



SPECTRA VAIL USER GUIDE



www.SpectraLogic.com

TABLE OF CONTENTS

Table Of Contents	2
Document Information	8
Copyright	9
Notices	9
Trademarks	9
Master License Agreement	10
Contacting Spectra Logic	16
Introduction	17
Related Publications	17
What's New	18
Chapter 1 - Understanding the Vail Application	21
Vail Application Features	22
Understanding BlackPearl Storage	24
Types of BlackPearl Storage	24
Understanding Vail Cloud Storage	25
Vail Application with AWS Cloud Storage	25
Vail Application with Microsoft Azure Cloud Storage Features	26
Vail Application with Google Cloud Storage	26
Understanding Vail Lifecycles	28
Considerations for Placement Rules	28
Considerations for Delete Rules	29
Understanding Object Storage	30
Understanding Vail Buckets	31
Considerations for Vail Bucket Creation	31
Understanding S3 with the Vail Application	32
Interfacing with Tape Libraries and Media	33
Tape Ingest	33
Spectra Logic Vail Application Naming Conventions	35
Vail Sphere Names	35
User Names	35
Group Names	35

Location Names	35
Portable Location Names	36
Storage Names	36
Spectra Vail Bucket Names	37
Cloud Bucket Names	37
BlackPearl Bucket Names	37
Lifecycle Names	38
Additional AWS Account Role Names	38
Vail Application Technical Information	39
Maximum Part and Object size	39
Chapter 2 - Important Information	40
Requirements	41
Spectra Logic Products Requirements	41
Supported Browsers	41
Vail Management Console Overview - Local Control	42
Main window	42
Taskbar	43
Toolbar	44
Embedded Dashboard	44
Icons	45
Vail Management Console Overview - Cloud Control	46
Main window	46
Taskbar	47
Toolbar	48
Icons	49
Chapter 3 - Configure BlackPearl System	50
Configure a BlackPearl Nearline Gateway	51
Register a BlackPearl Nearline Object Gateway with a Vail Sphere	52
Register A Vail Sphere - Local Control	52
Register A Vail Sphere - Cloud Control	56
Chapter 4 - Configure the Spectra Vail Application	61
Log In to the Vail Management Console	62
Create Storage	63
Create BlackPearl Storage	64
Create BlackPearl Standard Bucket Storage	64

Create BlackPearl Legacy Storage	67
Create BlackPearl Volume Pool Storage	70
Create Cloud Storage	74
Create Amazon S3 Cloud Storage	74
Create Microsoft Azure Cloud Storage	79
Create Google Cloud Platform Storage	83
Create S3 Compatible Cloud Storage	86
Create Storage Groups	90
Considerations for Expanding Storage	90
Create a Storage Group	90
Create a Lifecycle	91
Create a Vail Bucket	100
Configure an Object Storage Browser	108
Configure S3 Browser	108
Configure Cyberduck Object Storage Browser	109
Chapter 5 - Configure & Manage Users	111
Configure & Manage Sphere Administrator - Cloud Control	113
Create a Sphere Administrator	113
Change a Sphere Administrator Password	115
Edit Sphere Administrator Attributes	117
Delete a Sphere Administrator	119
Configure & Manage Vail Administrator - Local Control	120
Create a Vail Administrator	120
Change a Vail Administrator Password	123
Delete a Vail Administrator	123
Configure & Manage IAM Accounts	124
Add an IAM Account	124
View IAM Account Details	130
Edit an IAM Account	132
Delete an IAM Account Association	133
Configure & Manage IAM Users and Groups	134
Create an IAM User	134
View IAM User Details	135
Add an IAM User to an IAM Group	136
Remove an IAM User from an IAM Group	137

Delete an IAM User	138
Create an IAM Group	139
Delete an IAM Group	140
Create an IAM Group Policy	141
Edit an IAM Group Policy	142
Delete an IAM Group Policy	143
AWS Access Key Management	144
Create an Access Key	144
Enable an Access Key	145
Disable an Access Key	146
Delete an Access Key	147
Chapter 6 - Using the Spectra Vail Application	149
View Capacity Information	151
View Performance Metrics	154
View Vail Bucket Details	156
View Vail Bucket Contents	160
View Object Details	162
Create an Object Clone	166
Verify an Object Clone	169
Delete an Object Clone	171
Edit Global Settings	173
Change Lifecycle Rule Nightly Processing Time	173
Set Default Lifecycle	174
Enable Diagnostic Monitor	174
Configure AWS Infrastructure	176
Using Proxy Connections	177
Configure Proxy Connection	177
Edit Proxy Server	178
Delete Proxy Server	178
Edit a Vail Bucket	179
Delete a Vail Bucket	183
View Storage Details	184
Edit BlackPearl or Vail VM Endpoint	187
Change Endpoint Location	187
Add Additional Host Names	189

Change Endpoint URL	190
Configure Debug Logging	191
Edit Storage	193
Edit BlackPearl Bucket Storage	193
Edit BlackPearl Volume Pool Storage	196
Edit Vail VM Node Storage	198
Edit Google Cloud Platform Storage	201
Edit AWS S3 Cloud Storage	204
Edit Microsoft Azure Cloud Storage	208
Edit Other S3 Cloud Storage	212
Consolidate Storage	216
Delete Storage	217
View Lifecycle Details	221
Edit a Lifecycle	225
Delete a Lifecycle	228
Create a Location	229
Delete a Location	232
Clear the IAM Cache	233
View Reports	234
View Spectra Vail Application Messages	236
Message Details	238
Spectra Vail Application Logs	239
Update the Spectra Vail Application Software	240
Accessing the Technical Support Portal	243
Create an Account	243
Log Into the Portal	244
Opening a Support Ticket	245
Search for Help Online	245
Submit an Incident Online	248
Submit an Incident by Phone	250
Appendix A - BlackPearl Embedded Dashboard	251
Using the Embedded BlackPearl Dashboard	252
View the Status of the BlackPearl System	253
View System Overview	253
View Notifications	254

View Jobs	255
View Buckets	256
View Pools	257
View Volumes	258
View Tape Partitions - Main View	259
View Tape Partitions - Tape State View	260
View Tape Drives	261
View Tape Management	262
Dashboard Actions	263
Create a Volume Snapshot	263
Export a Tape Cartridge	263
Online a Tape Cartridge	264
Verify a Tape Cartridge	264
Change Job Priority	265
Create a Bucket	265
Start a Storage Pool Verification	266
Put a Tape Partition into Standby	266
Offline a Tape Drive	266
Appendix B - Create and Configure a Vail VM Node	267
Create a Vail VM Node	268
Vail VM Node Host Requirements	268
Create a Node Using VMWare vSphere	269
Create a Node Using Oracle VirtualBox	276
Configure the Vail VM Node Network Settings	283
Configure Network Settings	284
Configure the Vail VM Node Hostname	287
Configure the SSL Certificate	288
Register a Vail VM Node with a Vail Sphere	289
Frequently Asked Questions	295
Best Practices	296
Glossary	297

DOCUMENT INFORMATION

Document part number:

- 90990149

Document revision:

- Revision I

Document revision history:

Revision	Date	Description
A	June 2022	Initial Release
B	October 2022	Updated for Vail 2.0.0.
C	April 2023	Updated for Vail 2.3.0.
D	October 2023	Updated for Vail 2.5.4.
E	March 2024	Updated for Vail 3.0.1.
F	September 2024	Updated for Vail 3.1.3.
G	December 2024	Updated for Vail 3.2.0.
H	June 2025	Updated for Vail 3.4.0.
I	December 2025	Updated for Vail 3.5.0.

COPYRIGHT

Copyright © 2022-2025 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

ArcticBlue, BlackPearl, BlueScale, RioBroker, Spectra Cube, Spectra Logic, Spectra Vail, Spectra, SpectraGuard, 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.

MASTER LICENSE AGREEMENT

This Master License Agreement governs use of Spectra Logic Corporation stand-alone software such as StorCycle software ("Software"). Your organization has agreed to the license contained herein and terms and conditions of this Master License Agreement (the "MLA"). Use of the Software is affirmation of your acceptance and grants to your organization ("Licensee") the right to use the Software.

1. License.

1.1 Grant of License. Subject to all of the terms and conditions of this MLA, Spectra Logic Corporation and its wholly-owned subsidiaries ("Spectra") grant to Licensee a non-transferable, non-sublicensable, non-exclusive license during the applicable Term (as defined below) to use the object code form of the Software specified in the quote supplied either by Spectra or an authorized reseller internally and for operational use, and only in accordance with the technical specification documentation generally made available by Spectra to its licensees with regard to the Software ("Documentation"). The term "Software" will include any Documentation and any ordered Support and maintenance releases of the same specific Software product provided to Licensee under this MLA.

1.2 Term and Renewals. The Software is licensed under a subscription basis or is permanently licensed, as defined herein. Licensee's Software license is stated on the quote provided to Licensee.

(a) If the Software is ordered on a subscription basis ("Subscription"), the term of the software license will (i) commence upon receipt of a purchase order issued to Spectra directly from Licensee or from an authorized reseller issued on your behalf and will (ii) continue for the number of year(s) noted on the quote commencing on the date of activation of key(s) performed by Spectra Professional Services ("Subscription Term"). Unless terminated earlier in accordance with section 4, each Software Subscription Term will automatically renew upon expiration of the initial Software Subscription Term for additional successive terms unless either party gives the other prior written notice of cancellation at least thirty (30) days prior to expiration of the then-current term. Unless otherwise specified on the quote, the license fee for any Software Subscription Term renewal will be based on the then-current Subscription rates.

(b) If the Software is ordered on a permanent license basis ("Permanent"), the term of the software license will not expire except in accordance with section 4. The term of associated products such as support, user, server and storage elections will commence upon on the date of activation of key(s) performed by Spectra Professional Services and may be renewed at such time as the term of such quoted election(s) expire.

1.3 Installation. Software may be installed on Licensee's computers only by Licensee's employees, authorized resellers, or by Spectra Professional Services as requested by Licensee.

1.4 License Restrictions. Licensee shall not (and shall not allow any third party) to

(a) decompile, disassemble, or otherwise reverse engineer the Software or attempt to reconstruct or discover any source code, underlying ideas, algorithms, file formats or programming interfaces of the Software by any means whatsoever (except and only to the extent that applicable law prohibits or restricts reverse engineering restrictions, and then only with prior written notice to Spectra), (b) distribute, sell, sublicense, rent, lease or use the Software (or any portion thereof) for time sharing, hosting, service provider or like purposes, (c) remove any product identification, proprietary, copyright or other notices contained in the Software, (d) modify any part of the Software, create a derivative work of any part of the Software, or incorporate the Software into or with other software, except to the extent expressly authorized in writing by Spectra, or (e) publicly disseminate Software performance information or analysis (including, without limitation, benchmarks).

2. Ownership.

Notwithstanding anything to the contrary contained herein, except for the limited license rights expressly provided herein, Spectra retains all rights, title and interest in and to the Software (including, without limitation, all patent, copyright, trademark, trade secret and other intellectual property rights) and all copies, modifications and derivative works thereof. Licensee acknowledges that it is obtaining only a limited license right to the Software and that irrespective of any use of the words "purchase", "sale" or like terms hereunder no ownership rights are being conveyed to Licensee under this MLA or otherwise.

3. Payment and Delivery.

3.1 Payment. All payments, either to Spectra or an authorized reseller, are non-refundable (except as expressly set forth in this MLA). Unless otherwise specified on the applicable quote, all license fees, support and Professional Services fees, if any, are due within thirty (30) days of date of invoice. Licensee shall be responsible for all taxes, withholdings, duties and levies arising from the order (excluding taxes based on the net income of Spectra). Any late payments shall be subject to a service charge equal to 1.5% per month of the amount due or the maximum amount allowed by law, whichever is less.

3.2 Delivery. Immediately upon receipt of a purchase order from Licensee or on behalf of Licensee or from an authorized reseller on behalf of Licensee, Licensee will have the right to access the Software. Software will be delivered by electronic means unless otherwise specified on the applicable quote. Spectra will contact Licensee and request its server identification number(s) and provide Activation code(s).

4. Term of MLA.

4.1 Term.

(a)(i) If Licensee ordered a Software Subscription License, this MLA expires on the day the Term of the Software expires. However, the ability to retrieve/restore archived data will continue indefinitely.

(ii) If a Permanent Software License was ordered, the software license does not expire.

(b) Section 4.1(a) is subordinate to this section 4.1(b). Either party may terminate this MLA if the other party (a) fails to cure any material breach of this MLA within thirty (30) days after written notice of such breach, (b) ceases operation without a successor; or (c) seeks protection under any bankruptcy, receivership, trust deed, creditors arrangement, composition or comparable proceeding, or if any such proceeding is instituted against such party (and not dismissed within sixty (60) days thereafter). Termination is not an exclusive remedy and the exercise by either party of any remedy under this MLA will be without prejudice to any other remedies it may have under this MLA, by law, or otherwise.

4.2 Survival. Sections 1.4 (License Restrictions), 2 (Ownership), 3 (Payment and Delivery), 4 (Term of MLA), 5.3 (Disclaimer), 8 (Limitation of Remedies and Damages), 10 (Confidential Information), 11 (General), and Licensee's right to Work Product and ownership of Licensee Content described in Section 7 shall survive any termination or expiration of this MLA.

5. Limited Warranty and Disclaimer.

5.1 Limited Warranty. Spectra warrants to Licensee that for a period of ninety (90) days from the effective date (the "Warranty Period"), the Software shall operate in substantial conformity with the Documentation. In addition, Spectra warrants that (i) it has the right to enter into and perform all obligations under this MLA, (ii) no agreement exists that restricts or conflicts with the performance of Spectra's rights and obligation hereunder, (ii) the technical information provided to Licensee is accurate and complete, and (iv) the Software is free from any third-party intellectual property infringement claims. Spectra does not warrant that Licensee's use of the Software will be uninterrupted or error-free, will not result in data loss, or that any security mechanisms implemented by the Software will not have inherent limitations. Spectra's sole liability (and Licensee's exclusive remedy) for any breach of this warranty shall be, in Spectra's sole discretion, to use commercially reasonable efforts to provide Licensee with an error-correction or work-around which corrects the reported non-conformity, to replace the non-conforming Software with conforming Software, or if Spectra determines such remedies to be impracticable within a reasonable period of time, to terminate the Software license and refund the license fee and support fee, if any, paid for the non-conforming Software. Spectra shall have no obligation with respect to a warranty claim unless notified of such claim within the Warranty Period.

5.2 Exclusions. The above warranty will not apply (a) if the Software is used with hardware or software not specified in the Documentation, (b) if any modifications are made to the Software by Licensee or any third party, (c) to defects in the Software due to accident, abuse or improper use by Licensee, or (d) to items provided on a no charge or evaluation basis.

5.3 Disclaimer. THIS SECTION 5 CONTAINS A LIMITED WARRANTY AND EXCEPT AS EXPRESSLY SET FORTH IN THIS SECTION 5 THE SOFTWARE AND ALL SERVICES ARE PROVIDED "AS IS". NEITHER SPECTRA NOR ANY OF ITS SUPPLIERS MAKES ANY OTHER WARRANTIES, EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. LICENSEE MAY HAVE OTHER STATUTORY RIGHTS. HOWEVER, TO THE FULL EXTENT PERMITTED BY LAW, THE DURATION OF STATUTORILY REQUIRED WARRANTIES, IF ANY, SHALL BE LIMITED TO THE LIMITED WARRANTY PERIOD.

6. Support.

Spectra will provide the support services identified in the quote ("Support"). Support services for the Subscription License will coincide with the license term.

7. Professional Services.

7.1 Professional Services. Professional Services may be ordered by Licensee pursuant to a quote describing the work to be performed, fees and any applicable milestones, dependencies and other technical specifications or related information. The parties acknowledge that the scope of the Professional Services provided hereunder consists solely of either or both of (a) assistance with Software installation, deployment, and usage or (b) development or delivery of additional related Spectra copyrighted software or code. Spectra shall retain all right, title and interest in and to any such work product, code or software and any derivative, enhancement or modification thereof created by Spectra (or its agents) ("Work Product").

7.2 Licensee Content. Licensee grants Spectra a limited right to use any Licensee materials provided to Spectra in connection with the Professional Services (the "Licensee Content") solely for the purpose of performing the Professional Services for Licensee. Licensee owns and will retain ownership (including all intellectual property rights) in the Licensee Content.

8. Limitation of Remedies and Damages.

8.1 NEITHER PARTY SHALL BE LIABLE FOR ANY LOSS OF USE, LOST DATA, FAILURE OF SECURITY MECHANISMS, INTERRUPTION OF BUSINESS, OR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING LOST PROFITS), REGARDLESS OF THE FORM OF ACTION, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE), STRICT LIABILITY OR OTHERWISE, EVEN IF INFORMED OF THE POSSIBILITY OF SUCH DAMAGES IN ADVANCE.

8.2 NOTWITHSTANDING ANY OTHER PROVISION OF THIS MLA, SPECTRA'S AND AUTHORIZED RESELLER'S, IF ANY, ENTIRE LIABILITY TO LICENSEE SHALL NOT EXCEED THE AMOUNT ACTUALLY PAID BY LICENSEE UNDER THIS MLA.

8.3 THIS SECTION 8 SHALL NOT APPLY WITH RESPECT TO ANY CLAIM ARISING UNDER THE SECTIONS TITLED "GRANT OF LICENSE," "LICENSE RESTRICTIONS" OR "CONFIDENTIAL INFORMATION."

9. Indemnification.

(a) Spectra shall defend, indemnify and hold harmless Licensee from and against any claim of infringement of a patent, copyright, or trademark asserted against Licensee by a third party based upon Licensee's use of the Software in accordance with the terms of this MLA, provided that Spectra shall have received from Licensee (i) prompt written notice of such claim (but in any event notice in sufficient time for Spectra to respond without prejudice), (ii) the exclusive right to control and direct the investigation, defense, and settlement (if applicable) of such claim, and (iii) all reasonably necessary cooperation of Licensee.

(b) If Licensee's use of any of the Software is, or in Spectra's opinion is likely to be, enjoined due to the type of infringement specified above, or if required by settlement, Spectra may, in its sole discretion (i) substitute for the Software substantially functionally similar programs and documentation, (ii) procure for Licensee the right to continue using the Software, or if (i) and (ii) are commercially impracticable, (iii) terminate the MLA and refund to Licensee the license fee.

(c) The foregoing indemnification obligation of Spectra shall not apply if the Software is modified by any person other than Spectra, but solely to the extent the alleged infringement is caused by such modification, if the Software is combined with other non-Spectra products or process not authorized by Spectra, but solely to the extent the alleged infringement is caused by such combination, to any unauthorized use of the Software, to any unsupported release of the Software, or to any open source software or other third-party code contained within the Software. THIS SECTION 9 SETS FORTH SPECTRA'S AND RESELLER'S, IF ANY, SOLE LIABILITY AND LICENSEE'S SOLE AND EXCLUSIVE REMEDY WITH RESPECT TO ANY CLAIM OF INTELLECTUAL PROPERTY INFRINGEMENT.

10. Confidential Information.

Each party agrees that all code, inventions, know-how, business, technical and financial information it obtains ("Receiving Party") from the disclosing party ("Disclosing Party") constitute the confidential property of the Disclosing Party ("Confidential Information"), provided that it is identified as confidential at the time of disclosure or should be reasonably known by the Receiving Party to be Confidential Information due to the nature of the information disclosed and the circumstances surrounding the disclosure. Any software, documentation or technical information provided by Spectra (or its agents), performance information relating to the Software, and the terms of this MLA shall be deemed Confidential Information of Spectra without any marking or further designation. Except as expressly authorized herein, the Receiving Party will hold in confidence and not use or disclose any Confidential Information except as necessary to carry out the purpose of this MLA. The Receiving Party's nondisclosure obligation shall not apply to information which the Receiving Party can document (a) was rightfully in its possession or known to it prior to receipt of the Confidential Information, (b) is or has become public knowledge through no fault of the Receiving Party, (c) is rightfully obtained by the Receiving Party from a third party without breach of any confidentiality obligation, (d) is independently developed by employees of the Receiving Party who had no access to such information, or (e) is required to be disclosed pursuant to a regulation, law or court order (but only to the minimum extent required to comply with such regulation or order and with advance notice to the Disclosing Party). The Receiving Party acknowledges that disclosure of Confidential Information would cause substantial harm for which damages alone would not be a sufficient remedy, and therefore that upon any such disclosure by the Receiving Party the Disclosing Party shall be entitled to appropriate equitable relief in addition to whatever other remedies it might have at law.

11. General.

11.1 Assignment. This MLA will bind and inure to the benefit of each party's permitted successors and assigns. Neither party shall assign this MLA (or any part thereof) without the advance written consent of the other party, except that either party may assign this MLA in connection with a merger, reorganization, acquisition or other transfer of all or substantially all of such party's assets or voting securities. Any attempt to transfer or assign this MLA except as expressly authorized under this section 11.1 is null and void.

11.2 Severability. If any provision of this MLA shall be adjudged by any court of competent jurisdiction to be unenforceable or invalid, that provision shall be limited to the minimum extent necessary so that this MLA shall otherwise remain in effect.

11.3 Governing Law; Jurisdiction and Venue. This MLA shall be governed by the laws of the State of Colorado and the United States without regard to conflicts of laws provisions thereof, and without regard to the United Nations Convention on the International Sale of Goods. Except where statutory laws prohibit Licensee from entering into arbitration or choice of laws, any dispute or claim relating in any way to Licensee's use of the Software, or of a copyright issue, or to any associated support services, will be resolved by binding arbitration in Denver, Colorado. The prevailing party in any action to enforce this MLA will be entitled to recover its attorneys' fees and costs in connection with such action.

11.4 Amendments; Waivers. No supplement, modification, or amendment of this MLA shall be binding, unless executed in writing by an authorized representative of both parties. No waiver will be implied from conduct or failure to enforce or exercise rights under this MLA. No provision of any purchase order or other business form employed by Licensee will supersede the terms and conditions of this MLA, and any such document relating to this MLA shall be for administrative purposes only and shall have no legal effect.

11.5 Force Majeure. Neither party shall be liable to the other for any delay or failure to perform any obligation under this MLA (except for a failure to pay fees) if the delay or failure is due to events which are beyond the reasonable control of such party, including but not limited to any strike, blockade, war, act of terrorism, riot, natural disaster, failure or diminishment of power or of telecommunications or data networks or services, or refusal of approval or a license by a government agency.

11.6 Export Compliance. Licensee acknowledges that the Software is subject to export restrictions by the United States government and import restrictions by certain foreign governments. Licensee shall not and shall not allow any third-party to remove or export from the United States or allow the export or re-export of any part of the Software or any direct product thereof (a) into (or to a national or resident of) any embargoed or terrorist-supporting country, (b) to anyone on the U.S. Commerce Department's Table of Denial Orders or U.S. Treasury Department's list of Specially Designated Nationals, (c) to any country to which such export or re-export is restricted or prohibited, or as to which the United States government or any agency thereof requires an export license or other governmental approval at the time of export or re-export without first obtaining such license or approval, or (d) otherwise in violation of any export or import restrictions, laws or regulations of any United States or foreign agency or authority. Licensee agrees to the foregoing and warrants that it is not located in, under the control of, or a national or resident of any such prohibited country or on any such prohibited party list. The Software is further restricted from being used for the design or development of nuclear, chemical, or biological weapons or missile technology, or for terrorist activity, without the prior permission of the United States government.

11.7 Third-Party Code. The Software may contain or be provided with components subject to the terms and conditions of third party "open source" software licenses ("Open Source Software"). Open Source Software may be identified in the Documentation, or Spectra shall provide a list of the Open Source Software for a particular version of the Software to Licensee upon Licensee's written request. To the extent required by the license that accompanies the Open Source Software, the terms of such license will apply in lieu of the terms of this MLA with respect to such Open Source Software.

11.8 Entire Agreement. This MLA is the complete and exclusive statement of the mutual understanding of the parties and supersedes and cancels all previous written and oral agreements and communications relating to the subject matter contained herein.

Amazon Web Services

If Licensee has licensed Software for use in conjunction with Amazon Web Services, such web services will be provided by Amazon in accordance with its standard terms and conditions. SPECTRA MAKES NO WARRANTY REGARDING AMAZON SERVICES AND SUGGESTS THE USE OF AMAZON'S CONTINUOUS DATA BACK UP SERVICES.

CONTACTING SPECTRA LOGIC

To Obtain General Information - Spectra Logic Website: www.spectrallogic.com	
United States Headquarters	European Office
Spectra Logic Corporation 6285 Lookout Road Boulder, CO 80301 USA	Spectra Logic Europe Ltd. 329 Doncastle Road Bracknell Berks, RG12 8PE United Kingdom
Phone: 1.800.833.1132 or 1.303.449.6400 International: 1.303.449.6400 Fax: 1.303.939.8844	Phone: 44 (0) 870.112.2150 Fax: 44 (0) 870.112.2175
Spectra Logic Technical Support Technical Support Portal: support.spectrallogic.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
Additional international numbers available at support.spectrallogic.com/home If you have a Spectra Logic Portal account, please log in for country-specific numbers at support.spectrallogic.com/support-contact-info	
Spectra Logic Sales Website: shop.spectrallogic.com	
United States and Canada Phone: 1.800.833.1132 or 1.303.449.6400 Fax: 1.303.939.8844 Email: sales@spectrallogic.com	Europe Phone: 44 (0) 870.112.2150 Fax: 44 (0) 870.112.2175 Email: eurosales@spectrallogic.com
To Obtain Documents: support.spectrallogic.com/documentation	

INTRODUCTION

This guide describes the installation, configuration, and use, of the Spectra Vail® software. The guide helps you to optimize the software for best performance and data security.

This guide is intended for data center administrators and operators who maintain and operate object storage systems. This guide assumes a familiarity with large-scale data storage architecture, with installing, configuring, and use of data storage software, as well as with various data storage mediums, including cloud, disk, and tape.

RELATED PUBLICATIONS

Vail Online Help

This user guide is also available in web form, and can be accessed by clicking the question mark (?) icon in the Vail management console, or by entering the below URL into a web browser.

<https://support.spectrallogic.com/vail/vailonlinehelp.htm>

Spectra Logic Vail Release Notes

The Spectra Vail Software Releases Notes on the [Support Portal website](#) provide the most up-to-date information about the Spectra Vail application, including information about the latest software releases and documentation updates.

Spectra Logic BlackPearl Systems

The following documents related to the BlackPearl® Nearline Gateway and BlackPearl NAS systems are available from the Documentation screen on the BlackPearl user interface, and on the [Support Portal website](#) at: support.spectrallogic.com.

The [Spectra BlackPearl Nearline Gateway User Guide](#) provides detailed information about configuring, using, and maintaining your BlackPearl Nearline Gateway system.

The [Spectra BlackPearl DS3 API Reference](#) provides information on understanding and using the Spectra DS3 API.

The [BlackPearl NAS User Guide](#) provides detailed information about configuring, using, and maintaining your BlackPearl NAS system.

Spectra Logic Tape Libraries

User Guides for Spectra Logic tape libraries are posted on the [Support Portal website](#).

WHAT'S NEW

The Spectra Vail application is updated to Vail 3.5.0, which brings with it the following new features:

Current Version Storage

You are now able to use non-versioned buckets for cloud storage. This storage is identified as "Current-Version Only" storage, and can only be used for the current version of an object. Since objects disappear when a new version is created, Current Version storage is not allowed to be the only storage location of an object.

Storage Groups

Multiple storage definitions can be grouped and used as a single destination. In a Storage Group, data is distributed across the members of the group. This can significantly improve performance, and allows for expanded options when using multiple endpoints.

Application Database Improvements

The PostgreSQL database used by the Vail application now stores all partition data in a separate table, and all local user and user groups use a separate partition for each account. This simplifies indexes, reduces transaction overhead and storage requirements, and improves performance.

Note: This behavior only applies to new installations of Vail 3.5, or later. Older Vail installations that upgrade to Vail 3.5 may manually convert to the new database behavior. Contact Spectra Logic Technical Support for assistance.

Use With Buckets

The linked-bucket behavior used in previous Vail versions changes in Vail 3.5. The link bucket behavior has now been separated into a "Use With Buckets" feature and an "Import Content" feature. The previous linked bucket functionality is now achieved by selecting a single bucket to use, and selecting "Import Content". The "Use With Buckets" option allows you to make storage exclusive to a particular bucket or buckets, so objects from other buckets can never use it. This prevents storing inappropriate content, and also allows for faster deletion of the storage.

One use case of the previous linked bucket behavior was as a way to change the path for objects created on a cloud bucket. By default, the path on cloud storage includes the bucket name. Since linked bucket storage can only be used with the linked bucket, the bucket name is not included in the cloud bucket path. Now you can uncheck "Import Content" and configure a single "Use With Buckets" entry. When cloud storage can only be used with a single bucket, the cloud bucket path doesn't include the Vail bucket name. The cloud bucket path is now exposed as a modifiable value using the API. It cannot be changed using the Vail management console.

This change also makes it possible to modify the linked bucket attributes of storage. You can modify the "Import Content" checkbox to add this to existing storage. You cannot make storage more restrictive, but you can select additional "Use With Buckets" or remove the list entirely to allow previously restricted storage to be used by all buckets.

Performance Improvements

Lifecycle processing performance is significantly improved. The Vail application now uses more simultaneous operations and adjusts quickly to a large backlog of pending work. On start up, the Vail application no longer waits for initialization of the pending lifecycle queues before completing other initialization. Lifecycle queue initialization completes as a background task, and lifecycle processing for each storage destination begins after that initialization completes.

Volume Storage Write Throttling

You can now configure a write rate limit for a volume storage. This limit is the number of MiB per second the volume storage is allowed to process once the storage usage caution threshold has been reached. If storage usage is below the caution threshold, writes are unlimited. The configured rate limit is enforced at the caution threshold, and the rate is then reduced linearly as the volume usage increases. The rate limit applies to lifecycle processing and all data writes to the given storage.

Shared Packs for Efficient Tape Storage

The default maximum shared pack size is now 4 GiB. This allows for less fragmented space on tape storage as a smaller shared pack size increases the likelihood that all objects in the pack are deleted, allowing the pack itself to be deleted. Additionally, you can now configure the maximum shared pack size on tape storage.

Maximum Part and Object size

Starting with Vail 3.5.0, the maximum upload part size supported by the Vail application is 100 GiB, with a maximum object size of 1 PiB (10,000 parts of 100 GiB each). Not all cloud providers allow for objects of this size. An error is generated if an attempt is made to add cloud storage to an object that is larger than what the storage supports.

Automatic Update of Default BlackPearl Data Policies

The Vail application is updated to simplify the creation of tape storage. When you register a BlackPearl endpoint, the Vail application now updates the default data policies on the BlackPearl system for compatibility. The default data policies are updated to use Object Naming and CRC-32C checksums. This update only occurs if the default data policies are not already in use. Other data policies are not affected.

Setting Cloud Storage Class

When creating cloud storage, the Storage Class assigned to created cloud objects is now explicitly selected instead of being determined by the storage class of the storage definition. Additional cloud storage tier definitions for Microsoft Azure and Google Cloud storage have been added to the Vail application.

New Storage Class Compatibility

The Vail application now accepts FSX_OPENZFS as a valid storage class.

S3 Compatible Link Notifications

Linked bucket notifications are now supported on S3 compatible storage in addition to AWS linked buckets. The service must be configured to send SNS notifications in the required format.

Account Connection Limit

A connection limit can now be configured for a user account to prevent single account from using all available connections.

Default Lifecycle

After upgrading to Vail 3.5, all buckets must have an assigned lifecycle. If a lifecycle is not currently assigned to a bucket, the bucket owner's default is used. If the bucket owner does not have a configured default lifecycle, the global default is used. When updating to Vail 3.5, a new default lifecycle is created with basic settings that delete pending multi-part uploads and removes useless delete markers.

You can set any lifecycle as the default and can delete the automatically created default.

Clone Operations By Prefix

You can now create and delete clones for all objects using a provided prefix. Additionally, you can create restore requests for all objects with a shared prefix. This functionality is only available through the management API.

S3 Checksum and Signature Support Improvements

The Vail application now offers additional support for AWS checksum handling and signature Version 4a.

CHAPTER 1 - UNDERSTANDING THE VAIL APPLICATION

This chapter describes the concepts behind the Vail application and how it works with various storage technologies and mediums.

Vail Application Features	22
Understanding BlackPearl Storage	24
Types of BlackPearl Storage	24
Understanding Vail Cloud Storage	25
Vail Application with AWS Cloud Storage	25
Vail Application with Microsoft Azure Cloud Storage Features	26
Vail Application with Google Cloud Storage	26
Understanding Vail Lifecycles	28
Considerations for Placement Rules	28
Considerations for Delete Rules	29
Understanding Object Storage	30
Understanding Vail Buckets	31
Considerations for Vail Bucket Creation	31
Understanding S3 with the Vail Application	32
Interfacing with Tape Libraries and Media	33
Tape Ingest	33
Spectra Logic Vail Application Naming Conventions	35
Vail Sphere Names	35
User Names	35
Group Names	35
Location Names	35
Portable Location Names	36
Storage Names	36
Spectra Vail Bucket Names	37
Cloud Bucket Names	37
BlackPearl Bucket Names	37
Lifecycle Names	38
Additional AWS Account Role Names	38

Vail Application Technical Information	39
Maximum Part and Object size	39

VAIL APPLICATION FEATURES

The Spectra Logic Vail application is an advanced data management system providing a comprehensive solution for managing, protecting, and accessing data across hybrid cloud and on-premises environments. Spectra Vail application integrates various data storage technologies to streamline data workflows and ensure efficient data management and protection.

The major Spectra Vail application features include:

- Global access to a single name space, which allows for seamless data management across various locations and storage types.
- Command, control, and monitoring using a single management point, which simplifies the management process.
- An Amazon[®] S3 compatible interface for easy integration with applications and services using the Spectra Vail API.
- Data stored on-premises using flash, disk, tape, or third party storage.
- Data stored to and synchronized from cloud storage, allowing for flexible data management options. The Vail application supports cloud storage providers such as Amazon Web Services[®], Microsoft[®] Azure[®], Google Cloud Storage[®], and others.
- Once an object is uploaded anywhere in the Spectra Vail sphere, it is immediately available to anyone who is connected to the sphere, promoting collaboration and efficient data access.
- Rich policy engine provides time-based data placement for each storage type and geographic location, allowing for efficient data management and compliance with data governance policies.
- Unlimited number of data storage sites, offering scalability to meet growing data storage and security needs. Migrating data to multiple geographic locations around the world allows for increased data protection in the event of natural disasters or virus attacks, which normally would impact where and how a company accesses its data.
- Local or cloud-based control. Local control uses a standalone BlackPearl Nearline Object Gateway system without the need for cloud services, and includes support for multi-factor authentication. Cloud control allows organizations to log in from anywhere in the world to control and monitor data movement, enhancing accessibility and operational flexibility.

Other features of the Vail application include:

- A high-speed cache on the BlackPearl Nearline Object Gateway which enhances performance for frequently accessed data.
- System monitoring and tools to view capacity information and performance metrics, ensuring optimal data management.
- Supports Cross-Origin Resource Sharing (CORS) for S3 commands, enhancing web integration capabilities.
- The Vail application can use a HotPair BlackPearl system running BlackPearl OS 5.7.x or later, enhancing data protection and availability.
- An improved bucket object display includes information about the total number of objects, object size, and total size of all objects in the bucket, all within the Spectra Vail management console.
- Supports S3 bucket notifications for various bucket events and S3 tagging of objects, allowing IAM policies to use tags instead of prefixes.
- Includes the BlackPearl Embedded Dashboard for controlling common features of the BlackPearl system without accessing its user interface.

General Use Cases include:

- **Backup and Archive:** Serves as a target for backup applications and long-term data archiving.
- **Data Replication:** Enables replication of data from other storage environments, increasing data availability and redundancy.
- **Disaster Recovery:** Supports robust backup and disaster recovery solutions. Data can be migrated to multiple geographic locations to protect against natural disasters or cyber-attacks.
- **Data Archiving:** Provides cost-effective long-term archiving solutions.
- **Multi-Tenancy:** Suitable for environments requiring isolated data management within a shared infrastructure.

UNDERSTANDING BLACKPEARL STORAGE

The Vail application uses a BlackPearl Nearline Object Gateway as endpoint storage, which provides scalable, efficient, and flexible storage across multiple mediums, including disk, tape, and cloud. Data transfer operations on a BlackPearl system use the standard S3 interface, appearing to applications as AWS S3 storage.

The BlackPearl Nearline Object Gateway uses a system cache to enhance performance for frequently accessed data and uses the Advanced Bucket Management policy engine to manage data movement between different storage targets based on defined rules. The BlackPearl Nearline Object Gateway also supports integration with cloud storage providers, which allows for intelligent object placement and retrieval.

BlackPearl endpoint storage can be used as a path to tape storage. The BlackPearl Nearline Object Gateway uses the Amazon S3 Glacier compatible interface to store data on tape, offering a seamless backup and restore process using standard S3 commands.

Types of BlackPearl Storage

The Vail application supports using both bucket storage and NAS storage provided by a BlackPearl Nearline Object Gateway, each allowing for different use cases and operational needs.

- BlackPearl bucket storage is designed for object storage, using the S3 protocol, which is ideal for managing unstructured data such as backups, archives, and large data sets. This storage model is highly scalable, allowing for the efficient handling of large volumes of data with support for multipart uploads, object versioning, and lifecycle policies to automate data management tasks. Users interact with BlackPearl bucket storage using APIs, which allows for seamless integration with various applications and workflows that require object storage.
- BlackPearl NAS share storage is designed for direct file storage, providing a more traditional approach to storing and accessing data. NAS share storage supports common file protocols like SMB and NFS, making it suitable for environments where users and applications need shared access to files and directories. NAS share storage is particularly useful for collaborative work, as multiple users can read and write to the same files simultaneously. NAS share storage is often used for storing documents, media files, and other structured data that benefit from hierarchical organization and direct file access.

UNDERSTANDING VAIL CLOUD STORAGE

The Vail application uses cloud storage by seamlessly integrating with various cloud services, allowing you to manage and move data efficiently. The Vail application uses predefined policies and lifecycle rules to automate data migration, replication, and tiering within the cloud. This allows data to be optimally stored, balancing cost and accessibility.

The Vail application provides a unified management interface that allows you to oversee and control your data across different cloud environments. Data is stored in Vail buckets, which can link to cloud storage, allowing for flexible and scalable data management. The Vail software automates the process of moving data to different cloud storage tiers based on access frequency and other criteria, which allows for efficient use of cloud resources and cost savings.

Additionally, the Vail application supports object storage within the cloud, making it easy to store and retrieve large volumes of unstructured data. The integration with cloud storage services also enhances data durability and availability, leveraging the cloud's inherent redundancy and disaster recovery capabilities.

The Vail application supports bi-directional synchronization with cloud storage, allowing data to be synchronized between local BlackPearl buckets and cloud buckets. Any data placed in one location is automatically synchronized to the other location, facilitating seamless data access and management across different storage environments. This can be useful for distributed workflows where data needs to be accessed and processed in multiple locations.

The Vail application provides robust security measures to protect data during transfer to cloud targets. All network traffic between Vail nodes and cloud resources is secured using HTTPS with TLS encryption. This ensures that data is protected from unauthorized access during transit.

Vail Application with AWS Cloud Storage

The information below describes the features provided by Amazon AWS cloud storage when used with the Vail application.

- Bucket synchronization allows for seamless synchronization between AWS S3 buckets and Vail buckets. Synchronization occurs bi-directionally. This facilitates distributed and accelerated data ingestion by synchronizing data placed in an AWS bucket to a local Vail bucket and vice versa and allows you to apply AWS services to local data without permanently storing it in the cloud.
- The Vail application uses the multipart upload capabilities of AWS S3, which allows large objects to be uploaded in parts, with each part tracked to resume interrupted transfers seamlessly.
- AWS S3 supports immutable objects, preventing deletion during a specified retention period. The Vail application uses object locking so that objects stored in AWS S3 cloud storage comply with specified retention policies.

- AWS IAM policies are used for secure access and permissions management. AWS credentials (access keys and IAM roles) are used to authenticate and authorize access to S3 buckets.
- AWS S3 versioning is required for AWS buckets used with Vail, so that multiple versions of objects are maintained.
- AWS cloud storage supports creating lifecycle rules to transition objects between storage classes, such as moving infrequently accessed data to GLACIER or DEEP_ARCHIVE.
- Data security features include encrypting data using AWS Key Management Service (KMS) for secure key management, and using HTTPS and V4 authentication ensure data is protected in transit.

Vail Application with Microsoft Azure Cloud Storage Features

The information below describes the features provided by Microsoft Azure cloud storage when used with the Vail application.

- Data stored in Azure cloud pools is treated as native Azure objects, maintaining compatibility and allowing seamless access through Azure storage services.
- Azure cloud storage uses Azure block blobs for multipart uploads, determining the method based on object size. Multipart uploads on Azure do not provide a unique identifier. Part IDs must include a unique value in addition to the part number to prevent simultaneous uploads from interfering with each other.
- Vail object locking works with Azure immutable objects, allowing explicit control and expiration settings. The Vail application does not use Azure immutable storage settings directly, it manages immutable objects using special clone deletion processing.
- The Vail application uses the storage container shared secret key for authentication. Credentials must include permissions for deleting blobs and blob versions using Azure RBAC.
- Azure storage does not handle slashes and backslashes as ordinary characters, converting backslashes to slashes when writing blobs. Leading slashes and repeated slashes are ignored or compressed to a single slash respectively. The Vail application adjusts its behavior to accommodate these Azure storage constraints.

Note: Azure storage restricts user metadata keys to C# naming conventions, and unsupported characters in user metadata are not copied to Azure storage.

Vail Application with Google Cloud Storage

- Vail maps AWS storage classes to Google Cloud Storage tiers, but note that GLACIER and DEEP_ARCHIVE storage classes are not supported with Google Cloud Storage.
- The Vail application uses role-based authentication to connect to Google cloud storage.
- Google cloud storage supports multipart uploads using Google SDK, which simplifies and optimizes large uploads, using resumable upload to track the progress.

- Vail object locking works with immutable objects, allowing explicit control and expiration settings. The Vail application recognizes and handles locked objects, but does not use Google immutable storage settings directly.
- Data stored by the Vail application in Google Cloud storage is stored in packs, with each pack assigned a unique identifier. These packs are then listed in the pack list for the corresponding object version in the Vail application. The version ID and ETag are used to validate the data integrity of an object before it is retrieved and reconstructed.

UNDERSTANDING VAIL LIFECYCLES

A Vail lifecycle is a set of automated policies that manage the lifecycle of data from creation to deletion. These policies dictate how data is transitioned through various stages, such as moving data between different storage tiers based on age, access frequency, or other criteria, and eventually deleting data that is no longer needed. The purpose of implementing lifecycle policies is to optimize storage costs, improve data management efficiency, and provide compliance with regulatory and organizational data retention requirements.

Lifecycles are controlled using placement and deletion rules. These rules specify the conditions under which data is moved or deleted, such as after a certain period or if it has not been accessed for a specified duration. Administrators can set a specific time for lifecycle actions to minimize disruption to users and system performance. Logging and auditing features keep a record of all actions taken by the lifecycle policies, providing an audit trail for compliance and troubleshooting purposes.

By automating data management through well-defined lifecycle policies, organizations can ensure efficient storage resource use, data retention policy compliance, and data loss risk reduction.

Considerations for Placement Rules

When creating a Vail lifecycle placement rule, consider the below information to provide efficient and effective data management.

- Define the specific criteria for transitioning data between storage tiers based on factors such as the age of the data, access frequency, and data size. These criteria help you determine when and how data is moved from a higher-cost, high-performance tier to a more cost-effective, lower-performance tier, or even to archival storage. It is crucial to understand the characteristics and costs associated with each storage tier to optimize storage expenses while maintaining the necessary performance levels for your data.
- Consider the impact of data transitions on access and retrieval times. Data moved to an archival tier might result in longer retrieval times and potentially higher costs when accessed. Lifecycle rules should be designed in accordance with your organization data access patterns and compliance requirements.
- Define clear retention periods for different types of data to comply with legal and regulatory requirements. This includes setting rules for deleting data after it has been retained for a specified duration.

Additionally, Spectra Logic recommends testing the lifecycle placement rules in a controlled environment before applying them broadly. This helps to identify any unintended consequences and confirms that the rules function as expected, preventing potential disruptions in data availability or performance.

Considerations for Delete Rules

The primary goal of a delete rule is to automate the removal of data that is no longer needed, thereby optimizing storage costs and maintaining a clutter-free environment. When creating a Vail lifecycle delete rule, consider the below information to provide effective data management and compliance.

- Define the criteria that determine when data should be deleted, such as the age of the data, last access time, or specific metadata attributes. These criteria should align with the data retention policies and regulatory requirements of your organization for compliance with legal mandates and avoid unintentional loss of important data.
- Consider the retention periods for different types of data. Some data may need to be retained for longer periods due to regulatory requirements, legal holds, or business needs. It is crucial to ensure that these retention requirements are incorporated into your lifecycle delete rules to prevent premature deletion. Verify that data scheduled for deletion is not part of critical backup sets or disaster recovery plans, as its removal could impact the ability to restore important information.
- Consider the impact on data access and performance. Deleting large volumes of data may affect system performance or disrupt ongoing operations. Spectra Logic recommends scheduling delete operations during off-peak hours or in a staggered manner to minimize any potential impact.

Additionally, Spectra Logic recommends testing the lifecycle deletion rules in a controlled environment before applying them broadly. This helps to identify any unintended consequences and confirms that the rules function as expected, preventing potential disruptions in data availability or performance.

UNDERSTANDING OBJECT STORAGE

Objects may have their content stored in any number of locations across the Vail sphere. This flexibility optimizes access and provides data redundancy and recovery.

The Vail application uses lifecycle rules to add and remove object storage locations. Lifecycles also can be used to delete the entire object. Lifecycle rules automate data placement and retention across storage locations in the Vail sphere.

A common scenario involves immediately adding duplicate storage locations when objects are created. For example, a lifecycle policy can be configured to add AWS cloud storage immediately after an object is placed into on-premise BlackPearl storage. This allows for data availability in both local and cloud environments for enhanced data access, protection, and disaster recovery.

UNDERSTANDING VAIL BUCKETS

Vail buckets provide a versatile and scalable object storage solution. A Vail bucket is a top level component of Vail application platform, designed to streamline and unify data management across multiple storage environments, including on-premises and cloud-based systems. This integration allows organizations to create a cohesive data management strategy, leveraging the strengths of different storage systems while maintaining a unified view of their data.

You can create a Vail bucket and assign it a lifecycle to manage the placement, movement, and retention. Through automated rules, data can be tiered across different storage types based on access patterns, age, and other criteria, optimizing both cost and performance.

Vail buckets support bi-directional synchronization, allowing data to be mirrored between local storage and cloud targets, allowing for high availability and redundancy. This functionality is particularly advantageous for organizations implementing a hybrid or multi-cloud approach, as it provides data access and management across various storage locations.

Additionally, Vail buckets offer advanced features such as versioning, metadata management, and robust security controls, so that data is not only stored efficiently but also protected and easily accessible.

Considerations for Vail Bucket Creation

When creating a Vail lifecycle bucket, consider the below information to optimize data organization and bucket functionality.

- Plan a method to organize your data in the Vail bucket. Consider using a hierarchical structure with clear naming conventions to facilitate easy data retrieval and management. Plan for the use of metadata to tag and categorize data, which improves searchability and organization.
- Decide on versioning requirements for your data. Versioning allows you to maintain multiple versions of objects, which provides data protection against accidental deletions or modifications. If you chose to use versioning, object locking can be used, which allows you to protect the state of an object when the lock is applied.
- Decide on a data security strategy that includes using encryption, access controls, and bucket and object ownership permissions.
- Plan for data redundancy and disaster recovery. The Vail application supports bi-directional synchronization of buckets, allowing data to be mirrored between local and cloud storage targets.
- Verify compatibility with your existing storage infrastructure by verifying that your chosen storage targets support the necessary APIs and protocols for seamless integration.

UNDERSTANDING S3 WITH THE VAIL APPLICATION

The Vail application mimics the AWS S3 interface so that any client that uses S3 can interface with the Vail sphere without additional software. When integrating S3 with the Vail application, it is necessary to understand multiple technical aspects for consistent operation and optimal performance.

Configure AWS access keys (Access Key ID and Secret Access Key) for use with the Vail application, which uses the keys to authenticate API requests to S3 storage. Keys are configured in the Vail management console and provide programmatic access to S3 resources. Keys must be configured with sufficient permissions to access S3 buckets and perform required data storage operations. As a best practice, rotate access keys regularly and avoid embedding them in code. Store them securely using your security management tools.

Utilize IAM roles for secure and controlled access to data. IAM roles provide temporary security credentials, which reduce the risk of long-term credential exposure. The Vail application assumes an IAM role using AWS Security Token Service (STS). Use IAM policies that grant the minimum permissions necessary for Vail to function, following the principle of least privilege to minimize security risks. While less commonly used because of their complexity, ACLs can be configured to grant Vail access to individual objects in a bucket.

The Vail application supports multipart uploads for large files. For large file transfers, the Vail application breaks data into chunks to optimize upload and download performance. Configure chunk sizes appropriate for your network and storage performance requirements. If transferring large amounts of data over long distances, consider using AWS S3 Transfer Acceleration to speed up data transfers.

The Vail application uses AWS Signature Version 4 for signing API requests. Vail communicates with S3 over HTTPS so that data in transit is encrypted and secure. Additional security methods include using server-side encryption with AWS KMS (Key Management Service) for data at rest.

INTERFACING WITH TAPE LIBRARIES AND MEDIA

Tape storage is a highly reliable and cost-effective medium for long-term data archiving and backup. It offers substantial storage capacity at a lower cost per gigabyte compared to disk or solid-state storage, making it ideal for organizations managing large volumes of data. One of the key advantages of tape is its longevity, with a shelf life that can exceed 30 years, allowing for data preservation over extended periods.

Additionally, tape storage is known for its low energy consumption, as it does not require power when not in use, unlike spinning disks that consume power continuously. This can result in significant cost savings in terms of both energy and cooling requirements. Tapes also offer robust security features, including offline storage capabilities that provide an inherent protection against threats such as ransomware, which often target online systems.

The Vail application streamlines the data path to tape media using a BlackPearl Nearline Object Gateway, which provides the interface between the Vail application and tape storage. The BlackPearl system stores data in a local cache before writing it to tape media. The cache ensures that data is managed efficiently and tape drives are utilized effectively.

When data is requested, a client issues an S3 Glacier command to move data from tape media (Glacier) to the BlackPearl cache. A second S3 command retrieves the data from the BlackPearl cache.

For S3 clients that do not support S3 Glacier commands, Vail offers a compatibility mode. This mode automatically issues an object restore command for objects in Glacier-tape class, which allows for seamless data retrieval without client-side modifications.

The information below provides additional information on how the Vail application uses tape media storage.

Tape Ingest

Vail supports the ingest and synchronization of existing BlackPearl buckets currently on tape media. This is useful if you have been using a BlackPearl system and want to integrate your existing data with the Vail application. The ingestion process involves adding a BlackPearl ingest agent, which performs a HEAD operation to index the objects into its database, and applies lifecycle rules as needed. If a lifecycle rule is not applied, objects remain on the BlackPearl system but are accessible as Glacier objects through the Vail application, requiring an Object Restore command for retrieval.

Vail allows for universal tape ingest within a BlackPearl system managed by the same Vail Sphere. This means tapes exported from one BlackPearl system can be imported into another without database or object migration. Vail handles the metadata ingest and makes objects available across the sphere. The metadata on imported tapes is read and ingested into the new BlackPearl system database.

Glacier Tape Packing

The Vail application attempts to generate 64 GB packs of data before transferring it to tape media, which allows for increased performance both for write and read operations. When a file is requested from a pack, the entire 64 GB pack is restored to the BlackPearl cache.

When deleting data files, a pack is not deleted unless every object in a pack is deleted. Because of this, if only some files in a pack are deleted in the Vail application, the deleted files will no longer be available but will still count against the available storage space on the tape cartridge. When the remaining files in the pack are deleted, the entire pack is deleted and the previously used storage space is made available.

SPECTRA LOGIC VAIL APPLICATION NAMING CONVENTIONS

Before configuring the Spectra Vail sphere, Spectra Logic recommends establishing a naming convention for your data storage infrastructure. A well-considered naming convention allows for easier setup and configuration, as well as a roadmap for naming Spectra Vail components added at a later date.

Use the information below to develop a naming convention for the Spectra Vail sphere and resources before you install and configure the Spectra Vail application.

The Spectra Vail application uses the same naming restrictions as Amazon Web Services. For more information on allowed naming conventions, see [AWS User Documentation](#).

Vail Sphere Names

When using multiple Spectra Vail spheres, each sphere must have a unique name.

User Names

User names identify Spectra Vail sphere administrator users, as well as IAM users associated with the Vail sphere administrator's AWS account. Spectra Logic suggests using the same naming convention as your corporate email for user names.

For example, if associate Jane Smith uses the email address `janes@yourcompany.com`, use "janes" for the user name.

Group Names

Groups of users are typically configured when all users of the group share the same access policies. Spectra Logic suggests using self-explanatory names for groups.

For example, if your company's groups are assigned by department, use a naming convention that directly identifies each department such as Production, Engineering, or Accounting.

Location Names

Locations designate a physical location that contains Spectra Vail resources. Spectra Logic recommends naming each location after its physical location in the world.

For example, if Spectra Vail resources are located in Dallas, use that as the location name if there is only one Spectra Vail resource in that city. If there are multiple Spectra Vail resources consolidated in the same city, use suffixes to identify each group such as Dallas-HQ, Dallas-Research, or Dallas-Production.

Portable Location Names

Vail VM Node storage can be installed on a portable storage device such as a laptop or hardened disk array. Portable location names should be used for any device that only resides in a geographic location temporarily.

Storage Names

Storage names identify a storage target in the Spectra Vail sphere. Storage names are used for Spectra Vail VM node storage, VM pool storage, cloud storage, BlackPearl storage, and buckets or NAS shares configured on a BlackPearl system. Use the sections below to assist you in creating a storage name convention.

Vail VM Node Storage

The Spectra Vail VM node storage name is used as the top level name of the storage endpoint displayed in the Vail management console. Spectra Logic recommends using a name that includes both the location and type of storage.

For example, in the Dallas location, add the storage type as a suffix, such as Dallas-VM1 and Dallas-VM2.

VM Pool Storage

A Spectra Vail VM node uses one or more VM storage pools as data storage. Spectra Logic recommends using a name that includes the location, intended pool usage, and storage class.

For example, in the Dallas location, add suffixes for use and class such as Dallas_News_Standard and Dallas_Backup_Glacier.

BlackPearl Storage

BlackPearl storage includes disk pools and volumes configured on the BlackPearl system. Spectra Logic recommends using names that include the location, BlackPearl name, physical medium, and storage class.

For example, in the Dallas location, add suffixes for the BlackPearl name, physical medium, and storage class such as Dallas.BlackPearl1-Object-Standard and Dallas.BlackPearl2-Tape-Glacier.

Cloud Storage

Cloud storage includes storage locations provided by Amazon and other third-party vendors. Spectra Logic recommends using names that include the vendor, location, storage class, and intended usage.

For example, `AWS_USEast1_Standard_MarketingArchive`.

Spectra Vail Bucket Names

Spectra Vail buckets are the highest level of object grouping in the Spectra Vail sphere. Spectra Vail buckets are used with lifecycle rules, and buckets can include permissions by user, group, or role.

Spectra Logic recommends using names that either include the intended usage or user group name combined with intended usage. If you use a naming convention by groups, the associated group can be easily given access to all buckets sharing the group name prefix.

For example, use usage names such as `news-breaking` and `external-archive`, or group and usage names such as `eng-dev` and `eng-test`.

Note: Spectra Logic recommends following Amazon AWS documentation naming conventions and restrictions to maintain consistency and avoid potential issues. See [AWS Bucket Naming Rules](#) for more information.

Cloud Bucket Names

Cloud bucket names identify high-level containers in cloud storage and are not visible to end users. However, cloud bucket names are displayed in the configuration wizard. Spectra Logic recommends using names that include the type of cloud storage, location, and storage class.

Note: Cloud bucket names are restricted to lowercase characters, and do not allow underscores.

For example, use names for AWS cloud storage such as `vail-aws-uswest2-autotier` and `vail-aws-uswest2-S3glacier`.

Note: Do not create AWS cloud storage buckets with the prefix "spectra-logic-vail-". Buckets with that prefix do not display in the storage creation wizard and cannot be configured for use.

BlackPearl Bucket Names

BlackPearl bucket names identify high-level containers configured on BlackPearl systems and are not visible to end users. However, BlackPearl bucket names are displayed in the configuration wizard. Spectra Logic recommends using names that include the storage policy used by the BlackPearl bucket.

For example, `vail-singlecopytape` and `vail-dualcopytape`.

Note: Spectra Logic recommends following Amazon AWS documentation naming conventions and restrictions to maintain consistency and avoid potential issues. See [AWS Bucket Naming Rules](#) for more information.

Lifecycle Names

Lifecycles are policies that control where data is cloned, moved, or expired over time. Spectra Logic recommends using names that directly indicate the specific lifecycle rule configuration.

For example, use names such as `Copy_Everywhere_Keep4Days` and `Moveto_DallasNodeVM_After10Days`.

Lifecycle Rule Names

Lifecycle rules are used within a Lifecycle to specify the parameters for clone, move, or expiration rules.

Additional AWS Account Role Names

By default, the Spectra Vail sphere is configured with a master Administrator AWS account. Additional AWS accounts can be configured in the Spectra Vail sphere. Spectra Logic recommends using a name that indicates the intended role for the additional AWS account.

For example, use a name such as `VailSphereIAMreadandUserS3Control`.

VAIL APPLICATION TECHNICAL INFORMATION

This section provides information about various technical aspects of the Vail application.

Maximum Part and Object size

Starting with Vail 3.5.0, the maximum upload part size supported by the Vail application is 100 GiB, with a maximum object size of 1 PiB (10,000 parts of 100 GiB each). Not all cloud providers allow for objects of this size. An error is generated if an attempt is made to add cloud storage to an object that is larger than what the storage supports.

CHAPTER 2 - IMPORTANT INFORMATION

This chapter provides important information to know before using the Spectra Vail application.

Requirements	41
Spectra Logic Products Requirements	41
Supported Browsers	41
Vail Management Console Overview - Local Control	42
Main window	42
Taskbar	43
Toolbar	44
Embedded Dashboard	44
Icons	45
Vail Management Console Overview - Cloud Control	46
Main window	46
Taskbar	47
Toolbar	48
Icons	49

REQUIREMENTS

The following sections describe the requirements for using the Spectra Vail application.

Spectra Logic Products Requirements

- The Spectra Vail application version 3.1.x or later requires either a BlackPearl Nearline Gateway or a Vail VM node.
 - The BlackPearl Nearline Gateway requires BlackPearl OS 5.7.4 or later with a valid Vail Sphere activation key installed.
 - A Vail VM node requires a host machine with at a minimum 8 CPU cores, 16 GB of system memory, and a 10 GigE network connection.
- An S3 compatible client is required to access data stored in the Vail sphere.

Note: S3 clients communicating with the Vail sphere must use AWS v4 authentication.

- The Spectra Vail application uses the following ports for communication with BlackPearl systems and Vail VM Nodes. These ports must be open in your network infrastructure for the Spectra Vail application to function correctly.
 - **Inbound 80 and/or 443**
Inbound access is needed for these ports to access the BlackPearl user interface, and for S3 clients to transfer data to the gateway, using either the open (80) or secure port (443).
 - **Outbound 443**
Outbound access is needed for port 443 to allow data transfer to the Vail sphere, or other S3 endpoint nodes.

Supported Browsers

The Vail management console supports the Google® Chrome™ browser running on a Microsoft® Windows® or MacOS® system.

The browser versions listed below are supported.

Google Chrome:

- **Windows:** 88.0.4324.104 (Official Build) (x86_64), or later
- **MacOS:** 88.0.4324.96 (Official Build) (x86_64), or later

VAIL MANAGEMENT CONSOLE OVERVIEW - LOCAL CONTROL

The Vail management console provides browser-based configuration, management, and monitoring of the Vail sphere. The following sections describe the common features that appear in all screens in the management console when using Vail in a local control environment.




Figure 1 The Vail Sphere Dashboard screen.

Main window

The main window of the Vail management console displays the screen associated with each navigation link.

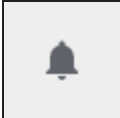


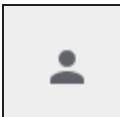
Taskbar

The taskbar (1) displays on the left side of all screens in the Vail management console. The following table provides a description of the selections in the taskbar.

