



SPECTRA VAIL USER GUIDE



www.SpectraLogic.com

TABLE OF CONTENTS

Table Of Contents	2
Document Information	6
Copyright	7
Notices	7
Trademarks	7
Master License Agreement	8
Contacting Spectra Logic	16
Introduction	17
Features	17
Related Publications	18
Chapter 1 - Things To Know	19
Requirements	20
Spectra Logic Products Requirements	20
Supported Browsers	20
Vail Naming Conventions	21
Vail Sphere Names	21
User Names	21
Group Names	21
Location Names	21
Portable Location Names	22
Storage Names	22
Vail Bucket Names	23
Cloud Bucket Names	23
BlackPearl Bucket Names	23
Lifecycle Names	24
Additional AWS Account Role Names	24
Vail Management Console Overview	25
Main window	25
Taskbar	26
Toolbar	26
Other Icons	28

Chapter 2 - Configure BlackPearl System	29
Configure a BlackPearl S3 Solution	30
Register a BlackPearl S3 with a Vail Sphere	31
Chapter 3 - Configure the Vail Application	39
Log In to the Vail Management Console	40
Vail Sphere Configuration Paths	41
Create Storage	42
Create BlackPearl Storage	43
Create BlackPearl Bucket Storage	43
Create Vail S3 NAS Storage	47
Create Cloud Storage	48
Create AWS S3 Cloud Storage	48
Create Microsoft Azure Cloud Storage	51
Create Google Cloud Platform Storage	53
Create Other S3 Cloud Storage	55
Create a Lifecycle	58
Create a Vail Bucket	66
Configure an Object Storage Browser	73
Configure S3 Browser	73
Configure Cyberduck Object Storage Browser	74
Chapter 4 - Configure & Manage Users	76
Configure & Manage Sphere Administrator Accounts	77
Create a Sphere Administrator	77
Change a Sphere Administrator Password	79
Edit Sphere Administrator Attributes	81
Delete a Sphere Administrator	83
Configure & Manage IAM Accounts	84
Add an IAM Account	84
Edit an IAM Account	87
Delete an IAM Account Association	89
Configure & Manage IAM Users and Groups	90
Create an IAM User	90
View IAM User Details	91
Add an IAM User to an IAM Group	92
Remove an IAM User from an IAM Group	93

Delete an IAM User	94
Create an IAM Group	95
Delete an IAM Group	96
AWS Access Key Management	97
Create an Access Key	97
Enable an Access Key	98
Disable an Access Key	99
Delete an Access Key	100
Chapter 5 - Create and Configure a Vail VM Node	102
Create Vail VM Node Storage	103
Vail VM Node Host Requirements	103
Create a Node Using VMWare vSphere	104
Create a Node Using Oracle VirtualBox	111
Configure the Vail VM Node Network Settings	119
Configure Network Settings	119
Configure the Vail VM Node Hostname	121
Configure the SSL Certificate	122
Register a Vail VM Node with a Vail Sphere	124
Chapter 6 - Using the Vail Application	130
View Capacity Information	132
View Performance Metrics	135
View Statistics	138
View Vail Bucket Details	141
View Vail Bucket Contents	145
View Object Details	147
Edit Global Settings	150
Change Lifecycle Rule Nightly Processing Time	150
Enable Diagnostic Monitor	150
Using Proxy Connections	152
Configure Proxy Connection	152
Edit Proxy Server	152
Delete Proxy Server	153
Edit a Vail Bucket	154
Delete a Vail Bucket	158
Create an Object Clone	159

Delete an Object Clone	161
View Storage Details	163
Edit BlackPearl or Vail VM Endpoint	165
Change Endpoint Location	165
Add Additional Host Names	166
Configure Debug Logging	167
Edit Storage	169
Edit BlackPearl Bucket Storage	169
Edit BlackPearl NAS Storage	172
Edit Vail VM Node Storage	173
Edit Google Cloud Platform Storage	174
Edit AWS S3 Cloud Storage	176
Edit Other Third-Party Cloud Storage	179
Delete Storage	180
View Lifecycle Details	184
Edit a Lifecycle	187
Delete a Lifecycle	190
Create a Location	191
Delete a Location	195
Clear the IAM Cache	196
View Reports	197
View Vail Application Messages	200
Message Details	202
Vail Application Logs	203
Update the Vail Application Software	204
Enable Diagnostic Monitor	206
Log Out of the Vail Management Console	207
Frequently Asked Questions	208
Glossary	209

DOCUMENT INFORMATION

Documentation part number:

- 90990149

Documentation revision:

- Rev. D - 10/9/2023 - Updated for Vail 2.5.4.

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.

COPYRIGHT

Copyright © 2022-2023 Spectra Logic Corporation. All rights reserved. This item and the information contained herein are the property of Spectra Logic Corporation.

NOTICES

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

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

TRADEMARKS

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

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.spectralogic.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.spectralogic.com	
United States and Canada - Phone Toll free US and Canada: 1.800.227.4637 International: 1.303.449.0160	Europe, Middle East, Africa Phone: 44 (0) 870.112.2185 Deutsch Sprechende Kunden Phone: 49 (0) 6028.9796.507 Email: spectralogic@stortrec.de
Mexico, Central and South America, Asia, Australia, and New Zealand Phone: 1.303.449.0160	
Spectra Logic Sales Website: shop.spectralogic.com	
United States and Canada Phone: 1.800.833.1132 or 1.303.449.6400 Fax: 1.303.939.8844 Email: sales@spectralogic.com	Europe Phone: 44 (0) 870.112.2150 Fax: 44 (0) 870.112.2175 Email: eurosales@spectralogic.com
To Obtain Documents - Spectra Logic Website: support.spectralogic.com/documentation	

INTRODUCTION

This guide describes the installation, configuration, and use, of the Spectra Logic® 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.

FEATURES

The Spectra Vail application software features:

- Global access to a single name space.
- Command, control, and monitoring using a single management point.
- Data stored on-premises using flash, disk, tape, or 3rd party storage.
- Data stored to and synchronized from cloud storage.
- Rich policy engine provides time-based data placement for each storage type and geographic location.
- Unlimited number of data storage sites.

The Spectra Vail application manages the movement and retention of data using cloud-based command and control, allowing an organization the ability to log in from anywhere in the world to control and monitor data movement.

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.

Once an object is uploaded anywhere in the Vail sphere, it is immediately available to anyone who is joined to the sphere.

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://developer.spectrallogic.com/doc/vail/ga/Default.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 User Guide](#) provides detailed information about configuring, using, and maintaining your BlackPearl 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](#).

CHAPTER 1 - THINGS TO KNOW

This chapter provides helpful things to know before using the Spectra Logic Spectra Vail application.

Requirements	20
Spectra Logic Products Requirements	20
Supported Browsers	20
Vail Naming Conventions	21
Vail Sphere Names	21
User Names	21
Group Names	21
Location Names	21
Portable Location Names	22
Storage Names	22
Vail Bucket Names	23
Cloud Bucket Names	23
BlackPearl Bucket Names	23
Lifecycle Names	24
Additional AWS Account Role Names	24
Vail Management Console Overview	25
Main window	25
Taskbar	26
Toolbar	26
Other Icons	28

REQUIREMENTS

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

Spectra Logic Products Requirements

- The Spectra Vail application version 2.0.x or later requires either a BlackPearl Nearline Gateway or a Vail VM node.
 - The BlackPearl Nearline Gateway requires BlackPearl 5.6.3 or later with a valid Vail Sphere activation key installed.
 - The 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 NAMING CONVENTIONS

Before configuring the 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 Vail components added at a later date.

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

The 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 Vail spheres, each sphere must have a unique name.

User Names

User names identify 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 Vail resources. 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.

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 Vail sphere. Storage names are used for 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 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 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`.

Vail Bucket Names

Vail buckets are the highest level of object grouping in the Vail sphere. 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 Vail sphere is configured with a master Administrator AWS account.

Additional AWS accounts can be configured in the 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 MANAGEMENT CONSOLE OVERVIEW

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.

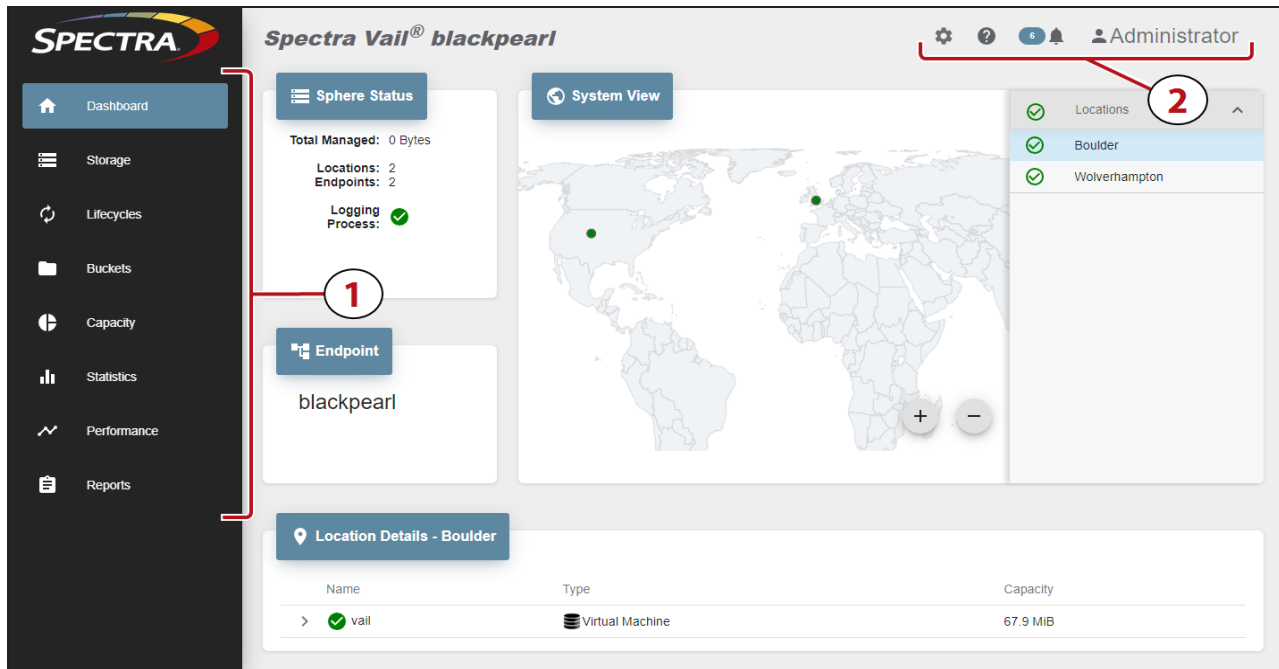


Figure 1 The Vail Sphere Dashboard screen.

Main window

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

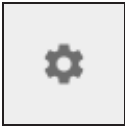

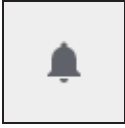
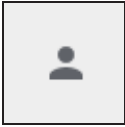
Taskbar

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	The Dashboard navigation link takes you to the Dashboard screen which provides an overview of the Vail sphere status, the quantity and location of configured storage, as well as the total size of managed data.
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 42 .
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 58
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 66 .
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 132 .
Statistics	The Statistics navigation link takes you to the Statistics screen which displays the data storage growth rate for endpoint and cloud storage. See View Statistics on page 138 .
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 135
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 197 .






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
	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 • Configure Network Settings • Configure Global Settings • Access Logs • Update the Vail application
	Online Help	<p>Opens a web browser to the Vail online help guide, a browser-based version of the Spectra Vail User Guide.</p>
	Messages	<p>Displays the number of unread messages generated by the Vail sphere. 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. • 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. • Error - An event which impacts data storage operations occurred. This may happen if the Vail sphere cannot communicate with storage endpoint. • OK - An issue that generated an error message is resolved.
	User	<p>Displays the user currently logged-in and provides access to the logout function.</p>

Other 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.
	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	Indicates an error of a component of the Vail sphere.

CHAPTER 2 - CONFIGURE BLACKPEARL SYSTEM

This chapter provides instructions for configuring a BlackPearl S3 solution for use with the Spectra Vail application.

Configure a BlackPearl S3 Solution	30
Register a BlackPearl S3 with a Vail Sphere	31

CONFIGURE A BLACKPEARL S3 SOLUTION

If you are using the Spectra Vail application in conjunction with a BlackPearl S3 solution, before you can access the Spectra Vail application, you must first install and configure the BlackPearl S3 solution.

Use the [*BlackPearl Nearline Gateway User Guide*](#) to configure your BlackPearl S3 solution.

Your BlackPearl S3 solution may have been installed and configured by Spectra Logic Professional Services.

Note: If you need assistance configuring your BlackPearl S3 solution, contact Spectra Logic. See [Contacting Spectra Logic on page 16](#).

REGISTER A BLACKPEARL S3 WITH A VAIL SPHERE

Here is how to register a BlackPearl S3 solution with a Vail sphere:

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 124.

1. Log in to the BlackPearl user interface.
2. If necessary, configure the IP addressing for the BlackPearl S3 solution. 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 S3 solution. The Spectra Vail application uses this name for the Vail node name.
 - a. Select **Status > Hardware**.

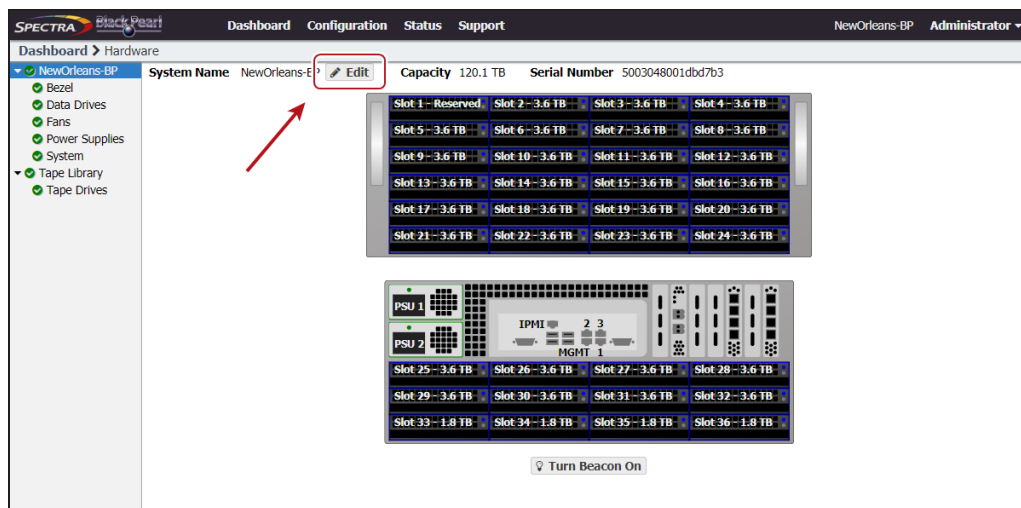


Figure 2 The BlackPearl user interface Hardware screen.

- 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**.

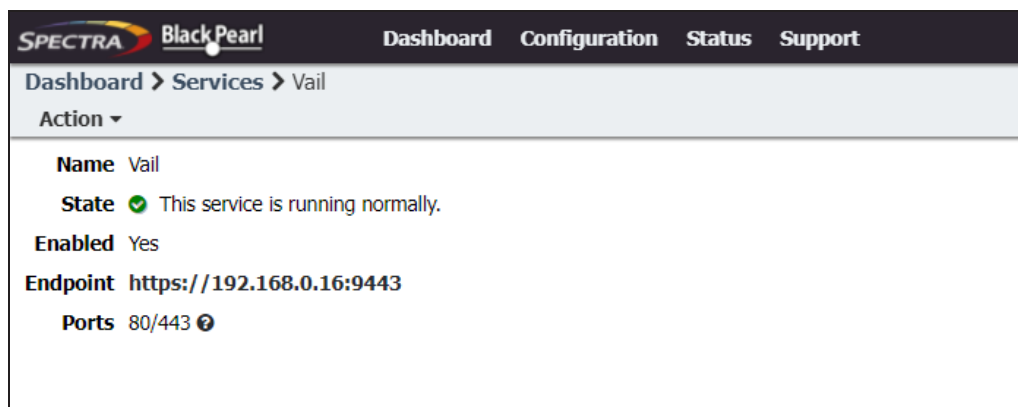


Figure 3 The Vail service details screen.

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 Vail application, the other pair is used by the BlackPearl DS3 service. If you change the pair of ports for the 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.

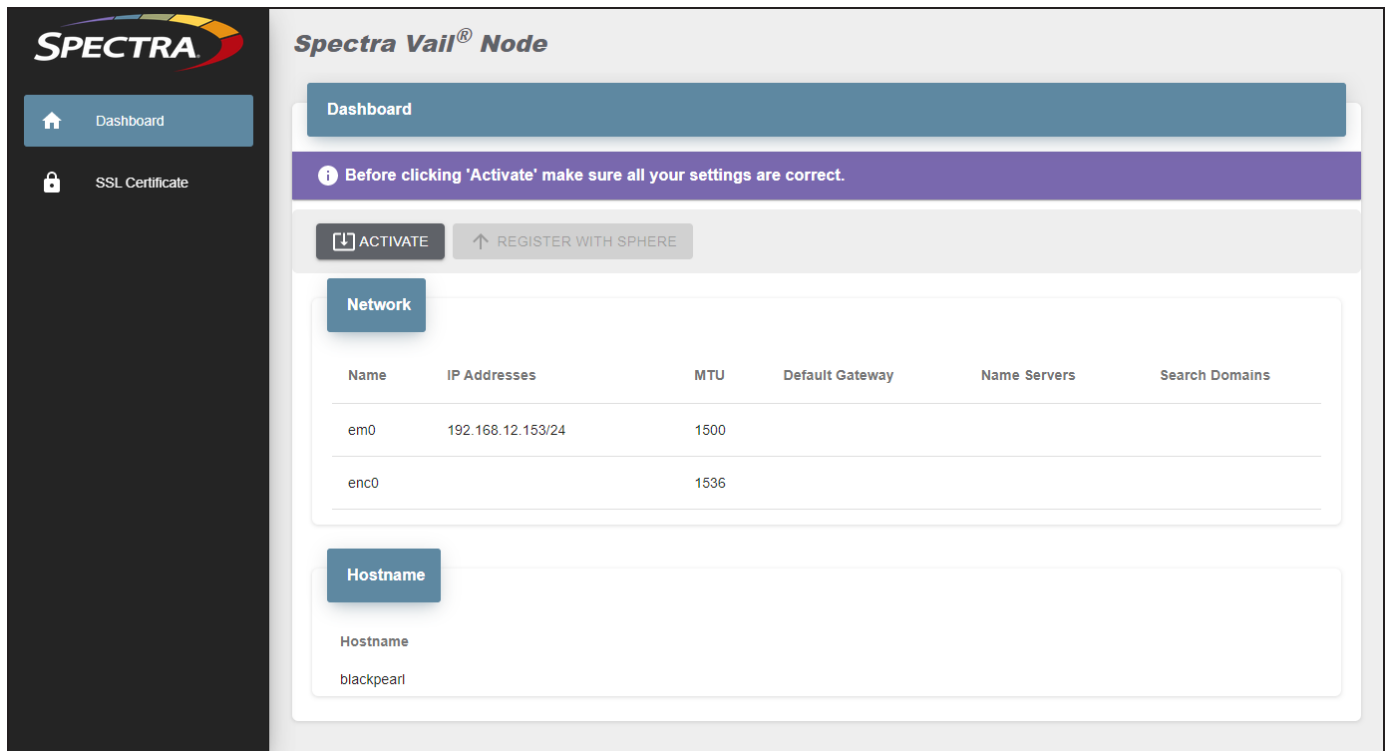


Figure 4 The Vail Node Dashboard - Activate and Register view.

9. If desired, update the SSL certificate before registering with the sphere:
 - a. In the taskbar of the Vail VM node management console, click **SSL Certificate**.
 - b. Under the **SSL Certificate** banner, click **Edit**.




Figure 5 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**.

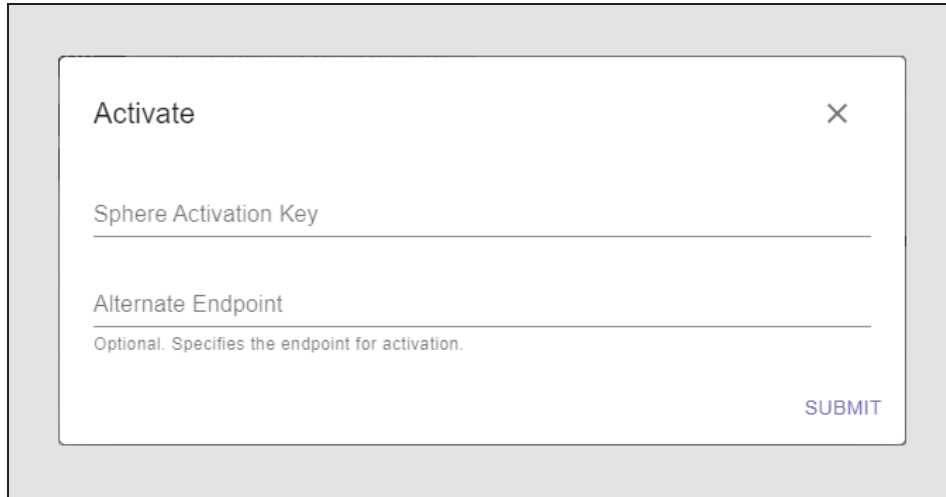


Figure 6 The Activate screen.

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.

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

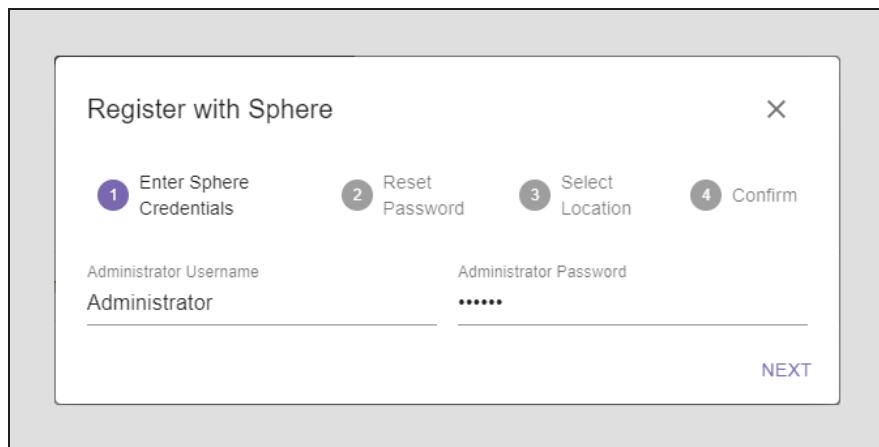


Figure 7 The Register with Sphere - Credentials screen.

14. 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.

15. 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 17](#).

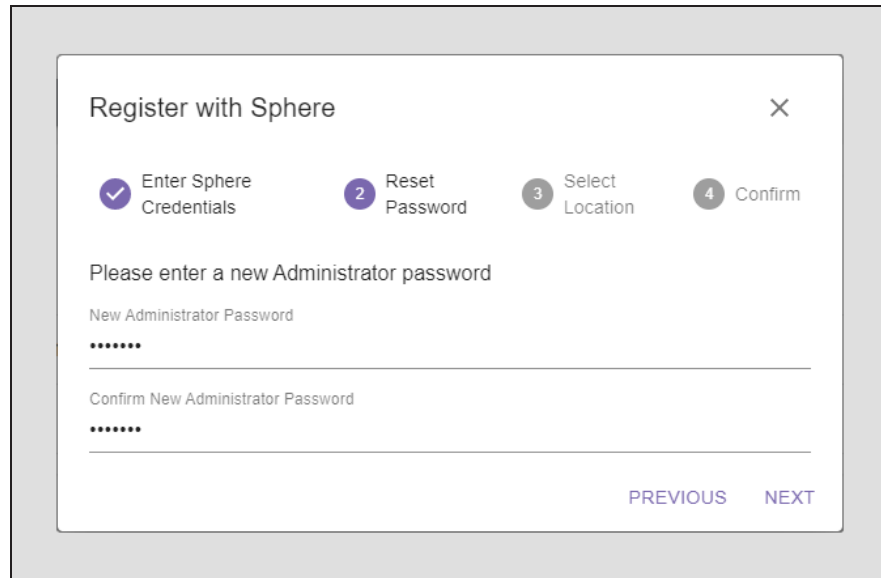
The screenshot shows a modal window titled "Register with Sphere" with a close button (X) in the top right corner. At the top, there is a progress bar with four steps: 1. Enter Sphere Credentials (checked), 2. Reset Password (active), 3. Select Location, and 4. Confirm. Below the progress bar, the text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

Figure 8 The Register with Sphere - Reset Password screen.

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

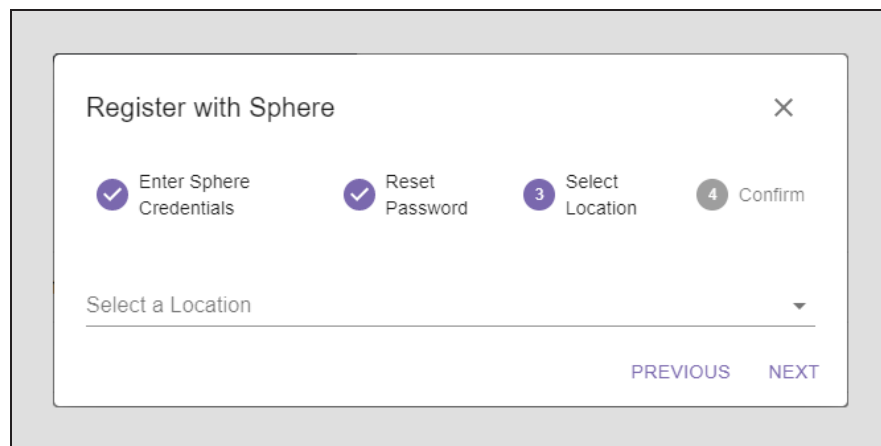
The screenshot shows the same "Register with Sphere" modal window. The progress bar now shows: 1. Enter Sphere Credentials (checked), 2. Reset Password (checked), 3. Select Location (active), and 4. Confirm. Below the progress bar, there is a dropdown menu labeled "Select a Location". At the bottom right, there are two buttons: "PREVIOUS" and "NEXT".

Figure 9 The Register with Sphere - Select Location screen.

17. On the Select Location screen, choose to create a new location, or to use an existing location:

- **Create a New Location below**
- **Select an Existing Location on page 38**

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 10 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**.

Figure 11 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 3 - CONFIGURE THE VAIL APPLICATION

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

Log In to the Vail Management Console	40
Vail Sphere Configuration Paths	41
Create Storage	42
Create BlackPearl Storage	43
Create BlackPearl Bucket Storage	43
Create Vail S3 NAS Storage	47
Create Cloud Storage	48
Create AWS S3 Cloud Storage	48
Create Microsoft Azure Cloud Storage	51
Create Google Cloud Platform Storage	53
Create Other S3 Cloud Storage	55
Create a Lifecycle	58
Create a Vail Bucket	66
Configure an Object Storage Browser	73
Configure S3 Browser	73
Configure Cyberduck Object Storage Browser	74

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 S3 solution management console, select **Configuration > Services**, then double-click the Vail service, and click the **Endpoint** URL displayed on the Vail Service screen.

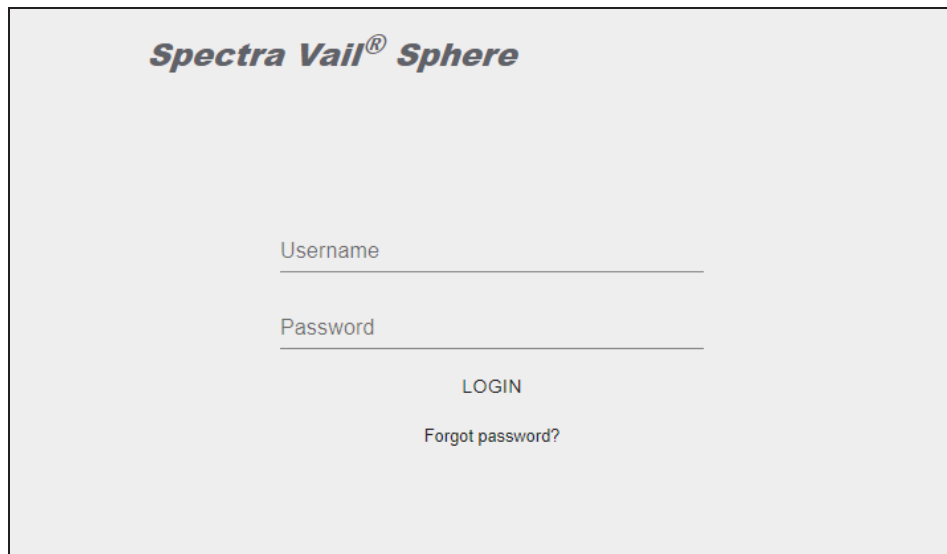


Figure 12 The Spectra Vail Sphere Login screen.

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

2. Enter the **Username** and **Password** you specified when you registered the first BlackPearl system with the Vail sphere.
3. Click **LOGIN**.

VAIL SPHERE CONFIGURATION PATHS

The Vail sphere has different configuration paths depending on if you are configuring a standalone Vail bucket, or if you plan to link Vail bucket to an existing bucket on BlackPearl or Cloud storage.

Use one of the checklists below to navigate through the two configuration paths.

Standalone Vail Bucket Path

Note: The Spectra Vail User Guide follows this configuration path in sequential page order.

Order		Action
1	<input type="checkbox"/>	Create Storage on the next page
2	<input type="checkbox"/>	Create a Lifecycle on page 58
3	<input type="checkbox"/>	Create a Vail Bucket on page 66

Linked Vail Bucket Path

Note: This configuration path requires you to jump between sections in the Spectra Vail User Guide out of sequential page order.

Order		Action
1	<input type="checkbox"/>	Create a Vail Bucket on page 66
2	<input type="checkbox"/>	Create Storage on the next page
3	<input type="checkbox"/>	Create a Lifecycle on page 58

CREATE STORAGE

Storage is used by the 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 S3 solution, or block VM storage such as a 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 48**
- **Create Vail VM Node Storage on page 103**

CREATE BLACKPEARL STORAGE

BlackPearl storage uses a bucket or Vail S3 share configured on a BlackPearl S3 solution. You can select the same BlackPearl S3 solution multiple times when creating BlackPearl storage, but each storage instance must use a unique bucket or Vail S3 share.

Before you can create BlackPearl storage in the Vail management console, you must register the BlackPearl S3 solution with the Spectra Vail application. See [Register a BlackPearl S3 with a Vail Sphere on page 31](#).

Use one of the sections below:

- **Create BlackPearl Bucket Storage below**
- **Create Vail S3 NAS Storage on page 47**

Create BlackPearl Bucket Storage

Before a BlackPearl storage endpoint can be configured in the Vail management console, a bucket must first be created on the BlackPearl S3 solution.

Create a BlackPearl Bucket

The instructions below assume a storage domain and data policy were previously configured on your BlackPearl S3 solution. For information on configuring a storage domain and data policy, see the [BlackPearl Nearline Gateway User Guide](#).

Here is how to create a bucket on a BlackPearl S3 solution:

1. Log in to your BlackPearl user interface.
2. Select **Configuration > Buckets**.
3. Select **Action > New**.
4. Enter the desired **Bucket Name**. 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.

5. Using the drop-down menus, select the bucket **Owner** and **Data Policy**.

- Notes:**
- In order to be used with the Vail application, buckets on the BlackPearl S3 solution must be configured to use just one type of storage. Buckets targeting different storage types are not supported.
 - Buckets must be associated with a data policy that does not use Versioning.
 - If the data policy for the bucket includes a tape storage domain, the storage domain must use LTFS Object Naming.
 - For additional information on configuring buckets, see the [Spectra BlackPearl Nearline Gateway User Guide](#).

6. Click **Create**.

7. Continue with [Create BlackPearl Storage in the Vail Sphere](#) below.

Create BlackPearl Storage in the Vail Sphere

Once you have created a bucket on your BlackPearl S3 solution, you can create BlackPearl storage in the Vail sphere.

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**.
3. Under the **Endpoint Storage** banner, click **Add**.

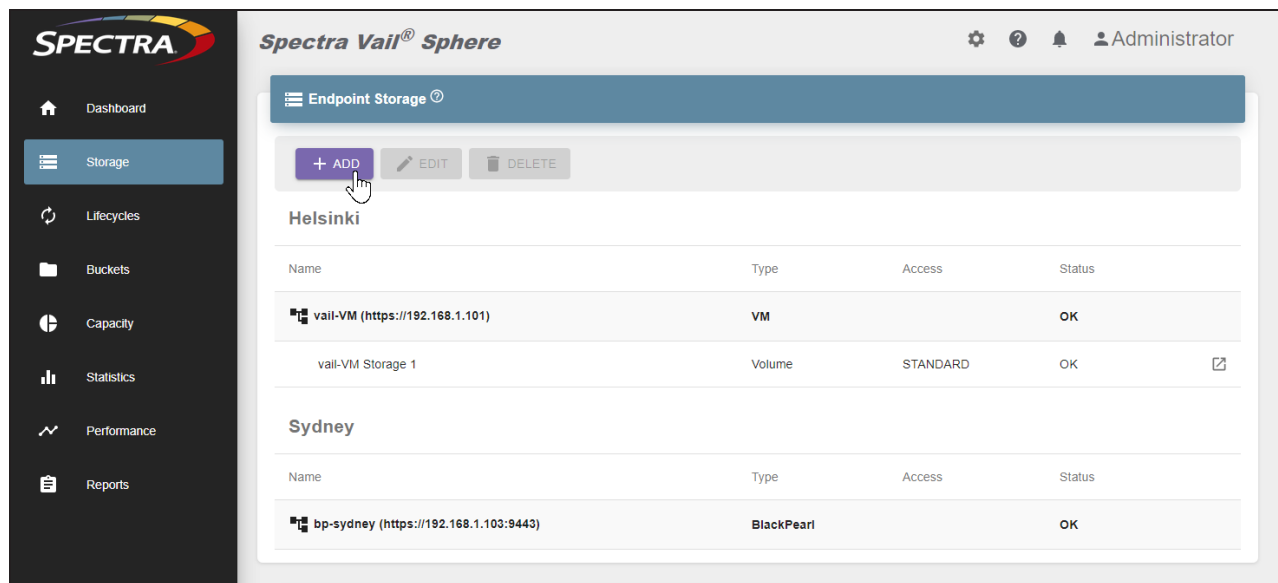


Figure 13 The Storage screen.

4. Using the drop-down menu, select the desired **BlackPearl** system and click **Next**.

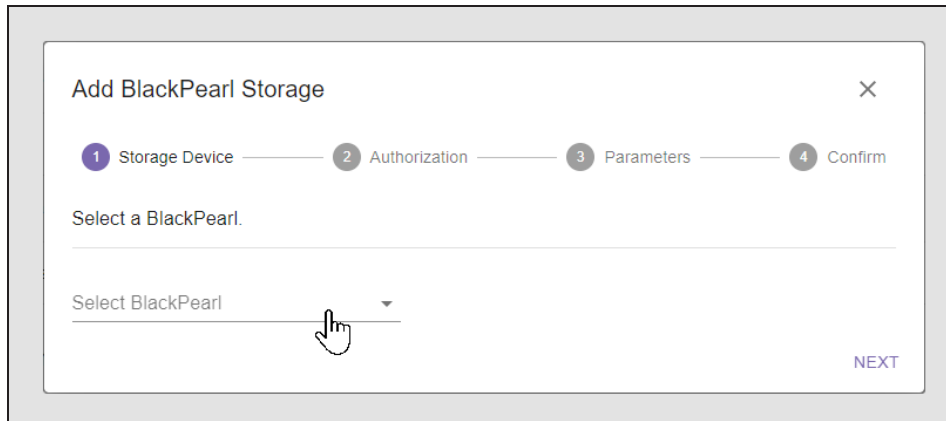
The screenshot shows a modal window titled "Add BlackPearl Storage" with a close button (X) in the top right. Below the title is a progress bar with four steps: 1. Storage Device (active), 2. Authorization, 3. Parameters, and 4. Confirm. The main content area says "Select a BlackPearl." followed by a text input field labeled "Select BlackPearl" with a dropdown arrow. A hand cursor is pointing at the dropdown arrow. A "NEXT" button is in the bottom right corner.

