) for more information.
 - The **Keep Latest** setting requires that the PUT job for the earlier version of the object complete before the PUT of the latest version of the object with the same name in order for the PUT job to succeed.

**CAUTION**

If you select **Keep Multiple Versions**, if the PUT of the earlier version is not complete before the PUT of the latest version, the BlackPearl gateway believes the latest version to be the same object as the earlier version and rejects it, and only the earlier version is retained.

9. If you plan to configure a data replication target, select or clear the **Always accept replicated PUT jobs** check box. This option controls whether all PUT jobs for this data policy are created even if one or more replication targets the gateway must PUT to are unavailable, or if there are global issues that would likely prevent the completion of the job.

Note: Using this parameter is discouraged, and using it for jobs on both the source and target gateways at the same time is extremely discouraged. Running jobs on both gateways when they are not able to communicate with each other can create replication conflicts that must be manually resolved.

10. Click **Create**. The new data policy displays on the Advanced Bucket Management screen.

Add Data Persistence Rules and Replication Rules to a Data Policy

Once a data policy is created, you must add persistence rules. A persistence rule is either permanent, meaning that data is kept in the specified storage domain at all times, or temporary, meaning that data is kept in the specified storage domain under certain circumstances, and then it can be deleted from that storage domain. Existing permanent and temporary persistence rules, and replication rules, may be retired so that the rule is not applied for any new incoming data, but will continue to retain data previously written. A data policy must include at least one permanent persistence rule.

Add a Data Persistence Rule to a Data Policy

1. If necessary, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6](#) on page 57).

2. Double-click the data policy in the Data Policies pane, or select the data policy and select **Action > Show Details** from the menu bar. The Data Policy details screen displays.

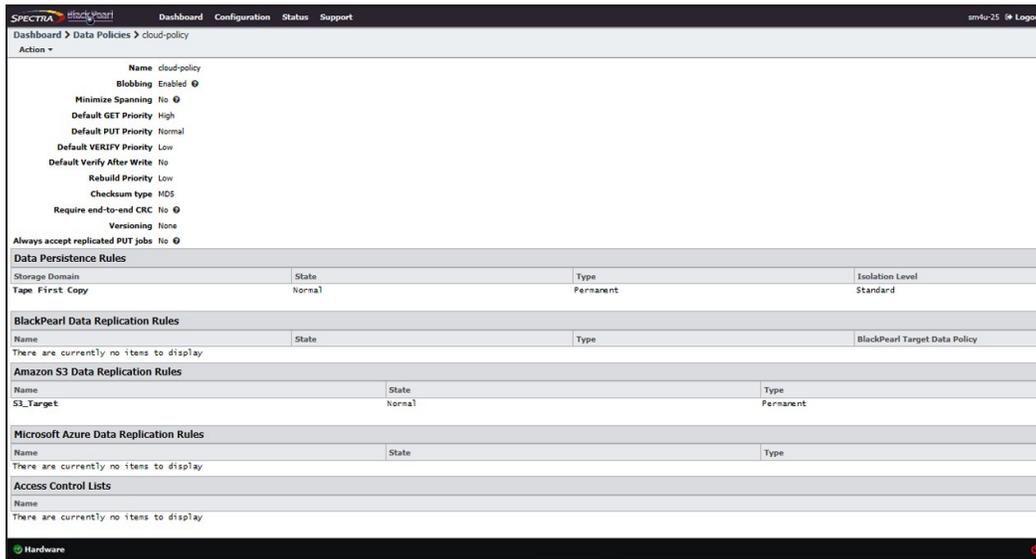


Figure 18 The Data Policy details screen.

3. Select **Action > New Data Persistence Rule**. The New Data Persistence Rule dialog box displays.

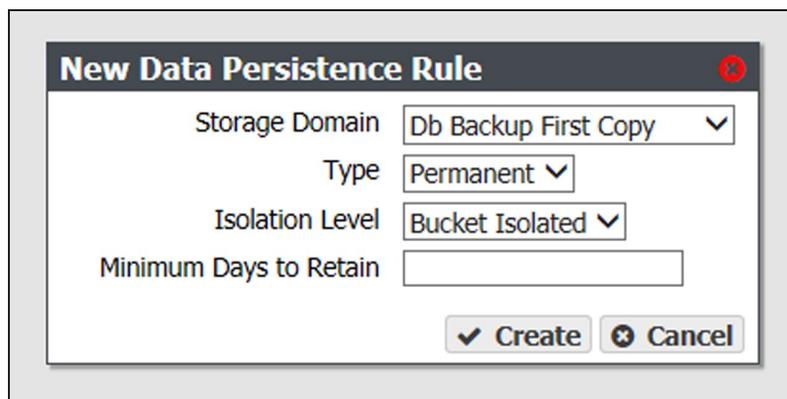


Figure 19 The New Data Persistence Rule dialog box.

4. Use the **Storage Domain** drop-down menu to select a storage domain from the list of previously created storage domains.

5. Use the **Type** drop-down menu to select whether the data persistence rule to use for the storage domain is **Temporary** or **Permanent**.

- **Temporary** - The data is kept in the specified storage domain under certain circumstances, and then it can be deleted from that storage domain.
- **Permanent** - The data is kept in the specified storage domain at all times.

Notes:

- The **Temporary** setting cannot be used for a storage domain that targets a tape library.
- When importing data, a **Temporary** persistence rule does not trigger copying data to a disk pool unless the data is staged with IOM (Intelligent Object Management) active and running. See for information on IOM.
- You cannot create a Data Persistence Rule with a setting of **Retired**. Existing persistence rules can be modified to be retired. See [Edit a Data Persistence Rule on page 118](#).

6. Use the **Isolation Level** drop-down menu to select the level of physical isolation required for the storage domain.

- **Standard** — This allows data from different buckets to reside on the same physical media, and may provide increased performance. This setting is recommended data policies configured to use disk storage.
- **Bucket Isolated** — Data from different buckets cannot be mixed on the same physical storage media.

Notes:

- The **Standard** isolation level provides the best capacity utilization and overall performance.
- **Bucket Isolated** allocates an entire disk pool to a bucket when needed. Allocating an entire disk pool to a bucket may use up resources quickly and is not recommended.

7. Enter the **Minimum Days to Retain** in the entry field to specify the minimum number of days the gateway should retain data written using a temporary persistence rule.

Notes:

- The **Minimum Days to Retain** for a persistence rule targeting a storage domain using a nearline pool (a 77-bay, 96-bay, or 107-bay expansion node) must be 90 days or greater.
- **Minimum Days to Retain** cannot be specified when using a **Type** of **Permanent**.

8. Click **Create**. The new data persistence rule displays on the Data Policy details screen.

Add a BlackPearl Data Replication Rule to a Data Policy

A BlackPearl replication target must be configured before adding a BlackPearl Data Replication Rule to the data policy. See [Create a BlackPearl Target on page 64](#).

1. If necessary, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the data policy in the Data Policies pane, or select the data policy and select **Action > Show Details** from the menu bar. The Data Policy details screen displays (see Figure 18 on page 87).
3. Select **Action > New BlackPearl Data Replication Rule**. The New BlackPearl Data Replication Rule dialog box displays.

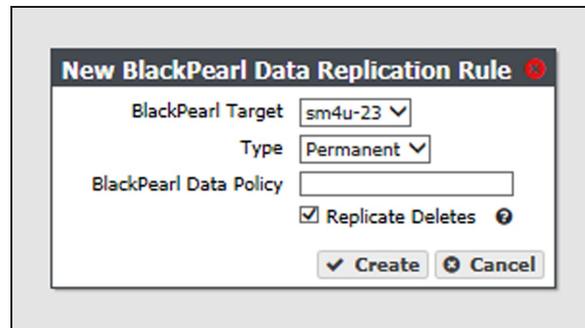


Figure 20 The New BlackPearl Data Replication Rule dialog box.

4. Use the **BlackPearl Target** drop-down menu to select a replication target from the list of previously created replication targets.
5. Use the **Type** drop-down menu to select whether the data persistence rule to use for the for the storage domain is **Permanent** or **Retired**.

Note: You cannot create a replication rule as **Retired**. You can only modify a rule from permanent to retired after the data policy is created.

6. In the **BlackPearl Data Policy** entry field, enter the name of the data policy on the target BlackPearl gateway to use when creating the bucket for replicated data. Alternatively, you can leave the field blank.

Notes:

- The data policy name is case sensitive.
- If the field is left blank and the BlackPearl target was configured with the setting “Replicated User Default Data Policy” enabled, and Access Control Replication was set to “Users”, the default data policy on the target gateway is used. If no default is set on the target gateway and the target gateway is configured with more than one data policy, the replication fails.

7. Select or clear the **Replicate Deletes** check box. When selected, any time a replicated file is deleted from the source gateway, it is also deleted from the target gateway.

Note: Replicated objects do not immediately delete. Objects are only deleted after running a verify operation on the bucket.

8. Click **Create**. The new BlackPearl data replication rule displays on the Data Policy details screen.

Add an Amazon S3 Data Replication Rule to a Data Policy

An Amazon S3 replication target must be configured before adding an Amazon S3 Data Replication Rule to the data policy. See [Create an Amazon S3 Target](#) on page 68.

1. If necessary, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6](#) on page 57).
2. Double-click the data policy in the Data Policies pane, or select the data policy and select **Action > Show Details** from the menu bar. The Data Policy details screen displays (see [Figure 18](#) on page 87).
3. Select **Action > New Amazon S3 Data Replication Rule**. The New Amazon S3 Data Persistence Rule dialog box displays.