Selection	Description
Dashboard	<p>The Dashboard navigation link takes you to the Dashboard screen which provides an overview of the Vail sphere status, the total managed data in the sphere, number and type of endpoints, and location details about each location and associated storage. Locations also display on the world map.</p> <p>The information displayed in the Location Details pane displays varies based on which type(s) of storage are configured for each location. For BlackPearl storage, the embedded dashboard displays allowing you to easily use the most common functions of a BlackPearl Nearline Object Gateway.</p>
Storage	The Storage navigation link takes you to the Storage screen which displays configured endpoint and cloud storage, and provides access to the wizard for configuring new storage, as well as editing and deleting storage. See Create Storage on page 63 .
Lifecycles	The Lifecycles navigation link takes you to the Lifecycles screen which displays configured lifecycles and provides access to the wizard to create new lifecycles, as well as editing and deleting lifecycles. See Create a Lifecycle on page 91
Buckets	The Buckets navigation link takes you to the Buckets screen which displays configured Vail buckets and provides access to the wizard to create new buckets, as well as editing and deleting buckets. See Create a Vail Bucket on page 100 .
Capacity	The Capacity navigation link takes you to the Capacity screen which displays endpoint and cloud storage capacity information. See View Capacity Information on page 151 .
Performance	The Performance navigation link takes you to the Performance screen which displays throughput and operation performance for both storage endpoints and the Vail sphere. See View Performance Metrics on page 154
Reports	The Reports navigation link takes you to the Reports screen which displays audit logs generated by the Vail sphere. Audit logs can be sorted by username or date. See View Reports on page 234 .

Toolbar

The toolbar (2) displays in the upper-right of the Vail management console. The following table provides an overview of the selections in the toolbar.








Icon	Meaning	Description
	Messages	<p>Displays the number of unread messages generated by the Vail sphere. The messages icon changes color depending on the highest severity message.</p> <p>The messages are categorized as:</p> <ul style="list-style-type: none"> • Info - An expected event occurred such as the completion of a software update. (Blue icon color). • Warning - An event that may impact the operation of the Vail sphere occurred. Determine the cause of the problem and remedy the issue if necessary. (Yellow icon color). • Error - An event which impacts data storage operations occurred. This may happen if the Vail sphere cannot communicate with storage endpoint. (Red icon color).
	Settings	<p>The settings menu allows you to:</p> <ul style="list-style-type: none"> • Configure IAM accounts, users, and groups • Update the Spectra Vail application • Access Logs • Configure Global Settings • Configure Network Settings • Configure Entitlements (license keys)
	Online Help	Opens a web browser to the Vail online help guide, a browser-based version of the Spectra Vail User Guide.
	User	Displays the user currently logged-in and provides access to the logout function.

Embedded Dashboard

The BlackPearl embedded dashboard (3) displays in the bottom of the Vail management console. See [BlackPearl Embedded Dashboard on page 251](#) for detailed information and instructions on using the features of the embedded dashboard.

Icons

The table below describes the icons that display on various screens in the Vail management console.

Icon	Meaning	Description
	View Details	Displays a detail screen for various aspects of the Vail sphere.
	Unknown Status - Group	Displays when the status of the subcomponents of a group are unknown.
	Good Status - Single	Indicates a good, working single component of the Vail sphere.
	Good Status - Group	Indicates a good, working group of components of the Vail sphere. This displays when all subcomponents of the group display good status.
	Warning Status	Indicates a problem with a component of the Vail sphere.
	Error Status - Single	Indicates an error of a component of the Vail sphere.
	Error Status - Group	Indicates an error with one or more subcomponents of the group.

VAIL MANAGEMENT CONSOLE OVERVIEW - CLOUD CONTROL

The Vail management console provides browser-based configuration, management, and monitoring of the Vail sphere. The following sections describe the common features that appear in all screens in the management console when using Vail in a cloud control environment.

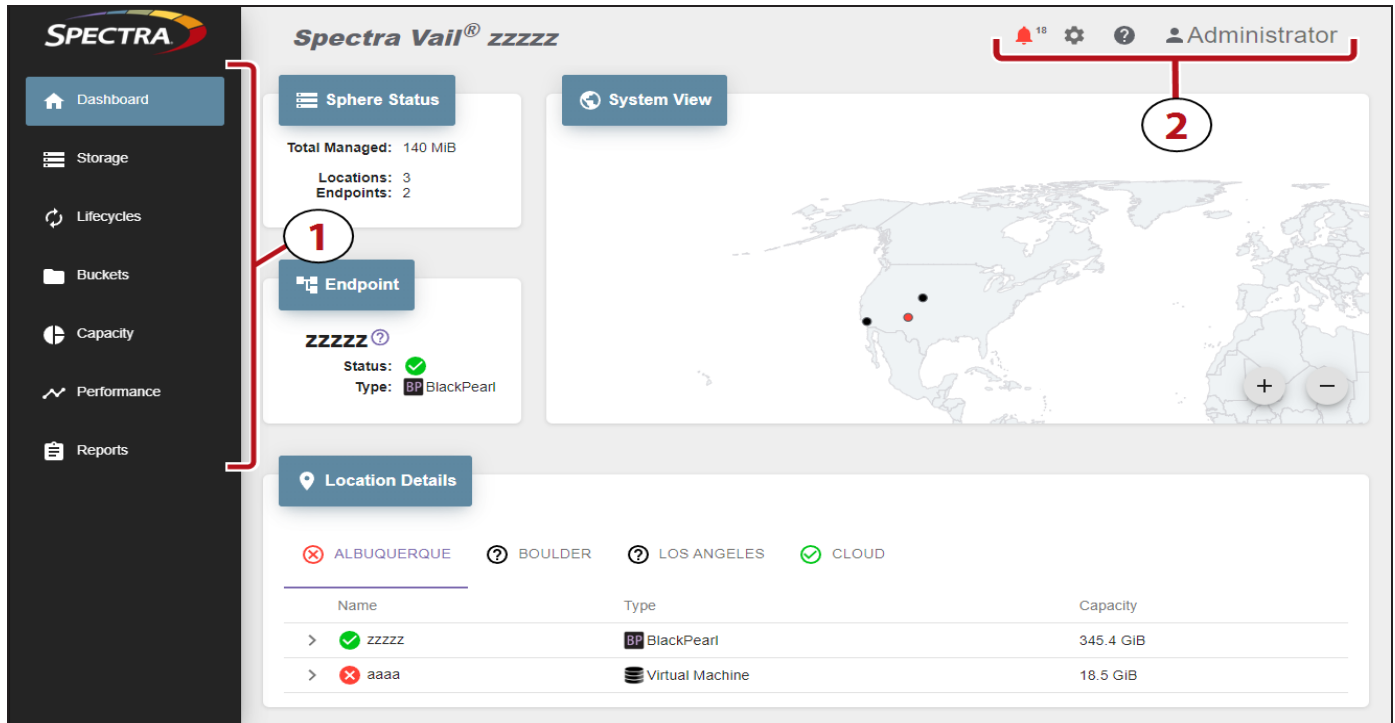


Figure 2 The Vail Sphere Dashboard screen.

Main window

The main window of the Vail management console displays the screen associated with each navigation link.

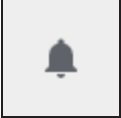

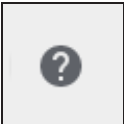
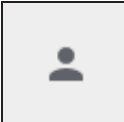
Taskbar

The taskbar (1) displays on the left side of all screens in the Vail management console. The following table provides a description of the selections in the taskbar.

Selection	Description
Dashboard	<p>The Dashboard navigation link takes you to the Dashboard screen which provides an overview of the Vail sphere status, the total managed data in the sphere, number and type of endpoints, and location details about each location and associated storage. Locations also display on the world map.</p> <p>The information displayed in the Location Details pane displays varies based on which type(s) of storage are configured for each location. For BlackPearl storage, the embedded dashboard displays allowing you to easily use the most common functions of a BlackPearl Nearline Object Gateway.</p>
Storage	The Storage navigation link takes you to the Storage screen which displays configured endpoint and cloud storage, and provides access to the wizard for configuring new storage, as well as editing and deleting storage. See Create Storage on page 63 .
Lifecycles	The Lifecycles navigation link takes you to the Lifecycles screen which displays configured lifecycles and provides access to the wizard to create new lifecycles, as well as editing and deleting lifecycles. See Create a Lifecycle on page 91
Buckets	The Buckets navigation link takes you to the Buckets screen which displays configured Vail buckets and provides access to the wizard to create new buckets, as well as editing and deleting buckets. See Create a Vail Bucket on page 100 .
Capacity	The Capacity navigation link takes you to the Capacity screen which displays endpoint and cloud storage capacity information. See View Capacity Information on page 151 .
Performance	The Performance navigation link takes you to the Performance screen which displays throughput and operation performance for both storage endpoints and the Vail sphere. See View Performance Metrics on page 154
Reports	The Reports navigation link takes you to the Reports screen which displays audit logs generated by the Vail sphere. Audit logs can be sorted by username or date. See View Reports on page 234 .








Toolbar

The toolbar (2) displays in the upper-right of the Vail management console. The following table provides an overview of the selections in the toolbar.

Icon	Meaning	Description
	Messages	<p>Displays the number of unread messages generated by the Vail sphere. The messages icon changes color depending on the highest severity message.</p> <p>The messages are categorized as:</p> <ul style="list-style-type: none"> • Info - An expected event occurred such as the completion of a software update. (Blue icon color). • Warning - An event that may impact the operation of the Vail sphere occurred. Determine the cause of the problem and remedy the issue if necessary. (Yellow icon color). • Error - An event which impacts data storage operations occurred. This may happen if the Vail sphere cannot communicate with storage endpoint. (Red icon color).
	Settings	<p>The settings menu allows you to:</p> <ul style="list-style-type: none"> • Configure Administrators • Configure IAM accounts, users, and groups • Configure Locations • Update the Spectra Vail application • Access Logs • Configure Global Settings • Configure Network Settings • Configure Entitlements (license keys)
	Online Help	Opens a web browser to the Vail online help guide, a browser-based version of the Spectra Vail User Guide.
	User	Displays the user currently logged-in and provides access to the logout function.

Icons

The table below describes the icons that display on various screens in the Vail management console.

Icon	Meaning	Description
	View Details	Displays a detail screen for various aspects of the Vail sphere.
	Unknown Status - Group	Displays when the status of the subcomponents of a group are unknown.
	Good Status - Single	Indicates a good, working single component of the Vail sphere.
	Good Status - Group	Indicates a good, working group of components of the Vail sphere. This displays when all subcomponents of the group display good status.
	Warning Status	Indicates a problem with a component of the Vail sphere.
	Error Status - Single	Indicates an error of a component of the Vail sphere.
	Error Status - Group	Indicates an error with one or more subcomponents of the group.

CHAPTER 3 - CONFIGURE BLACKPEARL SYSTEM

This chapter provides instructions for configuring a BlackPearl Nearline Object Gateway for use with the Spectra Vail application.

Configure a BlackPearl Nearline Gateway	51
Register a BlackPearl Nearline Object Gateway with a Vail Sphere	52
Register A Vail Sphere - Local Control	52
Register A Vail Sphere - Cloud Control	56

CONFIGURE A BLACKPEARL NEARLINE GATEWAY

If you are using the Spectra Vail application in conjunction with a BlackPearl Nearline Object Gateway, before you can access the Spectra Vail application, you must first install and configure the BlackPearl Nearline Object Gateway.

Use the [*BlackPearl Nearline Gateway User Guide*](#) to configure your BlackPearl Nearline Object Gateway.

Your BlackPearl Nearline Object Gateway may have been installed and configured by Spectra Logic Professional Services.

Note: If you need assistance configuring your BlackPearl Nearline Object Gateway, contact Spectra Logic. See [Contacting Spectra Logic](#) on page 16.

REGISTER A BLACKPEARL NEARLINE OBJECT GATEWAY WITH A VAIL SPHERE

Use one of the sections below to register a BlackPearl Nearline Object Gateway with the Spectra Vail application:

- **Register A Vail Sphere - Local Control below**
- **Register A Vail Sphere - Cloud Control on page 56**

Register A Vail Sphere - Local Control

Here is how to register a BlackPearl Nearline Object Gateway with a Vail sphere using local control:

Note: For instructions on registering a Vail VM node to a Vail sphere, see [Register a Vail VM Node with a Vail Sphere on page 289](#).

1. Log in to the BlackPearl user interface.
2. If necessary, configure the IP addressing for the BlackPearl Nearline Object Gateway. The Spectra Vail application node running on a BlackPearl system uses the IP address configured for the BlackPearl data port.
 - a. Select **Configuration > Network**.
 - b. Under **Network Interfaces**, select the row of the data connection.
 - c. Select **Action > Edit**. Configure the network settings as needed and click **Save**.

**IMPORTANT**

Spectra Logic recommends setting a static IP address.

3. If desired, change the system name of the BlackPearl Nearline Object Gateway. The Spectra Vail application uses this name for the Vail node name.
 - a. Select **Status > Hardware**.

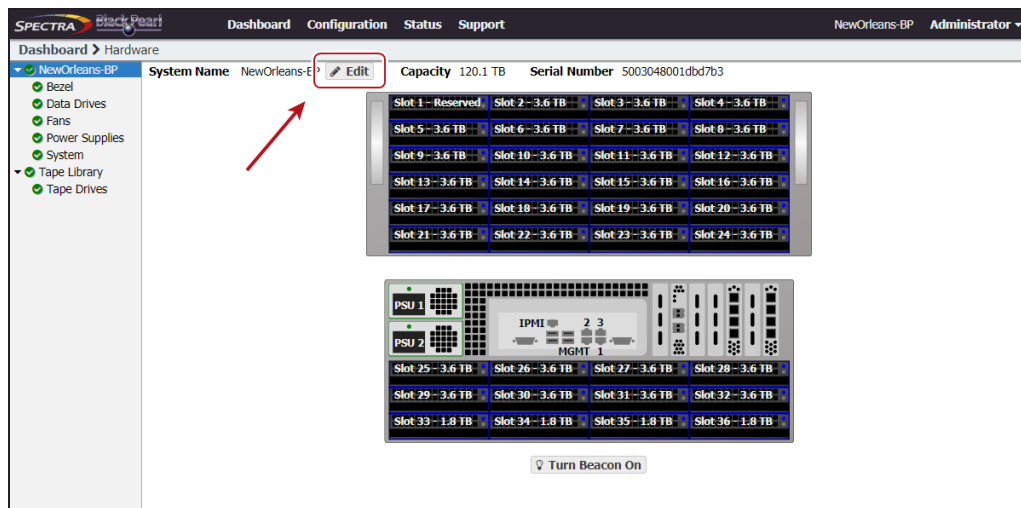


Figure 3 The BlackPearl user interface Hardware screen.

- b. Click **Edit**, enter the desired **Name**, and click **Save**.

Note: Spectra Logic recommends using the physical location of the BlackPearl system as the system name, for example Dallas.BlackPearl1-Object-Standard. The BlackPearl system name is limited to 15 characters before the first delimiter.

4. If necessary, add the Vail service key provided by Spectra Logic:
 - a. In the BlackPearl user interface, select **Support > Activation Keys**.
 - b. Select **Action > New**.
 - c. Enter the **Activation Key** and click **Save**.
5. In the BlackPearl user interface, select **Configuration > Services**.
6. Select the Vail service, then select **Action > Show Details**.

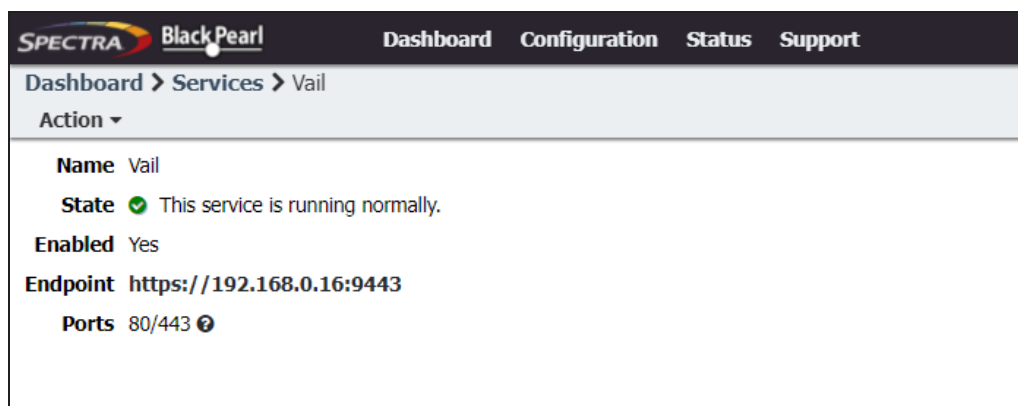


Figure 4 The Vail service details screen.

7. If desired, select **Action > Edit** to change the pair of ports used by the Spectra Vail application for HTTP and HTTPS connections. The ports automatically selected depend on if you have buckets created on your BlackPearl system.

- If buckets are configured on your BlackPearl system, the pair of ports selected is 80/443.
- If no buckets are configured, the pair of ports selected is 8080/8443.

Note: Whichever pair of ports is used by the Spectra Vail application, the other pair is used by the BlackPearl DS3 service. If you change the pair of ports for the Spectra Vail application, the DS3 service ports change to use the opposite pair of ports.

8. Click the **Endpoint** link in the Vail service details screen. A new web browser launches. The default web certificate is invalid, use your browser to bypass the certificate screen.

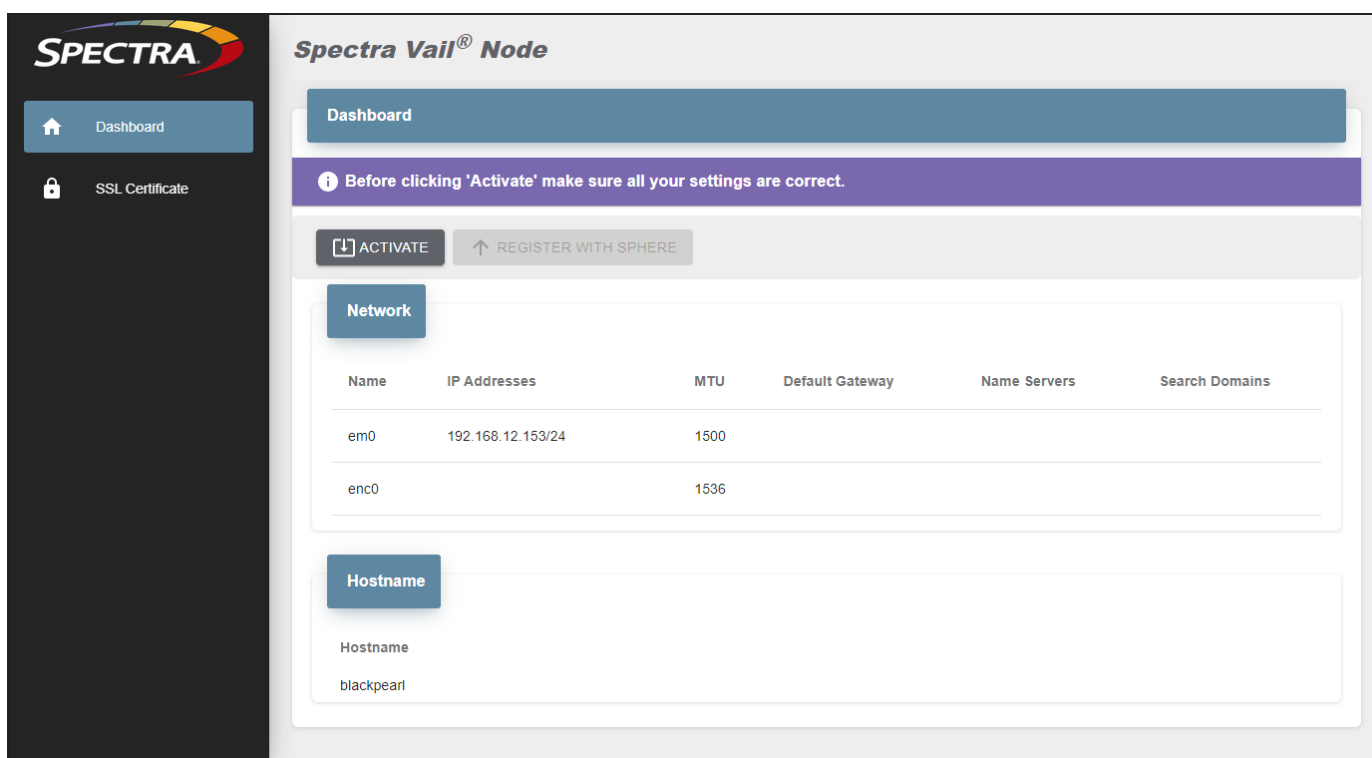


Figure 5 The Vail Node Dashboard - Activate and Register view.

9. If desired, update the SSL certificate before registering with the sphere:

- In the taskbar of the Vail VM node management console, click **SSL Certificate**.

- b. Under the **SSL Certificate** banner, click **Edit**.

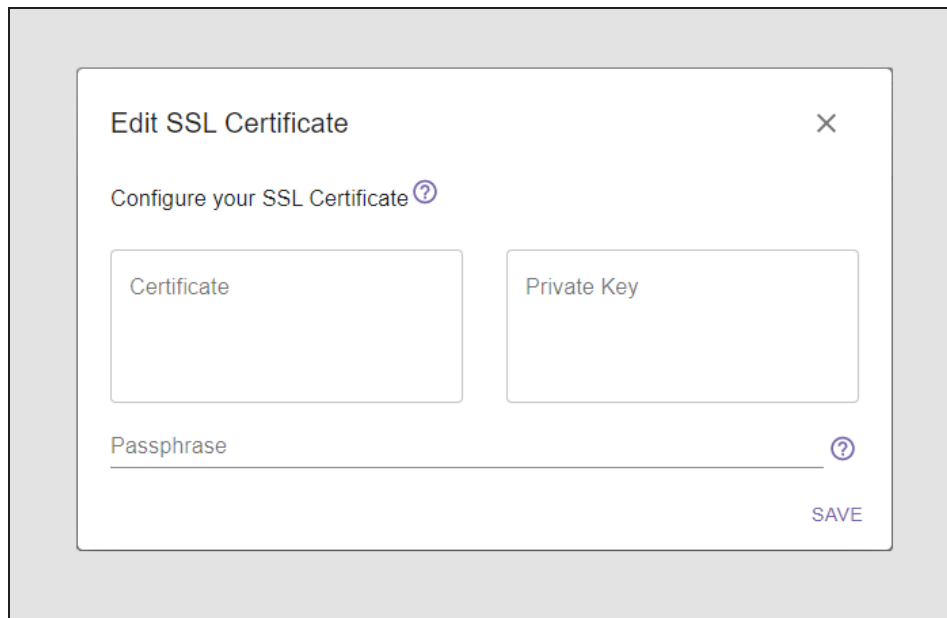


Figure 6 The Edit SSL Certificate screen.

- c. Enter the desired **Certificate** and **Private Key** in PEM format.
- d. If necessary, enter the **Passphrase** used to encrypt the private key.
- e. Click **Save**.

10. On the Vail dashboard screen, click **Activate**.

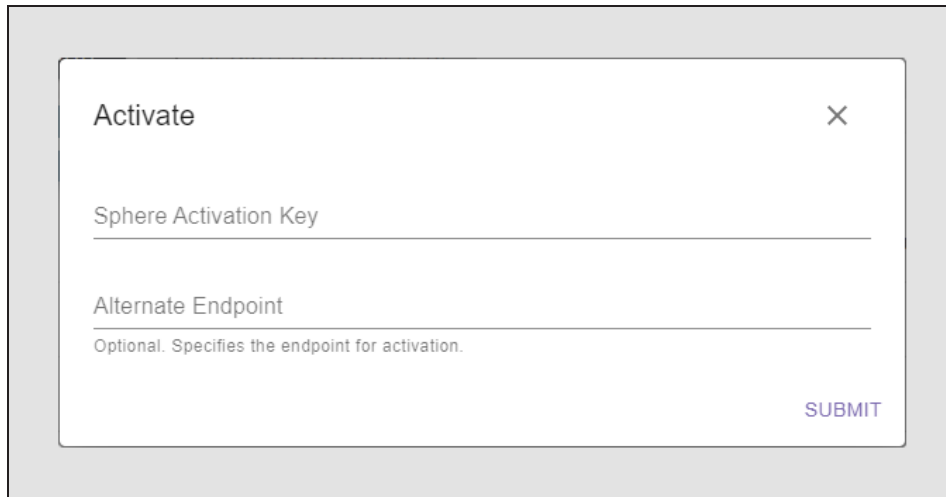


Figure 7 The Activate screen.

11. Enter the **Sphere Activation Key** and **Alternate Endpoint** provided by Spectra Logic.
12. Click **Submit**. Wait approximately 15 seconds while the Vail management console refreshes.

Register A Vail Sphere - Cloud Control

Here is how to register a BlackPearl Nearline Object Gateway with a Vail sphere using cloud control:

1. On the Vail dashboard screen, click **Register With Sphere**.

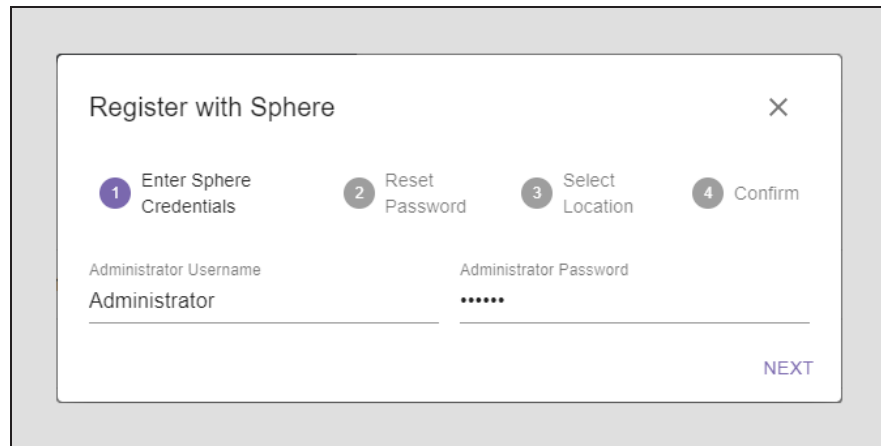


Figure 8 The Register with Sphere - Credentials screen.

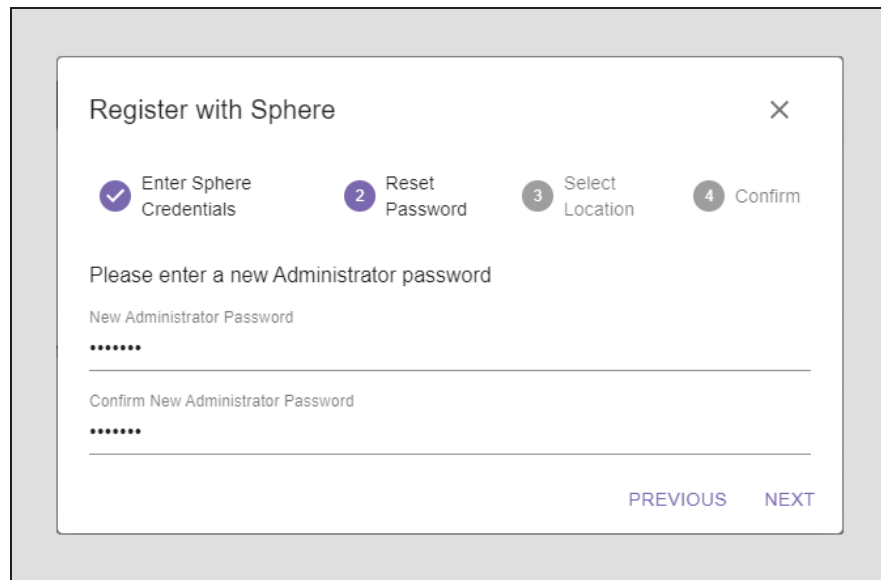
2. Enter the **Administrator Username** and **Administrator Password**.

- If this is the first BlackPearl system to register with a sphere, enter the credentials sent to the email address you provided to Spectra Logic when the sphere was created in AWS.

Note: You may need to set an email/MX rule to allow emails from AWS to the address entered when the sphere was created.

- Otherwise enter the credentials provided by your system administrator.

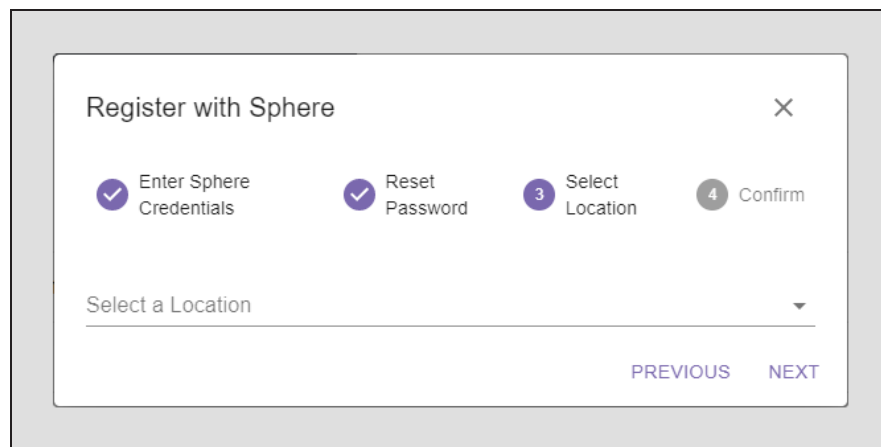
3. Click **Next**. If this is the first BlackPearl system to register with a sphere, you are prompted to set a new password. Otherwise, continue with [Step 5](#).



The screenshot shows a modal window titled "Register with Sphere" with a close button (X) in the top right corner. Below the title is a progress bar with four steps: 1. Enter Sphere Credentials (checked), 2. Reset Password (active), 3. Select Location, and 4. Confirm. The main content area prompts the user to "Please enter a new Administrator password" and provides two input fields: "New Administrator Password" and "Confirm New Administrator Password", both masked with dots. At the bottom right, there are "PREVIOUS" and "NEXT" buttons.

Figure 9 The Register with Sphere - Reset Password screen.

4. Enter a **New Administrator Password**, confirm the password, and click **Next**.



The screenshot shows the same "Register with Sphere" modal window. The progress bar now shows steps 1 and 2 as completed (checked), step 3 "Select Location" as active, and step 4 "Confirm" as next. The main content area prompts the user to "Select a Location" with a dropdown menu. At the bottom right, there are "PREVIOUS" and "NEXT" buttons.

Figure 10 The Register with Sphere - Select Location screen.

5. On the Select Location screen, chose to create a new location, or to use an existing location:
 - **Create a New Location below**
 - **Select an Existing Location on page 60**

Create a New Location

Here is how to create a new location:

1. To create a new location, use the drop-down to select **New Location**.
2. To map a location, you can search for the location, manually enter the latitude and longitude, or create a location with no corresponding geographic location.

Register with Sphere

Enter Sphere Credentials Reset Password **3 Select Location** 4 Confirm

Select a Location

New Location

Search and choose an address to use for your new location.
Note: You may skip this step if you wish to enter your location data manually.

Address Search

Please confirm the details below. If necessary, you may edit any pre-populated fields or execute another search.
Note: Latitude and Longitude values are used for the System View map on the dashboard.

Name

Latitude Longitude

PREVIOUS NEXT

Figure 11 The Register with Sphere - New Location screen.

- To search for a location...

- a. In the **Address Search** field, enter a geographic location. You can enter a full or partial postal address, city, county, province, or country.
- b. Select the correct match from the list.

Note: If no match is located, try changing the format of the address you entered. For example, use 9th Street in place of Ninth St.

- c. If desired, manually edit the **Name**.

Spectra Logic recommends naming each location after its physical location in the world.

For example, if Vail resources are located in Dallas, use that as the location name if there is only one Vail resource in that city. If there are multiple Vail resources consolidated in the same city, use suffixes to identify each group such as Dallas-HQ, Dallas-Research, or Dallas-Production.

- d. Confirm the information is correct and click **Next**.

- To manually enter a location...

- a. Enter the desired **Name**.

Spectra Logic recommends naming each location after its physical location in the world.

For example, if Vail resources are located in Dallas, use that as the location name if there is only one Vail resource in that city. If there are multiple Vail resources consolidated in the same city, use suffixes to identify each group such as Dallas-HQ, Dallas-Research, or Dallas-Production.

- b. Enter the **Latitude** and **Longitude** of the location.

- Notes:**
- When entering a value for **Latitude**, use positive values for locations north of the equator, and negative values for locations south of the equator.
 - When entering a value for **Longitude**, use positive values for locations east of the prime meridian, and negative values for locations west of the prime meridian.

- c. Click **Next**.

- To skip entering a location...

- a. Enter the desired **Name**.

Spectra Logic recommends naming each location after its physical location in the world.

For example, if Vail resources are located in Dallas, use that as the location name if there is only one Vail resource in that city. If there are multiple Vail resources consolidated in the same city, use suffixes to identify each group such as Dallas-HQ, Dallas-Research, or Dallas-Production.

- b. Click **Next**.

Note: If you do not enter an address or latitude and longitude, the location displays on the right-hand pane of the Vail dashboard, but does not display on the world map.

3. Confirm the information is correct, and click **Register**.

Wait while the BlackPearl system registers with the Vail sphere. This may take several minutes, during which time the Vail management console may display communication errors.

Select an Existing Location

Here is how to select an existing location:

1. Using the drop-down menu, **Select a Location** where you want to associate the BlackPearl Vail node and click **Next**.

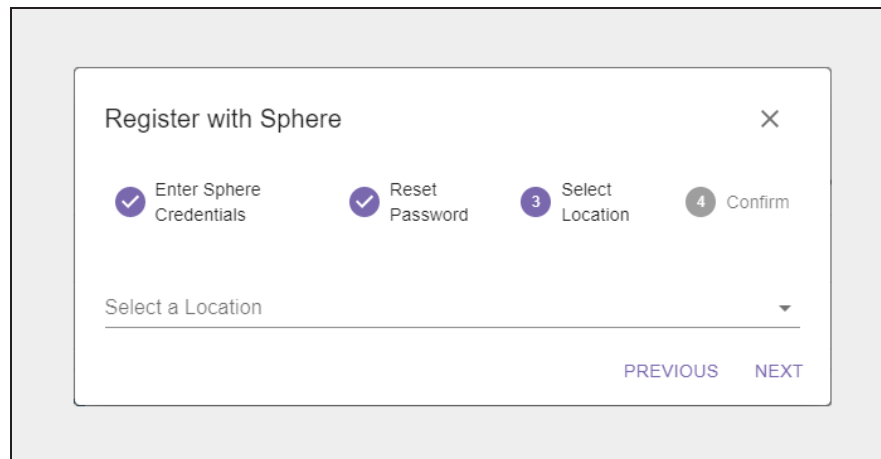
The image shows a screenshot of a web-based dialog box titled "Register with Sphere" with a close button (X) in the top right corner. The dialog box contains a progress indicator with four steps: "Enter Sphere Credentials" (marked with a checkmark), "Reset Password" (marked with a checkmark), "3 Select Location" (the current step, marked with a circle containing the number 3), and "4 Confirm" (marked with a circle containing the number 4). Below the progress indicator is a text input field with the placeholder text "Select a Location" and a downward arrow on the right side, indicating a dropdown menu. At the bottom right of the dialog box are two buttons: "PREVIOUS" and "NEXT".

Figure 12 The Register with Sphere - Select Location screen.

2. Confirm the information is correct, and click **Register**.

Wait while the BlackPearl system registers with the Vail sphere. This may take several minutes, during which time the Vail management console may display communication errors.

CHAPTER 4 - CONFIGURE THE SPECTRA VAIL APPLICATION

This chapter describes the configuration steps for the Spectra Vail application.

Log In to the Vail Management Console	62
Create Storage	63
Create BlackPearl Storage	64
Create BlackPearl Standard Bucket Storage	64
Create BlackPearl Legacy Storage	67
Create BlackPearl Volume Pool Storage	70
Create Cloud Storage	74
Create Amazon S3 Cloud Storage	74
Create Microsoft Azure Cloud Storage	79
Create Google Cloud Platform Storage	83
Create S3 Compatible Cloud Storage	86
Create Storage Groups	90
Considerations for Expanding Storage	90
Create a Storage Group	90
Create a Lifecycle	91
Create a Vail Bucket	100
Configure an Object Storage Browser	108
Configure S3 Browser	108
Configure Cyberduck Object Storage Browser	109

LOG IN TO THE VAIL MANAGEMENT CONSOLE

Use the instructions below to log in to the Vail management console.

1. Use one of the following methods:
 - Open a compatible web browser and enter the Vail management console URL into the address bar.
 - In the BlackPearl Nearline Object Gateway management console, select **Configuration > Services**, then double-click the Vail service, and click the **Endpoint** URL displayed on the Vail Service screen.

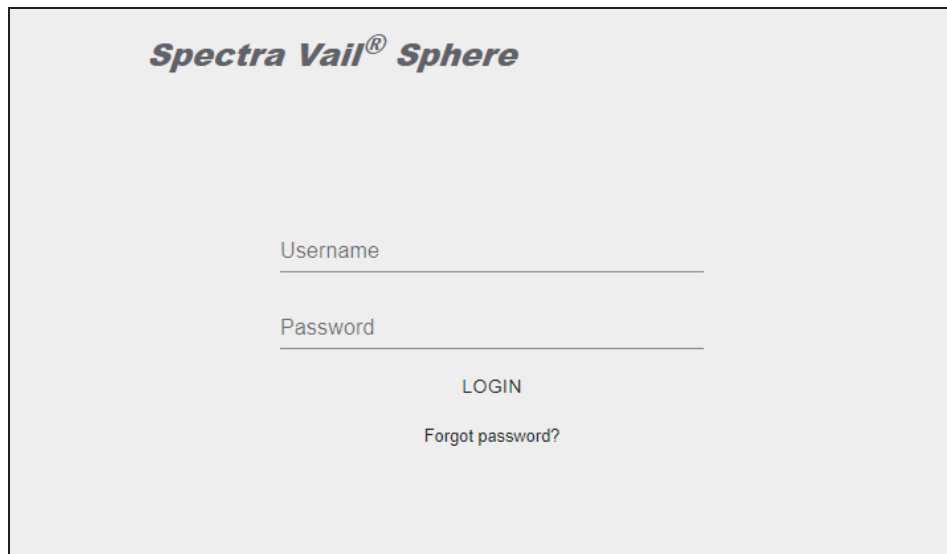


Figure 13 The Spectra Vail Sphere Login screen.

Note: Your web browser may display an invalid certificate warning page. Resolve the warning, and continue to [Step 2](#) below.

2. Use one of the below methods:
 - For a **cloud control** system - Enter the **Username** and **Password** you specified when you registered the first BlackPearl system with the Vail sphere.
 - For a **local control** system - Enter the **Username** and **Password** of the BlackPearl system administrator.
3. Click **LOGIN**.

CREATE STORAGE

Storage is used by the Spectra Vail application as targets for S3 clients and lifecycles to store data. There are two basic types of storage: endpoint storage and cloud storage. Endpoint storage includes a BlackPearl Nearline Object Gateway, or block VM storage such as a Spectra Vail VM node. Cloud storage is S3 object storage on AWS or other S3 cloud storage provider.

Use one of the sections below to create storage.

- **Create BlackPearl Storage on the next page**
- **Create Cloud Storage on page 74**
- **Create a Vail VM Node on page 268**

CREATE BLACKPEARL STORAGE

BlackPearl storage uses a bucket or volume pool configured on a BlackPearl Nearline Object Gateway. You can select the same BlackPearl Nearline Object Gateway multiple times when creating BlackPearl storage, but each storage instance must use a unique bucket or volume pool.

Note: Before you can create BlackPearl storage in the Vail management console, you must register the BlackPearl Nearline Object Gateway with the Spectra Vail application. See [Register a BlackPearl Nearline Object Gateway with a Vail Sphere on page 52](#).

When creating BlackPearl storage, you can select to create storage using a standard bucket, or to create storage using a linked bucket.

Use one of the sections below:

- **Create BlackPearl Standard Bucket Storage below**
- **Create BlackPearl Legacy Storage on page 67**
- **Create BlackPearl Volume Pool Storage on page 70**

Create BlackPearl Standard Bucket Storage

The instructions below assume a storage domain and data policy were previously configured on your BlackPearl Nearline Object Gateway. For information on configuring a storage domain and data policy, see the [BlackPearl Nearline Gateway User Guide](#).

A BlackPearl bucket does not need to be created before creating BlackPearl storage in the Spectra Vail application. A bucket is created automatically on the BlackPearl system during the process described below.

Here is how to create BlackPearl storage:

1. Log in to the Vail management console.
2. In the taskbar of the Vail management console, click **Storage**.

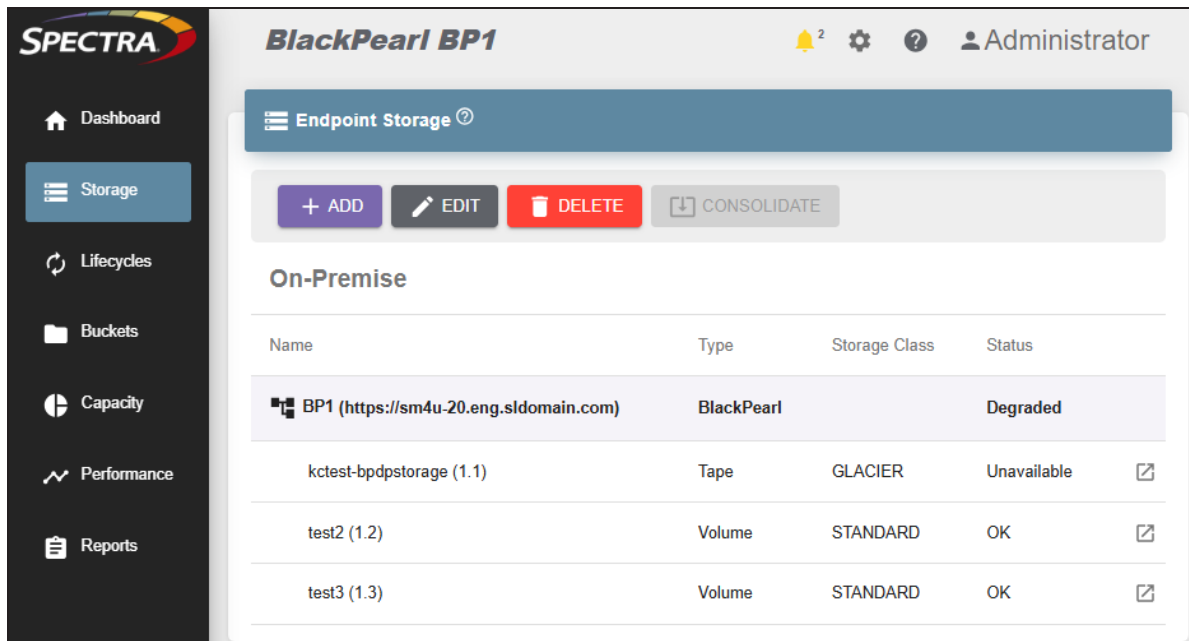


Figure 14 The Storage screen.

3. Select the row of the BlackPearl Nearline Object Gateway where you want to create storage.
4. Under the **Endpoint Storage** banner, click **Add**.
5. Use the **Select Storage Type** drop-down menu to select **BlackPearl Data Policy**, then click **Next**.

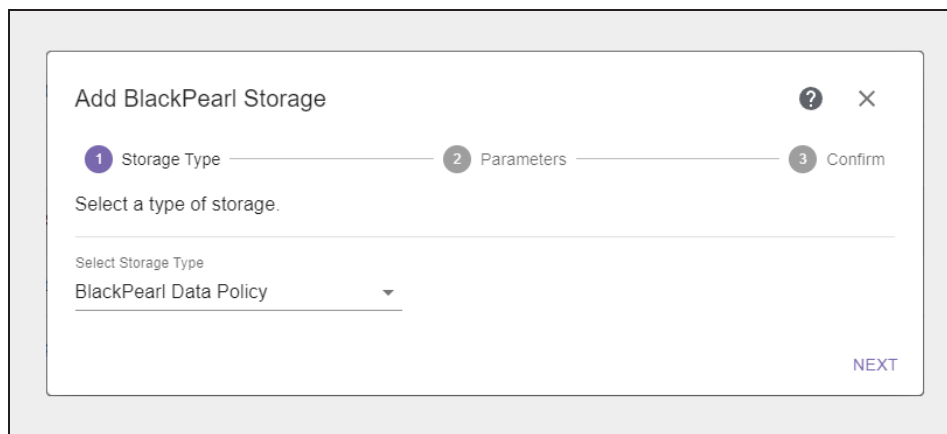


Figure 15 The Add BlackPearl Storage - Select Storage Type screen.

- Use the **Select BlackPearl Data Policy** drop-down menu to select a previously configured data policy on the BlackPearl Nearline Object Gateway. Only data policies configured to use Object Naming display in the drop-down menu.

Figure 16 The Add BlackPearl Storage - Parameters screen.

- Enter a **Storage Name** for the BlackPearl storage.

Spectra Logic recommends using names that include the location, BlackPearl name, physical medium, and storage class.

For example, in the Dallas location add suffixes for the BlackPearl name, physical medium and storage class such as Dallas-BlackPearl1-Object-SA and Dallas-BlackPearl2-Tape-Glacier.

- Use the **Select Storage Class** drop-down menu to select the storage class for the BlackPearl storage. The selections that display depend on the type of storage medium targeted by the BlackPearl data policy.
- If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.

10. If desired, select to enable **Third-party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also causes storage of full object metadata, enabling you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

11. If desired, select to enable **Restore In Place**.

- If enabled, the Vail application creates an additional data clone of objects on the same storage where the object exists.
- If you do not enable this option, an additional data clone of objects is created on different storage.

12. If desired, edit the **Caution Threshold** and **Warning Threshold**. These settings control when the Spectra Vail application sends a notification that the selected bucket capacity reaches the configured thresholds.

13. If desired, select **Disabled**. This creates the endpoint storage in a disabled state. The endpoint storage must be manually enabled before it can be used for data storage operations.

14. Click **Next**.

15. Review the configuration and click **Submit** to create the BlackPearl storage. The BlackPearl bucket to be used with this storage is automatically created on the BlackPearl system.

Create BlackPearl Legacy Storage

BlackPearl legacy storage allows you to link a BlackPearl bucket with a Spectra Vail bucket. When linking these buckets, changes made in the BlackPearl bucket are not automatically detected by the Spectra Vail application. The Spectra Vail bucket must be manually scanned to determine object changes in the BlackPearl bucket. See [Scan a Vail Bucket on page 1](#) for instructions on scanning a Spectra Vail bucket.

Here is how to create BlackPearl legacy storage:

1. Log in to the Vail management console.
2. If necessary, create a bucket as described in [Create a Vail Bucket on page 100](#), then return to this section.
3. In the taskbar of the Vail management console, click **Storage**.
4. Select the row of the BlackPearl Nearline Object Gateway where you want to create the storage endpoint.
5. Under the **Endpoint Storage** banner, click **Add**.

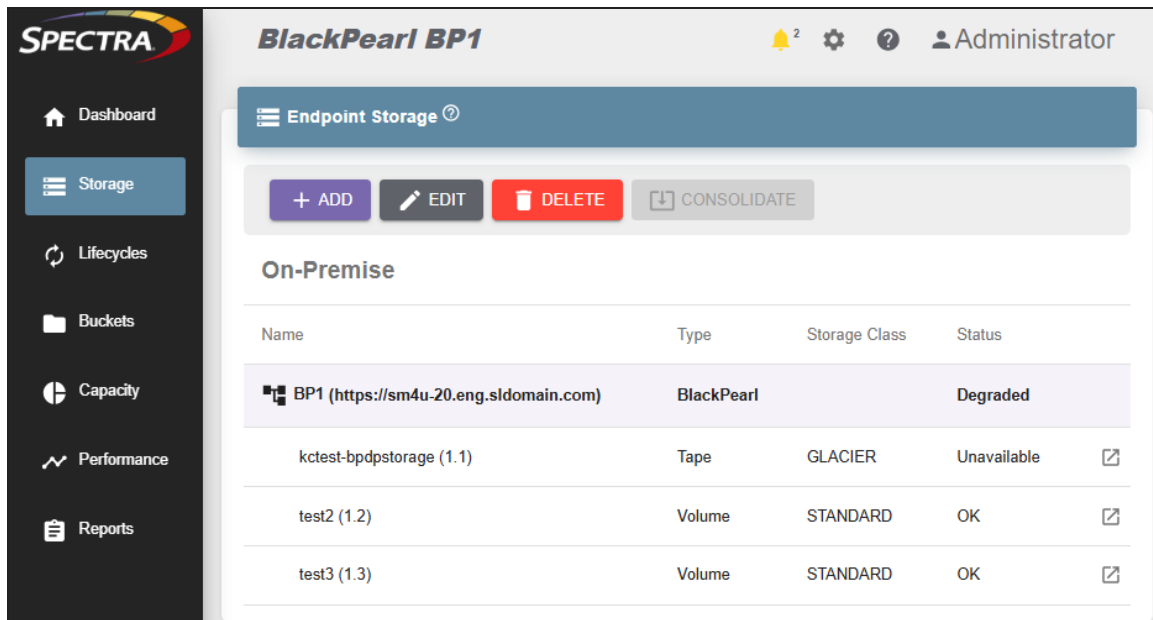


Figure 17 The Storage screen.

- Use the **Select Storage Type** drop-down menu to select **BlackPearl Legacy Storage**, then click **Next**.

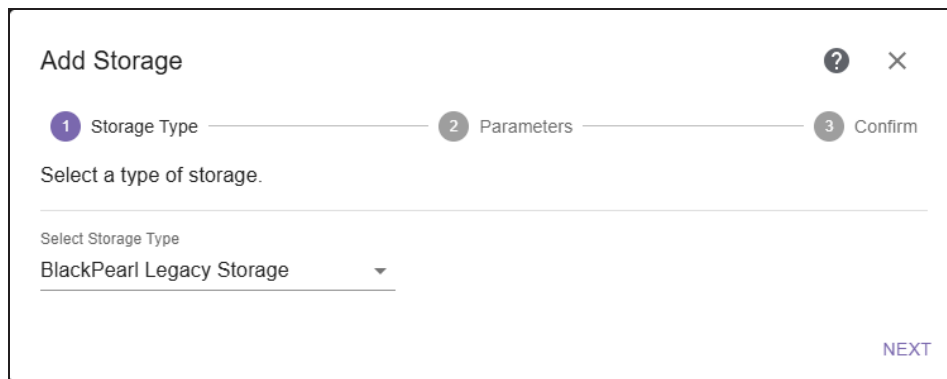


Figure 18 The Add BlackPearl Storage - Select Storage Type screen.

- Use the **Select BlackPearl Linked Bucket** drop-down menu to select a previously configured BlackPearl bucket.

Add Storage ? X

✓ Storage Type — 2 Parameters — 3 Confirm

Configure your storage below.

Select BlackPearl Linked Bucket
Spectra-BlackPearl-Backup-BP1-50... ? Storage Name

Select Storage Class
GLACIER

Use With Buckets
All buckets ?

☐ Restore In-place ?

☐ Disabled

PREVIOUS NEXT

Figure 19 The Add Storage - Parameters screen.

8. Enter a **Storage Name** for the BlackPearl storage.

Spectra Logic recommends using names that include the location, BlackPearl name, physical medium, and storage class.

For example, in the Dallas location add suffixes for the BlackPearl name, physical medium and storage class such as Dallas-BlackPearl1-Object-SA and Dallas-BlackPearl2-Tape-Glacier.

9. Use the **Select Storage Class** drop-down menu to select the storage class for the BlackPearl storage. The selections that display depend on the type of storage medium targeted by the BlackPearl data policy used by the BlackPearl bucket selected in [Step 7](#).
10. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.
11. If desired, select to enable **Restore In Place**.
 - If enabled, the Vail application creates an additional data clone of objects on the same storage where the object exists.
 - If you do not enable this option, an additional data clone of objects is created on different storage.

12. If desired, select **Disabled**. This creates the endpoint storage in a disabled state. The endpoint storage must be manually enabled before it can be used for data storage operations.

13. Click **Next**.

14. Review the configuration and click **Submit** to create the BlackPearl linked bucket storage.

Create BlackPearl Volume Pool Storage

BlackPearl volume pool storage is NAS storage provided by a BlackPearl system.

The instructions below assume a storage pool was previously configured on your BlackPearl Nearline Object Gateway. For information on configuring a storage pool, see the [BlackPearl Nearline Gateway User Guide](#).

Here is how to create BlackPearl volume pool storage:

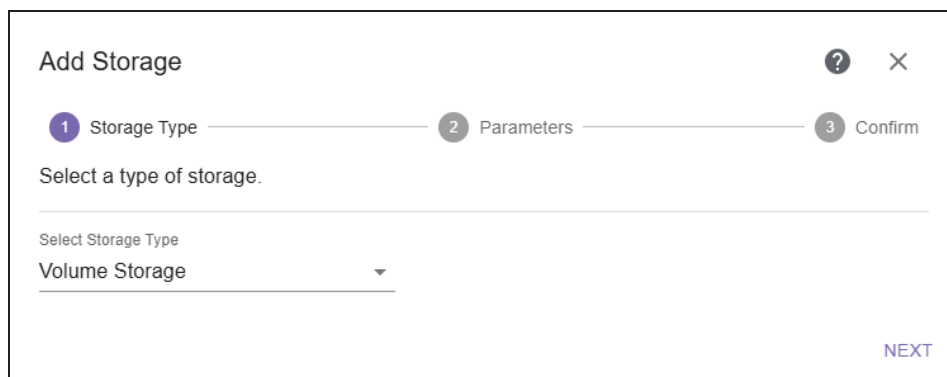
1. In the taskbar of the Vail management console, click **Storage**.
2. Select the row of the BlackPearl Nearline Object Gateway where you want to create storage.
3. Under the **Endpoint Storage** banner, click **Add**.

Name	Type	Storage Class	Status
BP1 (https://sm4u-20.eng.sldomain.com)	BlackPearl		Degraded
kctest-bdpstorage (1.1)	Tape	GLACIER	Unavailable
test2 (1.2)	Volume	STANDARD	OK
test3 (1.3)	Volume	STANDARD	OK

Figure 20 The Storage screen.

4. Use the **Select Storage Type** drop-down menu to select **Volume Storage**, then click **Next**.

Note: This option only displays if a volume is detected on the BlackPearl system.



The screenshot shows a dialog box titled "Add Storage" with a close button (X) and a help button (?). Below the title is a progress bar with three steps: 1. Storage Type (active), 2. Parameters, and 3. Confirm. Below the progress bar, it says "Select a type of storage." and "Select Storage Type". A dropdown menu shows "Volume Storage" selected. A "NEXT" button is in the bottom right corner.

Figure 21 The Add BlackPearl Storage - Select Storage Type screen.

5. Use the **Select BlackPearl Volume Pool** drop-down menu to select a previously configured storage pool on the BlackPearl system. The Parameters screen updates to show available options based on the storage pool selected.

Add Storage ? ×

✓ **Storage Type** 2 **Parameters** 3 **Confirm**

Configure your storage below.

Select BlackPearl Volume Pool
test ? **Storage Name**

Select Storage Class
STANDARD

Use With Buckets
All buckets ?

☐ Third-party Recovery ?

Caution Threshold 70 % ? Warning Threshold 80 % ?

Optional Data 70 % ?

Write Limit MiB/sec ?

Quota 3.5 TiB available Units MiB ?

☐ Disabled

[PREVIOUS](#) [NEXT](#)

Figure 22 The Add BlackPearl Volume Pool Storage - Parameters screen.

6. Enter a **Storage Name.**

Spectra Logic recommends using names that include the location, BlackPearl name, physical medium, and storage class.

For example, in the Dallas location add suffixes for the BlackPearl name, physical medium and storage class such as Dallas-BlackPearl1-Object-SA and Dallas-BlackPearl2-Tape-Glacier.

- 7. Using the **Select Storage Class** drop-down menu, select the storage class you want to use for this volume pool storage endpoint.**
- 8. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.**

9. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

10. If desired, edit the **Caution Threshold** and **Warning Threshold**. These settings control when the Spectra Vail application sends a notification that the selected bucket capacity reaches the configured thresholds.

11. If desired, edit the **Optional Data** percentage. This setting controls the amount of space used by optional data clones, which use available storage space to speed up data access. When the system reaches the percentage value specified, optional clones are deleted to maintain the percentage of used storage space under the specified value.

Note: If this field is left blank, no optional clones are stored and object access times are not recorded.

12. If desired, enter a value for **Write Limit**. This setting reduces the write speed when volume usage exceeds the caution threshold. If the storage usage is below the caution threshold, writes are unlimited. The configured rate limit is enforced at the caution threshold, and the rate is then reduced linearly as the volume usage increases. The rate limit applies to lifecycle processing and all data writes to the given storage.

13. If desired, enter a value for a **Quota**, and use the **Units** drop-down menu to select a unit size for the quota value. This setting controls the maximum amount of storage space on the storage pool that is used for the BlackPearl volume pool storage endpoint. When this percentage is reached, no additional data is added to the storage endpoint. If you do not want to use a quota limit, leave the fields blank.

Notes:

- Spectra Logic recommends setting a quota of 90% of volume storage space, or lower if desired.
- This setting can be modified after the BlackPearl volume pool storage is created.

14. If desired, select **Disabled**. This creates the endpoint storage in a disabled state. The endpoint storage must be manually enabled before it can be used for data storage operations.

15. Click **Next**.

16. Confirm all settings are correct and click **Submit**.

CREATE CLOUD STORAGE

Use one of the sections below to configure cloud storage:

- **Create Amazon S3 Cloud Storage below**
- **Create Microsoft Azure Cloud Storage on page 79**
- **Create Google Cloud Platform Storage on page 83**
- **Create S3 Compatible Cloud Storage on page 86**

Create Amazon S3 Cloud Storage

In Vail, Amazon S3 cloud storage uses a previously configured AWS endpoint target for object storage.

Here is how to create Amazon S3 cloud storage:

1. In the taskbar of the Vail management console, click **Storage**.
2. Under the **Cloud Storage** banner, click **Add**.

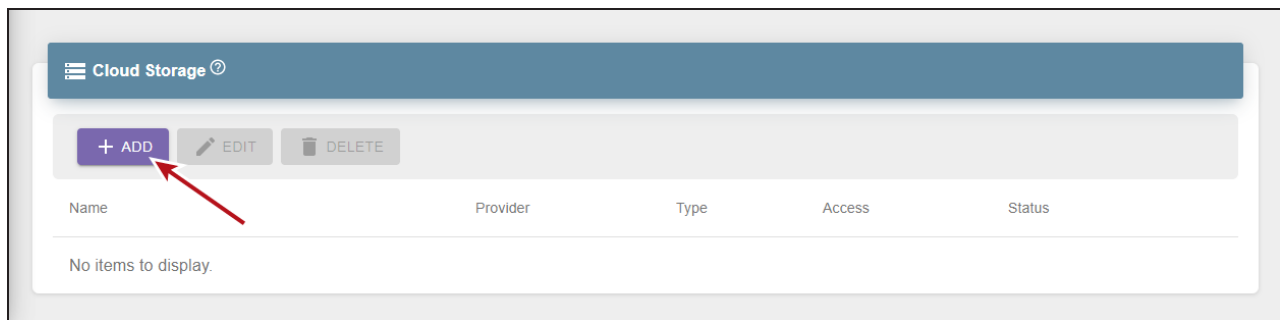


Figure 23 The Cloud Storage pane.

3. Use the **Select Storage Type** drop-down menu to select **Amazon S3 Storage**, and click **Next**.




Figure 24 The Add Cloud Storage - Storage Type screen.

Add Cloud Storage ? X

✓ Cloud Provider — 2 Authorization — 3 Parameters — 4 Confirm

Enter your authorization data below.

Region AWS Access Key ID

AWS Secret Access Key AWS IAM Role ARN ?

AWS IAM Role External ID
Optional. Specifies who can assume the role.

PREVIOUS NEXT

Figure 25 The Add Cloud Storage - AWS Authorization screen.

4. Enter the **Region**, **AWS Access Key ID**, **AWS Secret Access Key**, and the **AWS IAM Role ARN** of the account. Optionally, enter the **AWS IAM Role External ID**.
5. Click **Next**.

6. Use the **Select Amazon S3 Cloud Bucket** drop-down menu to select a cloud bucket associated with the AWS or IAM user configured for cloud storage.