Figure 14 The Add BlackPearl Storage - Storage Device screen.

5. Retrieve the S3 credentials of the owner of the BlackPearl bucket created in [Create a BlackPearl Bucket on page 43](#).
 - a. In the BlackPearl user interface, select **Configuration > Users**.
 - b. Select the user that was configured as the bucket owner, and select **Action > Show S3 Credentials**.
6. Enter the **BlackPearl S3 Access ID** and **BlackPearl S3 Secret Key** of the owner of the BlackPearl bucket created in [Create a BlackPearl Bucket on page 43](#) and click **Next**.

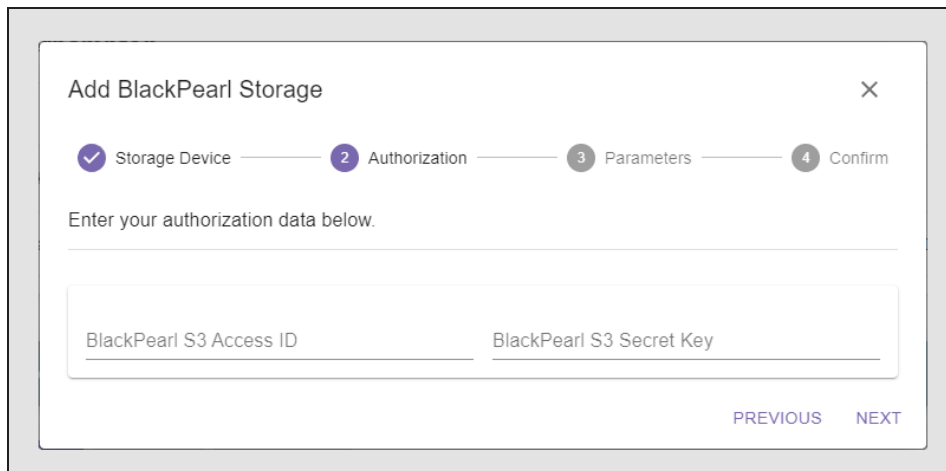
The screenshot shows the same modal window, now at Step 2: Authorization. The progress bar shows Step 1 as completed with a checkmark. The main content area says "Enter your authorization data below." followed by a text input field. Below this is a container with two text input fields: "BlackPearl S3 Access ID" and "BlackPearl S3 Secret Key". At the bottom right are "PREVIOUS" and "NEXT" buttons.

Figure 15 The Add BlackPearl Storage - Authorization screen.

7. Use the **Select BlackPearl Bucket** drop-down menu to select a previously configured bucket on the BlackPearl S3 solution.

Figure 16 The Add BlackPearl Storage - Parameters screen.

8. The **Storage Name** field is automatically populated with the bucket name. If desired you can edit the **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.

9. If you are creating a linked bucket, use the **Link to Bucket** drop-down menu to select the Vail bucket to which you want to link with the BlackPearl storage. When a Vail bucket is linked to a BlackPearl bucket, the Vail application discovers any existing data that was previously written to the BlackPearl bucket.

Note: If you are configuring a linked Vail bucket, the bucket must be created prior to creating the BlackPearl storage. See [Linked Vail Bucket Path on page 41](#).

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.

Note: These fields do not display if you selected linked storage in [Step 9](#).

11. Click **Next**.

12. Review the configuration, and click **Submit** to create the BlackPearl storage.

Create Vail S3 NAS Storage

The instructions below assume a storage pool and volume were previously configured on your BlackPearl S3 solution. For information on configuring a storage pool and volume, see the [BlackPearl Nearline Gateway User Guide](#).

Note: You cannot create Vail S3 NAS storage on a volume that currently has a share, or if the volume had a share assigned previously.

Here is how to create a Vail S3 share on a BlackPearl S3 solution:

1. Log in to your BlackPearl user interface.
2. Select **Configuration > NAS > Shares > Vail S3**.

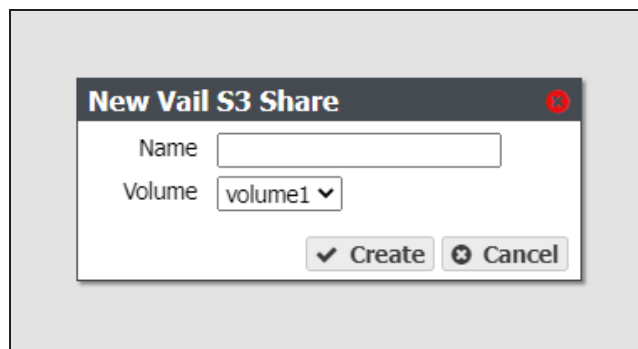


Figure 17 The New Vail S3 Share window.

3. Select **Action > New**.
4. Enter the desired **Name**.
5. Using the drop-down menu, select a **Volume** on the BlackPearl S3 solution.
6. Click **Create**.

The Vail S3 share immediately displays in the Vail management console and does not require manual configuration.

CREATE CLOUD STORAGE

Use one of the sections below to configure cloud storage:

- **Create AWS S3 Cloud Storage below**
- **Create Microsoft Azure Cloud Storage on page 51**
- **Create Google Cloud Platform Storage on page 53**
- **Create Other S3 Cloud Storage on page 55**

Create AWS S3 Cloud Storage

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

Here is how to create AWS S3 cloud storage:

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




Figure 18 The Cloud Storage pane.

3. Use the **Select Cloud Provider** drop-down menu to select **AWS S3**, and click **Next**.

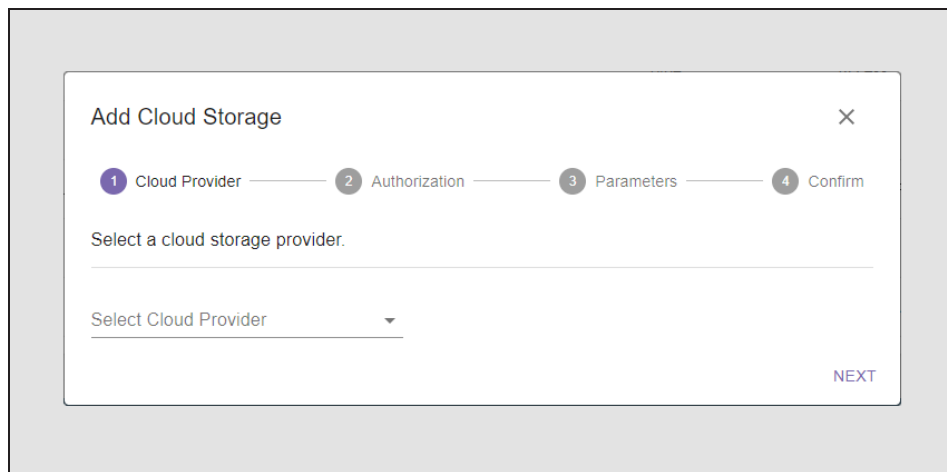


Figure 19 The Add Cloud Storage - Cloud Provider screen.

4. Select the desired authorization used to access cloud storage.

Figure 20 The Add Cloud Storage - AWS Authorization screen.

Option	Description
Use Credentials	<p>This option uses the AWS region and credentials of an AWS account to connect to the cloud storage.</p> <ul style="list-style-type: none"> To use the credentials of the AWS account associated with the Vail sphere administrator, leave the Region, AWS Access Key ID, and AWS Secret Access Key fields blank. To connect to cloud storage associated with a different AWS account, enter the Region, AWS Access Key ID, and AWS Secret Access Key of the account.
Use IAM Role	<p>This option uses the credentials of an IAM user under the main AWS account.</p> <ul style="list-style-type: none"> To use the credentials of an IAM user under the AWS account associated with the Vail sphere administrator account, leave the Region, AWS IAM Role ARN, and AWS IAM Role External ID fields blank. To connect the cloud endpoint to cloud storage using and IAM user under another AWS account, enter the Region and AWS IAM Role ARN. Optionally, enter the AWS IAM Role External ID of the account. In order to add the AWS cloud storage using an IAM role, a policy must be created in the other AWS account with a matching external ID to allow the IAM user to assume the role of the other AWS account.

5. Click **Next**.

6. Use the **Select 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 Vail bucket that has versioning disabled. Although the AWS bucket is capable of storing multiple versions of an object, if the Vail bucket does not have versioning enabled, only the latest version is preserved in the AWS bucket.

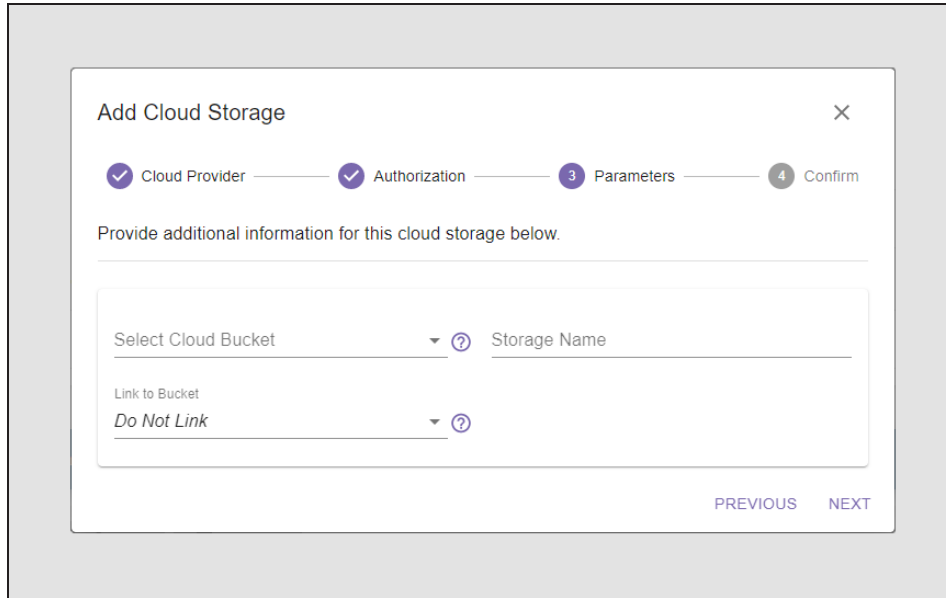


Figure 21 The Add Cloud Storage - Parameters screen.

7. The **Storage Name** field is automatically populated with the bucket name configured 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.

Spectra Logic recommends using names that include type of cloud storage, location, and storage class.

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

8. Use the **Link to Bucket** drop-down menu to select the Vail bucket which you want to link with the AWS S3 storage endpoint.

Note: If you want to link to a Vail bucket, the bucket must be created prior to creating the AWS S3 storage.



IMPORTANT

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

9. 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 Vail bucket. If the 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.
- By default the cloud storage target is created with the Storage Class set to Standard. If desired, you can edit the cloud storage target to change the Storage Class. See [Edit Google Cloud Platform Storage](#) on page 174.

Create Microsoft Azure Cloud Storage

In Vail, Microsoft Azure cloud storage uses a previously configured Azure endpoint target 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**.

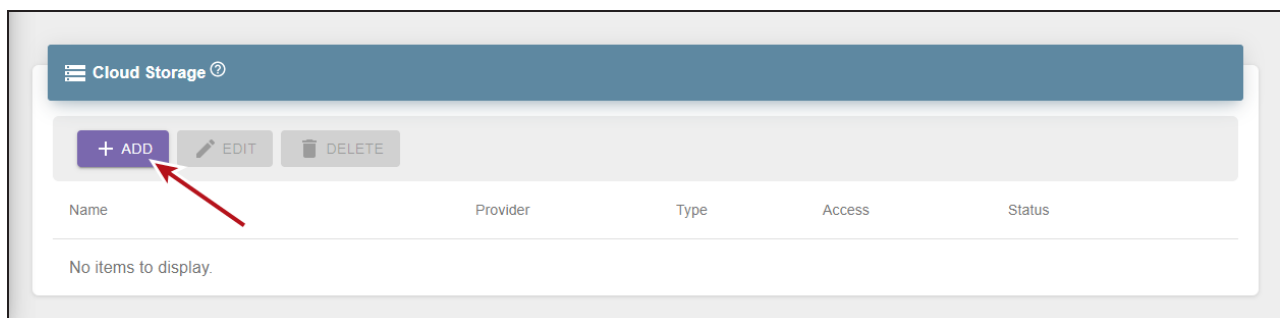


Figure 22 The Cloud Storage pane.

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

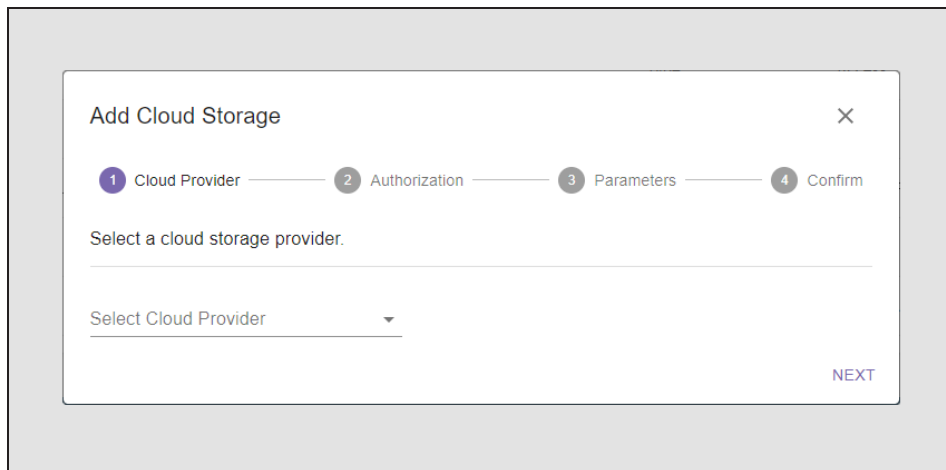
The screenshot shows a modal dialog titled "Add Cloud Storage" with a close button (X) in the top right corner. Below the title is a progress bar with four steps: 1. Cloud Provider (active, highlighted with a blue circle), 2. Authorization, 3. Parameters, and 4. Confirm. Below the progress bar, the text "Select a cloud storage provider." is displayed. Underneath is a drop-down menu labeled "Select Cloud Provider" with a downward arrow. In the bottom right corner, there is a blue "NEXT" button.

Figure 23 The Add Cloud Storage - Cloud Provider screen.

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

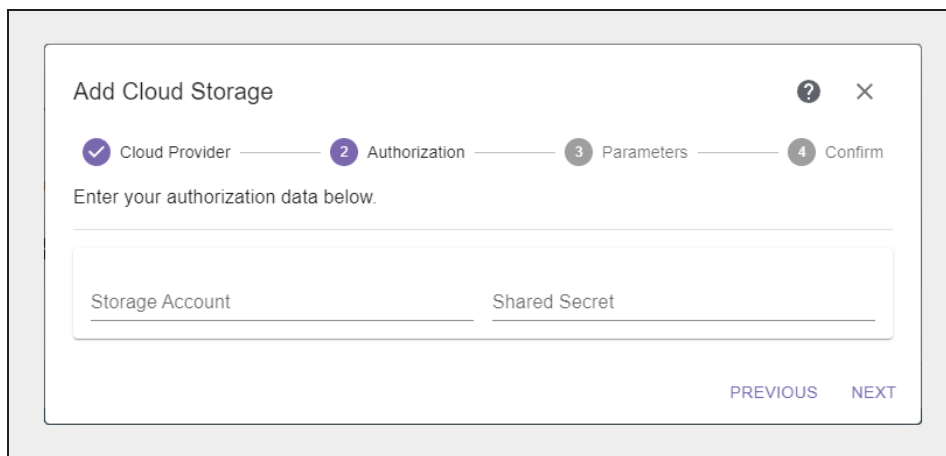
The screenshot shows the same "Add Cloud Storage" modal dialog, but now the "Authorization" step (2) is active, indicated by a blue checkmark in the progress bar. The "Cloud Provider" step is now marked with a question mark. Below the progress bar, the text "Enter your authorization data below." is displayed. Underneath is a form with two input fields: "Storage Account" and "Shared Secret". In the bottom right corner, there are two blue buttons: "PREVIOUS" and "NEXT".

Figure 24 The Add Cloud Storage - Authorization screen.

5. Click **Next**.

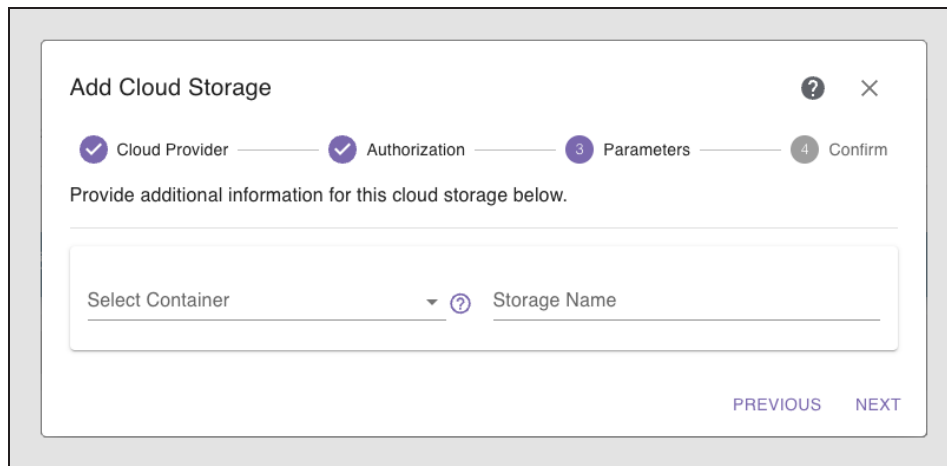


Figure 25 The Add Cloud Storage - Parameters screen.

6. Using the **Select Container** drop-down menu, select a previously created container on the Azure storage target.
7. Enter a **Storage Name**, then click **Next**.
8. 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**.

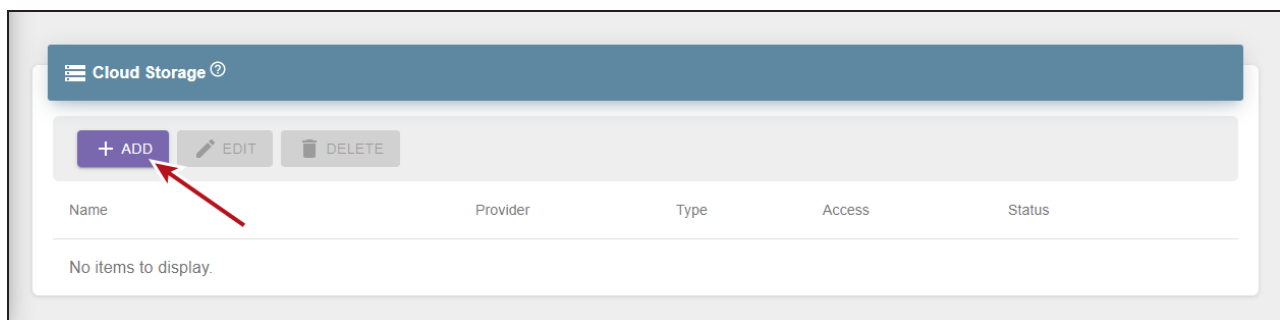
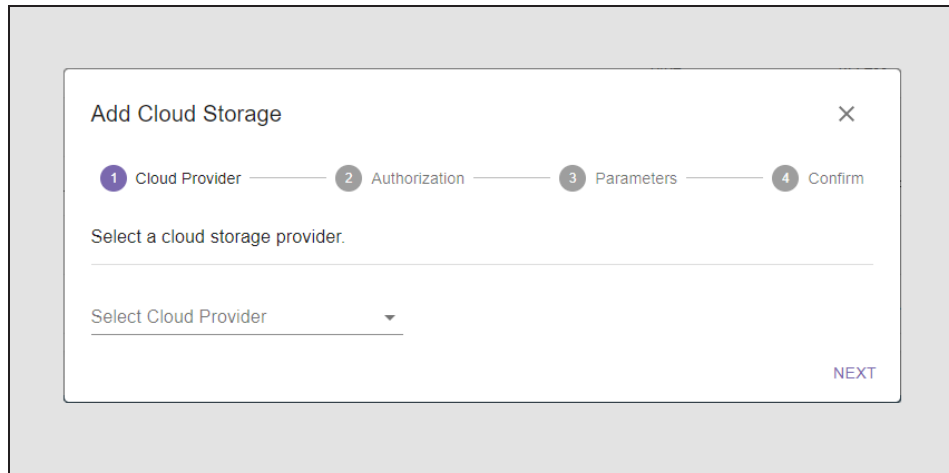


Figure 26 The Cloud Storage pane.

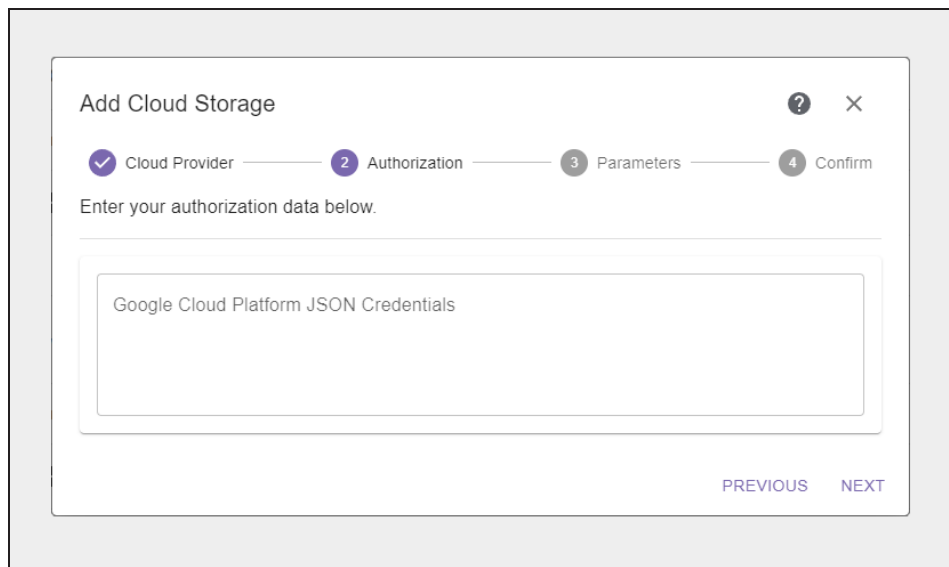
3. Use the **Select Cloud Provider** drop-down menu to select **Google Cloud Storage**, and click **Next**.



The screenshot shows a modal dialog titled "Add Cloud Storage" with a close button (X) in the top right corner. Below the title is a progress bar with four steps: 1. Cloud Provider (active, highlighted with a purple circle), 2. Authorization, 3. Parameters, and 4. Confirm. The main text says "Select a cloud storage provider." Below this is a dropdown menu labeled "Select Cloud Provider" with a downward arrow. In the bottom right corner, there is a purple "NEXT" button.

Figure 27 The Add Cloud Storage - Cloud Provider screen.

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



The screenshot shows the same "Add Cloud Storage" modal dialog, now at Step 2: Authorization. The progress bar shows Step 1 as completed (checked) and Step 2 as active. The main text says "Enter your authorization data below." Below this is a large text input field with the placeholder text "Google Cloud Platform JSON Credentials". In the bottom right corner, there are two purple buttons: "PREVIOUS" and "NEXT".

Figure 28 The Add Cloud Storage - Authorization screen.

5. Click **Next**.

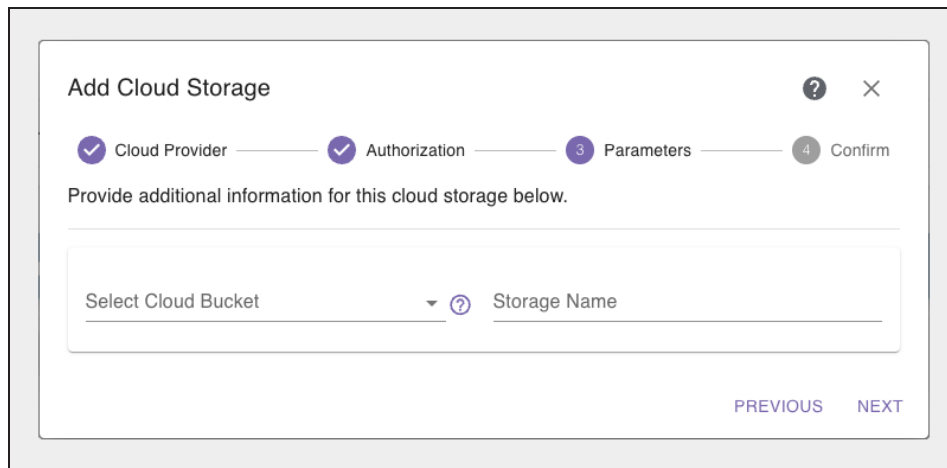


Figure 29 The Add Cloud Storage - Parameters screen.

6. Using the **Select Cloud Bucket** drop-down, select a previously created bucket in the Google Cloud Storage target.
7. Enter a **Storage Name**, then click **Next**.
8. Verify the information for the cloud storage is correct, and click **Submit**.

Create Other S3 Cloud Storage

Cloud storage that is not an AWS, Azure, or Google Cloud endpoint is configured as other third-party S3 cloud storage.

- Notes:**
- The bucket on the cloud storage target must be configured to use versioning.
 - You cannot create a linked bucket for use with other third-party S3 storage.

Here is how to create other third-party S3 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 Cloud Provider** drop-down menu to select **Other 3rd Party S3** and click **Next**.

The screenshot shows a modal window titled "Add Cloud Storage" with a close button (X) in the top right corner. Below the title is a progress bar with four steps: 1. Cloud Provider (active, highlighted with a blue circle), 2. Authorization, 3. Parameters, and 4. Confirm. The main text says "Select a cloud storage provider." Below this is a dropdown menu labeled "Select Cloud Provider" with a downward arrow. In the bottom right corner, there is a blue "NEXT" button.

Figure 30 The Add Cloud Storage - Cloud Provider screen.

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

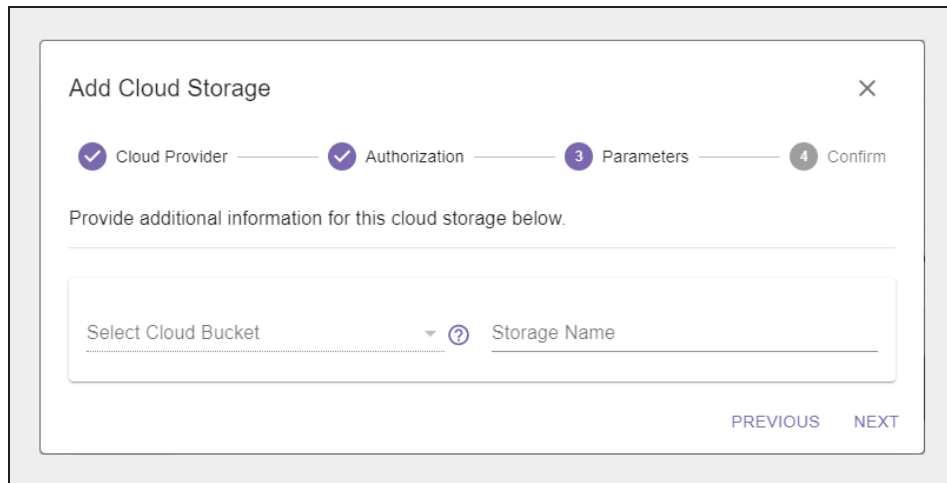
The screenshot shows the same "Add Cloud Storage" modal window, now at Step 2: Authorization. The progress bar shows Step 1 (Cloud Provider) as completed with a checkmark, and Step 2 (Authorization) as the current step. The main text says "Enter your authorization data below." Below this is a form with three input fields: "Data Path Endpoint", "Access Key", and "Secret Key". In the bottom right corner, there are blue "PREVIOUS" and "NEXT" buttons.

Figure 31 The Add Cloud Storage - 3rd Party Authorization screen.

5. Enter the **Access Key** and **Secret Key** for the administrator of the cloud endpoint.
6. Click **Next**.

- Using the **Select Cloud Bucket** drop-down menu, select a bucket previously configured on the cloud endpoint.

Note: Versioning must be enabled on the target bucket.



The screenshot shows a modal window titled "Add Cloud Storage" with a close button (X) in the top right corner. A progress bar at the top indicates four steps: "Cloud Provider" (completed with a checkmark), "Authorization" (completed with a checkmark), "Parameters" (active, highlighted with a blue circle and the number 3), and "Confirm" (disabled, greyed out with a circle and the number 4). Below the progress bar, the text "Provide additional information for this cloud storage below." is displayed. The main form area contains two fields: "Select Cloud Bucket" (a dropdown menu with a question mark icon) and "Storage Name" (a text input field). At the bottom right of the form, there are two buttons: "PREVIOUS" and "NEXT".

Figure 32 The Add Cloud Storage - 3rd Party Parameters screen.

- The **Storage Name** is automatically populated with the name of the bucket selected in [Step 7](#). If desired, you can change the **Storage Name**.
- Click **Next**.
- Verify the information for the cloud storage is correct, and 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.

Here is how to create a lifecycle:

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

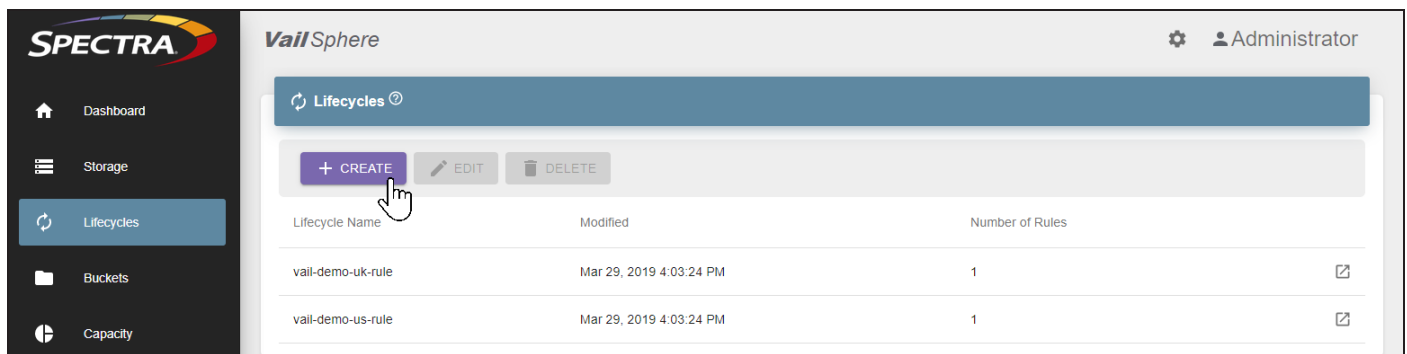


Figure 33 The Lifecycles screen.

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.

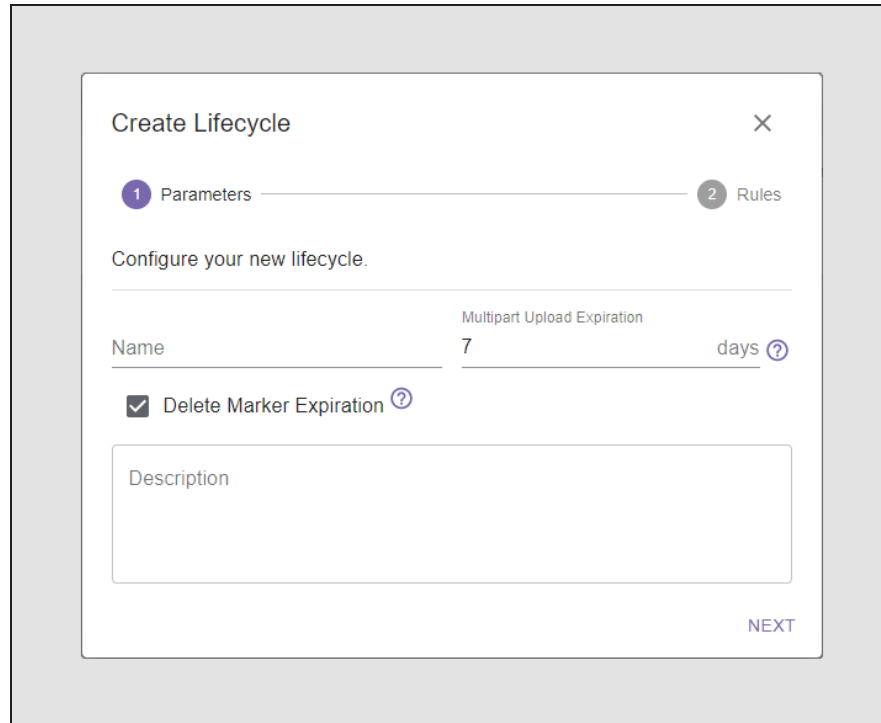


Figure 34 The Create Lifecycle - Parameters screen.

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. Select or clear **Delete Marker Expiration**. A delete marker keeps track of deletions of versioned objects so that S3 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.
6. Use the **Description** field to enter any additional information.
7. Click **Next**.
8. 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 63**

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**.

Create Lifecycle ? X

✓ Parameters 2 Rules

Define your rules for **Lifecycle2** ?

Note: Rules will be sorted by "Apply After" and "Versioning" values after submission. Maximum number of rules is 5.

Placement Rule

Name

Select Destination Storage ?

☐ Delete clones not on selected destination storage Destination Count All

Apply After days ? Versioning All

Include ?

Exclude ?

NEW PLACEMENT RULE | NEW DELETION RULE

PREVIOUS SUBMIT

Figure 35 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 to **Delete clones not on selected destination 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 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. Enter a number of days in the **Apply After** field. This setting controls how many days the Spectra Vail application waits before copying object data to storage destinations to maintain the specified number of copies. The timer for the **Apply After** setting takes effect at midnight UTC on the day data is written.

Note: For example, with a rule configured with **Apply After** set for two days, if you write data at 8 PM UTC, after four hours, the **Apply After** clock starts, and the data is copied 48 hours later. In this example, the total time before data is written is 52 hours.

Note: If you set **Apply After** to 0 (zero), data is copied immediately.

Note: The maximum value is 9999 days.

7. Using the **Versioning** drop-down menu, select what version of an object to copy when the rule executes.

Setting	Description
All	All versions of an object are copied.
Latest	Only the latest version of an object is copied.
Previous	All versions of an object, except for the latest version, are copied.

8. Optionally, enter a regular expression to **Include** in the clone job. Any object that matches this expression is included in the clone job.

For example:

- Entering `"/archive/images/production"` includes all objects that begin with the specified string.
- Entering `".txt$"` includes all objects that end with a .txt extension.

Note: Leave the field blank to include all objects.

9. Optionally, enter a regular expression to **Exclude** in the clone job. Any object that matches this expression is excluded from the clone job.

For example:

- Entering `"/archive/images/production"` excludes all objects that begin with the specified string.
- Entering `".txt$"` excludes all objects that end with a .txt extension.

Note: Leave the field blank to include all objects.

10. 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 **Lifecycle2**

Note: Rules will be sorted by "Apply After" and "Versioning" values after submission. Maximum number of rules is 5.

Deletion Rule

Name

Apply After days Versioning All

Noncurrent Versions to Keep count

Include

Exclude

NEW PLACEMENT RULE | NEW DELETION RULE

PREVIOUS SUBMIT

Figure 36 The Create Lifecycle - Deletion Rule screen.

2. Enter the desired **Name**.

3. Enter a number of days in the **Apply After** field. This setting controls how many days the Spectra Vail application waits before deleting data associated with the lifecycle. The timer for the **Apply After** setting takes effect at midnight UTC on the day data is written.

Note: For example, with a rule configured with Apply After set for two days, if you write data at 8 PM UTC, after four hours, the Apply After clock starts, and the data is deleted 48 hours later. In total the data is deleted 52 hours after it was written.

Note: You must enter a value greater than 0.

Note: The maximum value is 9999 days.

4. Using the **Versioning** drop-down menu, select what version of an object to delete when the rule executes.

Setting	Description
All	All versions of an object are deleted.
Latest	Only the latest version of an object is deleted.
Previous	All versions of an object, except for the latest version, are deleted.