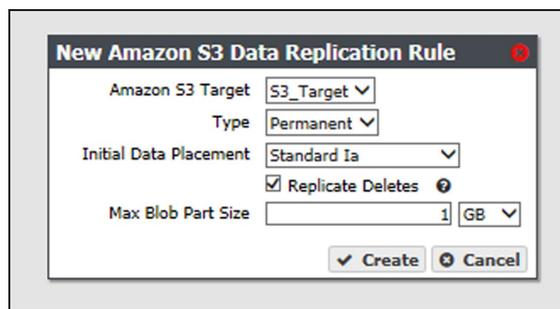


Figure 21 The New Amazon S3 Data Replication Rule dialog box.

4. Use the **Amazon S3 Target** drop-down menu to select an Amazon S3 replication target from the list of previously created replication targets.
5. Use the **Type** drop-down menu to select whether the data persistence rule to use for the for the replication target is **Permanent** or **Retired**.

Note: You cannot create a replication rule as **Retired**. You can only modify a rule from permanent to retired after creating the replication rule.
6. Use the **Initial Data Placement** drop-down menu to select the storage class for any objects transferred to the AWS S3 instance. See [Storage Classes](#) for detailed descriptions of the storage classes provided by AWS.

Note: The BlackPearl gateway uses "standard" restore for objects archived to Glacier and Glacier Deep Archive storage classes. Restore times are approximately 3-5 hours for Glacier, and 12 hours for Glacier Deep Archive, plus object download time.

 - **Standard** – Provides high availability and performance for frequently accessed data.

- **Reduced Redundancy** — Provides storage of objects on multiple devices across multiple facilities, but does not replicate objects as many times as Amazon S3 standard storage. The lower level of redundancy results in less durability and availability, but also lower storage costs.
- **Standard IA** (default) — Provides fast access to less frequently accessed data.
- **Glacier** — Provides secure, long-term archive for rarely accessed data.
- **Glacier Deep Archive** — Provides a low-cost, secure long-term archive for data that does not require quick retrieval.

Note: If you are configuring the replication to target a bucket that was previously created using the Amazon AWS interface, you must define a Lifecycle Management Rule in AWS to migrate data from the bucket default tier to the preferred tier, if necessary. Spectra Logic recommends using an immediate (0 days) move rule.

7. Select or clear the **Replicate Deletes** check box. When selected, any time a replicated file is deleted from the source gateway, it is also deleted from the target.
8. If desired, modify the **Max Blob Part Size**. This parameter defines the maximum object part size used when sending data to an Amazon S3 target. Larger blob sizes make public cloud workflows simpler, but may make it more difficult or impossible to reliably transmit blobs. Less reliable network connections to the public cloud require smaller blob sizes. The maximum blob size is 1 TB. The default maximum blob size is 1 GB.

Note: To prevent data transfer failures, it is important that this value not exceed the maximum blob size that the target is able to accept.

9. Click **Create**. The new Amazon S3 replication rule displays on the Data Policy details screen.

Add a Microsoft Azure Data Replication Rule to a Data Policy

A Microsoft Azure replication target must be configured before adding an Microsoft Azure Data Replication Rule to the data policy. See [Create a Microsoft Azure Target on page 72](#).

1. If necessary, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6 on page 57](#)).
2. Double-click the data policy in the Data Policies pane, or select the data policy and select **Action > Show Details** from the menu bar. The Data Policy details screen displays (see [Figure 18 on page 87](#)).

3. Select **Action > New Microsoft Azure Data Replication Rule**. The New Microsoft Azure Data Replication Rule dialog box displays.

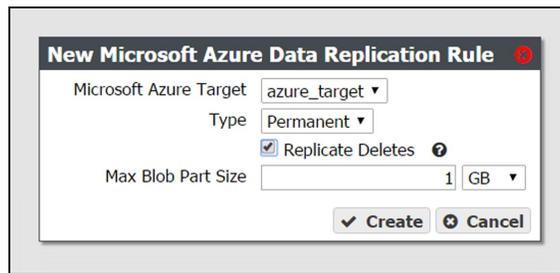


Figure 22 The New Microsoft Azure Data Replication Rule dialog box.

4. Use the **Microsoft Azure Target** drop-down menu to select a Microsoft Azure replication target from the list of previously created replication targets.
5. Use the **Type** drop-down menu to select whether the data persistence rule to use for the for the replication target is **Permanent** or **Retired**.

Note: You cannot create a replication rule as **Retired**. You can only modify a rule from permanent to retired after creating the replication rule.

6. Select or clear the **Replicate Deletes** check box. When selected, any time a replicated file is deleted from the source gateway, it is also deleted from the target.

Note: Replicated objects do not immediately delete. Objects are only deleted after running a verify operation on the bucket.

7. If desired, modify the **Max Blob Part Size**. This parameter defines the maximum object part size used when sending data to a Microsoft Azure target. Larger blob sizes make public cloud workflows simpler, but may make it more difficult or impossible to reliably transmit blobs. Less reliable network connections to the public cloud require smaller blob sizes. The maximum blob size is 1 TB. The default maximum blob size is 1 GB.

Note: To prevent data transfer failures, it is important that this value not exceed the maximum blob size that the target is able to accept.

8. Click **Create**. The new Microsoft Azure replication rule displays on the Data Policy details screen.

Create Data Policy ACLs

The sections below describe creating Access Control List (ACL) for both a group of user and an individual user.

New Data Policy ACL for a Group

Use the instruction in this section to create a new data policy ACL for a group.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Select the row for the data policy for which you want to create a new ACL for a group, then select **Action > Show Details**. The Data Policy details screen displays.
3. From the menu bar, select **Action > New Data Policy ACL For Group**. The New Data Policy ACL For Group dialog box displays.

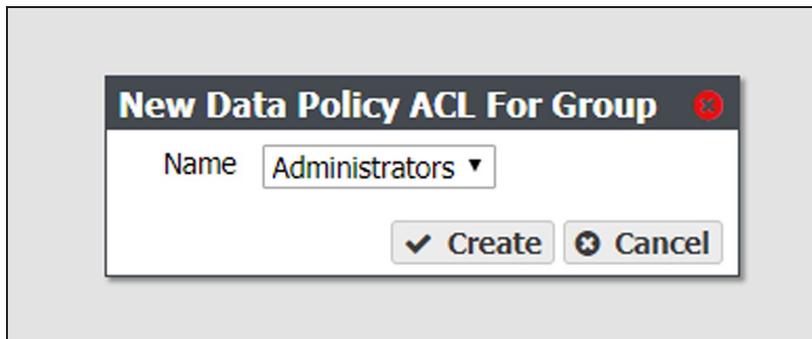


Figure 23 The New Data Policy ACL For Group dialog box.

4. Using the **Name** drop-down menu, select the group to be assigned to the data policy ACL.
5. Click **Create**.

New Data Policy ACL for a User

Use the instruction in this section to create a new data policy ACL for an individual user.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Select the row for the data policy for which you want to create a new ACL for a user, then select **Action > Show Details**. The Data Policy details screen displays.

3. From the menu bar, select **Action > New Data Policy ACL For User**. The New Data Policy ACL For User dialog box displays.

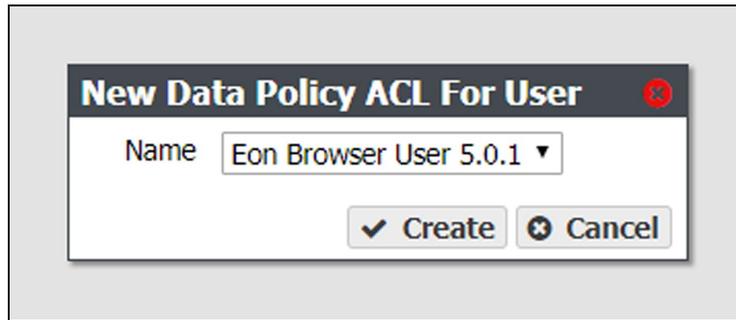


Figure 24 The New Data Policy ACL For User dialog box.

4. Using the **Name** drop-down menu, select the user to be assigned to the data policy ACL.
5. Click **Create**.

CREATE A BUCKET

Buckets are data transfer targets for read and write operations. The gateway stages data written to it on the cache and optimizes how it writes buckets to storage domains for best performance.

Clients write data to the gateway using a “bulk PUT” command, and read from the gateway with a “bulk GET” command. For more information on using these commands see the *Spectra BlackPearl DS3 API Reference*.

Note: Buckets can also be created using a DS3 client, or the DS3 API.



IMPORTANT

If you are creating a bucket that is used in a BlackPearl replication configuration, you must create the bucket on the source gateway, and the target gateway using identical names, or replication fails.

Use the instructions in this section to configure a bucket.