Note: AWS Buckets must be configured to use versioning before they can be used as cloud storage, even if they are assigned to a Spectra Vail bucket that has versioning disabled. Although the AWS bucket is capable of storing multiple versions of an object, if the Spectra Vail bucket does not have versioning enabled, only the latest version is preserved in the AWS bucket.

Add Cloud Storage ? ×

✓ Storage Type — ✓ Authorization — 3 Parameters — 4 Confirm

Provide additional information for this cloud storage below.

Select Amazon S3 Bucket: chrisr-test ?

Storage Name: chrisr-test

Select Storage Class: STANDARD

Destination Class: STANDARD ?

Use With Buckets: kctest1 ?

☐ Import Content ?

☐ Pause Notifications

☐ Export Objects ?

☐ Current Version Only ?

☐ Third-party Recovery ?

☐ Restore In-place ?

☐ Disabled

PREVIOUS NEXT

Figure 26 The Add Cloud Storage - Parameters screen.

7. The **Storage Name** field is automatically populated with the name of the bucket selected in Step 6. If desired, you can change the **Storage Name**. Spectra Logic recommends using names that include type of cloud storage, location, and storage class.

For example, use names for Amazon cloud storage such as AWS_uswest2_autotier and AWS_uswest2_S3glacier.

8. Use the **Select Storage Class** drop-down menu to select a storage class for the AWS S3 storage.
9. Use the drop-down menu to select the **Destination Class**.

If you select **Use Default**, objects are created without specifying a storage class value. Otherwise, Vail attempts to create objects on the storage with the specified storage class.

- Notes:**
- If a Lifecycle using this storage has **Ignore Storage Class** enabled, this setting is ignored.
 - The financial costs associated with each storage type are controlled by the cloud provider.

10. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.

11. If desired, select **Import Content**. This option requires you to select a single bucket with the **Use With Buckets** drop-down menu. The selected bucket stores object references without copying data. Lifecycles can be used to copy data if desired.

Note: Including this storage endpoint as a destination in a lifecycle allows you to synchronize the selected bucket with the storage endpoint. Synchronizing allows the bucket to maintain consistency with an external bucket regardless of whether objects are added locally or to the external bucket.

12. Select or clear **Pause Notifications** as desired. When notifications are paused, changes are only recognized when the bucket scan is manually triggered, either through the Vail management console or by API call. Otherwise changes are automatically recognized and linked bucket contents update automatically.

13. If desired, select to enable **Export Objects**. This option configures the cloud storage as write-only and directly migrates data to the cloud bucket instead of creating and managing clones of the object data. After the export completes, the Spectra Vail application does not retain any record of the exported data.

If you enable this option:

- The Third-Party Recovery option is enabled automatically.
- You cannot enable the Restore In-Place option.
- You cannot use the storage as part of a Lifecycle.
- The storage does not display capacity or usage information.

**IMPORTANT**

Once enabled, you cannot disable this setting for the cloud storage. Additionally, you cannot enable Export Objects at a later time, only when you are creating the cloud storage.

14. If desired, select **Current Version Only**. This option limits storing object clones on the storage endpoint to the most recent version only. This setting enables the use of cloud storage that does not support object versioning.

Note: This option can only be enabled during creation.

15. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

16. If you selected Glacier or Deep Archive as the storage class, you may select to enable **Restore In Place**. The restore in-place option uses internal resources on archival storage to create a clone in the same storage. If this setting is not enabled, new clones are created on different storage. Selecting this option may use additional cache space or incur additional cloud storage fees.

Note: You cannot enable this setting if you selected to enable Export Objects.

17. If desired, select **Disabled**. This creates the cloud storage in a disabled state. The cloud storage must be manually enabled before it can be used for data storage operations.

18. Click **Next**.

19. Verify the information for the cloud storage is correct, and click **Submit**.

Notes:

- There is a seven minute delay before the contents of the AWS bucket appear in the Spectra Vail bucket. If the Spectra Vail bucket is assigned to a lifecycle that is configured to run immediately, any data present in the AWS bucket is processed by the lifecycle after seven minutes.

Create Microsoft Azure Cloud Storage

In the Spectra Vail application, Microsoft Azure cloud storage uses a previously configured Azure container for storage.

Here is how to create Microsoft Azure cloud storage:

1. In the taskbar of the Vail management console, click **Storage**.
2. Under the **Cloud Storage** banner, click **Add**.




Figure 27 The Cloud Storage pane.

3. Use the **Select Storage Type** drop-down menu to select **Azure Cloud Storage**, and click **Next**.




Figure 28 The Add Cloud Storage - Storage Type screen.

4. Enter the **Storage Account** and **Shared Secret** information for the Azure container.

Note: You cannot change the Storage Account after you create the cloud storage.

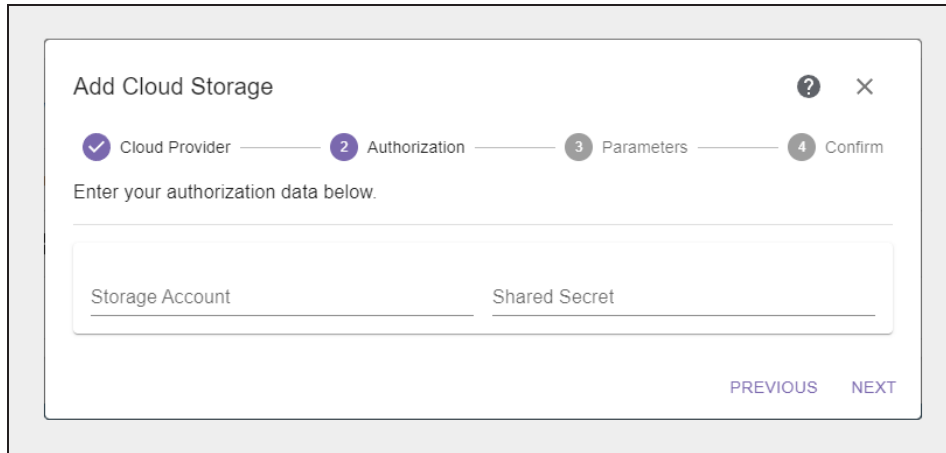


Figure 29 The Add Cloud Storage - Authorization screen.

5. Click **Next**.

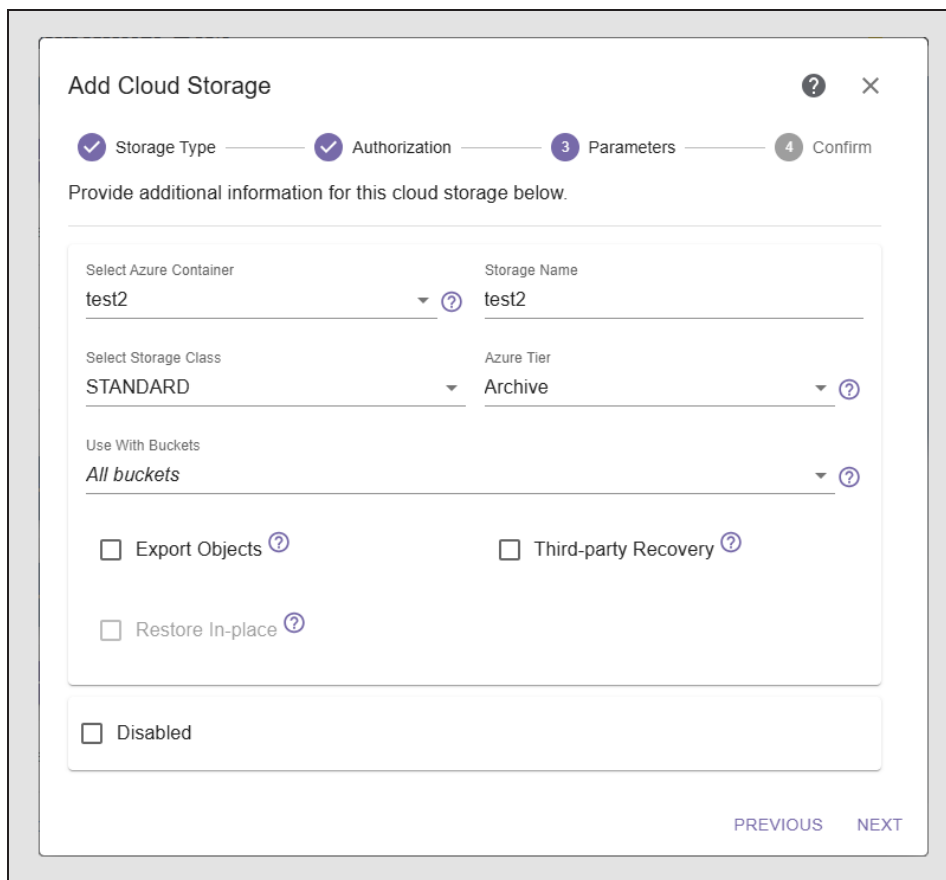


Figure 30 The Add Cloud Storage - Parameters screen.

6. Using the **Select Azure Container** drop-down menu, select a previously created container on the Azure storage target.

7. If desired, you can change the **Storage Name**. Spectra Logic recommends using names that include type of cloud storage, location, and storage class. For example, use names for Azure cloud storage such as Azure-Production-Denver.

8. Use the **Select Storage Class** drop-down menu to select a storage class for the Azure storage endpoint.

9. Use the drop-down menu to select the **Azure Tier**.

If you select **Use Default**, objects are created without specifying a storage class value. Otherwise, Vail attempts to create objects on the storage with the specified storage class.

Notes:

- If a Lifecycle using this storage has **Ignore Storage Class** enabled, this setting is ignored.

- The financial costs associated with each storage type are controlled by the cloud provider.

10. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.

11. If desired, select to enable **Export Objects**. This option configures the cloud storage as write-only and directly migrates data to the cloud bucket instead of creating and managing clones of the object data. After the export completes, the Spectra Vail application does not retain any record of the exported data.

If you enable this option:

- The Third-Party Recovery option is enabled automatically.
- You cannot enable the Restore In-Place option.
- You cannot use the storage as part of a Lifecycle.
- The storage does not display capacity or usage information.

**IMPORTANT**

Once enabled, you cannot disable this setting for the cloud storage. Additionally, you cannot enable Export Objects at a later time, only when you are creating the cloud storage.

12. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

13. If you selected Glacier or Deep Archive as the storage class, you may select to enable **Restore In Place**. The restore in-place option uses internal resources on archival storage to create a clone in the same storage. If this setting is not enabled, new clones are created on different storage. Selecting this option may use additional cache space or incur additional cloud storage fees.

Note: You cannot enable this setting if you selected to enable Export Objects.

14. If desired, select **Disabled**. This creates the cloud storage in a disabled state. The cloud storage must be manually enabled before it can be used for data storage operations.

15. Click **Next**.

16. Verify the information for the cloud storage is correct, and click **Submit**.

Create Google Cloud Platform Storage

In Vail, Google Cloud Platform storage uses a previously configured Google storage endpoint target for storage.

Here is how to create Google Cloud Platform storage:

1. In the taskbar of the Vail management console, click **Storage**.
2. Under the **Cloud Storage** banner, click **Add**.




Figure 31 The Cloud Storage pane.

3. Use the **Select Storage Type** drop-down menu to select **Google Cloud Storage**, and click **Next**.

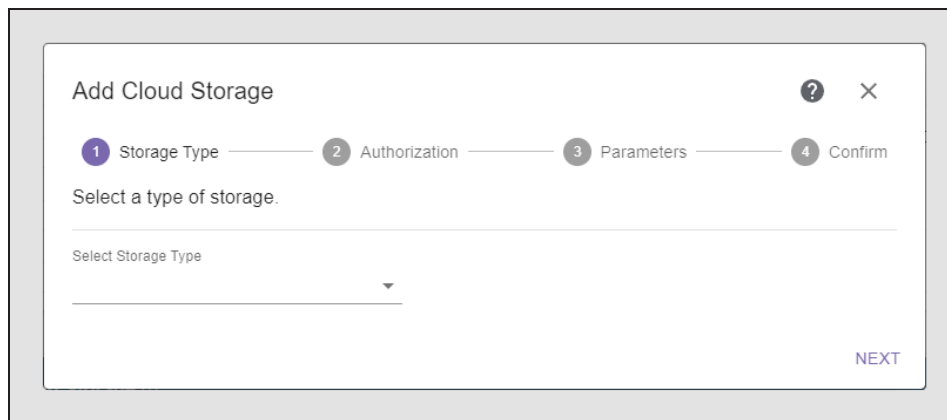


Figure 32 The Add Cloud Storage - Storage Type screen.

4. Enter the **Google Could Platform JSON Credentials** information for the endpoint.

Add Cloud Storage

1 Cloud Provider — 2 Authorization — 3 Parameters — 4 Confirm

Enter your authorization data below.

Google Cloud Platform JSON Credentials

PREVIOUS NEXT

Figure 33 The Add Cloud Storage - Authorization screen.

5. Click **Next**.

Add Cloud Storage

1 Storage Type — 2 Authorization — 3 Parameters — 4 Confirm

Provide additional information for this cloud storage below.

Select Google Cloud Bucket: bucket-for-chris

Storage Name: bucket-for-chris

Select Storage Class: STANDARD

Google Class: Use Default

Use With Buckets: All buckets

☐ Export Objects ☐ Third-party Recovery ☐ Disabled

PREVIOUS NEXT

Figure 34 The Add Cloud Storage - Parameters screen.

6. Using the **Select Google Cloud Bucket** drop-down, select a previously created bucket in the Google Cloud Storage target.

7. Enter a **Storage Name**. Spectra Logic recommends using names that include type of cloud storage, location, and storage class. For example, use names for Google cloud storage such as Google-Production-Denver.
8. Use the **Select Storage Class** drop-down menu to select a storage class for the Google cloud storage endpoint.
9. Use the drop-down menu to select the **Google Class**.

If you select **Use Default**, objects are created without specifying a storage class value. Otherwise, Vail attempts to create objects on the storage with the specified storage class.

- Notes:**
- If a Lifecycle using this storage has **Ignore Storage Class** enabled, this setting is ignored.
 - The financial costs associated with each storage type are controlled by the cloud provider.
10. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.
 11. If desired, select to enable **Export Objects**. This option configures the cloud storage as write-only and directly migrates data to the cloud bucket instead of creating and managing clones of the object data. After the export completes, the Spectra Vail application does not retain any record of the exported data.

If you enable this option:

- The Third-Party Recovery option is enabled automatically.
- You cannot enable the Restore In-Place option.
- You cannot use the storage as part of a Lifecycle.
- The storage does not display capacity or usage information.

**IMPORTANT**

Once enabled, you cannot disable this setting for the cloud storage. Additionally, you cannot enable Export Objects at a later time, only when you are creating the cloud storage.

12. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

13. If desired, select **Disabled**. This creates the cloud storage in a disabled state. The cloud storage must be manually enabled before it can be used for data storage operations.
 14. Click **Next**.
 15. Verify the information for the cloud storage is correct, and click **Submit**.
-

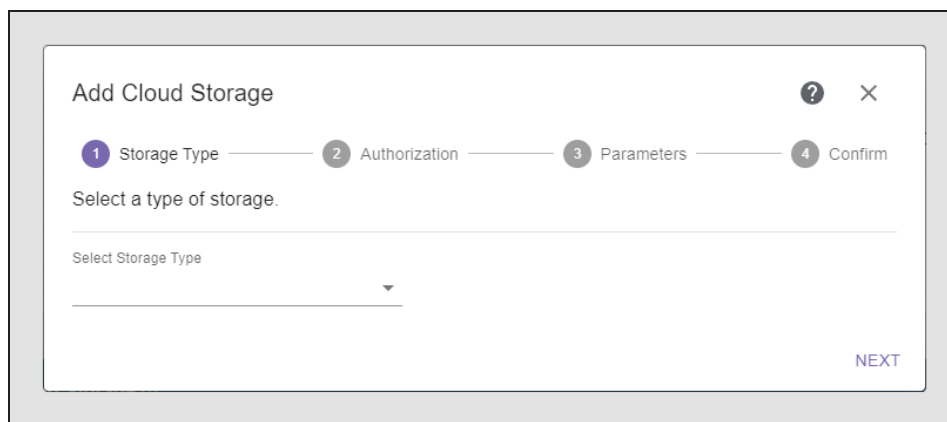
Create S3 Compatible Cloud Storage

Cloud storage that is not an AWS, Azure, or Google Cloud endpoint is configured as S3 compatible cloud storage.

Note: The bucket on the cloud storage target must be configured to use versioning.

Here is how to create S3 compatible cloud storage:

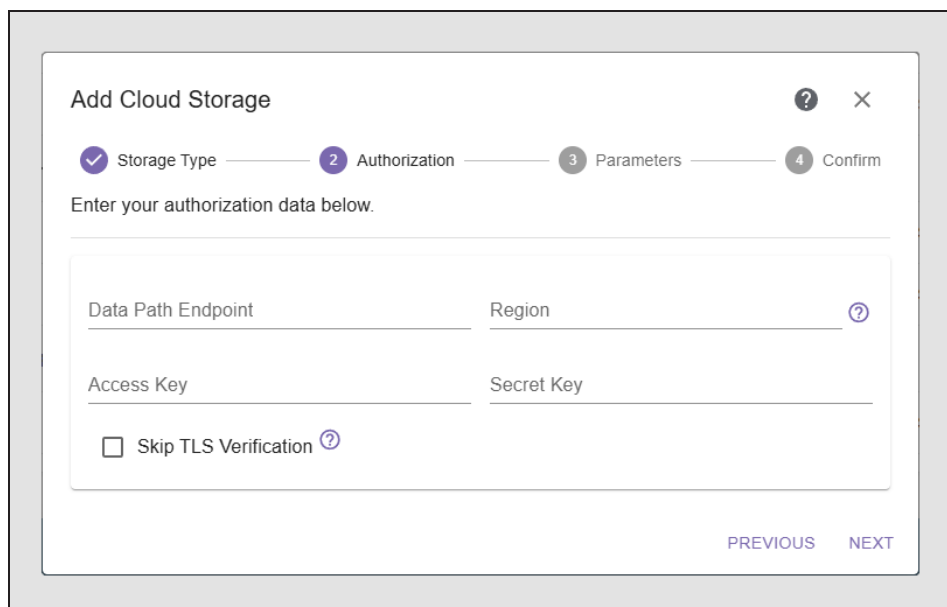
1. In the taskbar of the Vail management console, click **Storage**.
2. Under the **Cloud Storage** banner, click **Add**.
3. Use the **Select Storage Type** drop-down menu to select **S3 Compatible Storage** and click **Next**.



The screenshot shows a modal window titled "Add Cloud Storage" with a progress bar at the top indicating four steps: 1. Storage Type (active), 2. Authorization, 3. Parameters, and 4. Confirm. Below the progress bar, the text "Select a type of storage." is followed by a dropdown menu labeled "Select Storage Type". A "NEXT" button is located in the bottom right corner.

Figure 35 The Add Cloud Storage - Storage Type screen.

4. Enter the URL address for the **Data Path Endpoint**.



The screenshot shows the "Add Cloud Storage" dialog box at the "Authorization" step. The progress bar shows "Storage Type" as completed and "Authorization" as the current step. The text "Enter your authorization data below." is followed by a form with four input fields: "Data Path Endpoint", "Region", "Access Key", and "Secret Key". There is also a checkbox labeled "Skip TLS Verification". A "PREVIOUS" button and a "NEXT" button are at the bottom right.

Figure 36 The Add Cloud Storage - S3 Compatible Storage Authorization screen.

5. If required, enter a **Region** where the S3 compatible storage is located. If the region is not required, leave this field empty.
6. Enter the **Access Key** and **Secret Key** for the administrator of the cloud endpoint.
7. If desired, select **Skip TLS Verification**. This option disables TLS certificate verification for HTTPS endpoints.

Note: This setting does not apply to HTTP endpoints.

8. Click **Next**.

9. Using the **Select Cloud Bucket** drop-down menu, select a bucket previously configured on the cloud endpoint. The Parameters screen updates to show options applicable to the type of bucket selected.

Note: Versioning must be enabled on the target bucket.

Add Cloud Storage ? X

✓ Storage Type — ✓ Authorization — 3 Parameters — 4 Confirm

Provide additional information for this cloud storage below.

Select Cloud Bucket: robq-wasabi ?

Storage Name: robq-wasabi

Select Storage Class: STANDARD ?

Destination Class: Use Default ?

Use With Buckets: kctest1 ?

☐ Import Content ?

☐ Pause Notifications

☐ Export Objects ?

☐ Current Version Only ?

☐ Third-party Recovery ?

☐ Restore In-place ?

Addressing Style ?

☒ Path ☐ Virtual Hosted

☐ Disabled

PREVIOUS NEXT

Figure 37 The Add Cloud Storage - S3 Compatible Parameters screen displaying all possible settings. Your screen may appear different depending on the type of bucket selected.

10. The **Storage Name** is automatically populated with the name of the bucket selected in [Step 9](#). If desired, you can change the **Storage Name**.
11. Use the **Select Storage Class** drop-down menu to configure the storage class you want to use for this endpoint. The selected storage class is used when creating clones on the cloud bucket.

Note: The financial costs associated with each storage type are controlled by the cloud provider.

12. Use the drop-down menu to select the **Destination Class**.

If you select **Use Default**, objects are created without specifying a storage class value. Otherwise, Vail attempts to create objects on the storage with the specified storage class.

Notes:

- If a Lifecycle using this storage has **Ignore Storage Class** enabled, this setting is ignored.
- The financial costs associated with each storage type are controlled by the cloud provider.

13. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.

14. If desired, select **Import Content**. This option requires you to select a single bucket with the **Use With Buckets** drop-down menu. The selected bucket stores object references without copying data. Lifecycles can be used to copy data if desired.

Note: Including this storage endpoint as a destination in a lifecycle allows you to synchronize the selected bucket with the storage endpoint. Synchronizing allows the bucket to maintain consistency with an external bucket regardless of whether objects are added locally or to the external bucket.

15. Select or clear **Pause Notifications** as desired. When notifications are paused, changes are only recognized when the bucket scan is manually triggered, either through the Vail management console or by API call. Otherwise changes are automatically recognized and linked bucket contents update automatically.

16. If desired, select to enable **Export Objects**. This option configures the cloud storage as write-only and directly migrates data to the cloud bucket instead of creating and managing clones of the object data. After the export completes, the Spectra Vail application does not retain any record of the exported data.

If you enable this option:

- The Third-Party Recovery option is enabled automatically.
- You cannot enable the Restore In-Place option.
- You cannot use the storage as part of a Lifecycle.
- The storage does not display capacity or usage information.

**IMPORTANT**

Once enabled, you cannot disable this setting for the cloud storage. Additionally, you cannot enable Export Objects at a later time, only when you are creating the cloud storage.

17. If desired, select **Current Version Only**. This option limits storing object clones on the storage endpoint to the most recent version only. This setting enables the use of cloud storage that does not support object versioning.

Note: This option can only be enabled during creation.

18. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

19. If you selected Glacier or Deep Archive as the storage class, you may select to enable **Restore In Place**. The restore in-place option uses internal resources on archival storage to create a clone in the same storage. If this setting is not enabled, new clones are created on different storage. Selecting this option may use additional cache space or incur additional cloud storage fees.

Note: You cannot enable this setting if you selected to enable Export Objects.

20. Select the desired **Addressing Style**. This setting controls the URL format used when communicating with the cloud storage provider.

Selection	Description
Path Style	Path style formatting uses the bucket name as part of the URL path. Example: <i>http://endpoint/bucket-name/object-key</i>
Virtual-Hosted	Virtual-hosted style addressing uses the bucket as the prefix to the endpoint name Example: <i>http://bucket-name.endpoint/object-key</i>

21. If desired, select to **Disable** the cloud storage. This creates the cloud storage in a disabled state. The cloud storage must be manually enabled before it can be used for data storage operations.

22. Click **Next**.

23. Verify the information for the cloud storage is correct, and click **Submit**.

CREATE STORAGE GROUPS

Storage groups combine smaller existing storage pools into a common pool.

Considerations for Expanding Storage

- When you combine storage, if there are current clone processing jobs to storage, adding new storage to a storage group does not help clones complete if they are already pending when the storage is added. This occurs because the application does not change the destination storage once it has been selected. This is true even if the newly added storage is on the endpoint performing the operations.
- When objects are transferred to a storage group, a check is performed to determine if the storage group has enough space to land the objects. This check is only valid if all storage in the storage group is on the same storage endpoint. If you create a storage group with storage from different endpoints, the Vail application may select a storage destination that does not have sufficient space.

Create a Storage Group

Here is how to create a Storage group:

1. In the taskbar of the Vail management console, click **Storage**.
2. Under the Storage Groups banner, click **Create**.

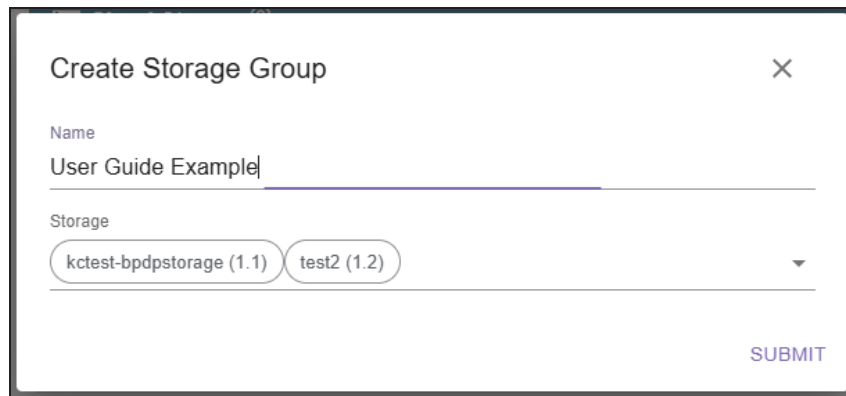


Figure 38 The Create Storage Group screen.

3. Enter a **Name** for the storage group.
4. Use the **Storage** drop-down menu to select the storage endpoints you want to group.
Note: All selected storage endpoints must be the same storage class.
5. Click **Submit**.

CREATE A LIFECYCLE

Lifecycles control where data is located, at what times, and for how long. When data is added to a Vail bucket, lifecycle rules determine where objects are initially placed, how data placement changes over time, and when to delete objects. Placement rules change data placement without altering the bucket contents. Deletion rules delete objects and should be used with caution.

Lifecycle rules are processed once per day. When this occurs, the Vail application generates a list of objects to be moved or expired and then processes the objects as a background process. The default processing time is midnight UTC, but processing time of day can be changed in the Global Settings. See [Change Lifecycle Rule Nightly Processing Time](#) on page 173.



IMPORTANT

The Vail application does not support aggregating storage pools that use the same storage class. You must configure separate Lifecycles that each target different storage to use multiple storage pools of the same storage class.

Here is how to create a lifecycle:

1. In the Vail management console taskbar, click **Lifecycles**.
2. Under the **Lifecycles** banner, click **Create**.

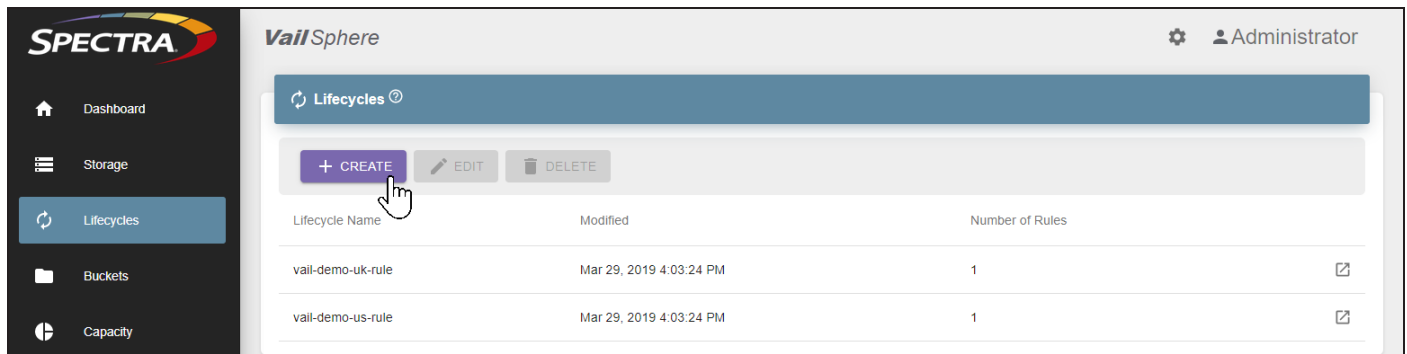


Figure 39 The Lifecycles screen.

3. Enter the desired **Name**.

Spectra Logic recommends using names that directly indicate the specific lifecycle rule configuration.

For example, use names such as Clone_Everywhere_Keep4Days and Moveto_DallasNodeVM_After10Days.

The screenshot shows the 'Create Lifecycle' dialog box with the 'Parameters' tab selected. The 'Name' field is empty. The 'Multipart Upload Expiration' is set to 7 days. The 'Restore To' dropdown is set to 'Use Default'. The 'Delete Marker Expiration' checkbox is checked. The 'Force Initial Copy' and 'Ignore Requested Storage Class' checkboxes are unchecked. There is a 'Description' text area and a 'NEXT' button at the bottom right.

Figure 40 The Create Lifecycle - Parameters screen.

4. Enter a value for **Multipart Upload Expiration** in days. This setting controls how long the Spectra Vail application waits before aborting multipart uploads. When the multipart upload aborts, all parts of the upload are deleted. This prevents retaining multiple incomplete uploads.

Note: To prevent multipart uploads from expiring, enter zero.

5. Use the **Restore To** drop-down menu to select a specific storage pool where you want to restore bucket objects. If you use the default setting, the Vail application decides which storage pool to restore objects.
6. Select or clear **Delete Marker Expiration**. A delete marker keeps track of deletions of versioned objects so that the application can determine if the object is missing. If enabled, the Spectra Vail application removes delete markers when they are the last remaining version of an object.

7. Select or clear **Force Initial Copy**. When enabled, the Spectra Vail application initially places data as STANDARD storage. Additional clones are created immediately as GLACIER storage. This may provide performance advantages as copying clones to GLACIER results in a clone that is ordered sequentially and more optimally packed.
8. Select or clear **Ignore Requested Storage Class**. When enabled, the Vail application does not consider the storage class requested in a PUT or upload operation and instead uses the storage class of the selected storage endpoint.
9. If desired, use the **Description** field to enter any additional information.
10. Click **Next**.
11. Add one or more placement or deletion rules. Placement rules add and remove clones from storage destinations, but do not change bucket contents. Deletion rules delete objects and should be used with caution.

Note: Each lifecycle is limited to five total rules.

- **Add a Placement Rule below**
- **Add a Deletion Rule on page 97**

Add a Placement Rule

Placement rules add object clones to the selected destination storage and optionally remove clones from storage destinations not specified in the placement rule. Placement rules do not alter bucket contents.

1. Click **New Placement Rule**.

Figure 41 The Create Lifecycle - Placement Rule screen.

2. Enter the desired **Name**.
3. Use the **Select Destination Storage** drop-down menu to select up to five previously configured storage destinations

Note: To remove a destination from the list, select the **Select Destination Storage** drop-down menu, and click on the **purple highlighted row** of the destination you want to remove.

4. If desired, select **Remove Other Storage**. This option removes clones from any destination storage not selected in [Step 3](#).

Note: This option only removes object clones. It does not change bucket contents.

5. Use the **Destination Count** drop-down menu to select the number of storage destinations you want to maintain a copy of the data when the rule executes, up to a maximum of five. If you have less than five storage endpoints, you are only able to select a number equal to or less than the number of storage endpoints. If you select **All**, every storage endpoint maintains a copy of object data

- Note:** If you select two destinations, but enter five storage destinations, then two copies of the object are maintained on any of the five specified destinations. The order in which you select destinations is the order the Spectra Vail sphere uses to determine where to store a copy of the data. If a storage destination is not available or busy when the rule executes, the Vail sphere selects the next destination.
6. Using the Filter drop-down menu (1) , select the desired filter, then click the **Add Filter Name** button (2). The screen expands to show the details of the selected filter.




Figure 42 The Select Filter drop-down menu (1) and Add Filter Name button (2).

7. Use the table below to complete configuring the filter.

If you selected...	Do the following...
Schedule Filter	<ul style="list-style-type: none"> Specify a Older Than number of days. When an object is older than this value, the Vail application applies the placement rule on the objects at the next daily processing time. Entering a value of zero applies the placement rule on the next daily processing time, which is between zero and 24 hours later. To apply the placement rule immediately after the lifecycle is created, do not configure a schedule filter. Click the trashcan icon to remove the schedule filter. <p>See Edit Global Settings on page 173 to view the currently configured daily processing time.</p> <p>Note: You can only configure one Schedule filter.</p>
Current Version Filter	<p>No actions are required. The Spectra Vail application applies the placement rule to the current version of objects.</p> <p>Note: If you select the Current Version filter, you cannot select the Non-Current Version filter.</p>
Non-Current Version Filter	<p>No actions are required. The Spectra Vail application applies the placement rule to the non-current version of objects.</p> <p>Note: If you select the Non-Current Version filter, you cannot select the Current Version filter.</p>
Include Name Filter	<p>Enter a regular expression. The placement rule applies to any object with a name that matches the provided expression.</p> <p>Note: If multiple Include Name filters are applied to the placement rule, the placement rule applies if <u>any</u> of the filters match the object name.</p>
Exclude Name Filter	<p>Enter a regular expression. The placement rule applies to any object with a name that matches the provided expression.</p> <p>Note: If multiple Exclude Name filters are applied to the placement rule, the placement rule applies if <u>any</u> of the filters match the object name.</p>
Include Object Tag Filter	<ul style="list-style-type: none"> Enter a Key and Value. The placement rule applies to any object with a matching object tag. The fields are case-sensitive. Object tag filtering is case-sensitive and object tags must be an exact match. If no Value setting is entered, the placement rule applies to all objects with an object tag that matches the specified Key. <p>Note: If multiple Include Object Tag filters are applied to the placement rule, the placement rule applies if <u>any</u> of the filters match the object name.</p>

If you selected...	Do the following...
Exclude Object Tag Filter	<ul style="list-style-type: none"> Enter a Key and Value. The placement rule applies to any object with a matching object tag. The fields are case-sensitive. Object tag filtering is case-sensitive and object tags must be an exact match. If no Value setting is entered, the placement rule applies to all objects with an object tag that matches the specified Key. <p>Note: If multiple Exclude Object Tag filters are applied to the placement rule, the placement rule applies if <u>any</u> of the filters match the object name.</p>

8. If desired, add additional Placement or Deletion rules.

9. Click **Submit**.

Add a Deletion Rule

Use deletion rules to delete objects at a specified interval. If a storage location uses versioning, deletion rules can be configured to delete the latest or previous version of an object, or all versions.

Note: Deletion rules always removes delete markers if the rule criteria are met.



CAUTION

A deletion rule deletes data from **all** storage locations configured in the lifecycle.

1. Click **New Deletion Rule**.

Create Lifecycle

Parameters Rules

Define your rules for **xzc**.

Note: Rules will be sorted based on schedule and version filters after submission. Maximum number of rules is 5.

Deletion Rule

Name

ADD SCHEDULE FILTER

NEW PLACEMENT RULE | NEW DELETION RULE

PREVIOUS SUBMIT

Figure 43 The Create Lifecycle - Deletion Rule screen.

2. Enter the desired **Name**.
3. Using the Filter drop-down menu (1), select the desired filter, then click the **Add Filter Name** button (2). The screen expands to show the details of the selected filter.




Figure 44 The Select Filter drop-down menu (1) and Add Filter Name button (2).

4. Use the table below to complete configuring the filter.

If you selected...	Do the following...
Schedule Filter	<ul style="list-style-type: none"> Specify a Older Than number of days. When an object is older than this value, the Vail application applies the deletion rule on the objects at the next daily processing time. Entering a value of zero applies the deletion rule on the next daily processing time, which is between zero and 24 hours later. To apply the deletion rule immediately after the lifecycle is created, do not configure a schedule filter. Click the trashcan icon to remove the schedule filter. <p>See Edit Global Settings on page 173 to view the currently configured daily processing time.</p> <p>Note: You can only configure one Schedule filter.</p>
Current Version Filter	<p>No actions are required. The Spectra Vail application applies the deletion rule to the current version of objects.</p> <p>Note: If you select the Current Version filter, you cannot select the Non-Current Version filter.</p>
Non-Current Version Filter	<p>Enter a number of non-concurrent versions of an object that should be kept and not expired. When this limit is reached, any excess non-concurrent versions are allowed to expire based on the configured schedule filter.</p> <p>Note: If you select the Non-Current Version filter, you cannot select the Current Version filter.</p>
Include Name Filter	<p>Enter a regular expression. The deletion rule applies to any object with a name that matches the provided expression.</p>

If you selected...	Do the following...
	<p>Note: If multiple Include Name filters are applied to the deletion rule, the deletion rule applies if <u>any</u> of the filters match the object name.</p>
Exclude Name Filter	<p>Enter a regular expression. The deletion rule applies to any object with a name that matches the provided expression.</p> <p>Note: If multiple Exclude Name filters are applied to the deletion rule, the deletion rule applies if <u>any</u> of the filters match the object name.</p>
Include Object Tag Filter	<ul style="list-style-type: none"> • Enter a Key and Value. The deletion rule applies to any object with an matching object tag. The fields are case-sensitive. Object tag filtering is case-sensitive and object tags must be an exact match. • If no Value setting is entered, the deletion rule applies to all objects with an object tag that matches the specified Key. <p>Note: If multiple Include Object Tag filters are applied to the deletion rule, the deletion rule applies if <u>any</u> of the filters match the object name.</p>
Exclude Object Tag Filter	<ul style="list-style-type: none"> • Enter a Key and Value. The deletion rule applies to any object with an matching object tag. The fields are case-sensitive. Object tag filtering is case-sensitive and object tags must be an exact match. • If no Value setting is entered, the deletion rule applies to all objects with an object tag that matches the specified Key. <p>Note: If multiple Exclude Object Tag filters are applied to the deletion rule, the deletion rule applies if <u>any</u> of the filters match the object name.</p>

5. If desired, add additional Placement or Deletion rules.

6. Click **Submit**.

CREATE A VAIL BUCKET

A Vail bucket is a logical target that is shared across the entire Vail sphere. Objects are placed and retrieved from a Vail bucket using an S3 compatible client. Data is then migrated to storage locations using the lifecycle associated with the bucket.

Note: A Vail sphere is limited to 1000 buckets.

Vail buckets can also be linked to an existing bucket on a BlackPearl system or S3 compatible storage endpoint. When buckets are linked, any changes to one bucket are propagated to the other bucket automatically. Only one linked bucket is allowed per storage location.

Note: When a Vail bucket is linked to an S3 cloud bucket, the Spectra Vail application synchronizes the buckets such that changes made on one bucket are propagated to the other bucket. In normal S3 operations, a very small object, such as a 0-length delete marker, is not cloned. However in a linked bucket configuration, small objects created on the linked cloud storage are represented by a clone in the Spectra Vail application because of the bucket synchronization. These clones display in the Vail management console and can be deleted. Deleting the clone of an object results in the object appearing that it was originally created on the Vail storage, not the linked cloud bucket.



IMPORTANT

If a Vail bucket is linked to a S3 cloud bucket, when an object is added to a S3 cloud bucket, the Spectra Vail application creates a version of the object with a clone that references the object in the S3 bucket. Because the objects are linked, if the object is deleted in the Spectra Vail application, the object on the S3 cloud bucket is deleted, even if no lifecycle is defined. If there are multiple versions of the object in the Spectra Vail application, when the object is deleted, only the object on the S3 cloud bucket that matches the version deleted in the Spectra Vail application is deleted from the S3 bucket.

Here is how to create a Vail bucket:

1. In the Vail management console taskbar, click **Buckets**.
2. Under the **Buckets** banner, click **Create**.

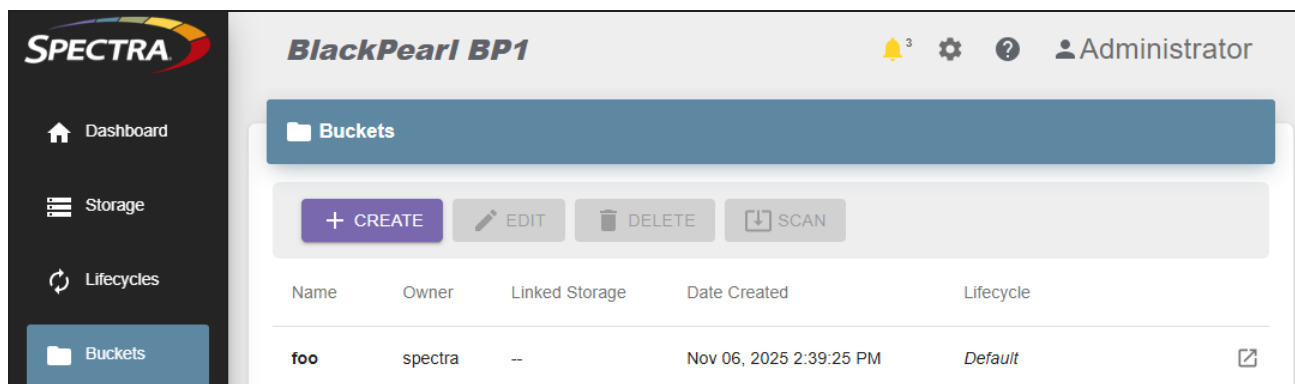


Figure 45 The Buckets screen.

3. Enter the desired **Bucket Name**. Spectra Logic recommends using names that either include the intended usage or use a group name combined with intended usage. If you use a naming convention by groups, the associated group can be easily given access to all buckets sharing the group name prefix.

For example, use usage names such as news-breaking and external-archive, or group and usage name such as eng-dev and eng-test.

Note: Vail bucket names must be between three and 63 characters, using only lowercase letters and numbers. The period (.) and dash (-) characters are valid in the middle of the bucket name, but are not valid as the first or last character of a bucket name.

Note: Spectra Logic recommends following Amazon AWS documentation naming conventions and restrictions to maintain consistency and avoid potential issues. See [AWS Bucket Naming Rules](#) for more information.

Create Bucket

1 Parameters 2 Policy 3 Lifecycle 4 Confirm

Configure your new bucket.

Bucket Name

☐ Enable Versioning ☐ Enable Object Locking ?

☐ Enable Encryption ☒ Enable Compression ☐ Hide Glacier Operations ?

Bucket Owner
vail.development

Object Ownership
ACLs Disabled (recommended) ?

NEXT

Figure 46 The Create Bucket - Parameters screen.

4. If desired, select **Enable Versioning** to allow the bucket to store multiple versions of an object.

5. If desired, select **Enable Object Locking**. This allows you to protect the state of an object when the lock is applied, while also allowing other versions non-locked versions to be modified, and allows new versions of an object to be added to the bucket.

There are two types of locks that can be used. A retention lock expires on a specific date and time. A legal lock must be manually removed.

- Notes:**
- Objects can be locked both when they are added to the bucket, and while they reside in the bucket using the Vail API.
 - Locked objects display a locked in the Vail management console.
 - This option is greyed-out unless you selected to enable versioning in [Step 4](#).

6. If desired, select **Enable Encryption** to encrypt data copied to the Vail bucket.

**IMPORTANT**

Files archived to an encrypted Vail bucket can only be decrypted by the Spectra Vail application.

Note: You must use the key provided by Spectra Logic when transferring data to a Vail bucket configured to use encryption, or data transfers to the bucket fail.

7. If desired, select **Enable Compression** to allow the Spectra Vail application to compact objects placed in the Vail bucket.

Note: Compression is not recommended if your workflow only uses files that are already compressed, such as ZIP files.

8. If desired, select **Hide Glacier Operations**. This option allows S3 clients that do not fully support restoring from AWS S3 Glacier tier storage by automatically requesting the object from Glacier storage when the client requests the object.

Note: Enabling this option changes the response from the Spectra Vail application to the S3 client when an object is not immediately available. Instead of a 403 invalid object state error, a 503 service unavailable error is returned.

**IMPORTANT**

This option is not compatible with S3 clients that fully support Glacier storage restores and may interfere with normal operation.

9. Use the **Bucket Owner** drop-down menu to select a user to own the bucket. The bucket owner sets permissions for the bucket.

10. Use the **Object Ownership** drop-down menu to select the type of ownership used for new objects, and how Access Control Lists (ACLs) are used.

Option	Description
ACLs Disabled	New objects written to this bucket are always owned by the bucket owner configured in Step 9 . <ul style="list-style-type: none">• Access to this bucket and its objects is specified using only policies. Note: This is the recommended setting.
Bucket Owner Preferred	If new objects written to this bucket specify the <i>bucket-owner-full-control</i> canned ACL, the objects are owned by the bucket owner. Otherwise they are owned by the object writer. <ul style="list-style-type: none">• Access to this bucket and its objects can be specified using ACLs or policies.
Object Writer	New objects written to the bucket are always owned by the object writer. <ul style="list-style-type: none">• Access to this bucket and its objects can be specified using either ACLs or policies.

Note: If Object Permissions is configured to use Object Writer, if an object is added to the bucket by a different account, that object is owned by the other account, but permissions for bucket operations are still controlled on the bucket owner

11. Click **Next**.

- If you selected **Enable Object Locking** in [Step 5 on page 102](#), continue with [Step 12](#) below
- Otherwise, skip to [Step 16 on page 105](#).

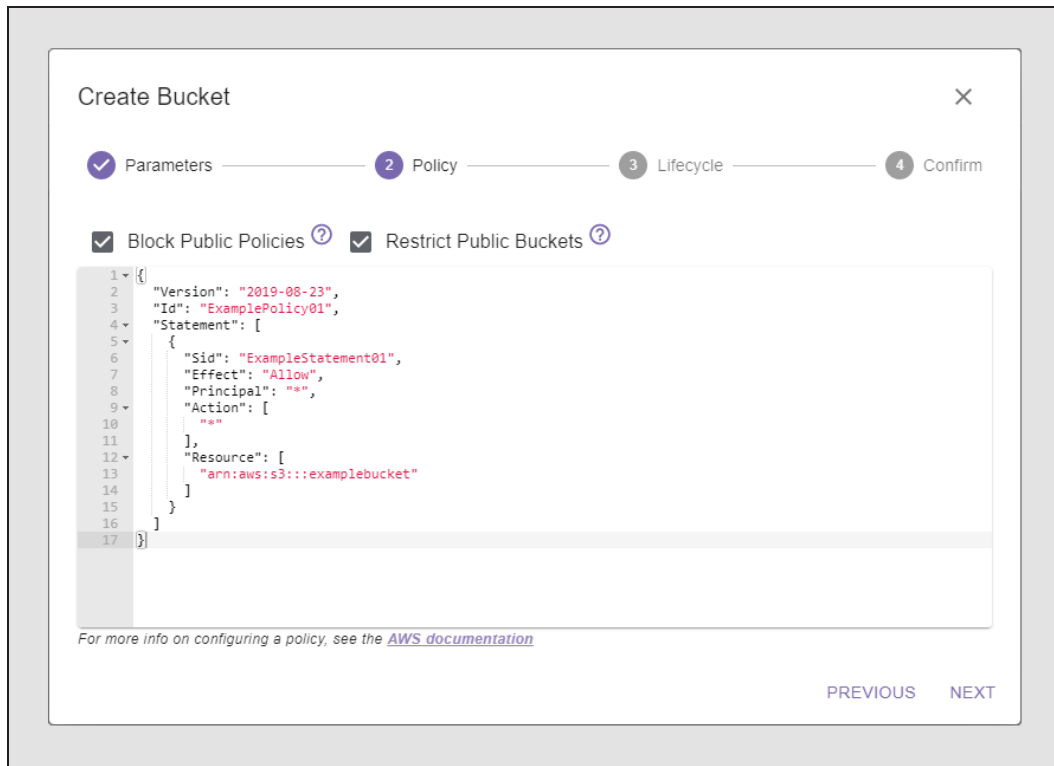
12. If desired, select **Use Default Retention** to configure a retention policy for objects to use if they are not uploaded to the bucket with a specified retention lock. To continue without specifying a default retention policy, click **Next** and skip to **Step 16 on page 105**.

Figure 47 The Create Bucket - Retention screen.

13. Use the **Retention Mode** drop-down menu to select the type of default retention lock. Retention locks have two modes that specify how the lock can be modified. Both Governance and Compliance mode locks can have the retention period extended.
- Retention locks in **Governance** mode can be reduced or removed if the user making the request has the correct permissions.
 - Retention locks in **Compliance** mode can only be extended, and the retention period cannot be removed or reduced. You must wait for the lock to expire.
14. Use the **Unit of Time** drop-down menu to select a unit of time for the default retention lock, then enter a value for **Number of Unit of Time**. The minimum value is 1 day and the maximum value is 36500 days (100 years).
15. Click **Next**.

16. Edit the example **Policy** code as required. Policy permissions are used if you want to exclude IAM user(s) under the main AWS account from accessing the Vail bucket.

Note: For additional information on configuring a policy, see the [Amazon S3 Actions](#) documentation.



Create Bucket

Parameters Policy Lifecycle Confirm

☒ Block Public Policies ☒ Restrict Public Buckets

```

1 {
2   "Version": "2019-08-23",
3   "Id": "ExamplePolicy01",
4   "Statement": [
5     {
6       "Sid": "ExampleStatement01",
7       "Effect": "Allow",
8       "Principal": "*",
9       "Action": [
10        "s3:*"
11      ],
12      "Resource": [
13        "arn:aws:s3:::examplebucket"
14      ]
15    }
16  ]
17 }

```

For more info on configuring a policy, see the [AWS documentation](#)

PREVIOUS NEXT

Figure 48 The Create Bucket - Policy screen.

17. If desired, select or clear **Block Public Policies**. Enabling this setting blocks new bucket policies that grant public access to buckets and objects. This setting does not change existing policies that allow public access.
18. If desired, select or clear **Ignore Public ACLs**. Enabling This setting ignores all ACLs that grant public access to objects or directories.
19. If desired, select or clear **Restrict Public Buckets**. Enabling this setting ignores public and cross-account access for buckets with policies that grant public access to buckets and objects.
20. If desired, select or clear **Block Public ACLs**. Enabling this setting this blocks public access to ACL permissions applied to newly added buckets or objects, and prevents the creation of new public access ACLs for existing buckets and objects. This setting does not change any existing permissions that allow public access to S3 resources using ACLs.

Figure 49 The Create Bucket - Access Control List screen.

21. Click Next.

- If you selected **ACLs Disabled** in [Step 10](#) on page 103, skip to [Step 26](#) on page 107.
- Otherwise continue with [Step 22](#) below.

22. Click Add ACL to configure ACL bucket permissions. ACL permissions are used when the bucket is shared across AWS accounts, and when older applications are being used that are not compatible with bucket policies.

23. Using the AWS Canonical ID drop-down menu, select an ID.

Note: The ID of the Vail sphere administrator is automatically configured in the Spectra Vail application. To add additional AWS accounts, see [Configure & Manage IAM Accounts](#).

24. Using the Permissions check boxes, set the permissions for the Vail bucket. If desired, you can assign multiple permissions.

Option	Description
Read	Allows the user to list the objects in a bucket.
Read ACP	Allows the user to read the bucket ACL information.
Write	Allows the user to create new objects in the bucket, and to overwrite existing objects.
Write ACP	Allows the user to write the ACL for the bucket.

If necessary, repeat [Step 22](#) through [Step 24](#) to configure additional ACLs.

Note: Use the trashcan icon to remove an ACL.

25. Click Next.

26. Using the **Select Lifecycle** drop-down menu, select a previously configured lifecycle and click **Next**.

Note: All buckets must be assigned a lifecycle. The Use Default option uses the lifecycle configured as the bucket owner's default lifecycle. If that option is not set, the global lifecycle default is used.



The screenshot shows a 'Create Bucket' dialog box with a title bar containing a question mark and a close button. Inside the dialog, there is a progress bar with four steps: 'Parameters' (checked), 'Policy' (checked), 'Lifecycle' (active, highlighted with a blue circle and the number 3), and 'Confirm' (disabled, greyed out with the number 4). Below the progress bar, there is a section titled 'Select Lifecycle' with a dropdown menu showing 'Use default'. At the bottom right of the dialog, there are two buttons: 'PREVIOUS' and 'NEXT'.

Figure 50 The Create Bucket - Lifecycle screen.

27. Review the configuration, then click **Submit** to create the bucket.

Note: A Vail sphere is limited to 1000 buckets.

CONFIGURE AN OBJECT STORAGE BROWSER

Before you can access and transfer data to a BlackPearl Nearline Object Gateway or Spectra Vail VM node, you must configure an object storage browser. The instructions in this section describe how to configure the S3 Browser and Cyberduck® cloud storage browser software.

Note: For other object browser programs compatible with the Spectra Vail application, refer to the documentation included with the software.

The instructions below assume you have previously installed the browser software.

Configure S3 Browser

Here is how to configure the S3 Browser:

1. Launch the S3 Browser software.

Note: You must use S3 Browser program version 9.0.8 or later.

2. Click **Accounts > Add New Account**.
3. Enter the desired **Account Name**.

Figure 51 The Add New Account wizard.

4. Using the **Account Type** drop-down menu, select **S3 Compatible Storage**.

5. Enter the IPv4 address of the BlackPearl Nearline Object Gateway or Spectra Vail VM node as the **REST Endpoint**.
6. Enter the **Access Key ID** and the **Secret Access Key** of an IAM user configured in the Spectra Vail application.
7. Clear the **Use secure transfer (SSL/TLS)** check box.
8. Click **Advanced S3-compatible storage settings**.




Figure 52 The Advanced S3-Compatible Storage Settings screen.

9. Using the **Signature Version** drop-down menu, select **Signature V4** and click **Close**.
Note: The Vail application also support AWS V4a signatures.
10. Click **Add new account**. The S3 Browser retrieves the list of buckets configured on the Spectra Vail sphere (see [View Vail Bucket Details](#) on page 156.)

Configure Cyberduck Object Storage Browser

Here is how to configure Cyberduck object storage browser:

1. Download and install the Cyberduck profile for third party S3 (HTTPS) connections.

The profile can be downloaded at:

[https://profiles.cyberduck.io/Spectra%20S3%20\(HTTPS\).cyberduckprofile](https://profiles.cyberduck.io/Spectra%20S3%20(HTTPS).cyberduckprofile)

Note: Use the Cyberduck user documentation for help installing the profile.

2. Launch the Cyberduck software.
3. Click **Open Connection**.

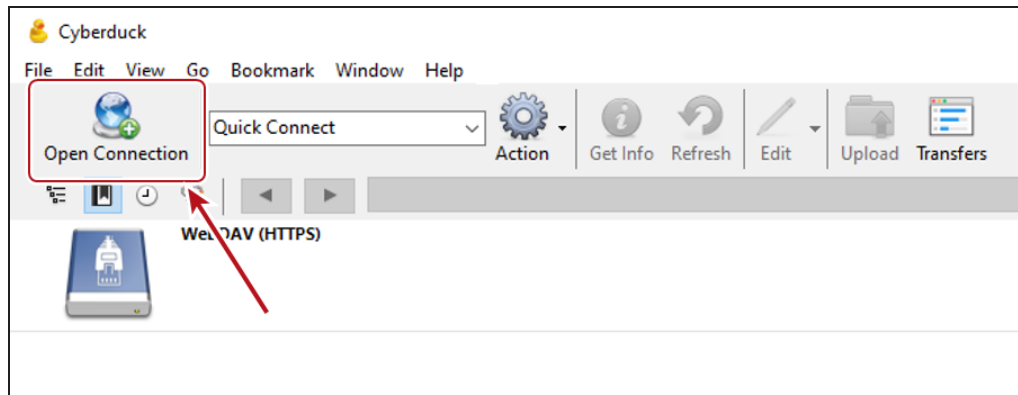


Figure 53 The Cyberduck Object Storage Browser home screen.

4. Using the drop-down menu, select **S3(HTTPS)**.
5. Using the **Server** entry field, enter the IP address of the BlackPearl Nearline Object Gateway or Spectra Vail VM node.

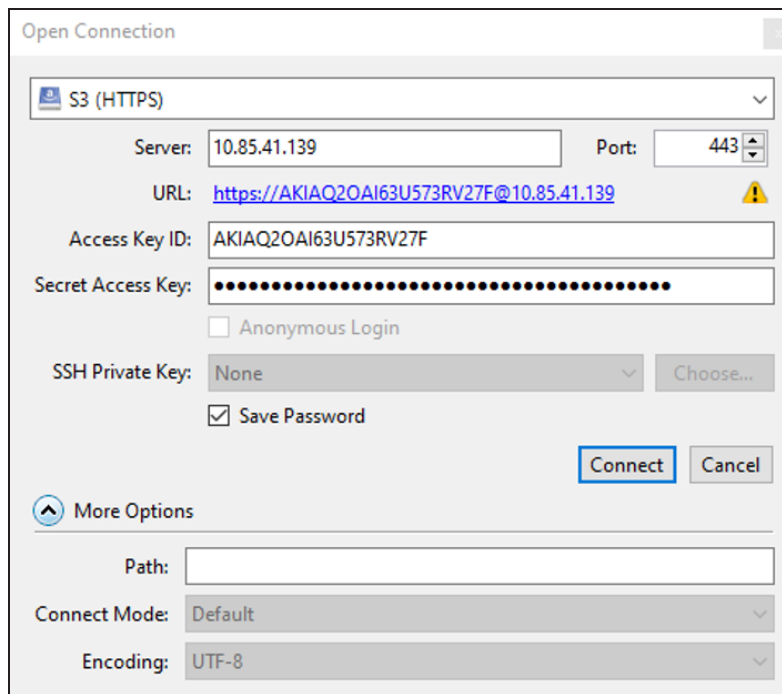


Figure 54 The Open Connection screen.

6. Enter the **Access Key ID** and the **Secret Access Key** of an IAM user configured in the Spectra Vail application.
7. Click **Connect**.

CHAPTER 5 - CONFIGURE & MANAGE USERS

This chapter describes the configuration and managing user accounts in the Spectra Logic Spectra Vail application. This chapter includes information about Vail sphere administrator accounts, IAM accounts, and IAM groups, as well as AWS access key management.

Configure & Manage Sphere Administrator - Cloud Control	113
Create a Sphere Administrator	113
Change a Sphere Administrator Password	115
Edit Sphere Administrator Attributes	117
Delete a Sphere Administrator	119
Configure & Manage Vail Administrator - Local Control	120
Create a Vail Administrator	120
Change a Vail Administrator Password	123
Delete a Vail Administrator	123
Configure & Manage IAM Accounts	124
Add an IAM Account	124
View IAM Account Details	130
Edit an IAM Account	132
Delete an IAM Account Association	133
Configure & Manage IAM Users and Groups	134
Create an IAM User	134
View IAM User Details	135
Add an IAM User to an IAM Group	136
Remove an IAM User from an IAM Group	137
Delete an IAM User	138
Create an IAM Group	139
Delete an IAM Group	140
Create an IAM Group Policy	141
Edit an IAM Group Policy	142
Delete an IAM Group Policy	143
AWS Access Key Management	144
Create an Access Key	144
Enable an Access Key	145
Disable an Access Key	146

Delete an Access Key	147
----------------------------	-----

CONFIGURE & MANAGE SPHERE ADMINISTRATOR - CLOUD CONTROL

Spectra Vail application sphere administrator accounts have full control over the entire sphere, with full access to configure and change all system settings. Use the information in this section to create, edit, or delete a sphere administrator when using a cloud controlled Vail application.

Note: The Spectra Vail application relies on the AWS Cognito server to manage sphere administrators. As a result, it is also possible to make sphere administrator level changes via the AWS management console.

Create a Sphere Administrator

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **Administrators (2)**.




Figure 55 The Dashboard screen - Navigation menu.

2. In the Sphere Administrator pane, click **Create**.

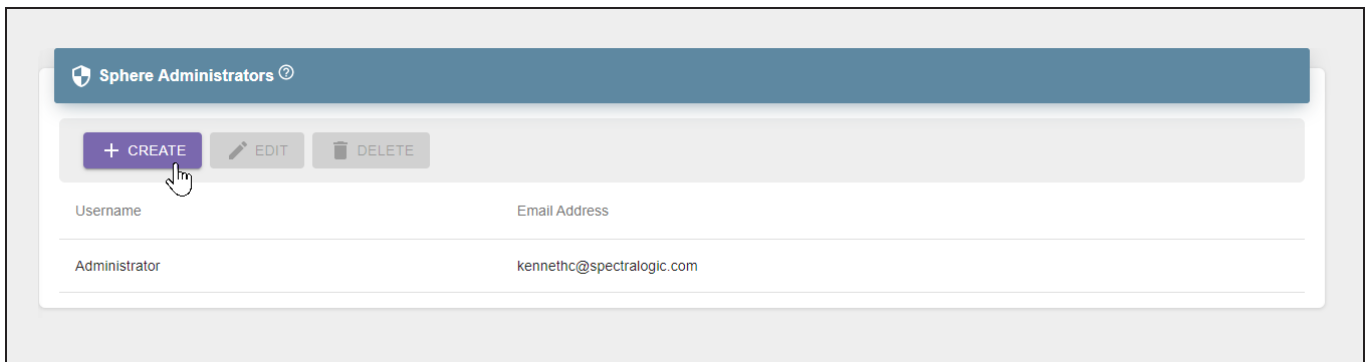


Figure 56 The Sphere Administrators pane.