5. If you selected **Previous** in [Step 4](#), enter a number for the **Noncurrent Versions to Keep** setting. This setting controls the number of non-current versions that are kept and not deleted. When this limit is reached, any excess non-current versions are deleted when the **Apply After** setting is reached.

Note: For example, if Apply After is set to 1, then the most recent version of an object is kept along with a the single most recent previous version of that object. All other versions are deleted.

6. Optionally, enter a regular expression to **Include** in the expiration job. Any object that matches this expression is included in the expiration job.

For example:

- Entering `"^/archive/images/production"` includes all objects that begin with the specified string.
- Entering `"\\.txt$"` includes all objects that end with a .txt extension.

Note: Leave the field blank to include all objects.

7. Optionally, enter a regular expression to **Exclude** in the expiration job. Any object that matches this expression is excluded from the expiration job.

For example:

- Entering `"^/archive/images/production"` excludes all objects that begin with the specified string.
- Entering `"\\.txt$"` excludes all objects that end with a .txt extension.

Note: Leave the field blank to include all objects.

8. 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.

Vail buckets can also be linked to an existing bucket on a BlackPearl system or AWS S3 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. You cannot link a bucket to non-AWS cloud storage endpoints.

Note: When a Vail bucket is linked to an AWS cloud bucket, the 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 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 an AWS cloud bucket, when an object is added to an AWS cloud bucket, the Vail application creates a version of the object with a clone that references the object in the AWS bucket. Because the objects are linked, if the object is deleted in the Vail application, the object on the AWS cloud bucket is deleted, even if no lifecycle is defined. If there are multiple versions of the object in Vail, when the object is deleted, only the object on the AWS cloud bucket that matches the version deleted in Vail is deleted from the AWS 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 37 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 38 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 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	<p>New objects written to this bucket are always owned by the bucket owner configured in Step 9.</p> <ul style="list-style-type: none"> Access to this bucket and its objects is specified using only policies. <p>Note: This is the recommended setting.</p>
Bucket Owner Preferred	<p>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.</p> <ul style="list-style-type: none"> Access to this bucket and its objects can be specified using ACLs or policies.
Object Writer	<p>New objects written to the bucket are always owned by the object writer.</p> <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 10 on page 69](#), continue with [Step 12](#) below. Otherwise, skip to [Step 16 on page 70](#).
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 70](#).

Figure 39 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 40 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 **Restrict Public Buckets**. Enabling this setting ignores public and cross-account access for buckets with policies that grant public access to buckets and objects.
19. Click **Next**. If you selected **ACLs Disabled** in [Step 10](#) on page 69, skip to [Step 26](#) on page 72. Otherwise continue with [Step 20](#) below.
20. 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.

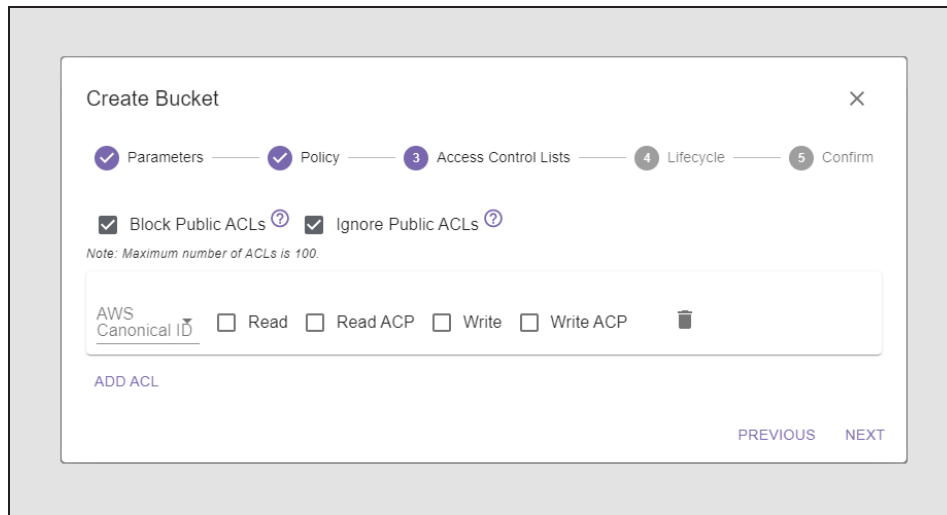


Figure 41 The Create Bucket - Access Control List screen.

21. 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](#).
22. 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](#) through [Step 22](#) to configure additional ACLs.

Note: Use the trashcan icon to remove an ACL.

- 23.** If desired, select or clear **Block Public ACLs**. Enabling this setting 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.
- 24.** If desired, select or clear **Ignore Public ACLs**. Enabling This setting ignores all ACLs that grant public access to objects or directories.
- 25.** Click **Next**.
- 26.** Using the **Select Lifecycle** drop-down menu, select a previously configured lifecycle and click **Next**.

Note: If you are creating a linked bucket and want to use the linked bucket as destination storage in a lifecycle, select None. After the linked storage is created and added to a lifecycle, you need to edit the bucket to select the desired lifecycle. See [Vail Sphere Configuration Paths on page 41](#) for more information.

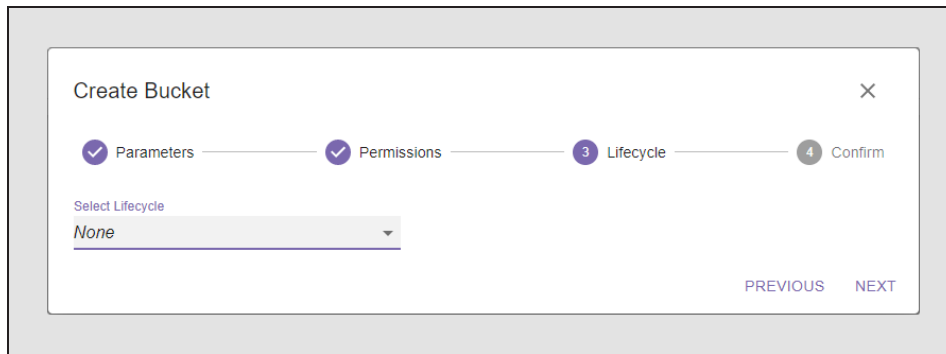


Figure 42 The Create Bucket - Lifecycle screen.

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

CONFIGURE AN OBJECT STORAGE BROWSER

Before you can access and transfer data to a BlackPearl S3 solution or 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**.

Add New Account online help

Enter new account details and click Add new account

Account Name:

 Assign any name to your account.

Account Type:

 Choose the storage you want to work with. Default is Amazon S3 Storage.

REST Endpoint:

 Specify S3-compatible API endpoint. It can be found in storage documentation. Example: rest.server.com:8080

Access Key ID:

 Required to sign the requests you send to Amazon S3, see more details at <https://s3browser.com/keys>

Secret Access Key:

 Required to sign the requests you send to Amazon S3, see more details at <https://s3browser.com/keys>

☐ Encrypt Access Keys with a password:

 Turn this option on if you want to protect your Access Keys with a master password.

☐ Use secure transfer (SSL/TLS)
 If checked, all communications with the storage will go through encrypted SSL/TLS channel

[Advanced S3-compatible storage settings](#)

Figure 43 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 S3 solution or 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**.

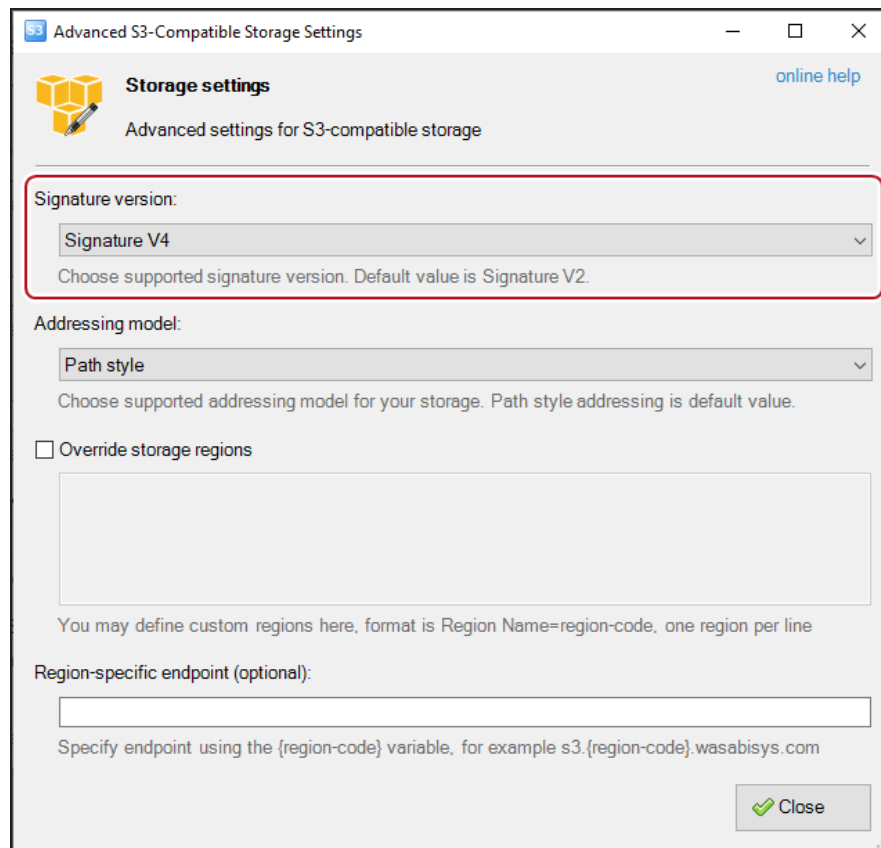


Figure 44 The Advanced S3-Compatible Storage Settings screen.

9. Using the **Signature Version** drop-down menu, select **Signature V4** and click **Close**.
10. Click **Add new account**. The S3 Browser retrieves the list of buckets configured on the Vail sphere (see [View Vail Bucket Details on page 141.](#))

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**.

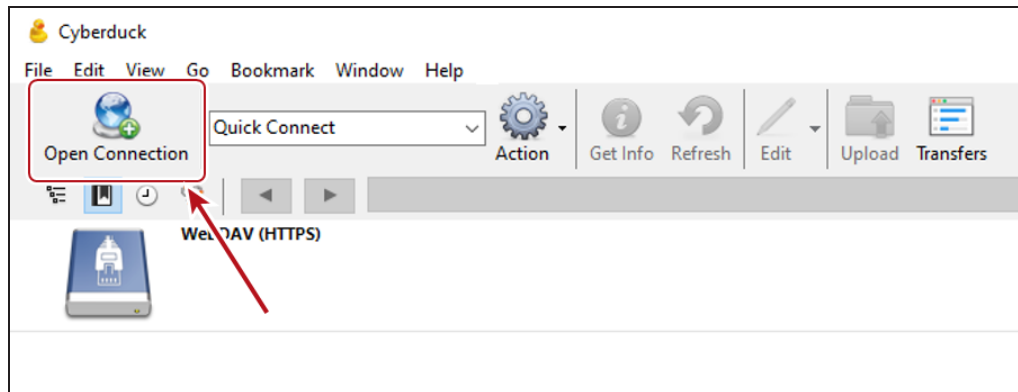


Figure 45 The Cyberduck Object Storage Browser home screen.

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




Figure 46 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 4 - 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 Accounts	77
Create a Sphere Administrator	77
Change a Sphere Administrator Password	79
Edit Sphere Administrator Attributes	81
Delete a Sphere Administrator	83
Configure & Manage IAM Accounts	84
Add an IAM Account	84
Edit an IAM Account	87
Delete an IAM Account Association	89
Configure & Manage IAM Users and Groups	90
Create an IAM User	90
View IAM User Details	91
Add an IAM User to an IAM Group	92
Remove an IAM User from an IAM Group	93
Delete an IAM User	94
Create an IAM Group	95
Delete an IAM Group	96
AWS Access Key Management	97
Create an Access Key	97
Enable an Access Key	98
Disable an Access Key	99
Delete an Access Key	100

CONFIGURE & MANAGE SPHERE ADMINISTRATOR ACCOUNTS

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.

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 47 The Dashboard screen - Navigation menu.

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

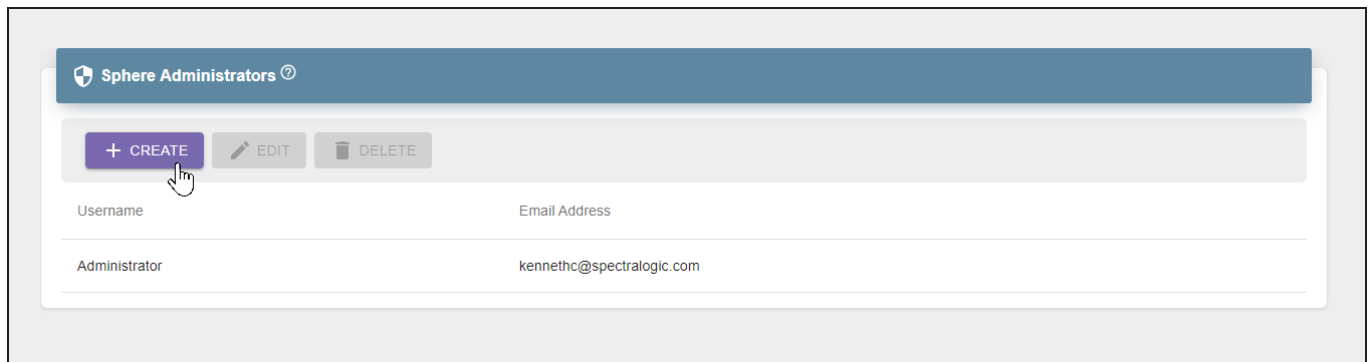


Figure 48 The Sphere Administrators pane.

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 the text 'Create Sphere Administrator' and a close button (X). Inside the dialog, there are two input fields: 'Username' and 'Email Address'. Below these fields is a section titled 'Select what types of emails this user wants to receive.' with three checkboxes: 'Info', 'Warning', and 'Error'. At the bottom right of the dialog is a purple 'SUBMIT' button.

Figure 49 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.
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)**.

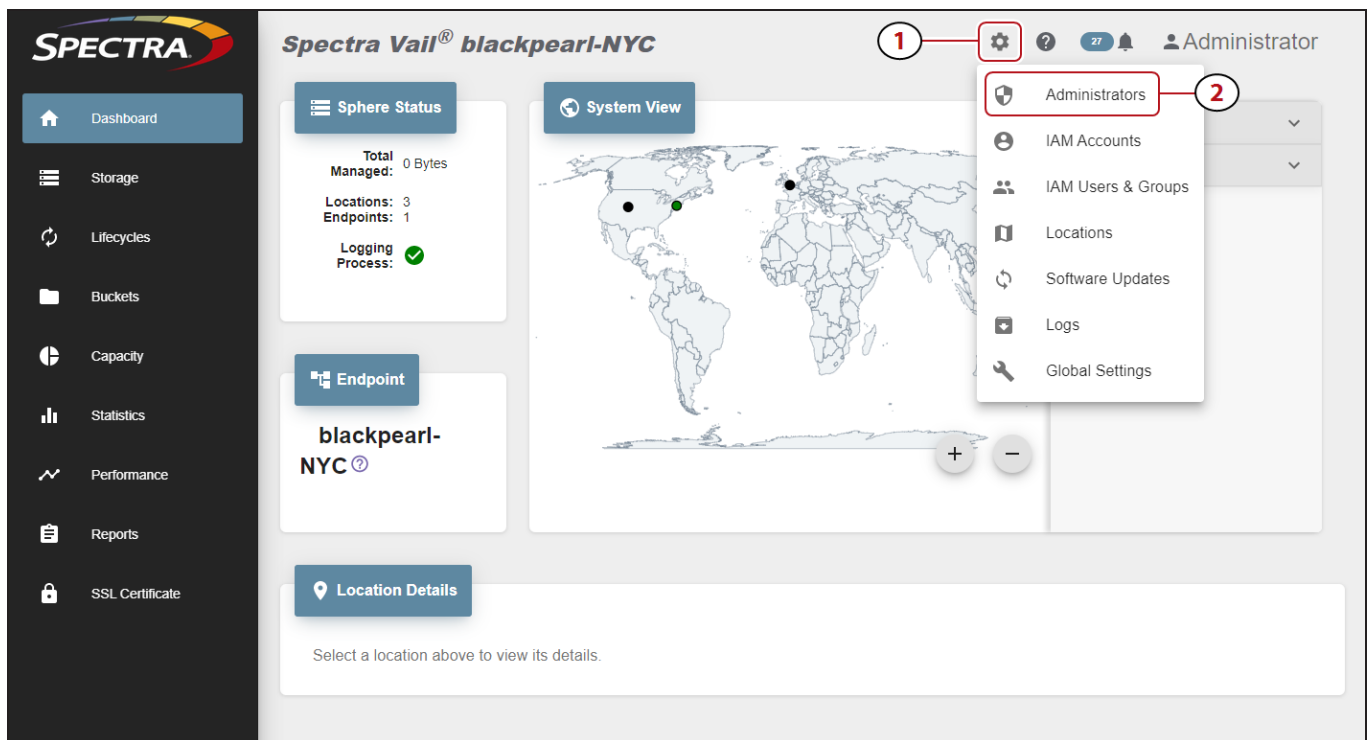


Figure 50 The Dashboard screen - Navigation menu.

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




Figure 51 The Sphere Administrators pane.

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

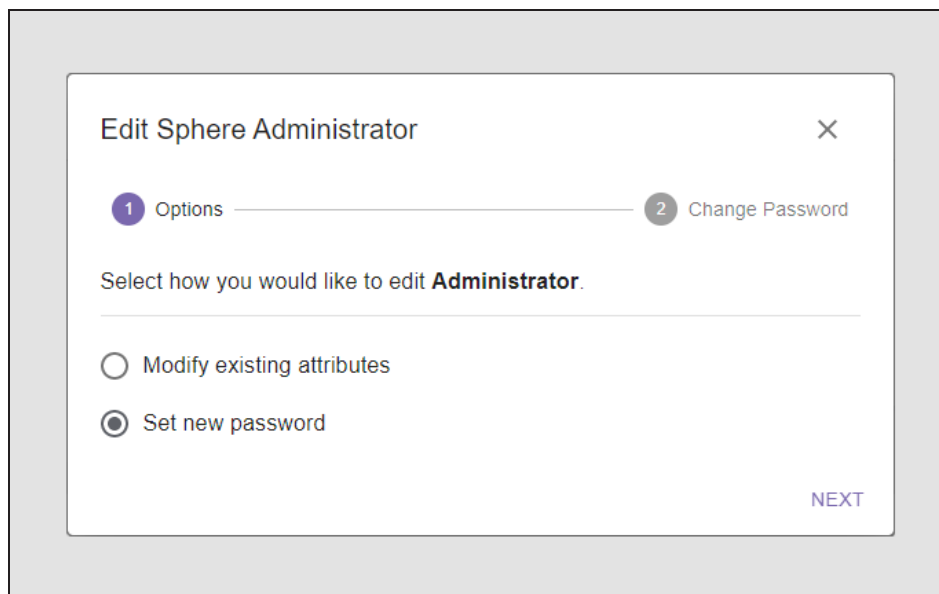


Figure 52 The Edit Sphere Administrator - Options screen.

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

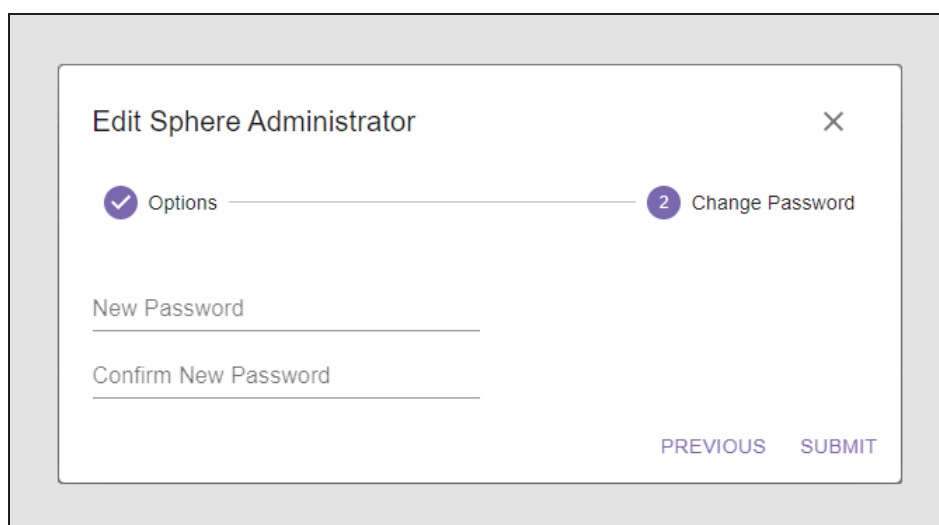


Figure 53 The Edit Sphere Administrator - Change Password screen.

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 54 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**.

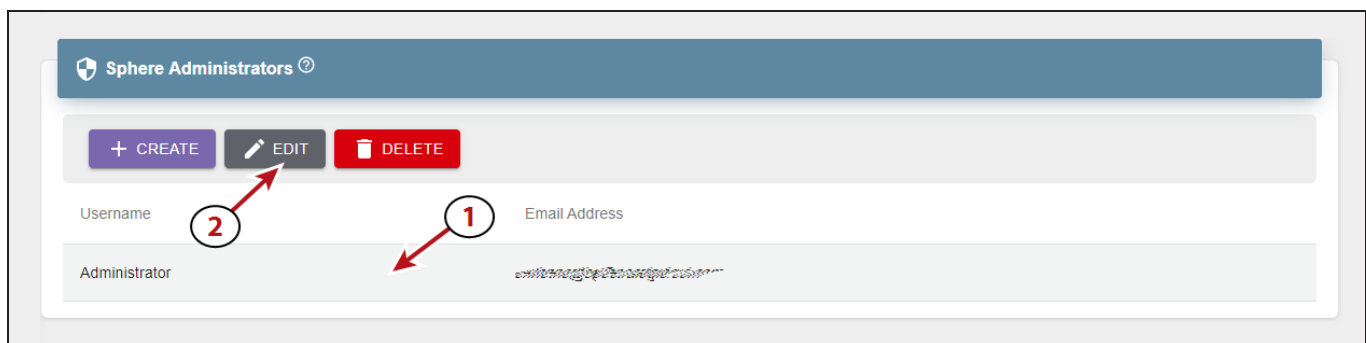
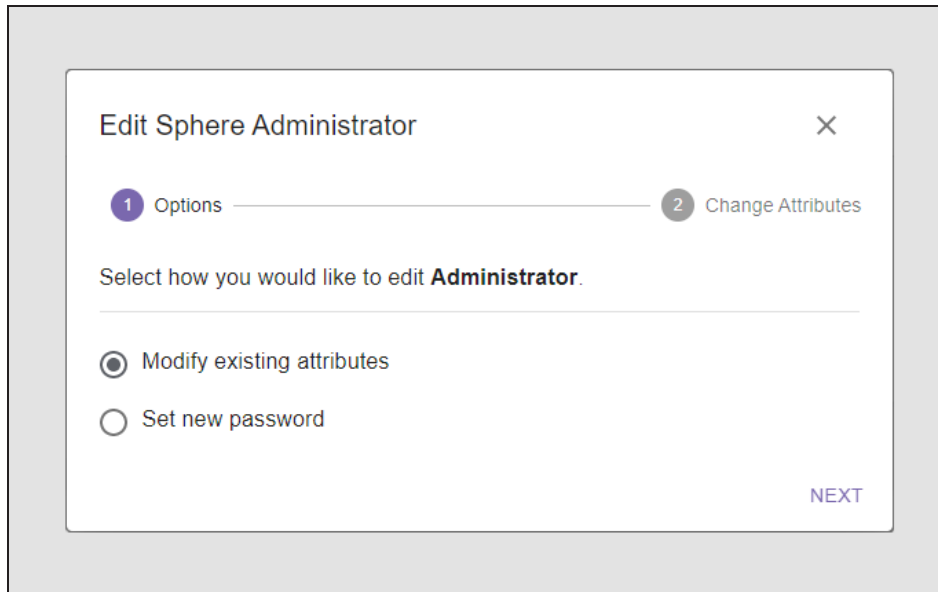


Figure 55 The Sphere Administrators pane.

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



Edit Sphere Administrator

1 Options ————— 2 Change Attributes

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

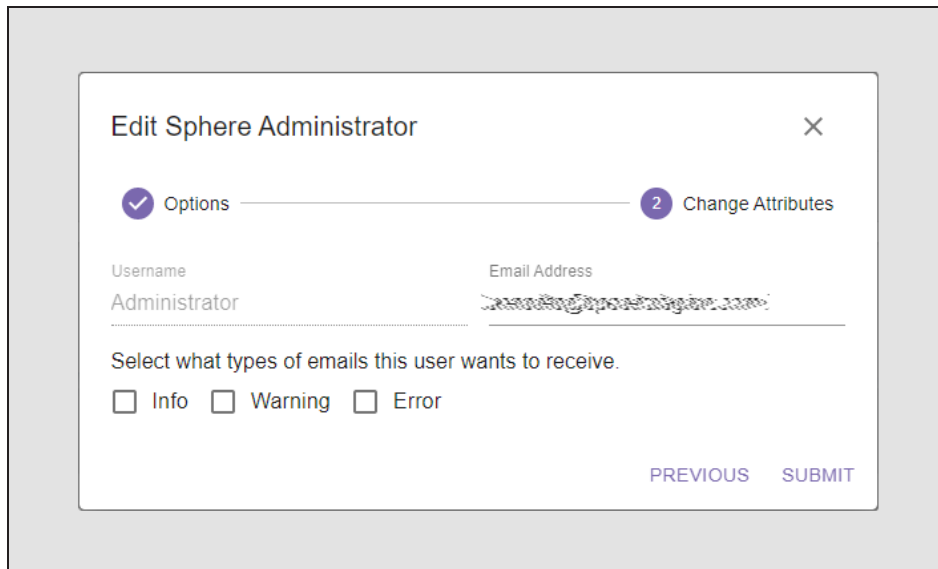
☒ Modify existing attributes

☐ Set new password

NEXT

Figure 56 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

✓ Options ————— 2 Change Attributes

Username Email Address

Administrator Administrator@spectravail.com

Select what types of emails this user wants to receive.

☐ Info ☐ Warning ☐ Error

PREVIOUS SUBMIT

Figure 57 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)**.

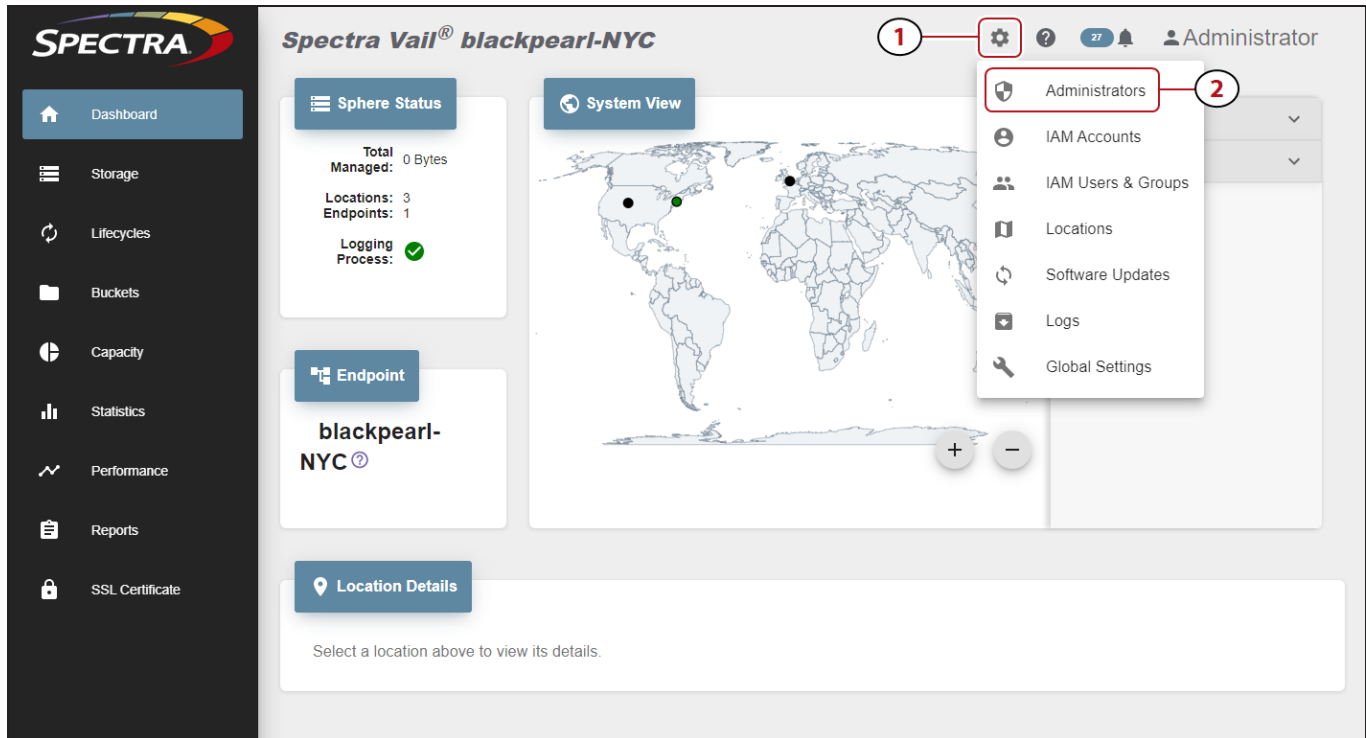


Figure 58 The Dashboard screen - Navigation menu.

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

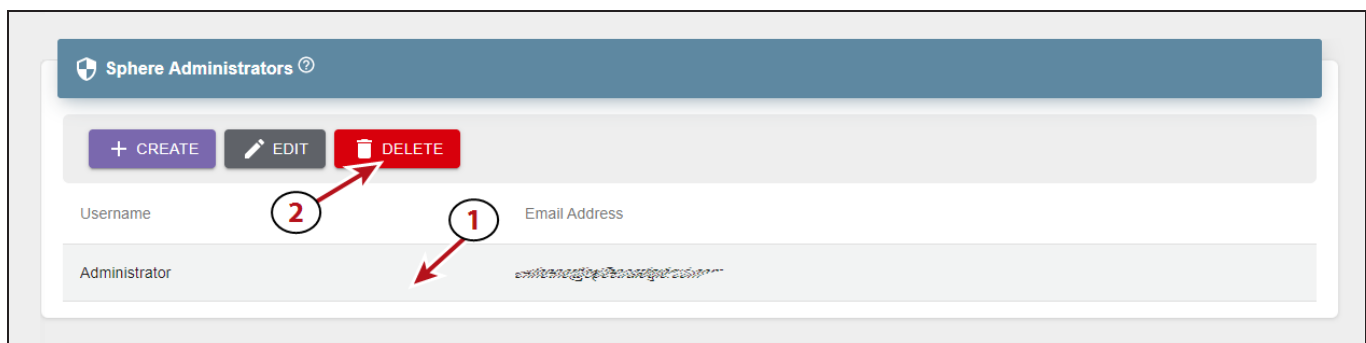


Figure 59 The Sphere Administrators pane.

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

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 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 Vail application.

Here is how to add 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 60 The Dashboard screen - Navigation menu.

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

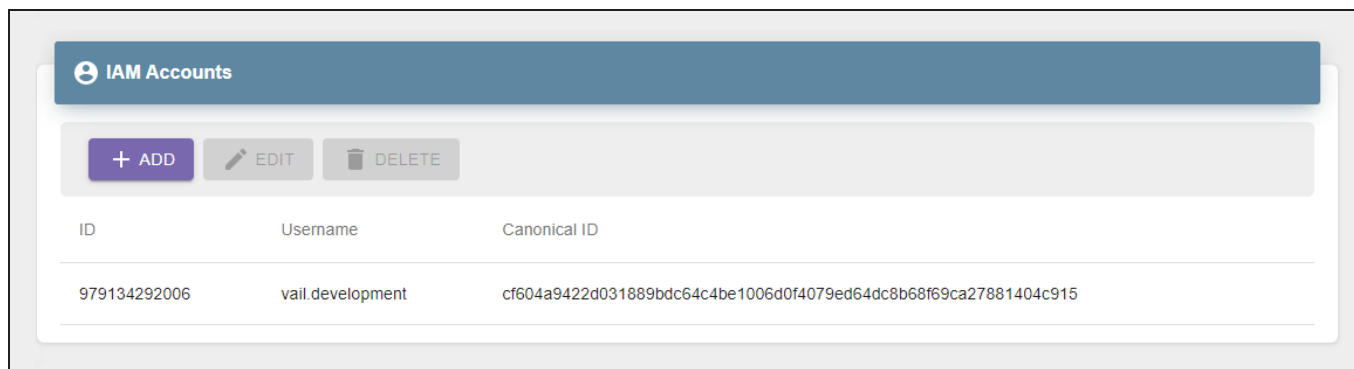


Figure 61 The IAM Accounts pane.

- 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 .
region	This parameter is not used when configuring an AWS user in the Spectra Vail application and must be left blank.
account	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. Note: You cannot use an AWS account ID alias when configuring an AWS user in the Spectra Vail application.
resource	The name of the specific resource.

The screenshot shows a modal window titled "Add IAM Account". It features a title bar with a help icon (?) and a close icon (X). The form contains the following fields:

- Role ARN**: A text input field.
- External ID**: A text input field with a note below it: "Optional. Specifies who can assume the role."
- Email**: A text input field.
- Description**: A text input field.

A blue **SUBMIT** button is positioned at the bottom right of the form area.

Figure 62 The Add IAM Account screen.

4. If desired, enter an **External ID**. The external ID is associated with the IAM role entered in [Step 3](#) and is configured when a role is created in an AWS account. The External ID is required to assume the role created in [Step 3](#). In the AWS management interface, the External ID can be found on the Roles section of the IAM screen, in the **Trusted relationships** tab.
5. 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**.
6. If desired, enter a **Description** for the IAM account.
7. Click **Submit**.

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)**.

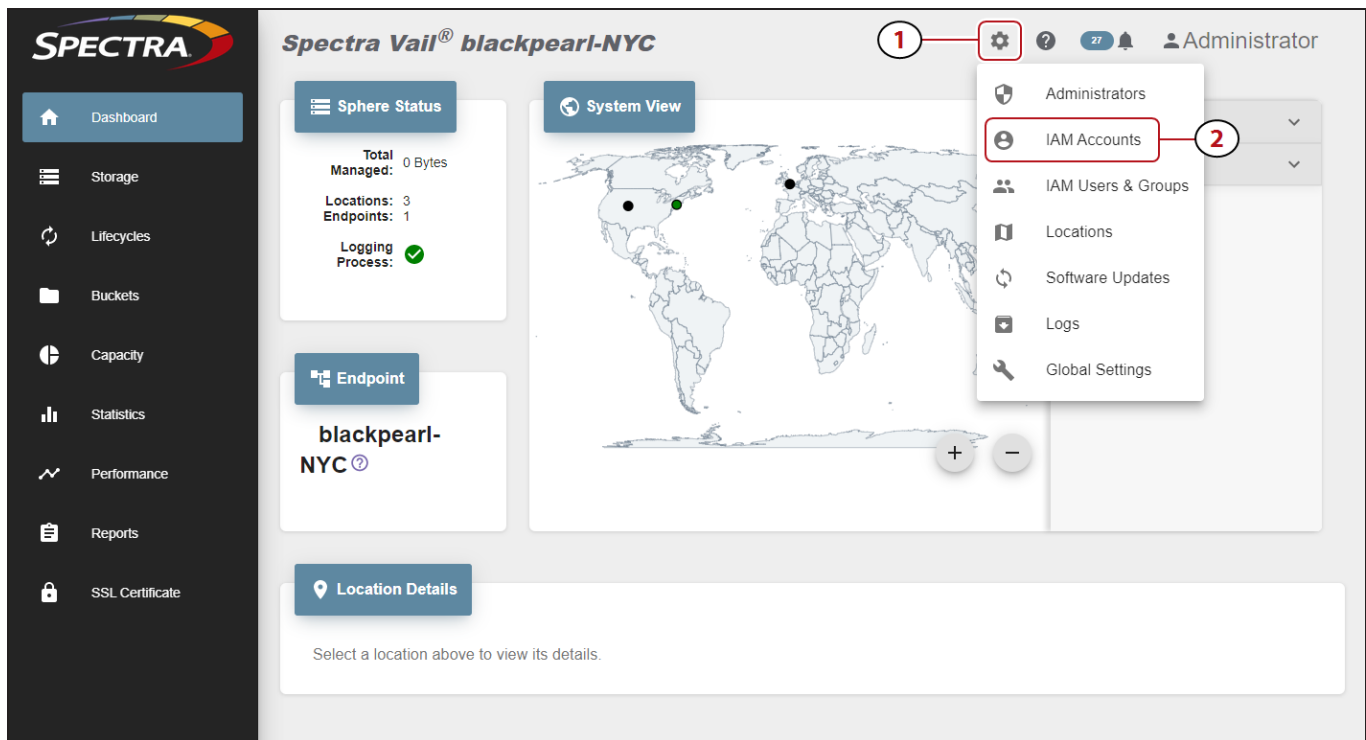


Figure 63 The Dashboard screen - Navigation menu.

