

IBM Spectrum Connect
Version 3.4.0

User Guide



Note

Before using this document and the product it supports, read the information in “Notices” on page 247.

Edition notice

Publication number: SC27-5999-24. This publication applies to version 3.4.0 of IBM Spectrum Connect and to all subsequent releases and modifications until otherwise indicated in a newer publication.

© **Copyright IBM Corporation 2013, 2018.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii
Tables	xi
About this guide	xiii
Who should use this guide	xiii
Conventions used in this guide	xiii
Related information and publications	xiii
Getting information, help, and service	xiv
IBM Publications Center	xiv
Sending or posting your comments	xiv
Chapter 1. IBM Spectrum Connect suite	1
Introduction	1
Included cloud interfaces	2
VMware virtual volumes	2
VMware Storage Policy Based Management (SPBM)	3
Storage space and service management	4
Management options	6
Installation	7
Compatibility and requirements	7
Downloading IBM Spectrum Connect software	7
Performing first-time installation of Spectrum Connect	8
Installing Spectrum Connect in the shared environment	11
Upgrading an existing installation	12
Uninstalling the Spectrum Connect software	17
Operation and management	18
Required and optional initial tasks	18
Configuring LDAP-based directory user access	21
Logging in	25
Web-based UI	27
Running initial setup	30
Managing high-availability groups	34
Managing server certificates	37
Managing Spectrum Connect users	40
Managing storage systems	43
Managing storage spaces and services	54
Administration	73
Checking and controlling the Spectrum Connect Linux service	74
Checking and modifying the configuration files	74
Changing the Spectrum Connect communication port	77
Management from the command-line interface	77
CLI – Switching to 'IBMSC' user mode	78
CLI – Managing Spectrum Connect users	78
CLI – Managing server certificates	80
CLI – Adding or removing storage system credentials	81
CLI – Managing storage systems	84
CLI – Setting the VASA credentials	87
CLI – Managing integration with vRealize Operations Manager	87
CLI – Backing up or restoring Spectrum Connect configuration	91
Troubleshooting IBM Spectrum Connect	92
Checking the log files	93
Checking the format of directory-based storage system credentials	95
Configuring event forwarding	96
Deleting virtual volumes and group pools via XCLI	97

Self-assist options for IBM Spectrum Connect	97
Chapter 2. Concept diagram for VMware environment	99
Chapter 3. IBM Storage Provider for VMware VASA	101
Managing and monitoring VASA access	101
Setting the VASA credentials	101
Managing VASA trusted certificates	103
Using the IBM Storage Provider for VMware VASA	103
Registering Spectrum Connect as a storage provider on vCenter server	103
Chapter 4. IBM Storage Enhancements for VMware vSphere Web Client	107
Managing integration with vSphere Web Client	107
Adding a vCenter server	107
Updating the credentials of a vCenter server	109
Removing a vCenter server.	110
Delegating storage services to a vCenter server	110
Canceling service delegation to a vCenter server.	112
Using the IBM Storage Enhancements for VMware vSphere Web Client	112
Required vSphere privileges	113
Viewing the IBM storage object information	114
Creating and mapping a new storage volume (LUN)	124
Extending a volume	131
Renaming a volume	132
Setting multipath policy enforcement for a volume	133
Unmapping a volume from one or more hosts	135
Deleting an unused volume	135
Managing consistency group attachment	136
Displaying the virtual volume information	138
Chapter 5. IBM Storage Plug-in for VMware vRealize Orchestrator	141
Managing integration with vRealize Orchestrator	141
Downloading and installing the plug-in package for vRO	142
Delegating storage services to the vRO server	146
Canceling service delegation to a vRO server.	147
Regenerating the vRO token	147
Using the IBM Storage Plug-in for VMware vRealize Orchestrator.	148
Chapter 6. IBM Storage Management Pack for VMware vRealize Operations Manager	151
Managing integration with vRealize Operations Manager.	151
Downloading the vROps management package	152
Deploying the management package on vROps	153
Connecting the vROps server to Spectrum Connect.	154
Controlling storage system monitoring on the vROps server.	155
Using the IBM Storage Management Pack for VMware vRealize Operations Manager	156
Overview dashboard	159
Performance dashboard	161
Top 10 dashboard	169
Displaying the overview of IBM storage objects	170
Defining thresholds and alerts for storage objects	171
Chapter 7. IBM Storage Enabler for Containers	177
Installing IBM Storage Enabler for Containers	179
Compatibility and requirements for IBM Storage Enabler for Containers.	179
Managing SSL certificates with IBM Storage Enabler for Containers	182
Downloading IBM Storage Enabler for Containers software	183
Performing installation of IBM Storage Enabler for Containers	183
Upgrading IBM Storage Enabler for Containers	188
Uninstalling the IBM Storage Enabler for Containers software	189