3. Enter the desired **Username**.

Spectra Logic suggests using the same naming convention as your corporate email for Vail sphere administrator names.

For example, if associate Jane Smith uses the email address `janes@yourcompany.com`, use "janes" for the user name.

The screenshot shows a 'Create Sphere Administrator' dialog box. It has a title bar with a question mark icon and a close 'X' button. The form contains two input fields: 'Username' and 'Email Address'. Below the 'Email Address' field, there is a note: 'A temporary password will be sent to this address'. Underneath, there is a section titled 'Select what types of emails this user wants to receive.' with four checkboxes: 'Info', 'Ok', 'Warning', and 'Error'. A 'SUBMIT' button is located at the bottom right of the dialog.

Figure 57 The Create Sphere Administrator screen.

4. Enter the **Email Address** for the sphere administrator. Emails sent to this address include system events and the temporary password for the account.

5. Select the type(s) of emails that the sphere administrator receives. The Spectra Vail application emails the administrator when an event of the selected type occurs.

Setting	Description
Info	An expected event occurred such as a job starting or completing successfully.
OK	A component of the Vail sphere reports an OK status.
Warning	Notifies the user of a failure that may adversely impact the Spectra Vail application.
Error	Notifies the user of a failure that caused significant adverse impact to the Spectra Vail application.

6. Click **Submit**.

A default password is emailed to the address entered in [Step 4](#)

Change a Sphere Administrator Password

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **Administrators (2)**.

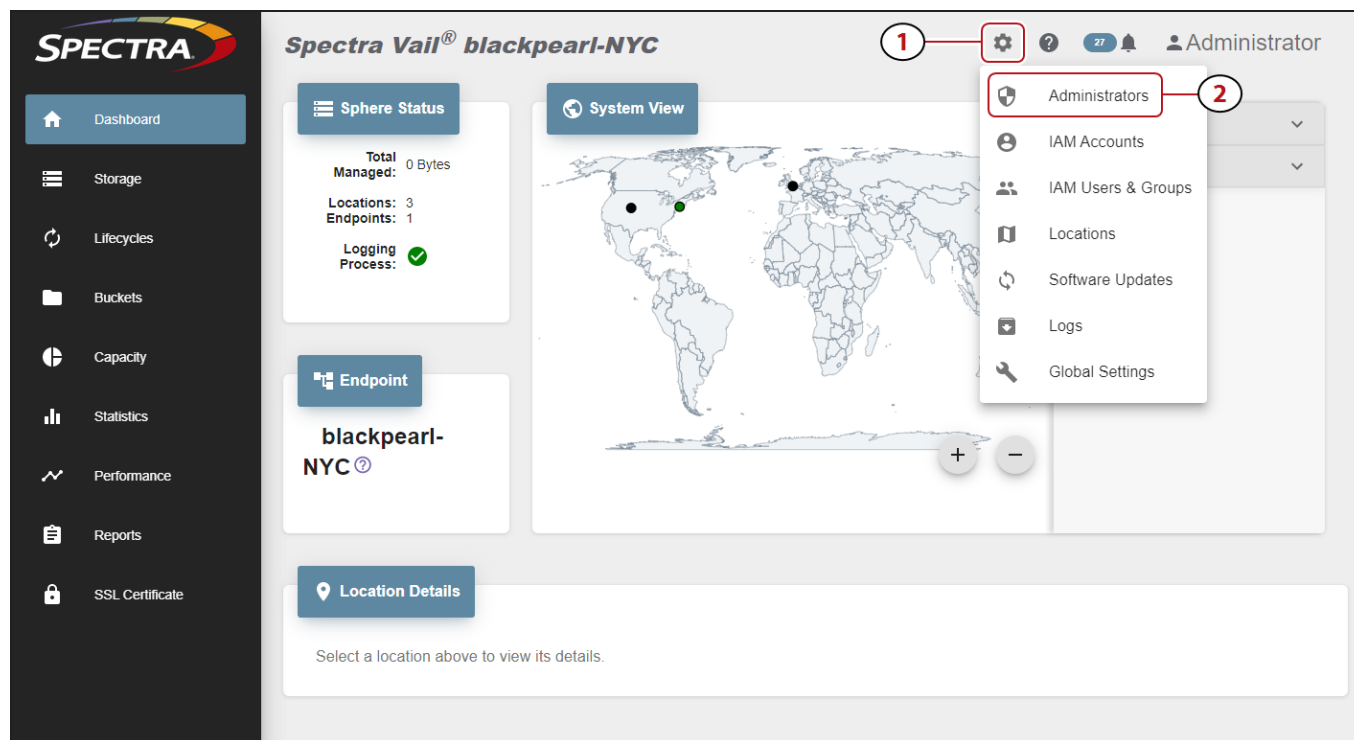


Figure 58 The Dashboard screen - Navigation menu.

2. Under the **Sphere Administrator** banner, (1) Select the row of the sphere administrator to edit, and (2) click **Edit**.




Figure 59 The Sphere Administrators pane.

3. Select **Set new password** and click **Next**.

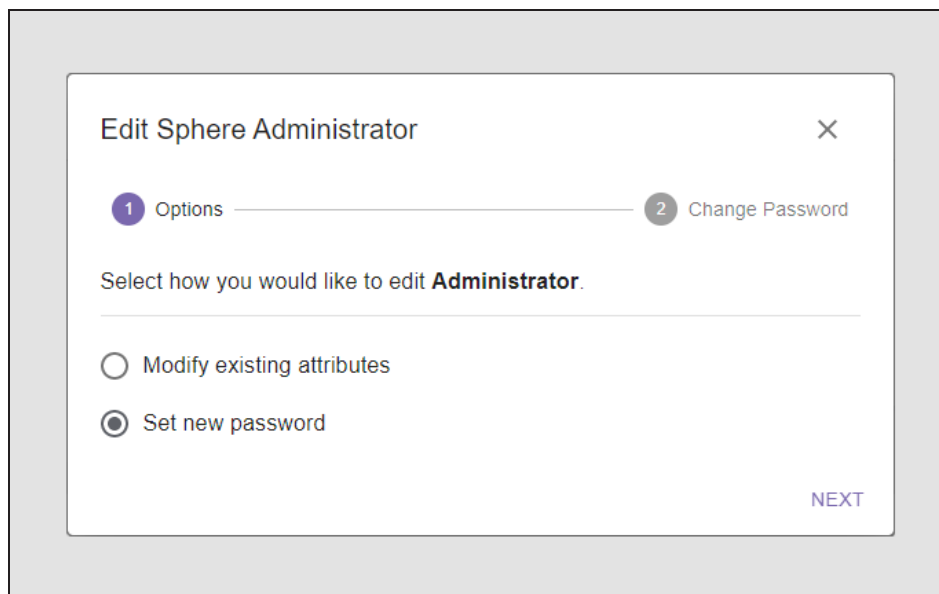


Figure 60 The Edit Sphere Administrator - Options screen.

4. Enter the desired **New Password**, then **Confirm New Password** and click **Submit**.

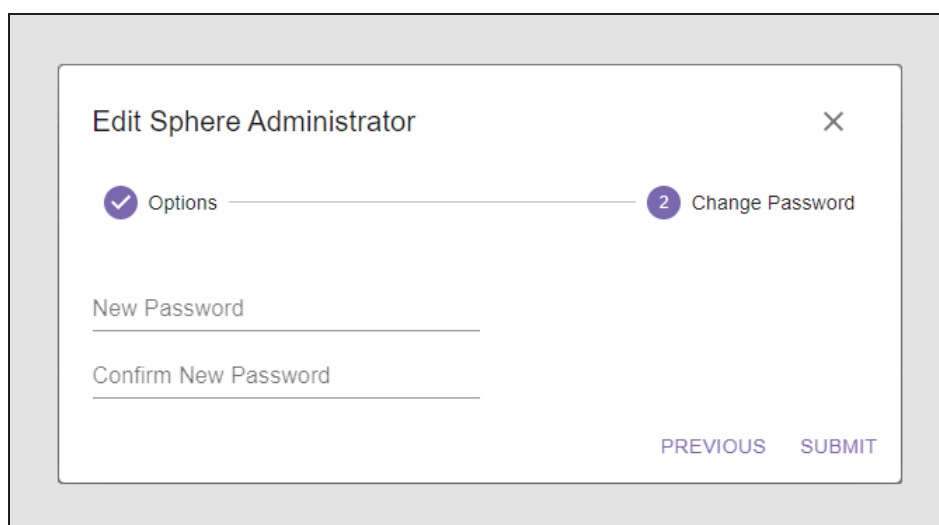


Figure 61 The Edit Sphere Administrator - Change Password screen.

Edit Sphere Administrator Attributes

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **Administrators (2)**.




Figure 62 The Dashboard screen - Navigation menu.

2. Under the **Sphere Administrator** banner, (1) Select the row of the sphere administrator to edit, and (2) click **Edit**.

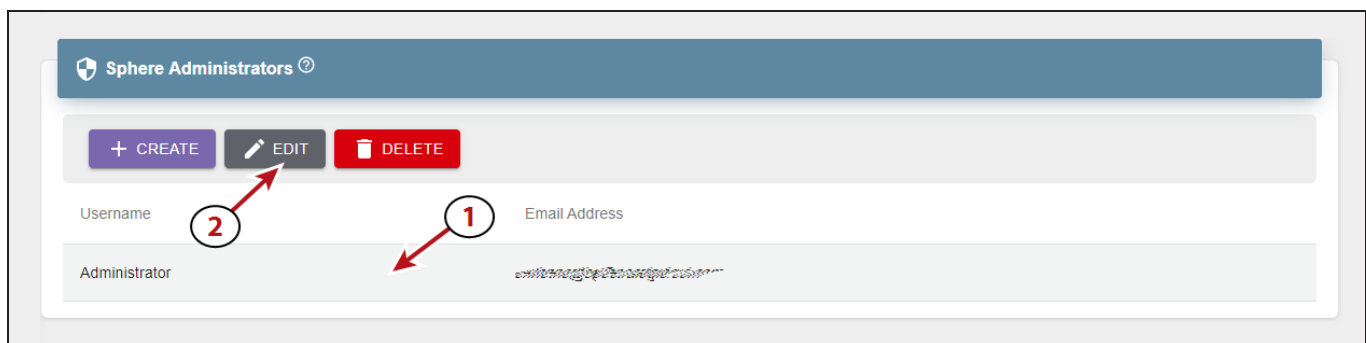
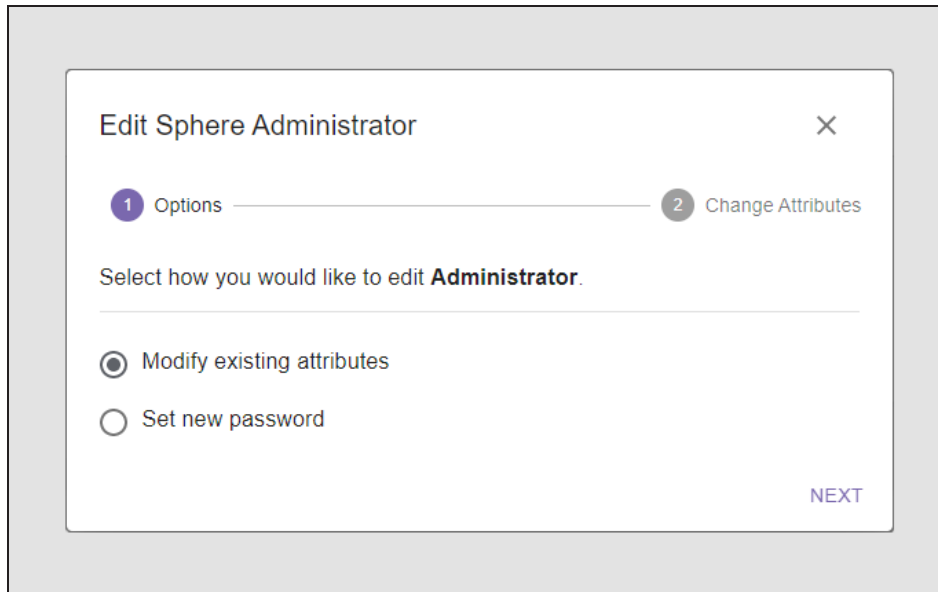


Figure 63 The Sphere Administrators pane.

3. Select **Modify existing attributes** and click **Next**.



Edit Sphere Administrator

1 Options 2 Change Attributes

Select how you would like to edit **Administrator**.

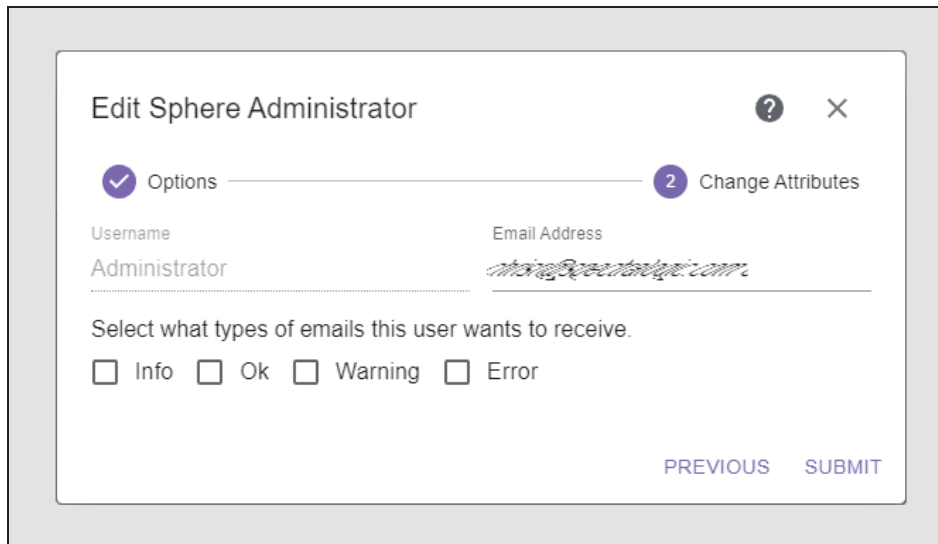
☒ Modify existing attributes

☐ Set new password

NEXT

Figure 64 The Edit Sphere Administrator - Options screen.

4. Change the **Email Address** or the types of email the sphere administrator receives, and click **Submit**. See [Step 5](#) for a description of email types.



Edit Sphere Administrator

1 Options 2 Change Attributes

Username Email Address

Administrator Administrator@spectravail.com

Select what types of emails this user wants to receive.

☐ Info ☐ Ok ☐ Warning ☐ Error

PREVIOUS SUBMIT

Figure 65 The Edit Sphere Administrator - Change Attributes screen.

Delete a Sphere Administrator

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **Administrators (2)**.

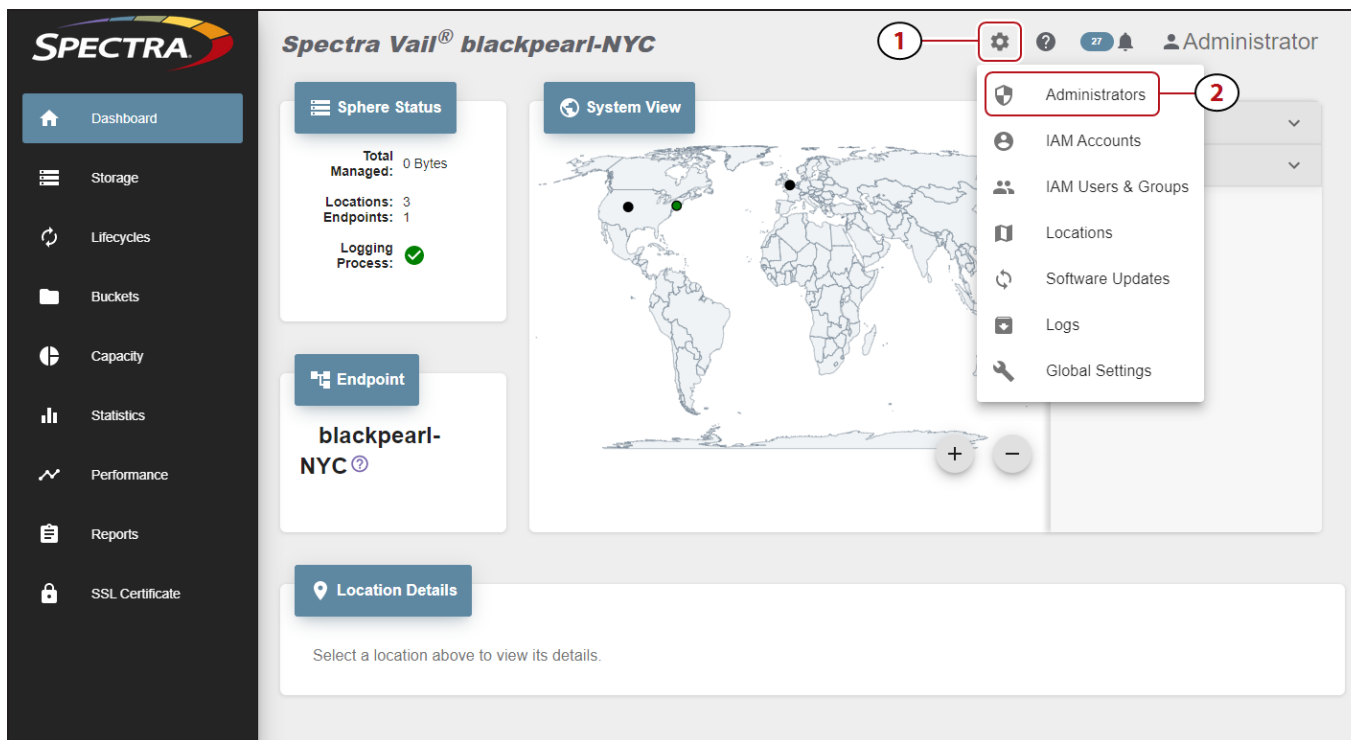


Figure 66 The Dashboard screen - Navigation menu.

2. Under the **Sphere Administrator** banner, (1) Select the row of the sphere administrator to delete, and (2) click **Delete**.

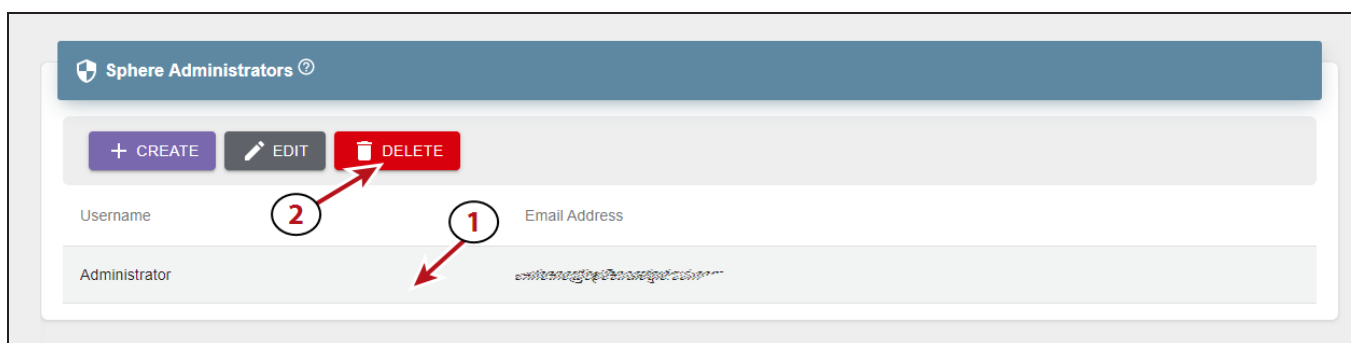


Figure 67 The Sphere Administrators pane.

3. Click **Delete** to permanently delete the sphere administrator.

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

New User

Username

Full Name

Password

Confirm Password

Session Timeout

User Access

☐ Administrator ☐ Monitor ☐ Login ☐ CIFS

S3 User Settings

Default Data Policy

Max Buckets

Global Bucket Access Control List

☐ List ☐ Read ☐ Write ☐ Delete ☐ Job ☐ Owner

Global Data Policy Access Control List

☐ Enabled

Figure 69 The New User dialog box.

3. Enter the desired **Username** for the user. The Username cannot contain capital letters or spaces and is limited to 16 characters. The Username is used to identify the user in the Vail environment.
4. Enter the user's **Full Name**.
5. Enter and confirm the desired **Password** for the user.
6. If desired, enter the value for the **Session Timeout** in minutes. This value cannot exceed 999 minutes.
7. Select Administrator and Login **User Access** permissions.
8. From the drop-down list, select a **Default Data Policy** for the user. If specified, the BlackPearl Nearline Object Gateway uses the selected data policy for all buckets created by the user, unless a different policy is specified during bucket creation.
9. Enter a value for the **Max Buckets** the user is allowed to create. The default value of 10000 is pre-entered.

- 10.** Select or clear options for the **Global Bucket Access Control List**. These options give or deny permission for the user being created to perform the action described in the table below, for all buckets present on the BlackPearl Nearline Object Gateway, as well as any buckets created at a future date.

Name	Description
List	The user can see the bucket and can list the objects in a bucket.
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 created by other users. The user can also see the details of jobs created by other users. Note: 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 all buckets, including all permissions listed above.

- 11.** If desired, under **Global Data Policy Access Control List**, select the check box to allow the user access to any data policy created on the BlackPearl Nearline Object Gateway.
- 12.** Click **Create** to create the new user. The BlackPearl Nearline Object Gateway generates a unique S3 Access ID and Secret Key for the user.

Change a Vail Administrator Password

1. From the right side of the menu bar, select **Current User > User Profile**. The User Profile screen displays.
2. Select **Action > Edit**. The Edit User Screen displays.

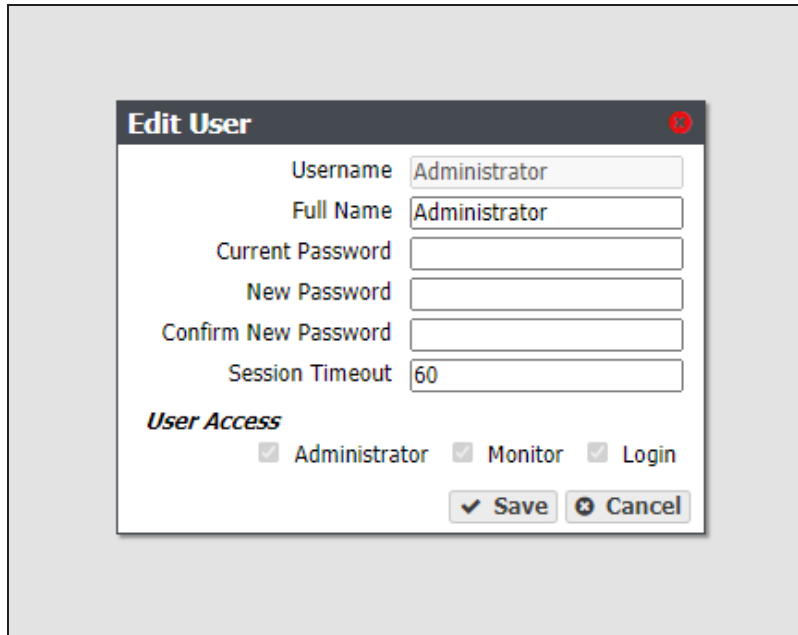


Figure 70 The Edit User dialog box.

3. If desired, edit the user's **Full Name**.
4. If you are changing the password, enter the desired **New Password**, then **Confirm New Password**.

Note: The new password does not take effect until after you log out of the BlackPearl user interface.

5. If desired, edit the value for the **Session Timeout** in minutes. This value cannot exceed 999 minutes.
6. Click **Save**.

Delete a Vail Administrator

1. From the menu bar, select **Configuration > Users**. The Users screen displays a list of all configured users and S3 groups (see [Figure 1](#)).
2. Select the user you want to delete, and then select **Action > Delete**. A confirmation window displays.
3. Click **Delete** to delete the user.

CONFIGURE & MANAGE IAM ACCOUNTS

Identity and Access Management (IAM) allows you to control access to resources by assigning permissions to users and groups that allow or deny access to a resource.

Note: When using IAM accounts, Spectra Logic recommends you carefully consider the security requirements associated with IAM accounts and IAM policies. See the following for more information.

<https://aws.amazon.com/blogs/security/category/security-identity-compliance/aws-identity-and-access-management-iam/>

Add an IAM Account

By default, an IAM account is created when the Spectra Vail application is configured and associated with the sphere. If you have additional IAM accounts and want the Spectra Vail application to access resources associated with other accounts, you can add them as IAM accounts in the Spectra Vail application.

There are two types of IAM accounts, AWS and Local. Use the sections below to add an IAM account:

- **Add an AWS IAM Account below**
- **Add a Local IAM Account on page 127**

Add an AWS IAM Account

Use the section below to add an AWS IAM account.

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Accounts (2)**.

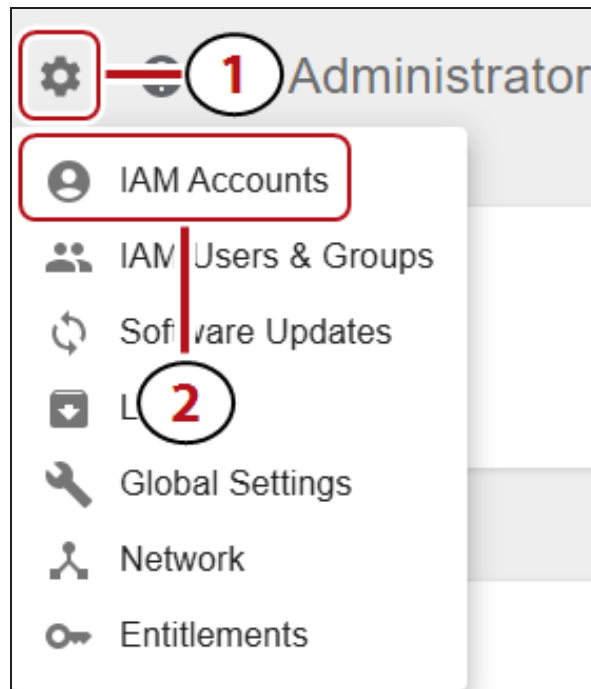


Figure 71 The Dashboard screen - Navigation menu.

2. Under the **IAM Accounts** banner, click **Add**.




Figure 72 The IAM Accounts pane.

3. Select **AWS** if necessary.

Figure 73 The Add IAM Account screen.

4. If desired, enter a **Description** for the IAM account.
5. Enter the **Role ARN**. The Role ARN is an IAM role that specifies what a user is allowed to do and is used by a user in one AWS account to assume a role in a different AWS account. The Role ARN can be found in the Role page of the AWS account to be added to the Vail sphere.

You must specify the AWS resource using the following format:

arn:partition:service:region:account:resource

Parameter	Description
partition	Identifies the partition containing the resource. For standard AWS regions, the partition is aws . For resources in other partitions, use aws-partitionname .
service	Identifies the AWS product. When configuring an AWS user in the Spectra Vail application, use the service name is iam .

Parameter	Description
region	This parameter is not used when configuring an AWS user in the Spectra Vail application and must be left blank.
account	<p>The full AWS account ID for the AWS account with no hyphens. This can be found on the My Account screen in the AWS management console.</p> <p>Note: You cannot use an AWS account ID alias when configuring an AWS user in the Spectra Vail application.</p>
resource	The name of the specific resource.

6. If desired, enter an **External ID**. The external ID is associated with the IAM role entered in [Step 5](#) and is configured when a role is created in an AWS account. The External ID is required to assume the role created in [Step 5](#). In the AWS management interface, the External ID can be found on the Roles section of the IAM screen, in the **Trusted relationships** tab.
 7. Enter the **Email** address of the owner of the AWS account. This email address can be found on the AWS Dashboard and is listed as the **Management Account Email Address**.
 8. If desired, set **Status** to Inactive. Inactive accounts cannot access the Vail sphere and all access keys for this account are marked as inactive.
 9. If desired, set a **Connection Limit** for the IAM account. This setting limits the number of connections that can be used by the IAM account as a percentage of total system connections.
- Note:** An empty value is treated as 100%.
10. Use the drop-down menu to select a **Default Lifecycle**.
 11. Click **Submit**.

Add a Local IAM Account

Use the section below to add a local IAM account.

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Accounts (2)**.




Figure 74 The Dashboard screen - Navigation menu.

2. Under the **IAM Accounts** banner, click **Add**.

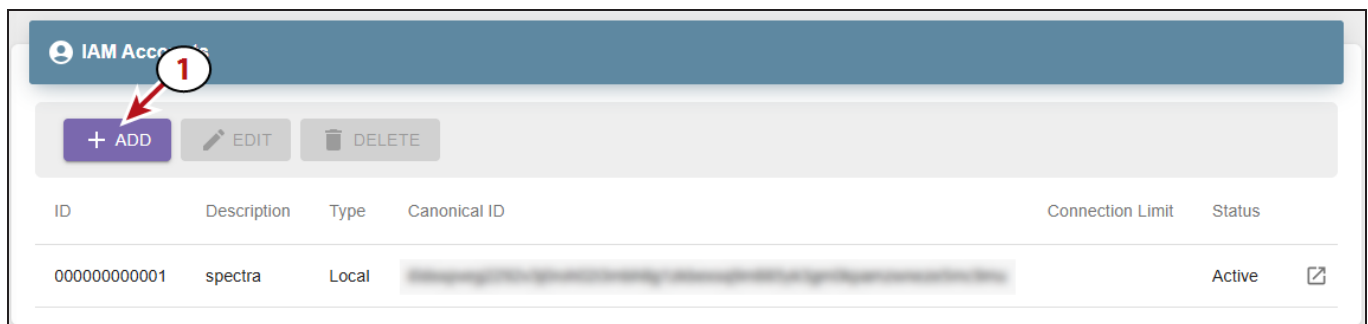


Figure 75 The IAM Accounts pane.

3. Select **Local**.

Add IAM Account ? X

☐ AWS ☒ Local

Description _____

Email _____

Status ? ☒ Active ☐ Inactive

Connection Limit _____ % ?

Default Lifecycle
Use global default ▼

SUBMIT

Figure 76 The Add IAM Account screen.

4. If desired, enter a **Description** for the IAM account.
5. Enter the **Email** address of the owner of the account.
6. If desired, set **Status** to Inactive. Inactive accounts cannot access the Vail sphere.
7. If desired, set a **Connection Limit** for the IAM account. This setting limits the number of connections that can be used by the IAM account as a percentage of total system connections.

Note: An empty value is treated as 100%.

8. Use the drop-down menu to select a **Default Lifecycle**.
9. Click **Submit**.

View IAM Account Details

After creating an IAM account, you can look at the details of the IAM account.

Here is how to view IAM account details:

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Accounts (2)**.




Figure 77 The Dashboard screen - Navigation menu.

2. Under the **IAM Accounts** banner, (1) select the row of the IAM account to view, and (2) click the **View Details** symbol.



Figure 78 The IAM Accounts pane.

3. The Properties window displays the configured settings for the selected IAM account.

spectra ×	
PROPERTIES	
ID	000000000001
Description	spectra
Email	spectra@spectralogic.com
Type	Local
Canonical ID	i0dxxpveg2292v3j0roh02i3mbh8g1zkbexs q9m885yk3gm0kpamzwneze5mc9mu
Connection Limit	
Default Lifecycle	<i>Global Default</i>
Status	Active

Figure 79 The IAM Accounts Properties window.

Edit an IAM Account

When editing an IAM account, only the email address and description can be changed.

Here is how to edit an IAM account:

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Accounts (2)**.

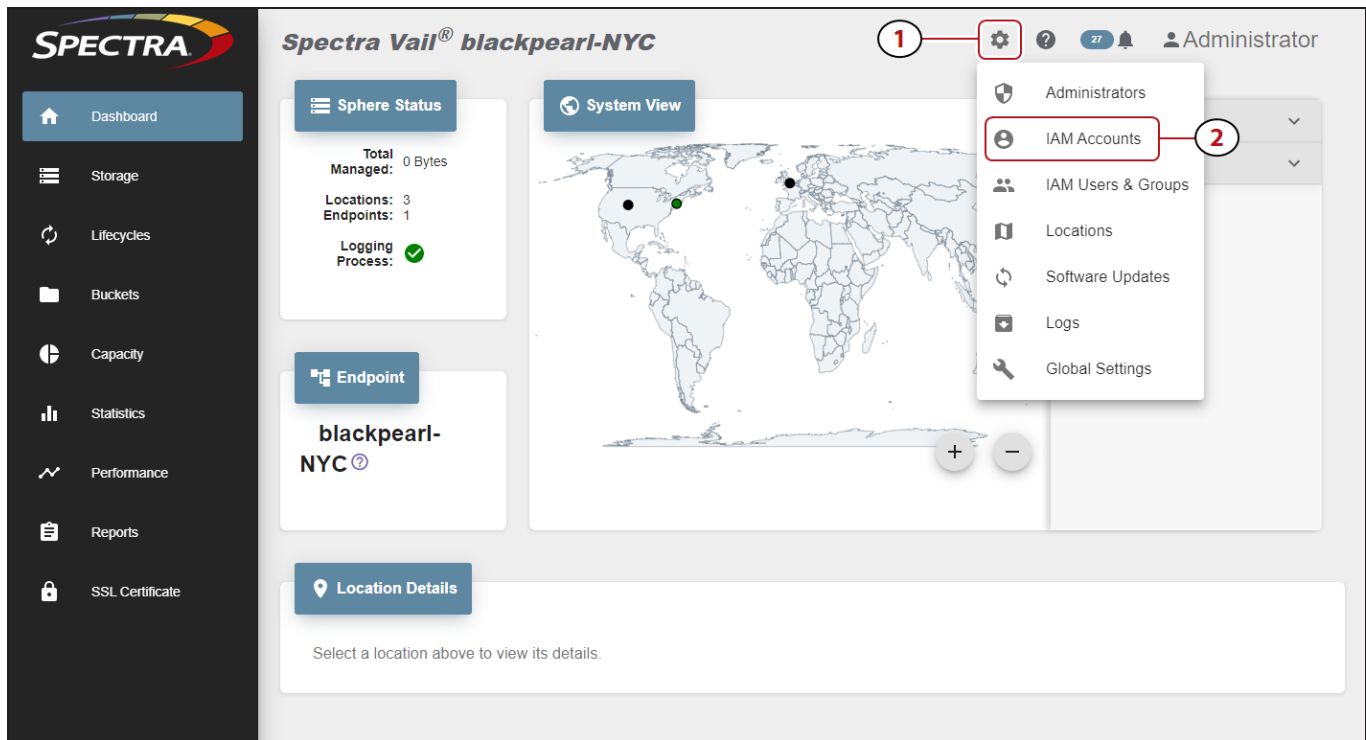


Figure 80 The Dashboard screen - Navigation menu. Editing Note - Leaving this image alone if you don't like the zoomed in look of the above.

2. Under the **IAM Accounts** banner, (1) select the row of the IAM account to edit, and (2) click **Edit**.



Figure 81 The IAM Accounts pane.

3. Change the **Email** address, **Description**, and **Status** as desired and click **Submit**.

Delete an IAM Account Association

If desired, you can delete an IAM account that is associated with the Vail sphere. You cannot delete an account association if that IAM account is being used by the Vail sphere.

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Accounts (2)**.




Figure 82 The Dashboard screen - Navigation menu.

2. Under the **IAM Accounts** banner, (1) select the row of the IAM account to delete, and (2) click **Delete**.




Figure 83 The IAM Accounts pane.

3. Click **Delete** to delete the IAM account association with the Spectra Vail application.

Note: The IAM account itself is not deleted.

CONFIGURE & MANAGE IAM USERS AND GROUPS

Create an IAM User

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Users & Groups (2)**.




Figure 84 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, click **Create**.
3. Enter the **Username** for the new IAM user. The username cannot exceed 64 characters.

Figure 85 The Create IAM User screen.

4. Click **Submit**.

Note: The username is converted to use all lower-case letters.

View IAM User Details

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Users & Groups (2)**.

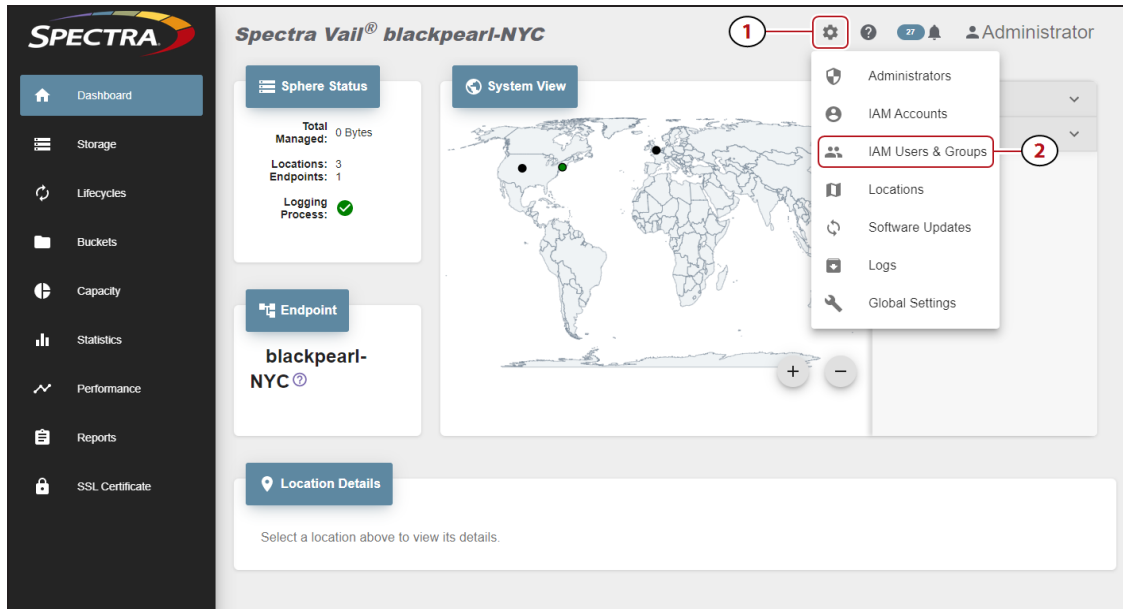


Figure 86 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, locate the row of the user for which you want to view details and click the **View Details** icon on the right end of the row.




Figure 87 The IAM Users banner - View Details button.

3. The IAM user details screen displays showing the **Properties**, **IAM Groups**, and **Access Keys** for the user.

Add an IAM User to an IAM Group

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Users & Groups (2)**.




Figure 88 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, locate the row of the user for which you want to add to an IAM group, and click the **View Details** icon on the right end of the row.




Figure 89 The IAM Users banner - View Details button.

3. Click **IAM Groups**.
4. Use the **Select Group** drop-down menu to select a previously created IAM group.
5. Click **Submit** to confirm adding the user to the IAM group.

Remove an IAM User from an IAM Group

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Users & Groups (2)**.




Figure 90 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, locate the row of the user for which you want to remove from an IAM group, and click the **View Details** icon on the right end of the row.




Figure 91 The IAM Users banner - View Details button.

3. Click **IAM Groups**.
4. **Select** the row of the group, then click **Remove**.

Delete an IAM User

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Users & Groups (2)**.




Figure 92 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, select the row of the user to delete, and click **Delete**.
3. Click **Delete** to confirm deleting the IAM user.

Note: When an IAM user is deleted, the AWS access key assigned to the user is also deleted.

Create an IAM Group

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Users & Groups (2)**.




Figure 93 The Dashboard screen - Navigation menu.

2. Under the **IAM Groups** banner, click **Create**.
3. Enter the **Name** for the new IAM Group.
4. Click **Submit**.

Delete an IAM Group

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and select **IAM Users & Groups (2)**.




Figure 94 The Dashboard screen - Navigation menu.

2. Under the **IAM Groups** banner, (1) select the row of the group to delete, and (2) click **Delete**.
3. Click **Delete** to confirm deleting the IAM group.

Create an IAM Group Policy

1. On the IAM Users and Groups screen and under the IAM Groups banner, click **Show Details** on the desired IAM group.




Figure 95 The IAM Users and Groups screen - IAM Groups.

2. On the IAM group details window, select the **Policies** tab at the top, then click **Create**.
3. Enter a **Policy Name** and enter the desired policy.
4. Click **Submit**.

Edit an IAM Group Policy

1. On the IAM Users and Groups screen and under the IAM Groups banner, click **Show Details** on the desired IAM group.

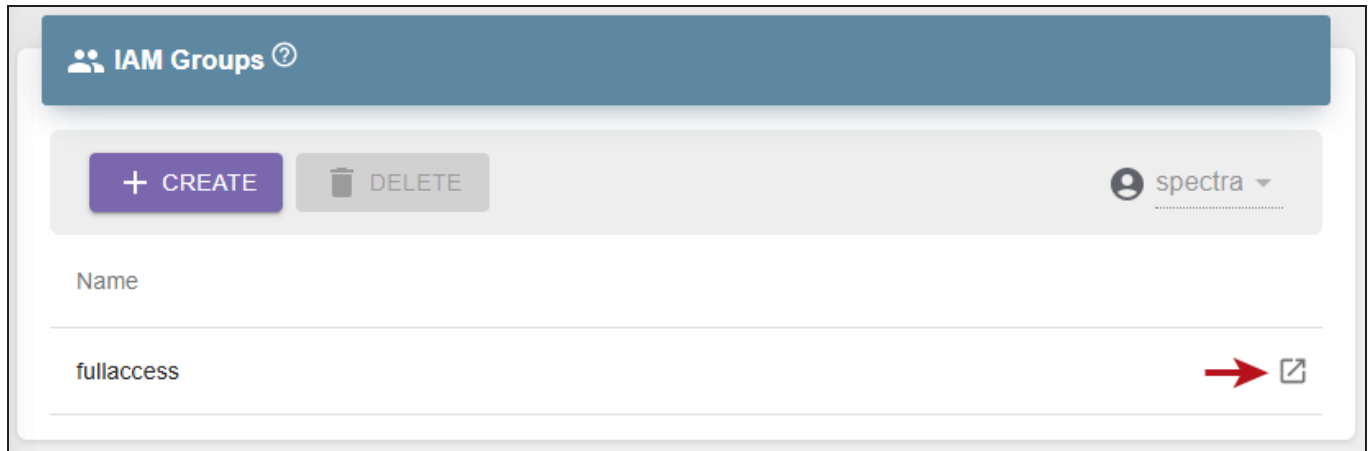


Figure 96 The IAM Users and Groups screen - IAM Groups.

2. On the IAM group details window, select the **Policies** tab at the top.
3. Select the desired policy, then click **Edit**.
4. Edit the policy as required, then click **Submit**.

Delete an IAM Group Policy

1. On the IAM Users and Groups screen and under the IAM Groups banner, click **Show Details** on the desired IAM group.



Figure 97 The IAM Users and Groups screen - IAM Groups.

2. On the IAM group details window, select the **Policies** tab at the top.
3. Select the desired policy, then click **Delete**.

AWS ACCESS KEY MANAGEMENT

Create an Access Key

If desired, you can create a new AWS access key for use by an IAM user.

1. In the upper right corner of the management console, click the **gear icon** and select **IAM Users & Groups**.

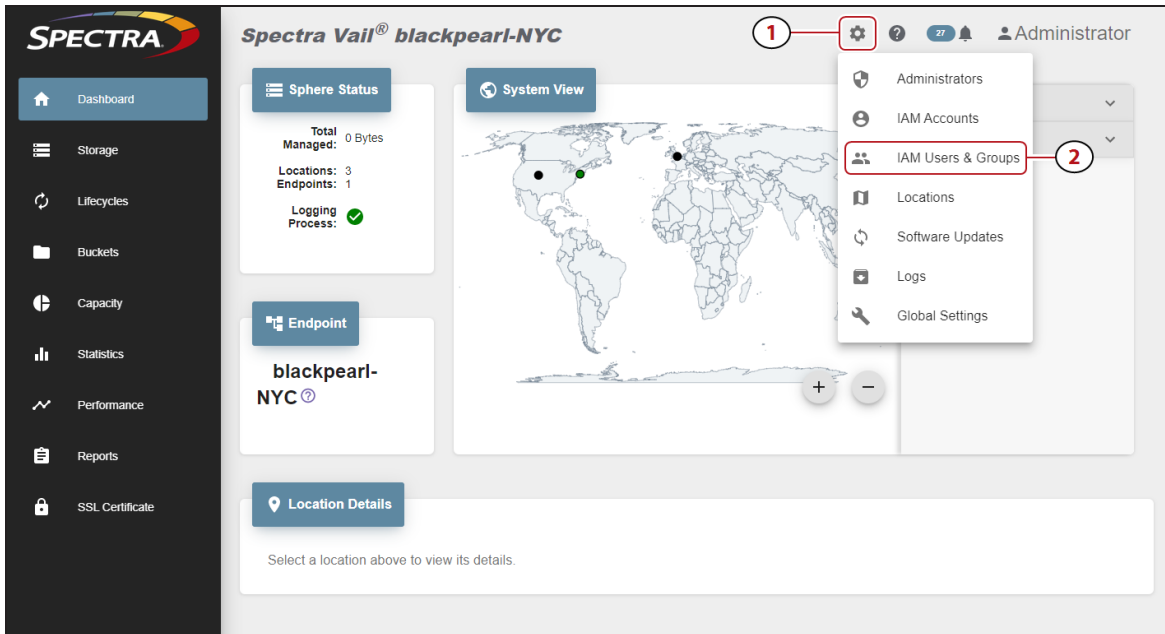


Figure 98 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, locate the row of the user for which you want to create an access key and click the **View Details** icon on the right end of the row.

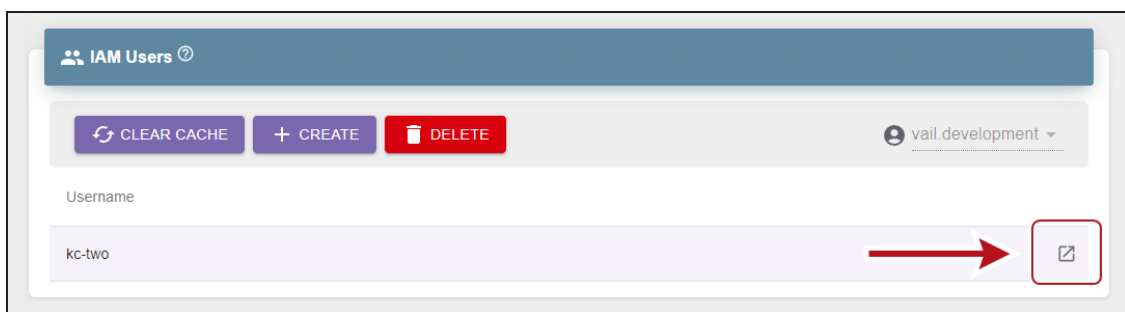


Figure 99 The IAM Users banner - View Details button.

3. Select **Access Keys**.
4. Click **Create**. The new access key displays in the list.

Enable an Access Key

If desired, you can enable a previously disabled AWS access key.

Note: New key(s) created through the Vail management console are automatically enabled.

1. In the upper right corner of the management console, click the **gear icon** and select **IAM Users & Groups**.

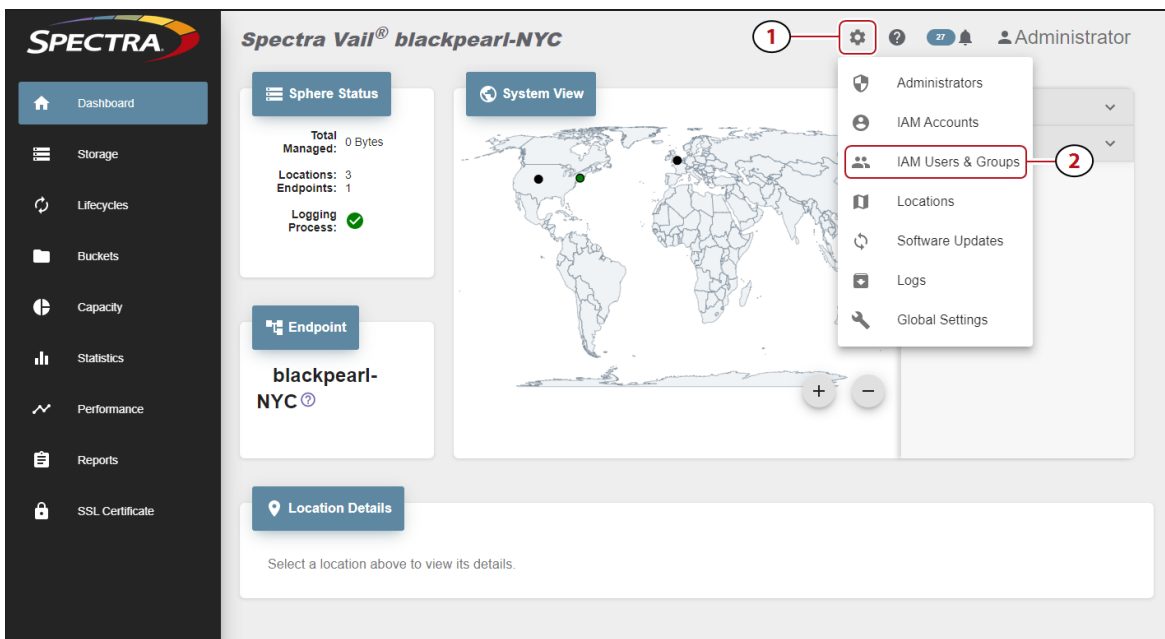


Figure 100 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, locate the row of the user for which you want to enable an access key and click the **View Details** icon on the right end of the row.

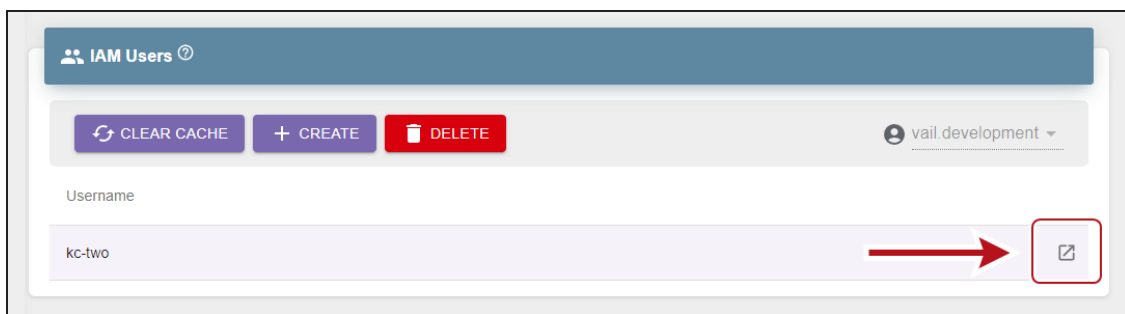


Figure 101 The IAM Users banner - View Details button.

3. Select **Access Keys**.
4. Select the row of the access key you want to enable and click **Enable**.
5. On the confirmation screen, click **Enable**.

Disable an Access Key

If desired, you can disable an access key. The access key is no longer able to be used with the Spectra Vail application, and is also disabled in the user's AWS account.

Note: The AWS access key can be re-enabled at a later date.

Here is how to disable a user access key:

1. In the upper right corner of the management console, click the **gear icon** and select **IAM Users & Groups**.




Figure 102 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, locate the row of the user for which you want to disable an access key and click the **View Details** icon on the right end of the row.




Figure 103 The IAM Users banner - View Details button.

3. Select **Access Keys**.




Figure 104 The User Properties - Access Keys screen.

4. Select the row of the key you want to disable and click **Disable**.
5. On the confirmation screen, click **Disable**.

Delete an Access Key

If desired, you can delete an AWS access key for an IAM user. This is helpful if the AWS access key credentials are compromised, or if required by your company security policy.

Here is how to delete an AWS access key:

1. In the upper right corner of the management console, click the **gear icon** and select **IAM Users & Groups**.

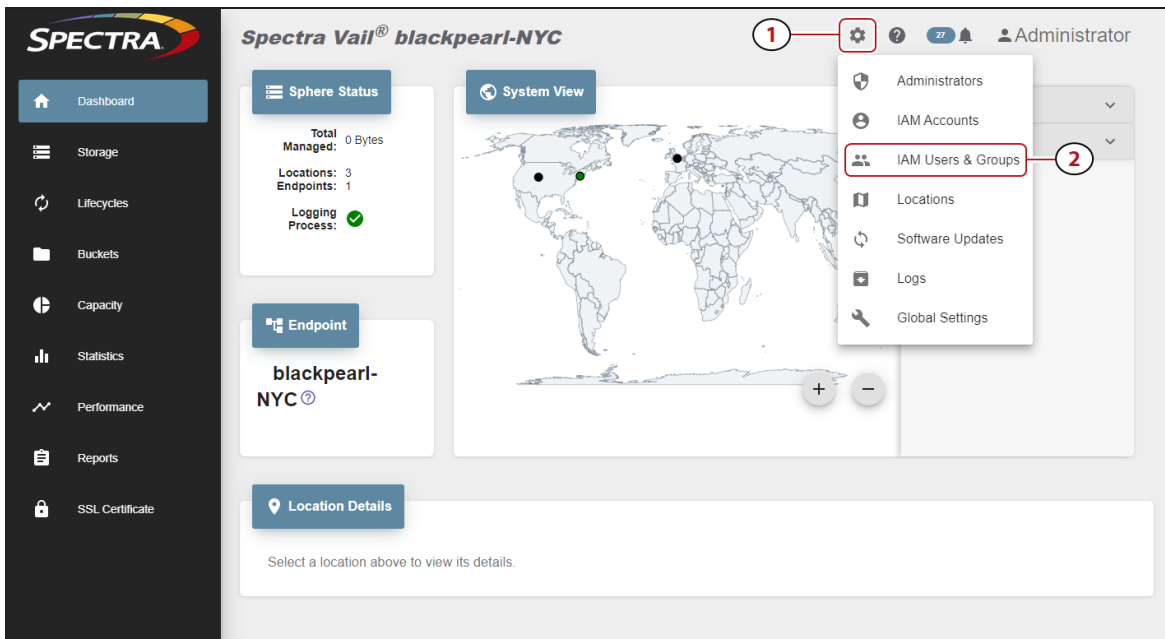


Figure 105 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, locate the row of the user for which you want to delete an access key and click the **View Details** icon on the right end of the row.




Figure 106 The IAM Users banner - View Details button.

3. Select **Access Keys**.
4. Select the row of the key you want to delete and click **Delete**.
5. Click **Delete** to confirm deleting the access key. The key is deleted from the IAM user account in the Spectra Vail application, and deleted from the associated AWS account.

CHAPTER 6 - USING THE SPECTRA VAIL APPLICATION

This chapter describes using the Spectra Vail application.

View Capacity Information	151
View Performance Metrics	154
View Vail Bucket Details	156
View Vail Bucket Contents	160
View Object Details	162
Create an Object Clone	166
Verify an Object Clone	169
Delete an Object Clone	171
Edit Global Settings	173
Change Lifecycle Rule Nightly Processing Time	173
Set Default Lifecycle	174
Enable Diagnostic Monitor	174
Configure AWS Infrastructure	176
Using Proxy Connections	177
Configure Proxy Connection	177
Edit Proxy Server	178
Delete Proxy Server	178
Edit a Vail Bucket	179
Delete a Vail Bucket	183
View Storage Details	184
Edit BlackPearl or Vail VM Endpoint	187
Change Endpoint Location	187
Add Additional Host Names	189
Change Endpoint URL	190
Configure Debug Logging	191
Edit Storage	193
Edit BlackPearl Bucket Storage	193
Edit BlackPearl Volume Pool Storage	196

Edit Vail VM Node Storage	198
Edit Google Cloud Platform Storage	201
Edit AWS S3 Cloud Storage	204
Edit Microsoft Azure Cloud Storage	208
Edit Other S3 Cloud Storage	212
Consolidate Storage	216
Delete Storage	217
View Lifecycle Details	221
Edit a Lifecycle	225
Delete a Lifecycle	228
Create a Location	229
Delete a Location	232
Clear the IAM Cache	233
View Reports	234
View Spectra Vail Application Messages	236
Message Details	238
Spectra Vail Application Logs	239
Update the Spectra Vail Application Software	240
Accessing the Technical Support Portal	243
Create an Account	243
Log Into the Portal	244
Opening a Support Ticket	245
Search for Help Online	245
Submit an Incident Online	248
Submit an Incident by Phone	250

VIEW CAPACITY INFORMATION

The Capacity page allows you to see data capacity information for the Spectra Vail sphere endpoints, each configured location, and cloud storage.

Note: Capacity values for BlackPearl storage display zeros until data is written to the storage.

In the Vail management console taskbar, click **Capacity**.

The Capacity screen is separated into three sections:

- The **Sphere Endpoint Physical Capacity** pane displays the combined total of all configured BlackPearl, Vail VM node, and cloud storage endpoints.

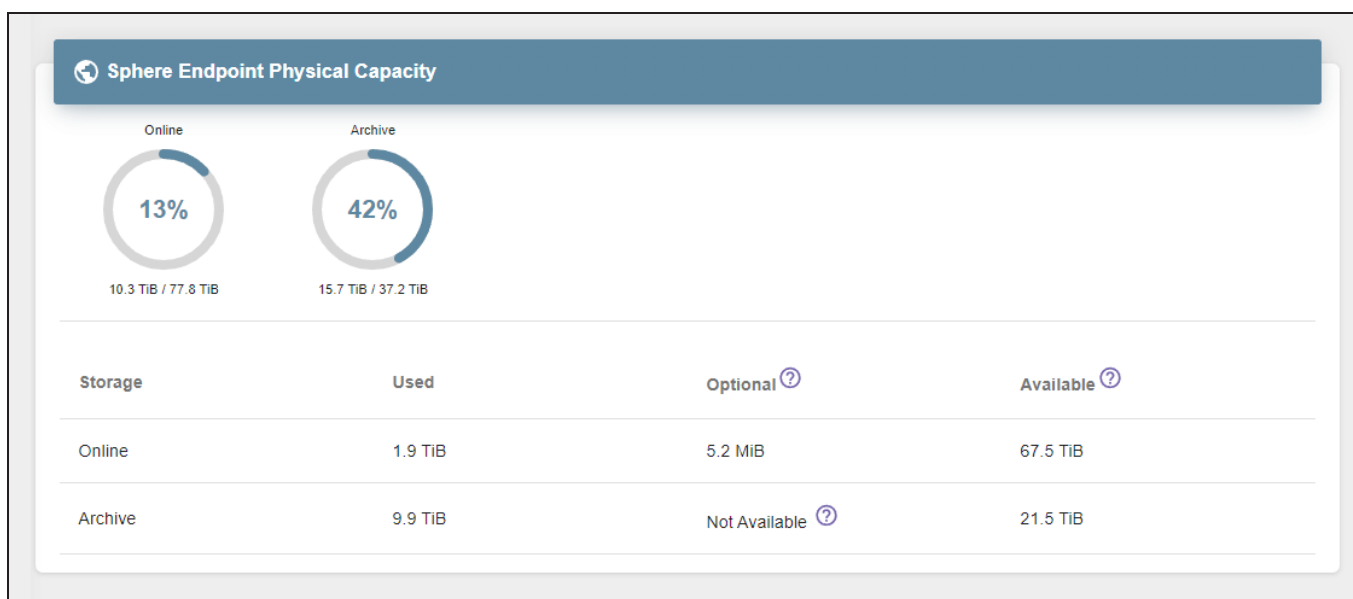


Figure 107 The Sphere Endpoint Physical Capacity pane.

Field	Description
Storage	The type of storage.
Used	The amount of space used for each storage type.
Optional	The amount of space used by the optional clones. There is a delay before this field is populated after creating storage.
Available	<p>The available space used for each storage type.</p> <p>Note: Available capacity does not account for capacity used by file system overhead.</p>

- The **Location Capacity** pane displays data capacity information for each configured location. Buttons in the top left of the pane allow you to view information for each location.