1. From the menu bar, select **Configuration > Buckets**. The Buckets screen displays.

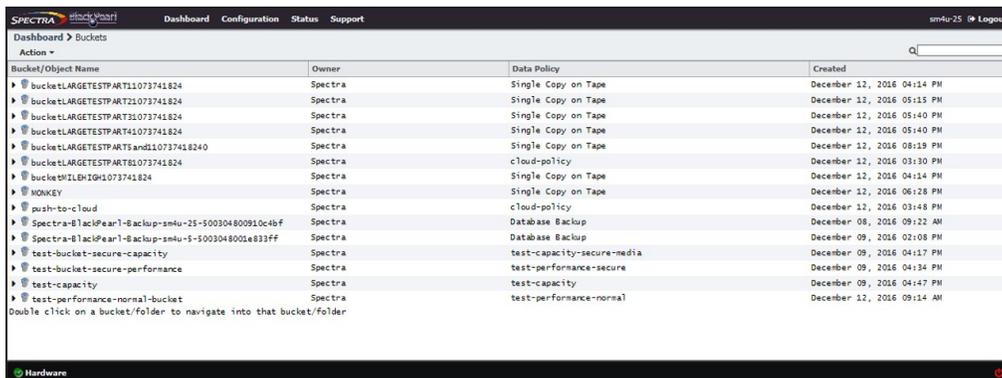


Figure 25 The Buckets screen.

2. Select **Action > New** from the menu bar. The New Bucket dialog box displays.

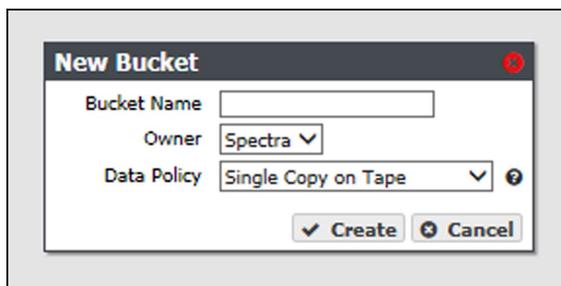


Figure 26 The New Bucket dialog box.

3. Enter a name for the bucket in the **Bucket Name** field.

When creating a bucket for use with an Amazon S3 or Microsoft Azure replication target, the bucket name must adhere to the cloud target naming requirements. The BlackPearl gateway attempts to create the bucket on the replication target using the name entered in Step 3 with the appended Cloud Bucket Prefix and Suffix, if applicable.



IMPORTANT

- For **BlackPearl OS 3.5.2 or earlier**, the BlackPearl gateway changes bucket names with upper case letters to all lower case letters when needed. If you are using bucket names that only differ by case, the buckets are combined on the cloud target causing possible data collision and bucket ownership/permission problems.
- For **BlackPearl OS 4.0 or later**, if the bucket name is incompatible with the naming requirements of the cloud target provider, bucket creation fails and an error message displays.

- Notes:**
- The bucket name cannot contain a colon (:), forward slash (/), or space.
 - The bucket name cannot exceed 255 characters.

4. Using the drop-down menu, select an **Owner** for the bucket from the list of users already created on the gateway.
5. Using the drop-down menu, select a **Data Policy** for the bucket from the list of previously created data policies on the gateway. The bucket uses this data policy when transferring data.
6. Click **Create**. The Buckets screen displays with the newly created bucket listed.

Create a New Bucket ACL for a Group

Use the instructions in this section to create a new Access Control List (ACL) using the specified group for a bucket.

1. From the menu bar, select **Configuration > Buckets**. The Buckets screen displays (see Figure 25 on page 95).
2. Select the bucket for which you want to create an ACL and select **Action > Show Details**. The bucket details screen displays.



Figure 27 The bucket details screen.

3. Select **Action > New Bucket ACL For Group**. The New Bucket ACL For Group dialog box displays.

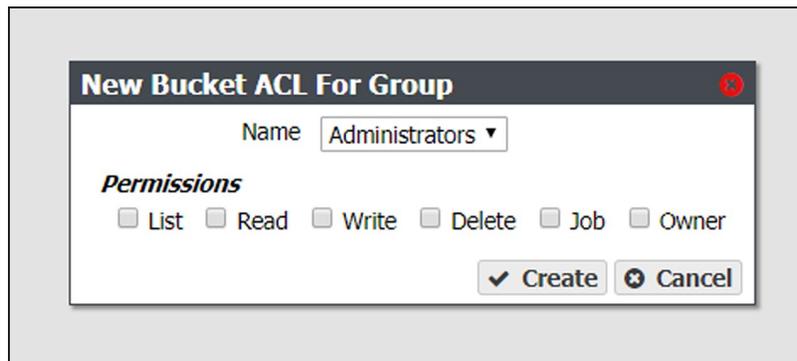


Figure 28 The New Bucket ACL For Group dialog box.

4. Using the **Name** drop-down list, select a group from the list of existing S3 groups on the BlackPearl gateway.
5. Select the desired **Permissions** for the group ACL.
 - **LIST** — The users in the group can see the bucket in a get buckets request and can list the objects in a bucket. The users can also perform any type of bucket or object get that does not involve returning the actual data for an object.
 - **READ** — The users can get objects and create GET jobs.
 - **WRITE** — The users can put objects and create PUT jobs.
 - **DELETE** — The users can delete objects, but cannot delete the bucket.
 - **JOB** — The group members can modify or cancel jobs that they did not create. The users can also see the details of jobs they did not create. Note that all users can view all jobs, but by default, only the initiator of the job can see the full details of a job.
 - **OWNER** — The users receives full access to the bucket, including all permissions listed above, and also receives permission to modify bucket ACLs for that bucket.
6. Click **Create**.

Create a New Bucket ACL for a User

Use the instructions in this section to create a new Access Control List (ACL) using the specified user for a bucket.

1. From the menu bar, select **Configuration > Buckets**. The Buckets screen displays (see Figure 25 on page 95).

2. Select the bucket for which you want to create an ACL and select **Action > Show Details**. The bucket details screen displays.

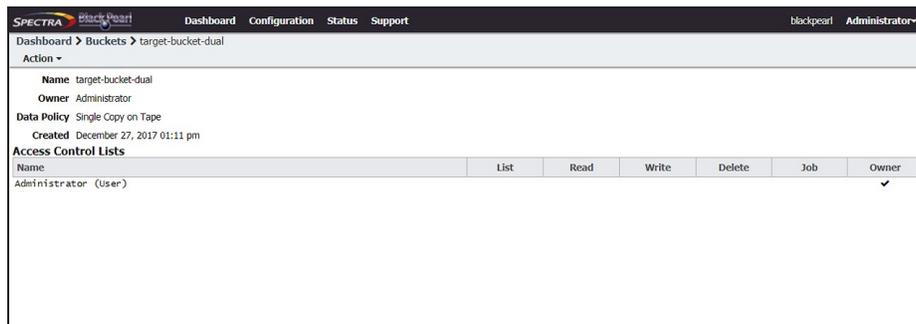


Figure 29 The bucket details screen.

3. Select **Action > New Bucket ACL For User**. The New Bucket ACL For User dialog box displays.

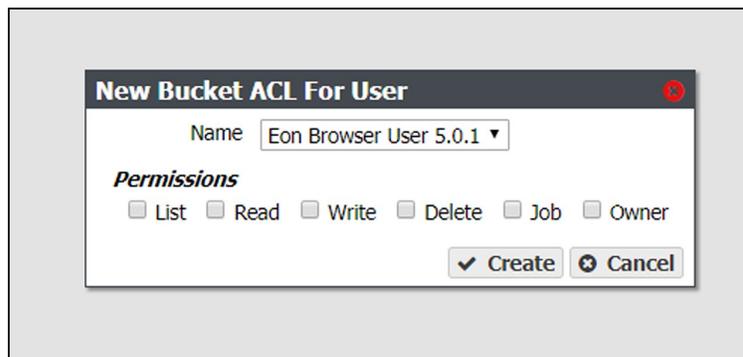


Figure 30 The New Bucket ACL For User dialog box.

4. Using the **Name** drop-down list, select a user from the list of existing users on the BlackPearl gateway.
5. Select the desired **Permissions** for the user ACL.
 - **LIST** — The user can see the bucket in a get buckets request and can list the objects in a bucket. The user can also perform any type of bucket or object get that does not involve returning the actual data for an object.
 - **READ** — The user can get objects and create GET jobs.
 - **WRITE** — The user can put objects and create PUT jobs.
 - **DELETE** — The user can delete objects, but cannot delete the bucket.
 - **JOB** — The user can modify or cancel jobs that they did not create. The user can also see the details of jobs they did not create. Note that all users can view all jobs, but by default, only the initiator of the job can see the full details of a job.
 - **OWNER** — The user receives full access to the bucket, including all permissions listed above, and also receives permission to modify bucket ACLs for that bucket.
6. Click **Create**.

TRANSFER DATA

After completing the above sections of this chapter, you are ready to begin transferring data to, and from, the BlackPearl gateway.

- To transfer data using the EON Browser, consult the *EON Browser User Guide*.
- To transfer data using the Spectra RioBroker[®] application, consult the *Spectra RioBroker User Guide*.
- To transfer data using the Spectra StorCycle[®] application, consult the *Spectra StorCycle User Guide*.
- To transfer data using a DS3 client, consult the *Spectra BlackPearl DS3 API Reference*, and documentation specific to each DS3 client.

**IMPORTANT**

Contact Spectra Logic Professional Services for assistance in setting up data transfer. See [Contacting Spectra Logic](#) on page 7.

**IMPORTANT**

The Spectra BlackPearl gateway limits the number of simultaneous active jobs. The number of simultaneous jobs differs depending on the software installed on the BlackPearl gateway.

- BlackPearl OS 3.x is limited to 1,000 simultaneous active jobs.
 - BlackPearl OS 4.x is limited to 1,000 simultaneous active jobs.
 - BlackPearl OS 5.x is limited to 10,000 simultaneous active jobs.
-

CHAPTER 3 - MANAGING ADVANCED BUCKET MANAGEMENT SETTINGS

This chapter describes using the BlackPearl user interface to manage storage domains, data policies, disk partitions, and buckets on the gateway after configuring Advanced Bucket Management. For initial Advanced Bucket Management configuration steps, see [Configuring Advanced Bucket Management on page 54](#).