Managing integration with IBM Storage Enabler for Containers	189
Delegating storage services to the IBM Storage Enabler for Containers interface	191
Canceling service delegation to IBM Storage Enabler for Containers	192
Using the IBM Storage Enabler for Containers	192
Configuring storage classes, PVCs and pods	193
Sample configuration for running a stateful container	195
Recovering a crashed Kubernetes node	198
Updating the Enabler for Containers configuration files	200
Troubleshooting the IBM Storage Enabler for Containers	203
Chapter 8. IBM Storage Automation Plug-in for PowerShell	207
Managing integration with Microsoft PowerShell	208
Downloading and installing the plug-in package for PowerShell	210
Delegating storage services to the PowerShell interface	210
Canceling service delegation to PowerShell	211
Using the IBM Storage Automation Plug-in for PowerShell	212
Chapter 9. Best practices and use cases	223
Handling datastores	223
Handling ESXi hosts that use storage volumes	223
Distributing volumes evenly on DS8000 systems	223
Setting the multipath policy for DS8000 and Storwize Family systems	224
Working with multiple storage systems.	224
Upgrading or installing Spectrum Connect with vSphere failover	225
Creating a VVol-enabled service	225
Creating a VVol-enabled service on XIV storage systems	226
Creating a VVol-enabled service on storage systems that run IBM Spectrum Virtualize	227
Configuring an LDAP user for a managed domain	228
Restoring VVol-based virtual machines	234
Additional reference information	236
Chapter 10. RESTful API for Spectrum Accelerate software-defined storage	237
RESTful API protocol.	237
Query request and response	238
Create request and response	238
Delete request and response	239
Update request and response	239
Action request and response	240
Storage system operations	240
Module operations	241
Disk operations.	242
Interface operations	243
Port operations.	244
Emergency shutdown	245
Notices	247
Trademarks	248