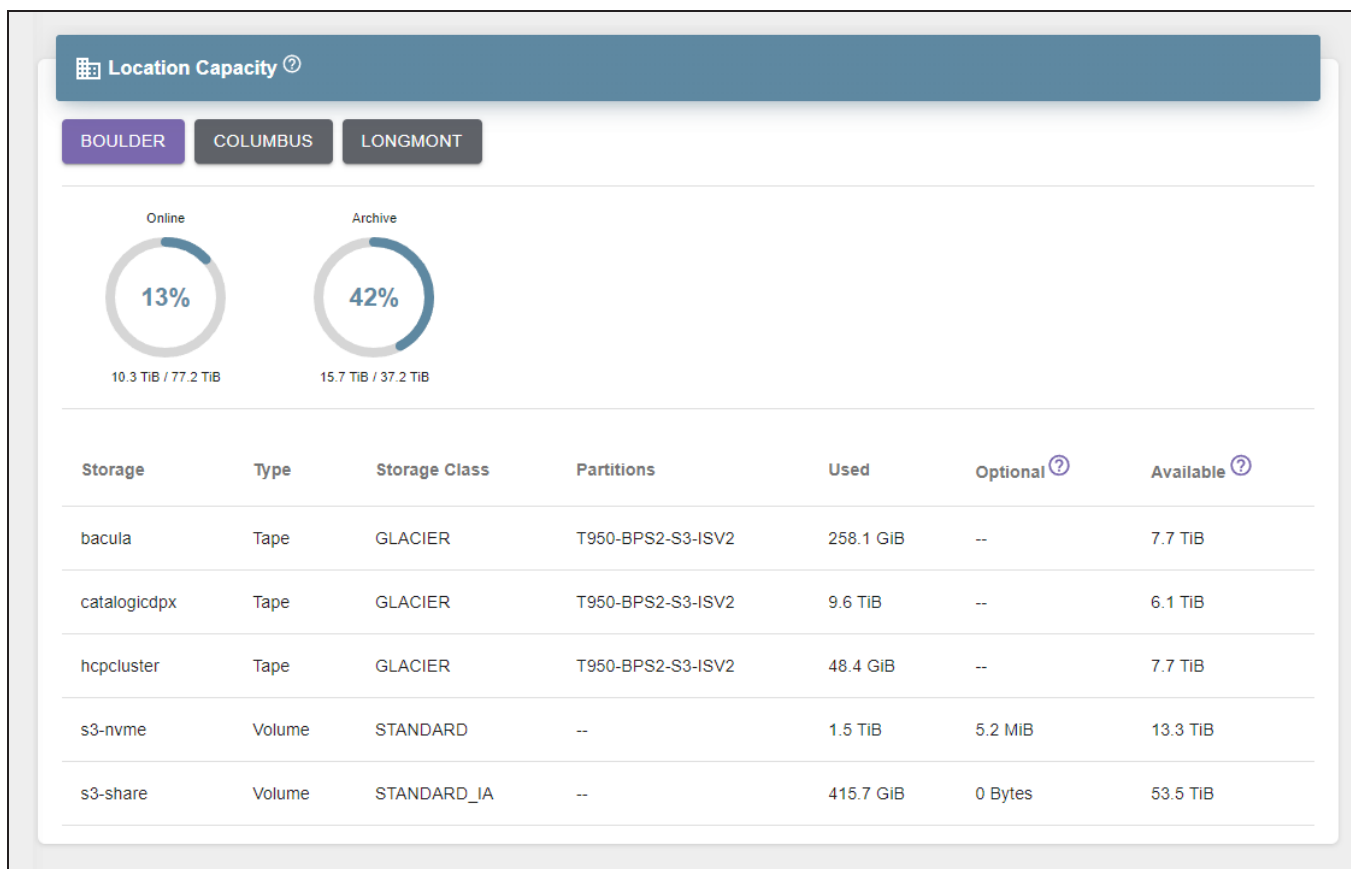


Figure 108 The Location Capacity pane.

Field	Description
Storage	The name of the location.
Type	The type of storage used for each location. Tape - Storage on tape media on a BlackPearl system. Volume - Storage on disk volume storage on a BlackPearl system.
Storage Class	The storage class used by the storage location.
Partitions	The BlackPearl data partition(s) that are used for storage.
Used	The amount of space used for each location.
Optional	The amount of space used for optional object clones.
Available	The available space used for each location. BlackPearl storage is over-provisioned, and may be used by multiple storage endpoints. Note: Available capacity does not account for capacity used by file system overhead.

- The **Cloud Capacity** pane displays aggregated data capacity information for each type of storage class used by cloud endpoints.




Figure 109 The Cloud Capacity pane.

VIEW PERFORMANCE METRICS

The Performance page displays data transfer and operation performance for the Vail sphere and all configured endpoints. The performance graphs display information in five minute or one day intervals.

In the Vail management console taskbar, click **Performance**.

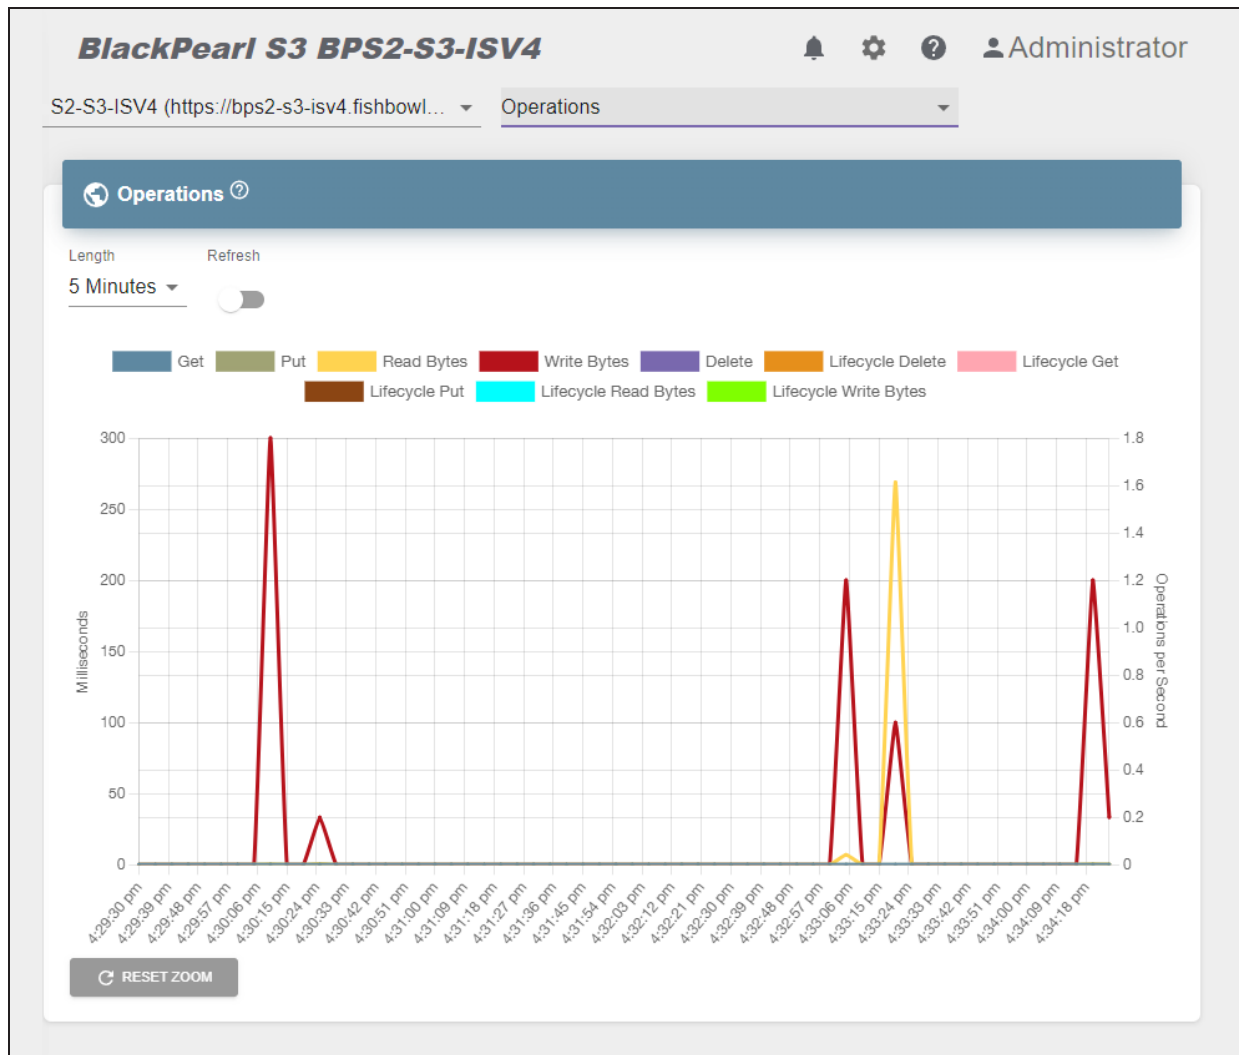


Figure 110 The Performance screen.

- Use the **Endpoint** drop-down menu to select an endpoint for any graph on the Performance screen.
- Use the **Graph Type** drop-down menu to select which graph to display.
- Use the **Length** drop-down menu to select between intervals of five minutes or one day.
- **Click the name** of a graph category to hide or reveal that line in the graph.
- Toggle the **Refresh** slider to refresh the display.

- To display the exact time and performance information, **mouseover** any point on a graph.

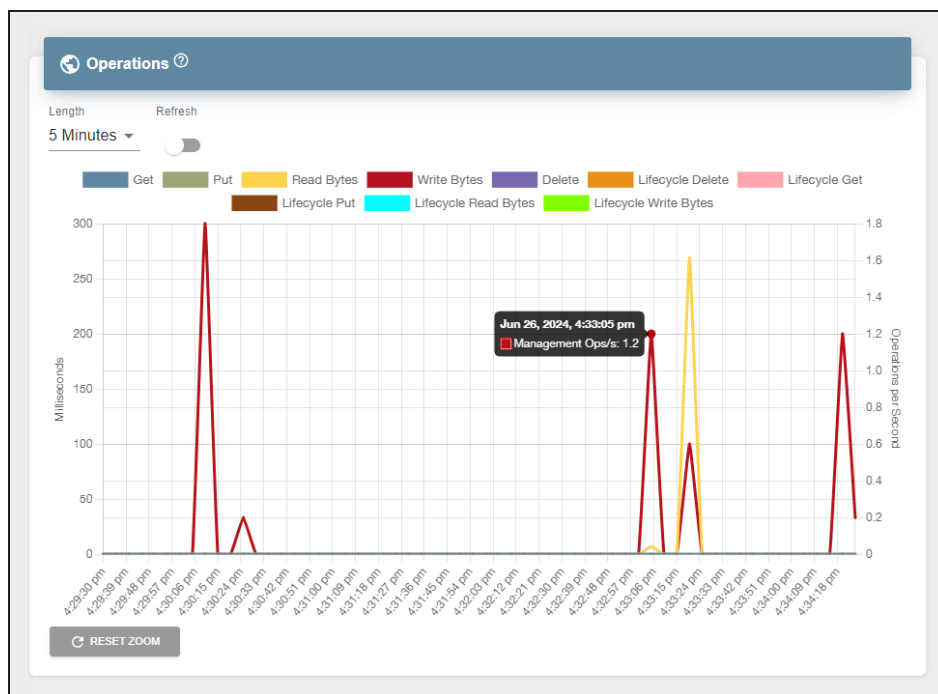


Figure 111 The Operations graph - mouseover.

VIEW VAIL BUCKET DETAILS

The buckets detail screen displays information about the selected Vail bucket, including bucket properties, ACLs, and policy.

Here is how to view the details of a Vail bucket:

1. In the Vail management console taskbar, click **Buckets**.
2. Under the **Buckets** banner, select a bucket row, then click the **View Details** icon on the right side of the pane.

Note: If you click the bucket name instead of the bucket row, the Bucket Contents pane displays. See [View Vail Bucket Contents](#) on page 160.

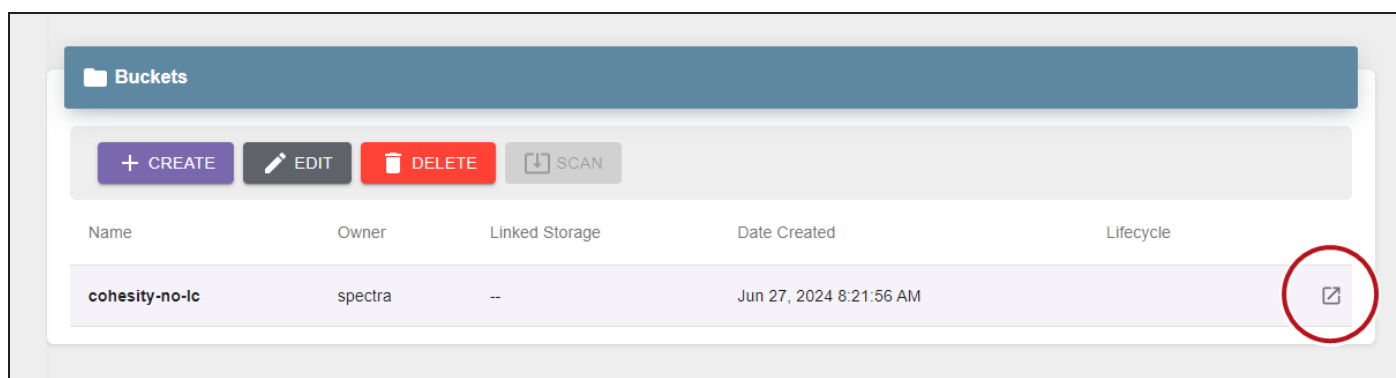


Figure 112 The Buckets pane.

3. Click **Properties**, **Usage**, **ACLs**, or **Policy** to view the current Vail bucket settings.




Figure 113 The Bucket Details - Properties screen.

- If you click **Properties...**

Field	Description
Owner	The AWS Canonical ID of the Vail bucket owner. By default the Vail sphere administrator is the bucket owner.
Linked Storage	The name of the bucket on the BlackPearl system or AWS cloud storage location to which the Vail bucket is linked, if applicable.
Versioning	Indicates if versioning is enabled or disabled for the Vail bucket.
Object Locking	Indicates if object locking is enabled or disabled for the Vail bucket.
Encryption	Indicates if encryption is enabled or disabled for the Vail bucket
Compression	Indicates if compression is enabled or disabled for the Vail bucket.
Hide Glacier Operations	Indicates if hiding glacier operations is enabled or disabled for the Vail bucket.
Object Ownership	Indicates the type of object ownership configured for the bucket
Lifecycle	The lifecycle associated with the Vail bucket.
Creation Date	The date the Vail bucket was created.
Default Retention Mode	Indicates if default retention mode is enabled or disabled for the Vail bucket
Default Retention Period	The retention time period configured for the bucket.

- If you click **Usage...**

Field	Description
Number of Objects	The number of objects currently in the bucket.
Total Size of Objects	The current size of all objects in the bucket, in GiB.
Average Object Size	The current average size of the objects in the bucket, in GiB.

- If you click **ACLs....**

Field	Description
Block Public ACLs	Indicates if the Vail bucket blocks public ACLs.
Ignore Public ACLs	Indicates if the Vail bucket allows public ACLs.
AWS Canonical ID	The ID of a users configured with ACL permissions for the Vail bucket.
Permissions	The ACL permission level for the user.

- If you click **Policy...**

Field	Description
Block Public Policy	Indicates if the Vail bucket blocks or allows public policies.
Restrict Public Buckets	Indicates if the Vail bucket blocks or allows public buckets.
Policy	The AWS policy information entered when the bucket was created displays.

4. Click the **X** in the upper-right corner to close the window.

VIEW VAIL BUCKET CONTENTS

The buckets contents screen displays all objects in a Vail bucket. If versioning is enabled for the bucket, other versions of the current object can also be viewed.

Here is how to view the contents of a Vail bucket:

1. In the Vail management console taskbar, click **Buckets**.




Figure 114 The Buckets screen.

2. Under the **Buckets** banner, click a **bucket name**.

Note: You must click the name directly. Clicking the row of the bucket does not display the bucket contents screen.

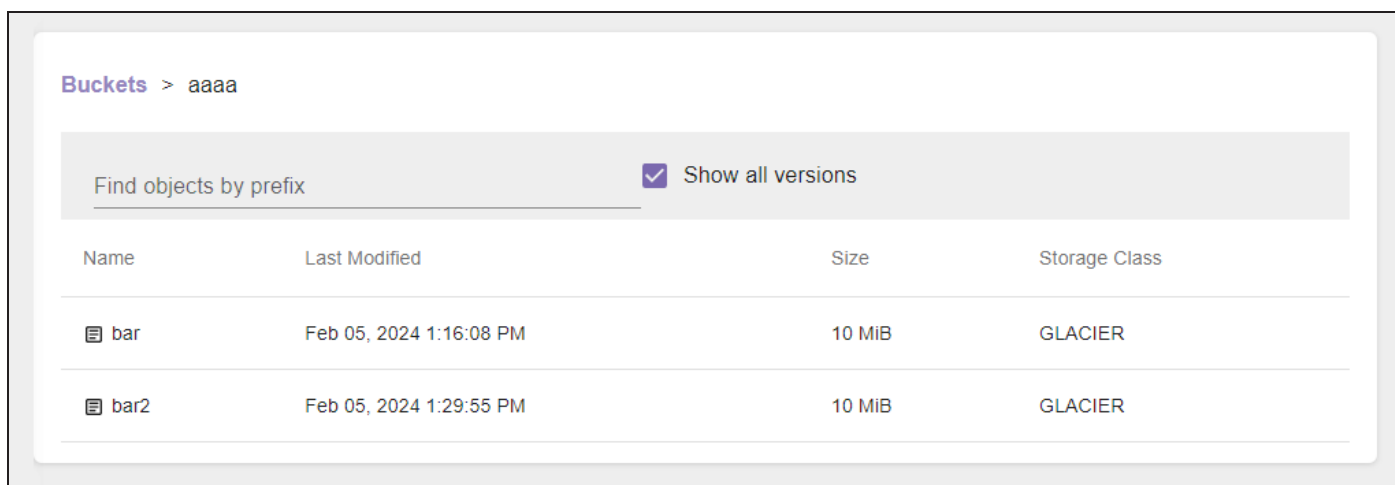
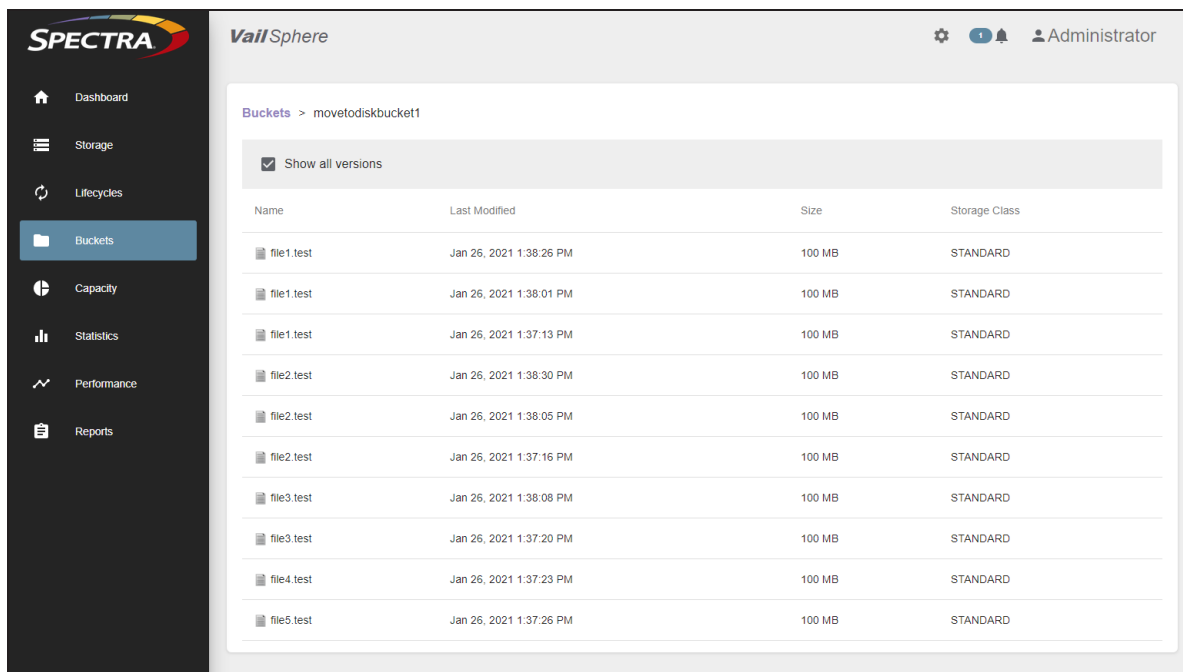


Figure 115 The Bucket Contents screen.

- Use the **Find objects by prefix** entry field to filter objects.

- Click **Show All Versions** to display every object version in the Vail bucket. The Last Modified field displays the day and time the object was uploaded.

Note: This option only displays if the bucket is configured for versioning.



SPECTRA VailSphere Administrator

Buckets > movetodiskbucket1

☒ Show all versions

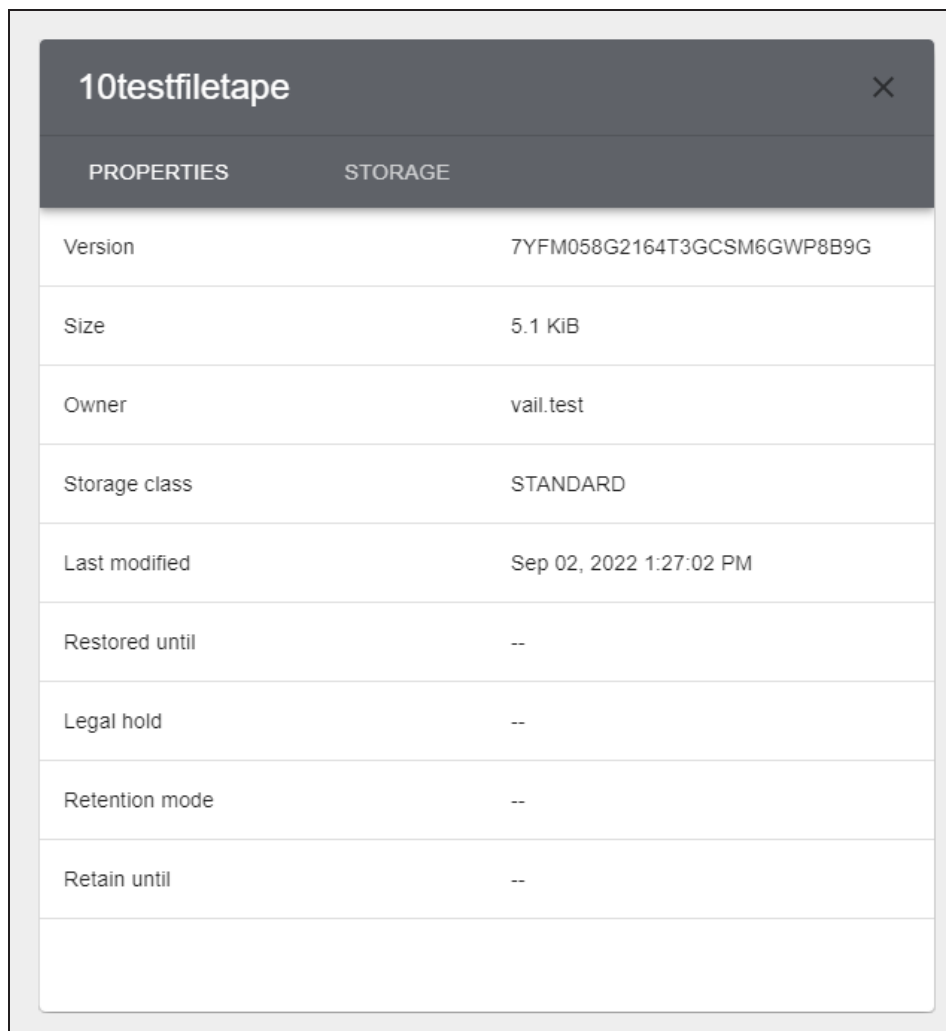
Name	Last Modified	Size	Storage Class
file1.test	Jan 26, 2021 1:38:26 PM	100 MB	STANDARD
file1.test	Jan 26, 2021 1:38:01 PM	100 MB	STANDARD
file1.test	Jan 26, 2021 1:37:13 PM	100 MB	STANDARD
file2.test	Jan 26, 2021 1:38:30 PM	100 MB	STANDARD
file2.test	Jan 26, 2021 1:38:05 PM	100 MB	STANDARD
file2.test	Jan 26, 2021 1:37:16 PM	100 MB	STANDARD
file3.test	Jan 26, 2021 1:38:08 PM	100 MB	STANDARD
file3.test	Jan 26, 2021 1:37:20 PM	100 MB	STANDARD
file4.test	Jan 26, 2021 1:37:23 PM	100 MB	STANDARD
file5.test	Jan 26, 2021 1:37:26 PM	100 MB	STANDARD

Figure 116 The Bucket Contents - Show All Versions screen.

- Click **Buckets** in the upper-left corner of the pane to return to the Buckets screen.

View Object Details

On the Bucket Details screen, **click the row** of an object to view its details. By default, the **Properties** pane displays.



The screenshot shows a modal window titled "10testfiletape" with a close button (X) in the top right corner. Below the title bar is a tabbed interface with two tabs: "PROPERTIES" (selected) and "STORAGE". The "PROPERTIES" tab displays a table with the following data:

PROPERTIES	STORAGE
Version	7YFM058G2164T3GCSM6GWP8B9G
Size	5.1 KiB
Owner	vail.test
Storage class	STANDARD
Last modified	Sep 02, 2022 1:27:02 PM
Restored until	--
Legal hold	--
Retention mode	--
Retain until	--

Figure 117 The Object Details - Properties screen.

Field	Description
Version	The UUID for the current version of the object.
Size	The object size on the storage target.
Owner	The AWS account name of the owner of the object.

Field	Description
Storage Class	The current storage class for the object. Note: The existence of a GLACIER clone does not necessarily cause the storage class of the object to change to GLACIER. If a non-GLACIER clone exists, (such as objects originally written to STANDARD storage) the object has a STANDARD storage class. This is true even if the STANDARD clone is optional.
Last Modified	The last modified date of the object.
Restored Until	The timestamp of when the object expires.
Legal Hold	Indicates if the object has a legal hold.
Retention Mode	Indicates the retention mode.
Retain Until	The duration that the object is retained by a legal hold.

Click **Storage** to display the current storage information for the object.

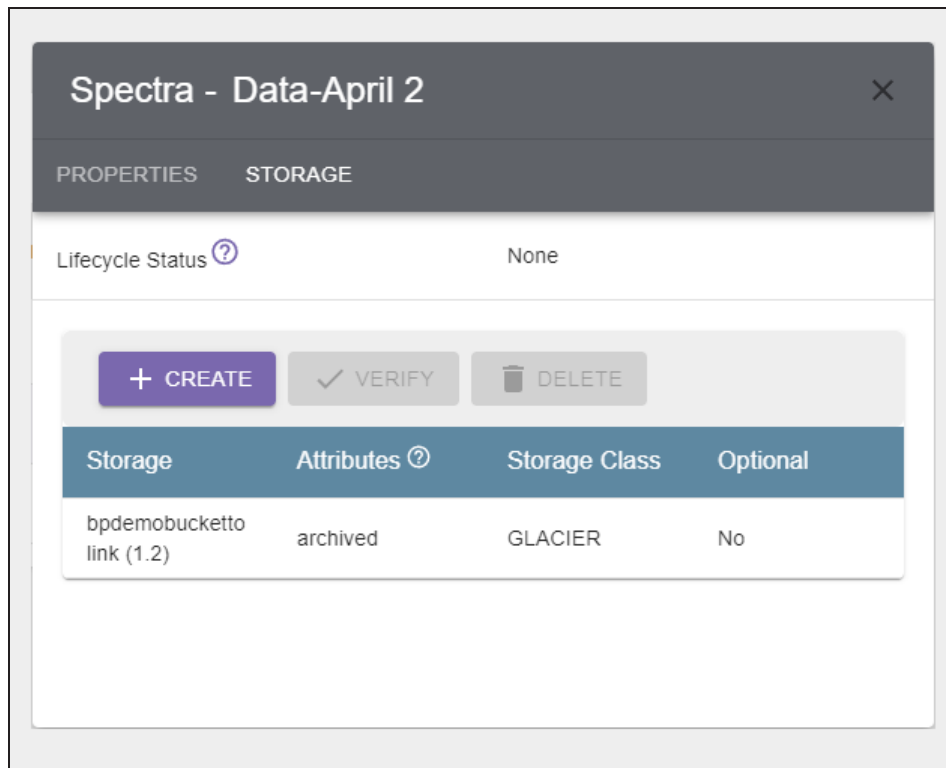


Figure 118 The Object Details - Storage screen.

Field	Description
Lifecycle Status	Indicates what Lifecycle-based changes are scheduled for the object.
Storage	<p>The name of the storage endpoint where the object is stored.</p> <p>If the object is 256 bytes or less after compression, it is stored in the application database and not on a storage endpoint. The storage field is blank when the object is stored in the database.</p> <p>Note: If the object is stored in the database but the Lifecycle targets a linked bucket storage endpoint, the application clones the object to the storage endpoint to ensure the contents of linked buckets are synchronized.</p>
Attributes	<p>Archived - The object is archived and must be restored in order to be accessed.</p> <p>Restored - The object is restored can be accessed.</p>
Storage Class	The current storage class for the object. See Storage Classes on page 297 for information on each storage class.

Field	Description
	Note: The existence of a GLACIER clone does not necessarily cause the storage class of the object to change to GLACIER. If a non-GLACIER clone exists, (such as objects originally written to STANDARD storage) the object has a STANDARD storage class. This is true even if the STANDARD clone is optional.
Optional	If yes, the clone is deleted when space is required.

CREATE AN OBJECT CLONE

If desired, you can delete a clone of an object in a Vail bucket using the Vail management console. You can only create an object clone if the object does not exist on all storage targets. You cannot have multiple clones on the same storage target.

Here is how to create an object clone using the Vail management console:

1. In the Vail management console taskbar, click **Buckets**.




Figure 119 The Buckets screen.

2. Under the **Buckets** banner, click a **bucket name**.

Note: You must click the name directly. Clicking the row of the bucket does not display the bucket contents screen.




Figure 120 The Bucket Contents screen.

3. If necessary, click **Show All Versions** to display every object version in the Vail bucket. The Last Modified field displays the day and time the object was uploaded.




Figure 121 The Bucket Contents - Show All Versions screen.

4. **Click** the row of the object you want to clone. The Object Properties window displays.

- Click the **Storage** tab.

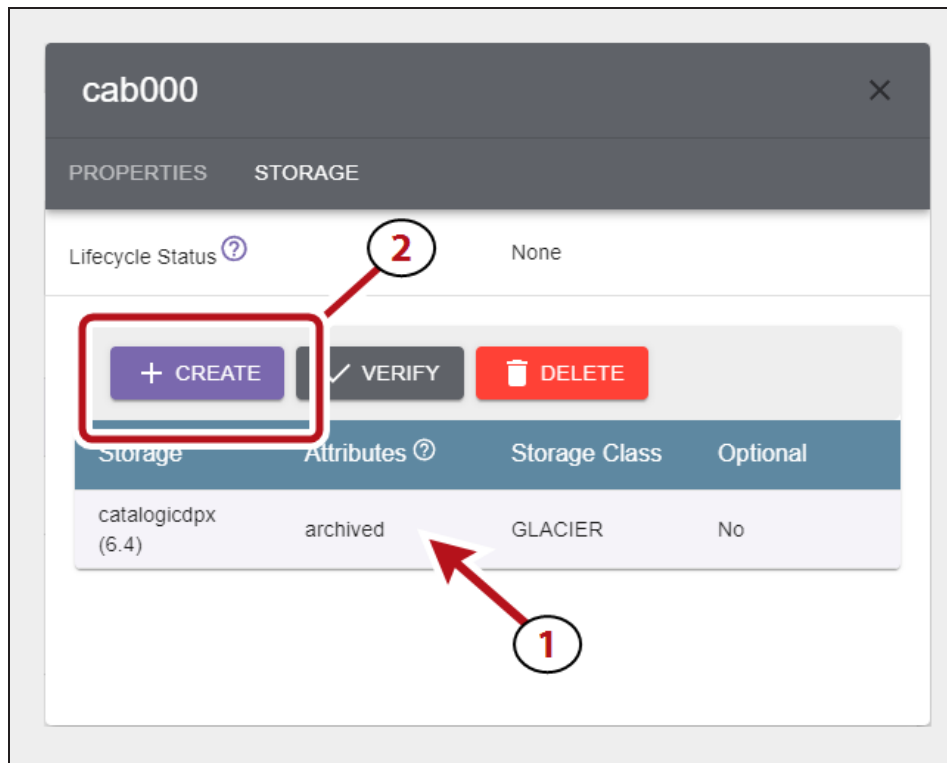


Figure 122 The Object Details - Storage screen.

- Select the row of the object (1), and click **Create** (2).
- Using the **Select Storage** drop-down menu, select a location to create the object clone.

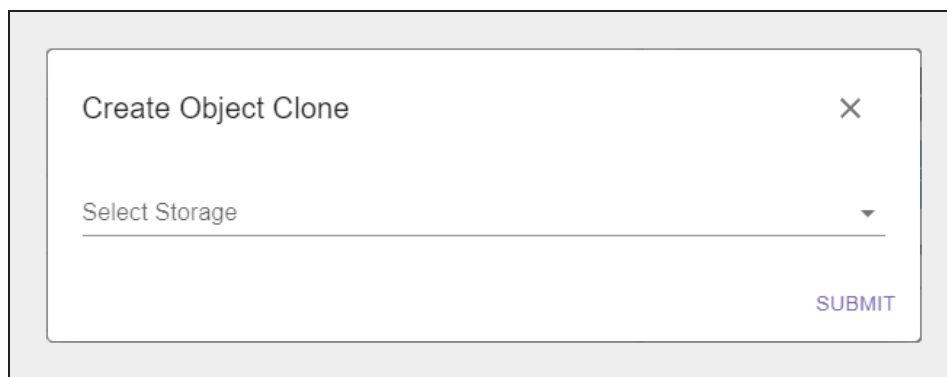


Figure 123 The Create Object Clone screen.

- Click **Submit** on the confirmation screen to create an object clone.

VERIFY AN OBJECT CLONE

Here is how to verify an object clone using the Vail management console:

1. In the Vail management console taskbar, click **Buckets**.

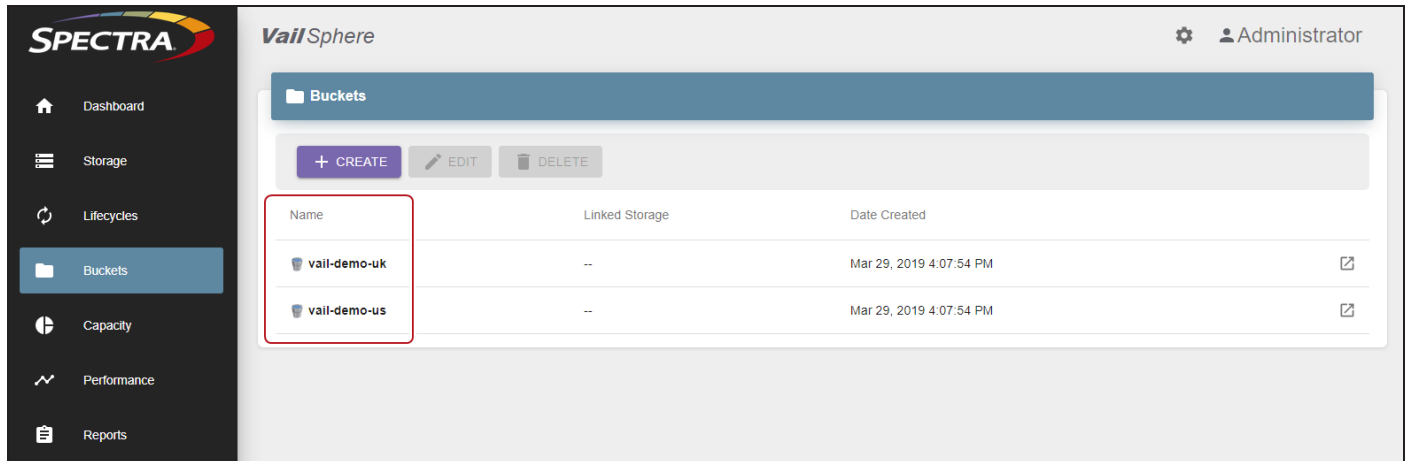


Figure 124 The Buckets screen.

2. Under the **Buckets** banner, click a **bucket name**.

Note: You must click the name directly. Clicking the row of the bucket does not display the bucket contents screen.

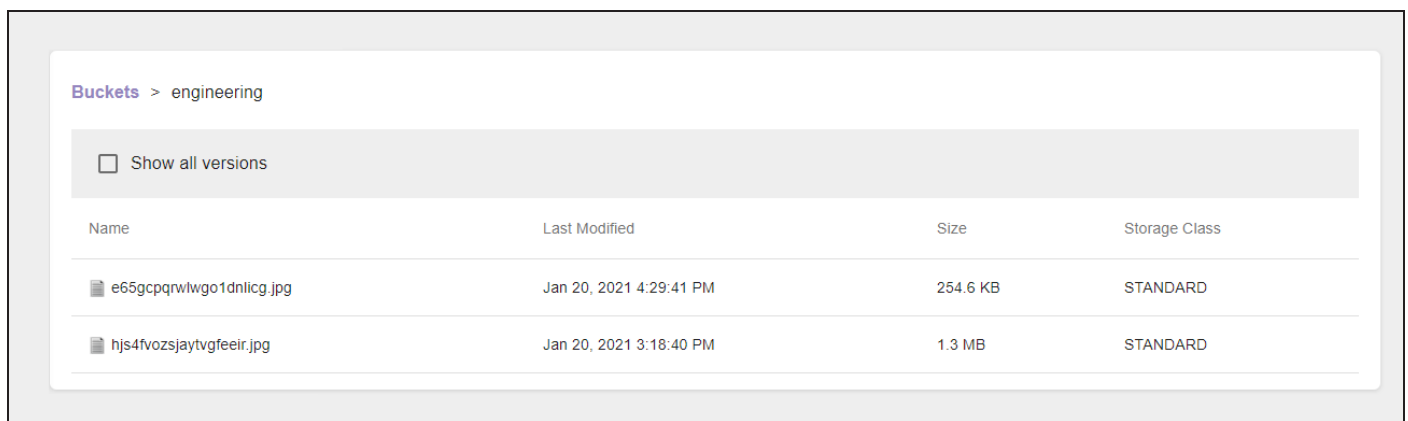


Figure 125 The Bucket Contents screen.

3. If necessary, click **Show All Versions** to display every object version in the Vail bucket. The Last Modified field displays the day and time the object was uploaded.




Figure 126 The Bucket Contents - Show All Versions screen.

4. Click the row of the clone you want to delete. The Object Properties window displays.
5. Click **Storage**.

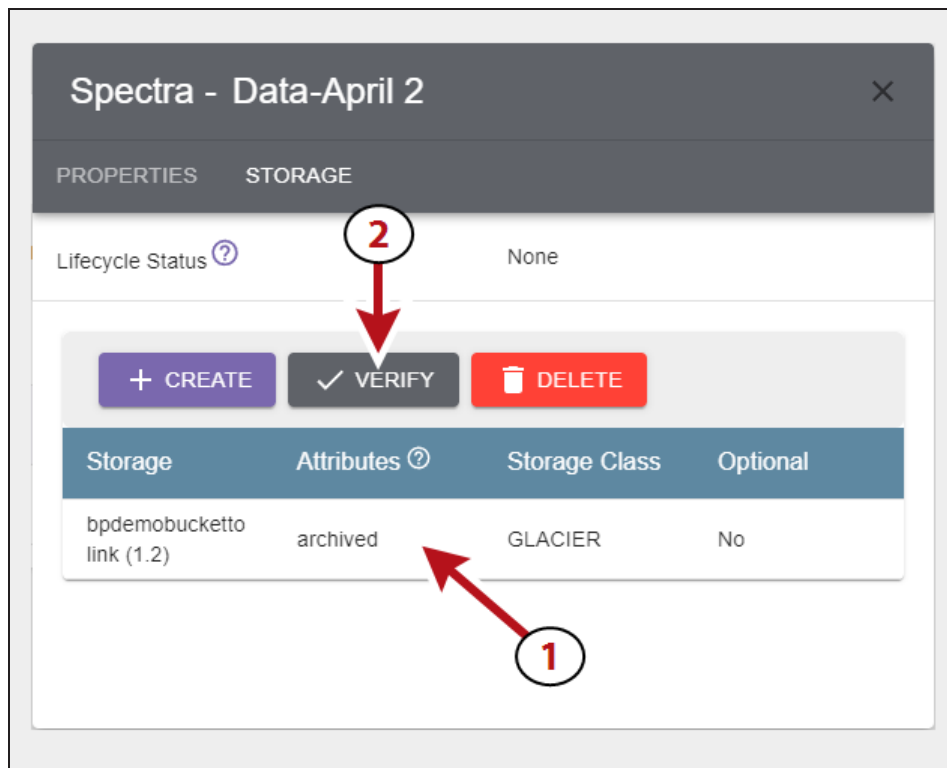


Figure 127 The Object Details - Storage screen.

6. Select the row of the clone (1), and click **Verify** (2).
7. Click **Submit** on the confirmation screen to verify the object clone.

DELETE AN OBJECT CLONE

If desired, you can delete a clone of an object in a Vail bucket using the Vail management console. You can only delete an object if another clone of the object exists elsewhere in the Vail sphere. If there is only one instance of the object in the sphere, it cannot be deleted.

Here is how to delete an object clone using the Vail management console:

1. In the Vail management console taskbar, click **Buckets**.

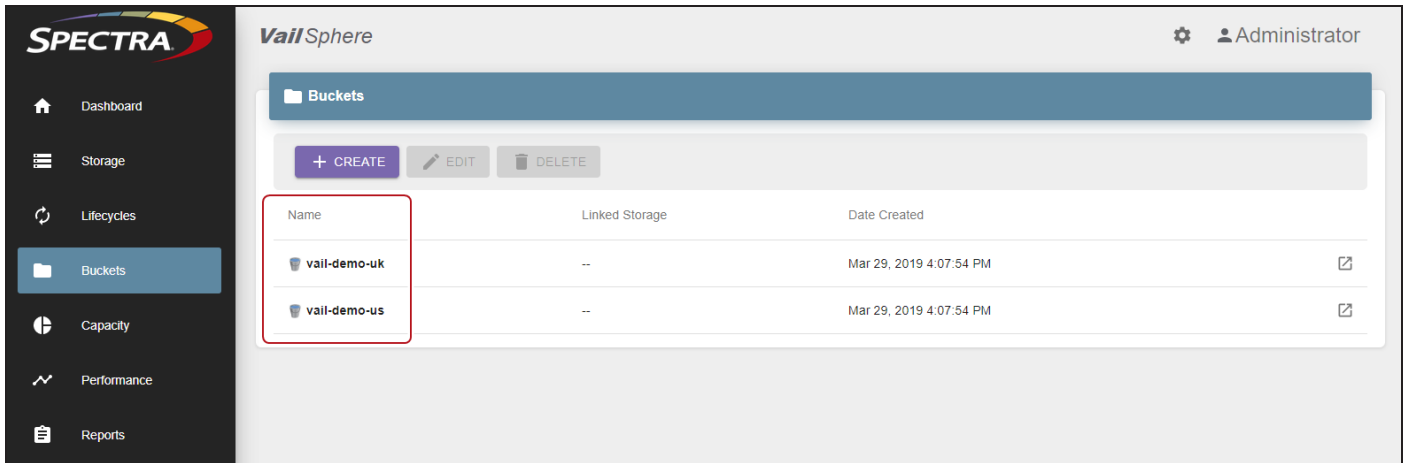


Figure 128 The Buckets screen.

2. Under the **Buckets** banner, click a **bucket name**.

Note: You must click the name directly. Clicking the row of the bucket does not display the bucket contents screen.

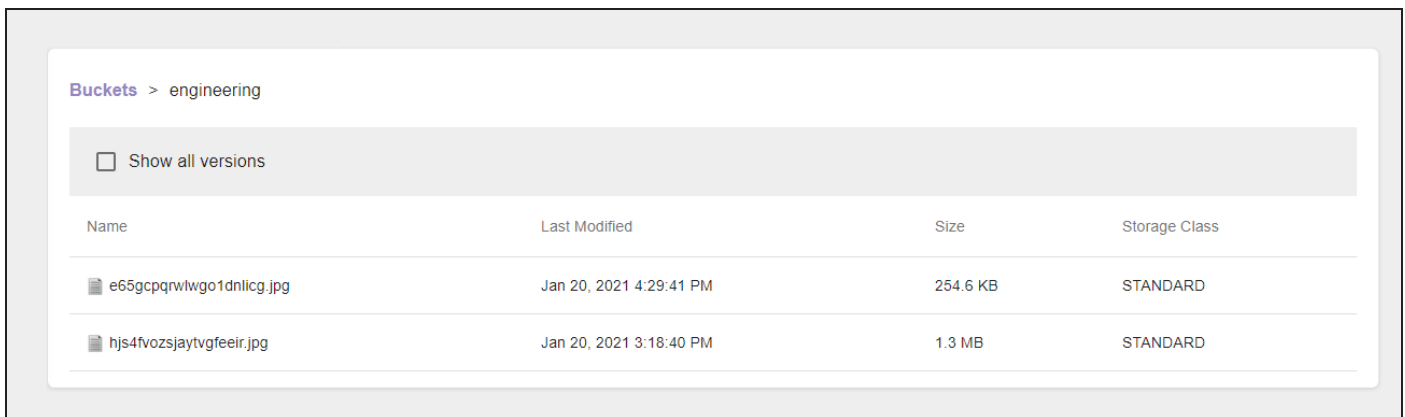


Figure 129 The Bucket Contents screen.

3. If necessary, click **Show All Versions** to display every object version in the Vail bucket. The Last Modified field displays the day and time the object was uploaded.




Figure 130 The Bucket Contents - Show All Versions screen.

4. Click the row of the clone you want to delete. The Object Properties window displays.
5. Click **Storage**.




Figure 131 The Object Details - Storage screen.

6. Select the row of the clone (1), and click **Delete** (2).
7. Click **Delete** on the confirmation screen to delete the object clone.

EDIT GLOBAL SETTINGS

If desired, you can edit the global settings of the Spectra Vail application to enable a diagnostic monitor or to change the nightly processing time used by the application.

Change Lifecycle Rule Nightly Processing Time



IMPORTANT

The Spectra Vail application restarts after changing the nightly processing time. Any data transfer operations fail when the application restarts. Internal operations, such as Lifecycles, automatically restart. External operations must be manually restarted.

Here is how to change the nightly processing time:

1. Discontinue storage operations. The Spectra Vail application restarts after changing the nightly processing time.
2. In the upper right corner of the Vail management console, click the **gear icon** and select **Global Settings**.
3. Under the **Global Settings** banner, click **Edit**.

Figure 132 The Edit Global Settings screen.

4. Enter the new UTC time for **Nightly Processing**.

Note: Changing this value does not affect any actions that are already scheduled.

5. Click **Submit**.

Set Default Lifecycle

The default lifecycle on the Vail Application automatically applies to buckets unless otherwise specified during bucket creation. Here is how to change the default Lifecycle:

1. In the upper right corner of the Vail Application, click the **gear icon** and select **Global Settings**.
2. Under the **Global Settings** banner, click **Edit**.

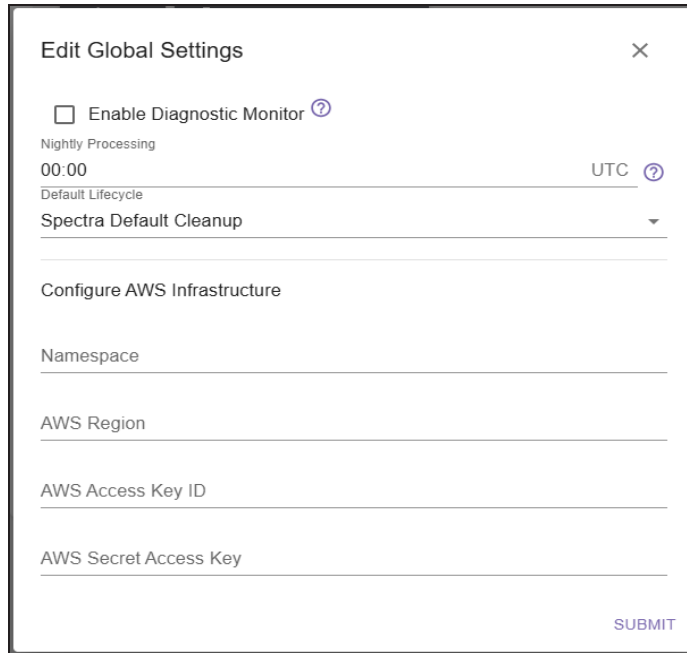


Figure 133 The Edit Global Settings screen.

3. Use the **Default Lifecycle** drop-down menu to select a previously created lifecycle, then click **Submit**.

Enable Diagnostic Monitor

The diagnostic monitor allows the Spectra Vail application to send diagnostic data to Spectra Logic.

Note: Contact Spectra Logic Technical Support before enabling the diagnostic monitor.

1. In the upper right corner of the Vail management console, click the **gear icon** and select **Global Settings**.
2. Under the **Global Settings** banner, click **Edit**.

Edit Global Settings

☐ Enable Diagnostic Monitor ?

Nightly Processing

00:00 UTC ?

Default Lifecycle

Spectra Default Cleanup

Configure AWS Infrastructure

Namespace

AWS Region

AWS Access Key ID

AWS Secret Access Key

SUBMIT

Figure 134 The Edit Global Settings screen.

3. Select **Enable Diagnostic Monitor**, then click **Submit**.

Configure AWS Infrastructure

For a local-control Vail sphere, configuring the AWS infrastructure settings is required in order to access AWS S3 buckets and to add IAM accounts in to the Vail application. Editing these settings assumes familiarity with your AWS environment.

Note: In a cloud-control Vail sphere, these settings are pre-populated and cannot be changed.

1. In the upper right corner of the Vail management console, click the **gear icon** and select **Global Settings**.
2. Under the **Global Settings** banner, click **Edit**.

Edit Global Settings

☐ Enable Diagnostic Monitor ?

Nightly Processing

00:00 UTC ?

Default Lifecycle

Spectra Default Cleanup

Configure AWS Infrastructure

Namespace

AWS Region

AWS Access Key ID

AWS Secret Access Key

SUBMIT

Figure 135 The Edit Global Settings screen.

3. Enter information for the **Namespace**, **AWS Region**, and **AWS Access** credentials.
4. Click **Submit**.

USING PROXY CONNECTIONS

If desired, you can configure the Spectra Vail application to use a proxy server to connect with external servers.

Configure Proxy Connection

Here is how to configure a proxy connection:

1. In the upper right corner of the Vail management console, click the **gear icon** and select **Network**.
2. Under the **Proxy Servers** banner, click **Create**.

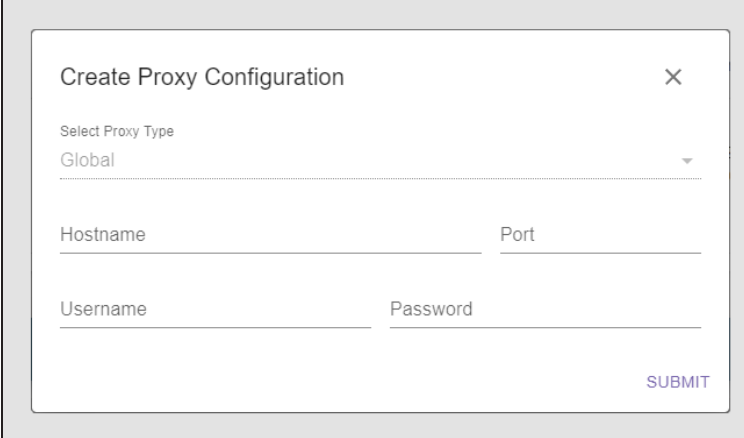


Figure 136 The Create Proxy Configuration screen.

Note: You can only configure a Global proxy type. The **Select Proxy Type** drop-down menu is grayed-out and not functional.

3. Enter the **Hostname** for the proxy server to use for external connections.
4. Enter the **Port** of the proxy server.
5. Enter the **Username** and **Password** to use when connecting through the proxy server.
6. Click **Submit**.

Edit Proxy Server

All options available when creating a proxy connection can be changed by editing the connection.

Here is how to edit a previously configured proxy configuration:

1. In the upper right corner of the Vail management console, click the **gear icon** and select **Network**.
2. Under the **Proxy Servers** banner, select the proxy connection and click **Edit**.
3. Update the proxy information as required, and click **Save**.

Delete Proxy Server

Here is how to delete a previously configured proxy configuration:

1. In the upper right corner of the Vail management console, click the **gear icon** and select **Network**.
2. Under the **Proxy Servers** banner, select the proxy connection and click **Delete**.
3. Update the proxy information as required, and click **Save**.

EDIT A VAIL BUCKET

If desired, you can edit Vail buckets to change various settings. You cannot change the bucket name, but all other settings used when creating a bucket are available when editing a Vail bucket, including encryption, versioning, access controls, and lifecycle selection.

Note: Prior to Vail 3.2.0, you cannot disable versioning if the bucket was initially configured to use versioning AND object locking when it was created. Starting with Vail 3.2.0, you are now able to change the versioning setting if the Vail bucket was created to use versioning.

Note: If you enable encryption on a bucket that is not currently configured to use encryption, only new data put to the bucket is encrypted. To encrypt existing data, you must use the PUT OBJECT copy command.

Here is how to edit a Vail bucket:

1. In the Vail management console taskbar, click **Buckets**.
2. Under the **Buckets** banner, select (1) the row of the bucket to edit, and (2) click **Edit**.

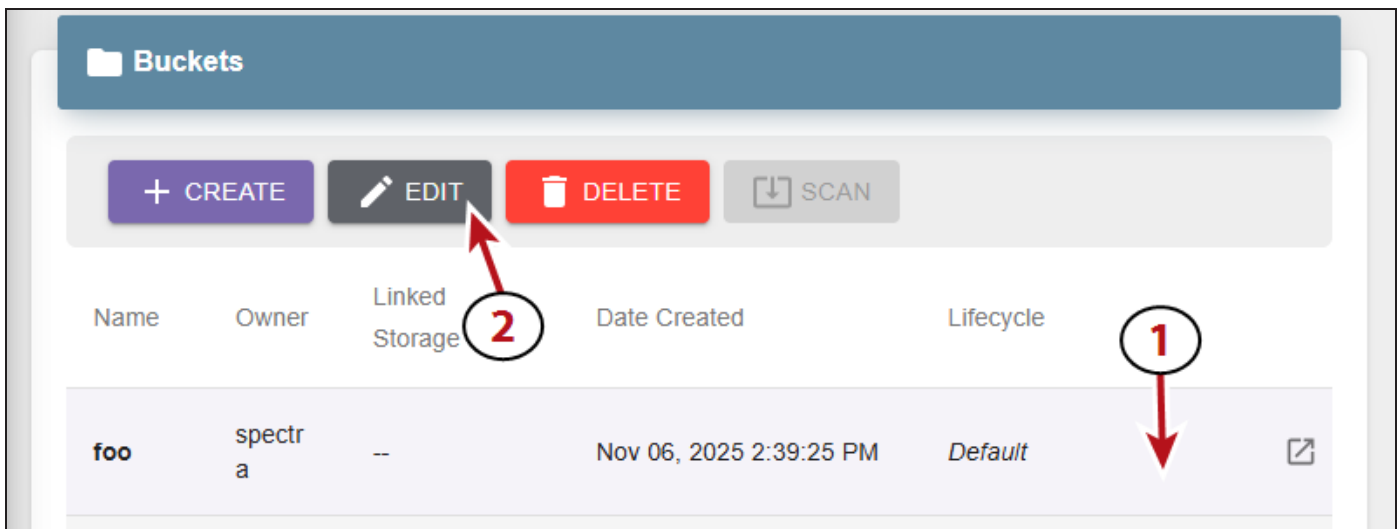


Figure 137 The Buckets pane.

3. Edit the settings on the Parameters screen as desired. See [Create a Vail Bucket on page 100](#) for information about each feature on the Parameters screen.

Note: Depending on the options selected when you created the bucket, the screens in this section may be different than what appears in the Vail management console.

Edit Bucket [Close]

1 Parameters 2 Policy 3 Lifecycle 4 Confirm

Edit bucket parameters.

Bucket Name
kctestbucket1

☐ Enable Versioning ☐ Enable Object Locking ?

☐ Enable Encryption ☒ Enable Compression ☐ Hide Glacier Operations ?

Bucket Owner
vail.development

Object Ownership
ACLs Disabled (recommended) ?

NEXT

Figure 138 The Edit Bucket - Parameters screen.

- Notes:**
- You are not able to change the Bucket Name.
 - If you disable versioning, any new objects are not versioned, but all previous versioned objects continue to be persisted.
4. Click **Next**. If you selected **Enable Object Locking** continue with [Step 5](#) below. Otherwise, skip to [Step 7 on page 181](#).

5. Edit the settings on the Retention screen as desired. See [Create a Vail Bucket on page 100](#) for information about each feature and option on the Retention screen.

Figure 139 The Edit Bucket - Retention screen.

6. Click **Next**.
7. Edit the settings on the Policy screen as desired. See [Create a Vail Bucket on page 100](#) for information about each feature and option on the Policy screen.

AWS documentation'. At the bottom right, there are 'PREVIOUS' and 'NEXT' buttons."/>

Figure 140 The Edit Bucket - Policy screen.

8. Click **Next**. If **Object Ownership** for this bucket is set to **ACL Disabled**, skip to Step 11 on page 182. Otherwise continue to Step 9 on page 182

9. Edit the settings on the Access Control Lists screen as desired. See [Create a Vail Bucket on page 100](#) for information about each feature and option on the Access Control List screen.

- Click **Add ACL** to add a new ACL to the bucket.
- Click the **trashcan icon** to delete an existing ACL.

Figure 141 The Edit Bucket - Access Control Lists screen.

10. Click **Next**.

11. If desired, use the **Select Lifecycle** drop-down menu to select a new lifecycle for the bucket, and click **Next**.

Figure 142 The Edit Bucket - Lifecycle screen.

12. Review the configuration, and click **Submit** to save the changes to the Vail bucket.

DELETE A VAIL BUCKET

If desired, you can delete an empty Vail bucket. To remove a bucket that contains objects, you must first delete all of the objects.

Here is how to delete a Vail bucket:

1. In the Vail management console taskbar, click **Buckets**.
2. Under the **Buckets** banner, (1) select the bucket and (2) select **Delete**.

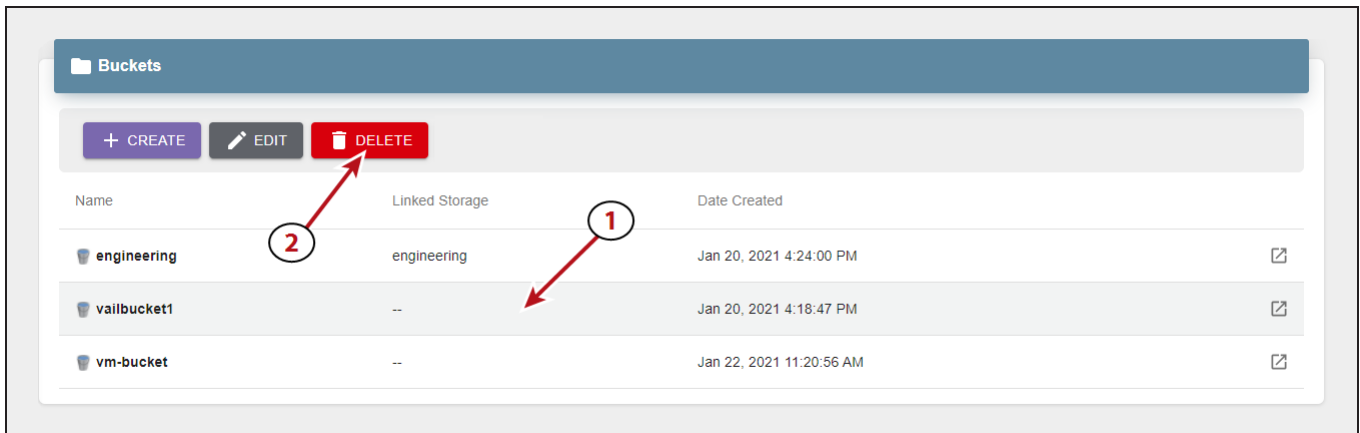


Figure 143 The Buckets pane.