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

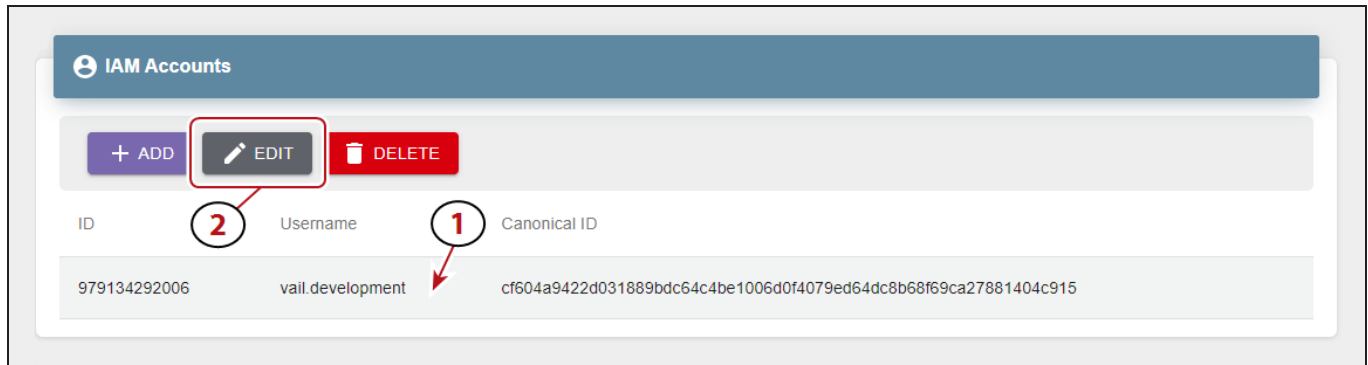


Figure 64 The IAM Accounts pane.

3. Change the **Email** address and **Description** 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 65 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 66 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 67 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 68 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 69 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.

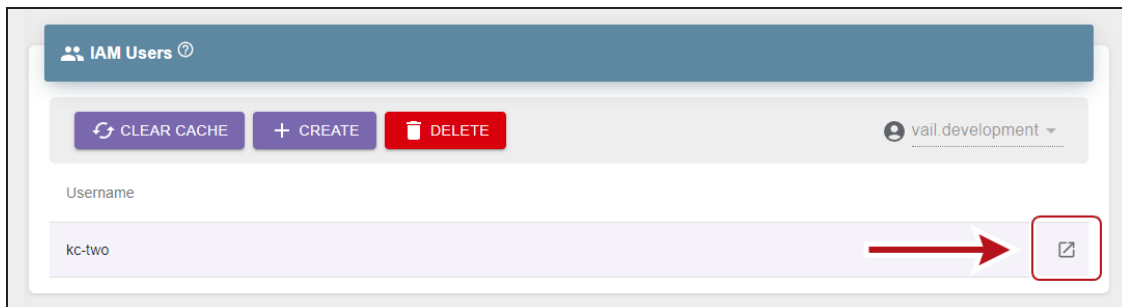


Figure 70 The IAM Users banner - View Details button.

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 71 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.

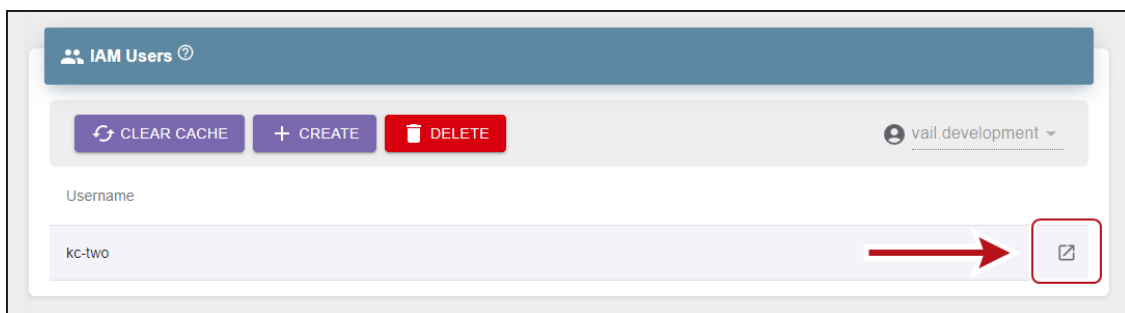


Figure 72 The IAM Users banner - View Details button.

3. Click **IAM Groups**.
4. **Select** the row of the group , then click **Add**.
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 73 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 74 The IAM Users banner - View Details button.

3. Click **IAM Groups**.
4. **Select** the row of the group , then click **Remove**.
5. Click **Remove** to confirm removing the user from the IAM group

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 75 The Dashboard screen - Navigation menu.

2. Under the **IAM Users** banner, (1) select the row of the user to delete, and (2) 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 76 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 77 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.

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 78 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 79 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 80 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.

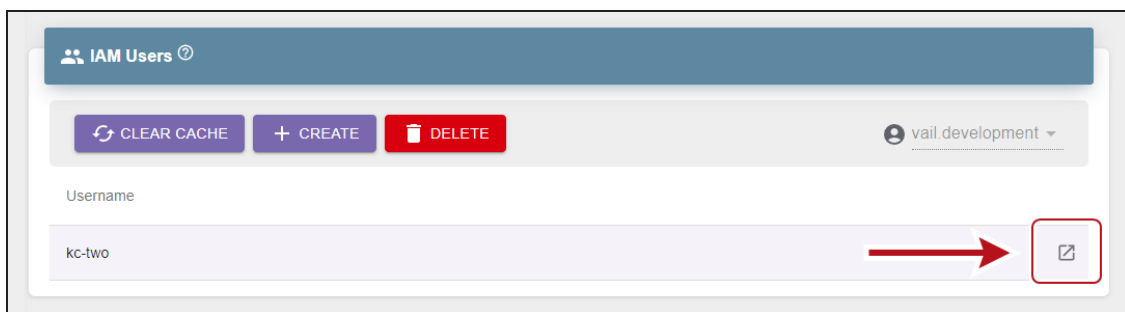


Figure 81 The IAM Users banner - View Details button.

3. Select **Access Keys**.
4. 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**.

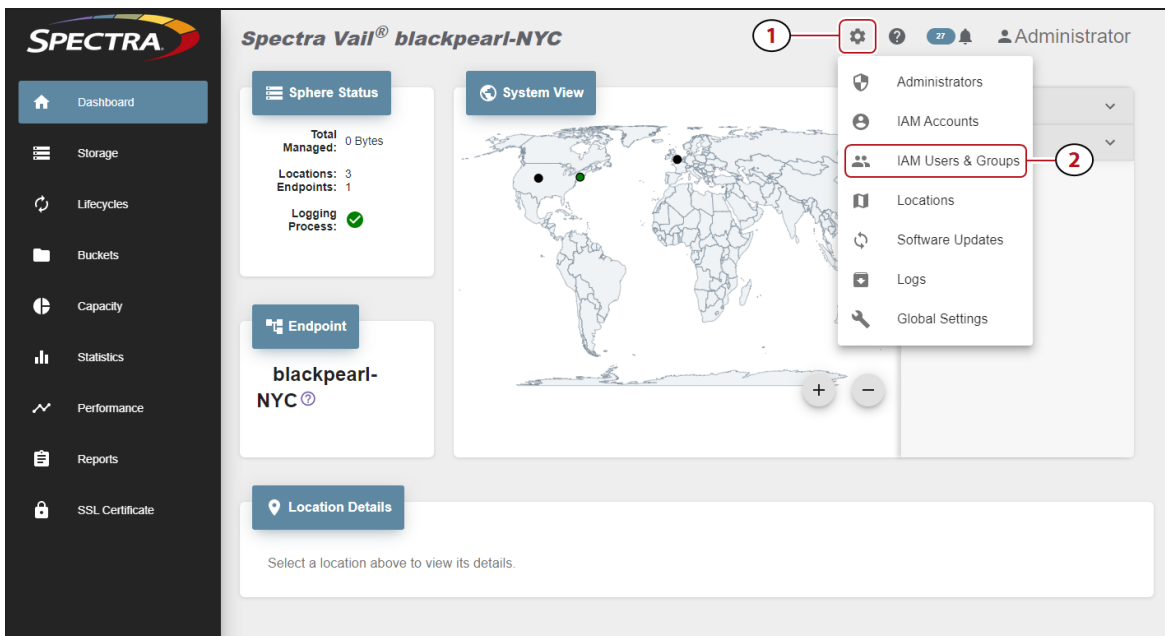


Figure 82 The Dashboard screen - Navigation menu.

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 83 The IAM Users banner - View Details button.

3. Select **Access Keys**.

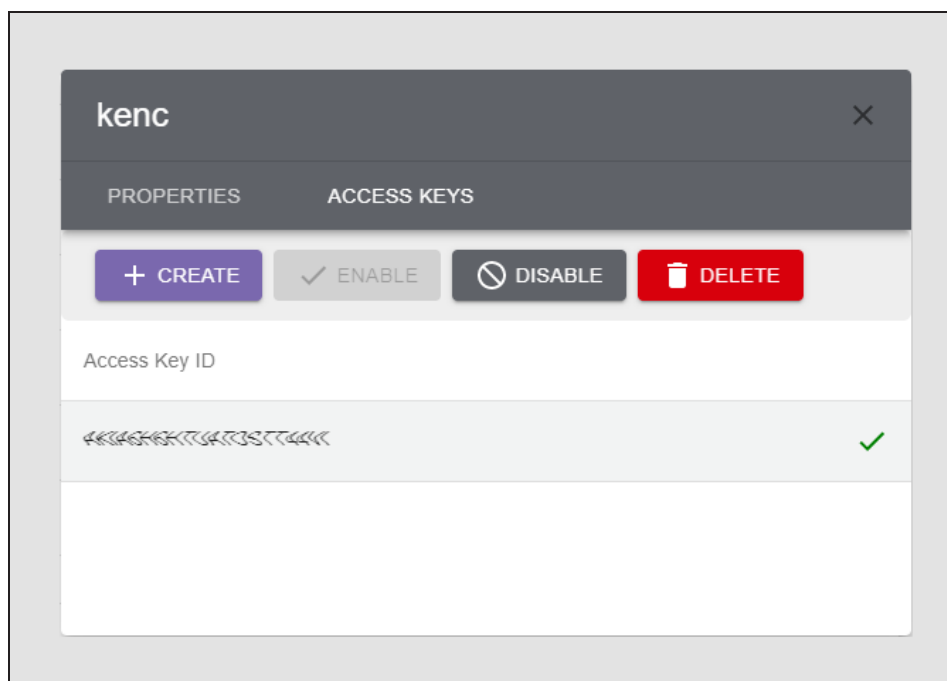


Figure 84 The User Properties - Access Keys screen.

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 85 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.

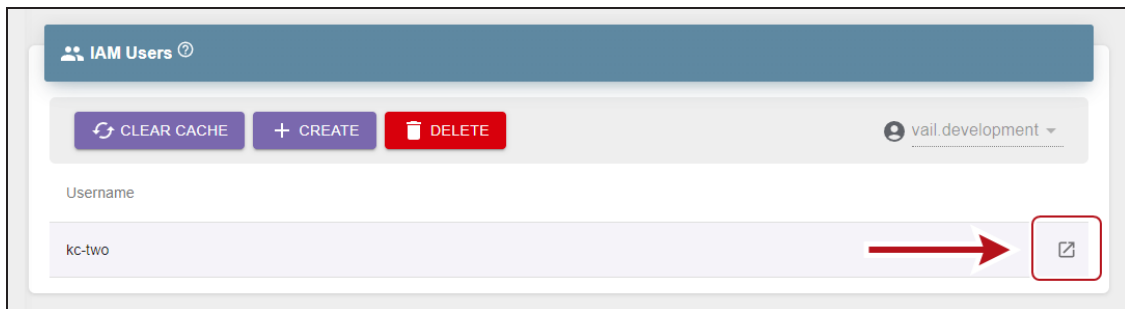


Figure 86 The IAM Users banner - View Details button.

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 5 - CREATE AND CONFIGURE A VAIL VM NODE

This chapter describes the creation and configuration steps for a Vail VM Node.

Create Vail VM Node Storage	103
Vail VM Node Host Requirements	103
Create a Node Using VMWare vSphere	104
Create a Node Using Oracle VirtualBox	111
Configure the Vail VM Node Network Settings	119
Configure Network Settings	119
Configure the Vail VM Node Hostname	121
Configure the SSL Certificate	122
Register a Vail VM Node with a Vail Sphere	124

CREATE VAIL VM NODE STORAGE

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.

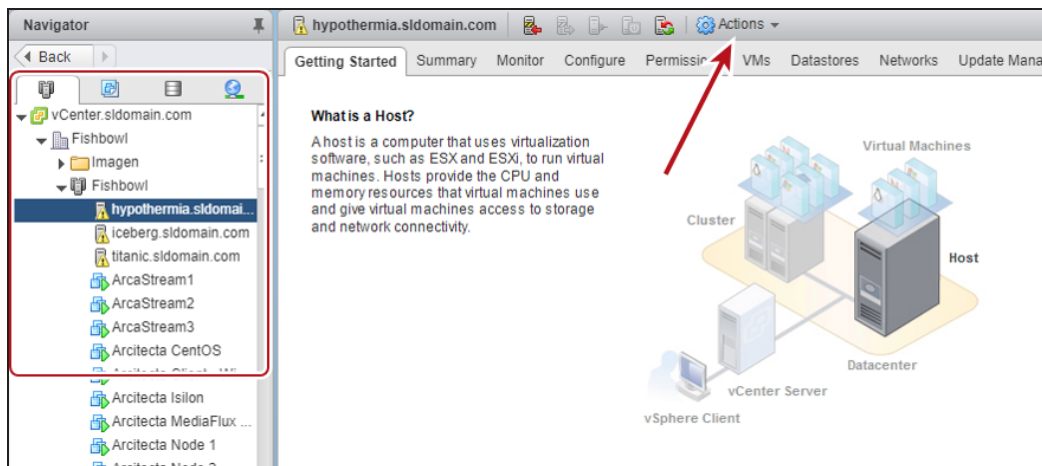


Figure 87 The VMWare vSphere home screen.

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**.

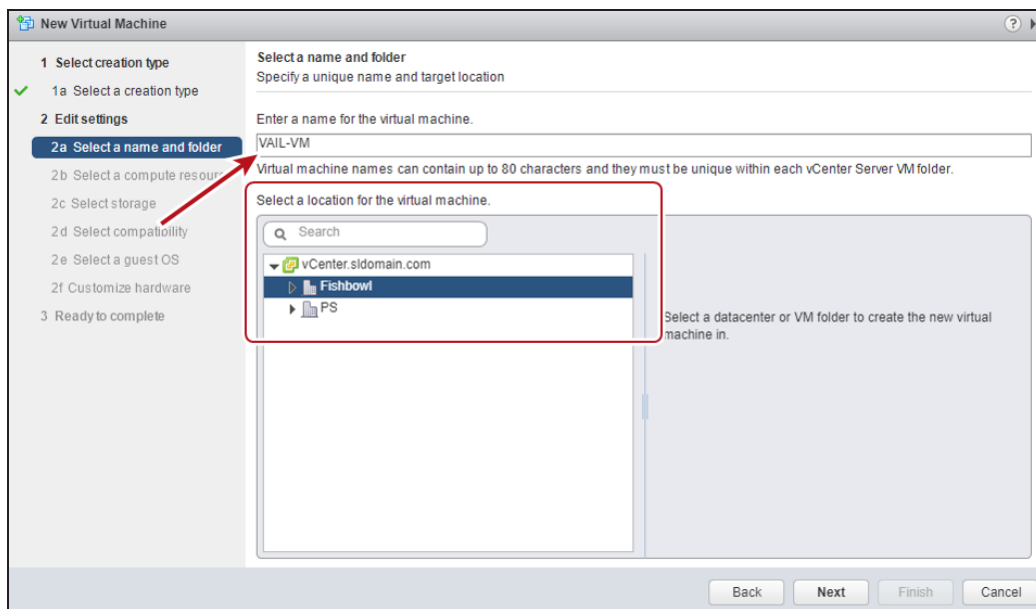


Figure 88 The New Virtual Machine - Select Name and Folder screen.

8. Using the **Select a compute resource** network browser, select an ESXi-based host in your network environment, and click **Next**.

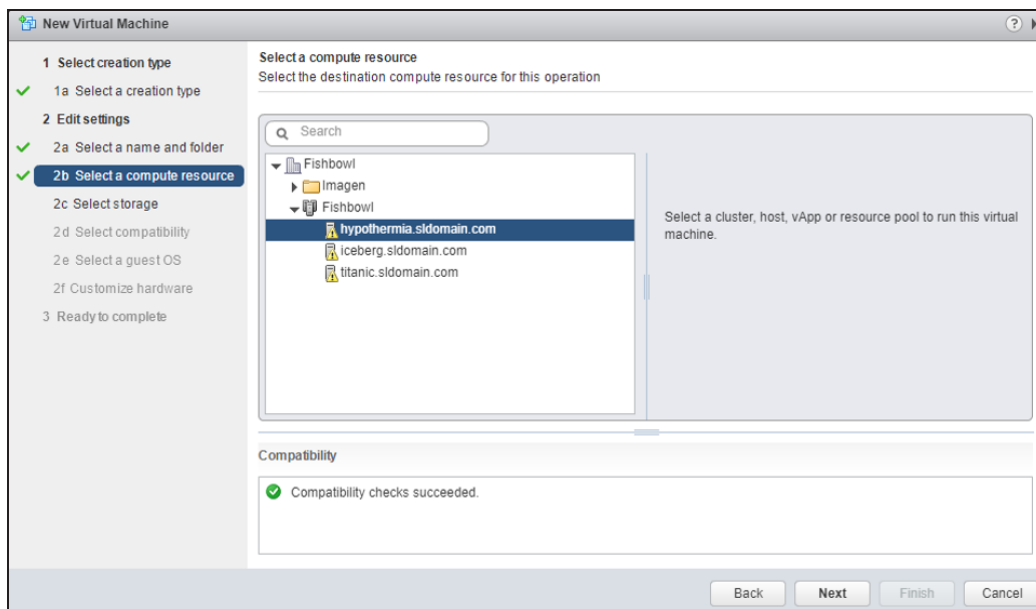


Figure 89 The New Virtual Machine - Select Compute Resource screen.

9. Using the **Select storage** table, select where to store the VM configuration files and virtual disks, and click **Next**.

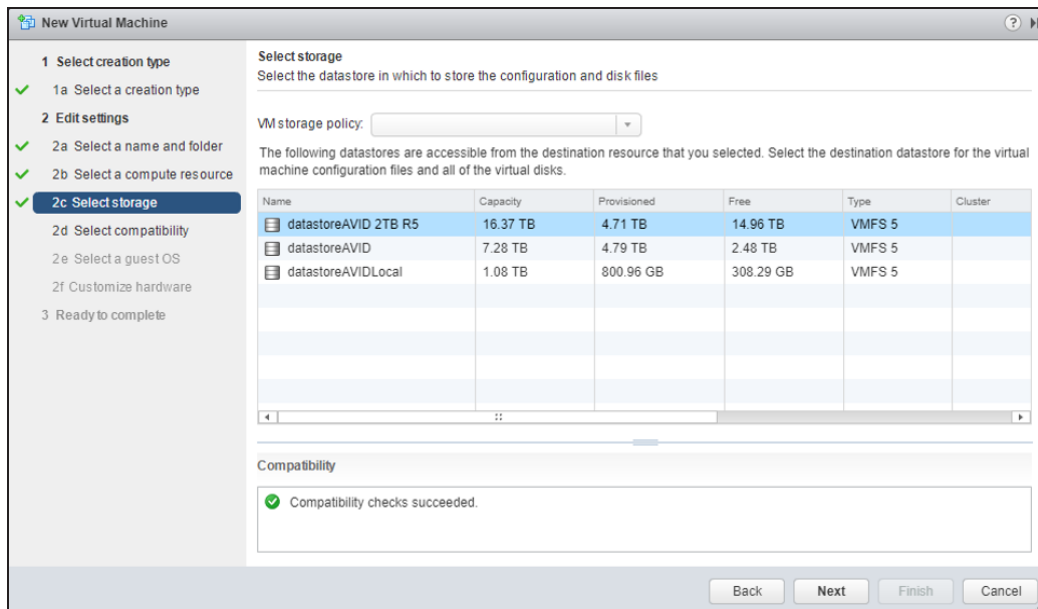


Figure 90 The New Virtual Machine - Select Storage screen.

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**

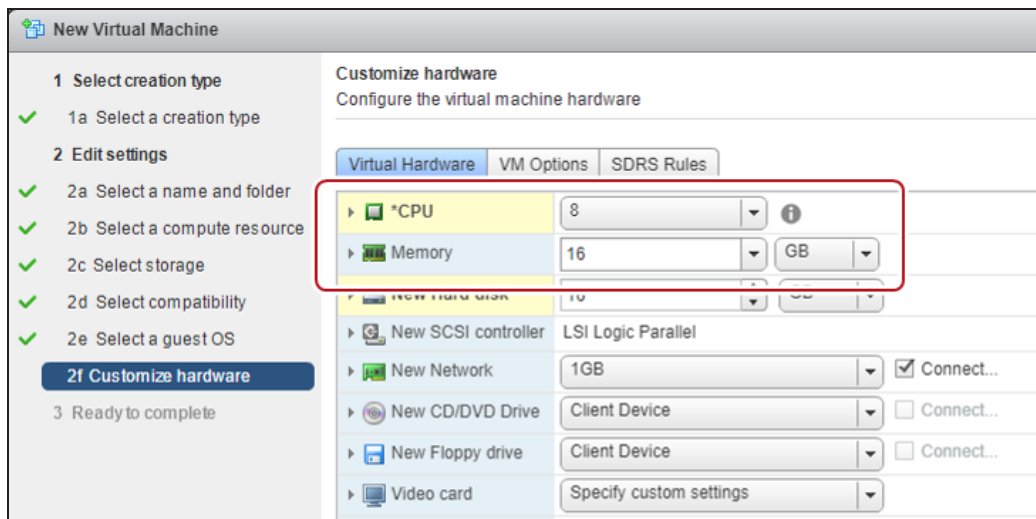


Figure 91 The New Virtual Machine - Customize Hardware screen.

14. On the right-hand side of the **New Hard disk** row, click the **X icon** to delete the default hard disk.




Figure 92 The New Virtual Machine - Customize Hardware screen.

15. Using the **New device** drop-down menu, select **Existing Hard Disk**, then click **Add**.

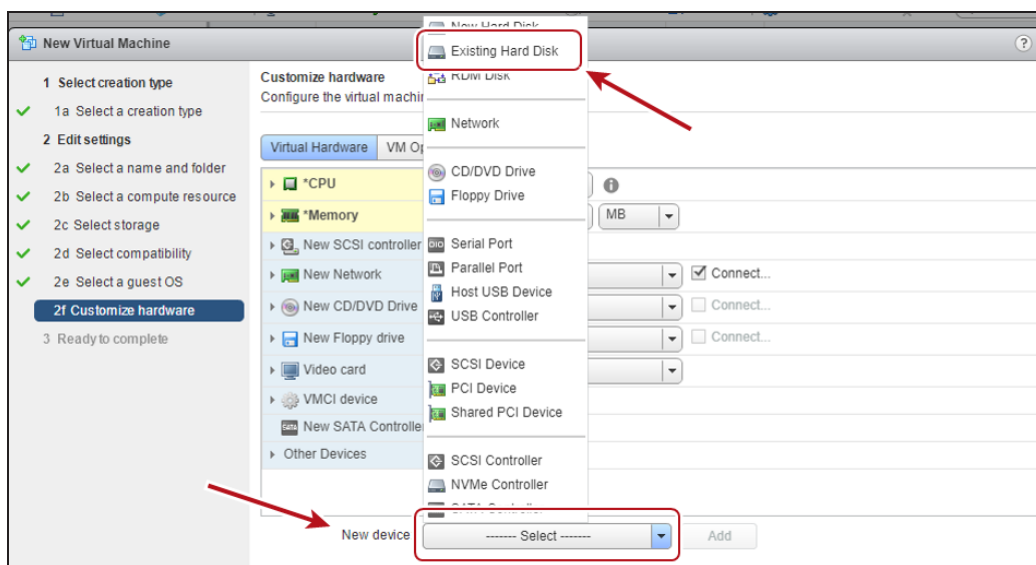


Figure 93 The New Virtual Machine - Customize Hardware screen.

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**.

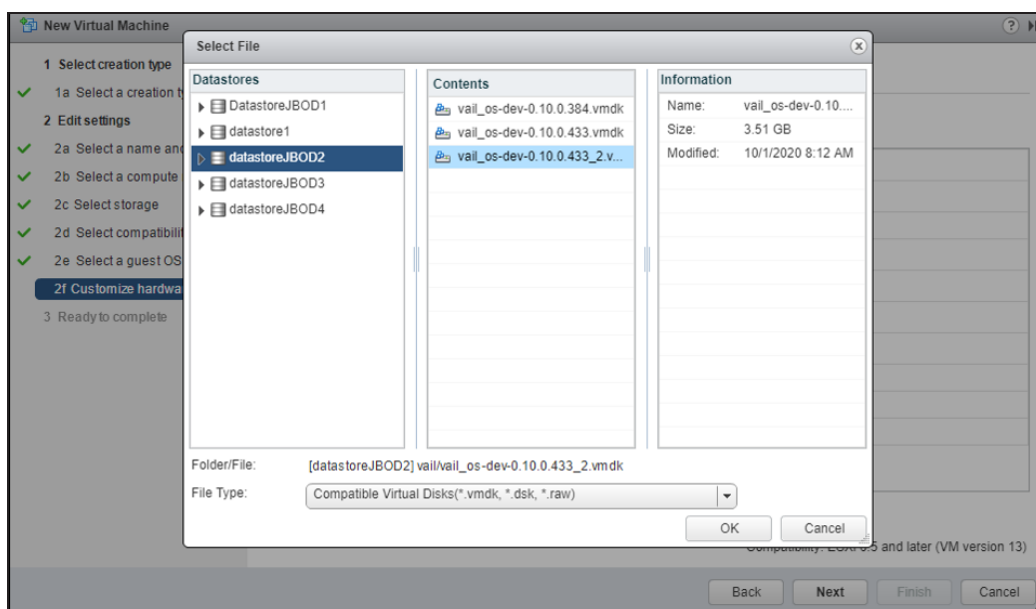


Figure 94 The New Virtual Machine - Customize Hardware - Select File screen.

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.

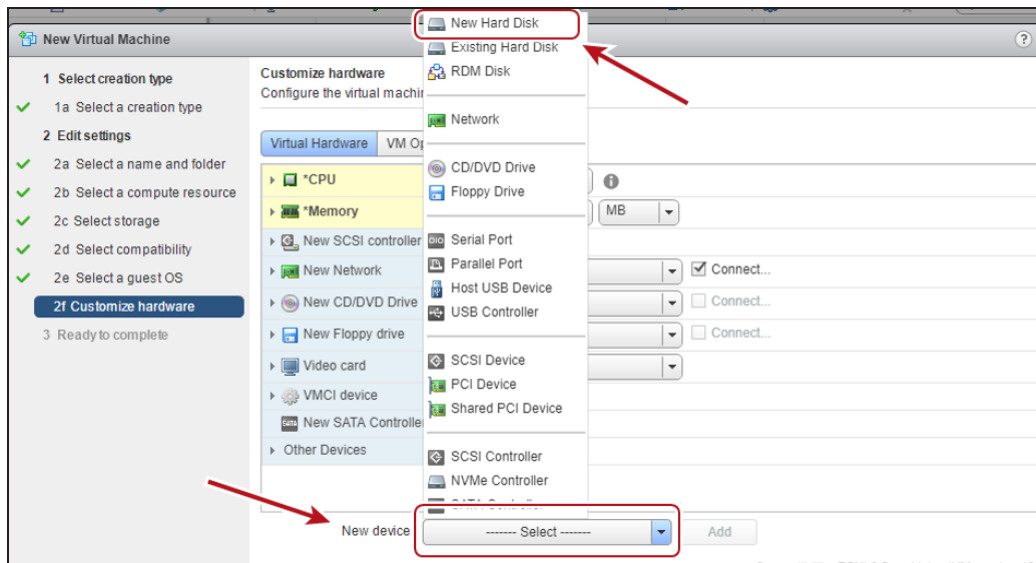


Figure 95 The New Virtual Machine - Customize Hardware screen.

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.

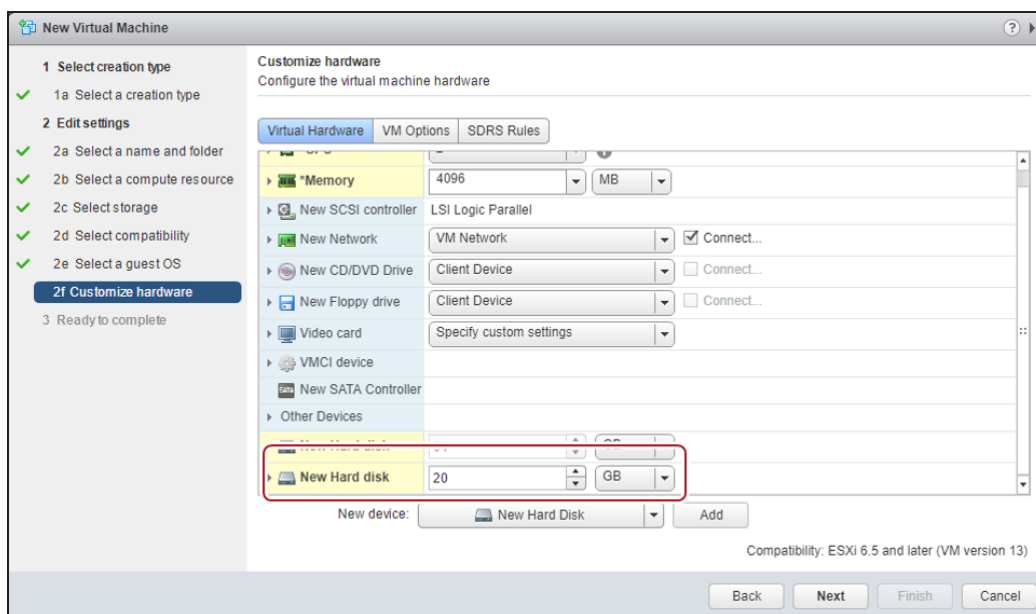


Figure 96 The New Virtual Machine - Customize Hardware screen.

19. Using the **New Network** drop-down menu, select **VM Network** and click **Next**.

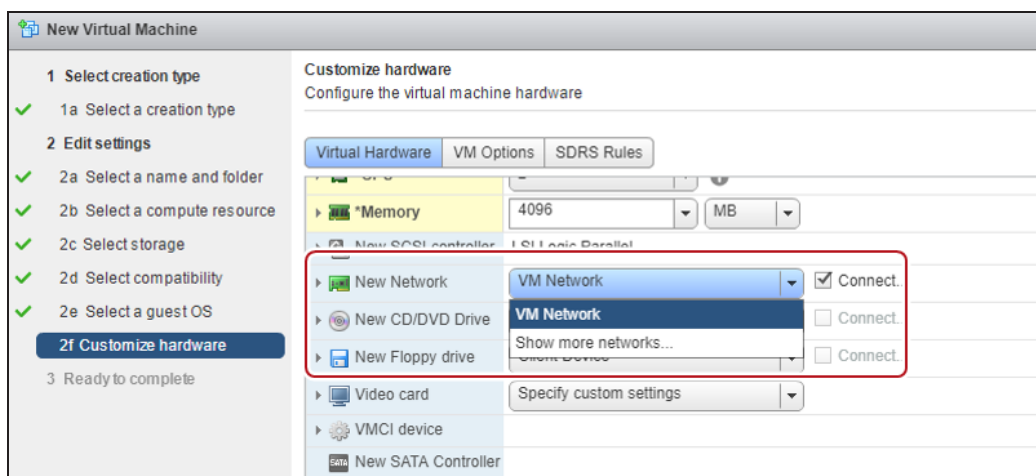


Figure 97 The New Virtual Machine - Customize Hardware screen.

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.

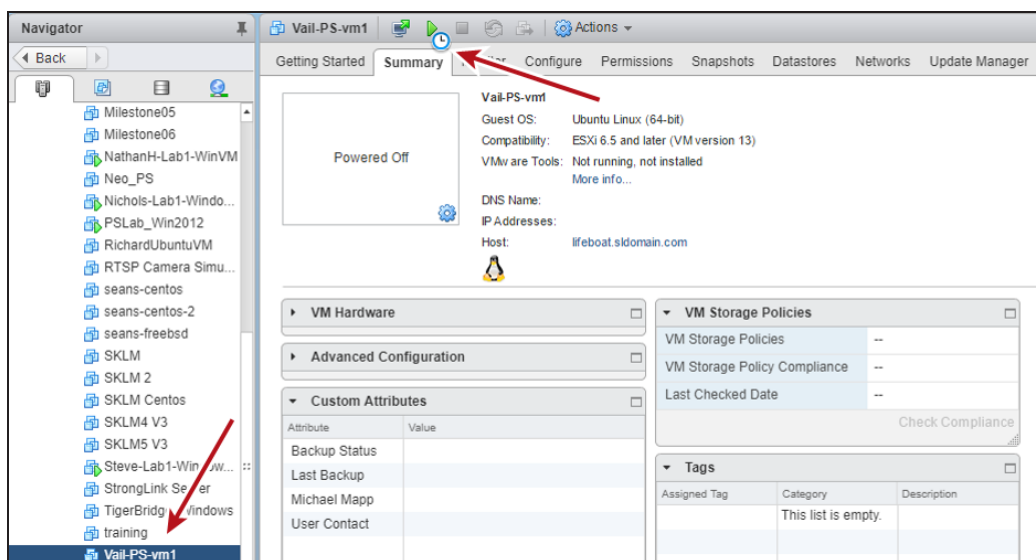
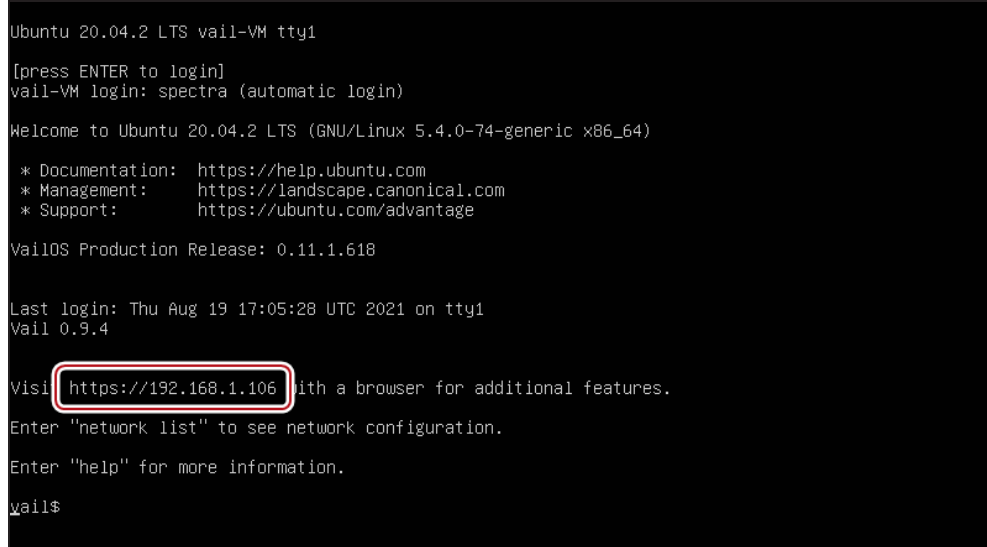


Figure 98 The New Virtual Machine - Summary screen.