Figures

1. Integration of IBM storage systems with different virtualization and automation environments	2
2. Using virtual volumes with IBM Spectrum Connect	3
3. Storage Policy Based Management (SPBM) concept.	4
4. Storage elements without VVol utilization.	5
5. Single service per space and storage system	5
6. Multiple services per storage space and single service per storage system	6
7. Multiple services per storage space and storage systems	6
8. Spectrum Connect version number.	16
9. Hyper-Scale Manager welcome page	26
10. Spectrum Connect login screen in a standard web browser.	27
11. Spaces/Storage Services and Storage Systems panes	28
12. Interfaces and Spaces/Services panes	28
13. Storage Systems and Monitoring panes	29
14. Initial setup wizard, defining HA group	31
15. Initial setup wizard, defining SSL certificate.	32
16. Initial setup wizard, defining storage system credentials	33
17. Initial setup wizard, defining Spectrum Connect credentials	34
18. High-availability group concept.	35
19. General Settings option on the Settings menu	35
20. General Settings dialog box	36
21. Connection security warning in the Mozilla FireFox web browser	38
22. Generate option on Server Certificate dialog box	39
23. Upload files option on Server Certificate dialog box	40
24. Users option in the Setting menu	41
25. Add option in the Users dialog box	42
26. Edit button in user account row.	43
27. Storage Systems pane	44
28. Settings button	45
29. Current storage system username (for all storage systems).	45
30. Add button	46
31. Add New IBM Storage System dialog box	47
32. Storage Systems pane, bar view	48
33. Storage Systems pane, table view	49
34. Storage Systems pane with storage elements arranged according to descending space usage	50
35. Storage Systems pane with highlighted search string results	51
36. Advanced filtering tool	51
37. Storage Systems pane with storage elements filtered according to disk size	52
38. Array Settings dialog box	53
39. Spaces/Storage Services pane	54
40. New Storage Space dialog box	55
41. List of storage spaces	56
42. New Storage Service dialog box.	57
43. Add New Resource dialog box	61
44. Storage Resources table	64
45. Service Storage Resources table	65
46. Add New Resource for replicated service dialog box, Create New option	66
47. Add New Resource for replicated service dialog box, Choose Existing option	67
48. Paired pools	69
49. Storage Resources table	70
50. Service Storage Resources table	71
51. Edit Resource dialog box	72
52. Edit Settings for replicated pools dialog box	72
53. Manage Resources dialog box	73
54. Connection security warning in the Mozilla FireFox web browser	80
55. Collect log option	94

56. XIV role mapping attributes for directory (LDAP) users	96
57. Integration of IBM storage systems with a VMware environment	99
58. VASA Provider settings dialog box	102
59. Registered VASA servers (vCenter servers that employ VASA services)	103
60. vSphere Web Client – Storage Providers list	104
61. New Storage Provider dialog box for VASA 2.0	105
62. vCenter certificate thumbprint dialog box	105
63. Storage Providers list displaying Spectrum Connect.	106
64. Add vCenter Server for vWC dialog box	108
65. Interfaces pane	109
66. vCenter Server Settings dialog box	110
67. vCenter server with delegated service	111
68. VMware vSphere Web Client – Create Role dialog box	114
69. IBM Storage categories in vSphere Web Client	115
70. IBM Storage Service information	116
71. IBM Storage Space information	117
72. IBM Storage Volume information	118
73. IBM Consistency Group information	118
74. IBM Storage VVol information	119
75. IBM Storage Service summary	120
76. IBM Storage Space summary	121
77. IBM Storage Volume summary.	122
78. IBM Consistency Group summary	123
79. IBM Storage VVol summary.	124
80. IBM storage service view – Clicking Create New Volume	125
81. Top Level Objects view – Clicking Create a new IBM Storage volume.	125
82. Create New Volume wizard.	126
83. Creating multiple volumes	127
84. Selecting storage service	127
85. Advanced Host Mapping dialog box.	128
86. Selecting LUN	129
87. Selecting consistency group	130
88. Naming a new consistency group	130
89. Create New Volume dialog box with all fields filled out	131
90. Clicking Extend on the pop-up menu	132
91. Extend Volume dialog box	132
92. Rename volume option	133
93. Rename Volume dialog box	133
94. Set Multipath Policy Enforcement option	134
95. Change Multipath Policy Enforcement dialog box	134
96. Unmap volume	135
97. Delete volume	136
98. Adding volume to consistency group	137
99. Selecting a consistency group	137
100. Removing volume from consistency group.	138
101. VVol Summary tab.	139
102. VVol Related Objects tab	139
103. vRO server on the Interfaces pane	142
104. Download plug-in package button	143
105. Current vRO Token	143
106. Successful installation of IBM Storage plug-in for vRO 6.x	144
107. Successful installation of IBM Storage plug-in for vRO 7.x	144
108. Restarting vRO Server service (vRO 6.x)	145
109. Restarting vRO Server service (vRO 7.x)	145
110. Start Workflow: Set Server and Token dialog box	146
111. vRO server with delegated services	147
112. vRealize Orchestrator – available workflows – General tab	149
113. vRealize Orchestrator – available objects	150
114. Monitoring pane	152
115. Download PAK File button	153
116. Deploying the management package on the vROps	154