3. On the confirmation screen, click **Delete**.

VIEW STORAGE DETAILS

The storage detail screen displays advanced information about endpoint or cloud storage, as well as data usage information.

Here is how to view the details of storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** or **Cloud Storage** banner, click the **View Details** icon to the far right of the storage for which you want to view details.

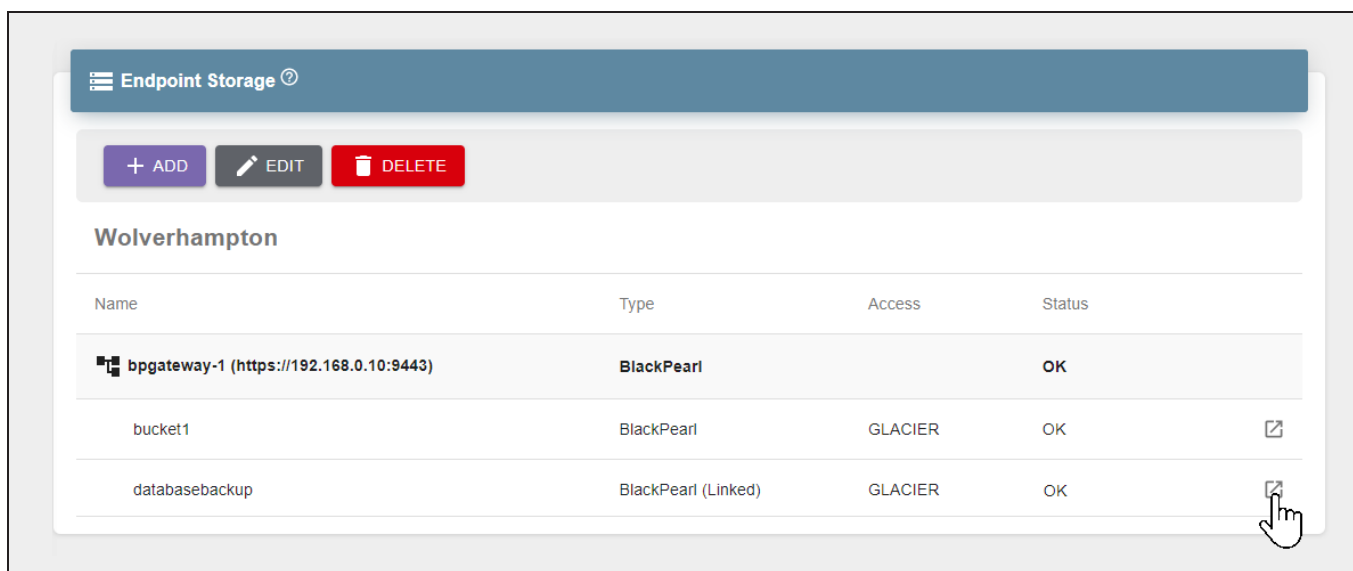


Figure 144 The Endpoint Storage pane.

- If you select **Properties** the following fields display.

Note: Not all fields listed below display for all storage endpoints. Some of the entries below may not apply to the storage endpoint for which you are viewing detailed information.

Field	Description
ID	The assigned ID of the storage which is used to identify the storage in certain error messages.
Type	The type of storage used on a BlackPearl or cloud storage endpoint.
Export	Indicates if the data in the storage was exported when the storage was created.
Storage Class	The storage class used by the storage endpoint.
Status	The current status of the storage endpoint.

Field	Description
Third-Party Recovery	Indicates if the option to allow third party recovery is enabled or disabled. This option writes additional data per object which allows full objects to be generated from the storage endpoint.
Restore In Place	Indicates if the storage endpoint creates clone data on the same endpoint where an object resides, or a different endpoint.
Disabled	Indicates if the storage is in an enabled or disabled state.
Bucket	The bucket used by the BlackPearl or cloud storage endpoint.
Link to Bucket	The Vail bucket linked to the storage endpoint.
Bucket Owner	The owner of the cloud bucket used by the storage endpoint.
Data Path Endpoint	The URL of the data path used by the cloud storage endpoint.
Skip TLS Verification	Indicates if the cloud storage endpoint uses or skips TLS verification.
Addressing Style	The method used to address the cloud storage endpoint.
Data Policy	The data policy on the BlackPearl Nearline Object Gateway used by the storage endpoint.
Data Partitions	The tape or disk partitions on the BlackPearl Nearline Object Gateway used by the storage endpoint.
Caution Threshold	The percentage of used space before the Vail application generates a caution system message.
Warning Threshold	The percentage of used space before the Vail application generates a warning message.
Optional Data	The maximum percentage of storage available for optional data clones.
Quota	The maximum percentage of storage available for use.

- If you select **Usage...**

Note: Not all fields listed below display for all storage endpoints. Some of the entries below may not apply to the storage endpoint for which you are viewing detailed information.

Field	Description
Number of clones	The number of clones kept by the storage endpoint for each object.

Field	Description
Total size of clone content	The total size of clones on the storage endpoint.
Total stored size of all clones	The amount of data used by the clones on the storage endpoint. Because clones are compressed before they are written to storage, this value may be different from the original content size.
Total size of optional data	<p>The amount of optional data stored on the storage endpoint. The Vail application only uses optional data to improve performance when there is sufficient storage space available. Optional data is automatically deleted when additional space on the storage endpoint is required to store non-optional data.</p> <p>Note: This option only displays for Volume storage endpoints.</p>
Average clone size	The average size of all clones on the storage endpoint.

EDIT BLACKPEARL OR VAIL VM ENDPOINT

If desired, you can edit the BlackPearl Nearline Object Gateway or Vail VM Node endpoint to change the location of the system in the Vail sphere, enable debug logging, or adding additional host names that can be used to access the endpoint.

Note: The images below show editing a BlackPearl endpoint. The processes are the same for a Vail VM node endpoint.

Change Endpoint Location

Here is how to change the regional location of an endpoint:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** banner, (1) select the row of the endpoint and (2) click **Edit**.

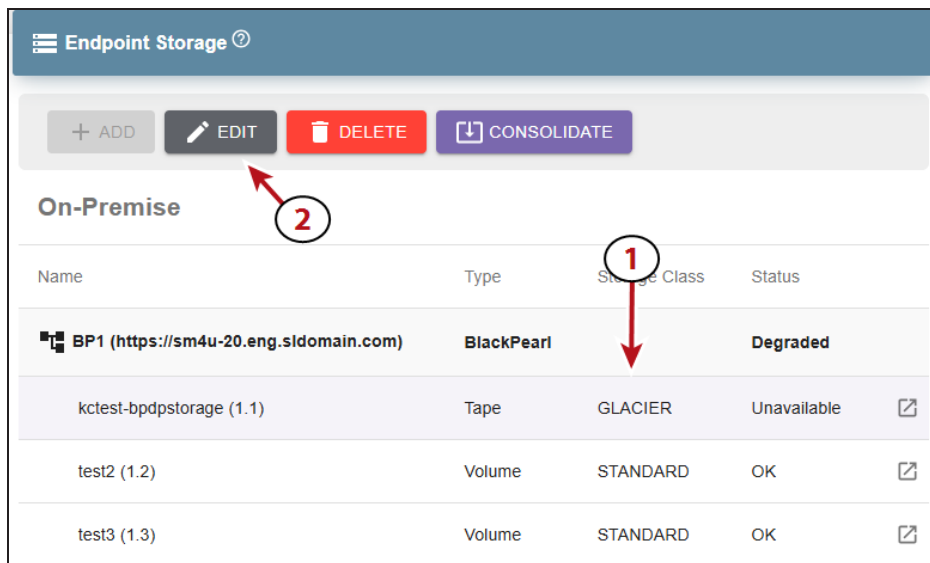
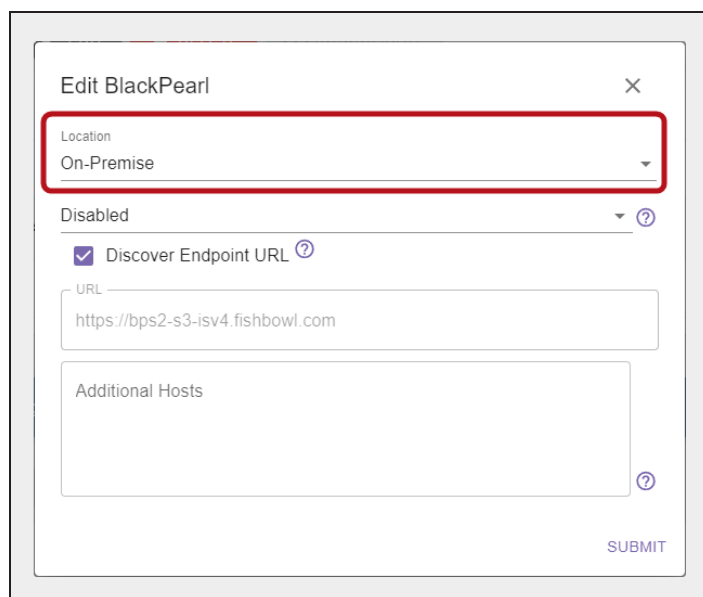


Figure 145 The Endpoint Storage pane.

3. Using the drop-down menu, select a new **Location** for the endpoint.



Edit BlackPearl

Location
On-Premise

Disabled

☒ Discover Endpoint URL

URL
https://bps2-s3-isv4.fishbowl.com

Additional Hosts

SUBMIT

Figure 146 Edit *Endpoint* - Location screen.

4. Click **Submit**.

Add Additional Host Names

Host names are used to access the endpoint. Here is how to add additional host names for the endpoint:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** banner, (1) select the row of the endpoint and (2) click **Edit**.

The screenshot shows the 'Endpoint Storage' pane. At the top, there are buttons: '+ ADD', 'EDIT', 'DELETE', and 'CONSOLIDATE'. Below these is the 'On-Premise' section. A table lists storage endpoints. The first row is highlighted, and a red arrow labeled '1' points to it. A second red arrow labeled '2' points to the 'EDIT' button.

Name	Type	Storage Class	Status
BP1 (https://sm4u-20.eng.sldomain.com)	BlackPearl		Degraded
kctest-bdpstorage (1.1)	Tape	GLACIER	Unavailable
test2 (1.2)	Volume	STANDARD	OK
test3 (1.3)	Volume	STANDARD	OK

Figure 147 The Endpoint Storage pane.

3. In the **Additional Hosts** dialog box, enter the desired host name(s).

Figure 148 Edit *Endpoint*- Additional Hosts screen.

4. Click **Submit**.

Change Endpoint URL

The URL listed on the Edit *Endpoint* screen is the address that other systems use when communicating with the storage endpoint. Typically this URL is discovered using name recognition sources.



IMPORTANT

The Spectra Vail application restarts after changing the endpoint URL. Any data transfer operations fail when the application restarts. Internal operations, such as Lifecycles, automatically restart. External operations must be manually restarted.

1. Discontinue storage operations. The Spectra Vail application restarts after changing the endpoint URL.
2. In the Vail management console taskbar, click **Storage**.
3. Under the **Endpoint Storage** banner, (1) select the row of the endpoint and (2) click **Edit**.

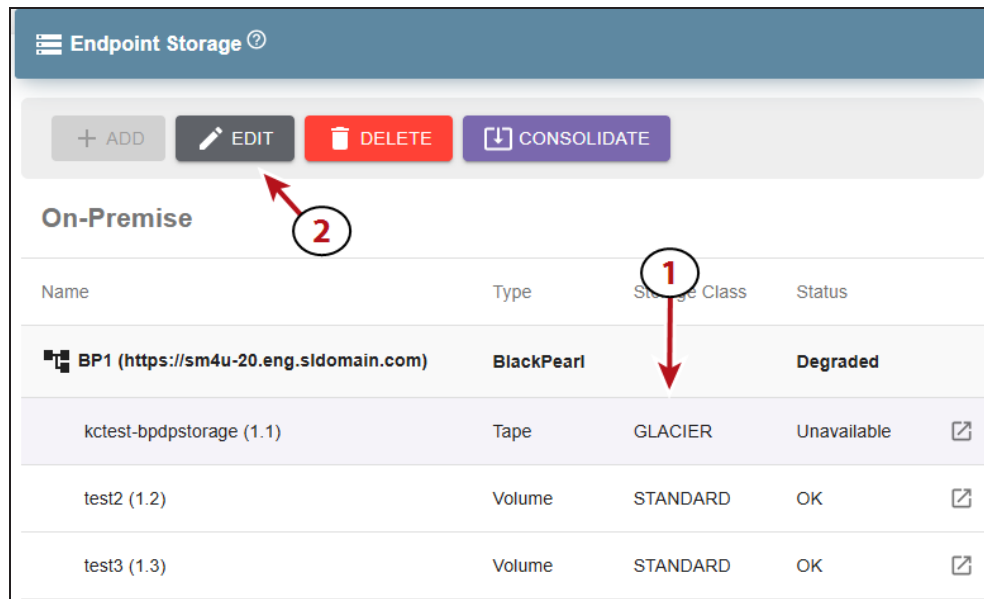


Figure 149 The Endpoint Storage pane.

- To determine the URL automatically, select **Discover Endpoint URL** checkbox.
- To set the URL manually, clear the **Discover Endpoint URL** checkbox, enter the **URL** in the entry field.

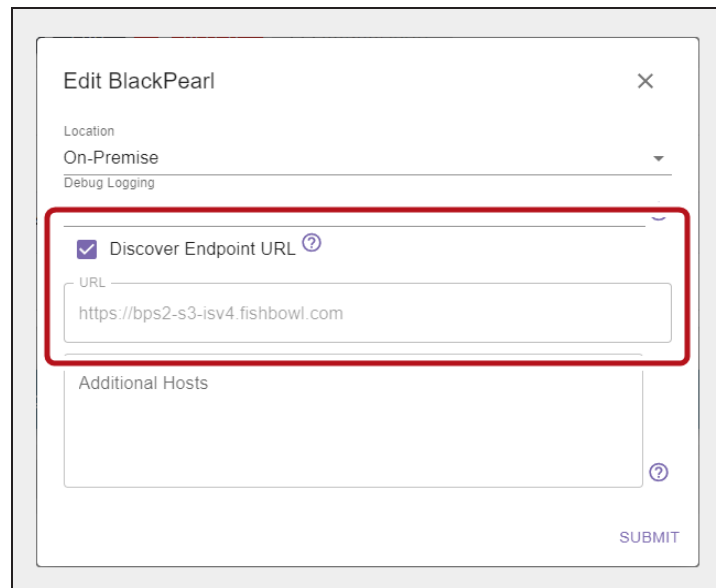


Figure 150 Edit *Endpoint* screen.

4. Click **Submit**.

Configure Debug Logging

The Spectra Vail application allows you to set the level of information included in system logs.

**IMPORTANT**

Contact Spectra Logic Technical Support before modifying this setting.

Here is how to edit the debug logging level for the endpoint:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** banner, (1) select the row of the endpoint and (2) click **Edit**.

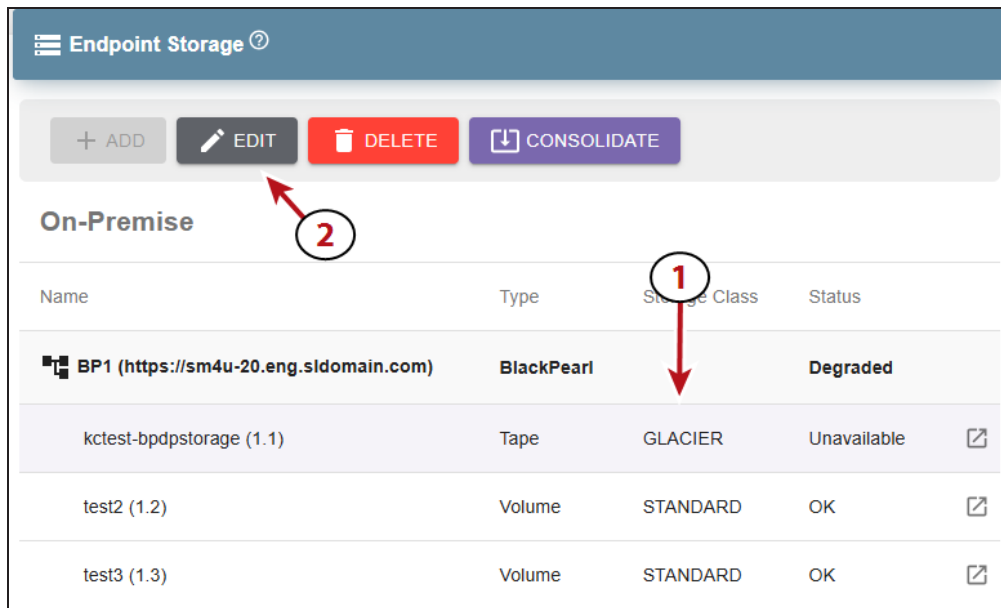


Figure 151 The Endpoint Storage pane.

3. Using the drop-down menu, select the **Debug Logging** level.

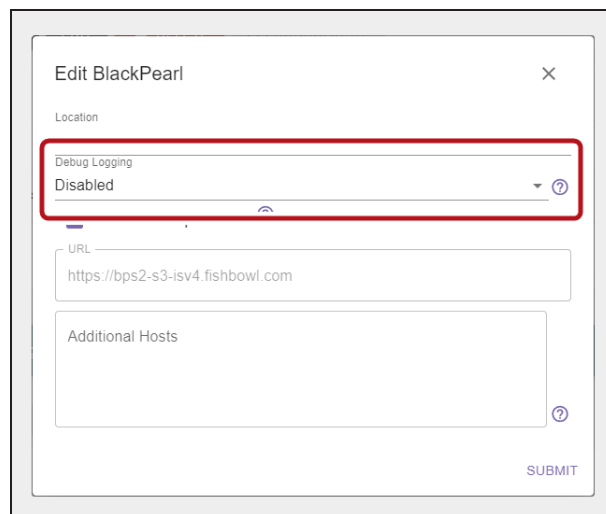


Figure 152 Edit *Endpoint* screen.

4. Click **Submit**.

EDIT STORAGE

If desired, you can edit storage to change various settings. The settings you can change are different for each type of storage.

Use one of the sections below to edit storage.

- **Edit BlackPearl Bucket Storage below**
- **Edit BlackPearl Volume Pool Storage on page 196**
- **Edit Vail VM Node Storage on page 198**
- **Edit Google Cloud Platform Storage on page 201**
- **Edit AWS S3 Cloud Storage on page 204**
- **Edit Microsoft Azure Cloud Storage on page 208**
- **Edit Other S3 Cloud Storage on page 212**

Edit BlackPearl Bucket Storage

Here is how to edit BlackPearl bucket storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** banner, (1) select the row of the storage and (2) click **Edit**.

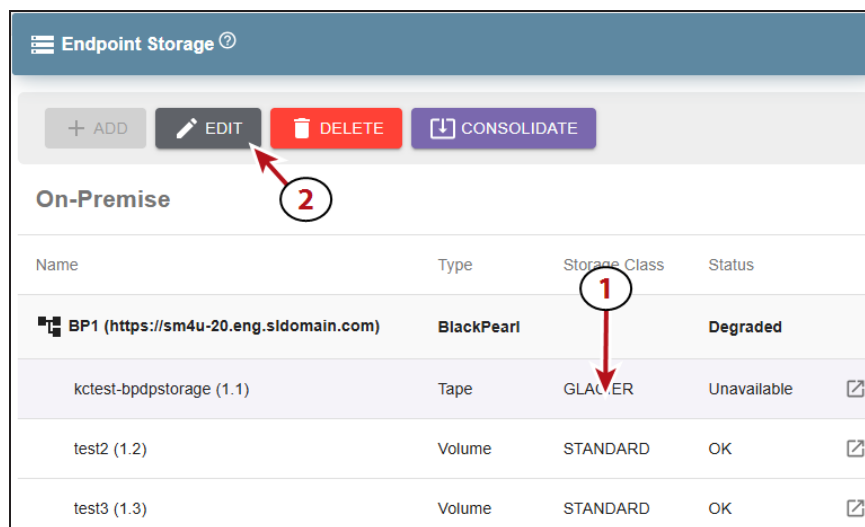


Figure 153 The Endpoint Storage pane.

3. If desired, edit the **Storage Name**, **Storage Class**, and the **Caution** and **Warning** thresholds.

Note: If you are editing a linked bucket, the fields for Caution and Warning Thresholds do not display.

Edit Endpoint Storage [X]

1 Parameters 2 Confirm

Configure your storage below.

BlackPearl Bucket
vail-1-1

Storage Name
kctest-bdpstorage

Select Storage Class
GLACIER

Use With Buckets
kctest1

☐ Third-party Recovery ? ☐ Restore In-place ?

Caution Threshold 70% Warning Threshold 80%

☐ Disabled

NEXT

Figure 154 The Edit Endpoint Storage - Parameters - BlackPearl screen.

4. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.
5. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

6. If you selected Glacier or Deep Archive as the storage class, you may select to enable **Restore In Place**. The restore in-place option uses internal resources on archival storage to create a clone in the same storage. If this setting is not enabled, new clones are created on different storage. Selecting this option may use additional cache space or incur additional cloud storage fees.

7. If desired, select **Disabled**. This sets the endpoint storage to a disabled state. The endpoint storage must be manually enabled before it can be used for data storage operations.
8. Click **Next**.
9. Review the configuration and click **Submit**.

Edit BlackPearl Volume Pool Storage

Here is how to edit BlackPearl volume pool storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** banner, (1) select the row of the storage and (2) click **Edit**.

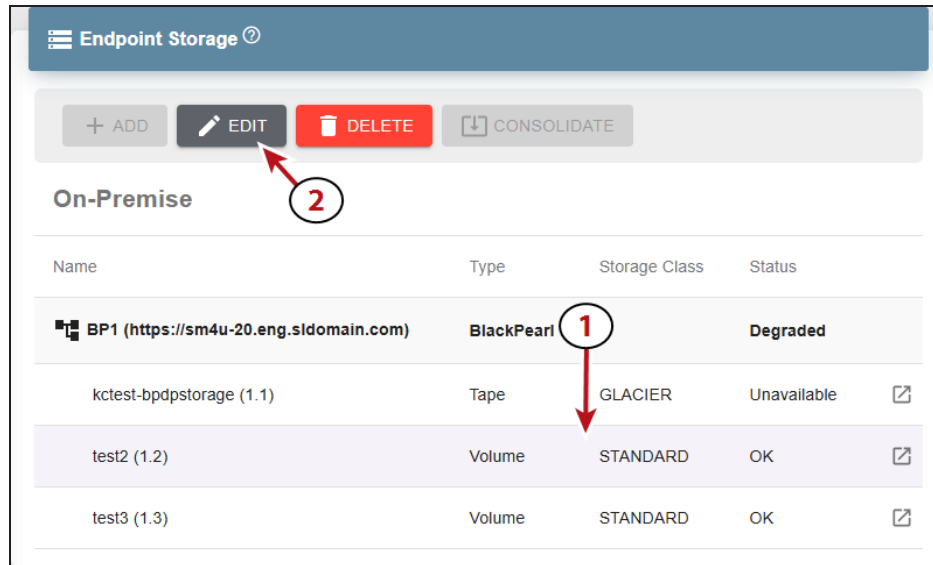


Figure 155 The Endpoint Storage pane.

3. If desired, you can change the **Storage Name**, **Storage Class**, **Caution Threshold**, or **Warning Threshold**.

Edit Endpoint Storage [X]

1 Parameters 2 Confirm

Configure your storage below.

Volume: /export/vail-1-2 Storage Name: test2

Select Storage Class: STANDARD

Use With Buckets: All buckets ?

☐ Third-party Recovery ?

Caution Threshold: 70 % Warning Threshold: 80 % ?

Optional Data: 10 % ?

Write Limit: MiB/sec ?

Quota: 3.5 TiB available Units: ?

☐ Disabled

NEXT

Figure 156 The Edit Endpoint Storage - Parameters screen.

4. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.

5. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

6. If desired, you can set the **Optional Data** threshold, which specifies the percentage of storage space to be used for optional clones to speed up data access. Optional clones are deleted as necessary to maintain space used below this percentage. If this field is left blank, no optional clones are stored and object access times are not tracked.
7. If desired, enter a value for a **Quota**, and use the **Units** drop-down menu to select a unit size for the quota value. This setting controls the maximum amount of storage space on the storage pool that is used for the BlackPearl volume pool storage endpoint. When this percentage is reached, no additional data is added to the storage endpoint. If you do not want to use a quota limit, leave the fields blank.
8. If desired, select **Disabled**. This sets the endpoint storage to a disabled state. The endpoint storage must be manually enabled before it can be used for data storage operations.
9. Click **Next**.
10. Review the configuration, and click **Submit** to save the changes to the BlackPearl storage.

Edit Vail VM Node Storage

Here is how to edit Vail VM node storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** banner, select the row of the storage and click **Edit**.
3. If desired, edit the **Storage Name**.

Figure 157 The Edit Endpoint Storage - Parameters - Vail VM Node screen.

4. If desired, use the **Select Storage Class** drop-down menu to change the storage class for the Vail VM storage endpoint.
5. If desired, use the **Use With Buckets** drop-down menu (not pictured) to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.
6. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

7. If desired, you can change the **Caution Threshold** or **Warning Threshold**.
8. If desired, change the **Optional Data** threshold, which specifies the percentage of storage space to be used for optional clones of objects that are no longer required to be present on the storage.
9. Click **Next**.

10. Review the configuration, and click **Submit** to save the changes to the Vail VM storage.

Edit Google Cloud Platform Storage

Here is how to edit Google Cloud Platform storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Cloud Storage** banner, (1) select the row of the storage and (2) click **Edit**.




Figure 158 The Endpoint Storage pane.

3. Select either **Modify Parameters** or **Modify Authorization**, then click **Next**.

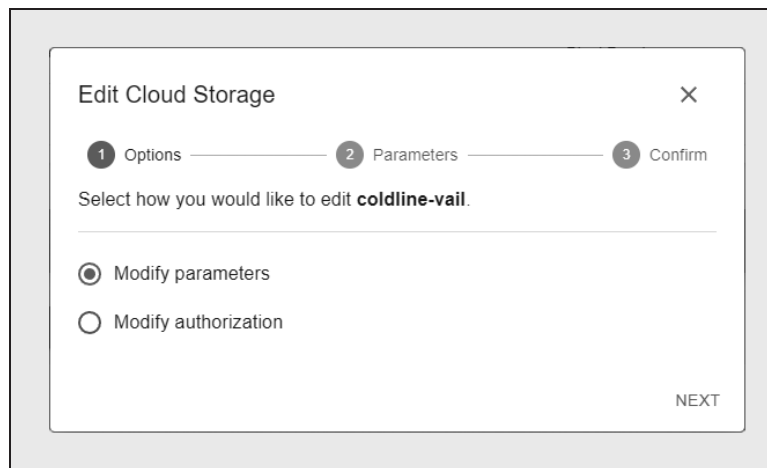


Figure 159 The Edit Cloud Storage - Options screen.

If you selected **Modify Parameters...**

- a. If desired, you can change the **Storage Name** and **Storage Class**.

Edit Cloud Storage [X]

1 Options — 2 **Parameters** — 3 Confirm

Provide additional information for this cloud storage below.

Google Cloud Bucket: bucket-for-chris [?] Storage Name: kctest-googlecloud

Select Storage Class: STANDARD_IA [?] Google Class: STANDARD [?]

Use With Buckets: All buckets [?]

☐ Export Objects [?] ☐ Third-party Recovery [?]

☒ Disabled

PREVIOUS NEXT

Figure 160 The Edit Cloud Storage - Parameters - Google Cloud Platform screen.

- b. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.
- c. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

Note: This option uses additional storage space and negatively affects performance.

- d. If desired, select **Disabled**. This sets the cloud storage to a disabled state. The cloud storage must be manually enabled before it can be used for data storage operations.
- e. Click **Next**.
- f. Review the configuration, and click **Submit** to save the changes to the cloud storage.

If you selected Modify Authorization...

- a. If desired, you can enter new **Google Cloud Platform JSON Credentials**.

Note: If you change your credentials in the Google Cloud Platform system, you must update the Spectra Vail application with the new credentials.

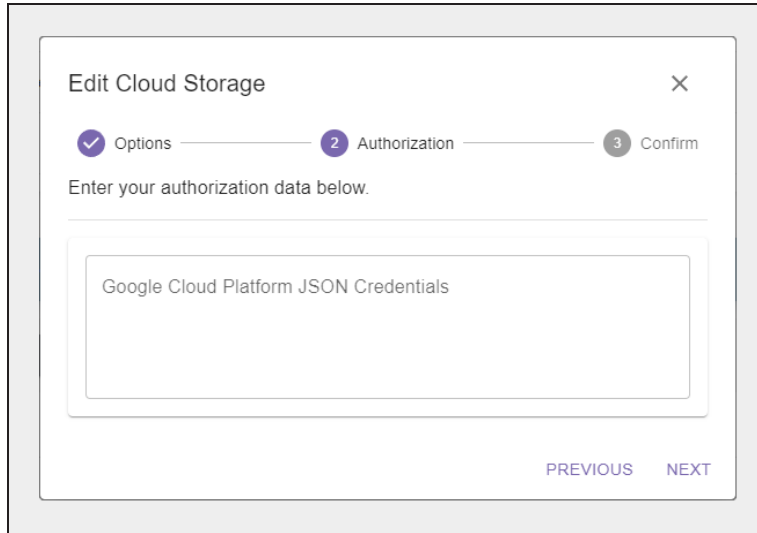


Figure 161 The Edit Cloud Storage - Authorization - Google Cloud Platform screen.

- b. Click **Next**.
- c. Review the configuration, and click **Submit** to save the changes to the cloud storage.

Edit AWS S3 Cloud Storage

Here is how to edit Amazon AWS S3 cloud storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Cloud Storage** banner, (1) select the row of the storage and (2) click **Edit**.




Figure 162 The Endpoint Storage pane.

3. Select either **Modify Parameters** or **Modify Authorization**, then click **Next**.




Figure 163 The Edit Cloud Storage - Options screen.

If you selected **Modify Parameters...**

- a. If desired, you can change the **Storage Name**.

Edit Cloud Storage [X]

✓ Options — 2 Parameters — 3 Confirm

Provide additional information for this cloud storage below.

Amazon S3 Bucket: chrisr-test [?] Storage Name: chrisr-test

Select Storage Class: STANDARD [?] Destination Class: Use Default [?]

Use With Buckets: All buckets [?]

☐ Import Content [?] ☐ Pause Notifications

☐ Export Objects [?] ☐ Current Version Only [?]

☐ Third-party Recovery [?] ☐ Restore In-place [?]

☐ Disabled

PREVIOUS NEXT

Figure 164 The Edit Cloud Storage - Parameters - AWS S3 Storage screen.

- b. If you are editing a linked bucket storage endpoint, if desired, select **Pause Notifications** to stop receiving notifications when the contents of the AWS bucket have changed. If you are editing a standard bucket storage endpoint, this setting is greyed out and cannot be changed.
- c. Change the **Storage Class** if desired.
- d. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.

- e. If you selected Glacier or Deep Archive as the storage class, you may select to enable **Restore In Place**. The restore in-place option uses internal resources on archival storage to create a clone in the same storage. If this setting is not enabled, new clones are created on different storage. Selecting this option may use additional cache space or incur additional cloud storage fees.
- f. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

- Notes:**
- This option uses additional storage space and negatively affects performance.
 - You cannot change this setting if you selected to Export Objects when you created the cloud storage.
- g. If desired, select **Disabled**. This sets the cloud storage to a disabled state. The cloud storage must be manually enabled before it can be used for data storage operations.
 - h. Click **Next**.
 - i. Review the configuration, and click **Submit** to save the changes to the cloud storage.

If you selected Modify Authorization...

- a. If desired, you can change the **Region**, **AWS Access Key** information and **AWS IAM** role settings.

The screenshot shows a modal window titled "Edit Cloud Storage" with a close button (X) in the top right corner. Below the title is a progress bar with three steps: "Options" (checked), "Authorization" (active), and "Confirm" (disabled). The main heading is "Enter your authorization data below." Below this is a form with five input fields: "Region", "AWS Access Key ID", "AWS Secret Access Key", "AWS IAM Role ARN" (with a help icon), and "AWS IAM Role External ID" (with a description: "Optional. Specifies who can assume the role."). At the bottom right of the form are two buttons: "PREVIOUS" and "NEXT".

Figure 165 The Edit Cloud Storage - Authorization - AWS S3 Storage screen.

- b. Click **Next**.
- c. Review the configuration, and click **Submit** to save the changes to the cloud storage.

Edit Microsoft Azure Cloud Storage

Here is how to edit Microsoft Azure cloud storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Cloud Storage** banner, (1) select the row of the storage and (2) click **Edit**.




Figure 166 The Endpoint Storage pane.

3. Select either **Modify Parameters** or **Modify Authorization**, then click **Next**.




Figure 167 The Edit Cloud Storage - Options screen.

If you selected Modify Parameters...

- a. If desired, you can change the **Storage Name** and **Storage Class**.

The screenshot shows the 'Edit Cloud Storage' dialog box with the 'Parameters' tab selected. The dialog has a progress bar at the top with three steps: 'Options' (checked), 'Parameters' (active), and 'Confirm'. Below the progress bar, it says 'Provide additional information for this cloud storage below.' The main form area contains several fields: 'Azure Container' with a dropdown menu showing 'test1', 'Storage Name' with a text input field containing 'kctest-azurecloud', 'Select Storage Class' with a dropdown menu showing 'STANDARD', 'Azure Tier' with a dropdown menu showing 'Cool', and 'Use With Buckets' with a dropdown menu showing 'kctest1'. There are also checkboxes for 'Export Objects', 'Third-party Recovery', and 'Disabled'. Each field has a question mark icon for help. At the bottom right, there are 'PREVIOUS' and 'NEXT' buttons.

Figure 168 The Edit Cloud Storage - Parameters - Azure Storage screen.

- b. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.
- c. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

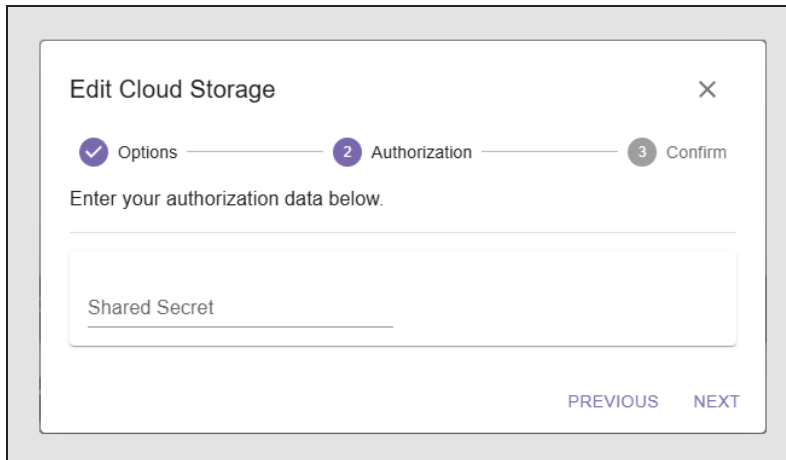
- Notes:**
- This option uses additional storage space and negatively affects performance.
 - You cannot change this setting if you selected to Export Objects when you created the cloud storage.
- d. If you selected Glacier or Deep Archive as the storage class, you may select to enable **Restore In Place**. The restore in-place option uses internal resources on archival storage to create a clone in the same storage. If this setting is not enabled, new clones are created on different storage. Selecting this option may use additional cache space or incur additional cloud storage fees.

Note: You cannot enable this setting if you selected to Export Objects when you created the cloud storage.

- e. If desired, select to **Disable** the endpoint storage. This sets the cloud storage to a disabled state. The cloud storage must be manually enabled before it can be used for data storage operations.
- f. Click **Next**.
- g. Review the configuration, and click **Submit** to save the changes to the cloud storage.

If you selected **Modify Authorization**...

- a. If desired, you can change the **Shared Secret** information.



The screenshot shows a modal window titled "Edit Cloud Storage" with a close button (X) in the top right corner. Below the title is a progress indicator with three steps: "Options" (marked with a checkmark), "2 Authorization" (the current step), and "3 Confirm". Below the progress bar, the text "Enter your authorization data below." is displayed. A text input field is present, with the placeholder text "Shared Secret". At the bottom right of the modal, there are two buttons: "PREVIOUS" and "NEXT".

Figure 169 The Edit Cloud Storage - Authorization - Azure Storage screen.

- b. Click **Next**.
- c. Review the configuration, and click **Submit** to save the changes to the cloud storage.

Edit Other S3 Cloud Storage

Here is how to edit Other S3 cloud storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Cloud Storage** banner, (1) select the row of the storage and (2) click **Edit**.




Figure 170 The Endpoint Storage pane.

3. Select either **Modify Parameters** or **Modify Authorization**, then click **Next**.

If you selected **Modify Parameters**...

- a. If desired, you can change the **Storage Name**.

Edit Cloud Storage [X]

1 Options — 2 **Parameters** — 3 Confirm

Provide additional information for this cloud storage below.

Cloud Bucket: joshbucket [?] Storage Name: joshbucket

Select Storage Class: STANDARD [?] Destination Class: Use Default [?]

Use With Buckets: foo kctest1 [?]

☐ Import Content [?] ☐ Pause Notifications

☐ Export Objects [?] ☒ Current Version Only [?]

☐ Third-party Recovery [?] ☐ Restore In-place [?]

Addressing Style [?]
☒ Path ☐ Virtual Hosted

☐ Disabled

PREVIOUS NEXT

Figure 171 The Edit Cloud Storage - Parameters - Other S3 Storage screen.

- b. If you are editing a linked bucket storage endpoint, if desired, select **Pause Notifications** to stop receiving notifications when the contents of the AWS bucket have changed. If you are editing a standard bucket storage endpoint, this setting is greyed out and cannot be changed.
- c. Edit the **Storage Class** if desired.
- d. Edit the **Destination Class** if desired.
- e. If desired, use the **Use With Buckets** drop-down menu to select buckets for the storage endpoint. This option configures the storage endpoint to only store data object clones on the selected buckets.

- f. If desired, select to enable **Third-Party Recovery**. This option creates a clone for delete markers, 0-length objects, and tiny objects, even though those clones are unnecessary. It also enables storage of full object metadata, which allows you to generate full objects from the content on the storage endpoint.

- Notes:**
- This option uses additional storage space and negatively affects performance.
 - You cannot enable this setting if you selected to Export Objects when you created the cloud storage.
- g. If you selected Glacier or Deep Archive as the storage class, you may select to enable **Restore In Place**. The restore in-place option uses internal resources on archival storage to create a clone in the same storage. If this setting is not enabled, new clones are created on different storage. Selecting this option may use additional cache space or incur additional cloud storage fees.
- h. Select the desired **Addressing Style**. This setting controls the URL format used when communicating with the cloud storage provider.

Selection	Description
Path Style	Path style formatting uses the bucket name as part of the URL path. Example: <i>http://endpoint/bucket-name/object-key</i>
Virtual-Hosted	Virtual-hosted style addressing uses the bucket as the prefix to the endpoint name Example: <i>http://bucket-name.endpoint/object-key</i>

- i. If desired, select **Disabled**. This sets the cloud storage to a disabled state. The cloud storage must be manually enabled before it can be used for data storage operations.
- j. Click **Next**.
- k. Review the configuration and click **Submit**.

If you selected **Modify Authorization...**

- a. If desired, edit URL address of the **Data Path Endpoint**.
- b. If desired, change the **Region** of the Other S3 endpoint.

Edit Cloud Storage [X]

✓ Options — 2 Authorization — 3 Confirm

Enter your authorization data below.

Data Path Endpoint
 Region

Access Key Secret Key

☐ Skip TLS Verification

PREVIOUS NEXT

Figure 172 The Edit Cloud Storage - Authorization - Other S3 Storage screen.

- c. If desired, you can change the **Access Key** and **Secret Key** information.
- d. If desired, select **Skip TLS Verification**. This option disables TLS certificate verification for HTTPS endpoints.

Note: This setting does not apply to HTTP endpoints.

- e. Click **Next**.
- f. Review the configuration, and click **Submit** to save the changes to the cloud storage.

CONSOLIDATE STORAGE

The consolidate storage function performs two tasks, consolidation of data packs and consolidation of metadata packs. Both tasks run when you consolidate storage, you cannot run one task separately.

Consolidate Storage Pack

This option is useful if you have deleted a large number of object clones and want to consolidate the partial data packs. The consolidate storage pack task runs everyday automatically at the scheduled daily processing time. You only need to consolidate storage packs manually if you do not want to wait for the daily processing schedule.

Consolidate Metadata Packs

This option is useful if you have third-party recovery enabled. The third-party recovery option writes daily metadata packs for use in recovering your data outside of the Vail environment. These metadata packs accumulate over time, so the consolidation of metadata packs merge these packs into the smallest number of metadata packs possible.

Note: The consolidate storage feature may take a long time depending on the number of objects.

Here is how you consolidate storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** or banner, (1) select the row of the storage, and (2) click **Consolidate**.




Figure 173 The Endpoint Storage pane.

3. On the confirmation screen, click **Consolidate**.

Note: The consolidate storage feature may take a long time depending on the number of objects.

DELETE STORAGE

When you delete storage, you can select to delete all data on the storage, or to move data to alternative storage.



CAUTION

If you select **Delete All Data**, any object clone that is **only** persisted on the storage is permanently deleted and cannot be recovered.

If you select **Choose Alternative Storage**, any object clone that exists only on the storage to be deleted is moved to the specified alternate storage. After all unique objects are moved, the storage is deleted.

To make sure you do not lose any data unintentionally, it is recommended to select **Choose Alternative Storage** and allow the Spectra Vail application to migrate any necessary data to alternative storage.

Note: You cannot delete storage that contains the only clone of a locked object.

Here is how to delete endpoint storage or cloud storage and optionally move data to alternative storage:

1. In the Vail management console taskbar, click **Storage**.
2. Under the **Endpoint Storage** or **Cloud Storage** banner, (1) select the row of the storage, and (2) click **Delete**.

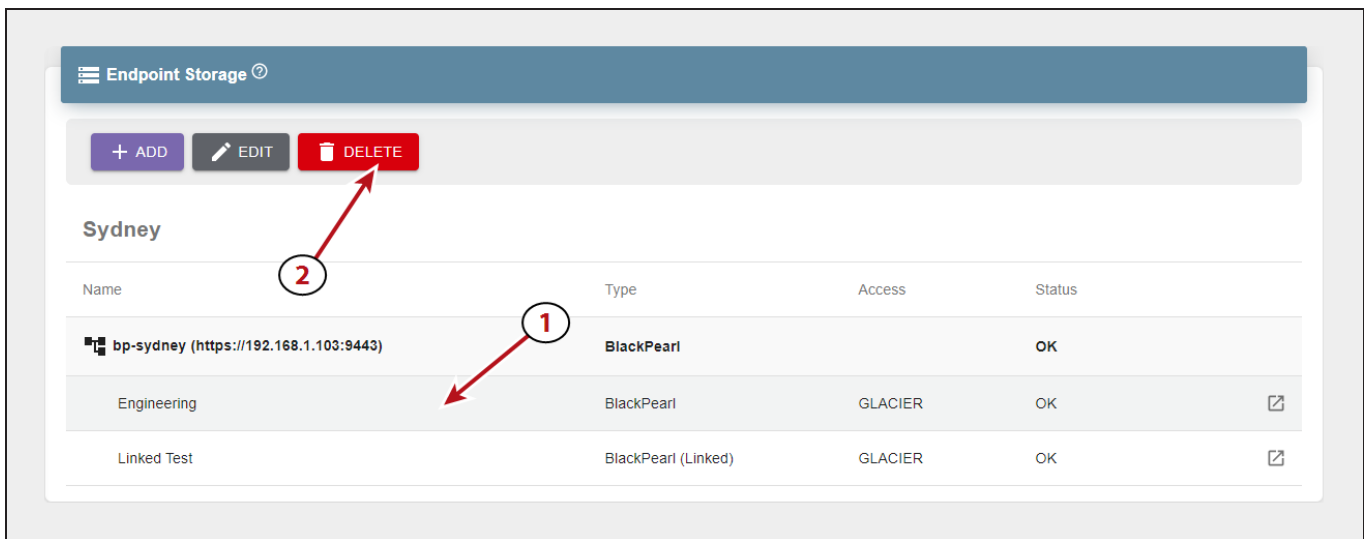


Figure 174 The Endpoint Storage pane.

To move unique object data to alternative storage:

1. Select **Choose Alternative Storage**.




Figure 175 The Delete Storage - Delete Options screen.

2. Using the **Select Storage** drop-down menu, select the storage to use as alternative storage.
3. Click **Next**.

4. Select the **check box** confirming you understand the storage is permanently deleted after moving unique object data to the alternative storage.




Figure 176 The Delete Storage - Confirm Alternative screen.

5. Click **Delete**.

To delete all data:

1. Select **Delete All Data** and click **Next**.



CAUTION

If you select **Delete All Data**, any object clone that is **only** persisted on the storage is permanently deleted and cannot be recovered.

2. Enter the name of the storage in **Confirmed Name** field.

The screenshot shows a 'Delete Storage' dialog box with a close button (X) in the top right corner. It features a progress bar with two steps: 'Delete Options' (completed, marked with a checkmark) and '2 Confirm Delete' (active). The main text asks, 'Are you sure you want to permanently delete 'Accounting'?'. Below this, a red warning message states: 'ALL DATA on storage 'Accounting' will be deleted. If there are objects unique to this storage, this action results in data loss.' A text box prompts the user to 'Enter the storage name below to verify your intent to permanently delete 'Accounting''. Below the text box is a field labeled 'Confirmed Name'. At the bottom right, there are two buttons: 'PREVIOUS' and 'DELETE'.

Figure 177 The Delete Storage - Confirm Delete screen.

3. Click **Delete**.

VIEW LIFECYCLE DETAILS

The lifecycles detail screen displays information about the selected lifecycle, including all lifecycle properties and rules.

Here is how to view the details of a lifecycle:

1. In the Vail management console taskbar, click **Lifecycles**.
2. Under the **Lifecycles** banner, click the **View Details** icon on the right side of the pane for the lifecycle which you want to view details.




Figure 178 The Lifecycles pane.

3. Click **Properties** or **Rules** to view the current lifecycle settings. Click the **X** in the upper-right corner to close the window.

- The Properties screen:

kctest-lifecycle1 ×	
PROPERTIES	RULES
Description	
Upload Expiration	3
Marker Expiration	Enabled
Restore To	kctest-azurecloud (0.1)
Force Initial Copy	Disabled
Ignore Storage Class	Disabled
Use With Buckets	kctest1
Modified	Nov 05, 2025 1:26:38 PM

Figure 179 The Lifecycle Rule Details - Properties screen.

Field	Description
Description	The text, if any, entered in the Description field when creating the bucket.
Upload Expiration	The number of days that must pass before a multipart upload is aborted. When a multipart upload is aborted, it deletes all parts associated with the upload, which prevents remaining incomplete uploads from being stored.
Marker Expiration	Indicates if the Delete Marker Expiration option is Enabled or Disabled .
Restore To	The configured location where clones are restored.
Force Initial Copy	Indicates if the lifecycle is configured to initially place data as STANDARD storage. Additional clones are created immediately as GLACIER storage.
Ignore Storage Class	Indicates if the lifecycle is configured to ignore the storage class requested in a PUT or upload operation and instead use the configured storage class of the selected storage endpoint.

Field	Description
Use With Buckets	Indicates if the lifecycle is configured to use specific buckets.
Modified	The date and time the lifecycle was last modified.

- The Rules screen:

Figure 180 The Lifecycle Rule Details - Rules screen.

Field	Description
Name	The name of the lifecycle.
Type	The type of lifecycle rule. Values: Clone, Move, Expiration.
Apply After	The number of days before the lifecycle rule is applied.
Destinations - Count	The number of destinations configured for the lifecycle. Values: 1-5, All.
Destinations - Storage	The storage endpoint(s) used by the lifecycle.

Field	Description
Destinations - Delete	Whether or not the lifecycle is configured to delete clones on storage destinations that are not configured in the lifecycle.
Filters - Versions	The versioning setting configured for the lifecycle.
Filters - Include	The text string used to filter objects to include in storage operations.
Filters - Exclude	The text string used to filter objects to exclude from storage operations.

EDIT A LIFECYCLE

If desired, you can edit a lifecycle to change how it controls data movement and retention. All settings used when creating a lifecycle are available when editing a lifecycle.

Here is how to edit a lifecycle:

1. In the Vail management console taskbar, click **Lifecycles**.
2. Under the **Lifecycles** banner, (1) select the lifecycle to edit, and (2) click **Edit**.

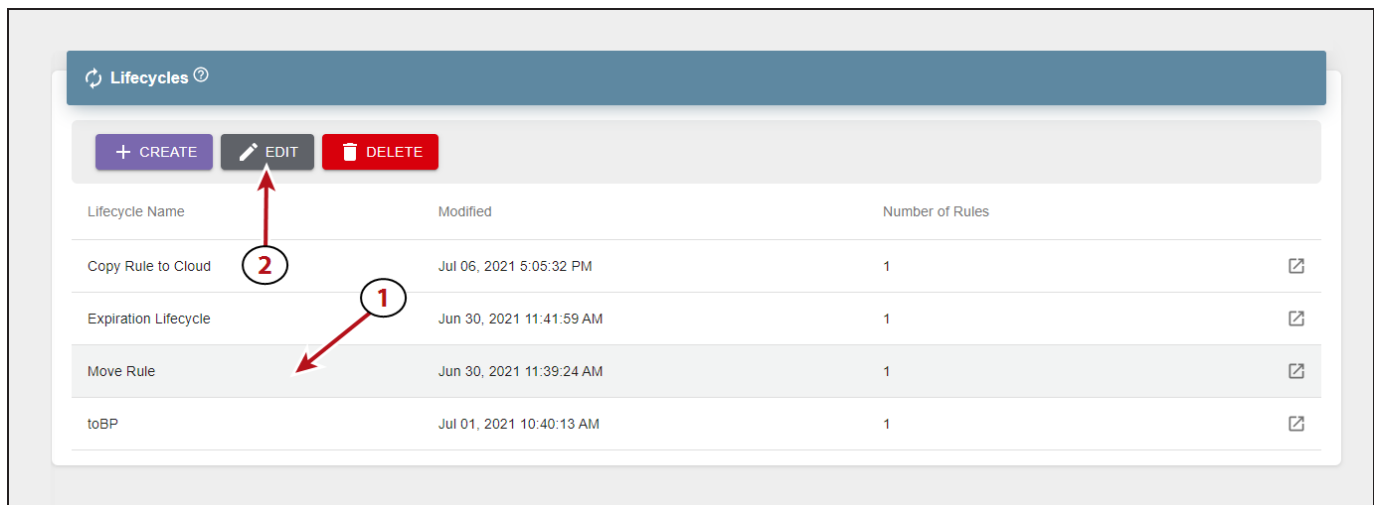


Figure 181 The Lifecycles screen.

3. If desired, edit the lifecycle **Name** and **Description**.

Edit Lifecycle [X]

1 Parameters 2 Rules

Edit lifecycle parameters.

Name: to_cloud Multipart Upload Expiration: 7 days

Restore To: Use Default

☒ Delete Marker Expiration

☐ Force Initial Copy

☐ Ignore Requested Storage Class

Description

NEXT

Figure 182 The Edit Lifecycle - Parameters screen.

4. If desired, enter a value for **Multipart Upload Expiration** in days. This setting controls how long the Spectra Vail application waits before aborting multipart uploads. When the multipart upload aborts, all parts of the upload are deleted. This prevents retaining multiple incomplete uploads.
5. If desired, use the **Restore To** drop-down menu to select a specific storage pool where you want to restore bucket objects. If you use the default setting, the Vail application decides which storage pool to restore objects.
6. If desired, select or clear **Delete Marker Expiration**. A delete marker keeps track of deletions of versioned objects so that the application can determine if the object is missing. If enabled, the Spectra Vail application removes delete markers when they are the last remaining version of an object.

7. If desired, select or clear **Force Initial Copy**. When enabled, the Vail application initially places data as STANDARD storage. Additional clones are created immediately as GLACIER storage. This may provide performance advantages as copying clones to GLACIER results in a clone that is ordered sequentially and more optimally packed.
8. If desired, select or clear **Ignore Requested Storage Class**. When enabled, the Vail application does not consider the storage class requested in a PUT or upload operation and instead uses the storage class of the selected storage endpoint.
9. Click **Next**.

Create Lifecycle ? X

✓ Parameters ————— 2 Rules

Define your rules for **test2**. ?

Note: Rules will be sorted based on schedule and version filters after submission. Maximum number of rules is 5.

Placement Rule

Name

Select Destination Storage ▼ ?

☐ Delete clones not on selected destination storage

Destination Count: All ▼

ADD SCHEDULE FILTER ▼

🗑

[NEW PLACEMENT RULE](#) | [NEW DELETION RULE](#)

PREVIOUS SUBMIT

Figure 183 The Edit Lifecycle - Rules screen.

- Use the links below to create or edit lifecycle rules.
 - **Add a Placement Rule on page 93**
 - **Add a Deletion Rule on page 97**
- To delete a lifecycle rule, click the **trash can icon**.

Note: There is no confirmation step for this action.

10. After making the desired changes, click **Submit**.

DELETE A LIFECYCLE

If desired, you can delete a lifecycle when its data placement schema is no longer needed.

Note: You cannot delete a lifecycle currently being used by a Vail bucket.

Here is how to delete a lifecycle:

1. In the Vail management console taskbar, click **Lifecycles**.
2. Under the **Lifecycles** banner, (1) select the lifecycle to delete, and (2) click **Delete**.

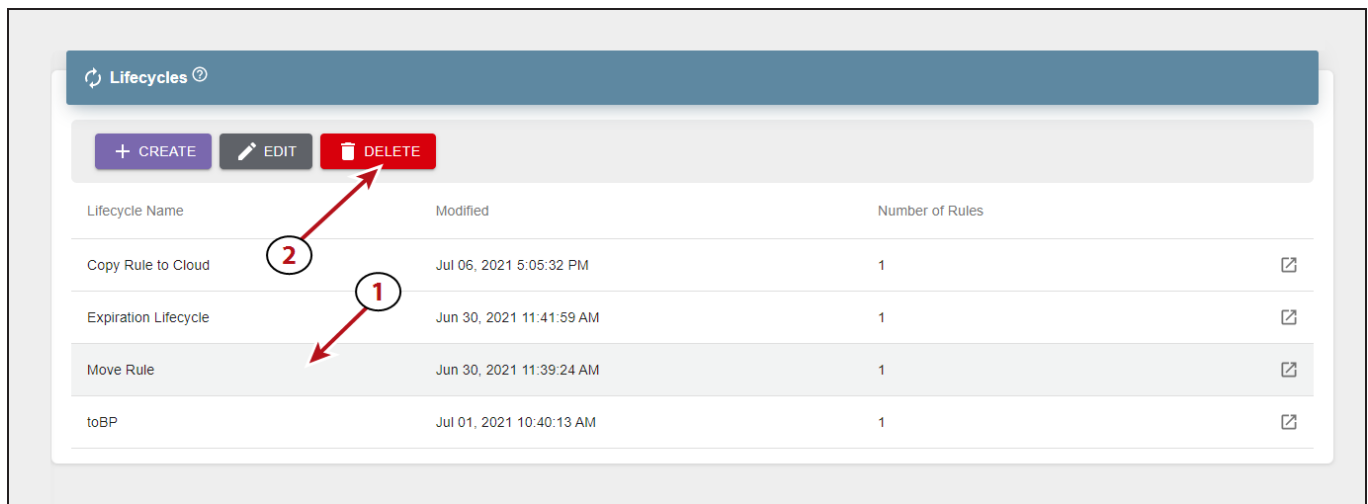


Figure 184 The Lifecycles pane.

3. A confirmation screen displays. Click **Delete** to confirm deleting the lifecycle.

CREATE A LOCATION

Locations are used to identify sites on the dashboard world map as well as to group storage endpoints by their physical location. Locations are only used in a cloud-control Vail sphere. These screens do not display in a local-control Vail sphere.

Note: You can also create a location when registering a BlackPearl node, or Vail VM node with the Vail sphere.

Here is how to create a location:

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and click **Locations (2)**.

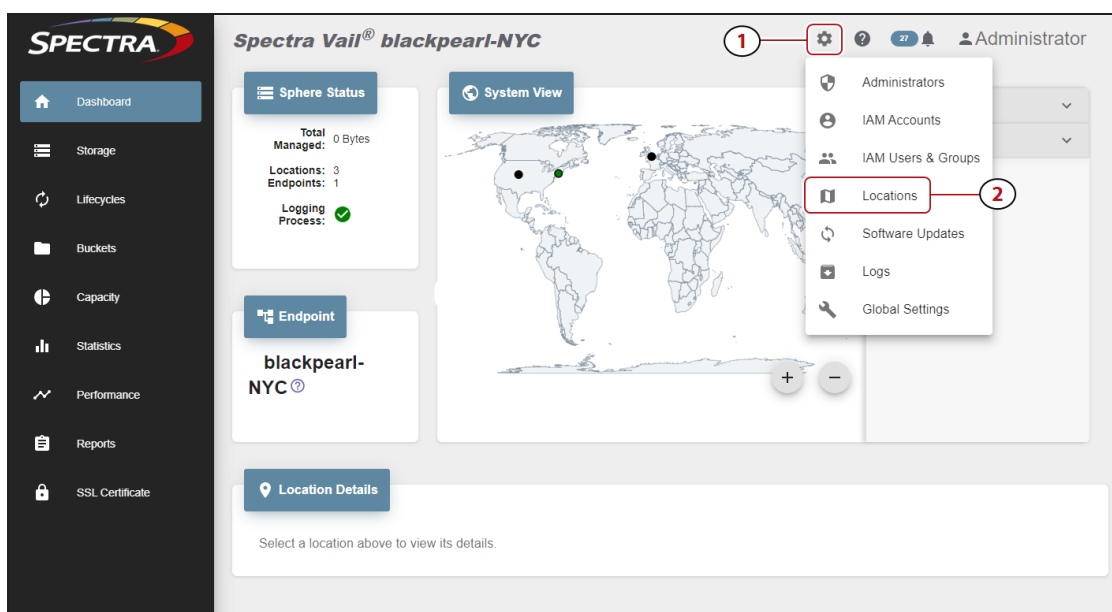


Figure 185 The Dashboard screen - Navigation menu.

2. Click **Create**.

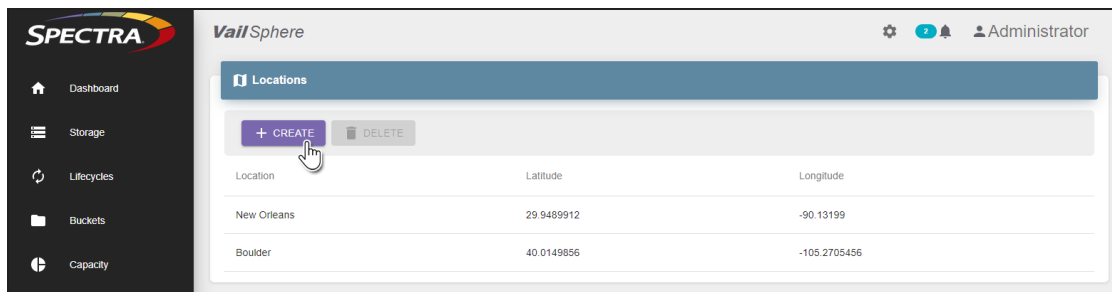


Figure 186 The Locations screen.

3. To map a location, you either search for the location, manually enter the latitude and longitude, or create a location with no corresponding geographic location.

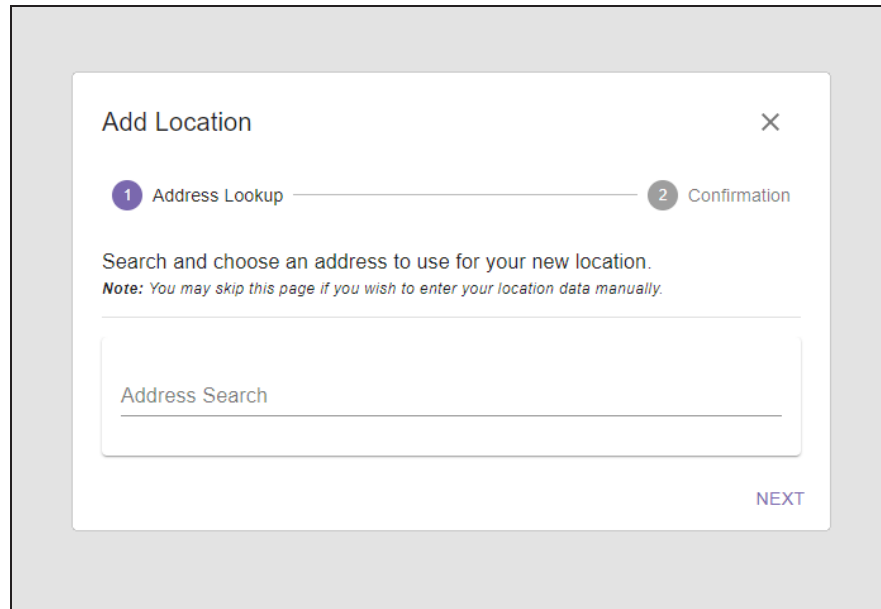


Figure 187 The Add Location - Address Lookup screen.