Manage Buckets	102
Show Bucket Physical Placement	102
Edit a Bucket ACL	102
Delete a Bucket ACL	103
Edit a Bucket	103
Delete a Bucket	104
Manage a Storage Domain	106
Edit a Storage Domain Member	106
Exclude a Storage Domain Member	107
Cancel Storage Domain Member Exclusion	108
Delete a Storage Domain Member	108
Edit a Storage Domain	109
Delete a Storage Domain	110
Manage Data Replication Rules	112
Edit a BlackPearl Data Replication Rule	112
Edit an Amazon S3 Data Replication Rule	113
Edit a Microsoft Azure Data Replication Rule	114
Delete a Data Replication Rule	115
Manage a Data Policy	116
Edit a Data Policy	116
Delete a Data Policy	117
Delete a Data Policy ACL	118
Edit a Data Persistence Rule	118
Delete a Data Persistence Rule	119
Manage Replication Targets	121

Edit a BlackPearl Replication Target	121
Edit an Amazon Replication Target	122
Edit an Azure Replication Target	123
Verify a Replication Target	123
Put a Replication Target in Standby State	124
Activate a Replication Target	125
Delete a Replication Target	126
Manage a Disk Partition	127
Edit a Disk Partition	127
Delete a Disk Partition	128
Manage Online and Nearline Disk Pools	129
Edit a Nearline or Online Disk Pool	129
Import a Nearline or Online Disk Pool	130
Delete a Nearline or Online Disk Pool	131
Change the BlackPearl Gateway Cache Configuration	133

MANAGE BUCKETS

Use the instructions in this section to configure ACLs for a bucket, edit, or delete a bucket. For instructions on creating a new bucket, see [Create a Bucket on page 95](#).

Show Bucket Physical Placement

Once data is transferred to the BlackPearl gateway, you can view the physical placement of the data. The BlackPearl user interface displays data placement on disk pools, tapes, and replication targets. Use the instructions in this section to view physical placement of a specified bucket.

1. From the menu bar, select **Configuration > Buckets**. The Buckets screen displays (see [Figure 25 on page 95](#)).
2. Select the bucket for which you want to view physical placement and select **Action > Show Physical Placement**. The *Bucket Physical Placement* screen displays.

SPECTRA BlackPearl												
Dashboard Configuration Status Support											neoengtest Administrator	
Dashboard > Buckets > TestBlackPearlCopyMultipleFromZipPacking_155726881296202... > Physical Placement												
Disk Pools												
Name	Health	S3 Pool Type			Standby	Partition						
Arctic_Blue_1_11060724571267...	Online	Nearline			No	neo-disk-partition						
Double click on a row to see more detailed information												
Tapes												
Barcode	Serial Number	Type	State	Write Protected	Available	Used	Tape Library Partition	Storage Domain	Bucket	Last Modified	Last Verified	Loaded In Drive
There are currently no items to display												
Microsoft Azure Targets												
Name	Account Name			State	Standby							
There are currently no items to display												
BlackPearl Targets												
Name	Data Path End Point			State	Standby							
There are currently no items to display												
Amazon S3 Targets												
Name	Data Path End Point	Region	State	Standby	S3 File Naming							
There are currently no items to display												

Figure 31 The *Bucket Physical Placement* screen.

Edit a Bucket ACL

Use the instructions in this section to edit an Access Control List (ACL) for a bucket.

1. From the menu bar, select **Configuration > Buckets**. The Buckets screen displays (see [Figure 25 on page 95](#)).
2. Select the bucket for which you want to edit an ACL and select **Action > Show Details**. The bucket details screen displays (see [Figure 27 on page 96](#)).

3. Select the Access Control List you want to edit and select **Action > Edit Bucket ACL**. The Edit Bucket ACL dialog box displays.

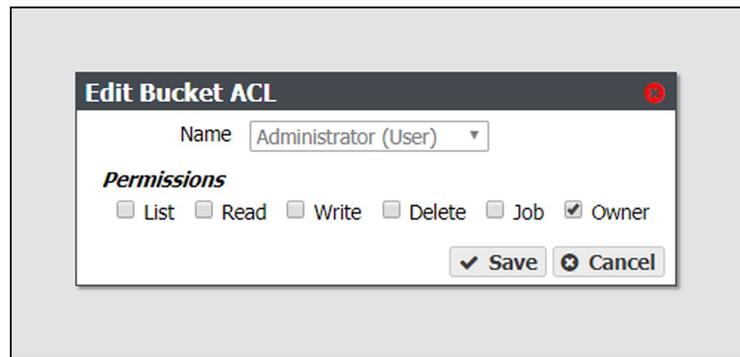


Figure 32 The Edit Bucket ACL dialog box.

Note: You cannot change the user of an existing bucket ACL list. The **Name** drop-down list is unavailable.

4. Select or clear the desired **Permissions** for the bucket ACL.
5. Click **Save**.

Delete a Bucket ACL

Use the instructions in this section to delete an Access Control List (ACL) for a bucket.

1. From the menu bar, select **Configuration > Buckets**. The Buckets screen displays (see Figure 25 on page 95).
2. Select the bucket for which you want to delete an ACL and select **Action > Show Details**. The bucket details screen displays (see Figure 27 on page 96).
3. Select the Access Control List you want to delete and select **Action > Delete Bucket ACL**. The Delete Bucket ACL confirmation dialog box displays.
4. Click **Delete**.

Edit a Bucket

Use the instructions in this section to change the owner of a bucket or to change the data policy used by the bucket.

1. From the menu bar, select **Configuration > Buckets**. The Buckets screen displays (see Figure 25 on page 95).

2. Select the bucket you want to edit and select **Action > Edit**. The Edit Bucket dialog box displays.

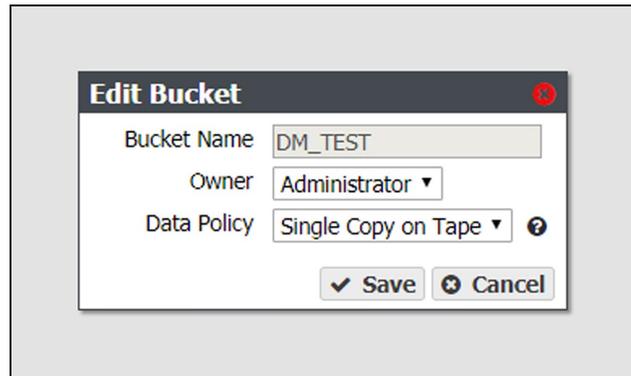


Figure 33 The Edit Bucket dialog box.

3. The **Bucket Name** is unavailable and cannot be changed.
4. From the drop-down list, change the **Owner** assigned to the bucket.



IMPORTANT

After changing bucket owners, you may need to reconfigure your client application to access the bucket.

5. Use the **Data Policy** drop down menu to change the data policy assigned to the bucket. For more information on data policies, see [Data Policies on page 20](#).
6. Click **Save**.

Delete a Bucket

Use the instructions in this section to delete a bucket.



CAUTION

When you delete a bucket, all data contained in the bucket is lost. Any tapes associated with the bucket are marked as Free, and are available to the gateway for other storage operations immediately. Any bucket data that was written to tape media is retained until the tape is loaded into a drive and new data is written.

Note: You cannot delete a bucket created by the Spectra Vail, StorCycle, or RioBroker applications if the bucket contains data. To delete the bucket, use application that created the bucket to delete all data, then delete the bucket using the BlackPearl user interface.

1. From the menu bar, select **Configuration > Buckets**. The Buckets screen displays (see [Figure 25 on page 95](#)).

2. Select the bucket you want to delete and select **Action > Delete**. A confirmation dialog box displays.

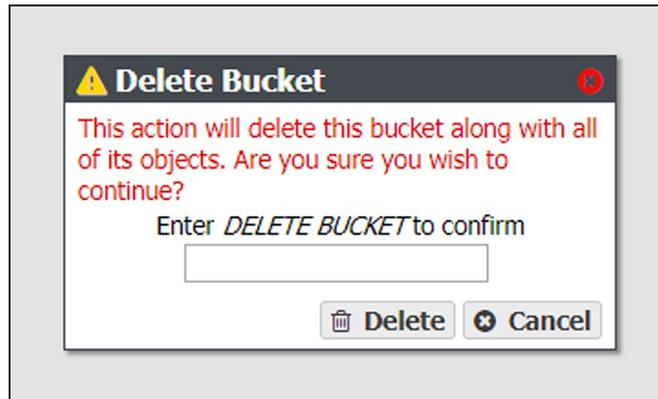


Figure 34 The bucket delete confirmation dialog box.

3. Type `DELETE BUCKET` into the entry field, and then click **Delete**.

MANAGE A STORAGE DOMAIN

Use the instructions in this section to edit or delete a storage domain member or a storage domain.

Edit a Storage Domain Member

Use the instructions in this section to change the write preference for a storage domain member.

Note: For every storage domain, at least one storage domain member must have a write preference other than **Never_Select**.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the row for the storage domain with the storage domain member that you want to edit, or select the storage domain row and then select **Action > Show Details**. The Storage Domain details screen displays (see Figure 15 on page 81).
3. Select the storage domain member row and then select **Action > Edit Storage Domain Member**.