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 99 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 100 Oracle VM VirtualBox Manager.


5. Enter the desired **Name**.

Figure 101 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**.

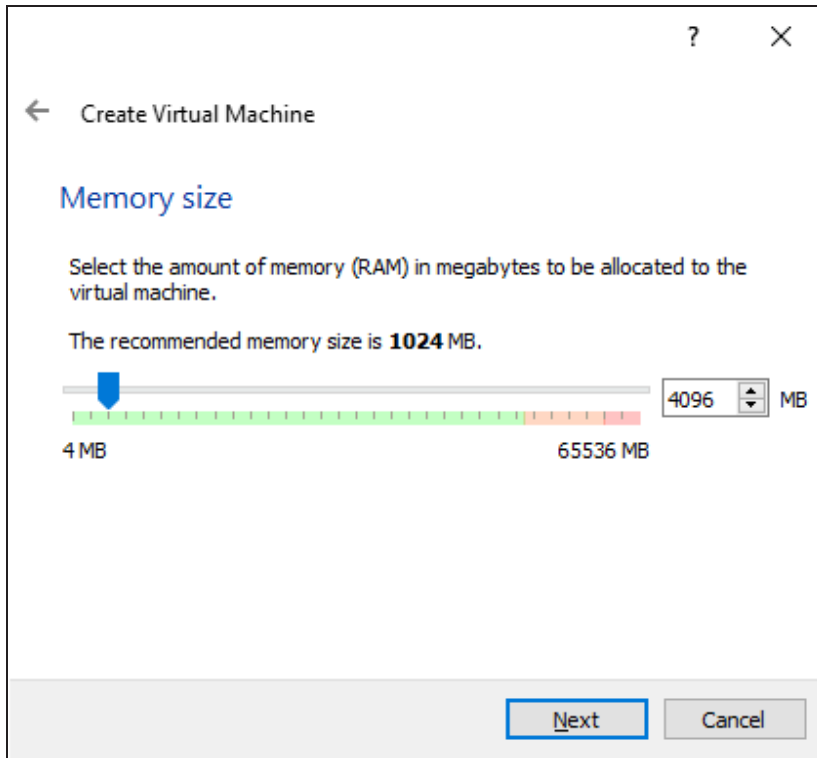


Figure 102 The Create Virtual Machine - Memory Size screen.

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 103 The Hard Disk Selector screen.

15. On the Create Virtual Machine - Hard disk screen, click **Create**.

16. After the VM is created, click **Settings**.

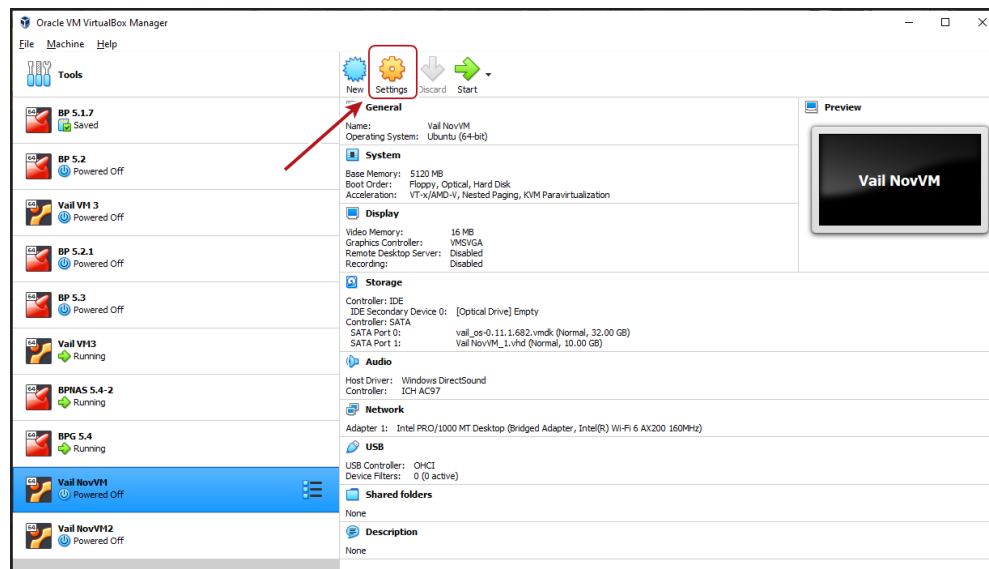


Figure 104 Oracle VM VirtualBox Manager.

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.

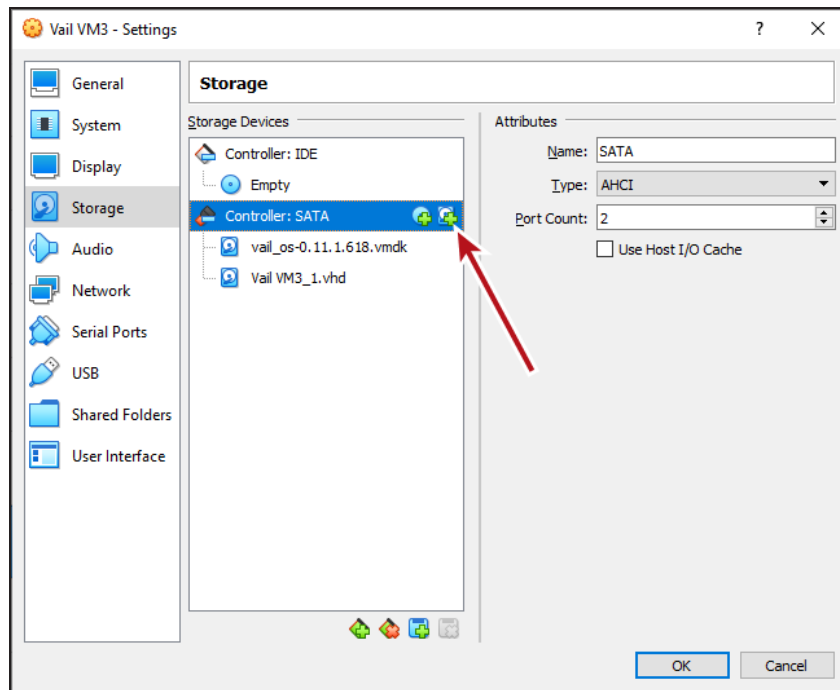


Figure 105 The VM Settings - Storage screen.

19. Select **Create**.




Figure 106 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.

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 107 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**.

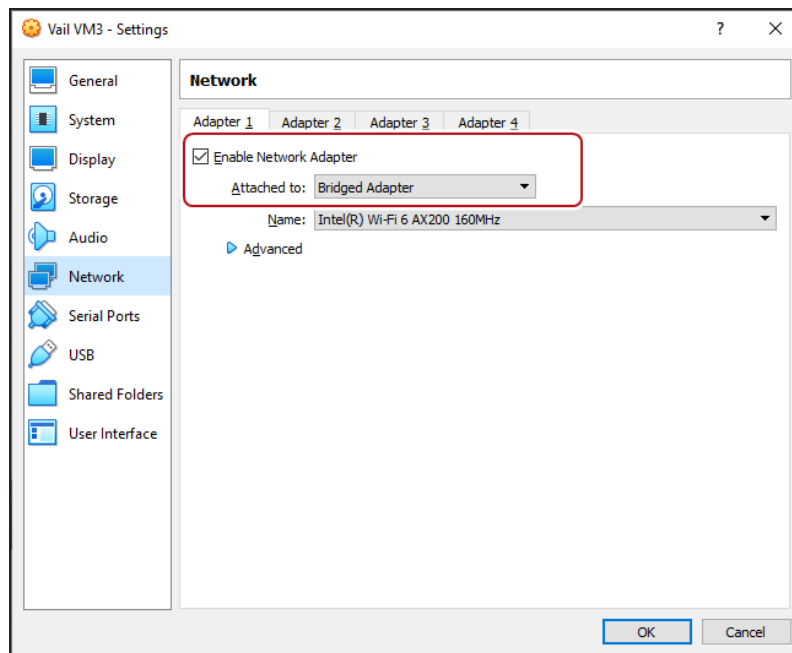


Figure 108 The VM Settings - Network screen.

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 109 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 110 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 below**
- **Configure the Vail VM Node Hostname on page 121**
- **Configure the SSL Certificate on page 122**

Configure Network Settings

Here is how to configure the Vail VM node IP address:

1. In the upper right corner of the Vail management console, click the **gear icon** and select **Network**.

2. 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 111 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**.

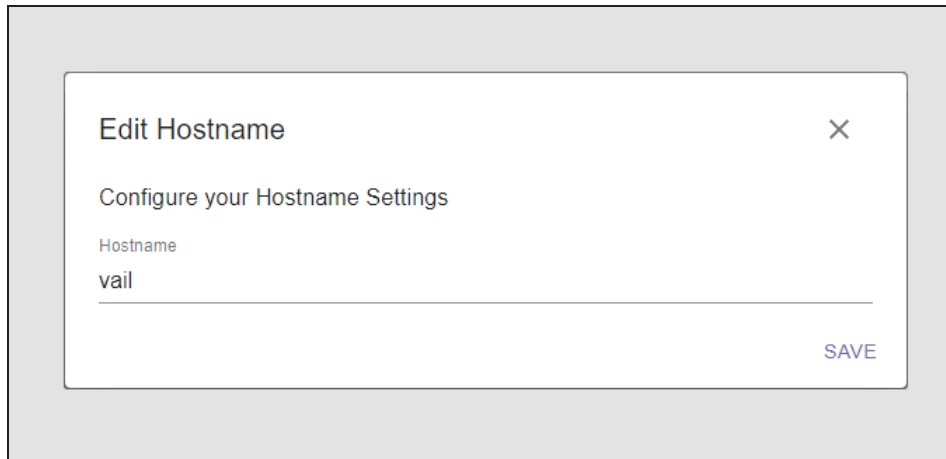


Figure 112 The Vail VM Node Edit Hostname screen.

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 S3 solution 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**.

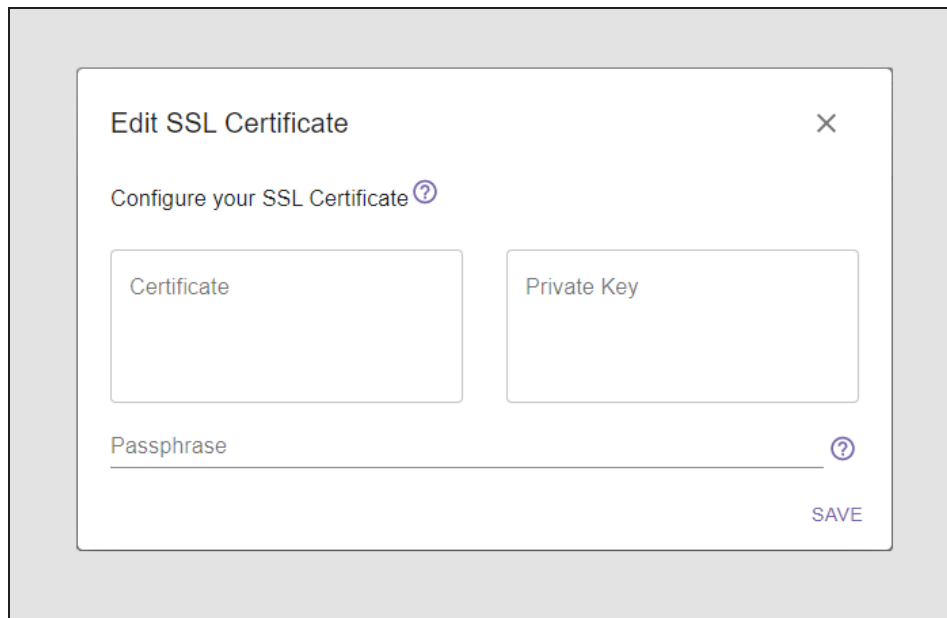
The image shows a modal dialog box titled "Edit SSL Certificate" with a close button (X) in the top right corner. Below the title is a subtitle "Configure your SSL Certificate" followed by a help icon (?). The dialog contains three input fields: "Certificate" and "Private Key" are side-by-side, and "Passphrase" is below them. A help icon (?) is next to the "Passphrase" field. A "SAVE" button is located at the bottom right of the dialog box.

Figure 113 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**.

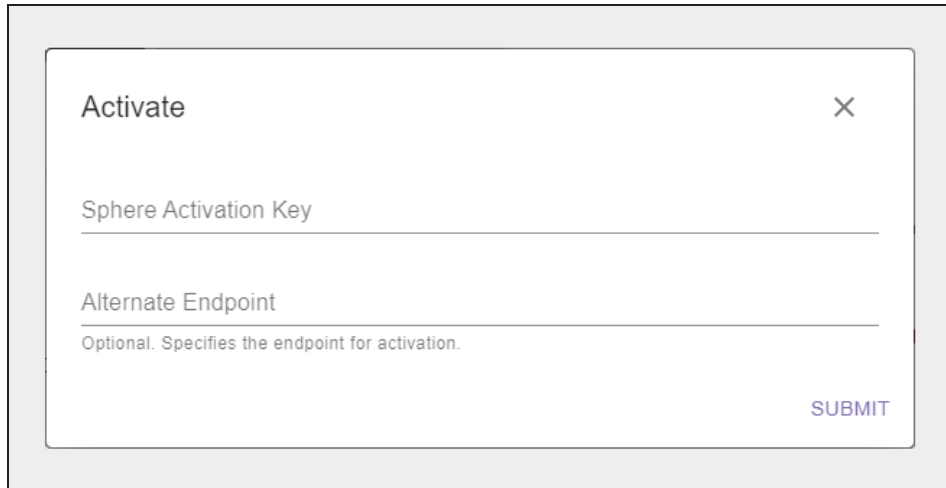


Figure 114 The Vail VM Node - Activate screen.

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**.

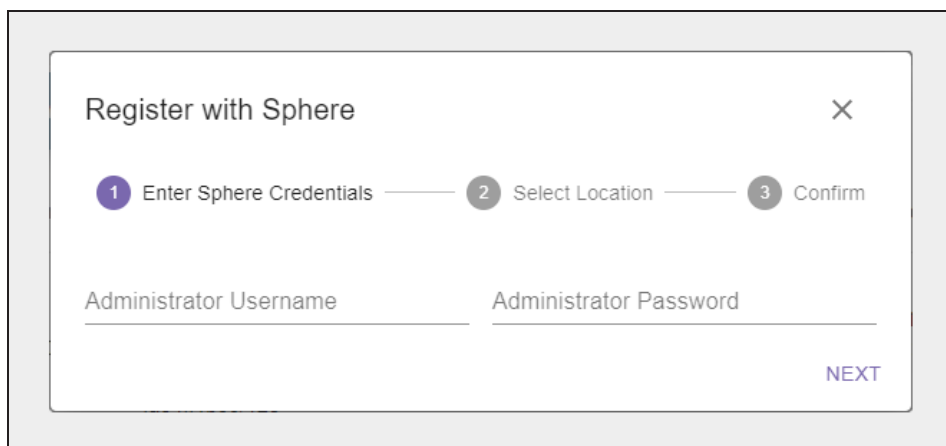


Figure 115 The Register With Sphere - Credentials screen.

7. Enter the Spectra Vail application **Administrator Username** and **Password**.

8. Click **Next**.

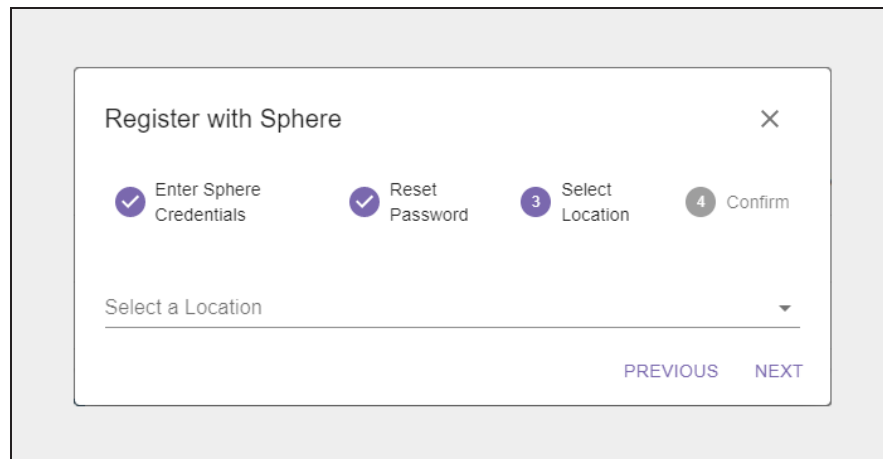


Figure 116 The Register with Sphere - Select Location screen.

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 129**

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 117 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.

Note: After the Vail VM node is registered with the Vail sphere, it is automatically added as endpoint storage and does not need to be configured in the Vail management console.

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 118 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.

Note: After the Vail VM node is registered with the Vail sphere, it is automatically added as endpoint storage and does not need to be configured in the Vail management console.

CHAPTER 6 - USING THE VAIL APPLICATION

This chapter describes using the Spectra Logic Spectra Vail application.

View Capacity Information	132
View Performance Metrics	135
View Statistics	138
View Vail Bucket Details	141
View Vail Bucket Contents	145
View Object Details	147
Edit Global Settings	150
Change Lifecycle Rule Nightly Processing Time	150
Enable Diagnostic Monitor	150
Using Proxy Connections	152
Configure Proxy Connection	152
Edit Proxy Server	152
Delete Proxy Server	153
Edit a Vail Bucket	154
Delete a Vail Bucket	158
Create an Object Clone	159
Delete an Object Clone	161
View Storage Details	163
Edit BlackPearl or Vail VM Endpoint	165
Change Endpoint Location	165
Add Additional Host Names	166
Configure Debug Logging	167
Edit Storage	169
Edit BlackPearl Bucket Storage	169
Edit BlackPearl NAS Storage	172
Edit Vail VM Node Storage	173
Edit Google Cloud Platform Storage	174

Edit AWS S3 Cloud Storage	176
Edit Other Third-Party Cloud Storage	179
Delete Storage	180
View Lifecycle Details	184
Edit a Lifecycle	187
Delete a Lifecycle	190
Create a Location	191
Delete a Location	195
Clear the IAM Cache	196
View Reports	197
View Vail Application Messages	200
Message Details	202
Vail Application Logs	203
Update the Vail Application Software	204
Enable Diagnostic Monitor	206
Log Out of the Vail Management Console	207

VIEW CAPACITY INFORMATION

The Capacity page allows you to see data capacity information for the 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.

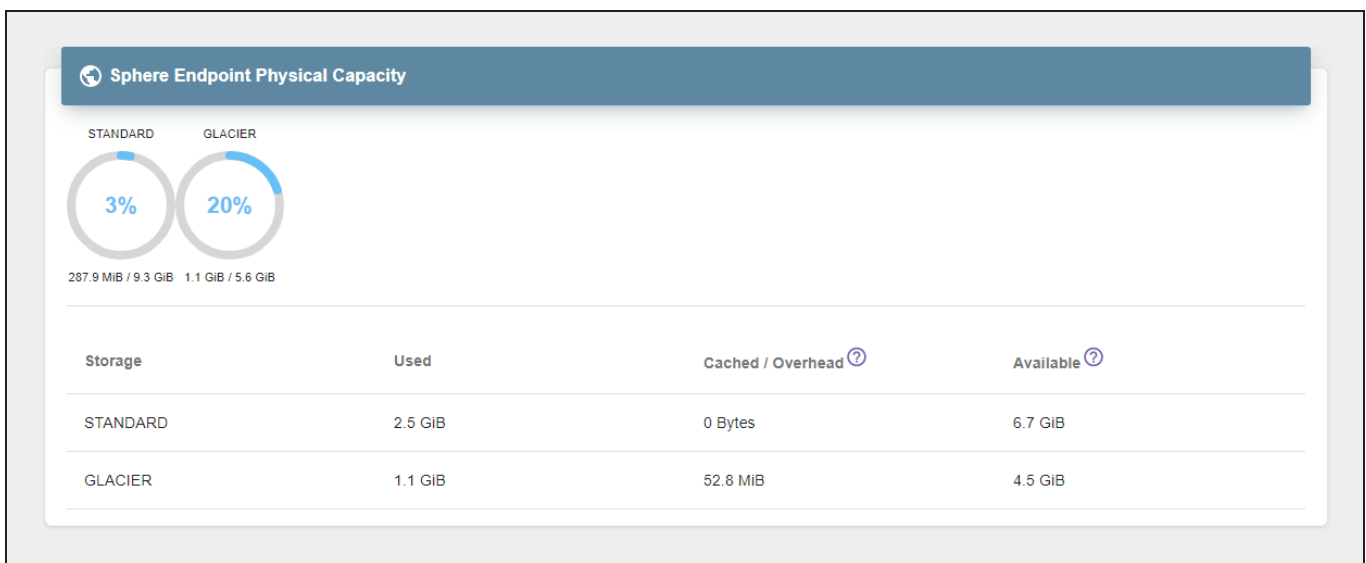


Figure 119 The Sphere Endpoint Physical Capacity pane.

Field	Description
Storage	The AWS storage class. Each type of storage class has its own row. See Storage Classes on page 209 for information about each storage class.
Used	The amount of space used for each storage type.
Cached/Overhead	The amount of space used by the Spectra Vail application for temporary storage and file system overhead.
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.

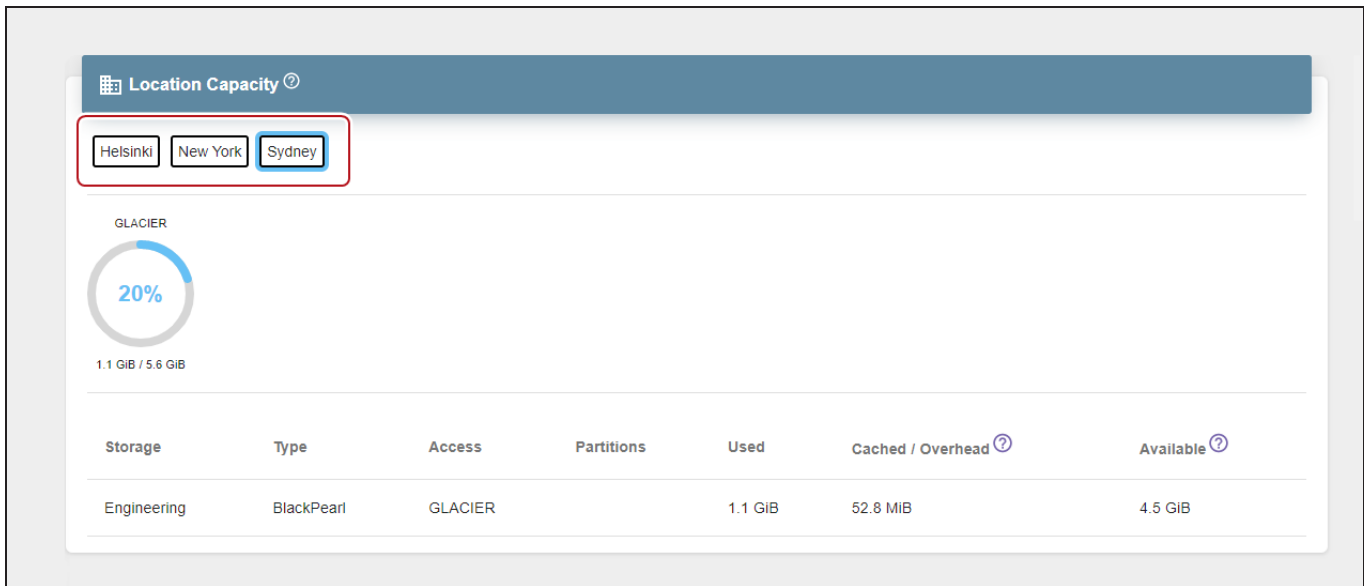


Figure 120 The Location Capacity pane.

Field	Description
Storage	The name of the location.
Type	The type of storage used for each location. BlackPearl - Storage on a BlackPearl system. Volume - Storage on a Vail VM node or Cloud endpoint.
Access	The access type for each location. These correspond to the storage class configured for each location.
Partitions	The BlackPearl data partition(s) that are used for storage. Note: This information does not display in Vail 2.0.0.
Used	The amount of space used for each location.
Cached/Overhead	The amount of space used by the Spectra Vail application for temporary storage and file system overhead.
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 121 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 show the last 24 hours of information in five minute intervals.

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

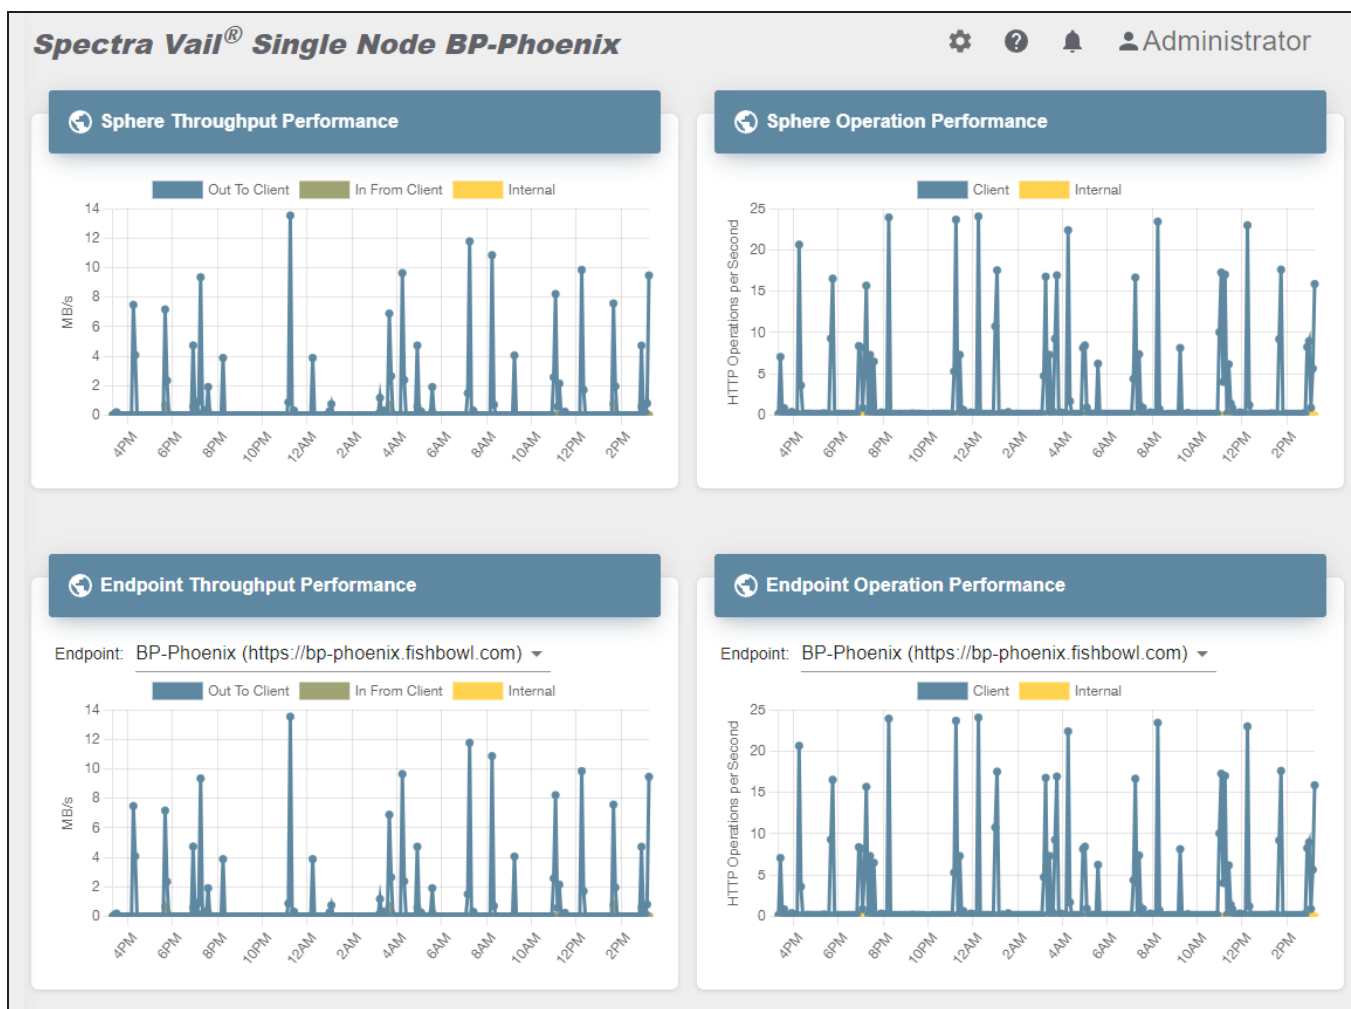


Figure 122 The Performance screen.

- The **Throughput Performance** graphs display the following information.

Option	Description
Out To Client	The data transfer rate of objects being read via an S3 client or other reads from outside the Vail sphere.

Option	Description
In From Client	The data transfer rate of objects being written via an S3 client or other writes from outside the Vail sphere.
Internal	The data transfer rate of objects being read or written inside the Vail sphere network, such as data transfers via lifecycles.

- The **Operation Performance** graphs display the following information.

Option	Description
Client	The number of GET/PUT/POST commands from an S3 client or from outside the Vail sphere.
Internal	The number of GET/PUT/POST commands from inside the Vail sphere, for example from lifecycles.

- The Endpoint Throughput Performance and Endpoint Operation Performance graphs display same information as the Throughput Performance and Operation Performance graphs for a single endpoint. Use the **Endpoint** drop-down menu to switch between endpoints.

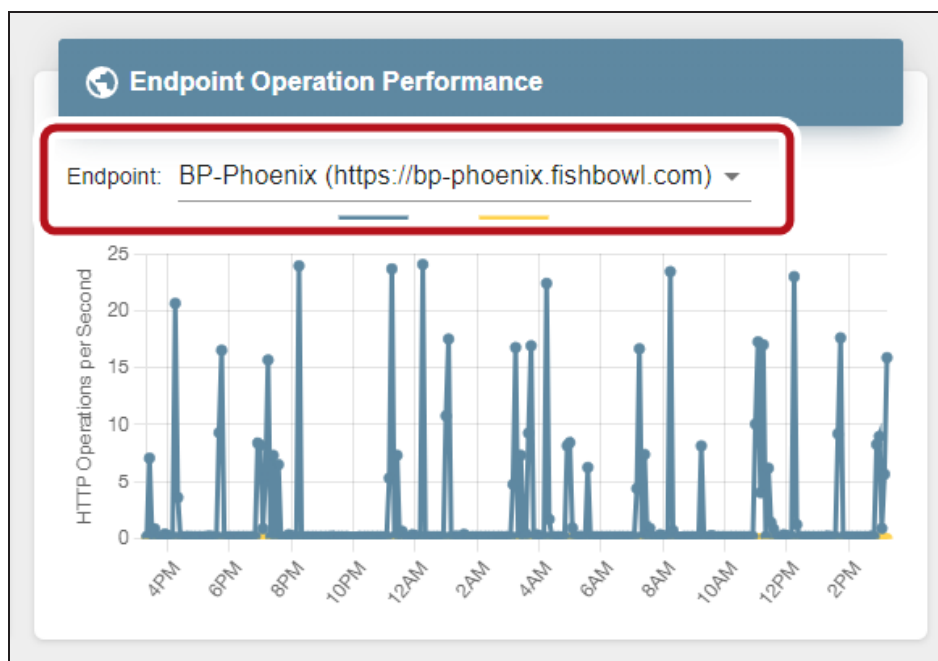


Figure 123 The Endpoint Throughput Performance pane.

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

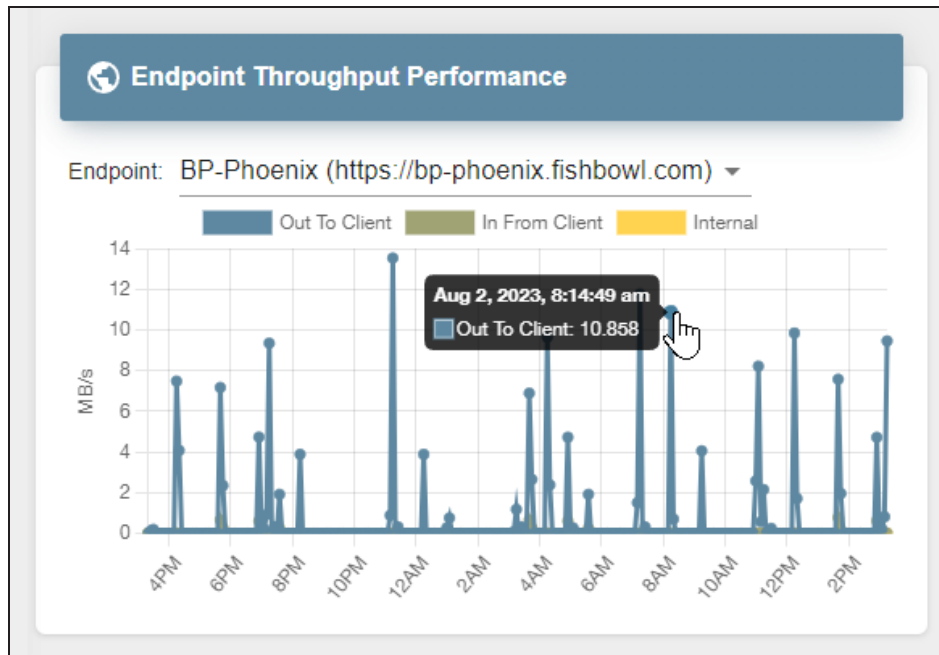


Figure 124 The Endpoint Throughput Performance pane - mouseover.

VIEW STATISTICS

The Statistics page displays the data storage growth rate over time for both online and archive class storage.

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

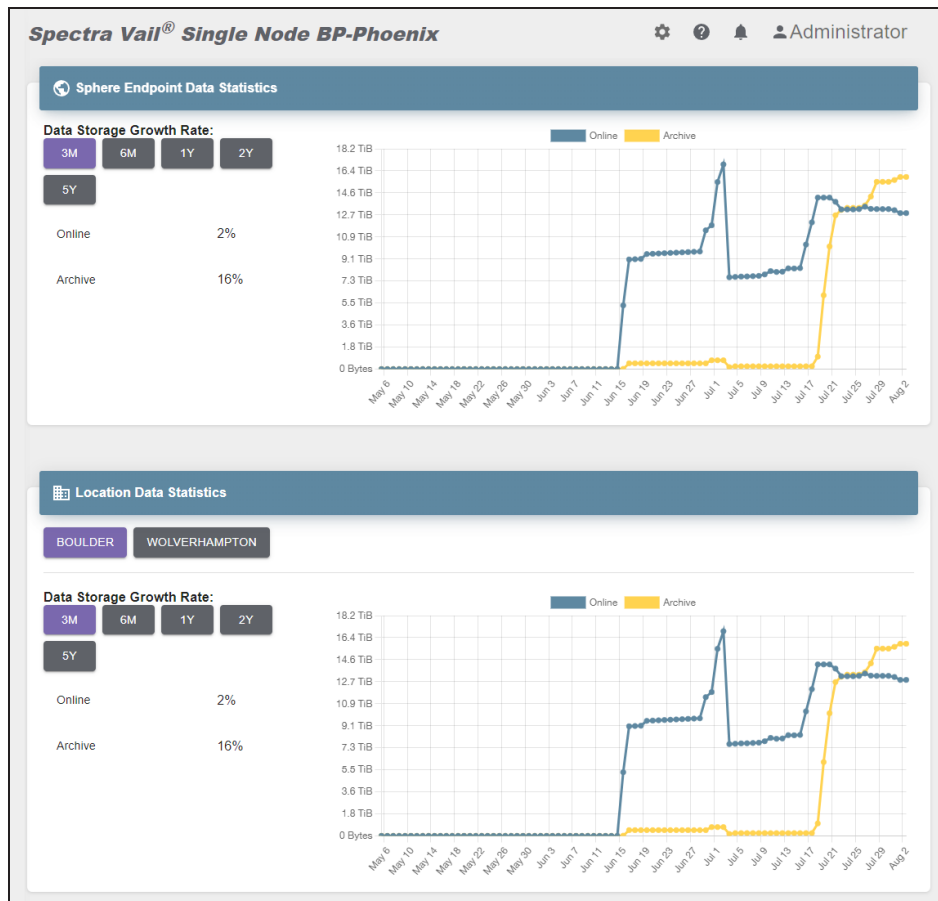


Figure 125 The Statistics screen.