- To search for a location...
 - a. In the **Address Search** field, enter a geographic location. You can enter a full or partial postal address, city, county, province, or country. Select the correct match from the list and click **Next**.

Note: If no match is located, try changing the format of the address you entered. For example, use 9th Street in place of Ninth St.

 - b. Confirm the information is correct, or edit as needed, and click **Submit**.

- To manually enter a location...
 - a. Click **Next**.

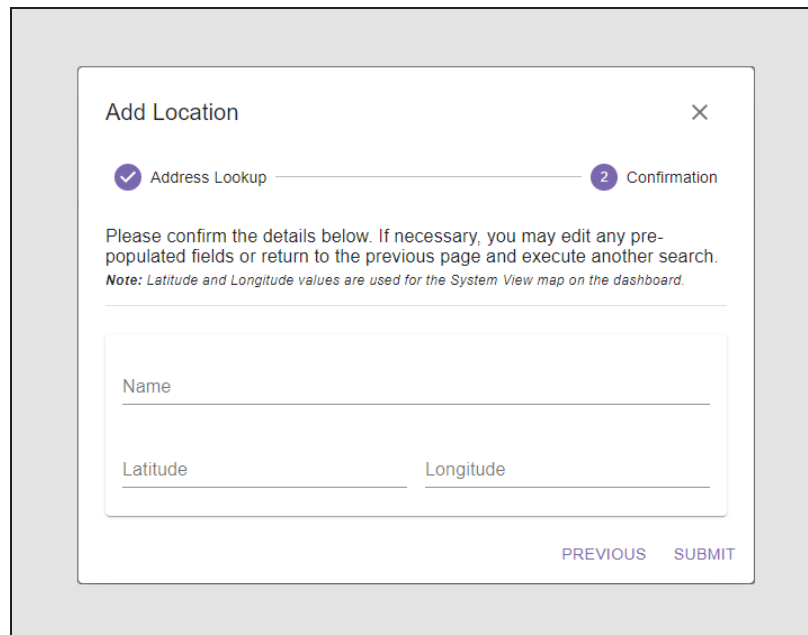


Figure 188 The Add Location - Manual Entry screen.

- a. Enter the desired **Name**.

Spectra Logic recommends naming each location after its physical location in the world.

For example, if Vail resources are located in Dallas, use that as the location name if there is only one Vail resource in that city. If there are multiple Vail resources consolidated in the same city, use suffixes to identify each group such as Dallas-HQ, Dallas-Research, or Dallas-Production.

- b. Enter the **Latitude** and **Longitude** of the location.

- Notes:**
- When entering a value for **Latitude**, use positive values for locations north of the equator, and negative values for locations south of the equator.
 - When entering a value for **Longitude**, use positive values for locations east of the prime meridian, and negative values for locations west of the prime meridian.

- c. Click **Submit**.

- To skip entering a location...
 - a. Click **Next**.
 - b. Enter the desired **Name** and click **Submit**.

The new location now displays on the world map on the Dashboard.

Note: If you do not enter an address or latitude and longitude, the location displays on the right-hand pane of the dashboard, but does not display on the world map.

DELETE A LOCATION

Locations are used to identify sites on the dashboard world map as well as to group storage endpoints by their physical location. If desired, you can delete a location that is no longer in use. Locations are only used in a cloud-control Vail sphere. These screens do not display in a local-control Vail sphere.

Here is how to delete a location:

1. In the upper right corner of the Vail management console, click the **gear icon (1)** and click **Locations (2)**.

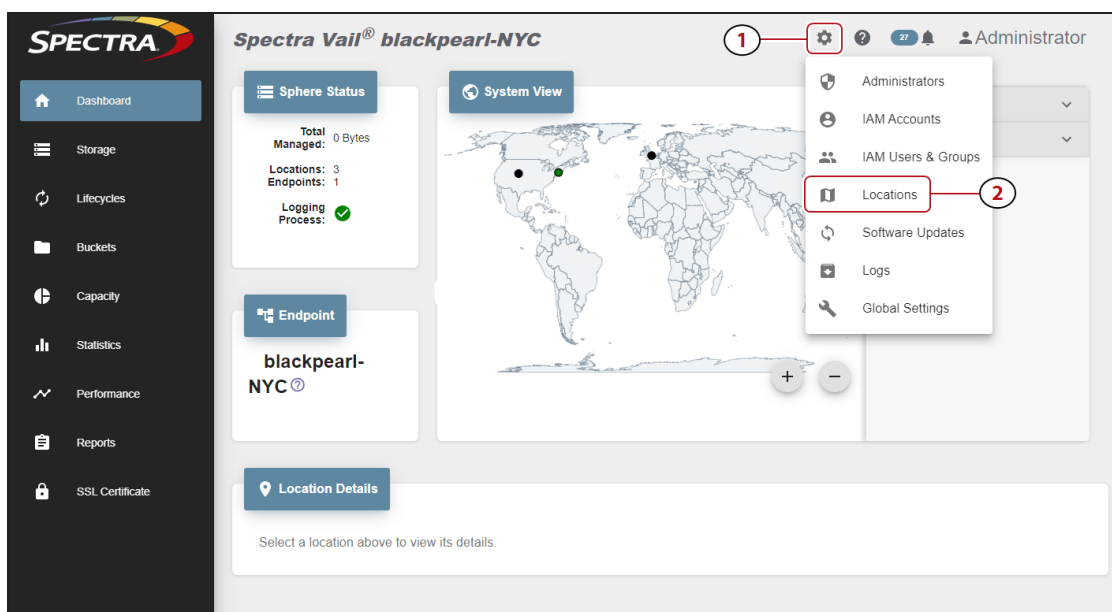


Figure 189 The Dashboard screen - Navigation menu.

2. Select the row of the location you want to delete and click **Delete**.
3. On the confirmation screen, click **Delete** to delete the location.

CLEAR THE IAM CACHE

The Spectra Vail application maintains an IAM (Identity and Access Management) cache independent from the cache maintained by Amazon Web Services. When users make IAM changes, AWS sends a notification to the Spectra Vail application, but the Vail management console may not update immediately. Clearing the Spectra Vail application IAM cache deletes the current information and causes the Spectra Vail application to retrieve all IAM information from AWS.

Additionally, clear the IAM Cache if you make security changes to or create a new set of IAM credentials in the AWS management console and want them to immediately display in the Vail management console.

Note: It may take several minutes for AWS security changes to take effect. Spectra Logic recommends waiting approximately 3-5 minutes after making changes before clearing the IAM cache, or updated settings may not display.

Here is how to clear the IAM cache:

1. In the upper right corner of the management console, click the **gear icon** and select **IAM Users & Groups**.
2. In the **IAM Users** pane, click **Clear Cache**.
3. In the confirmation window, click **Clear Cache**.

VIEW REPORTS

The Reports screen allows you to view any existing audit logs for the Spectra Vail application, and detailed information for each audit log.

- In the Vail management console taskbar, click **Reports**.

SPECTRA Vail Sphere Administrator

Reports

▼ Audit Log Show audit logs for a given period of time.

User Name Start Date 01/19/2021 End Date 01/26/2021

Description	User	Time
Endpoint SKISLQ4MY6MYX3JFL1T1 credentials were added		Jan 26, 2021 4:35:19 PM
Endpoint SKISOH3M5K5UES81SO5G credentials were added		Jan 26, 2021 4:33:55 PM
Endpoint SKISPO6CGZWBNENRZKQV credentials were added		Jan 26, 2021 3:40:19 PM
Endpoint SKISTUI2RJSLKZN7JXDN credentials were added		Jan 26, 2021 3:38:55 PM
Endpoint SKISG5FARZGYZSJ29RYW credentials were added		Jan 26, 2021 2:45:19 PM
Endpoint SKISTRKOOAFOAQDZH4IS credentials were added		Jan 26, 2021 2:43:55 PM
Endpoint SKISLXOIHOL6WIYODPJE credentials were added		Jan 26, 2021 1:50:19 PM
Endpoint SKIS8CVNKTJDFJKYSVK credentials were added		Jan 26, 2021 1:48:55 PM
Endpoint SKISSRGAZA2U9P5CEU43 credentials were added		Jan 26, 2021 12:55:19 PM
Endpoint SKIS7KUHGARJHFAVU7TT credentials were added		Jan 26, 2021 12:53:55 PM

Figure 190 The Reports screen.

- Use the **User Name**, **Start Date**, or **End Date** menus to refine the list of audit logs.

Note: Not all audit logs contain a User Name.

- Click the **View Details** icon on the right end of each audit log row to view details about the audit log.




Figure 191 The Audit Logs details screen.

Option	Description
Description	The description of the audit log.
User	The user associated with the log.
Request Path	The API path for the log.
Request Method	The method by which the log was generated.
Request Data	The contents of the log.
Request Timestamp	The time and date the log was generated.
Resource ID	The ID of the resource associated with the log.
Resource Name	The name of the resource associated with the log.
Host IP	The IP address of the Vail sphere.
Client IP	The IP address of the BlackPearl system or Vail VM node associated with the log.
Server	The name of the resource within the Vail sphere.

VIEW SPECTRA VAIL APPLICATION MESSAGES

Spectra Vail application messages provide important information about the status and current functionality of the Vail sphere. If desired, you can configured sphere administrators to receive messages automatically.

Note: The Spectra Vail application does not generate a message when an AWS cloud storage target is unavailable for backup operations. Some third-party applications may report this event as a warning message in their user interface.

Here is how to view messages:

In the upper right corner of the management console, click the **bell icon**. The value to the left of the icon indicate the number of unread messages.




Figure 192 The Bell icon.

The messages screen displays. Any unread messages are shown in bold font.

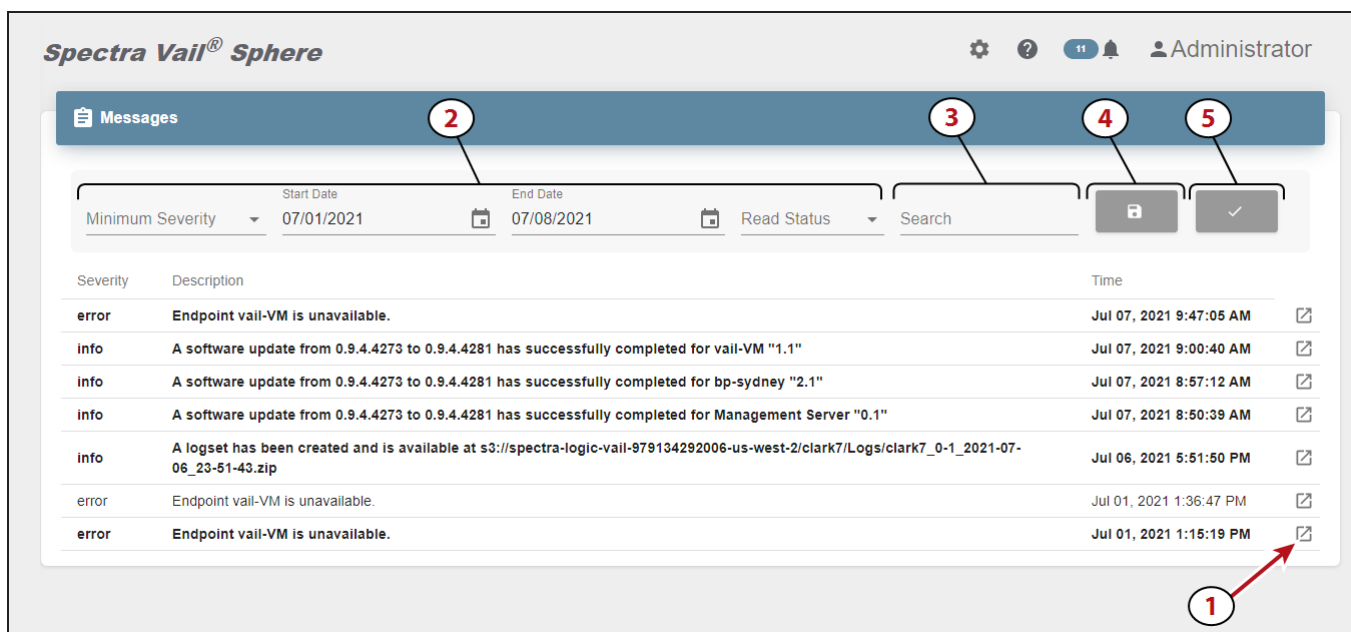


Figure 193 The Messages screen.

- To view message details, on the right end of the message row, click the **View Details** icon (1).
- You can sort messages using the **Minimum Severity**, **Start Date**, **End Date**, and **Read Status** drop-down menus (2).
- You can search messages for a text string by typing in the **Search** field (3).
- To download messages to your local host, in the upper-right corner of the Messages pane, click the **disk icon** (4).
- To mark all messages as read, in the upper-right corner of the Messages pane, click the **check mark icon** (5).

Message Details

In addition to the information on the Messages screen, the message details pane also displays the message key.

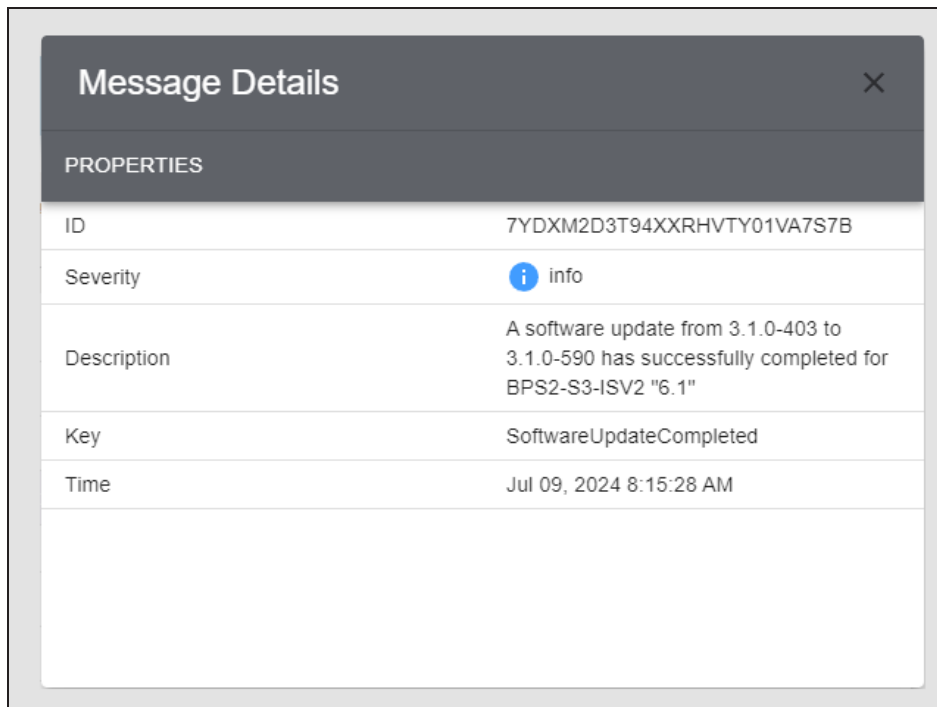


Figure 194 The Message Details screen.

Field	Description
ID	The ID value of the message.
Severity	<p>The severity of the message.</p> <p>Info - an event occurred such as a successful firmware update of the Vail sphere.</p> <p>Warning - An event that may affect data transfers occurred, such as the Vail sphere detects a down-level firmware version.</p> <p>Error - An event that prevents data transfers occurred, such as the nonavailability of a storage endpoint.</p>
Description	The message description.
Key	The message key. This value is useful when using the REST API to gather messages.
Time	The date and time the message was generated.

SPECTRA VAIL APPLICATION LOGS

Use the Logs page to generate and download logs for use in troubleshooting problems with the Vail sphere.

Note: If you delete the logs bucket in your AWS account, the bucket is recreated the next time you generate a log set in the Spectra Vail application.

In the upper right corner of the Vail management console, click the **gear icon** and select **Logs**.




Figure 195 The Logs screen.

- To generate a new logset, click **Create** and use the **Select Endpoint** drop-down menu to select the storage for which you want to generate a logset.
- To generate a new logset, click **Create**.
- To download an existing logset, select the row of the logset and click **Download**.
- To delete an existing logset, select the row of the logset and click **Delete**.

UPDATE THE SPECTRA VAIL APPLICATION SOFTWARE

Use the instructions below to update the Vail sphere software, and the software that storage endpoints use to communicate with the Vail sphere.

Each component must be updated separately, and each component update must be initiated manually. Components include the Vail sphere, BlackPearl Nearline Object Gateway and Vail VM nodes.

Note: The software running on the BlackPearl system is not updated using this process. See the [BlackPearl Nearline Gateway User Guide](#) for instructions on updating BlackPearl software.

In general, update the Spectra Vail application software in the following order:

- Vail sphere software
- BlackPearl software
- Vail VM Node software

Note: Spectra Logic recommends checking the Vail Release Notes for any changes to the update order that may be required for specific Vail release versions.

Here is how to update the Vail sphere or endpoint software:



IMPORTANT

The Spectra Vail application restarts after updating the application software. Any data transfer operations fail when the application restarts. Internal operations, such as Lifecycles, automatically restart. External operations must be manually restarted.

1. Discontinue storage operations. The Spectra Vail application restarts after updating the application software.
2. In the upper right corner of the management console, click the **gear icon** and select **Software Updates**.




Figure 196 The Software Updates screen.

To update using the online package sever...

- Select the row of the component you want to update and click **Update**.




Figure 197 The Update screen.

- Click **Update**. The update process for the selected component begins.



IMPORTANT

Do not reboot or power-cycle the BlackPearl Nearline Object Gateway or Vail VM node during the update process or the BlackPearl Nearline Object Gateway or Vail VM node fails to initialize.

Note: Depending on what component is being updated, the Vail management console may display a lost communication error while the component updates.

To update using a local file...

- a. Select the row of the component you want to update and click **Upload**.




Figure 198 The Upload screen.

- b. Click **Choose File**, then browse to the archive update file.
- c. Click **Upload**. The file is uploaded and the update process for the selected component begins.



IMPORTANT

Do not reboot or power-cycle the BlackPearl Nearline Object Gateway or Vail VM node during the update process or the BlackPearl Nearline Object Gateway or Vail VM node fails to initialize.

Note: Depending on what component is being updated, the Vail management console may display a lost communication error while the component updates.

ACCESSING THE TECHNICAL SUPPORT PORTAL

The Spectra Logic Technical Support portal provides access to the Knowledge Base, the current version of Vail software, and additional service and support tools. You can also open or update a support incident and upload log files.

Create an Account

Access to *User Guides* and compatibility matrices does not require you to create an account. You must create a user account and log in to access *Release Notes*, to download the latest version of Vail software, or to open a support incident.

Note: If you have multiple Spectra Logic products, the serial numbers for all products will be associated with your account. If you do not see the serial numbers for all of your products when you log in, contact Technical Support (see [Contacting Spectra Logic](#)).

1. Access the Technical Support portal login page at support.spectralogic.com.
2. On the home page, click **Register Now**.

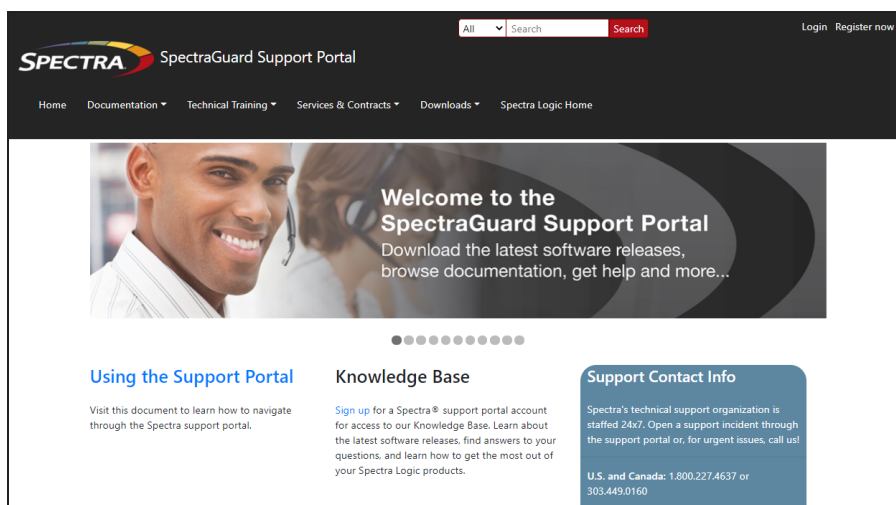


Figure 199 The Spectra Logic Technical Support portal home page.

3. Enter your registration information. Your account is automatically associated with the serial numbers of all Spectra Logic products owned by your site.
 - If you have an invitation, follow the link and enter the invitation code.

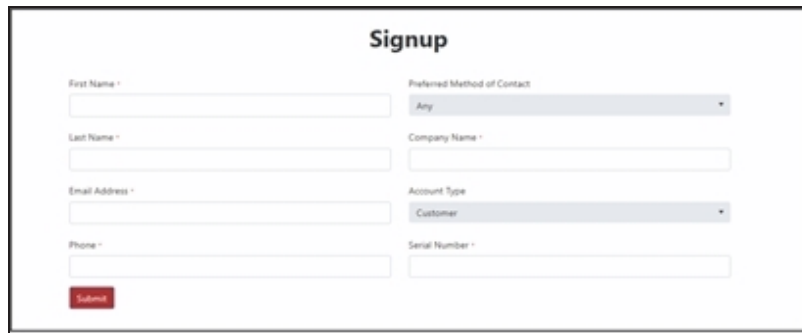
The image shows a web form titled "Signup". It contains two columns of input fields. The left column has "First Name", "Last Name", "Email Address", and "Phone". The right column has "Preferred Method of Contact" (a dropdown menu with "Any" selected), "Company Name", "Account Type" (a dropdown menu with "Customer" selected), and "Serial Number". A red "Submit" button is located at the bottom left of the form.

Figure 200 The Signup screen.

- If you do not have an invitation, enter the requested information to create your account. When you are finished, click **Submit**.

When the account is approved, you receive an email with an initial password. Use your email address and the password provided in the email to log in to your account. After you log in, you can change your password if desired.

Log Into the Portal

1. Access the Technical Support portal login page at support.spectralogic.com.
2. Use your email address and password to log into the Technical Support Portal.

OPENING A SUPPORT TICKET

You can open a support incident using the Spectra Logic Technical Support portal or telephone.

Search for Help Online




Figure 201 The Spectra Logic Technical Support portal home page.

1. Make notes about the problem, including what happened just before the problem occurred.
2. Gather the following information:
 - Your Spectra Logic customer number
 - Company name, contact name, phone number, and email address
 - The library serial number on the **Configuration>Settings** screen.
 - Type of host system being used
 - Type and version of host operating system being used
 - Type and version of host storage management software being used
3. If necessary, log in to the Support Portal by clicking **Login**, enter your **email address** and **password**, and click **Log in**.

Note: See [Accessing the Technical Support Portal](#) on page 243 if you have not previously created an account on the Technical Support portal.
4. From any page, select **Incident>Incidents & Inventory**.




Figure 202 Select **Incidents>Incidents & Inventory**.

5. Select **Open or View Incidents**.

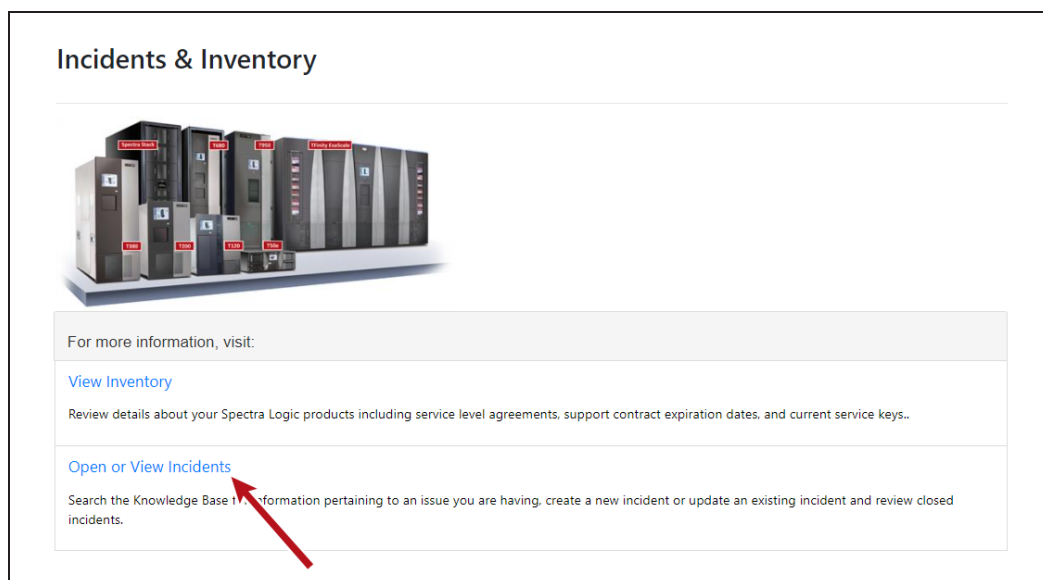


Figure 203 Select **Open or View Incidents**.

6. In the Search dialog box, enter a term or phrase about your problem (1) and click **Search** (2).

Open or View Incidents

What can we help you with?

Search Search

Serial Number: Status: Created Date:

Product: Active mm/dd/yyyy

☒ Export to Excel

Incident N...	Product	Account	Date Open...	Subject	Description	Modified On	Problem R...
294900	TLB0421021 - Spectra T120 Library	Spectra Logic Corporation - Boulder	03-05-2021	Open by ASL parser	This is a test - please disregard. Subject = Manual T120 AutoSupport Log Set for Library TLB0421021 on 03-05-2021 11:58:34	4/20/2021 5:05 PM	

Figure 204 Enter a search phrase and click **Search**.

7. If the search does not provide an answer, click **Open a New Incident**.

Open or View Incidents

What can we help you with?

export Search

search results for export Sort by Relevance

[T120 Import/Export Modes](#) 20200709194431340

T120 Import/Export Modes, by the operator's selection. **Exporting** tapes When a tape is selected by the host software to be **exported**. Moves" button. **Exporting** tapes Like the Shared mode, all **exports** initiated by the host software, is process for importing and **exporting** tapes, to and from the library. The different modes also, available for single partition libraries. The 8 EE slots in the EE port are available for import and **export**, port, are physical moves. As a result, the number of tapes to be imported or **exported** in a single

[Operators cannot import or export media](#) 20201130201216990

User with operator privileges cannot import or export tapes to a partition's storage slots.

[Spectra TFinity Library User Guide](#) 20200827222441320

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. . In particular, but without limitation, the Software Product may not be **exported** or re-**exported**. Modules (LCM and RCM) 50 **Exporting** Controllers, 230 IMPORT AND **EXPORT**

[Next](#)

[Open a New Incident](#)

Figure 205 Click **Open a New Incident**.

8. On the Create Incident page, enter the requested information providing as much detail as possible. When you are finished, click **Submit**.

Create Incident

Severity *

Problem Description *

Email addresses to include in correspondence

Customer *

Product *

Select files...

DELIVERY Address For Shipping Parts

☐ Confirm The Ship To Address

Submit

Figure 206 Enter information about your incident and click **Submit**.

Submit an Incident Online

1. Make notes about the problem, including what happened just before the problem occurred.
2. Gather the following information:
 - Your Spectra Logic customer number
 - Company name, contact name, phone number, and email address
 - The library serial number on the **Configuration>Settings** screen.
 - Type of host system being used
 - Type and version of host operating system being used
 - Type and version of host storage management software being used
3. If necessary, log in to the Support Portal by clicking **Login**, enter your **email address** and **password**, and click **Log in**.

Note: See [Accessing the Technical Support Portal on page 243](#) if you have not previously created an account on the Technical Support portal.

4. From any page, select **Inventory>My Inventory**.

5. Locate the row of the product for which you want to submit an incident and click **Create Incident**.

View Inventory

Click on the view icon for additional inventory details where you can update product nickname, firmware version, operating and software systems.

To edit column filters, click to the right of the column name. If your serial number is not listed below, click [here](#).

Find by Model, Serial # or Account Search

	Prod... ⌵	Product Ni... ⌵	Account ⌵	SLA ⌵	ASM ⌵	Supp... ⌵	Servi... ⌵	Action ⌵
<input checked="" type="checkbox"/>	0906802 - Spectra T680 Library	Fishbowl	Spectra Logic Corporation - Boulder	Next Business Day Replacem...	None	31/12/2050	3BB 3HN BB7 DNB 2AZ	<div style="background-color: #007bff; color: white; padding: 2px 5px; border-radius: 3px;">Renew Contract</div> <div style="background-color: #28a745; color: white; padding: 2px 5px; border-radius: 3px;">Create Incident</div>
<input checked="" type="checkbox"/>	1724A05 - Spectra TFinity Library	Training Room	Spectra Logic Corporation - Boulder	Next Business Day Replacem...	None	31/12/2050	WXY YCG L4X TT4 HVS	<div style="background-color: #007bff; color: white; padding: 2px 5px; border-radius: 3px;">Renew Contract</div> <div style="background-color: #28a745; color: white; padding: 2px 5px; border-radius: 3px;">Create Incident</div>
<input checked="" type="checkbox"/>	1311A06 - Spectra TFinity Library		Spectra Logic Corporation - Boulder	Next Business Day Replacem...	None	31/12/2050	FTJ 4DV ZLC YHB Z6B	<div style="background-color: #007bff; color: white; padding: 2px 5px; border-radius: 3px;">Renew Contract</div> <div style="background-color: #28a745; color: white; padding: 2px 5px; border-radius: 3px;">Create Incident</div>

Figure 207 Click **Create Incident**.

6. On the Create Incident page, enter the requested information providing as much detail as possible. When you are finished, click **Submit**.

Create Incident

Severity *

Problem Description *

Email addresses to include in correspondence

Customer *

Spectra Logic Corporation - Boulder

Product *

0906802 - Spectra T680 Library

Select files...

DELIVERY Address For Shipping Parts

6101 Lookout Rd, Boulder, CO 80301-3580 UNITED STATES

☐ Confirm The Ship To Address

Submit

Figure 208 Enter information about your incident and click **Submit**.

Submit an Incident by Phone

Contact Spectra Logic Technical Support by phone using the information below.

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
Additional international numbers available at support.spectralogic.com/home If you have a Spectra Logic Portal account, please log in for country-specific numbers at support.spectralogic.com/support-contact-info	

APPENDIX A - BLACKPEARL EMBEDDED DASHBOARD

This chapter describes the use of the BlackPearl Embedded Dashboard in the Vail application. The embedded dashboard allows you to view information about each BlackPearl system configured in the Vail sphere. It also allows you to easily configure and manage commonly used functions of the BlackPearl system.

Using the Embedded BlackPearl Dashboard	252
View the Status of the BlackPearl System	253
View System Overview	253
View Notifications	254
View Jobs	255
View Buckets	256
View Pools	257
View Volumes	258
View Tape Partitions - Main View	259
View Tape Partitions - Tape State View	260
View Tape Drives	261
View Tape Management	262
Dashboard Actions	263
Create a Volume Snapshot	263
Export a Tape Cartridge	263
Online a Tape Cartridge	264
Verify a Tape Cartridge	264
Change Job Priority	265
Create a Bucket	265
Start a Storage Pool Verification	266
Put a Tape Partition into Standby	266
Offline a Tape Drive	266

USING THE EMBEDDED BLACKPEARL DASHBOARD

The embedded BlackPearl dashboard allows you to quickly view the status of critical aspects of a BlackPearl system in the Vail sphere, and easily perform commonly used functions of the system.

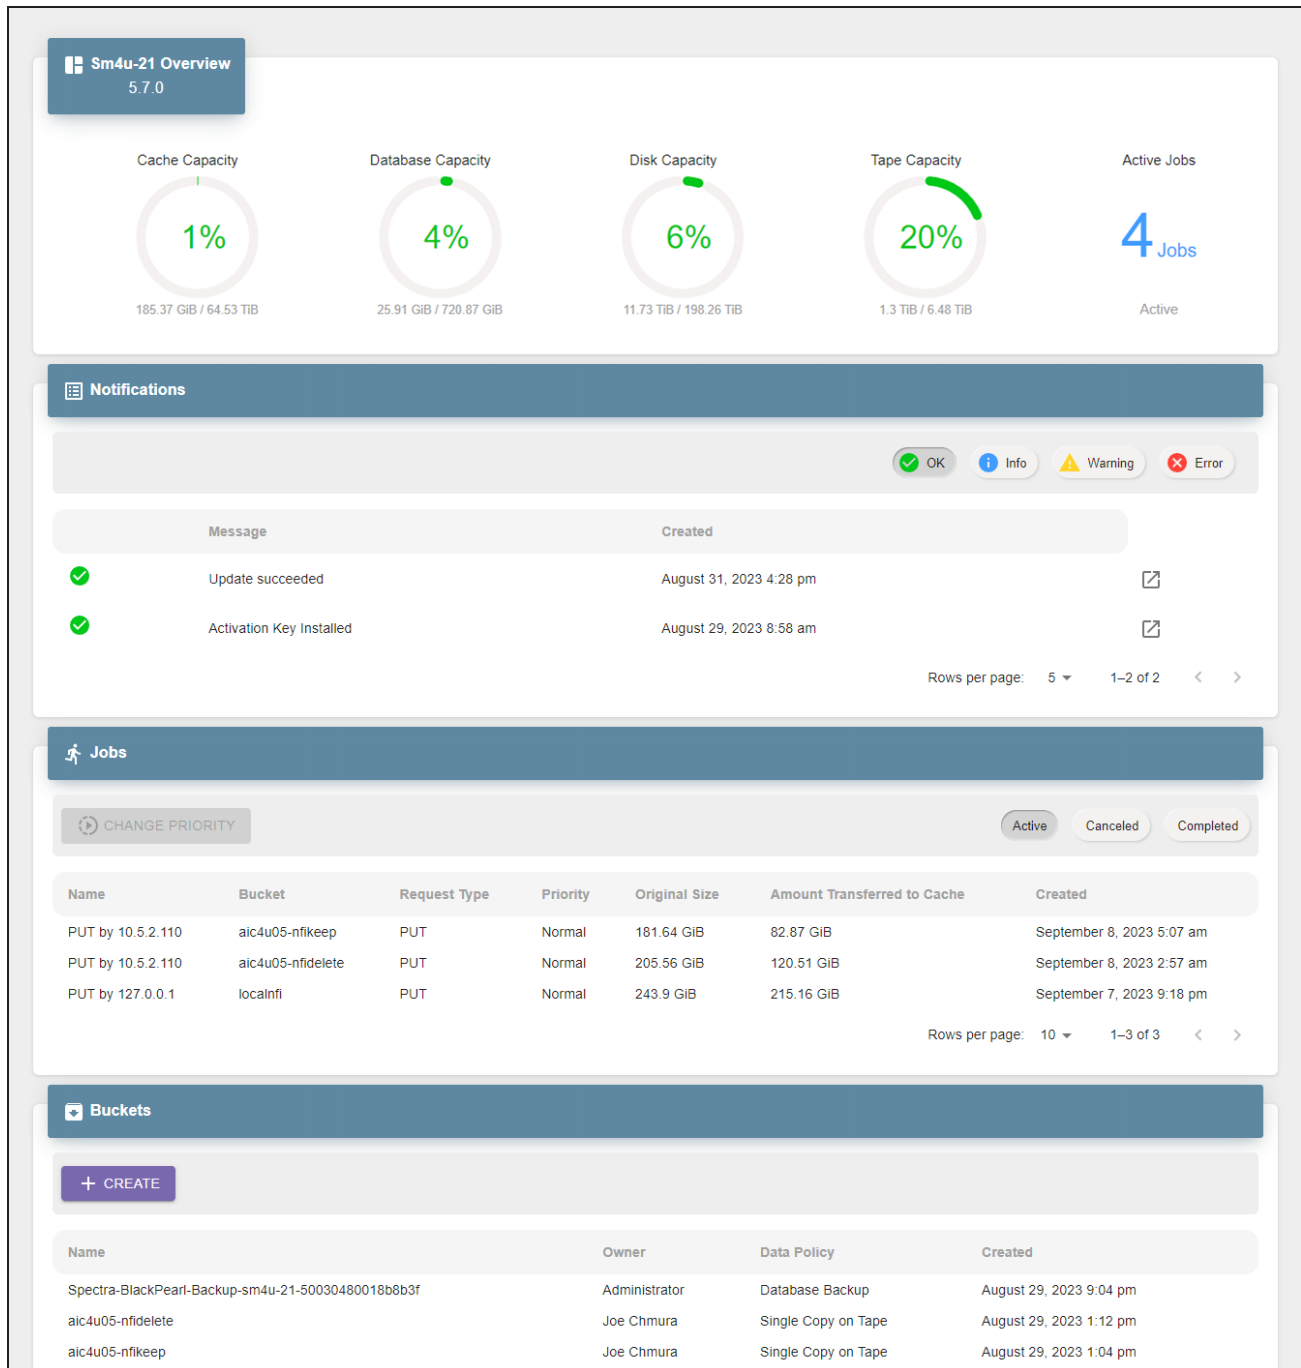


Figure 209 The Embedded Dashboard.

VIEW THE STATUS OF THE BLACKPEARL SYSTEM

Use the sections below to view the status of multiple aspects of the BlackPearl system.

View System Overview

The Overview pane provides a quick look at the most critical aspects of the BlackPearl system.

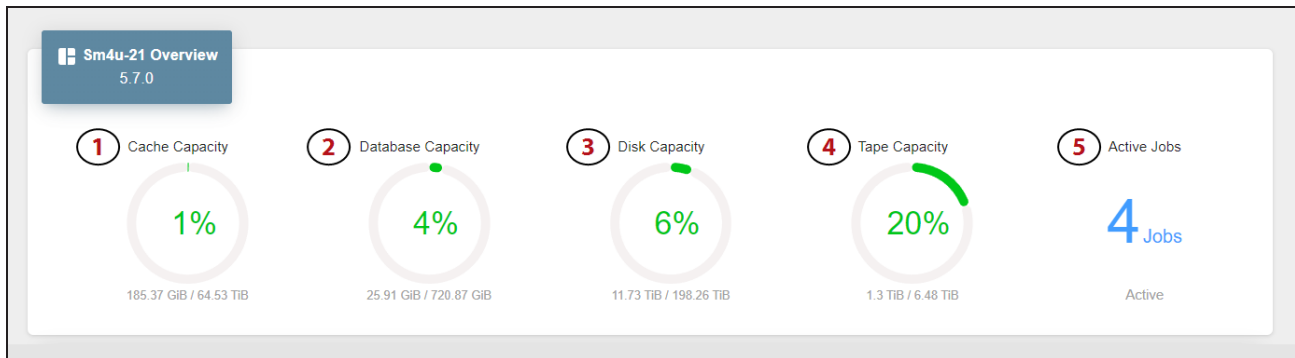


Figure 210 The Overview pane.

1. The BlackPearl system cache capacity and percentage of used cache space.
2. The capacity of the BlackPearl system database and percentage of used space.
3. The capacity of all disk-based storage connected to the BlackPearl system and percentage of used space.
4. The capacity of all tape-based storage in the tape library connected to the BlackPearl system and percentage of used space.
5. The number of active jobs running on the BlackPearl system.

Mouse-over the green section of any percentage graph to display the amount of used space, and over the gray section to display the amount of remaining space.




Figure 211 Mouse-over a graph to view specific details.

View Notifications

Notifications provide information about errors that occur on the system, caution messages that alert you to issues that may impact your workflow, and informational messages. Additionally, notifications may provide troubleshooting advice to help you resolve issues that may occur.

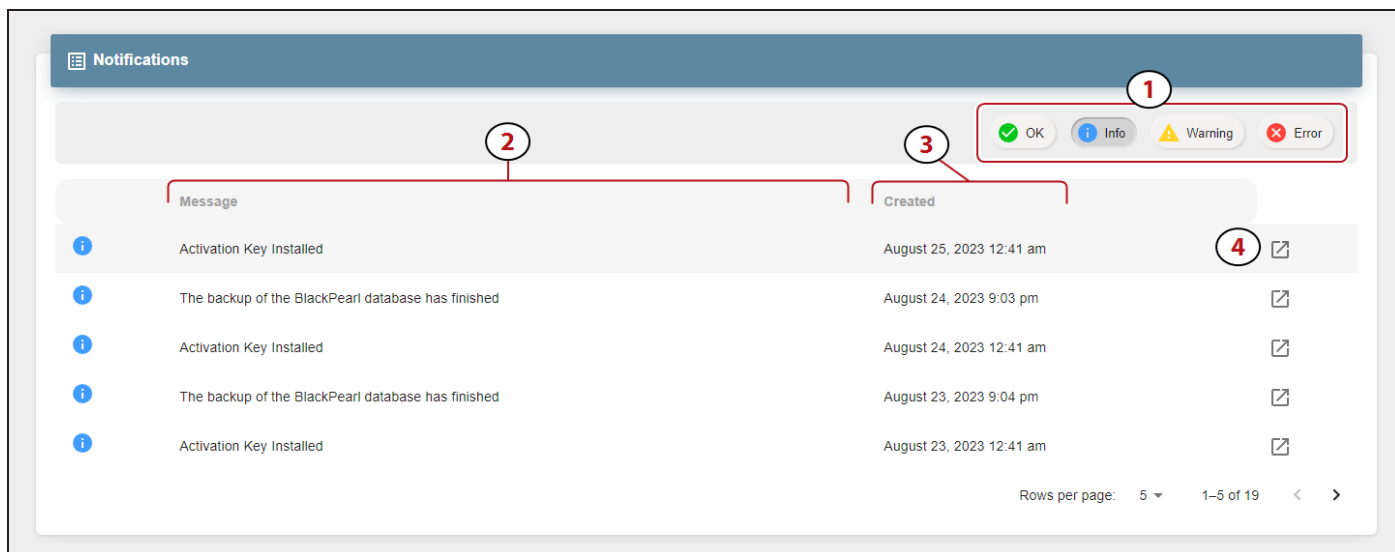


Figure 212 The Notifications pane.

1. Use the **Notification Type** buttons to switch between OK, Info, Warning, and Error messages.
2. Contains a brief description of the notification.
3. Displays the timestamp the notification was generated.
4. Click the **Details Button** to view additional message **Details** and **Troubleshooting Advice**.

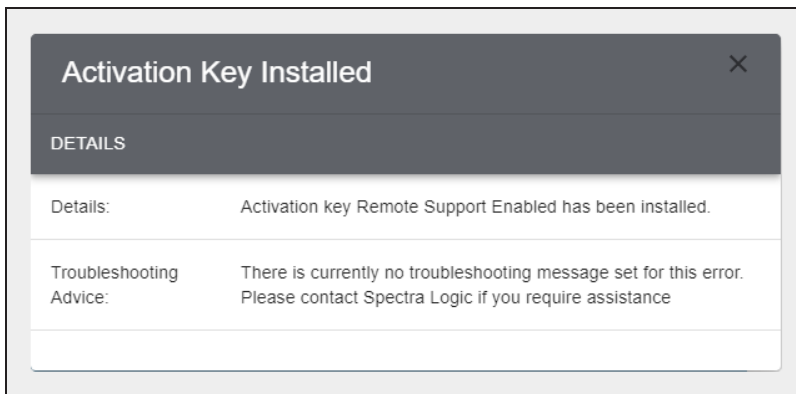


Figure 213 The Notification details dialog box.

View Jobs

The Jobs pane provides information on each Active, Canceled, or Completed job processed by the BlackPearl system.

2	3	4	5	6	6	7
Name	Bucket	Request Type	Priority	Original Size	Amount Transferred to Cache	Created
PUT by 127.0.0.1	Spectra-BlackPearl-Backup-sm4u-21-50030480018b8b3f	PUT	Normal	6.61 MiB	6.61 MiB	August 24, 2023 9:00 pm
PUT by 127.0.0.1	Spectra-BlackPearl-Backup-sm4u-21-50030480018b8b3f	PUT	Normal	5.25 MiB	5.25 MiB	August 23, 2023 9:00 pm
PUT by 127.0.0.1	Spectra-BlackPearl-Backup-sm4u-21-50030480018b8b3f	PUT	Normal	6.1 MiB	6.1 MiB	August 22, 2023 9:00 pm
PUT by 10.5.2.130	g	PUT	Normal	1000 MiB	1000 MiB	August 22, 2023 12:38 pm
PUT by 127.0.0.1	Spectra-BlackPearl-Backup-sm4u-21-50030480018b8b3f	PUT	Normal	5.4 MiB	5.4 MiB	August 21, 2023 9:00 pm
PUT by 10.5.2.130	SpectraApp1	PUT	Normal	4.35 MiB	4.35 MiB	August 21, 2023 2:46 pm
PUT by 10.5.2.130	d	PUT	Normal	3 GiB	3 GiB	August 21, 2023 2:46 pm
PUT by 10.5.2.130	e	PUT	Normal	25 MiB	25 MiB	August 21, 2023 2:46 pm
PUT by 10.5.2.130	SpectraApp2	PUT	Normal	300 MiB	300 MiB	August 21, 2023 2:46 pm
PUT by 10.5.2.130	c	PUT	Normal	75 MiB	75 MiB	August 21, 2023 2:46 pm

Rows per page: 10 1-10 of 15 < >

Figure 214 The Jobs pane.

1. Use the **Job Type** buttons to switch between Active, Canceled, and Completed jobs.
2. The name of the job includes the job type and the IP address of the job initiator.
3. The bucket used in the PUT or GET operation.
4. The type of job request.
5. The assigned priority of the job.
6. The original size and amount of data transferred to the BlackPearl system cache.
7. Displays the timestamp of when the job was initiated.

Use the **Change Priority** button to change the priority of an active job. See [Change Job Priority on page 265](#) for more information.

View Buckets

The Buckets pane provides information about all buckets configured on the BlackPearl system.

Name	Owner	Data Policy	Created
Spectra-BlackPearl-Backup-sm4u-21-50030480018b8b3f	Administrator	Database Backup	August 21, 2023 9:00 pm
SpectraApp	SpectraApp	Single Copy on Tape	August 21, 2023 2:33 pm
SpectraApp1	SpectraApp	Single Copy on Tape	August 21, 2023 2:45 pm
SpectraApp2	SpectraApp	Single Copy on Tape	August 21, 2023 2:45 pm
a	SpectraApp	Single Copy on Tape	August 21, 2023 2:45 pm
b	SpectraApp	Single Copy on Tape	August 21, 2023 2:45 pm
c	SpectraApp	Single Copy on Tape	August 21, 2023 2:45 pm
d	SpectraApp	Single Copy on Tape	August 21, 2023 2:45 pm
e	SpectraApp	Single Copy on Tape	August 21, 2023 2:45 pm
f	SpectraApp	Single Copy on Tape	August 21, 2023 2:46 pm

Rows per page: 10 1-10 of 12 < >

Figure 215 The Buckets pane.

1. Displays the name of the bucket.
2. The bucket owner configured on the BlackPearl system.
3. The data policy used by the bucket.
4. Displays the timestamp of when the bucket was created.

Use the **Create** button to create a new bucket. See [Create a Bucket](#) on page 265 for instructions.

View Pools

The Pools pane displays information about all disk storage pools configured on the BlackPearl system including dedicated BlackPearl system cache and database pools.

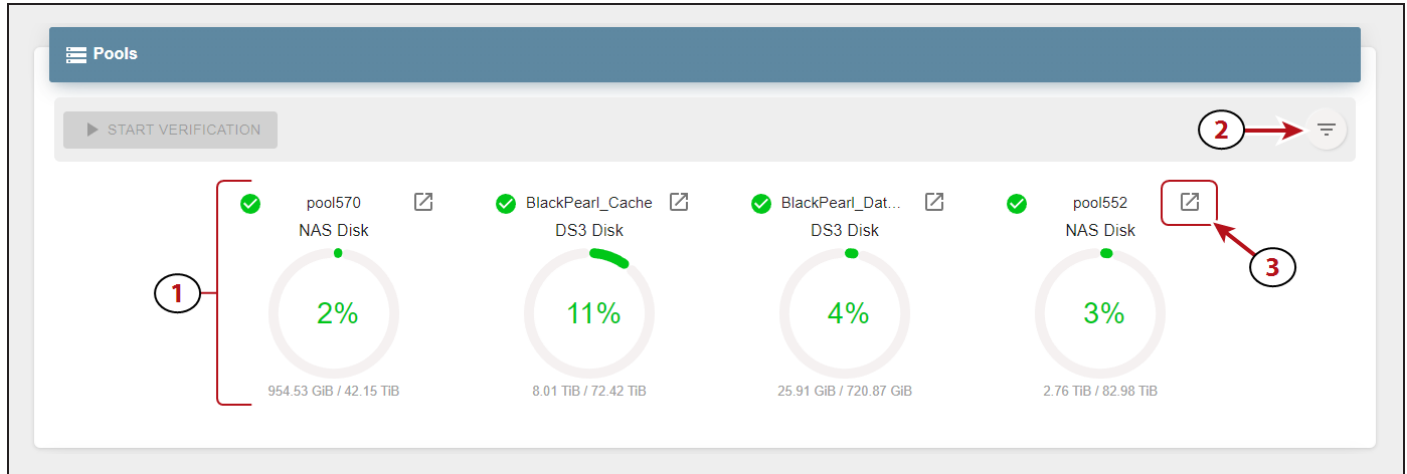


Figure 216 The Pools pane.

1. Each percentage graph displays both the used and remaining space for the associated pool.
2. Use the **Filter** button to select which pools to display on the Pools pane.
3. Click the **Details** button to view additional information about a specified pool.




Figure 217 The pool details dialog box.

Use the **Start Verification** button to verify the data contained on the pool. See [Start a Storage Pool Verification](#) on page 266 for more information.

View Volumes

The Volumes pane displays information about all volumes configured on the BlackPearl system.

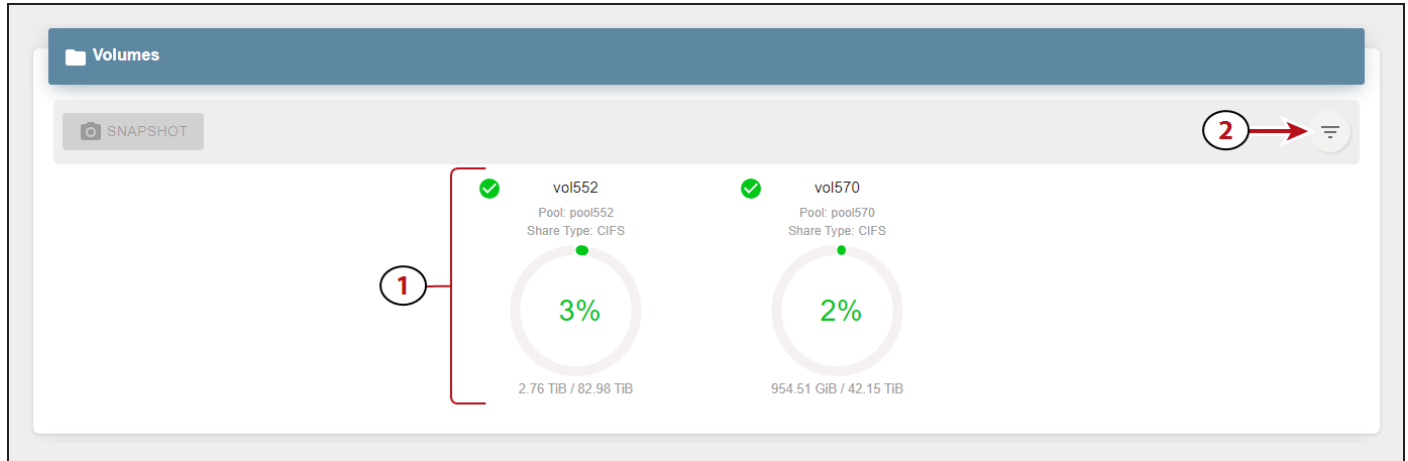


Figure 218 The Volumes pane.

1. Each percentage graph displays both the used and remaining space for the associated pool.
2. Use the **filter button** to select which pools to display on the Pools pane.

Use the **Snapshot** button to create a snapshot. For more information see [Create a Volume Snapshot](#) on page 263.

View Tape Partitions - Main View

The Tape Partitions pane displays information about the tape partitions configured on the tape library attached to the BlackPearl system. The Tape Partitions pane features both a main view and a tape cartridge state view.

To display the main view, manipulate the slider (2) to the left position.

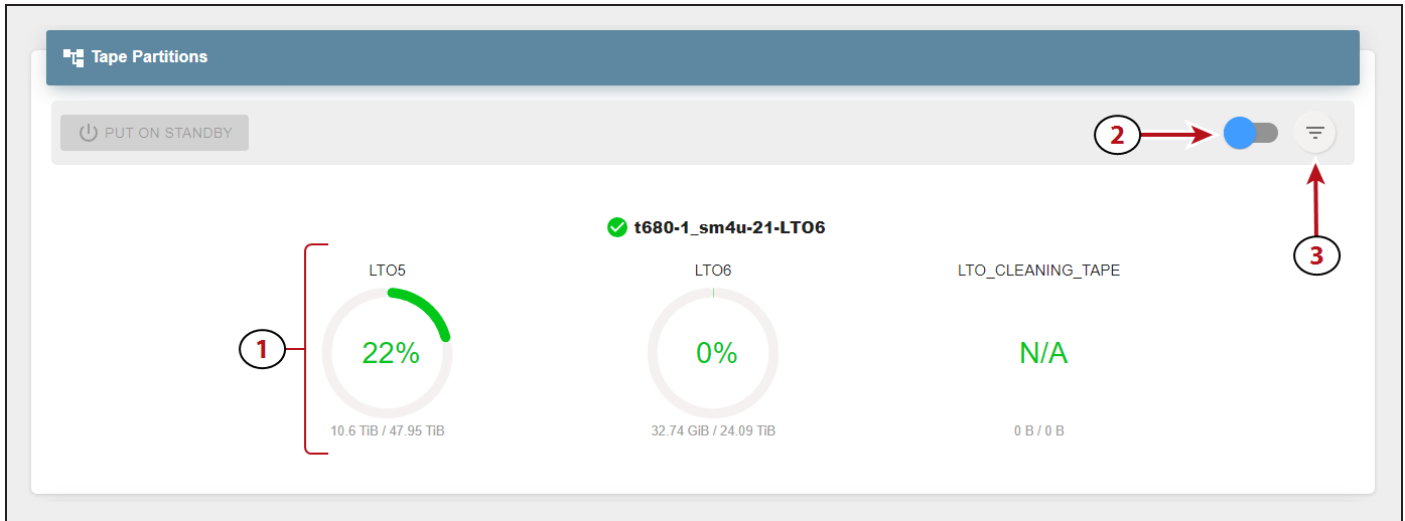


Figure 219 The Tape Partitions pane - main view.

1. Each percentage graph displays both the used and remaining space for the associated type and generation of media present in the tape partition. Mouse-over the green section of any percentage graph to display the amount of used space, and over the gray section to display the amount of remaining space.
2. Use the slider to change the display a graph of the current state of each tape cartridge present in the partition.
3. Use the **Filter** button to select which pools to display on the Tape Partitions pane.

If you need to service the tape library, you can put a tape partition into a standby state. See [Put a Tape Partition into Standby on page 266](#) for more information.

View Tape Partitions - Tape State View

The Tape Partitions pane displays information about the tape partitions configured on the tape library attached to the BlackPearl system. The Tape Partitions pane features both a main view and a tape cartridge state view.

To display the tape cartridge state view, manipulate the slider (2) to the right position.

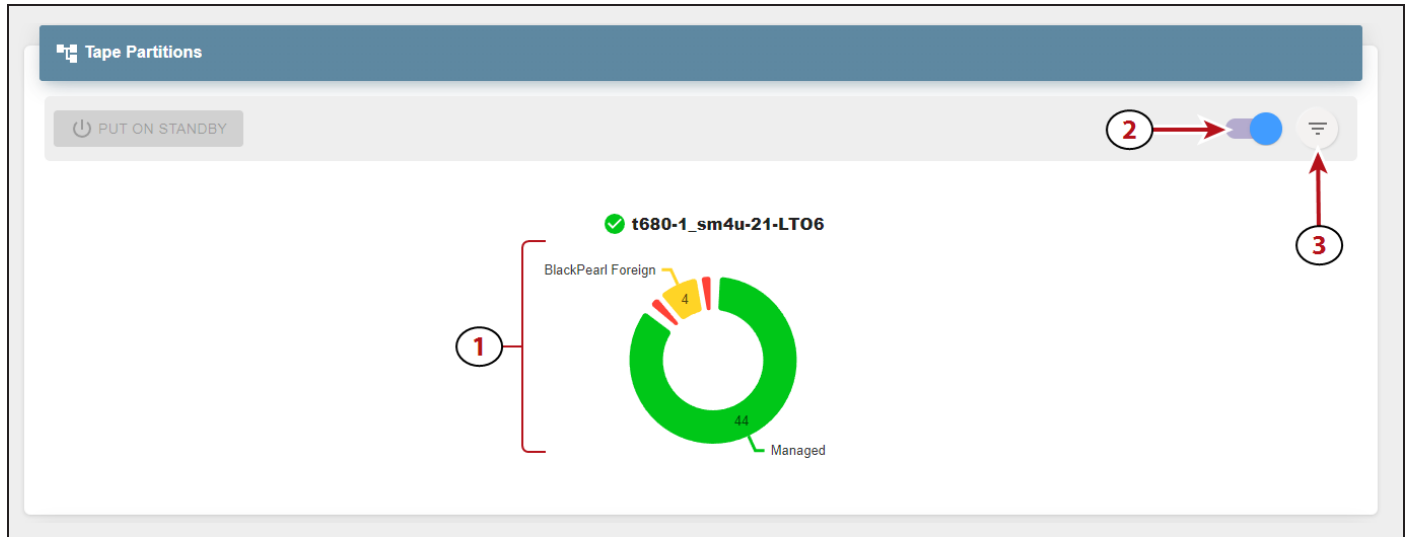


Figure 220 The Tape Partitions pane - main view.

1. The state of all tape cartridges in the partition. Each state combines different generations of tape media if present. Mouse-over any part of the graph to display more detailed information.
2. Use the slider to change the display a graph of the current state of each tape cartridge present in the partition.
3. Use the **Filter** button to select which pools to display on the Tape Partitions pane.

If you need to service the tape library, you can put a tape partition into a standby state. See [Put a Tape Partition into Standby on page 266](#) for more information.

View Tape Drives

The Tape Drives pane displays information about all tape drives installed in the tape library connected to the BlackPearl system.

1 Status	2 Type	3 Serial Number	4 Tape Barcode	5 Current Task	6 Cleaning Required	7 Online	8 Reserved Task Type
Normal	LTO6	1023003646	519815L5	WriteChunkToTapeTask	No	Yes	ANY
Normal	LTO6	1024003646	503887L5	WriteChunkToTapeTask	No	Yes	ANY

Figure 221 The Tape Drives pane.

1. The current status of the tape drive.
2. The drive type and generation.
3. The drive serial number as assigned by the tape library.
4. The physical barcode of the tape cartridge loaded into the tape drive. This field is blank when no tape is loaded.
5. The current task being performed by the drive. This field is blank when no task is in progress.
6. Indicates if the tape drive requires cleaning.
7. Indicates if the tape drive is online or offline.
8. The reserved task type, if configured. The default setting is Any.

Use the **Take Offline** button to take the drive offline. See [Offline a Tape Drive on page 266](#) for more information.

View Tape Management

The Tape Management pane displays the status of all managed tapes in the tape library connected to the BlackPearl gateway.