Note: Do not click the partition name when selecting the data persistence rule row or the Partition details screen will open.

The Edit Storage Domain Member dialog box displays.

Figure 35 The Edit Storage Domain Member dialog box.

4. If desired, edit the percentage for the **Automatic Compaction Threshold**. Automatic compaction occurs when the percentage of deleted objects on a tape cartridge exceeds this value.
5. If desired, using the drop down menu, edit the **Write Preference** field as necessary. See [Add a Storage Domain Member to a Storage Domain on page 81](#) for a description of each field.

Note: When editing a storage domain member, the **Partition** and **Tape Type** settings are unavailable.

6. Click **Save**. The edited storage domain member displays on the Storage Domain details screen.

Exclude a Storage Domain Member

Use the instructions in this section to exclude a storage domain member. This command migrates data off of the selected storage domain member before deleting the storage domain member.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the row for the storage domain with the storage domain member that you want to exclude, or select the storage domain row and then select **Action > Show Details**. The Storage Domain details screen displays (see Figure 15 on page 81).
3. Select the desired storage domain member row and then select **Action > Exclude Storage Domain Member**.

Note: Do not click the partition name when selecting the data persistence rule row or the Partition details screen will open.

The Exclude Storage Domain Member dialog box displays.

4. Type `EXCLUDE` in the entry field and then click **Exclude**.

Cancel Storage Domain Member Exclusion

Use the instructions in this section to cancel a storage domain member exclusion in process.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the row for the storage domain with the storage domain member that is currently being excluded, or select the storage domain row and then select **Action > Show Details**. The Storage Domain details screen displays (see Figure 15 on page 81).
3. Select the desired storage domain member row and then select **Action > Cancel Storage Domain Member Exclusion**.

Note: Do not click the partition name when selecting the data persistence rule row or the Partition details screen will open.

The Cancel Storage Domain Member Exclusion dialog box displays.

4. Type `CANCEL` in the entry field and then click **Cancel**.

Delete a Storage Domain Member

Use the instructions in this section to delete a storage domain.

- Notes:**
- You cannot delete a storage domain member that has a tape or pool assigned to the storage domain.
 - You cannot delete the last storage domain member of a storage domain assigned to a data policy.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the row for the storage domain with the storage domain member that you want to edit, or select the storage domain row and then select **Action > Show Details**. The Storage Domain details screen displays (see Figure 15 on page 81).
3. Select the storage domain member row and then select **Action > Delete Storage Domain Member**.

Note: Do not click the partition name when selecting the data persistence rule row or the Partition details screen will open.

A confirmation dialog box displays.

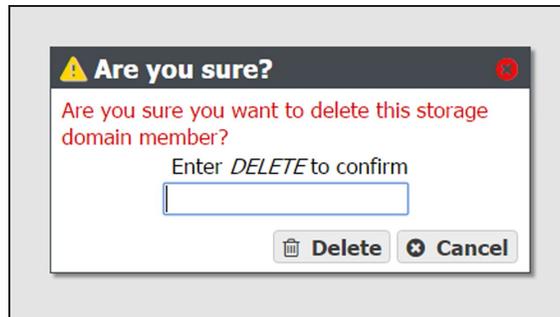


Figure 36 The delete storage domain member confirmation dialog box.

4. Type `DELETE` into the entry field, and then click **Delete**.

Edit a Storage Domain

Use the instructions in this section to change the parameters for a storage domain.



IMPORTANT

If this storage domain is assigned to a data policy that uses versioning, after data is persisted, you cannot change this setting from Object ID to Object Name.

Note: If an edit you select is not allowed, the gateway generates an error message when you click **Save**, explaining why the edit is not allowed.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).

2. Select the row for the storage domain that you want to edit and then select **Action > Edit**. The Edit storage domain screen displays

Note: Alternatively, select **Action > Edit Storage Domain** from the storage domain details screen.

Figure 37 The Edit Storage Domain dialog box.

3. Edit the fields as necessary. See [Create a Storage Domain on page 75](#) for a description of each field.
4. Click **Save**. The edited storage domain displays on the Advanced Bucket Management screen.

Delete a Storage Domain

Use the instructions in this section to delete a storage domain.

Note: You cannot delete a storage domain if it is used by a data policy. See [Delete a Data Replication Rule on page 115](#) for instructions on removing the storage domain from a data policy.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6 on page 57](#)).

2. Select the row for the storage domain that you want to delete and then select **Action > Delete**. A confirmation dialog box displays.

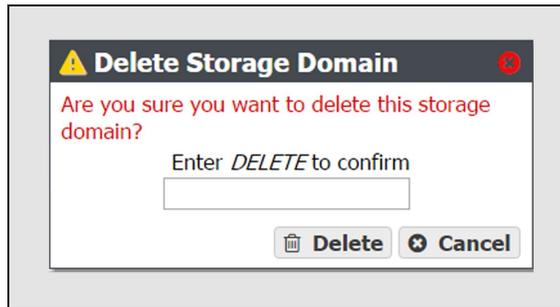


Figure 38 The delete storage domain confirmation dialog box.

3. Type `DELETE` into the entry field, and then click **Delete**.

MANAGE DATA REPLICATION RULES

Edit a BlackPearl Data Replication Rule

Use the instructions in this section to change the type, data policy, or whether to replicate deletes, for a BlackPearl data replication rule.

Note: Some edits are restricted based on whether a bucket is currently using the data policy. If an edit is not allowed, an error message displays when you click **Save**.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the data policy with the BlackPearl data replication rule that you want to edit, or select the data policy and then select **Action > Show Details**. The Data Policy details screen displays (see Figure 18 on page 87).
3. Double-click the row of the BlackPearl data replication rule that you want to edit, or select the row for the replication rule and then select **Action > Edit Rule**.

Note: Do not click the target name when selecting the data persistence rule row or the Replication Targets screen will open.

The Edit BlackPearl Data Replication Rule dialog box displays.

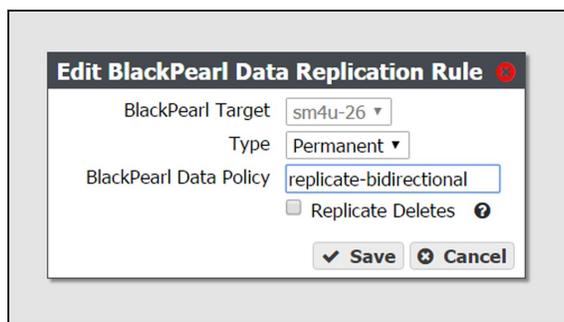


Figure 39 The Edit BlackPearl Data Replication Rule dialog box.

4. The BlackPearl target name is unavailable and cannot be changed.
5. Edit the fields as necessary. See [Add a BlackPearl Data Replication Rule to a Data Policy on page 88](#) for a description of each field.
6. Click **Save**. The edited BlackPearl data replication rule displays on the Data Policy details screen.

Edit an Amazon S3 Data Replication Rule

Use the instructions in this section to change the type, Initial Data Placement, whether to Replicate Deletes, or the Max Blob Size for an Amazon S3 data replication rule.

Note: Some edits are restricted based on whether a bucket is currently using the data policy. If an edit is not allowed, an error message displays when you click **Save**.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the data policy with the Amazon S3 replication rule that you want to edit, or select the data policy and then select **Action > Show Details**. The Data Policy details screen displays (see Figure 18 on page 87).
3. Double-click the row of the Amazon S3 data replication rule that you want to edit, or select the row for the replication rule and then select **Action > Edit Rule**.

Note: Do not click the target name when selecting the data persistence rule row or the Replication Targets screen will open.

The Edit Amazon S3 Data Replication Rule dialog box displays.

Figure 40 The Edit Amazon S3 Data Replication Rule dialog box.

4. The Amazon S3 target name is unavailable and cannot be changed.
5. Edit the fields as necessary. See [Add an Amazon S3 Data Replication Rule to a Data Policy](#) on page 90 for a description of each field.

Note: If you are configuring the replication to target a bucket that was previously created using the Amazon AWS interface, you must define a Lifecycle Management Rule in AWS to migrate data from the bucket default tier to the preferred tier, if necessary. Spectra Logic recommends using an immediate (0 days) move rule.

6. Click **Save**. The edited Amazon S3 data replication rule displays on the Data Policy details screen.

Edit a Microsoft Azure Data Replication Rule

Use the instructions in this section to change the type, whether to Replicate Deletes, or the Max Blob Size for a Microsoft Azure data replication rule.

Note: Some edits are restricted based on whether a bucket is currently using the data policy. If an edit is not allowed, an error message displays when you click **Save**.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the data policy with the Microsoft Azure replication rule that you want to edit, or select the data policy and then select **Action > Show Details**. The Data Policy details screen displays (see Figure 18 on page 87).
3. Double-click the row of the Microsoft Azure data replication rule that you want to edit, or select the row for the replication rule and then select **Action > Edit Rule**.

Note: Do not click the target name when selecting the data persistence rule row or the Replication Targets screen will open.

The Edit Microsoft Azure Data Replication Rule dialog box displays.

Figure 41 The Edit Microsoft Azure Data Replication Rule dialog box.

4. The Microsoft Azure target name is unavailable and cannot be changed.
5. Edit the fields as necessary. See [Add a Microsoft Azure Data Replication Rule to a Data Policy on page 91](#) for a description of each field.
6. Click **Save**. The edited Microsoft Azure data replication rule displays on the Data Policy details screen.