117.	Adding the vROps server to Spectrum Connect	155
118.	Storage system monitored by the vROps server	156
119.	vROps GUI – IBM STORAGE option.	157
120.	IBM Spectrum Accelerate Storage Systems pane	159
121.	IBM Spectrum Accelerate volume health status	160
122.	Alert widget	161
123.	Virtual Machines and XIV and vCenter Relations panes	162
124.	HEALTH TREE pane	162
125.	METRIC SELECTOR and METRIC GRAPH panes of the Performance dashboard	163
126.	IBM XIV Top 10 dashboards	170
127.	Overview of an IBM FlashSystem A9000 flash canister	171
128.	Add Symptom Definition dialog box.	172
129.	Adding name and description for alert definition	172
130.	Selecting base object type for alert definition	173
131.	Selecting impact for alert definition	173
132.	Adding symptom for alert definition.	174
133.	Editing default monitoring policy.	175
134.	Enabling KPI for a storage object metrics	175
135.	Integration of IBM storage systems in Kubernetes environment	178
136.	IBM Storage Enabler for Containers post-installation status	187
137.	Add New Enabler for Containers Interface dialog box	190
138.	IBM Storage Enabler for Containers interface on the Interfaces pane	190
139.	Enabler for Containers interface with a delegated service	191
140.	Integration of IBM storage systems in PowerShell environment	208
141.	Add New PowerShell Interface dialog box	209
142.	PowerShell interface on the Interfaces pane	209
143.	PowerShell interface with a delegated service	211
144.	General tab, LDAP dialog box	229
145.	User Credentials tab, LDAP dialog box	230
146.	Role Mapping tab, LDAP dialog box.	230
147.	Group configuration on Active Directory server	232
148.	User configuration on Active Directory server.	233
149.	Storage Credentials dialog box.	234
150.	Defining vv01-db, as HA group name	235

Tables

1.	Configuration files renamed during Spectrum Connect installation	11
2.	Configuration files renamed during Spectrum Connect upgrade	17
3.	Required tasks in sequential order	18
4.	Optional tasks	19
5.	Arguments for sc_ldap	22
6.	ldap.ini configuration parameters.	24
7.	Spectrum Connect UI elements	29
8.	Service parameters	58
9.	Storage resource parameters	61
10.	Paired pool parameters	67
11.	Configuration files	75
12.	Arguments for sc_users	79
13.	User-related arguments for sc_setting	79
14.	Arguments for sc_ssl	81
15.	Arguments for sc_storage_credentials	83
16.	Arguments for sc_storage_array	85
17.	Storage resource-related arguments for sc_setting	86
18.	Arguments for sc_vasa_admin	87
19.	Arguments for sc_vrops_server	88
20.	Arguments for sc_vrops_adapter	90
21.	Arguments for sc_configuration	92
22.	VMware log file locations.	94
23.	Required vSphere privileges	113
24.	IBM storage objects and events in vRO	148
25.	IBM Storage Icons in vROps	157
26.	Capacity metrics	164
27.	Health metrics	164
28.	Counter metrics.	165
29.	Performance metrics	166
30.	Configuration parameters in ubiquity_installer.conf.	184
31.	Configuration parameters in storage-class-template.yml	193
32.	Configuration parameters in pvc-template.yml	194
33.	Configuration parameters in sanity-pod.yml	195
34.	Configuration parameters in ubiquity-configmap.yml	200
35.	Configuration parameters in scbe-credentials-secret.yml	202
36.	Configuration parameters in ubiquity-db-credentials-secret.yml	202
37.	Configuration parameters in storage-class.yml, ubiquity-db-pvc.yml, sanity-pvc.yml	203
38.	Troubleshooting for IBM Storage Enabler for Containers	205
39.	Cmdlets available via IBM Storage Automation Plug-in for PowerShell	212

About this guide

This guide describes how to install, configure, and use IBM Spectrum Connect and its solution components.