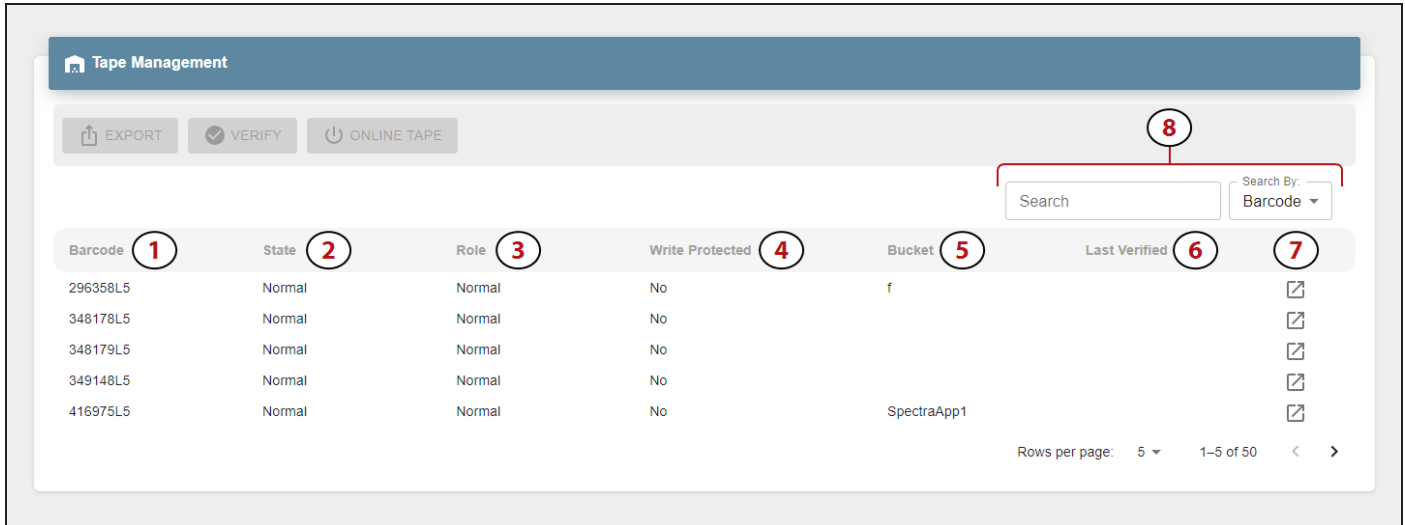


Figure 222 The Tape Management pane.

1. The physical barcode label on the tape cartridge.
2. The current state of the tape cartridge.
3. Indicates if the tape is configured for use as a **Normal** or **Test** tape.
4. The physical **Write Protected** status of the tape cartridge.
5. The name of any BlackPearl system bucket(s) present on the tape cartridge.
6. Displays the timestamp of the last tape verification.
7. Click the **Details** button to display additional information about the selected tape cartridge.
8. Use the **Search** entry field and **Search By** drop-down menu to find a specific tape cartridge.

See one of the following sections for instructions to export, verify, or online a tape cartridge:

- [Export a Tape Cartridge on the next page](#)
- [Verify a Tape Cartridge on page 264](#)
- [Online a Tape Cartridge on page 264](#)

DASHBOARD ACTIONS

In addition to displaying information about the BlackPearl system, the embedded dashboard allows you to perform the most frequently-used actions as described in the sections below.

Create a Volume Snapshot

A volume snapshot is an image of a volume's configuration and data makeup as they were when the snapshot was generated. Restoring to a previously created snapshot allows you to go “back in time” and restore the volume to the state it was in when the snapshot was created.

See [Volume Snapshots on page 1](#) for more information.

Here is how to create a volume snapshot:

1. In the BlackPearl dashboard, navigate to the **Volumes** pane.
2. **Select** the volume for which you want to create a snapshot.
3. Click **Snapshot**.
4. If desired, edit the pre-generated **Snapshot** name.

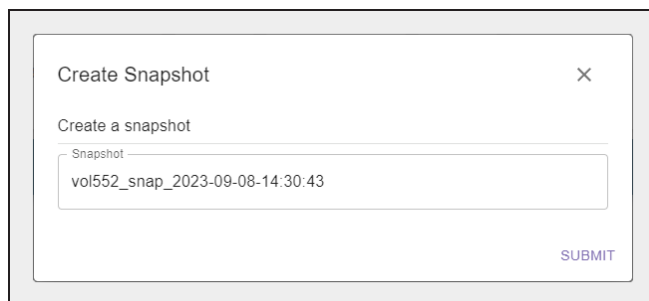


Figure 223 The Export Tape dialog box.

5. Click **Submit**.

Export a Tape Cartridge

Exporting a tape cartridge prepares it for physical removal from the attached tape library. In a Spectra Logic tape library, the cartridge is moved from the storage pool to the Entry/Exit pool, before it is physically exported from the library at the library front panel.

1. In the BlackPearl system dashboard, navigate to the **Tape Management** pane.
2. **Select** the tape you want to export.
3. Click **Export**.

4. If desired, edit the **Export Label** and **Export Location**.

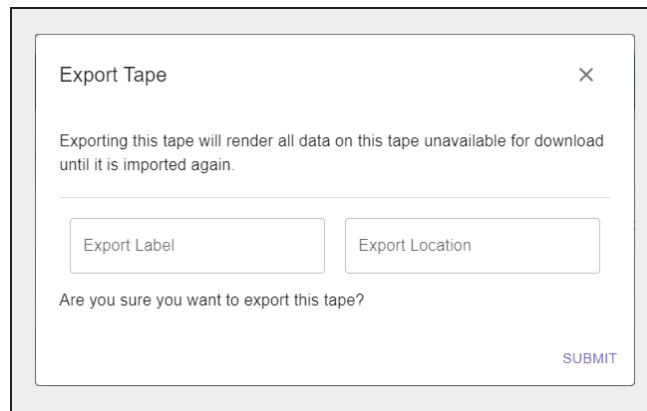


Figure 224 The Export Tape dialog box.

5. Click **Submit**.

Online a Tape Cartridge

Setting a tape cartridge to "online" prepares the cartridge for use by the BlackPearl system. This allows the system to use the tape cartridge for data storage operations.

Here is how to online a tape cartridge:

1. In the BlackPearl system dashboard, navigate to **Tape Management**.
2. Select a tape in the **Offline** state.
3. Click **Online Tape**.
4. Click **Submit**.

Verify a Tape Cartridge

The BlackPearl system can perform a data integrity verification of all data on a selected tape cartridge to confirm it is still viable. While the verification is in progress, client access has priority over the data integrity scan.

Here is how to verify a tape cartridge:

1. In the BlackPearl system dashboard, navigate to **Tape Management**.
2. **Select** the tape you want to verify.
3. Click **Verify Tape**.
4. Click **Submit**.

Change Job Priority

If desired, you can change the priority of an active job on the BlackPearl system.

Here is how you change the priority of a job:

1. In the BlackPearl system dashboard, navigate to the **Jobs** pane.
2. If necessary, click **Active** to display the list of active jobs.
3. **Select** the job for which you want to change priority and click **Change Priority**.
4. Use the **drop-down** menu to select a new priority for the job.
5. Click **Submit**.

Create a Bucket

Buckets on the BlackPearl system are data transfer targets for read and write operations. When you create a new bucket on the system, you assign it a owner and a data policy. You can then use the new bucket in your other Spectra software applications as a target for data storage on the BlackPearl system.

Here is how you create a new bucket:

1. In the BlackPearl system dashboard, navigate to the **Buckets** pane.
2. Click **Create**.
3. Enter a **Bucket Name**.

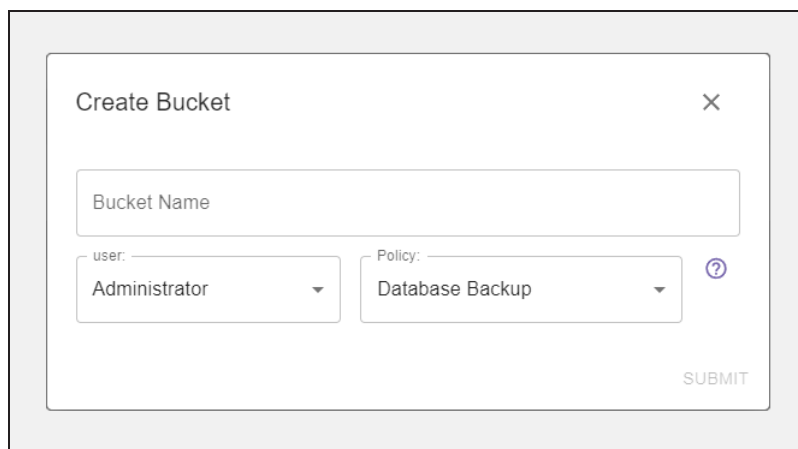


Figure 225 The Create Bucket dialog box.

4. Using the **User** drop-down menu, select an owner for the bucket.
5. Using the **Policy** drop-down menu, select a data policy for the bucket.
6. Click **Submit**.

Start a Storage Pool Verification

The BlackPearl system can perform a data integrity verification of all data on a selected storage pool to confirm it is still viable.

Here is how to start data verification on a storage pool:

1. In the BlackPearl system dashboard, navigate to the **Pools** pane.
2. **Select** the pool that you want to verify.
3. Click **Start Verification**.
4. Click **Submit**.

Put a Tape Partition into Standby

If you need to perform service on the tape library associated with your BlackPearl gateway, or with the BlackPearl gateway itself, you must first put the tape library into a standby state. Otherwise, the BlackPearl gateway may attempt to use the tape library while it is in service.

Note: After the tape partition is placed in standby, any currently running tape operations continue until they are complete, which may take 30 minutes or longer.

Here is how to out a tape partition into standby:

1. In the BlackPearl dashboard, navigate to the Tape Partitions pane.
2. Select the partition you want to set to standby.
3. Click **Put On Standby**.
4. Click **Submit**.

Offline a Tape Drive

If a tape drive is experiencing errors and needs to be physically replaced, the drive can be taken offline to prevent the BlackPearl system from using the drive for data storage operations until the replacement is complete.

Here is how to offline a tape drive:

1. In the BlackPearl system dashboard, navigate to the **Tape Drives** pane.
2. **Select** the drive you want to offline.
3. Click **Take Offline**.
4. Click **Submit**.

APPENDIX B - CREATE AND CONFIGURE A VAIL VM NODE

This chapter describes the creation and configuration steps for a Vail VM Node.

Create a Vail VM Node	268
Vail VM Node Host Requirements	268
Create a Node Using VMWare vSphere	269
Create a Node Using Oracle VirtualBox	276
Configure the Vail VM Node Network Settings	283
Configure Network Settings	284
Configure the Vail VM Node Hostname	287
Configure the SSL Certificate	288
Register a Vail VM Node with a Vail Sphere	289

CREATE A VAIL VM NODE

Using a Vail VM node is useful when you want on-premise Standard or Standard_IA class storage.

The instructions in this section describe setup of a Vail VM node using a VMDK file. A Vail VM node can also be created using an OVA file. Contact Spectra Logic for assistance.

Note: Contact Spectra Logic for assistance configuring a Vail VM node with other virtual machine software such as Fusion, or Synology.

Vail VM Node Host Requirements

A Vail VM node requires the following:

- 8 core CPU or higher
- 16 GB RAM or higher
- 10 GigE Ethernet network connection or higher
- A network that allows access to port 443 to allow for data transfer

Router Requirements

**IMPORTANT**

All Vail VM nodes must be able to see each other using their announced IP address or hostname.

You may need to adjust the settings of any firewalls or proxy servers in your environment for the Vail VM nodes to communicate with each other. Contact your system administrator for assistance.

Port Requirements

All Vail VM nodes must be on a network that allows access on port 443.

VM Instance Protection

Spectra Logic recommends creating Vail VM nodes on reliable host computers and establishing a strong data protection system for your VM instances including regular snapshots to be used in the event of disaster recovery.

Create a Node Using VMWare vSphere

Here is how to create a Vail VM node using a VMDK file using VMWare® vSphere. These instructions are specific to vSphere and require familiarity with VM software.

1. If the Vail VM image file was provided to you by Spectra Logic, skip to [Step 2](#). Otherwise, download the latest Vail VM node image:
 - a. In the Vail management console, click the **gear icon**, then **Software Updates**.
 - b. Click **Download VM Image**.

Note: The file size is approximately 800 MB.

2. After the download completes, unpack the file.
3. Launch the VMWare vSphere application.
4. In the **Navigator** pane, select the host on which to create the VM node.




Figure 226 The VMWare vSphere home screen.

5. From the title bar, select **Actions > New Virtual Machine**.
6. In the New Virtual Machine wizard, select **Create a new virtual machine** and click **Next**.

7. Enter a **Name** for the VM node, select a **Location** , and click **Next**.




Figure 227 The New Virtual Machine - Select Name and Folder screen.

8. Using the **Select a compute resource** network browser, select an ESXi-based host in your network environment, and click **Next**.

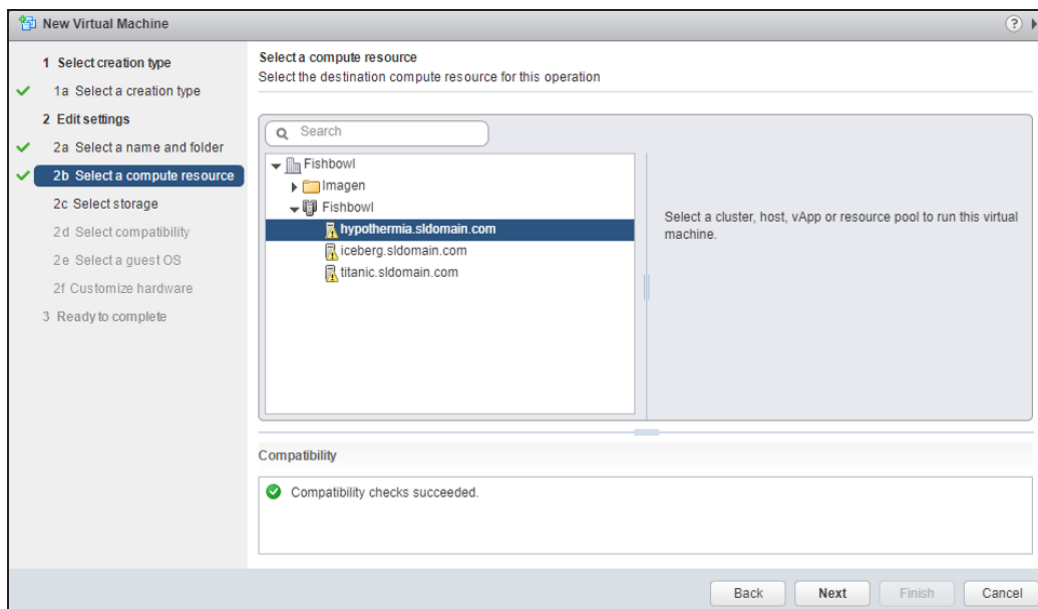


Figure 228 The New Virtual Machine - Select Compute Resource screen.

9. Using the **Select storage** table, select where to store the VM configuration files and virtual disks, and click **Next**.

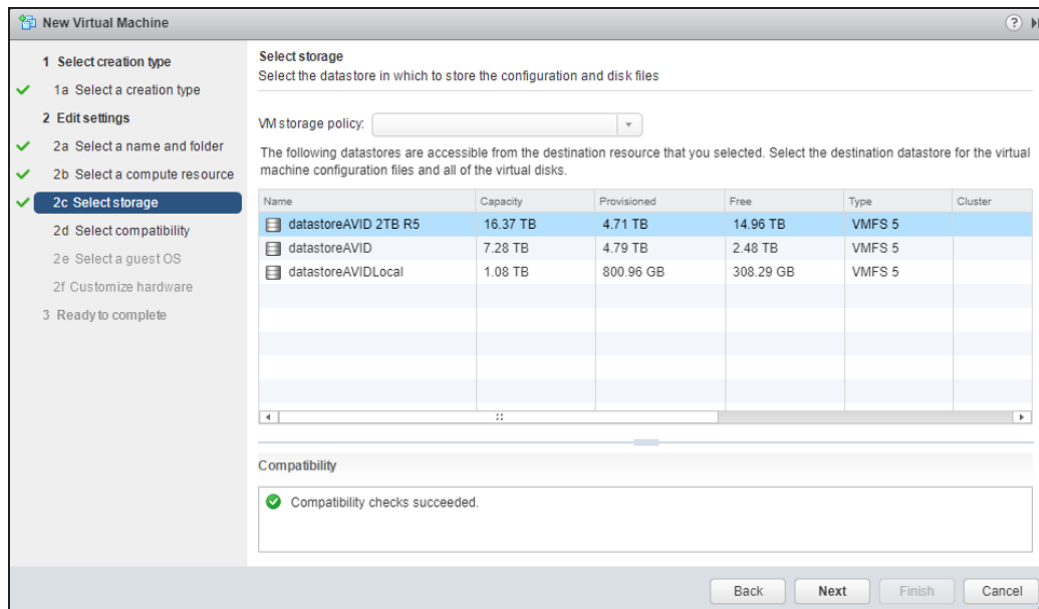


Figure 229 The New Virtual Machine - Select Storage screen.

10. Using the **Compatible with** drop-down menu, select **ESXi 6.5 and later**, and click **Next**.
11. Using the **Guest OS** drop-down menus, select the following:
- Guest OS Family: **Linux**
 - Guest OS Version: **Ubuntu Linux (64-bit)**
12. Click **Next**.

13. Using the **Customize hardware** screen, select the following:

- CPU: **8**
- Memory: **16 GB**




Figure 230 The New Virtual Machine - Customize Hardware screen.

14. On the right-hand side of the **New Hard disk** row, click the **X icon** to delete the default hard disk.

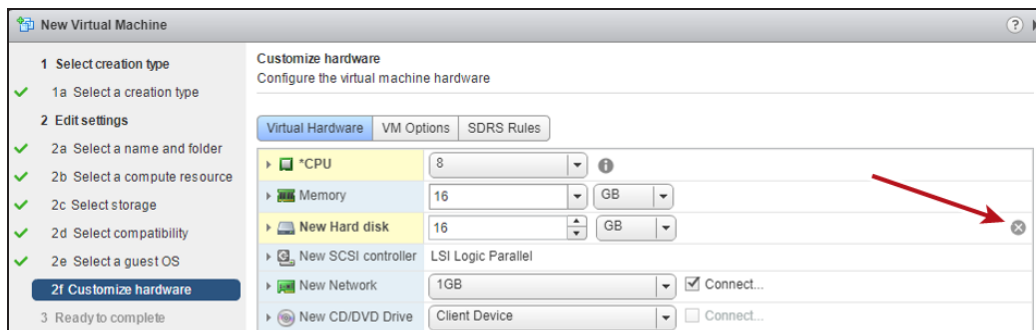


Figure 231 The New Virtual Machine - Customize Hardware screen.

15. Using the **New device** drop-down menu, select **Existing Hard Disk**, then click **Add**.




Figure 232 The New Virtual Machine - Customize Hardware screen.

16. Select the storage location of the VMDK file in the **Datastores** pane, then select the VMDK file to use in the **Contents** pane, and click **OK**.




Figure 233 The New Virtual Machine - Customize Hardware - Select File screen.

17. Using the **New device** drop-down menu, select **New Hard Disk**, then click **Add**. This creates the drive that the Vail VM node uses for data storage.

Note: If you increase the size of the drive after creating the Vail VM, the Vail application recognizes this change and allows you to use the newly available storage space.

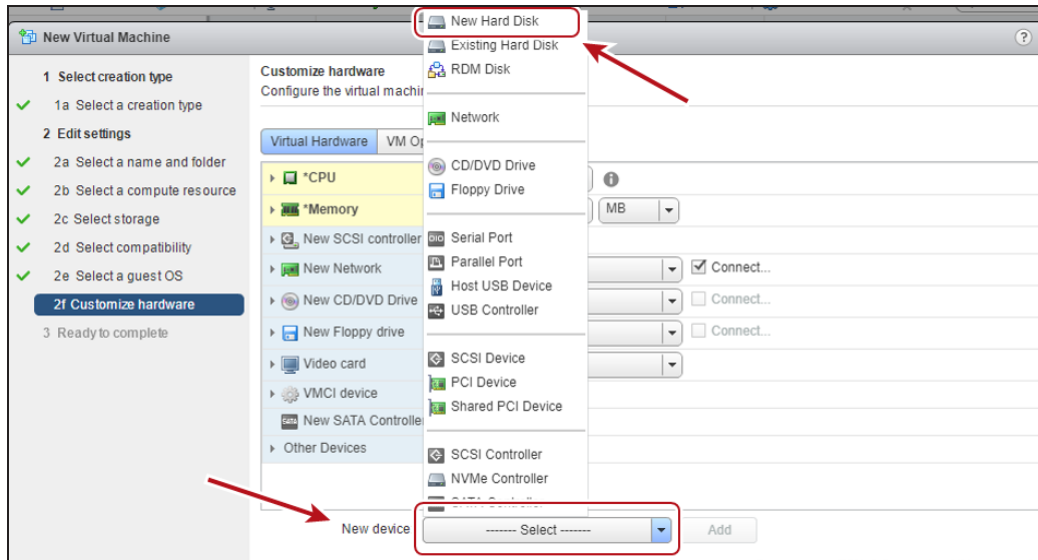


Figure 234 The New Virtual Machine - Customize Hardware screen.

18. Adjust the **Size** of the hard disk as required for your data storage environment.

Note: The size displays as GiB in the Vail management console.

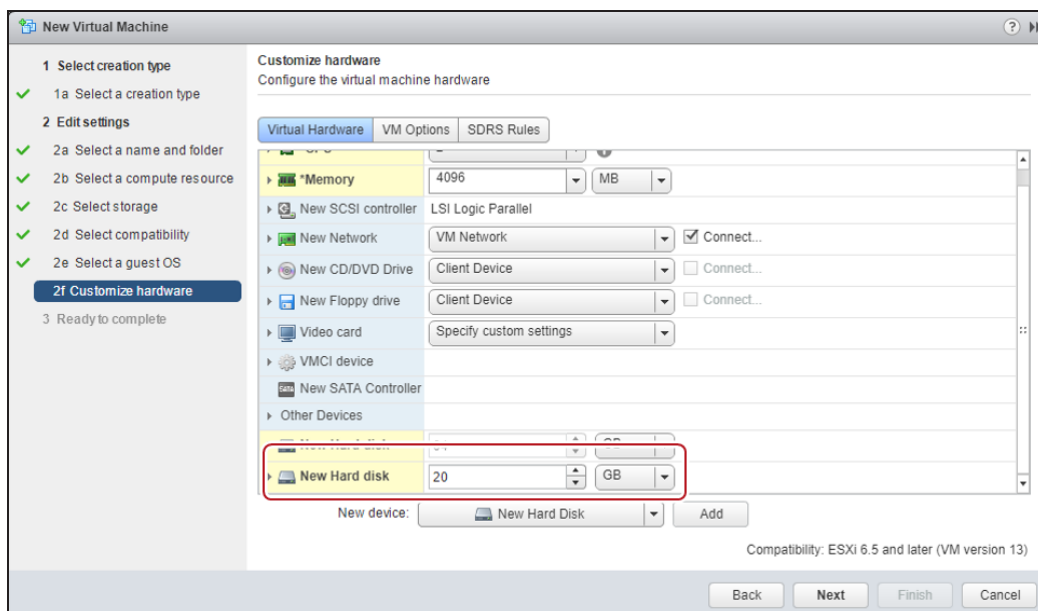


Figure 235 The New Virtual Machine - Customize Hardware screen.

19. Using the **New Network** drop-down menu, select **VM Network** and click **Next**.

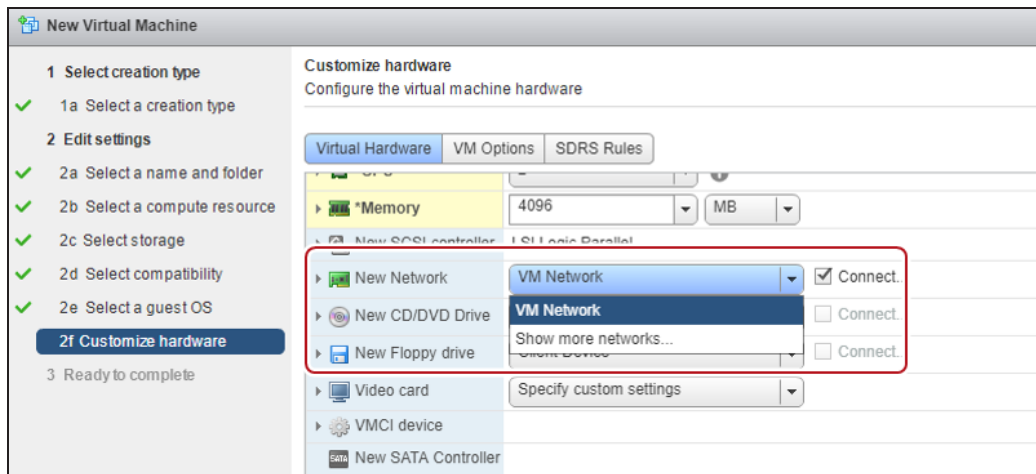


Figure 236 The New Virtual Machine - Customize Hardware screen.

20. Verify all settings are correct and click **Finish**.

21. In the **Navigator** pane, select the VM you just created, and on the title bar, click the **green Play triangle** to power-on the VM node.

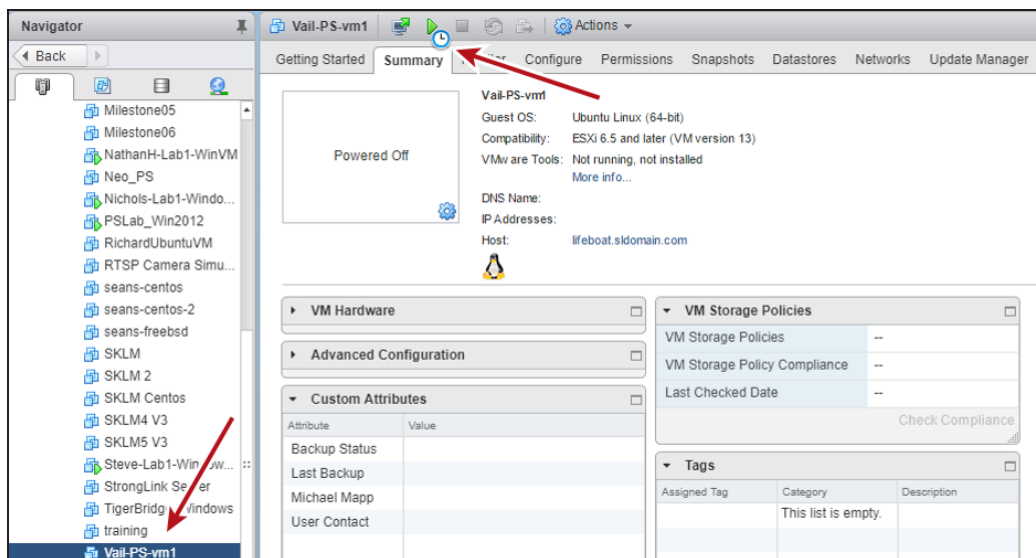
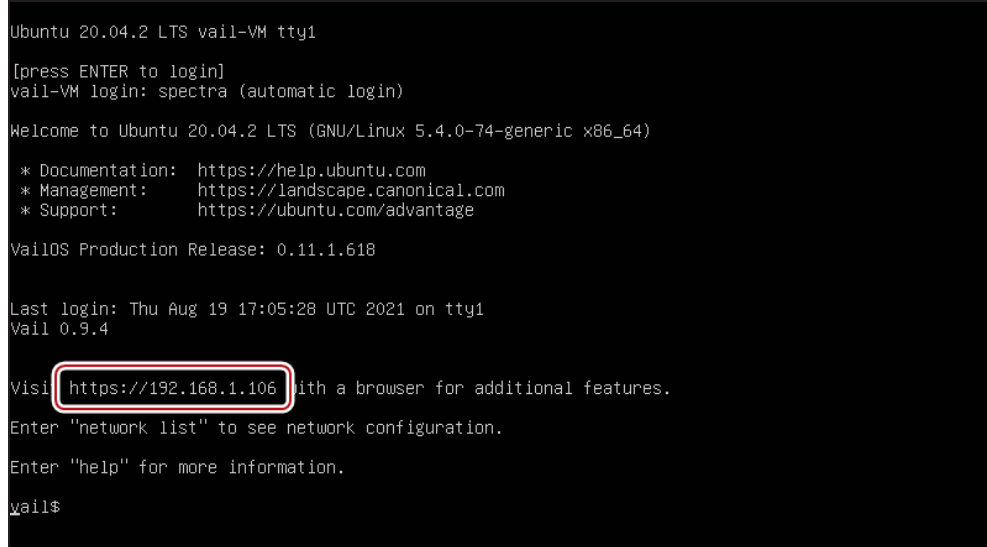


Figure 237 The New Virtual Machine - Summary screen.

22. When the VM boot completes, press **Enter**. If a DHCP server is configured, the IP address of the Vail VM node displays.

- Notes:**
- Do not close the VM window.
 - If no DHCP server is configured, contact Spectra Logic Professional Services to set a manual IP address.
 - You can change the network configuration of the Vail VM node after logging into the Vail VM management console.



```
Ubuntu 20.04.2 LTS vail-VM tty1
[press ENTER to login]
vail-VM login: spectra (automatic login)

Welcome to Ubuntu 20.04.2 LTS (GNU/Linux 5.4.0-74-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage

VailOS Production Release: 0.11.1.618

Last login: Thu Aug 19 17:05:28 UTC 2021 on tty1
Vail 0.9.4

Vail https://192.168.1.106 with a browser for additional features.
Enter "network list" to see network configuration.
Enter "help" for more information.
vail$
```

Figure 238 The Vail VM command line screen.

23. Open a web browser and enter the IP address. You are automatically logged in to the Vail VM user interface.

Note: The Vail VM node management console does not require any login credentials at this time.

Create a Node Using Oracle VirtualBox

Here is how to create a Vail VM node using a VMDK file using Oracle VirtualBox. These instructions are specific to the Windows version of Oracle VirtualBox and require familiarity with VM software.

1. If the Vail VM image file was provided to you by Spectra Logic, skip to [Step 2](#). Otherwise, download the latest Vail VM node image:
 - a. In the Vail management console, click the **gear icon**, then **Software Updates**.
 - b. Click **Download VM Image**.

Note: The file size is approximately 800 MB.

2. After the download completes, unpack the file.
3. Launch Oracle VirtualBox.


4. Click **New**.

Figure 239 Oracle VM VirtualBox Manager.

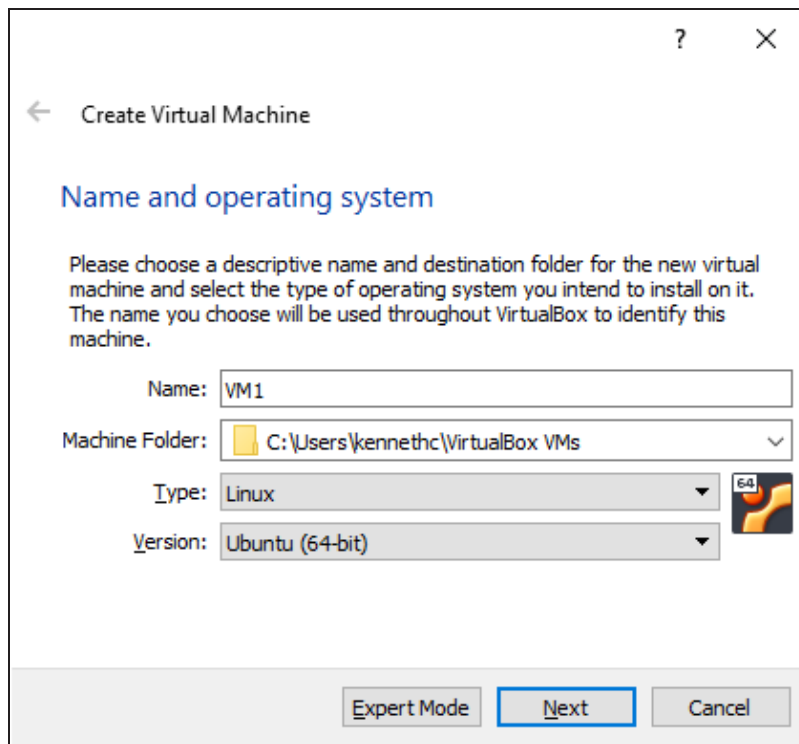
5. Enter the desired **Name**.

Figure 240 The Create Virtual Machine - Name & OS screen.

6. If desired, change the **Machine Folder** location.
7. Using the **Type** drop-down menu, select **Linux**.
8. Using the **Version** drop-down menu, select **Ubuntu 64-bit**.

9. Click **Next**.

Note: If you are asked to select the number of CPUs to use for the Vail VM, use the default setting.

10. Set the **Memory size** to **4096 MB** and click **Next**.

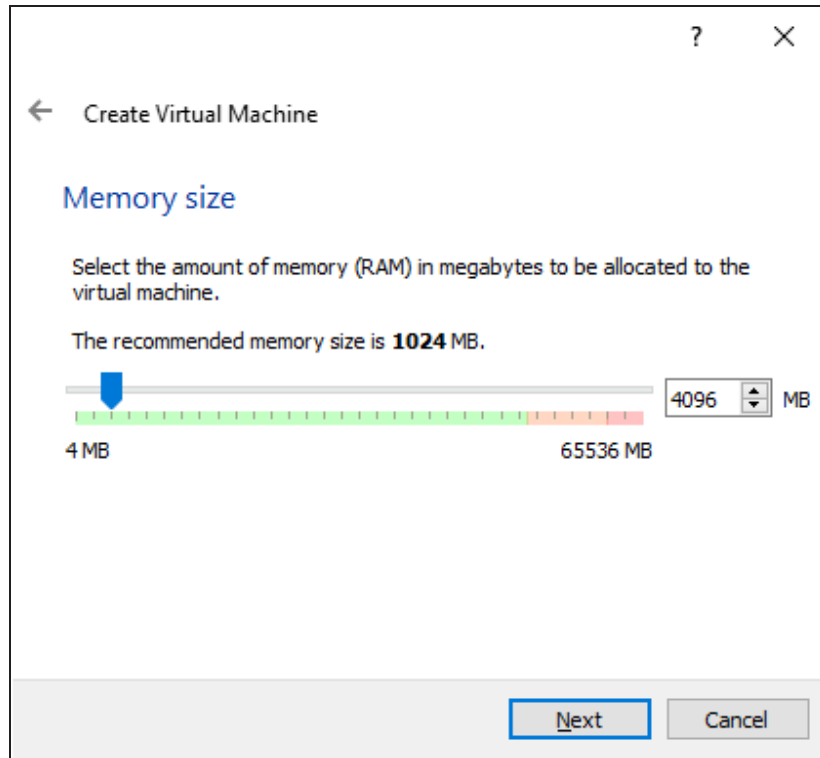


Figure 241 The Create Virtual Machine - Memory Size screen.

11. Select **Use an existing virtual hard disk file**, and click the folder icon to the right of the drop-down menu.

12. In the Hard Disk Selector screen, click **Add**, and browse to the VMDK you unpacked in [Step 2](#).

13. Select the file and click **Open**.

14. Under the **Not Attached** header, select the row of the new hard drive, then click **Choose**.




Figure 242 The Hard Disk Selector screen.

15. On the Create Virtual Machine - Hard disk screen, click **Create**.

16. After the VM is created, click **Settings**.

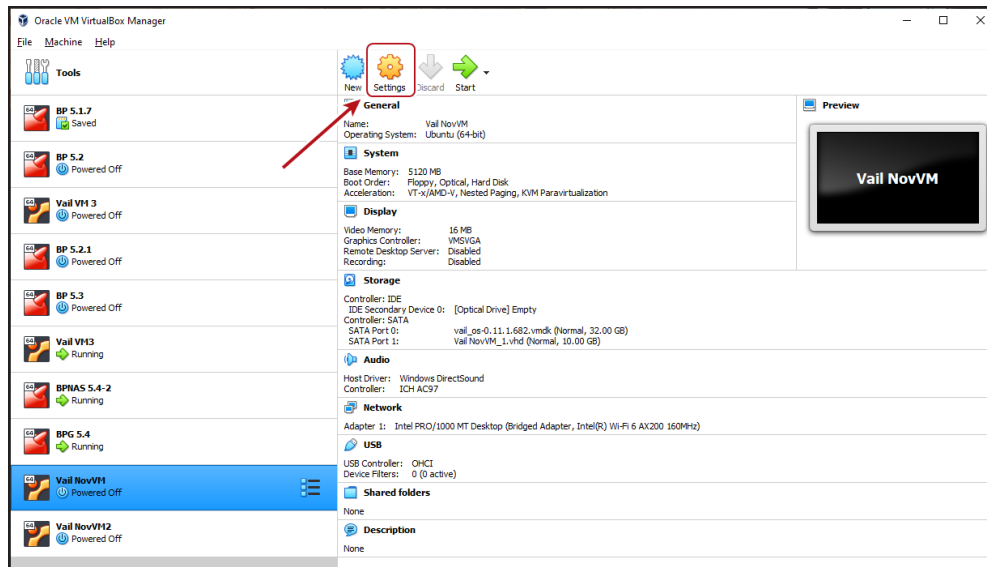


Figure 243 Oracle VM VirtualBox Manager.

17. In the left-hand pane of the Setting screen, click **Storage**.

18. Select the **Controller: SATA** row, and click the **Add hard disk** icon.




Figure 244 The VM Settings - Storage screen.

19. Select **Create**.




Figure 245 The Hard Disk Selector screen.

20. Select **VHD (Virtual Hard Disk)** and click **Next**. This is the disk the Vail VM node uses for data storage.

Note: If you increase the size of the drive after creating the Vail VM, the Vail application recognizes this change and allows you to use the newly available storage space.

21. Choose to allow the virtual hard disk to be **Dynamically allocated**, or to have a **Fixed size**, and click **Next**.

22. Configure the VHD file and size in GB, then click **Create**.

Note: The size displays as GiB in the Vail management console.

23. In the **Not Attached** list, select the row of the new hard drive, then click **Choose**.




Figure 246 The Hard Disk Selector screen.

24. In the left-hand pane of the Settings screen, click **Network**.

25. Using the **Attached to:** drop-down menu, select **Bridged Adapter**.




Figure 247 The VM Settings - Network screen.

26. If necessary, click the blue **Advanced** arrow to configure additional settings for your network environment.

27. Click **OK** to close the Settings window.

28. In the Oracle VM Manager main window, select the VM, and click **Start**.

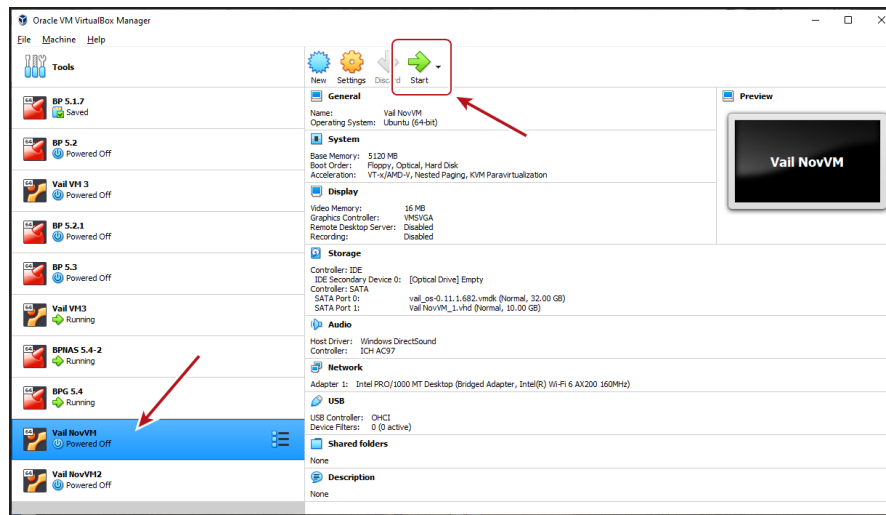


Figure 248 The VM Settings - Storage screen.

29. When the VM boot completes, press **Enter**. If a DHCP server is configured, the IP address of the Vail VM node displays.

- Notes:**
- Do not close the VM window.
 - If no DHCP server is configured, contact Spectra Logic Professional Services to set a manual IP address.
 - You can change the network configuration of the Vail VM node after logging into the Vail VM management console.

```

Ubuntu 20.04.2 LTS vail-VM tty1
[press ENTER to login]
vail-VM login: spectra (automatic login)

Welcome to Ubuntu 20.04.2 LTS (GNU/Linux 5.4.0-74-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage

VailOS Production Release: 0.11.1.618

Last login: Thu Aug 19 17:05:28 UTC 2021 on tty1
Vail 0.9.4

Visit https://192.168.1.106 with a browser for additional features.
Enter "network list" to see network configuration.
Enter "help" for more information.
vail$

```

Figure 249 The Vail VM command line screen.

30. Open a web browser and enter the IP address. You are automatically logged in to the Vail VM user interface.

Note: The Vail VM node management console does not require any login credentials at this time.

CONFIGURE THE VAIL VM NODE NETWORK SETTINGS

If desired, use the instructions in this section to edit the Vail VM node IP address, hostname, and SSL certificate.

If your Spectra Vail application is running on a BlackPearl system, the network settings for IP addressing, SSL certificates, and hostname are controlled by the BlackPearl system. See the [*BlackPearl Nearline Gateway User Guide*](#) for information.

**IMPORTANT**

Spectra Logic recommends setting a static IP address and changing the hostname as described in the sections below.

Use one of the sections below to configure network settings.

- **Configure Network Settings on the next page**
- **Configure the Vail VM Node Hostname on page 287**
- **Configure the SSL Certificate on page 288**

Configure Network Settings



IMPORTANT

The Spectra Vail application restarts after changing the VM network settings. Any data transfer operations fail when the application restarts. Internal operations, such as Lifecycles, automatically restart. External operations must be manually restarted.

Here is how to configure the Vail VM node IP address:

1. Discontinue storage operations. The Spectra Vail application restarts after changing the network settings.
2. In the upper right corner of the Vail management console, click the **gear icon** and select **Network**.

3. Select the interface adapter row and click **Edit**.

The screenshot shows the 'Edit Network' window with the following settings:

- IPv4 Mode:** Manual (selected)
- IPv6 Mode:** Automatic
- Static Addresses:** A table with one entry: IP Address '192.168.12.231' and Prefix Length '24'.
- IPv4 Default Gateway:** 192.168.12.1
- IPv6 Default Gateway:** fe80::1a58:80ff:fec2:c57f
- MTU:** 1500
- Mode:** Manual (selected over DHCP)
- Name Servers:** 192.168.12.1 and fe80::1a58:80ff:fec2:c57f
- Search Domains:** lan
- SAVE** button at the bottom right.

Figure 250 The Vail VM Node Edit Network screen.

By default, DHCP is selected on the Edit Network screen to provide the IPv4 address. However, Spectra Logic recommends configuring a static IPv4 address.

Note: If you require the Vail VM node to be configured using a DHCP address, Spectra Logic recommends you use your DHCP server to bind the IP address to the Vail VM node.

- To configure the IP address manually,
 - a. Using the **IPv4 Mode** and **IPv6 Mode** drop-down menus, select **Manual**.
 - a. Edit the IPv4 and IPv6 **IP Addresses** as required.
 - b. Enter a value for the **Prefix Length**.

Note: To add a new IP address, click the + sign. To remove an IP address, click the **garbage can** icon.

- c. Edit the **IPv4 Default Gateway**.
- d. If desired, enter the **IPv6 Default Gateway**.
- e. Change the **MTU** value as desired.
- f. Enter one or more **Name Server(s)** and **Search Domain(s)**.
- g. Click **Save**.

Note: The Vail VM node interface refreshes after the node changes network settings. The interface may display a lost communication error for several seconds.

- To use DHCP to set the IP address,

Note: If you require the Vail VM node to be configured using a DHCP address, Spectra Logic recommends you use your DHCP server to bind the IP address to the Vail VM node.

- a. If necessary, using the **IPv4 Mode** drop-down menu, select **DHCP**.
- b. If necessary, using the **IPv6 Mode** drop-down menu, select **Automatic**.
- c. Configure the DNS settings:
 - To configure the DNS settings automatically, select **DHCP** and click **Save**.
 - To configure DNS settings manually, select **Manual**. Enter one or more **Name Server(s)** and **Search Domain(s)** and click **Save**.
- d. Click **Save**.

Note: The Vail VM node interface refreshes after the node changes network settings. The interface may display a lost communication error for several seconds.

Configure the Vail VM Node Hostname

The Vail VM node hostname is used as the top level name of the storage endpoint displayed in the Vail management console. Spectra Logic recommends using a name that includes both the location and type of storage.

For example, in the Dallas location, add the storage type as a suffix such as, Dallas-VM1 and Dallas-VM2.

Here is how to configure the hostname:

1. In the upper right corner of the Vail management console, click the **gear icon** and select **Hostname**.
2. Under the **Hostname** banner, click **Edit**.

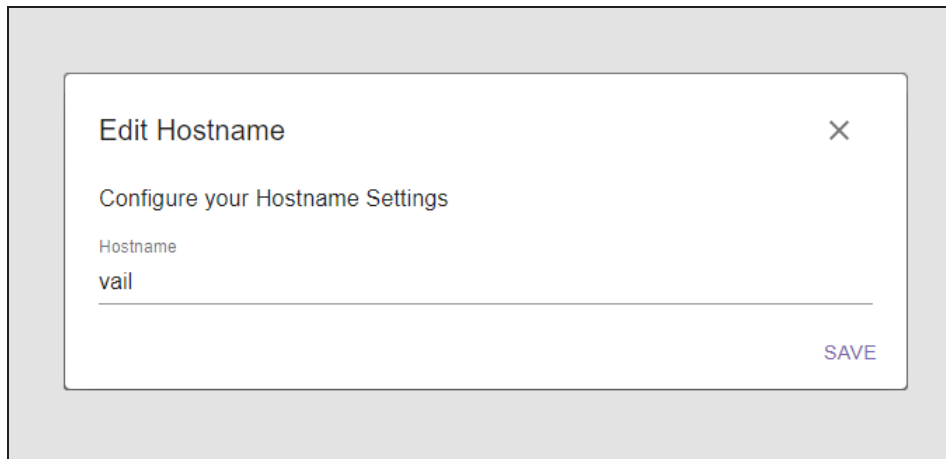


Figure 251 The Vail VM Node Edit Hostname screen.

3. Edit the desired **Hostname** and click **Save**.

Note: Only alphanumeric and the dash (-) character are allowed. The hostname is case sensitive.

Configure the SSL Certificate



IMPORTANT

The Spectra Vail application requires that SSL certificate for the Spectra Vail application and the BlackPearl Nearline Object Gateway are recognized as valid by clients on your DNS network servers.

Here is how to configure SSL certificate:

1. In the upper right corner of the Vail management console, click the **gear icon** and select **SSL Certificate**.
2. Under the **SSL Certificate** banner, click **Edit**.

The screenshot shows a modal dialog titled "Edit SSL Certificate". Inside the dialog, there is a subtitle "Configure your SSL Certificate" followed by a question mark icon. Below this, there are three text input fields: "Certificate", "Private Key", and "Passphrase". The "Passphrase" field also has a question mark icon to its right. At the bottom right of the dialog, there is a blue "SAVE" button.

Figure 252 The Vail VM Node Edit SSL Certificate screen.

3. Enter the desired **Certificate** and **Private Key** in PEM format.
4. If necessary, enter the **Passphrase** that was used to encrypt the private key.
5. Click **Save**.

REGISTER A VAIL VM NODE WITH A VAIL SPHERE

Registering a Vail VM node with a Vail sphere allows you to use the node for data storage.

Here is how to register a Vail VM node with a Vail sphere:

1. In the Vail VM node management console taskbar, click **Dashboard**.
2. Under the **Dashboard** banner, click **Activate**.

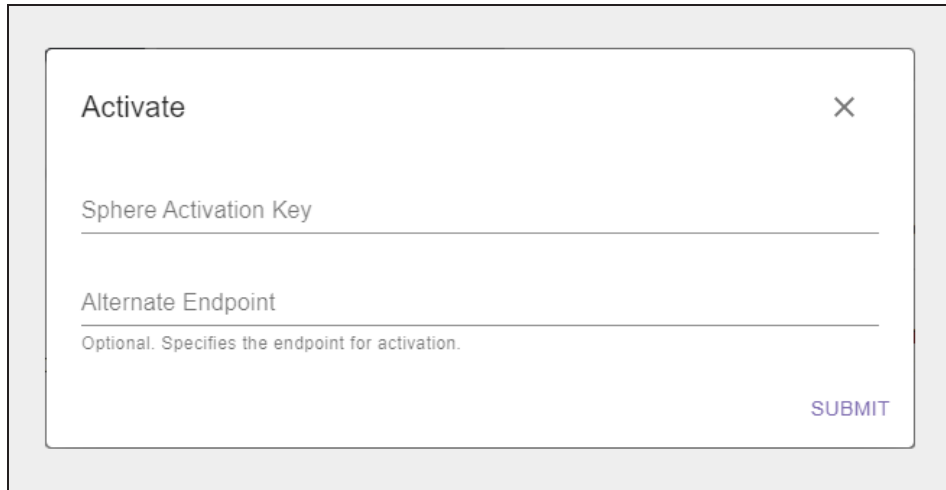


Figure 253 The Vail VM Node - Activate screen.

3. Enter the **Sphere Activation Key** provided by Spectra Logic.
4. If necessary, enter the **Alternate Endpoint**.
5. Click **Submit**. After a few moments the Dashboard screen refreshes once activation completes.
6. Under the **Dashboard** banner, click **Register With Sphere**.

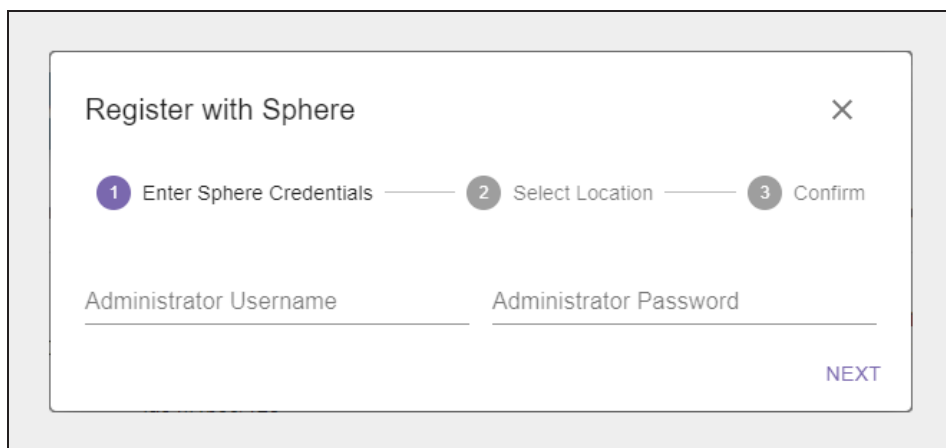


Figure 254 The Register With Sphere - Credentials screen.

7. Enter the Spectra Vail application **Administrator Username** and **Password**.

8. Click **Next**.




Figure 255 The Register with Sphere - Select Location screen.

9. On the Select Location screen, chose to create a new location, or to use an existing location:
- **Create a New Location below**
 - **Select an Existing Location on page 294**

Create a New Location

Here is how to create a new location:

1. To create a new location, use the drop-down to select **New Location**.

2. To map a location, you can search for the location, manually enter the latitude and longitude, or create a location with no corresponding geographic location.

Register with Sphere

✓ Enter Sphere Credentials ✓ Reset Password 3 Select Location 4 Confirm

Select a Location
New Location

Search and choose an address to use for your new location.
Note: You may skip this step if you wish to enter your location data manually.

Address Search

Please confirm the details below. If necessary, you may edit any pre-populated fields or execute another search.
Note: Latitude and Longitude values are used for the System View map on the dashboard.

Name

Latitude Longitude

PREVIOUS NEXT

Figure 256 The Register with Sphere - New Location screen.

- To search for a location...

- a. In the **Address Search** field, enter a geographic location. You can enter a full or partial postal address, city, county, province, or country.
- b. Select the correct match from the list.

Note: If no match is located, try changing the format of the address you entered. For example, use 9th Street in place of Ninth St.

- c. If desired, manually edit the **Name**.

Spectra Logic recommends naming each location after its physical location in the world.

For example, if Vail resources are located in Dallas, use that as the location name if there is only one Vail resource in that city. If there are multiple Vail resources consolidated in the same city, use suffixes to identify each group such as Dallas-HQ, Dallas-Research, or Dallas-Production.

- d. Confirm the information is correct and click **Next**.

- To manually enter a location...

- a. Enter the desired **Name**.

Spectra Logic recommends naming each location after its physical location in the world.

For example, if Vail resources are located in Dallas, use that as the location name if there is only one Vail resource in that city. If there are multiple Vail resources consolidated in the same city, use suffixes to identify each group such as Dallas-HQ, Dallas-Research, or Dallas-Production.

- b. Enter the **Latitude** and **Longitude** of the location.

- Notes:**
- When entering a value for **Latitude**, use positive values for locations north of the equator, and negative values for locations south of the equator.
 - When entering a value for **Longitude**, use positive values for locations east of the prime meridian, and negative values for locations west of the prime meridian.

- c. Click **Next**.

- To skip entering a location...

- a. Enter the desired **Name**.

Spectra Logic recommends naming each location after its physical location in the world.

For example, if Vail resources are located in Dallas, use that as the location name if there is only one Vail resource in that city. If there are multiple Vail resources consolidated in the same city, use suffixes to identify each group such as Dallas-HQ, Dallas-Research, or Dallas-Production.

- b. Click **Next**.

Note: If you do not enter an address or latitude and longitude, the location displays on the right-hand pane of the Vail dashboard, but does not display on the world map.

3. Confirm the information is correct, and click **Register**.

Wait while the Vail VM node registers with the Vail sphere. This may take several minutes, during which time the Vail VM node interface changes to the Vail management console, and may display communication errors.

Select an Existing Location

Here is how to select an existing location:

1. Using the drop-down menu, **Select a Location** where you want to associate the Vail VM node and click **Next**.




Figure 257 The Register with Sphere - Select Location screen.

2. Confirm the information is correct, and click **Register**.

Wait while the Vail VM node registers with the Vail sphere. This may take several minutes, during which time the Vail VM node interface changes to the Vail management console, and may display communication errors.

FREQUENTLY ASKED QUESTIONS

This section covers frequently asked questions that help you understand how the Spectra Vail application operates.

Why Do Vail Jobs Show as Canceled in the BlackPearl User Interface?

When the Spectra Vail application requests an object(s) from a BlackPearl Nearline Object Gateway, it initiates a Start Bulk Get job in the BlackPearl Nearline Object Gateway. However, the Spectra Vail application has a back-door path to read objects from the BlackPearl cache. The BlackPearl Nearline Object Gateway is only aware of when objects are read through the front door path. When the Spectra Vail application completes reading the requested object(s) from the BlackPearl cache, it cancels the job on the BlackPearl Nearline Object Gateway.

What is the Difference Between AWS Linked Buckets and BlackPearl Linked Buckets?

Vail linked buckets allow the Spectra Vail application to connect to an AWS or BlackPearl bucket, and link to the objects in that bucket. These linked buckets are connected to a Vail bucket. With both AWS and BlackPearl linked buckets, any objects that are currently in the bucket become part of the associated Vail bucket when they are linked. Additionally, any objects added to the AWS or BlackPearl bucket after it is linked to Vail also becomes part of the Vail linked bucket.

AWS linked buckets additionally allow objects added to the Vail bucket to be copied to the linked AWS bucket.

Note: A BlackPearl system does not support this feature.

Who Owns Objects Managed by the Vail Sphere?

Objects copied from an external bucket to a Vail bucket are owned by the owner of the Vail bucket, while objects copied to an external bucket are owned by the user with the credentials used when creating the BlackPearl storage for the bucket.

At What Size Must a PUT Job be a Multi-Part Upload?

The upper size limit before an object must be PUT using multi-part upload is 5 GB.

Note: Spectra Logic recommends using multi-part upload for any object over 1 GB.

Why Do I Receive AWS Connectivity Error Messages From Third-Party Software But Not From Vail?

The Spectra Vail application and the BlackPearl Nearline Gateway do not generate error messages when an AWS connection is unavailable. However some third-party applications, such as Rubrik, may generate an error message when this occurs. In most cases, no user action is necessary.

BEST PRACTICES

Note: This section will be updated in future releases of this guide.

Using Spectra Vail and a BlackPearl System Simultaneously

The best practice for using a BlackPearl system with the Spectra Vail application is to allow the Spectra Vail application to control all aspects of the BlackPearl system.

If you need to use the BlackPearl system for data storage outside of the Vail application control, keep logical separation of as many elements on the BlackPearl system as possible. All buckets, users, data policies, and storage domains should be separate and unique between the BlackPearl system side and the Vail application side.

GLOSSARY

BlackPearl System

A BlackPearl Nearline Object Gateway is used to provide the interface between the Spectra Vail application and tape storage. A BlackPearl system stores data in a local cache before writing to tape media. When data is requested by the Spectra Vail application the BlackPearl system copies data from tape storage to the cache so it can be accessed by the Spectra Vail application. A BlackPearl system can additionally provide storage to disk media, using Online and NAS storage.

Lifecycle

A lifecycle consists of one or more rules that dictate where objects data is stored and the length of time it is stored in each specified storage location. Users control the data placement using placement and deletion rules, and the storage endpoint where those clones are placed. Lifecycle rules are interpreted on a once per day basis, thereby producing a list of content to move. Data is then moved as a background process.

The available storage targets consist of Vail VM nodes, S3 buckets, and BlackPearl® systems that are associated with the Spectra Vail application. Users can create up to five rules per lifecycle to govern the movement and location of data. Users can delete rules at any time, and any data movement in progress completes based on the known rules at the time the transaction started.

Storage

A storage destination consists of either disk-based storage provided by a Vail cluster, block storage provided by a Vail VM node, a BlackPearl bucket, a BlackPearl NAS share, or an AWS® S3 repository. Disk-based and block storage can utilize the Standard or Standard-Infrequent Access storage classes, while BlackPearl bucket storage on tape can only use the Glacier storage class. AWS repositories can use any storage class.

Storage Classes

Amazon S3 provides multiple storage classes for different use cases. The Spectra Vail application recognizes all storage classes supported by AWS, but only uses storage class types Standard, Standard-Infrequent Access, and Glacier.

The Spectra Vail application makes a best guess regarding where to place data if any other storage class is specified. Lifecycles can be used to transition data from one storage class to another.

Standard (SA)

This storage class is best for frequently accessed data, as it offers high performance, availability, and data durability, as well as low latency and high throughput.

Standard (SA) is fast access storage such as disk, flash, or block storage, as well as Amazon S3 or third-party S3 object storage.

Standard-Infrequent Access (IA)

This storage class is best suited for data that does not need to be accessed frequently, but needs to be retrieved immediately when access is requested. The Standard-IA storage class offers the same low latency, high performance and durability of the Standard storage class, but at lower cost.

Glacier

This storage class is best suited for long-term storage and archiving, as it offers high security and durability at the lowest cost. This storage class is fundamentally different in that in order to access data in Glacier storage, the data must first be retrieved, and this retrieval can take many hours to complete. In order to use this storage class, S3 clients must be able to issue an "object restore" command to move the object from Glacier storage to Standard storage. After the object is available on Standard storage, a GET command is used to access the object.

Vail Bucket

A Vail bucket is the highest-level logical storage container for S3 objects. Each Vail bucket is a unique endpoint and displays a single view of all objects in the bucket, which can have managed copies at multiple sites, in multiple clouds, and in multiple storage classes or tiers.

Vail buckets may be assigned a lifecycle to control the movement of data, but do not require a lifecycle. Multiple Vail buckets can use the same lifecycle. Vail buckets can also be configured to use encryption.

Linked Bucket

The Spectra Vail application is able to link to existing AWS S3 or BlackPearl buckets and create a linked bucket. When this is done the Spectra Vail application is immediately aware of the existing data which allows for ongoing synchronization with external storage targets in the Vail sphere, while still allowing for the application of lifecycle rules.

Only one linked bucket is allowed per storage location.

Location

A Location denotes a physical location in the world that consists of a set of storage targets or physical storage such as a BlackPearl Nearline Object Gateway, tape storage connected to a BlackPearl Nearline Object Gateway, Vail VM node storage, and Vail clusters that share the same physical location.

BlackPearl Storage

The Spectra Vail application uses a BlackPearl Nearline Object Gateway to provide disk storage as an S3 Standard (SA) target, and to optionally provide On-Prem Glacial storage using Spectra Logic tape libraries.

On-Prem Glacier Storage

A BlackPearl Nearline Object Gateway with On-Prem Glacier storage allows data to move seamlessly into tape storage in a way not previously possible. It enables users to deploy a tier of deep storage that is cost effective, easy to manage, and scalable to exabytes of data.