Delete a Data Replication Rule

Use the instructions in this section to delete a data replication rule for any type of target from a data policy.

Note: You cannot delete a data replication rule if the data policy is used by a bucket. See [Delete a Bucket on page 104](#) for instructions for deleting buckets using a data policy.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6 on page 57](#)).
2. Double-click the data policy with the replication rule that you want to delete, or select the data policy and then select **Action > Show Details**. The Data Policy details screen displays (see [Figure 18 on page 87](#)).
3. Select the row for the replication rule and then select **Action > Delete Rule**.

Note: Do not click the target name when selecting the data persistence rule row or the Replication Targets screen will open.

A confirmation dialog box displays.

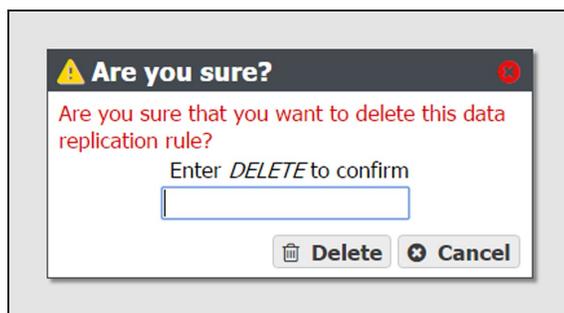


Figure 42 The delete data replication rule confirmation dialog box.

4. Type `DELETE` into the entry field, and then click **Delete**.

MANAGE A DATA POLICY

Use the instructions in this section to manage access control lists (ACLs), and to edit or delete persistence rules, replication rules, and data policies.

Edit a Data Policy

Use the instructions in this section to edit a data policy.

**IMPORTANT**

If you previously configured the data policy Versioning setting to **Keep Multiple Versions**, you are not able to change this setting after the data policy was created.

Note: If an edit you selected is not allowed, the gateway generates an error message, when you click **Save**, explaining why the edit is not allowed. For example, you cannot change the **Checksum type** if the data policy is used by a bucket.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6 on page 57](#)).

2. Select the row for the data policy that you want to edit and then select **Action > Edit**. The Edit Data Policy screen displays.

The screenshot shows the 'Edit' dialog box for a data policy. The dialog box is titled 'Edit' and contains the following fields and options:

- Data Policy Name:** Dual Copy on Tape
- Blobbing Enabled
- Minimize Spanning
- Performance:**
 - Default GET Priority: High
 - Default PUT Priority: Normal
 - Default VERIFY Priority: Low
 - Default Verify After Write
 - Rebuild Priority: Low
- Data Security:**
 - Checksum type: MD5
 - Require end-to-end CRC
- Versioning:**
 - None
 - Keep Latest
 - Keep Multiple Versions
- Data Replication Rules:**
 - Always accept replicated PUT jobs

At the bottom right, there are 'Save' and 'Cancel' buttons.

Figure 43 The Edit Data Policy dialog box.

3. Edit the fields as necessary. See [Create a Data Policy on page 83](#) for a description of each field.
4. Click **Save**. The edited data policy displays on the Advanced Bucket Management screen.

Delete a Data Policy

Use the instructions in this section to delete a data policy.

Note: You cannot delete a data policy if any buckets exist that use the specified data policy. See [Delete a Bucket on page 104](#) for instructions for deleting buckets using a data policy.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6 on page 57](#)).

2. Select the row for the data policy that you want to delete and then select **Action > Delete**. A confirmation dialog box displays.

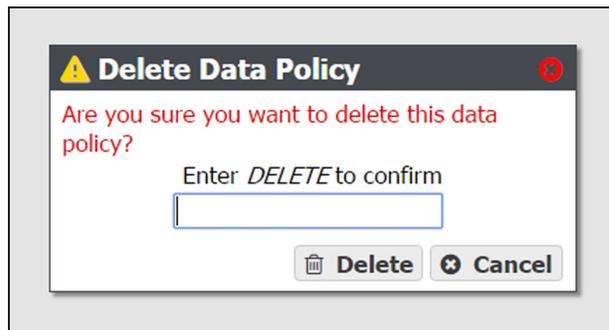


Figure 44 The delete data policy confirmation dialog box.

3. Type `DELETE` into the entry field, and then click **Delete**.

Delete a Data Policy ACL

Use the instruction in this section to delete a data policy ACL.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Select the row for the data policy for which you want to delete an ACL, then select **Action > Show Details**. The Data Policy details screen displays.
3. Select the row for the data policy ACL which you want to delete.
4. From the menu bar, select **Action > Delete Data Policy ACL**. A confirmation window displays.
5. Click **Delete**.

Edit a Data Persistence Rule

Use the instructions in this section to change the type, Isolation Level, or Minimum Days to Retain for a data persistence rule.

Note: Some edits are restricted based on whether a bucket is currently using the data policy. If an edit is not allowed, an error message displays when you click **Save**.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Double-click the data policy with the persistence rule that you want to edit, or select the data policy and then select **Action > Show Details**. The Data Policy details screen displays (see Figure 18 on page 87).

3. Double-click the row of the persistence rule that you want to edit, or select the row for the persistence rule and then select **Action > Edit Rule**.

Note: Do not click the storage domain name when selecting the data persistence rule row or the Storage Domain details screen will open.

The Edit Data Persistence Rule dialog box displays.

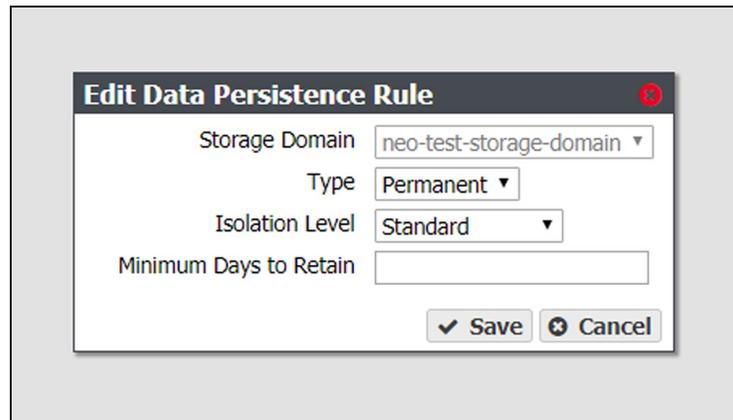


Figure 45 The Edit Data Persistence Rule dialog box.

4. The **Storage Domain** setting is unavailable and cannot be changed.
5. Edit the fields as necessary. See [Add Data Persistence Rules and Replication Rules to a Data Policy](#) on page 86 for a description of each field.
6. Click **Save**. The edited data persistence rule displays on the Data Policy details screen.

Delete a Data Persistence Rule

Use the instructions in this section to delete a data persistence rule from a data policy.



CAUTION Deleting a persistence rule deletes all data associated with the rule.

- Notes:**
- You cannot delete a data persistence rule if the data policy is used by a bucket. See [Delete a Bucket](#) on page 104 for instructions for deleting buckets using a data policy.
 - You cannot delete a persistence rule if it is the last permanent persistence rule for a data policy.
1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6](#) on page 57).

2. Double-click the data policy with the persistence rule that you want to delete, or select the data policy and then select **Action > Show Details**. The Data Policy details screen displays (see Figure 18 on page 87).
3. Select the row for the persistence rule and then select **Action > Delete Rule**.

Note: Do not click the storage domain name when selecting the data persistence rule row or the Storage Domain details screen will open.

A confirmation dialog box displays.

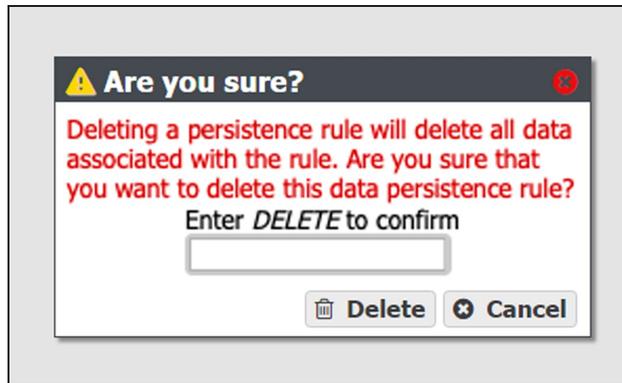


Figure 46 The delete data persistence rule confirmation dialog box.

4. Type `DELETE` into the entry field, and then click **Delete**.

MANAGE REPLICATION TARGETS

Use the instruction in this section to manage existing replication targets.

Edit a BlackPearl Replication Target

Use the instructions in this section to modify the configuration of an existing BlackPearl replication target.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen (see [Figure 6](#) on page 57).
2. Double-click the BlackPearl replication target that you want to edit, or select the replication target and then select **Action > Edit**. The Edit BlackPearl Target dialog box displays.

The screenshot shows the 'Edit BlackPearl Target' dialog box with the following fields and values:

- Name:** BlackPearl2
- Data Path End Point:** sm2u-10.eng
- Data Path Port:** (dropdown menu)
- BlackPearl Settings:**
 - Administrator Username or S3 Access ID:** QWRtaW5pc3RyYXRvcg==
 - Administrator S3 Secret Key:** bfitest
 - Default Read Preference:** Last Resort
 - Access Control Replication:** None
 - Replicated User Default Data Policy:** (empty field)
- Proxy Server:**
 - Data Path Proxy Server:** (empty field)

Buttons at the bottom right: Save, Cancel.

Figure 47 The Edit BlackPearl Target dialog box.