The storage growth rate graphs display the amount of online and archive storage used on both endpoint and cloud storage, as well as storage per configured location.

- Click a **date range** button to view statistics for the selected time period.




Figure 126 The Sphere Endpoint Data Statistics pane.

- Click a **location** button to see data for a configured location.

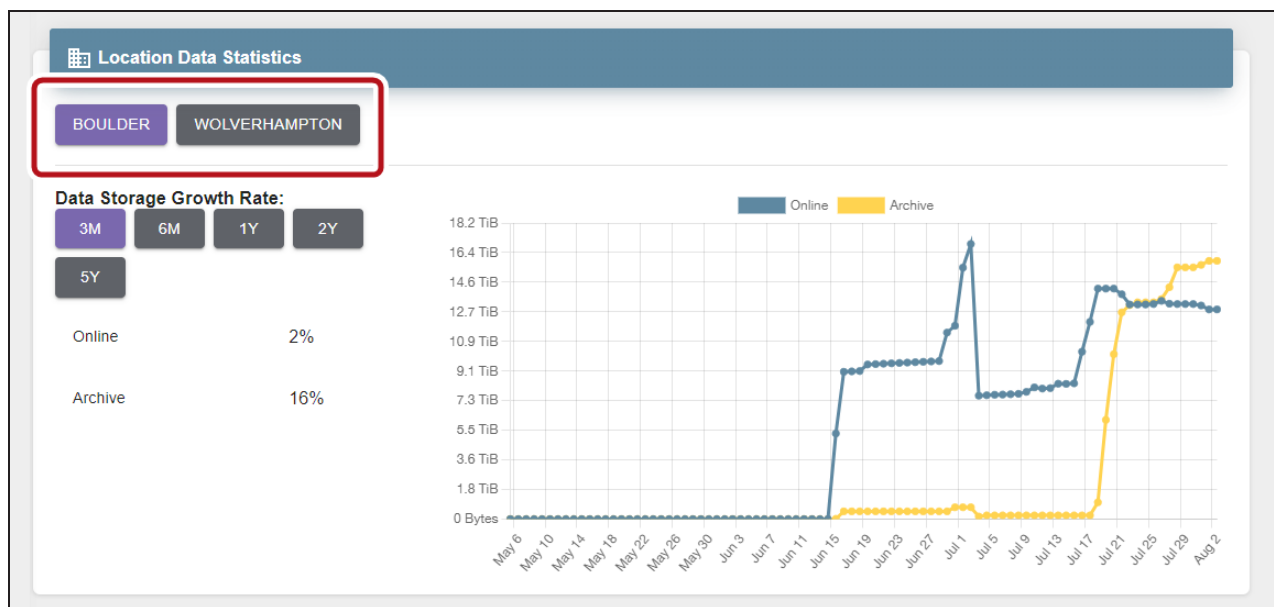


Figure 127 The Location Data Statistics pane.

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

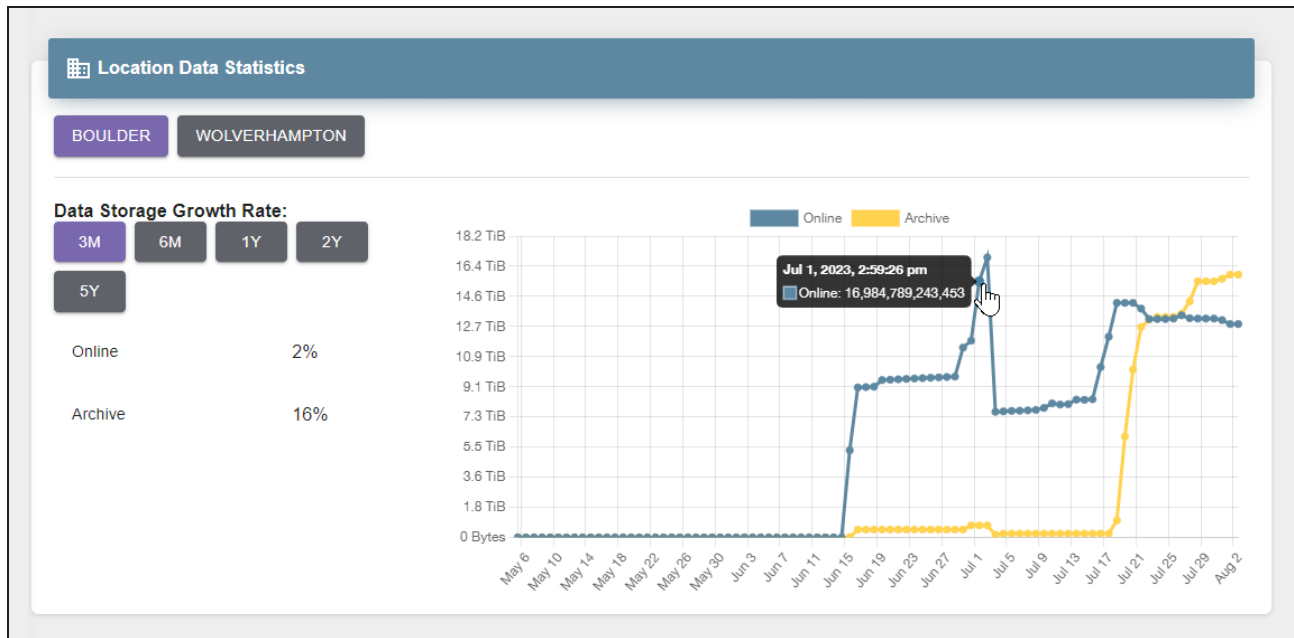


Figure 128 The Location Data Statistics pane - mouseover.

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 145.

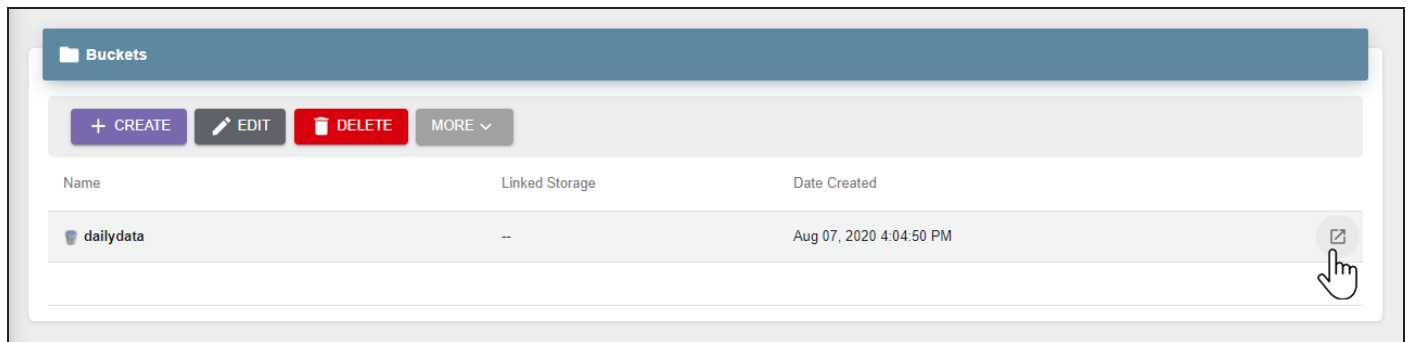
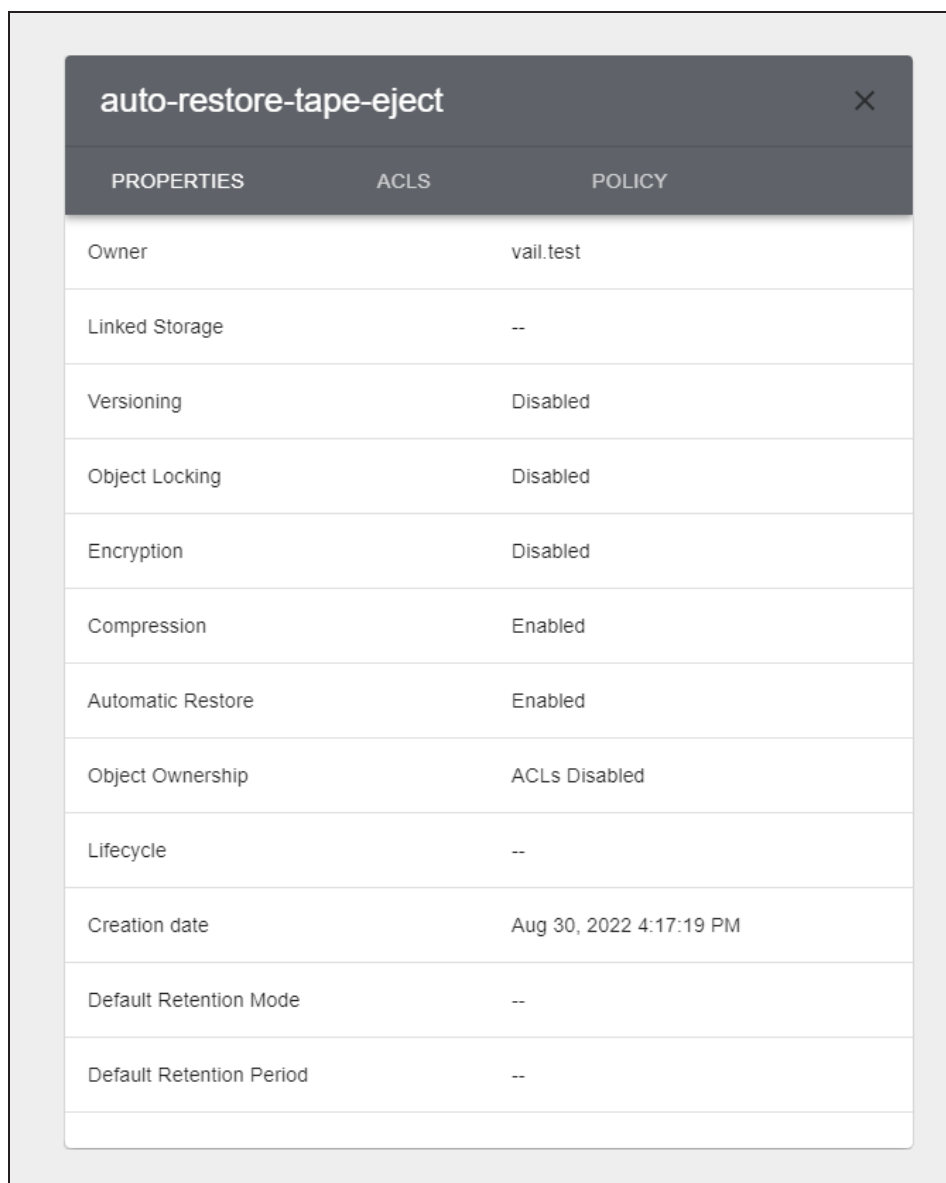


Figure 129 The Buckets pane.

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



auto-restore-tape-eject	
PROPERTIES	POLICY
Owner	vail.test
Linked Storage	--
Versioning	Disabled
Object Locking	Disabled
Encryption	Disabled
Compression	Enabled
Automatic Restore	Enabled
Object Ownership	ACLs Disabled
Lifecycle	--
Creation date	Aug 30, 2022 4:17:19 PM
Default Retention Mode	--
Default Retention Period	--

Figure 130 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.
Automatic Restore	Indicates if automatic restore 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 **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 select **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**.

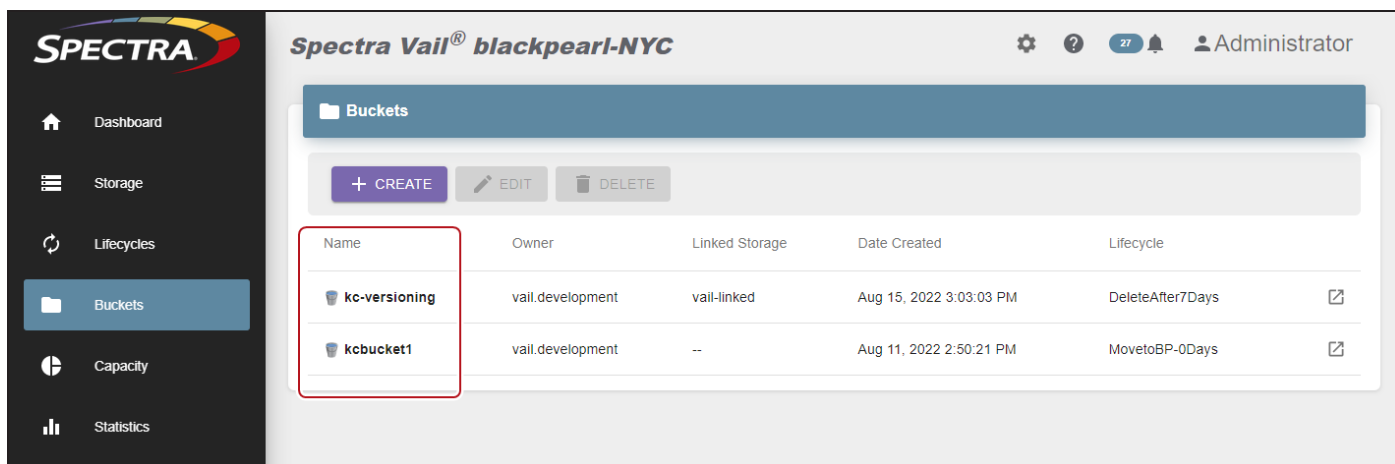


Figure 131 The Buckets screen.

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.

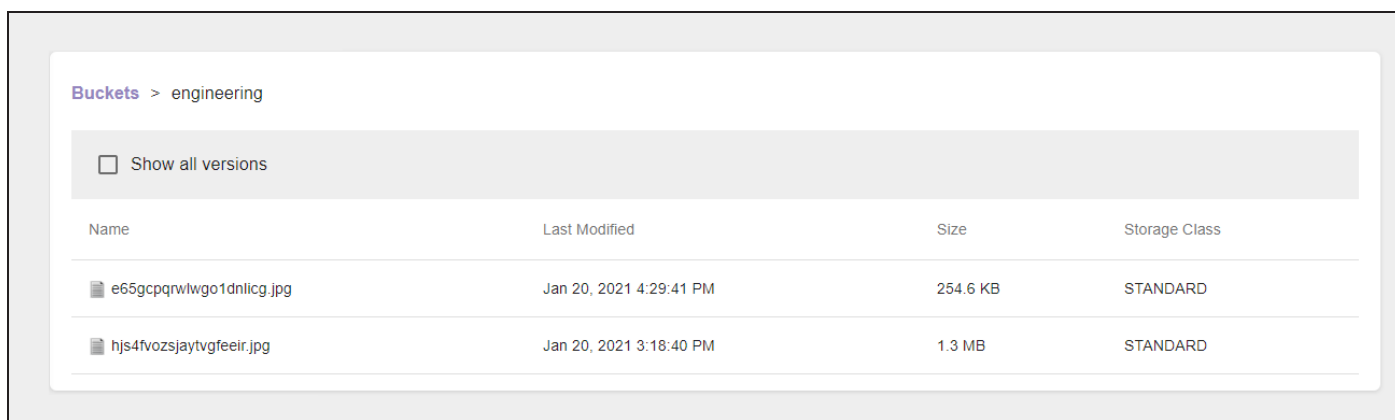
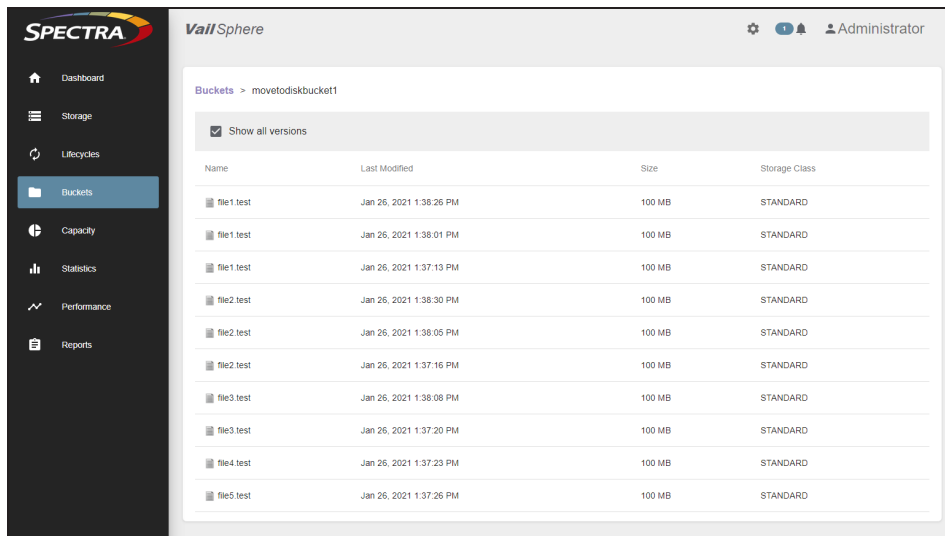


Figure 132 The Bucket Contents screen.

3. 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.



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 133 The Bucket Contents - Show All Versions screen.

4. 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.

10testfiletape

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 134 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.
Storage Class	The current storage class for the object.

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.
Restored Until	
Legal Hold	Indicates if the object has a legal hold.
Retention Mode	Indicates the retention mode.
Last Modified	The last modified date of the object.
Retain Until	The duration that the object is retained by a legal hold.

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

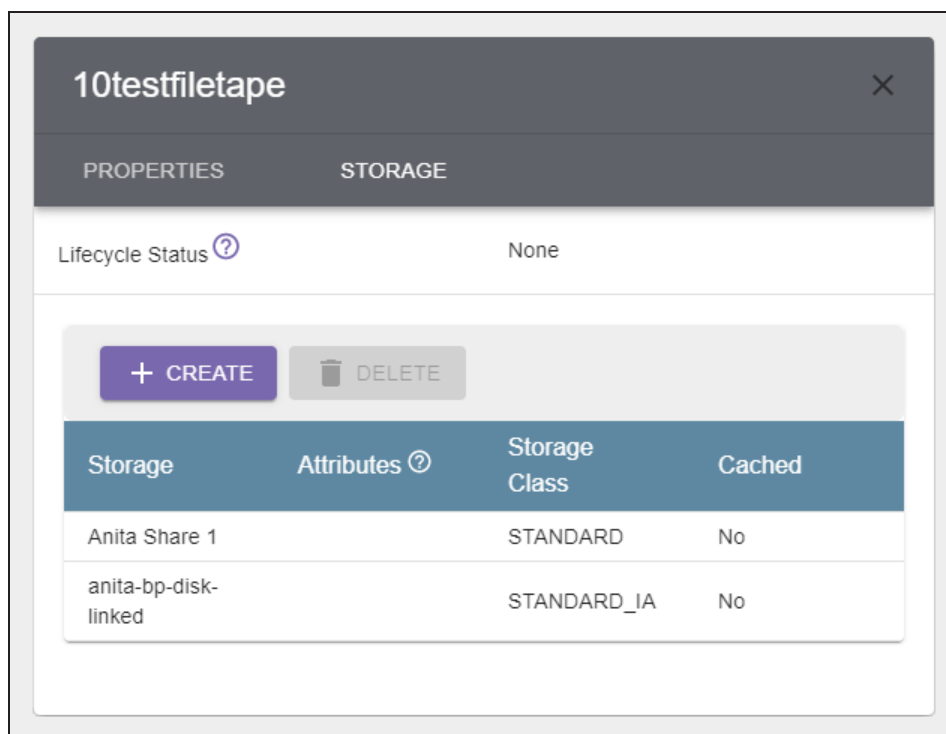


Figure 135 The Object Details - Storage screen.

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>Partial - Lifecycle processing is partially completed.</p> <p>Nearline - The object must be restored in order to access the object.</p>
Storage Class	<p>The current storage class for the object. See Storage Classes on page 209 for information on each storage class.</p> <p>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.</p>
Cached	If cached, the clone is deleted when space is required.

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

Here is how to change the nightly processing time:

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**.

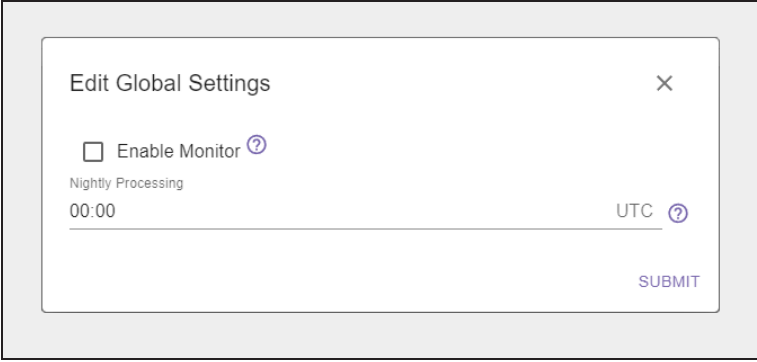


Figure 136 The Edit Global Settings screen.

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

Notes:

- Changing this value does not affect any actions that are already scheduled.
- All nodes must be rebooted after changing the Nightly Processing time.

4. 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**.

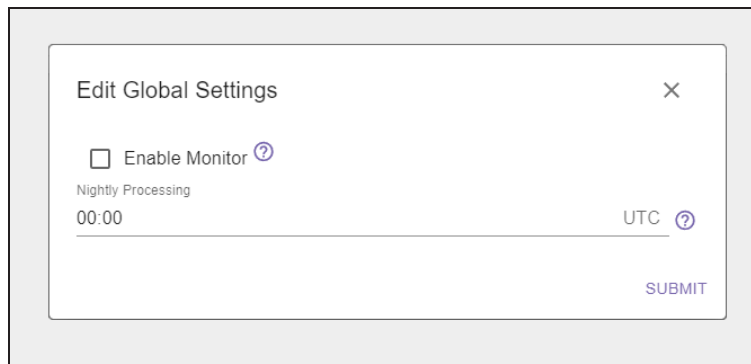


Figure 137 The Edit Global Settings screen.

3. Select **Enable Monitor**, then 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**.

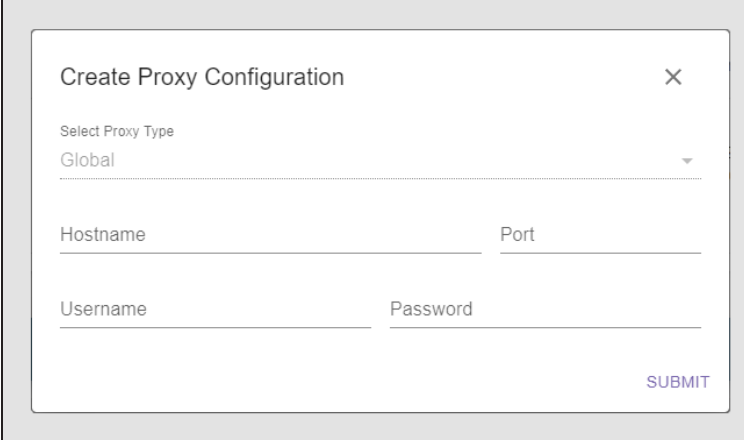


Figure 138 The Create Proxy Configuration screen.

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: You cannot disable versioning if the bucket was initially configured to use versioning AND object locking when it was created.

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**.

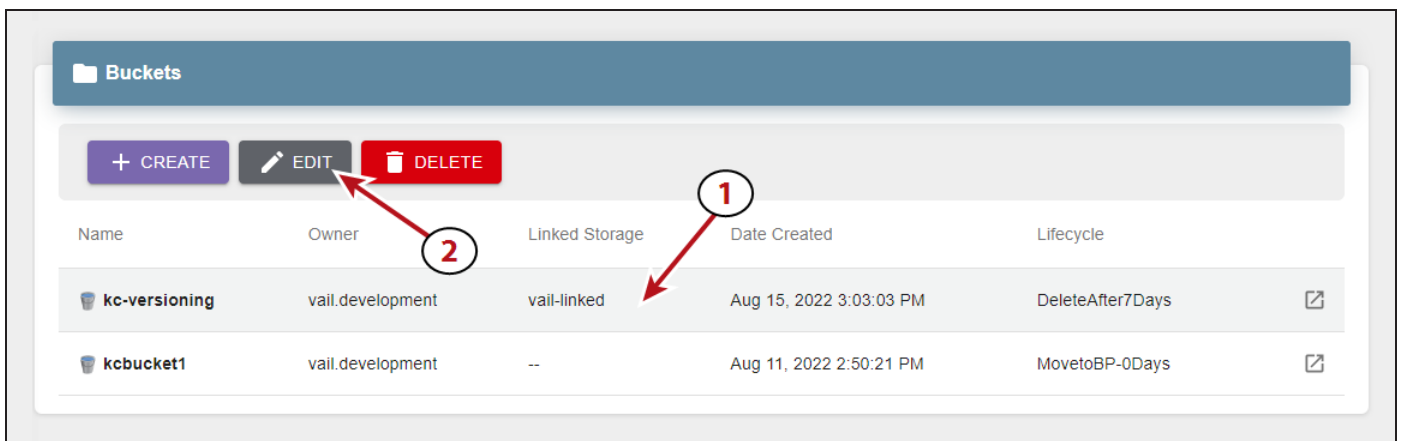


Figure 139 The Buckets pane.

3. Edit the settings on the Parameters screen as desired. See [Create a Vail Bucket on page 66](#) 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.

Figure 140 The Edit Bucket - Parameters screen.

- Notes:**
- You are not able to change the Bucket Name or Bucket Owner.
 - You cannot disable versioning if the bucket was initially configured to use versioning AND object locking when it was created.
 - 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 156](#).

5. Edit the settings on the Retention screen as desired. See [Create a Vail Bucket on page 66](#) for information about each feature and option on the Retention screen.

Figure 141 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 66](#) for information about each feature and option on the Policy screen.

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

Figure 142 The Edit Bucket - Policy screen.

8. Click **Next**.

9. Edit the settings on the Access Control Lists screen as desired. See [Create a Vail Bucket on page 66](#) 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.

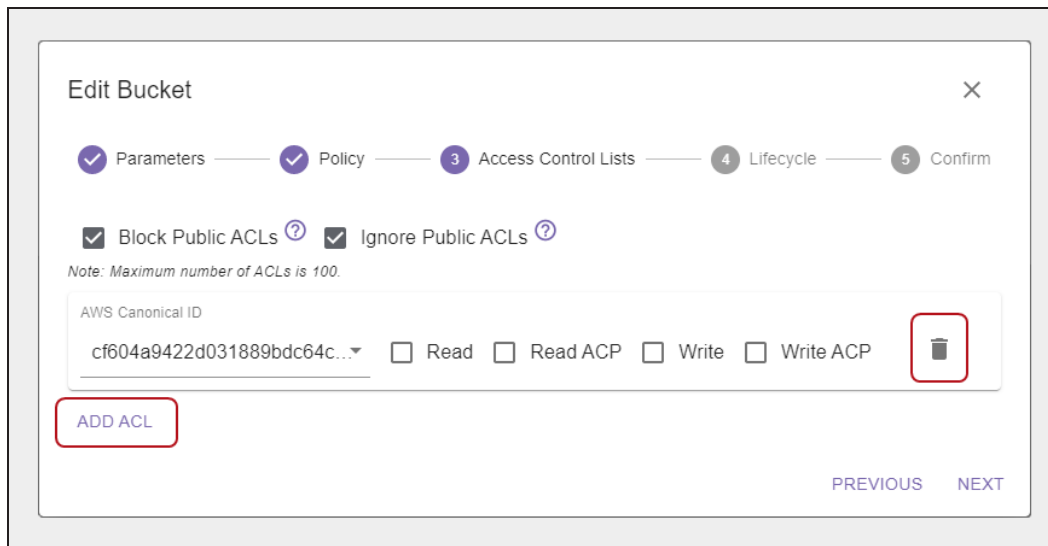


Figure 143 The Edit Bucket - Access Control Lists screen.

10. Click **Next**.

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

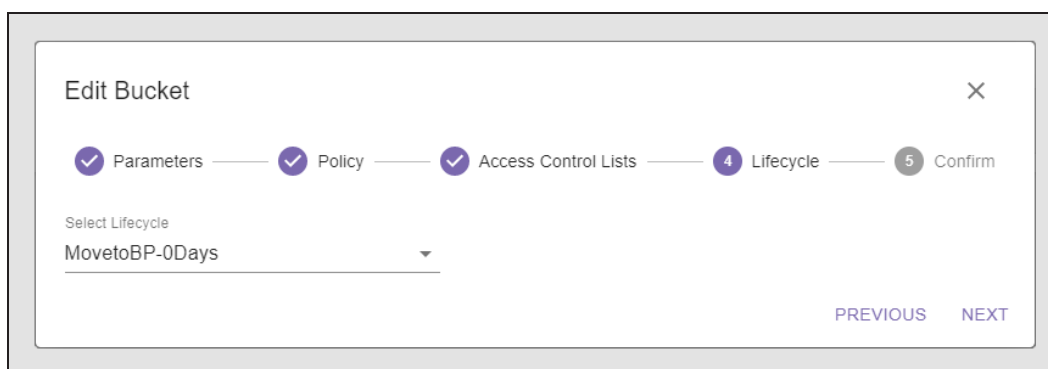


Figure 144 The Edit Bucket - Lifecycle screen.

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**.

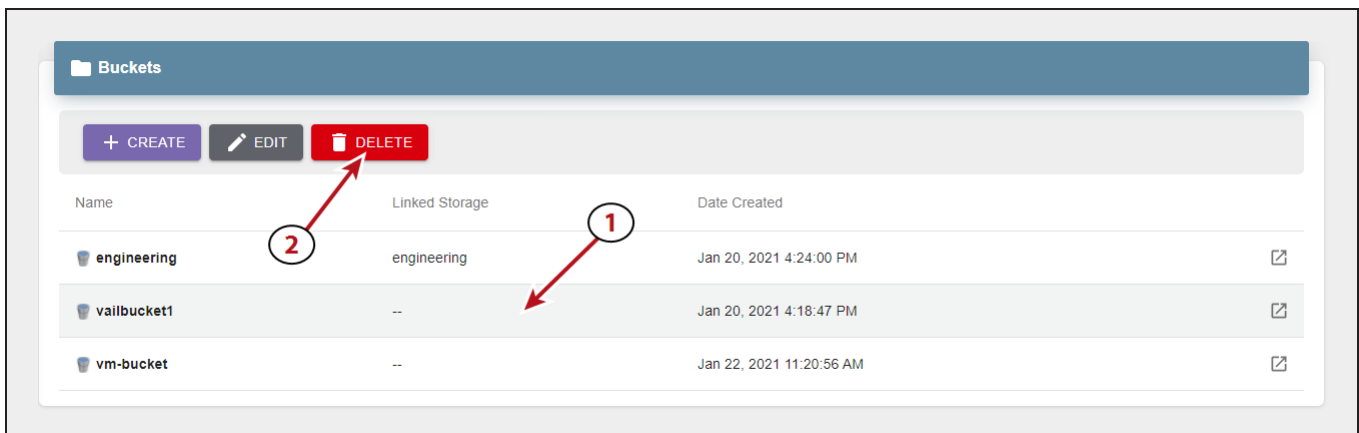


Figure 145 The Buckets pane.

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

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**.

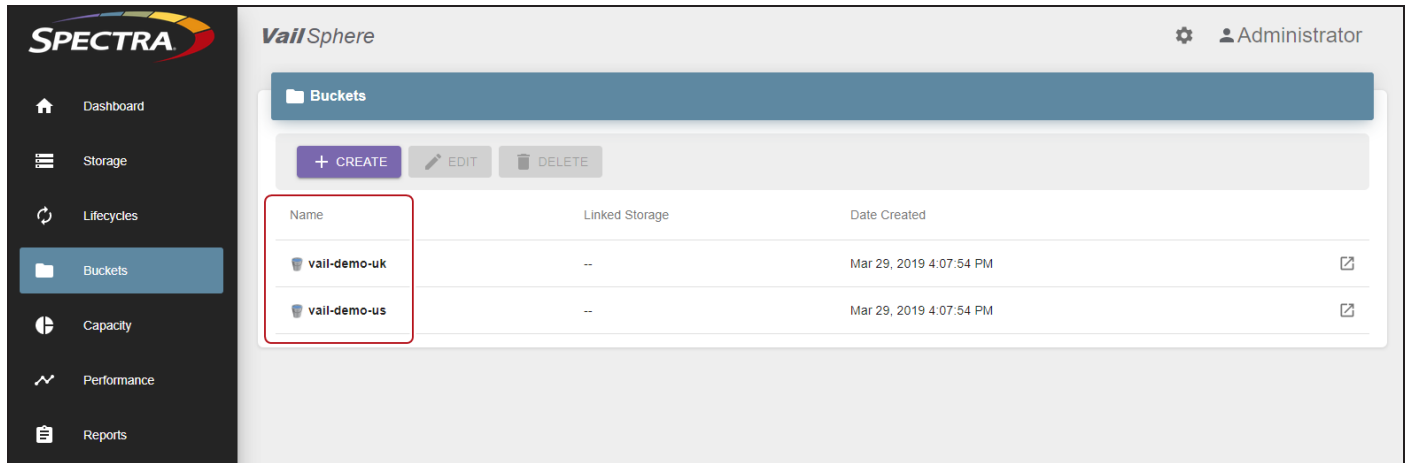


Figure 146 The Buckets screen.

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.

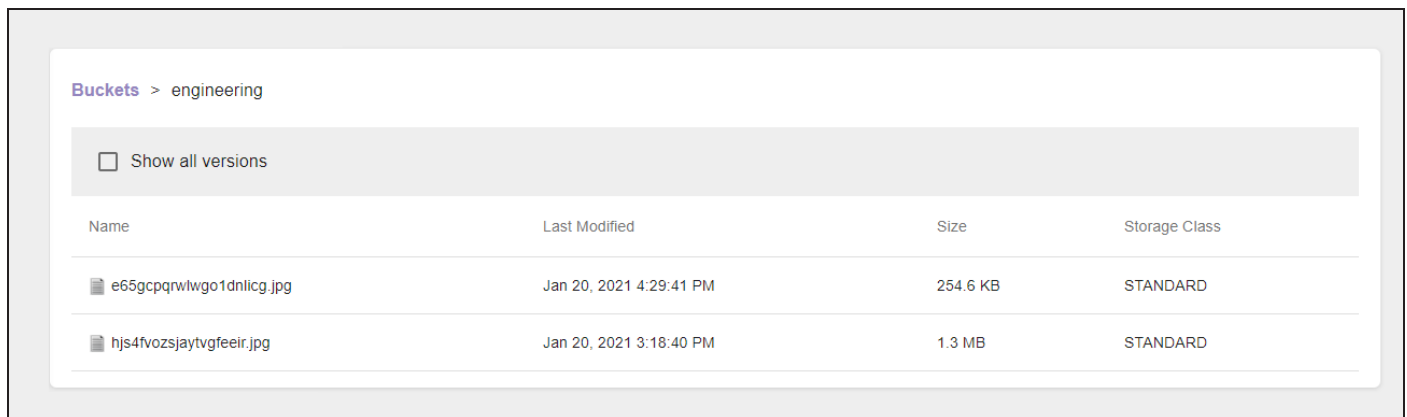


Figure 147 The Bucket Contents screen.

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.