3. Edit the fields as necessary. See [Create a BlackPearl Target](#) on page 64 for a description of each field.
4. Click **Save**. The edited BlackPearl replication target displays on the Replication Targets screen.

Edit an Amazon Replication Target

Use the instructions in this section to modify the configuration of an existing Amazon replication target.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen (see Figure 6 on page 57).
2. Double-click the Amazon replication target that you want to edit, or select the replication target and then select **Action > Edit**. The Edit Amazon S3 Target dialog box displays.

Edit Amazon S3 Target

Name: Bills AWS

Data Path End Point:

Region: US_WEST_2

HTTPS

Amazon S3 Settings

Access Key: AKIAI4H2EXZ2HOWJVDTQ

Secret Key: iZrPqAB/2BgSyXjgDzHbjms

Cloud Bucket Prefix:

Cloud Bucket Suffix:

Staged Data Expiration: 30 Days

BlackPearl Settings

Default Read Preference: Minimum Latency

Proxy Server

Proxy Domain:

Proxy Host:

Proxy Port:

Proxy Username:

Proxy Password:

Save Cancel

Figure 48 The Edit Amazon S3 Target dialog box.

3. Edit the fields as necessary. See [Create an Amazon S3 Target](#) on page 68 for a description of each field.
4. Click **Save**. The edited Amazon replication target displays on the Replication Targets screen.

Edit an Azure Replication Target

Use the instructions in this section to modify the configuration of an existing Azure replication target.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen (see Figure 6 on page 57).
2. Double-click the Azure replication target that you want to edit, or select the replication target and then select **Action > Edit**. The Edit Microsoft Azure Target dialog box displays.

The screenshot shows a dialog box titled "Edit Microsoft Azure Target". It contains the following fields and controls:

- Name:** Bills Azure
- HTTPS:**
- Microsoft Azure Settings:**
 - Account Name:** spectradvt
 - Account Key:** e7aRkpYLGFNIEIKkB33D/s
 - Cloud Bucket Prefix:** (empty)
 - Cloud Bucket Suffix:** (empty)
- BlackPearl Settings:**
 - Default Read Preference:** Minimum Latency

At the bottom right, there are two buttons: "Save" and "Cancel".

Figure 49 The Edit Microsoft Azure Target dialog box.

3. Edit the fields as necessary. See [Create a Microsoft Azure Target](#) on page 72 for a description of each field.
4. Click **Save**. The edited Azure replication target displays on the Replication Targets screen.

Verify a Replication Target

Use the instructions in this section to verify connectivity to the target and optionally verify replicated data on the replication target.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen (see Figure 6 on page 57).

2. Double-click the replication target for which you want to verify connectivity, or select the replication target and then select **Action > Verify Data**. The *Verify target type* Target dialog box displays.

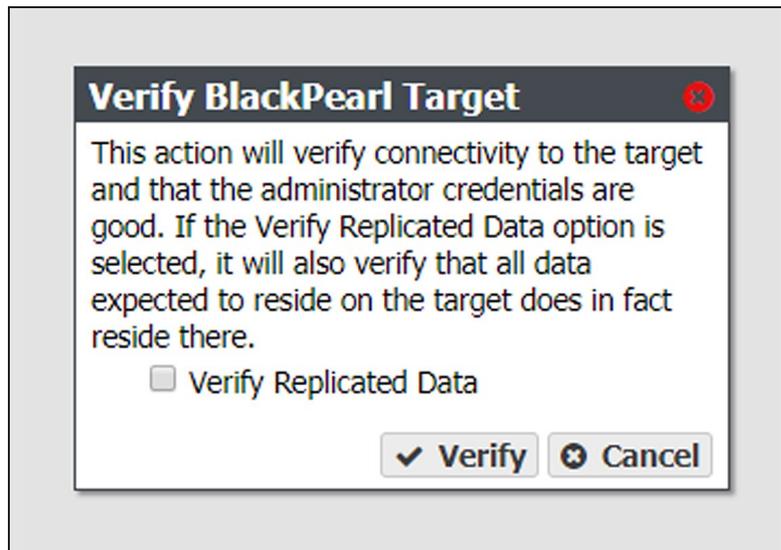


Figure 50 The *Verify target type* Target dialog box.

3. If desired, select **Verify Replicated Data** to confirm that the expected data resides on the replication target.
Note: Depending on the amount of data on the replication target, this process may take a long time to complete.
4. Click **Verify**. The gateway confirms connectivity to the target and optionally verifies the replicated data.

Put a Replication Target in Standby State

If you need to perform service on a replication target, it is recommended that you first put the replication target into a standby state. Otherwise, the BlackPearl gateway may attempt to use the target while it is in service.

Use the instructions in this section to place a replication target into a standby state. No data is transferred to the replication target while in standby.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen (see [Figure 6](#) on page 57).

2. Select the replication target that you want to put into standby and then select **Action > Put Target in Standby**. The Put Target in Standby dialog box displays.



Figure 51 The Put Target in Standby dialog box.

3. Click **Deactivate**. The target is now in standby.

Activate a Replication Target

Use the instructions in this section to activate a replication target currently in standby. Once activated, data transfers are allowed to the replication target.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen (see Figure 6 on page 57).
2. Select the replication target that you want to activate and then select **Action > Activate Target**. The Activate Target dialog box displays.

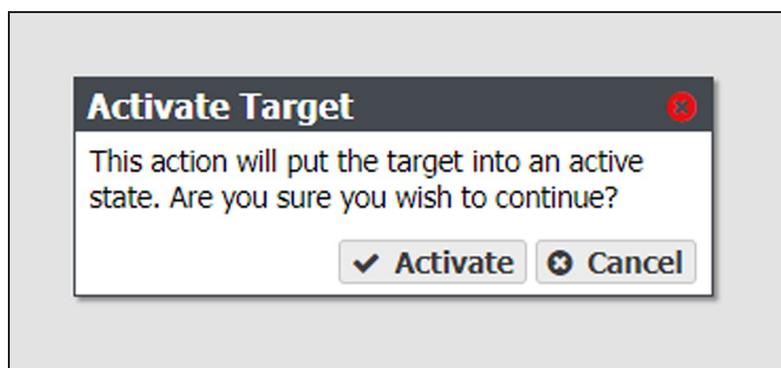


Figure 52 The Activate Target dialog box.

3. Click **Activate**. The target is now in an active state.

Delete a Replication Target

Use the instructions to delete an existing replication target.



CAUTION If you delete a replication target, all data on the target is deleted.

Note: You cannot delete a replication target if it is used by a data policy. See [Delete a Data Replication Rule on page 115](#) for instructions on removing a replication target from a data policy.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Replication Targets** to display the Replication Targets screen (see [Figure 6 on page 57](#)).
2. Select the replication target that you want to delete and then select **Action > Delete**. The Delete *target type* Target dialog box displays.

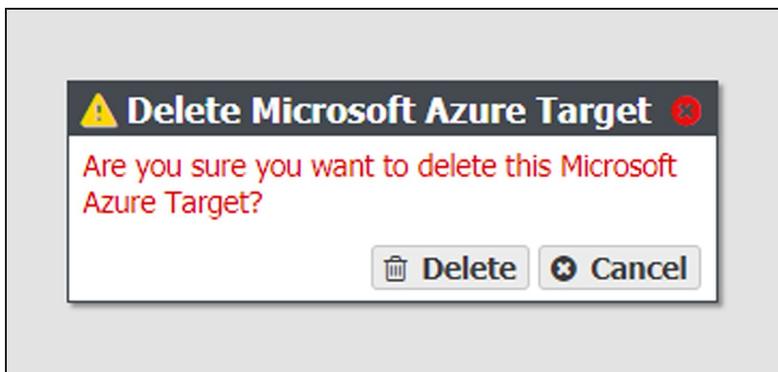


Figure 53 The Delete *target type* Target dialog box.

3. Click **Delete** to delete the selected replication target.

MANAGE A DISK PARTITION

Use the instructions in this section to edit or delete a disk partition.

Edit a Disk Partition

Use the instructions in this section to change the parameters for a disk partition.

- Notes:**
- You cannot remove a member pool that contains data.
 - You cannot remove the last member pool of a disk partition assigned to a storage domain.
1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
 2. Select the disk partition row and then select **Action > Edit**. The Edit Disk Partition dialog box displays.

Name	Capacity	S3 Pool Type	Health	
Arctic_Blue_1_32499795877	129.2 TB	Nearline	Online	

Figure 54 The Edit Disk Partition dialog box.

3. Edit the fields as necessary. See [Create a Disk Pool](#) on page 55 for a description of each field.
4. Click **Save**. The edited disk partition displays on the Advanced Bucket Management screen.

Delete a Disk Partition

Use the instructions in this section to delete a disk partition.

Note: You cannot delete a disk partition that is a member of any storage domain. See [Delete a Storage Domain Member on page 108](#) for information on deleting the disk partition from a storage domain.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Data Policy Management** to display the Advanced Bucket Management screen (see [Figure 6 on page 57](#)).
2. Select the row of the disk partition you want to delete and then select **Action > Delete**. A confirmation dialog box displays.



Figure 55 The Delete Disk Partition confirmation dialog box.

3. Type `DELETE` into the entry field, and then click **Delete**.

MANAGE ONLINE AND NEARLINE DISK POOLS

Use the instructions in this section to manage nearline and online disk pools. For information on managing network attached storage (NAS) disk pools, see .