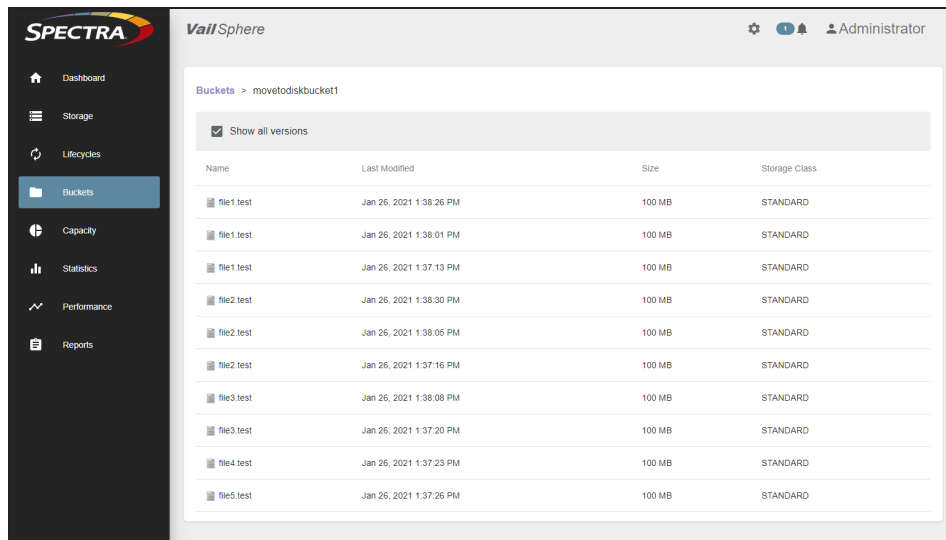


Figure 148 The Bucket Contents - Show All Versions screen.

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

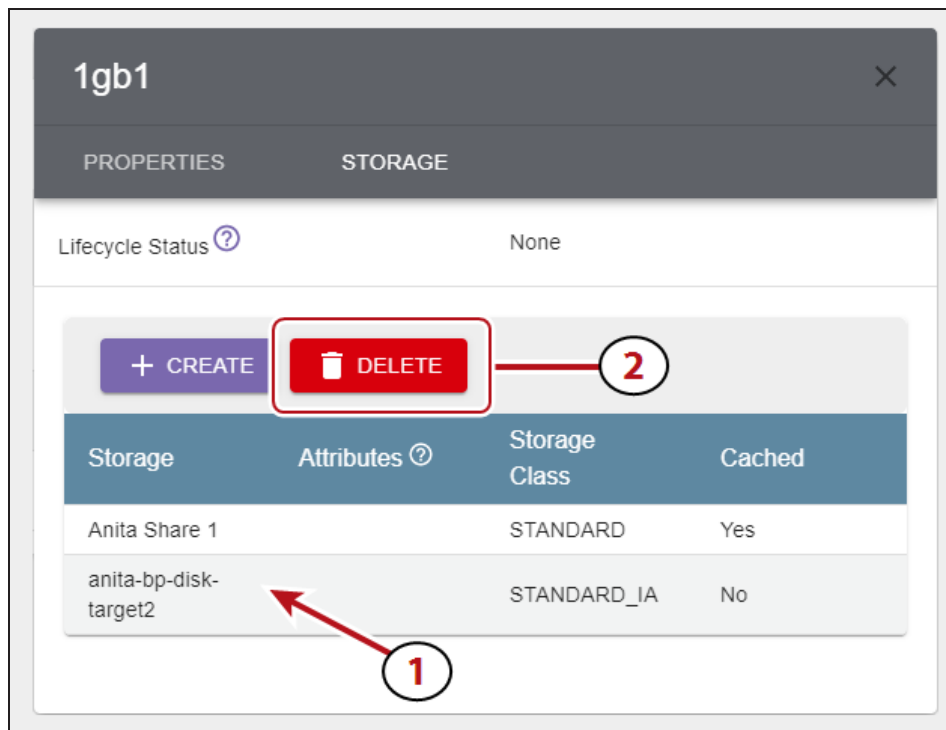


Figure 149 The Object Details - Storage screen.

6. Select the row of the object (1), and click **Create** (2).
7. Click **Create** on the confirmation screen to create an 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**.

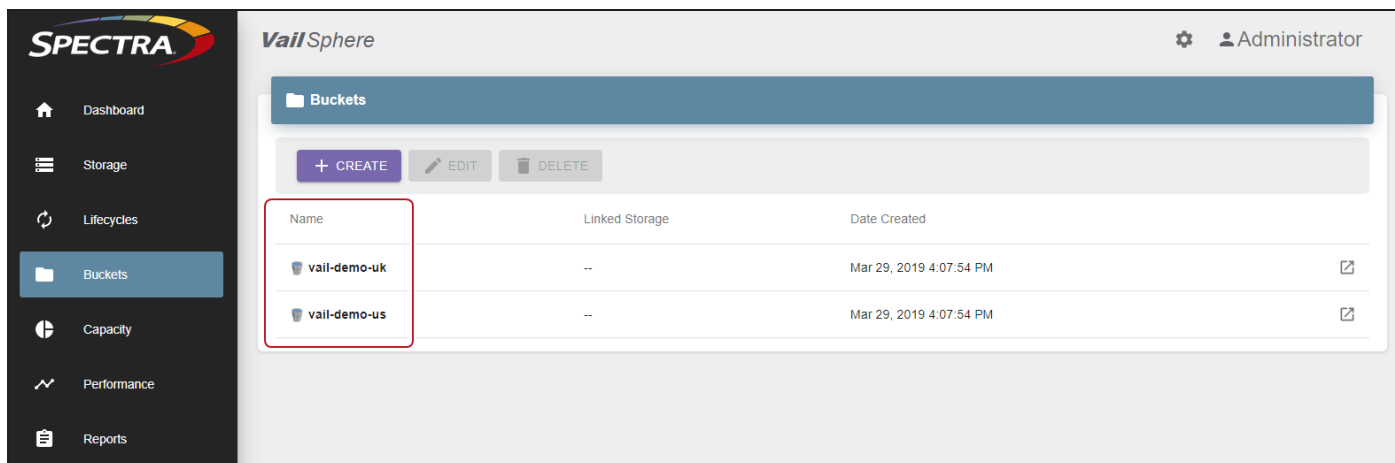


Figure 150 The Buckets screen.

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.

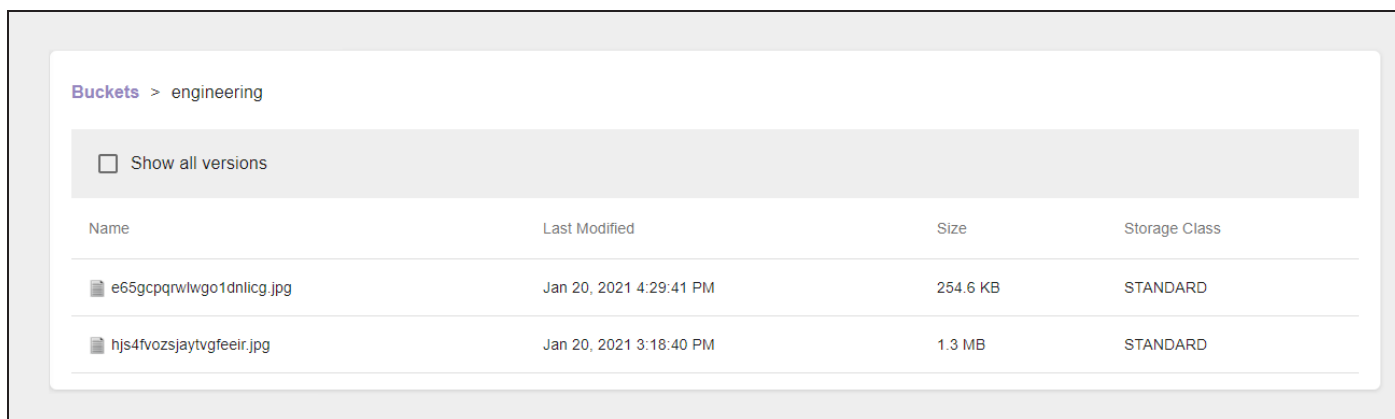


Figure 151 The Bucket Contents screen.

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 152 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 153 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.

VIEW STORAGE DETAILS

The storage detail screen displays the storage ID, as well as caution and warning thresholds of the selected storage.

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.

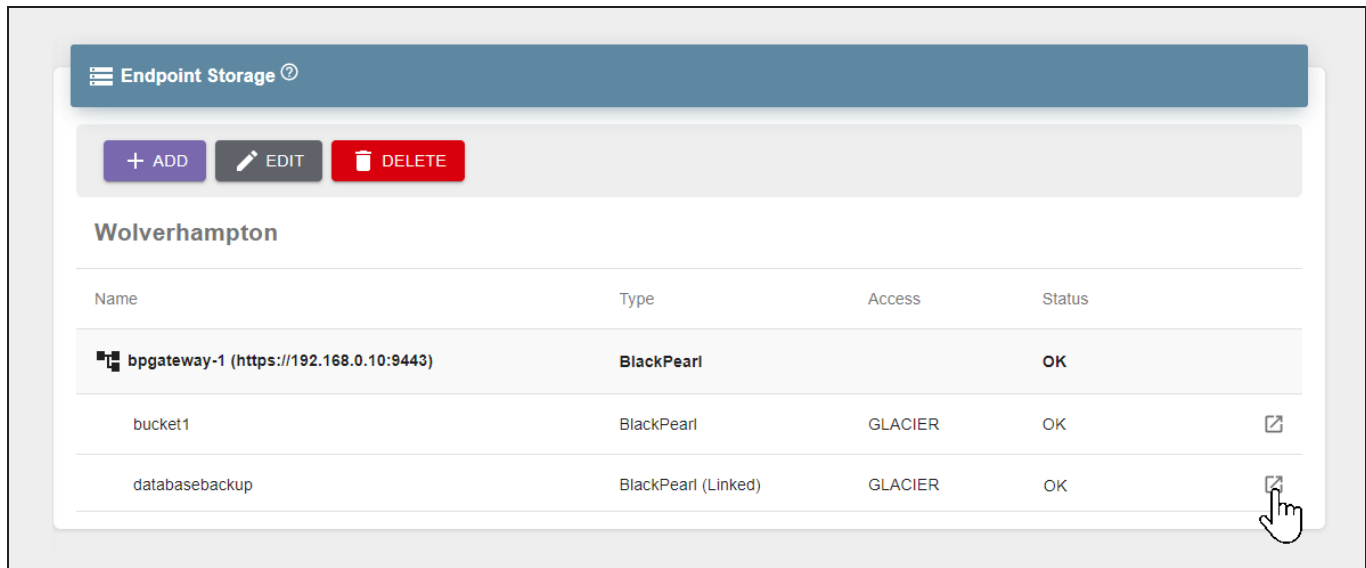


Figure 154 The Endpoint Storage pane.

3. The **Properties** screen displays the current storage endpoint settings.

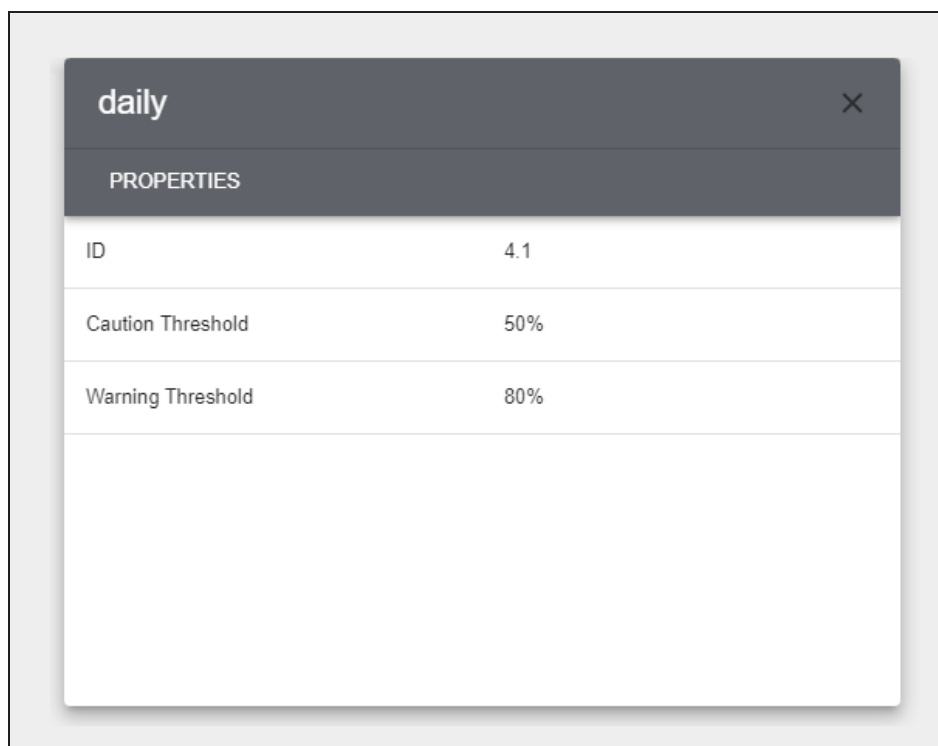


Figure 155 The Storage Details - Properties screen.

Note: Caution and Warning Thresholds do not display for linked storage.

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

EDIT BLACKPEARL OR VAIL VM ENDPOINT

If desired, you can edit the BlackPearl S3 solution 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**.

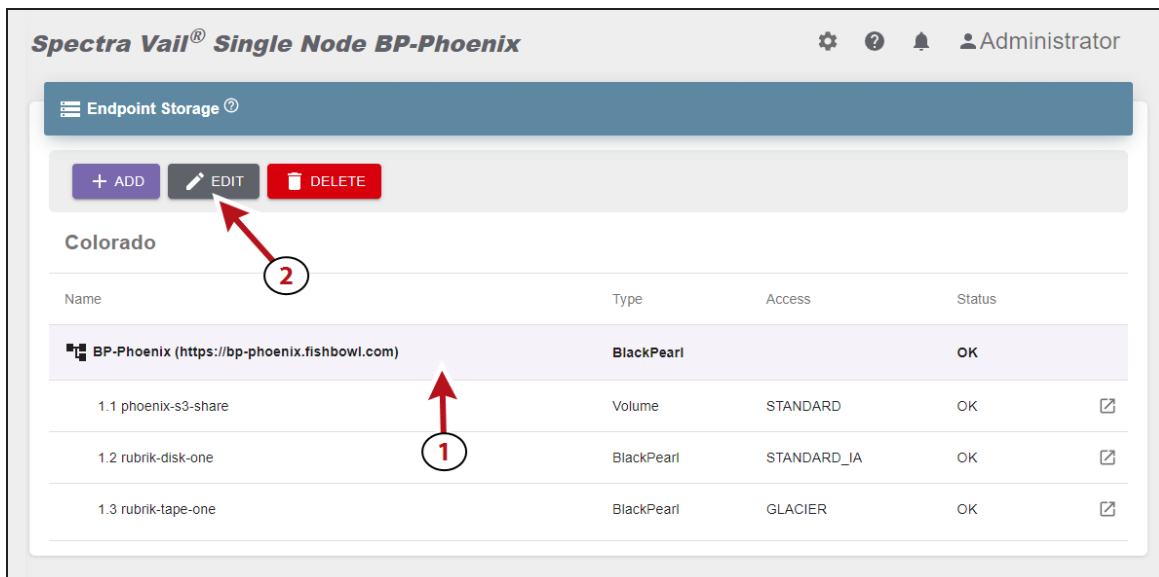


Figure 156 The Endpoint Storage pane.

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

Figure 157 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**.

Name	Type	Access	Status
BP-Phoenix (https://bp-phoenix.fishbowl.com)	BlackPearl		OK
1.1 phoenix-s3-share	Volume	STANDARD	OK
1.2 rubrik-disk-one	BlackPearl	STANDARD_IA	OK
1.3 rubrik-tape-one	BlackPearl	GLACIER	OK

Figure 158 The Endpoint Storage pane.

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

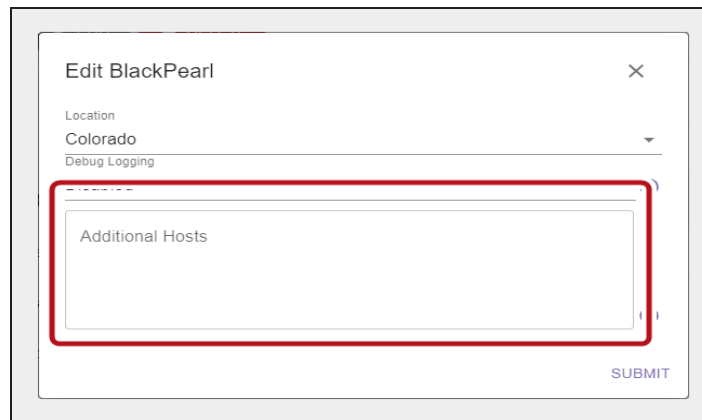


Figure 159 Edit *Endpoint*- Additional Hosts screen.

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**.

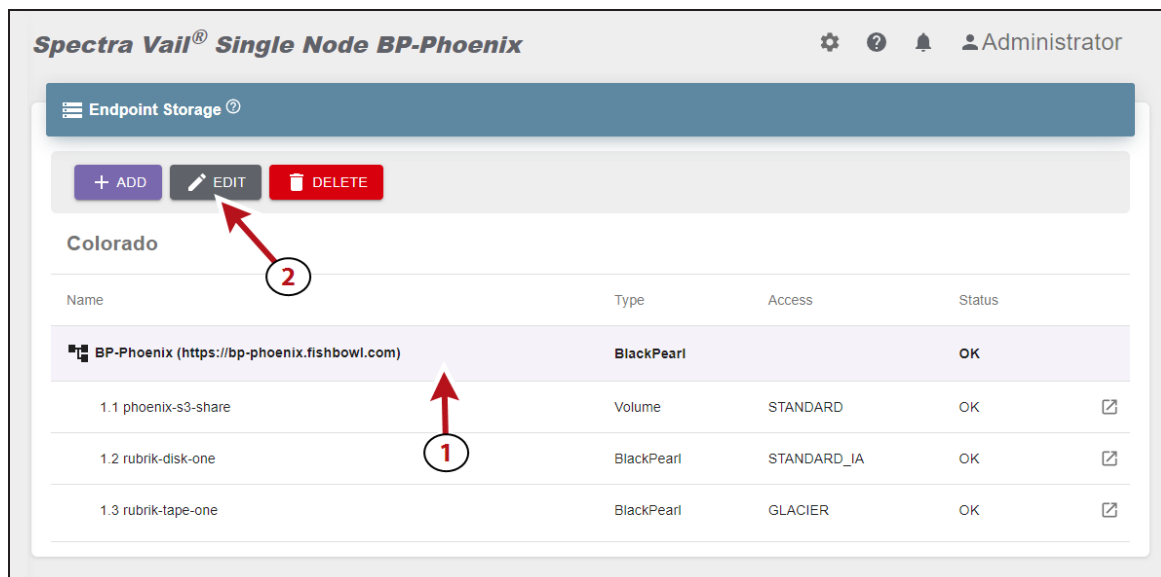


Figure 160 The Endpoint Storage pane.

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

Edit BlackPearl

Location

Debug Logging
Disabled

Additional Hosts

SUBMIT

Figure 161 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 Bucket Storage below**
- **Edit Vail VM Node Storage on page 173**
- **Edit Google Cloud Platform Storage on page 174**
- **Edit AWS S3 Cloud Storage on page 176**
- **Edit Other Third-Party Cloud Storage on page 179**

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**.

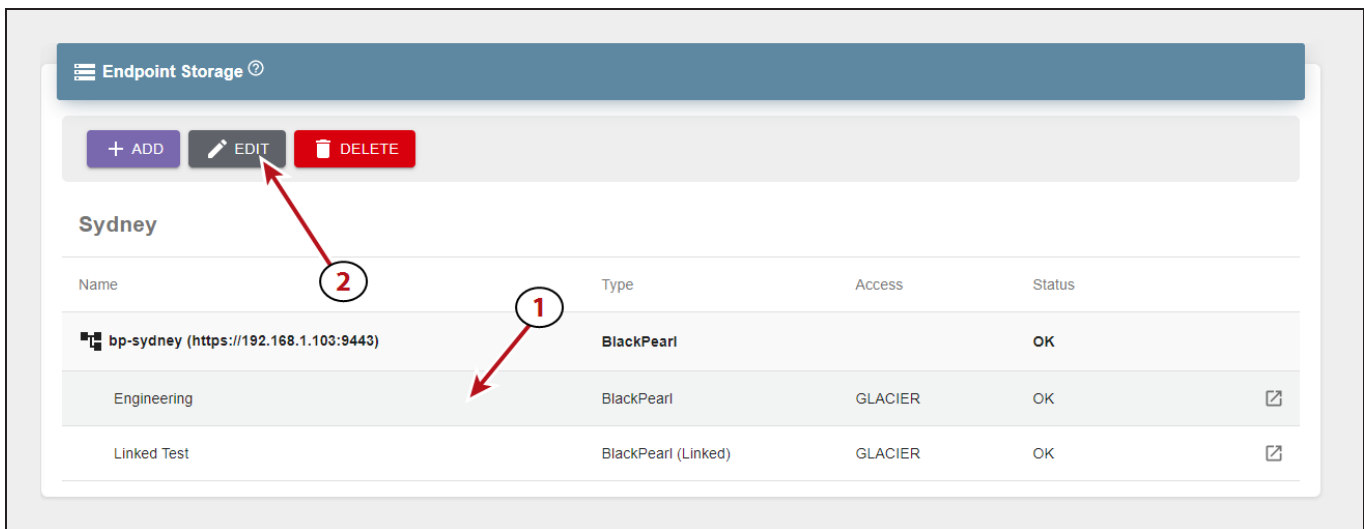


Figure 162 The Endpoint Storage pane.

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

The screenshot shows a dialog box titled "Edit Endpoint Storage" with a close button (X) in the top right corner. Below the title is a progress indicator with three steps: "1 Options" (active), "2 Parameters", and "3 Confirm". The main text asks the user to "Select how you would like to edit 1.2 rubrik-disk-one." Below this text are two radio button options: "Modify parameters" (selected) and "Modify authorization". A "NEXT" button is located in the bottom right corner of the dialog box.

Figure 163 Edit Endpoint Storage - Options screen.

If you selected Modify Parameters...

- a. If desired, you can change the **Storage Name**, **Storage Class**, **Caution Threshold**, or **Warning Threshold**. Additionally, you can configure the storage to **Restore To New Clone**, which creates an additional data clone on different storage.

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

The screenshot shows a modal window titled "Edit Endpoint Storage" with a close button (X) in the top right. Below the title is a progress bar with three steps: "Options" (checked), "Parameters" (active), and "Confirm". The main instruction is "Configure your BlackPearl storage below." The form contains several fields: "BlackPearl Bucket" with a dropdown menu showing "rubrik-tape-one"; "Storage Name" with a text input showing "1.3 rubrik-tape-one"; "Link to Bucket" with a dropdown menu showing "Do Not Link"; "Select Storage Class" with a dropdown menu showing "GLACIER"; a checkbox labeled "Restore To New Clone"; "Caution Threshold" with a text input showing "70" and a percentage symbol; and "Warning Threshold" with a text input showing "80" and a percentage symbol. There are question mark icons next to several fields. At the bottom right are "PREVIOUS" and "NEXT" buttons.

Figure 164 The Edit Endpoint Storage - Parameters - BlackPearl screen.

- b. Click **Next**.
- c. Review the configuration, and click **Submit** to save the changes to the BlackPearl storage.

If you selected Modify Authorization...

- a. If desired, you can enter a new **BlackPearl S3 Access ID** and **BlackPearl S3 Secret Key**, then click **Next**.

Figure 165 The Edit Endpoint Storage - Authorization- BlackPearl screen.

- b. Review the configuration, and click **Submit** to save the changes to the BlackPearl storage.

Edit BlackPearl NAS Storage

Here is how to edit BlackPearl NAS 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**.

Name	Type	Access	Status
bp-sydney (https://192.168.1.103:9443)	BlackPearl		OK
Engineering	BlackPearl	GLACIER	OK
Linked Test	BlackPearl (Linked)	GLACIER	OK

Figure 166 The Endpoint Storage pane.

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

The screenshot shows a modal window titled "Edit Endpoint Storage" with a close button (X) in the top right. Below the title is a progress indicator with two steps: "1 Parameters" (active) and "2 Confirm". The instruction "Configure your storage below." is displayed. The form contains three sections:

- Storage Name: 1.1 phoenix-s3-share
- Select Storage Class: STANDARD (dropdown menu)
- Caution Threshold: 70 %
- Warning Threshold: 80 % (with a help icon ?)
- Optional Data: 80 % (with a help icon ?)

 A "NEXT" button is located at the bottom right of the form.

Figure 167 The Edit Endpoint Storage - Parameters - NAS screen.

4. Click **Next**.
5. 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, you can change the **Storage Name**, **Caution Threshold**, **Warning Threshold**, or **Storage Class**.

Edit Endpoint Storage [X]

1 Parameters ————— 2 Confirm

Configure your storage below.

Storage Name vaiVM1 Storage 1	Select Storage Class STANDARD ▼
Caution Threshold 70 %	Warning Threshold 80 % ⓘ

NEXT

Figure 168 The Edit Endpoint Storage - Parameters - Vail VM Node screen.

4. Click **Next**.
5. 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**.

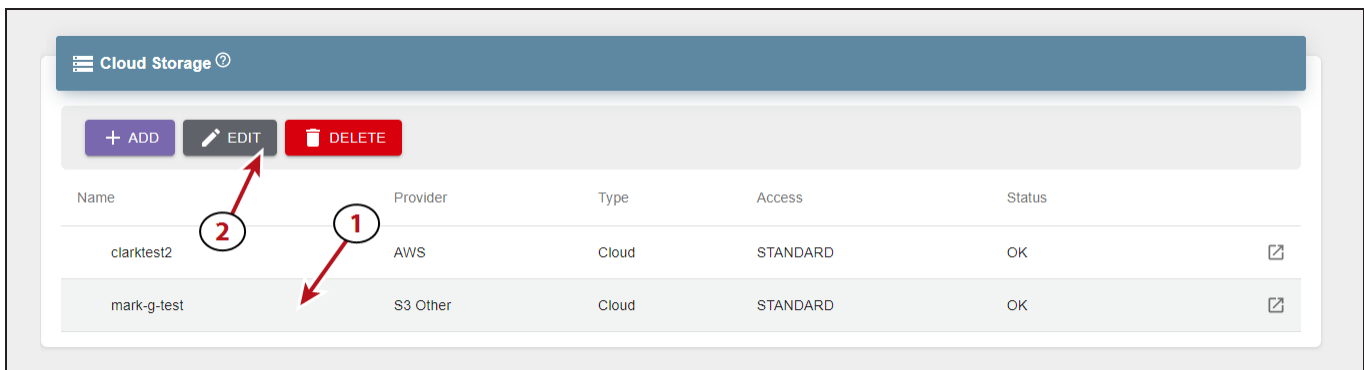


Figure 169 The Endpoint Storage pane.

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

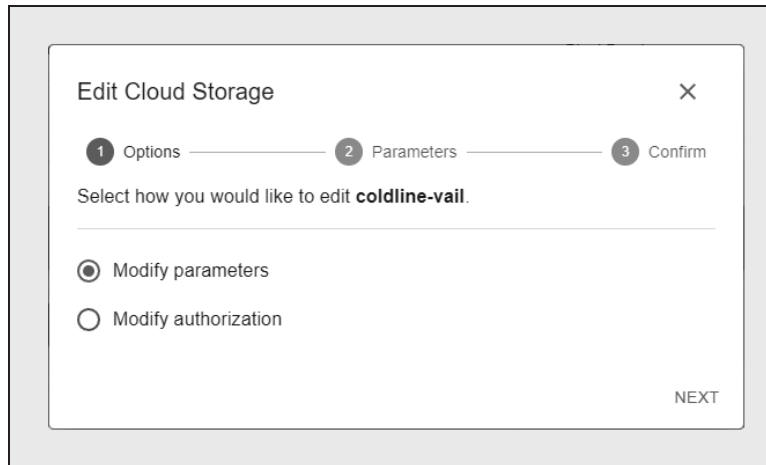
The screenshot shows a dialog box titled "Edit Cloud Storage" with a close button (X) in the top right corner. Below the title is a progress bar with three steps: "1 Options" (active), "2 Parameters", and "3 Confirm". The main text asks, "Select how you would like to edit coldline-vail." Below this are two radio button options: "Modify parameters" (selected) and "Modify authorization". A "NEXT" button is located in the bottom right corner.

Figure 170 The Edit Cloud Storage - Options screen.

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

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

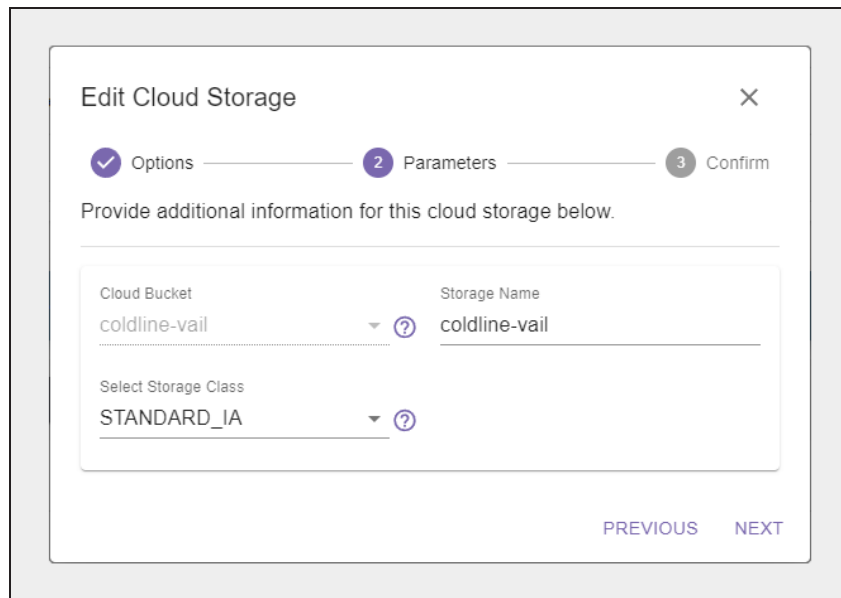
The screenshot shows the "Edit Cloud Storage" dialog box at the "Parameters" step. The progress bar now shows "1 Options" as completed with a checkmark, and "2 Parameters" as the active step. The main text says, "Provide additional information for this cloud storage below." Below this is a form with three fields: "Cloud Bucket" (a dropdown menu showing "coldline-vail"), "Storage Name" (a text input field containing "coldline-vail"), and "Select Storage Class" (a dropdown menu showing "STANDARD_IA"). Each field has a help icon (question mark in a circle). At the bottom right, there are "PREVIOUS" and "NEXT" buttons.

Figure 171 The Edit Cloud Storage - Parameters - Google Cloud Platform screen.

- b. Click **Next**.
- c. 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.

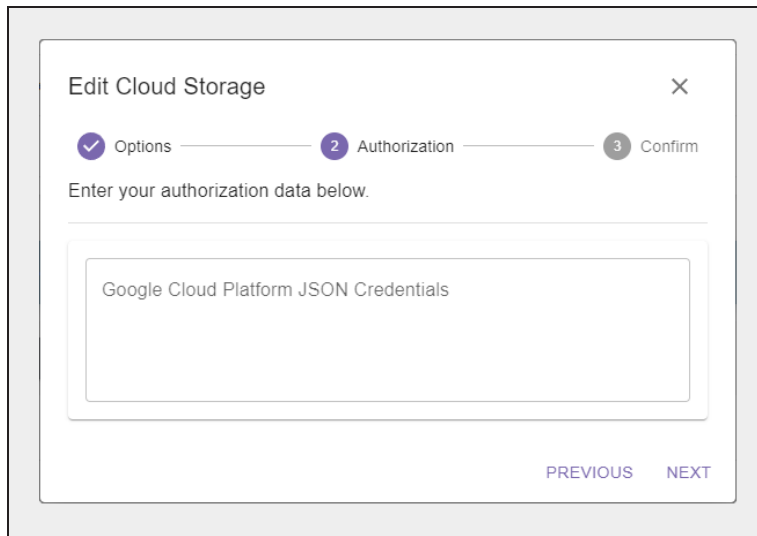


Figure 172 The Edit Cloud Storage - Authorization - Google Cloud Platform screen.

- 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**.

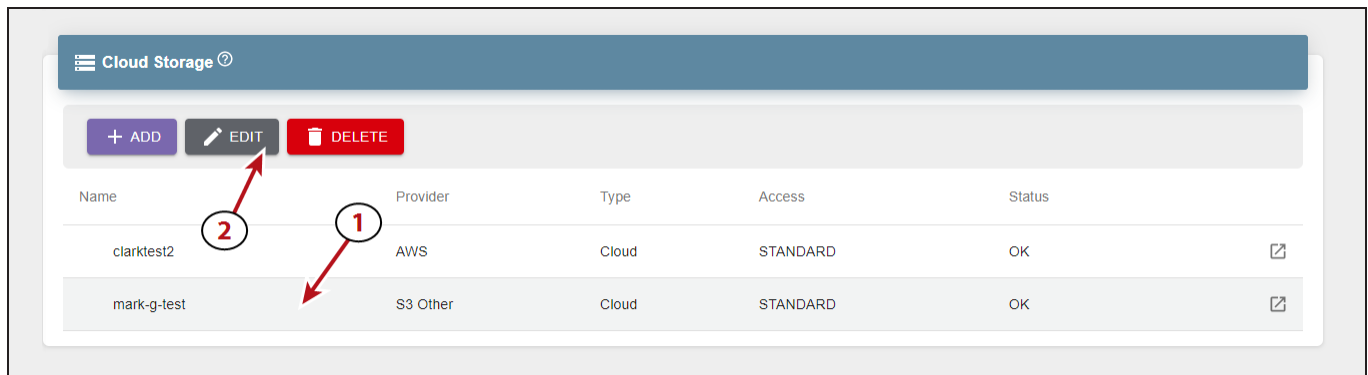


Figure 173 The Endpoint Storage pane.

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




Figure 174 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]

✓ Options — 2 Parameters — 3 Confirm

Provide additional information for this cloud storage below.

Cloud Bucket: clarktest2 [?] Storage Name: clarktest2

Link to Bucket: Do Not Link [?] Select Storage Class: STANDARD [?]

☐ Restore To New Clone [?]

PREVIOUS NEXT

Figure 175 The Edit Cloud Storage - Parameters - AWS S3 Storage screen.

- b. If you selected **Deep_Archive** in [Step a](#), select **Restore to New Clone**, if desired. This option creates a new clone on different storage, instead of using the existing archival storage.
- c. Click **Next**.
- d. Review the configuration, and click **Submit** to save the changes to the cloud storage.

Edit Cloud Storage [X]

✓ Options — 2 Authorization — 3 Confirm

Enter your authorization data below.

☒ Use Credentials ☐ Use IAM Role [?]

AWS Access Key ID: _____ AWS Secret Access Key: _____

PREVIOUS NEXT

Figure 176 The Edit Cloud Storage - Authorization - AWS S3 Storage screen.

- e. Click **Next**.

- f. Review the configuration, and click **Submit** to save the changes to the cloud storage.

Edit Other Third-Party Cloud Storage

Here is how to edit other third-party 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**.

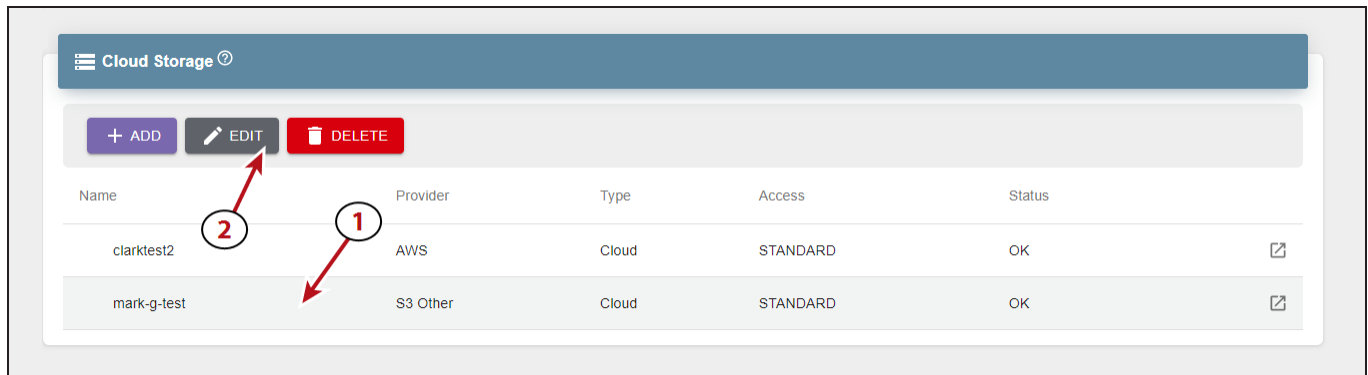


Figure 177 The Endpoint Storage pane.