Edit a Nearline or Online Disk Pool

If desired, you can edit a previously configured nearline or online disk pool to enable or disable power saving mode, encryption, or to add write performance drives to the disk pool.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
2. Select the disk pool you want to edit and select **Action > Edit**. The Edit *disk pool name* dialog box displays.

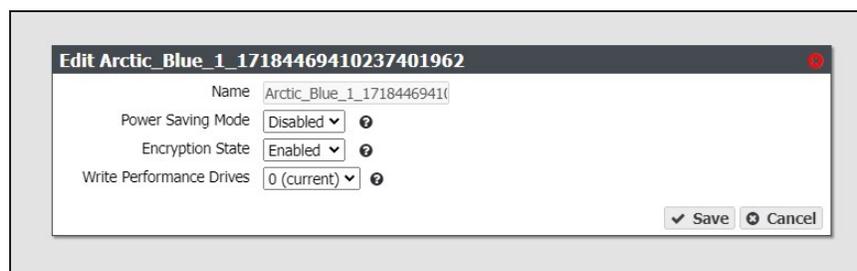


Figure 56 The Edit *disk pool* dialog box.

3. If desired, enable or disable **Power Saving Mode** for the disk pool. Enabling this feature configures the standby timer to 60 minutes. When there is no I/O to the storage pool for 60 minutes, the drives in the pool spin down and use minimal power.

Notes:

- Spectra Logic recommends leaving power saving mode **disabled**.
- To use this feature, all drives in the storage pool must be power-saving compatible.

4. If desired, using the **Encryption State** drop-down menu, select **Enabled** to encrypt the drives in the pool or **Disabled** to decrypt the drives.
5. If desired, using the **Write Performance Drives** drop-down menu to select the number of write performance drives you want to add to the disk pool.

Note: These drives are permanently part of the storage pool and cannot be removed.

6. Click **Save**.

Import a Nearline or Online Disk Pool

Data present on storage pools on expansion nodes can be moved from one BlackPearl gateway to another, either to increase name space redundancy, or as part of disaster prevention. Use the instructions in this section to import a pool.

1. Ensure the expansion node containing the pool you want to import is cabled to the BlackPearl gateway master node.
2. Ensure the settings for the DS3 service **Default Conflict Resolution Mode** are set to your preference. See for more information.
3. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Policy Management** to display the Advanced Bucket Management screen (see Figure 6 on page 57).
4. Disk pools eligible for import display with a health of **Foreign**. Select the nearline pool you want to import in the Storage Pools pane, and select **Action > Import Pool**. The Import dialog box displays.

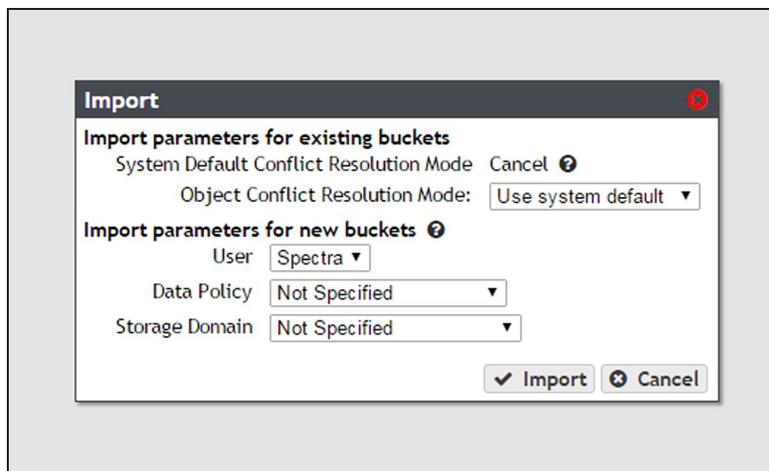


Figure 57 The Import dialog box.

5. Configure the import parameters for existing buckets present on the storage pool.
 - a. The **System Default Conflict Resolution Mode** is configured when editing the DS3 service and is not editable in this dialog box. See for more information.
 - b. Using the drop-down menu, select a behavior for **Object Conflict Resolution Mode**. This setting configures how the gateway handles a file name conflict when importing foreign objects on the storage pool to an existing bucket.

Value	Description
Use system default	The gateway uses the behavior displayed in the System Default Conflict Resolution Mode row of the Import dialog box.
Cancel	Abort the import process if a file name conflict is discovered.
Accept Most Recent	Keep the file with the most recent creation date.
Accept Existing	Keep the file currently in the BlackPearl database.

6. Configure the import parameters for new buckets present on the storage pool.
 - a. Using the drop-down menu, select the **User** to be the owner of the imported bucket (s).
 - b. Using the drop-down menu, select the **Data Policy** to use for the imported bucket (s).
 - c. Using the drop-down menu, select the **Storage Domain** to use for the imported bucket(s).
7. Click **Import**.

Delete a Nearline or Online Disk Pool

If you want to delete an online or nearline disk pool (for example, to create a larger pool after adding additional disks to the gateway) use the instructions in this section to delete a disk pool.

Note: You must delete all objects contained in the disk pool before you can delete the pool.

1. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Policy Management** to display the Advanced Bucket Management screen (see [Figure 6](#) on page 57).

2. Select the pool you want to delete in the Storage Pools pane, and select **Action > Delete**. A confirmation dialog box displays.



Figure 58 The Delete Nearline Disk Pool confirmation dialog box.

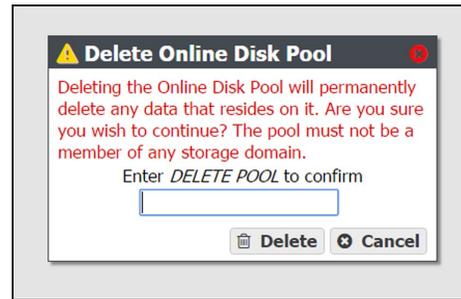


Figure 59 The Delete Online Disk Pool confirmation dialog box.

3. Enter `DELETE POOL` into the entry field and click **Delete**. The disk pool is deleted.

CHANGE THE BLACKPEARL GATEWAY CACHE CONFIGURATION

If desired, you can change the configuration of the BlackPearl gateway cache to allow for better performance, or create a larger or smaller cache suitable for your workflow environment. You can also add new disks installed in the gateway to the system cache.

Note: Changing the system cache should only be done after careful consideration of your desired workflow.

After changing the BlackPearl cache configuration, all data in the cache is deleted, and the system reboots.

**CAUTION**

All data currently in the system cache is deleted. As long as all jobs are complete before you change the system cache configuration, no data is lost.

Use the instructions in this section to change the system cache configuration.

1. Discontinue data transfers to the BlackPearl gateway and ensure that all jobs complete (see). Migrate any data you want to keep off of the system cache.
2. From the menu bar, select **Configuration > Advanced Bucket Management > Storage & Policy Management** to display the Advanced Bucket Management screen (see [Figure 6](#) on [page 57](#)).

3. In the Storage Pools pane, select the BlackPearl_Cache, and then select **Action > Edit**. The Change Cache Configuration dialog box displays.

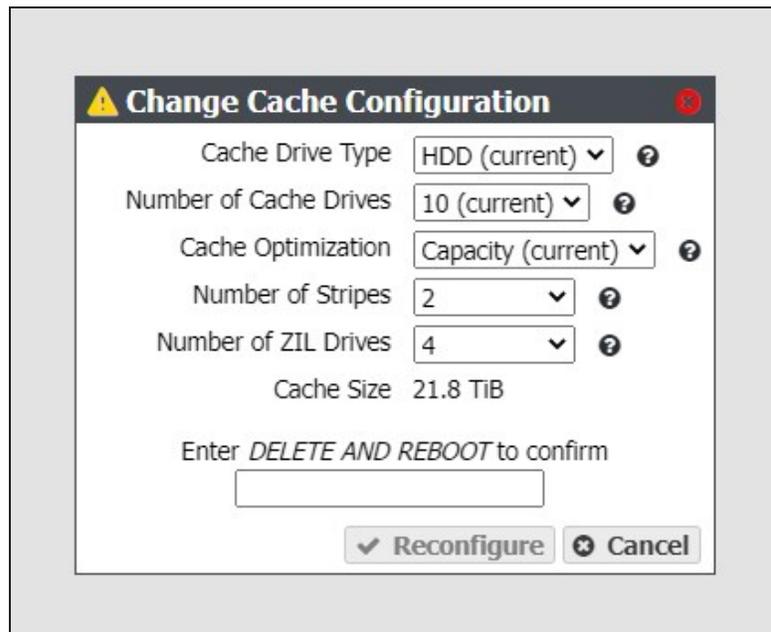


Figure 60 The Change Cache Configuration dialog box.

4. If desired, change the **Cache Drive Type**. You cannot mix drive types in the system cache.
5. If desired, change the **Number of Cache Drives**. The number of drives selected affects the number of stripes that can be used in Capacity mode.
6. Select the method of **Cache Optimization**. The default is Performance mode. Use Capacity mode to create a larger cache capable of storing a greater number of objects.
7. If you selected Capacity mode in [Step 6](#), select the **Number of Stripes** to use for the cache.
8. Select the **Number of ZIL Drives** (high-performance drives) to assign to the BlackPearl cache.
9. Enter `DELETE AND REBOOT` in the dialog box, and click **Reconfigure**.



CAUTION

All data currently in the system cache is deleted. As long as all jobs are complete before you change the system cache configuration, no data is lost.

10. The BlackPearl gateway creates a new system cache and reboots.