3. Proceed through the Edit Cloud Storage wizard, then click **Save**.

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 ensure 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**.

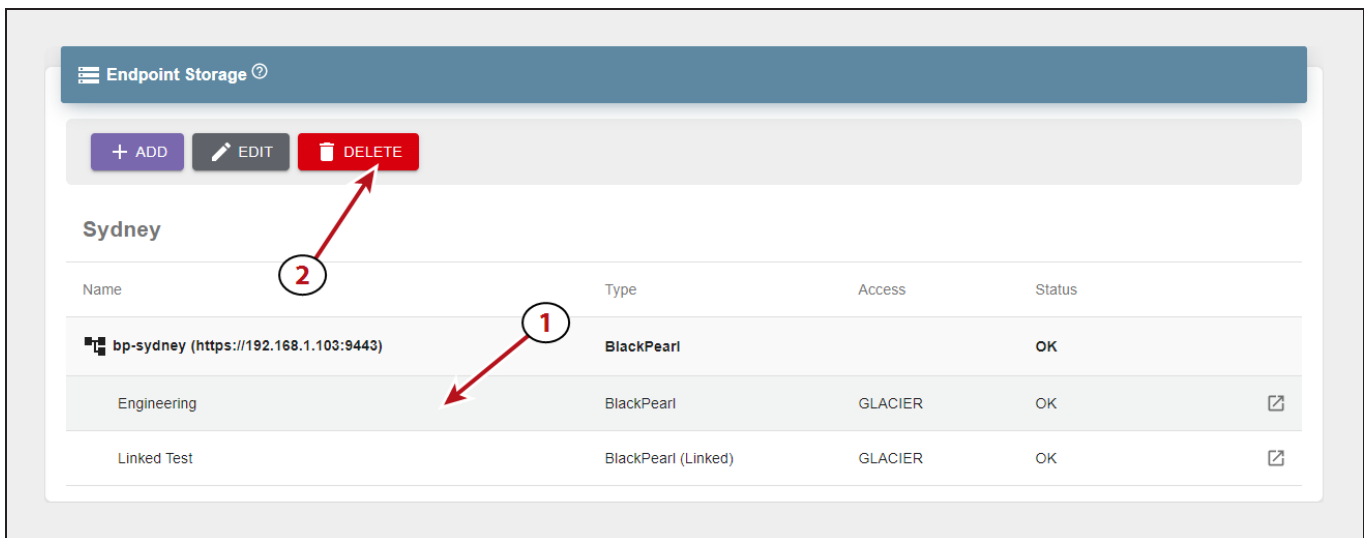
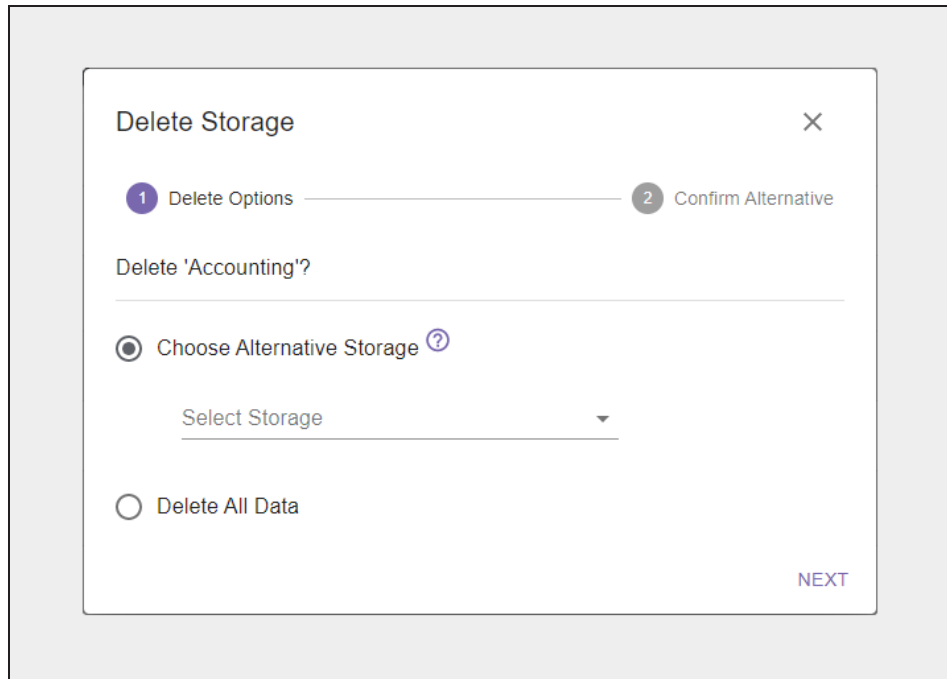


Figure 178 The Endpoint Storage pane.

- To move unique object data to alternative storage:

1. Select **Choose Alternative Storage**.



The screenshot shows a 'Delete Storage' dialog box with a close button (X) in the top right corner. The dialog has a progress bar with two steps: '1 Delete Options' (active) and '2 Confirm Alternative'. Below the progress bar, the text 'Delete 'Accounting'?' is followed by a horizontal line. Under this line, there are two radio button options. The first option is 'Choose Alternative Storage' with a question mark icon, and it is selected. Below this option is a 'Select Storage' drop-down menu. The second option is 'Delete All Data'. In the bottom right corner of the dialog, there is a 'NEXT' button.

Figure 179 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.

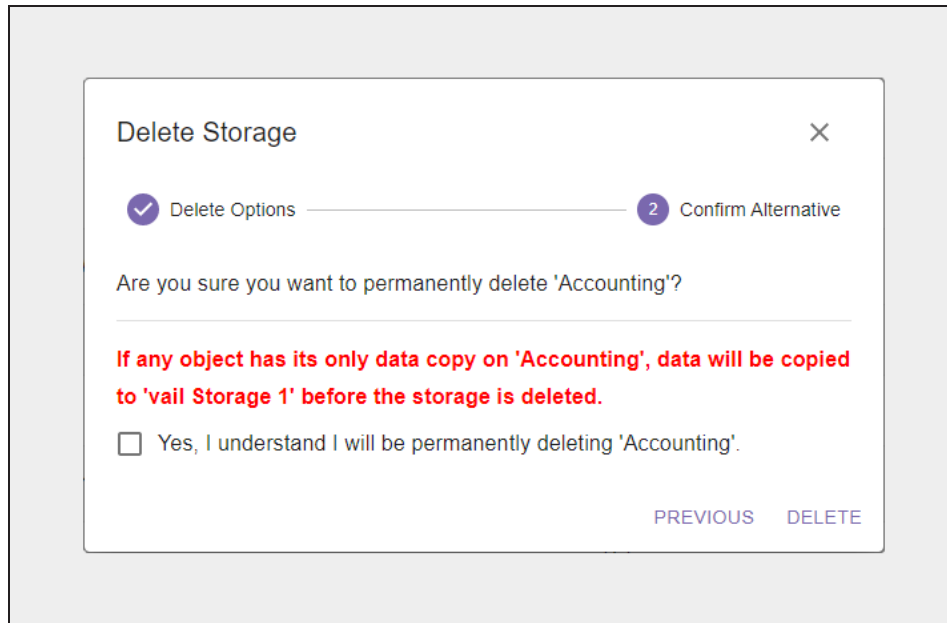


Figure 180 The Delete Storage - Confirm Alternative screen.

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 exists only on the storage is permanently deleted and cannot be recovered.

2. Enter the name of the storage in **Confirmed Name** field.

Figure 181 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 182 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:

Copy Rule to Cloud	
PROPERTIES	RULES
Description	Text entered in the Description field when creating a Lifecycle.
Upload Expiration	3
Marker Expiration	Enabled
Modified	Jul 06, 2021 5:05:32 PM

Figure 183 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 .
Modified	The date and time the lifecycle was last modified.

- The Rules screen:

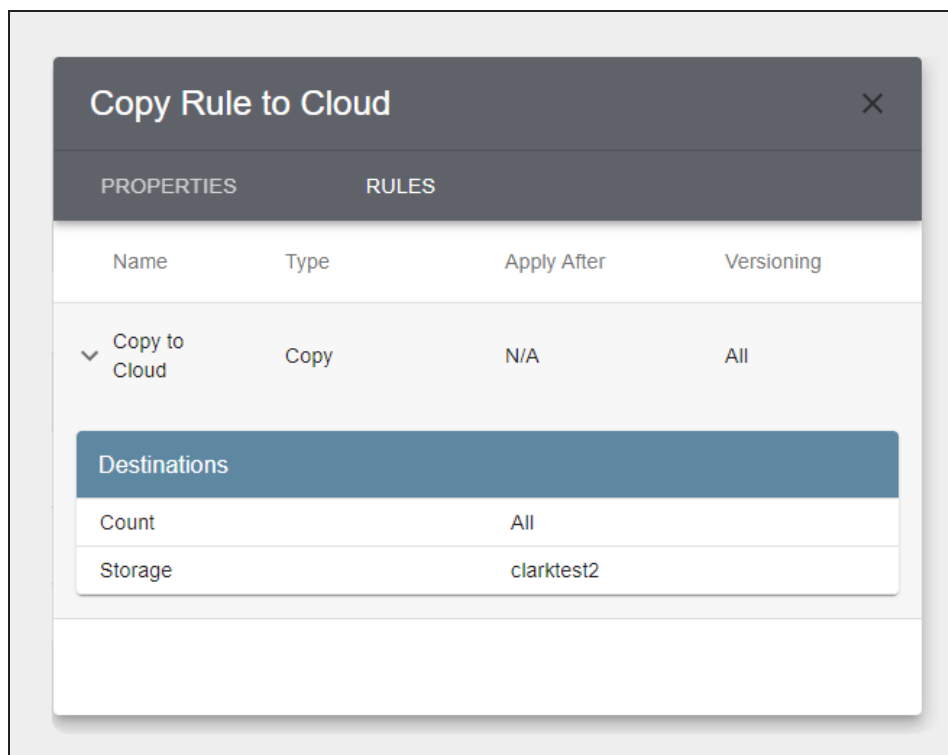


Figure 184 The Lifecycle Rule Details - Rules screen.

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.
Versioning	Indicates the level of Versioning used by the lifecycle.
Destinations - Count	The number of destinations configured for the lifecycle. Values: 1-5, All.
Destinations - Storage	The storage endpoint(s) used by the lifecycle.

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**.

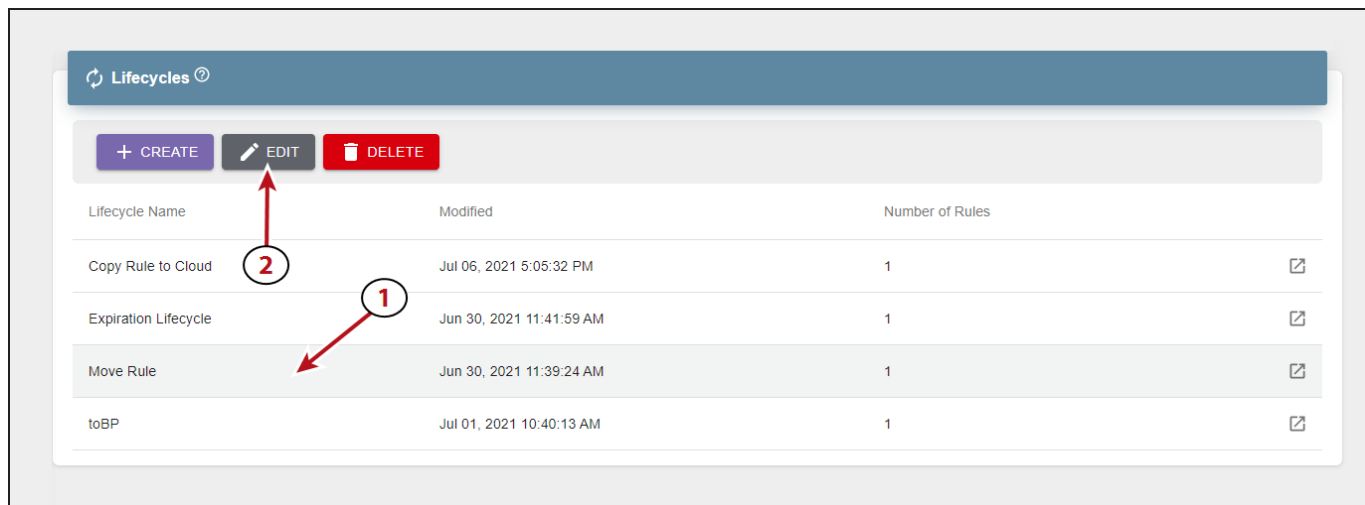
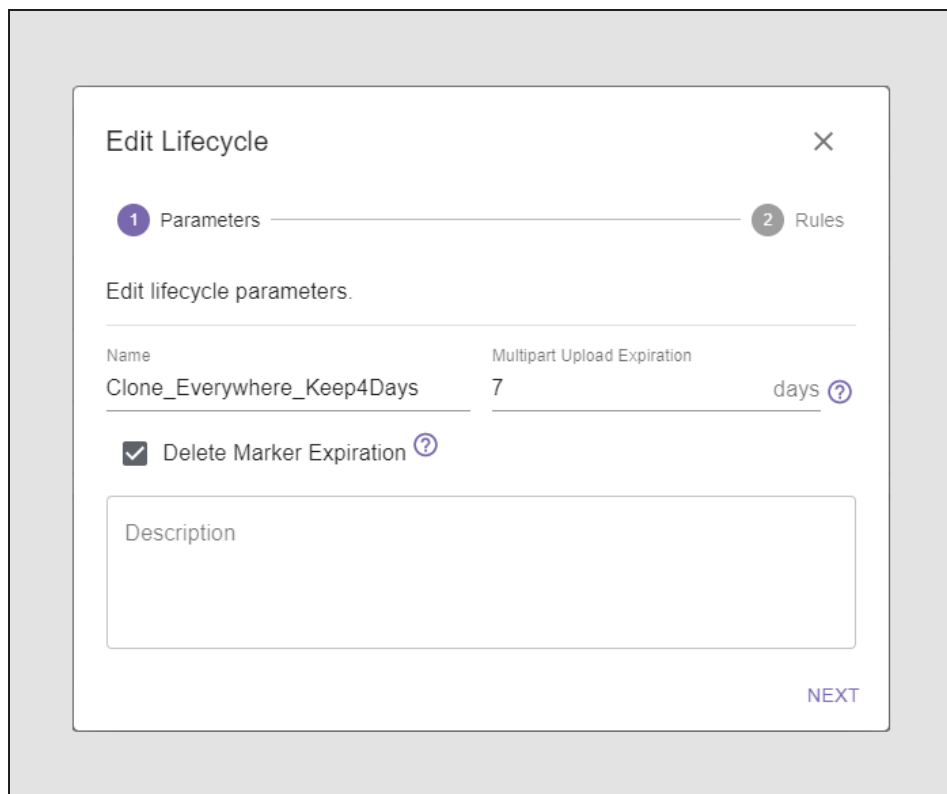


Figure 185 The Lifecycles screen.

3. If desired, edit the lifecycle **Name**, **Multipart Upload Expiration**, **Description**, and select or clear **Delete Marker Expiration**.



The screenshot shows a modal dialog titled "Edit Lifecycle" with a close button (X) in the top right corner. Below the title is a progress indicator with two steps: "1 Parameters" (active) and "2 Rules". The main content area is titled "Edit lifecycle parameters." and contains the following fields:

- Name:** A text input field containing "Clone_Everywhere_Keep4Days".
- Multipart Upload Expiration:** A text input field containing "7", followed by a "days" label and a help icon (?).
- Delete Marker Expiration:** A checkbox that is checked, followed by a help icon (?).
- Description:** A large text area for entering a description.

A "NEXT" button is located in the bottom right corner of the dialog.

Figure 186 The Edit Lifecycle - Parameters screen.

4. Click **Next**.

Create Lifecycle ? ✕

✓ Parameters ————— 2 Rules

Define your rules for **Lifecycle2** ?

Note: Rules will be sorted by "Apply After" and "Versioning" values after submission. Maximum number of rules is 5.

Placement Rule

Name

Select Destination Storage ?

☐ Delete clones not on selected destination storage

Destination Count: All

Apply After: days ? Versioning: All

Include ?

Exclude ?

NEW PLACEMENT RULE | NEW DELETION RULE

PREVIOUS SUBMIT

Figure 187 The Edit Lifecycle - Rules screen.

- Use the links below to create or edit lifecycle rules.
 - **Add a Placement Rule on page 60**
 - **Add a Deletion Rule on page 63**
- To delete a lifecycle rule, click the **trash can icon**.

Note: There is no confirmation step for this action.

5. 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 188 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.

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 189 The Dashboard screen - Navigation menu.

2. Click **Create**.



Figure 190 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.

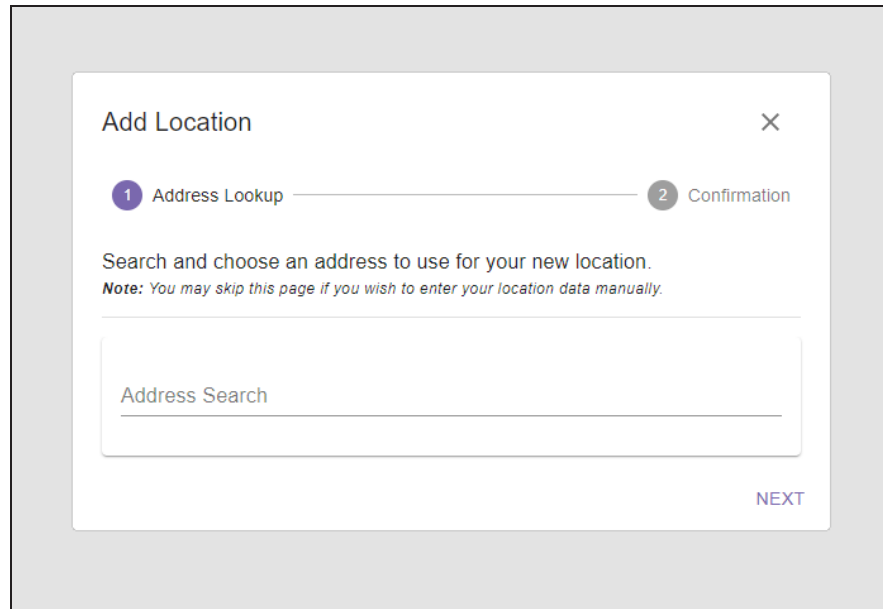


Figure 191 The Add Location - Address Lookup screen.

- 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**.

The screenshot shows a modal window titled "Add Location" with a close button (X) in the top right corner. Below the title is a progress bar with two steps: "1 Address Lookup" (marked with a checkmark) and "2 Confirmation" (marked with a circle containing the number 2). Below the progress bar, there is a text block: "Please confirm the details below. If necessary, you may edit any pre-populated fields or return to the previous page and execute another search." followed by a note: "Note: Latitude and Longitude values are used for the System View map on the dashboard." Below this text is a form with three input fields: "Name", "Latitude", and "Longitude". At the bottom right of the form are two buttons: "PREVIOUS" and "SUBMIT".

Figure 192 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.

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)**.

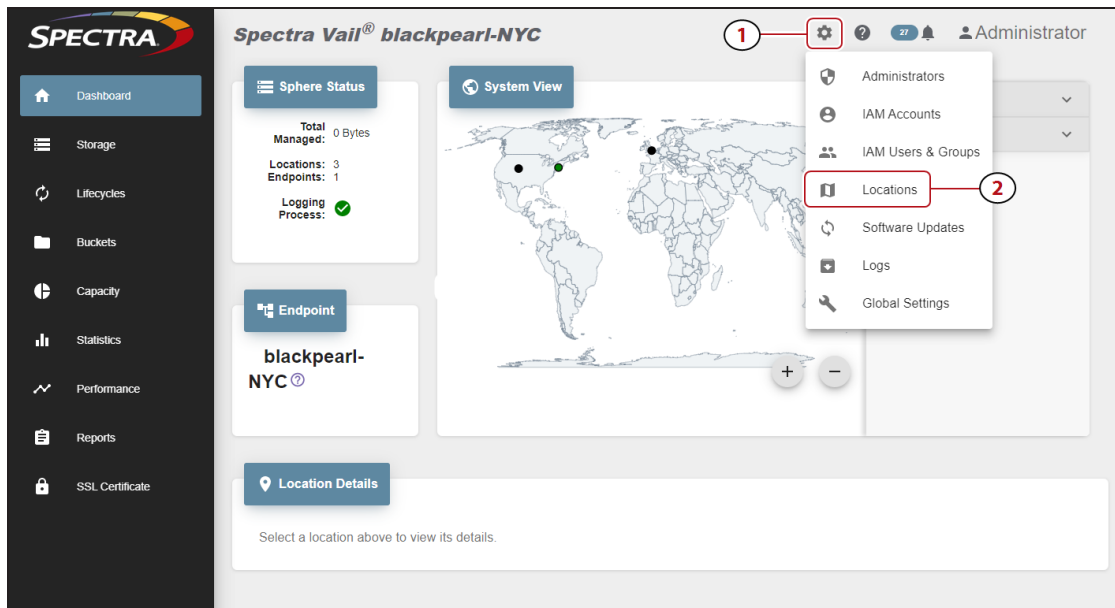


Figure 193 The Dashboard screen - Navigation menu.

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**.


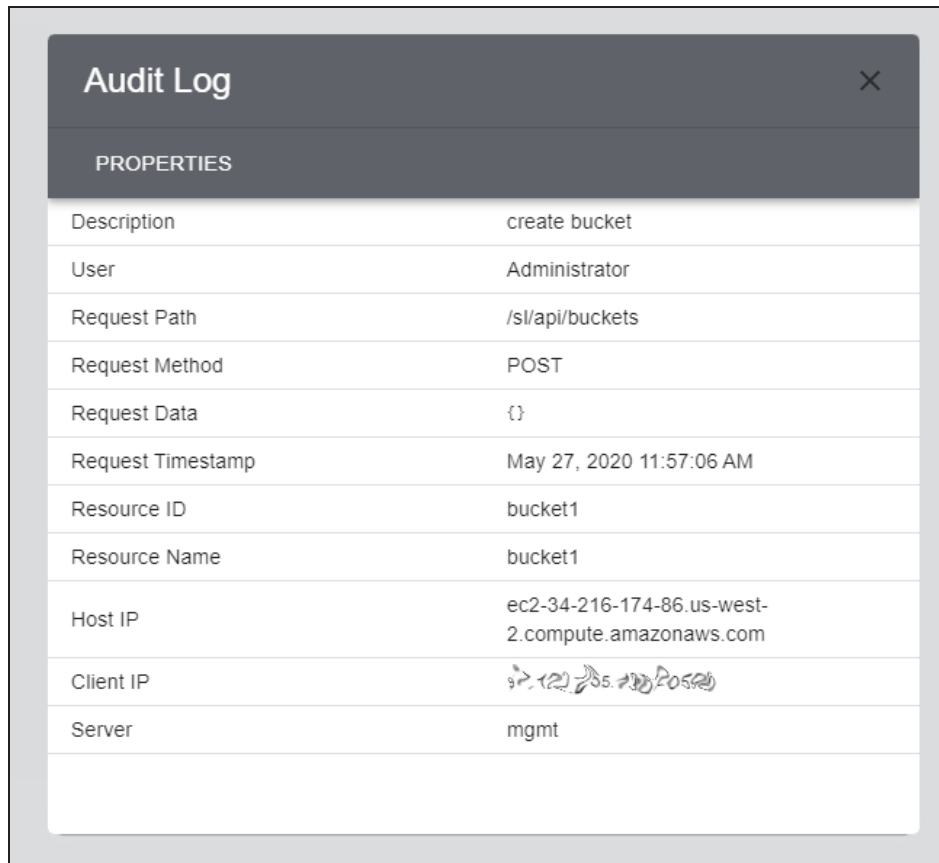


Figure 194 The Reports screen.

- 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.



Audit Log		×
PROPERTIES		
Description	create bucket	
User	Administrator	
Request Path	/sl/api/buckets	
Request Method	POST	
Request Data	{}	
Request Timestamp	May 27, 2020 11:57:06 AM	
Resource ID	bucket1	
Resource Name	bucket1	
Host IP	ec2-34-216-174-86.us-west-2.compute.amazonaws.com	
Client IP	127.0.0.1	
Server	mgmt	

Figure 195 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.

Option	Description
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 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 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 196 The Bell icon with unread messages.

The messages screen displays. Any unread messages are shown in bold font.

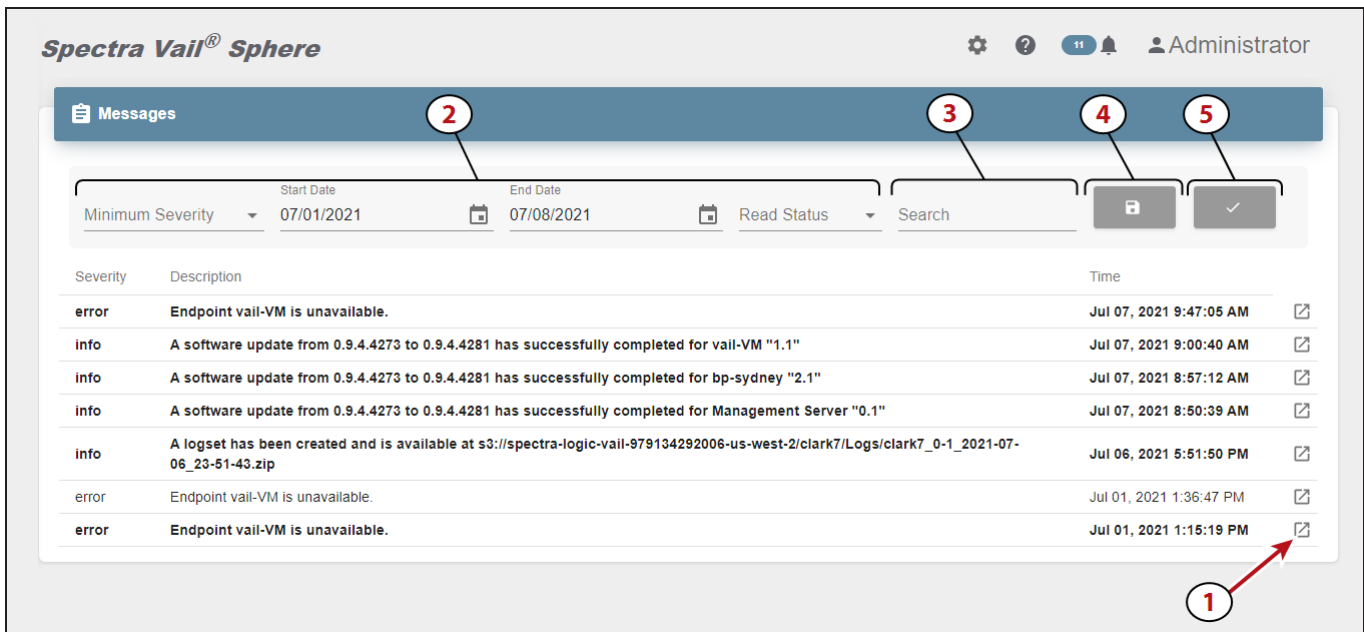


Figure 197 The Messages screen.

- 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 **checkmark icon** (5).

Message Details

In addition to the information on the Messages screen, the message details pane also displays the message key.

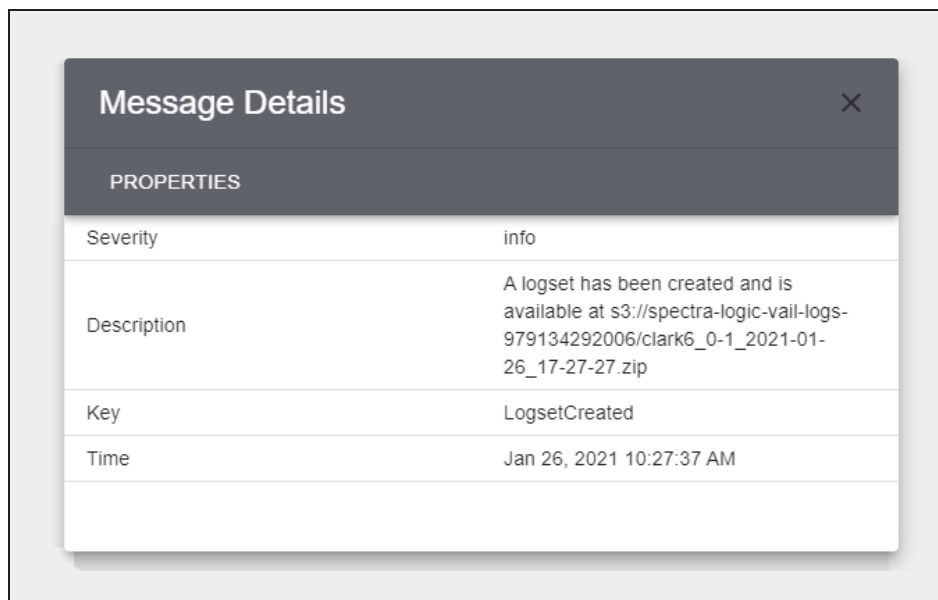


Figure 198 The Message Details screen.

Field	Description
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.

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**.

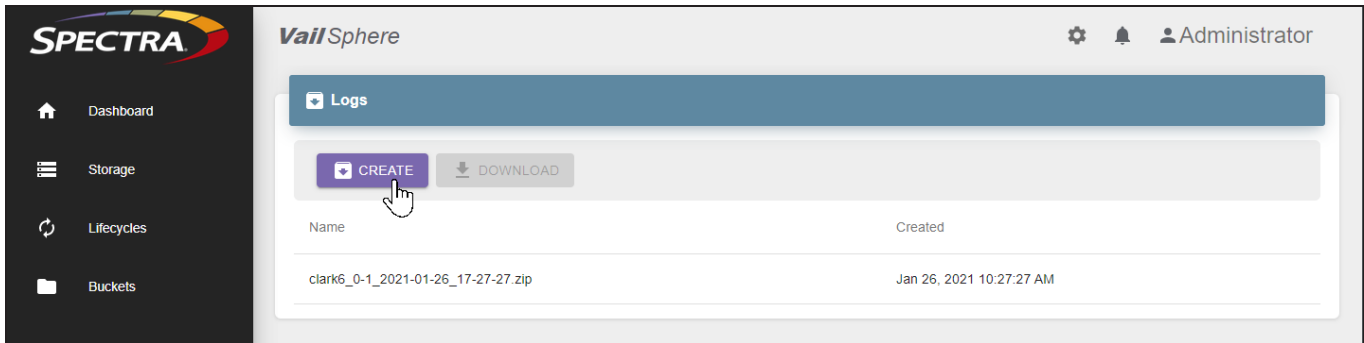


Figure 199 The Logs screen.

- 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 download an existing logset, select the row of the logset and click **Download**.

UPDATE THE VAIL APPLICATION SOFTWARE

Use the instructions below to update 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 BlackPearl S3 solution 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.

Update the Vail application software in the following order:

- BlackPearl software.
- Vail VM Node software.

Here is how to update the storage endpoint software:

1. In the upper right corner of the management console, click the **gear icon** and select **Software Updates**.

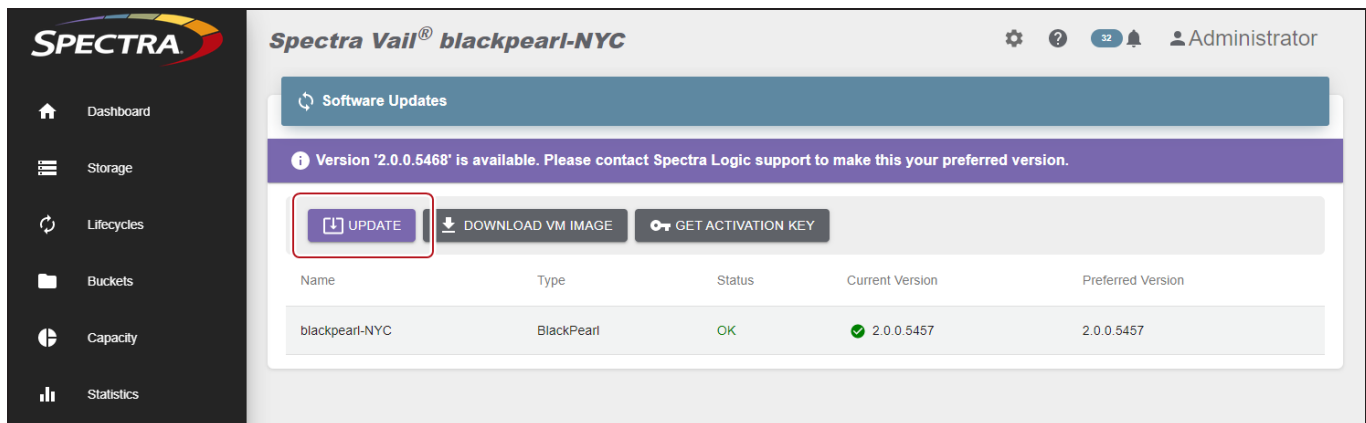


Figure 200 The Software Updates screen.

2. Select the row of the component you want to update and click **Update**.

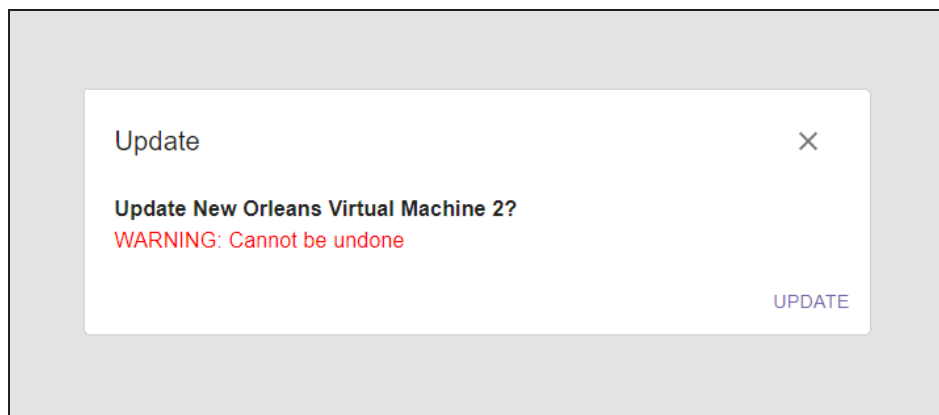


Figure 201 The Update confirmation screen.

3. Click **Update**. The update process for the selected component begins.



IMPORTANT

Do not reboot or power-cycle the BlackPearl S3 solution or Vail VM node VM during the update process or the BlackPearl S3 solution 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.

ENABLE DIAGNOSTIC MONITOR

If desired, you can enable the Spectra Vail application to send diagnostic information to Spectra Logic.

Here is how to enable the diagnostic monitor:

1. In the upper right corner of the management console, click the **gear icon** and select **Global Settings**.
2. Under the **Global Settings** banner, click **Edit**.
3. Select **Enable Monitor** and click **Submit**.

LOG OUT OF THE VAIL MANAGEMENT CONSOLE

To log out of the Vail management console, in the upper-right corner, click the name of the current user, and then click **Logout**.




Figure 202 The Logout function.

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 S3 solution, it initiates a Start Bulk Get job in the BlackPearl S3 solution. However, the Vail application has a back-door path to read objects from the BlackPearl cache. The BlackPearl S3 solution 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 S3 solution.

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 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.

GLOSSARY

BlackPearl System

A BlackPearl S3 solution 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 S3 solution, tape storage connected to a BlackPearl S3 solution, Vail VM node storage, and Vail clusters that share the same physical location.

BlackPearl Storage

The Spectra Vail application uses a BlackPearl S3 solution 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 S3 solution 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.