Who should use this guide

This guide is intended for system administrators who are familiar with the VMware vCenter and vSphere environments, Microsoft PowerShell, container technology, and with the specific IBM® storage system that is in use.

Conventions used in this guide

These notices are used in this guide to highlight key information.

Note: These notices provide important tips, guidance, or advice.

Important: These notices provide information or advice that might help you avoid inconvenient or difficult situations.

Attention: These notices indicate possible damage to programs, devices, or data. An attention notice appears before the instruction or situation in which damage can occur.

Related information and publications

You can find additional information and publications related to IBM Spectrum Connect on the following information sources.

- IBM Knowledge Center(ibm.com/support/knowledgecenter)
- IBM DS8000® on IBM Knowledge Center (ibm.com/support/knowledgecenter/STUVMB)
- IBM DS8800 on IBM Knowledge Center (ibm.com/support/knowledgecenter/STXN8P)
- IBM DS8870 on IBM Knowledge Center (ibm.com/support/knowledgecenter/ST8NCA)
- IBM FlashSystem 900 on IBM Knowledge Center(ibm.com/support/knowledgecenter/STKMQB)
- IBM SAN Volume Controller on IBM Knowledge Center (ibm.com/support/knowledgecenter/STPVGU)
- IBM Storwize® V3500 on IBM Knowledge Center(ibm.com/support/knowledgecenter/STLM6B)
- IBM Storwize V3700 on IBM Knowledge Center(ibm.com/support/knowledgecenter/STLM5A)
- IBM Storwize V5000 on IBM Knowledge Center(ibm.com/support/knowledgecenter/STHGJ)
- IBM Storwize V7000 on IBM Knowledge Center(ibm.com/support/knowledgecenter/ST3FR7)

- IBM Storwize V7000 Unified on IBM Knowledge Center(ibm.com/support/knowledgecenter/ST5Q4U)
- IBM XIV® Storage System on IBM Knowledge Center(ibm.com/support/knowledgecenter/STJTAG)
- IBM Spectrum Accelerate on IBM Knowledge Center (ibm.com/support/knowledgecenter/STZSWD)
- IBM FlashSystem® A9000 on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKMM)
- IBM FlashSystem A9000R on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKN5)
- VMware Documentation (vmware.com/support/pubs)
- VMware Product Support (vmware.com/support)
- VMware Knowledge Base (kb.vmware.com)
- Microsoft PowerShell (msdn.microsoft.com/en-us/powershell)
- Persistent volumes on Kubernetes (kubernetes.io/docs/concepts/storage/volumes)

Getting information, help, and service

If you need help, service, technical assistance, or want more information about IBM products, you can find various sources to assist you. You can view the following websites to get information about IBM products and services and to find the latest technical information and support.

- IBM website (ibm.com)
- IBM Support Portal website (ibm.com/support/entry/portal/support?brandind=Hardware~System_Storage)
- IBM Directory of Worldwide Contacts website (ibm.com/planetwide)

Use the Directory of Worldwide Contacts to find the appropriate phone number for initiating voice call support. Select the Software option, when using voice response system.

When asked, provide your Internal Customer Number (ICN) and/or the serial number of the storage system that requires support. Your call will then be routed to the relevant support team, to whom you can provide the specifics of your problem.

IBM Publications Center

The IBM Publications Center is a worldwide central repository for IBM product publications and marketing material.

The IBM Publications Center website (ibm.com/shop/publications/order) offers customized search functions to help you find the publications that you need. You can view or download publications at no charge.

Sending or posting your comments


Your feedback is important in helping to provide the most accurate and highest quality information.

Procedure

To submit any comments about this guide:

- Go to IBM Spectrum Connect on IBM Knowledge Center (ibm.com/support/knowledgecenter/SS6JWS), drill down to the relevant page, and then click the **Feedback** link that is located at the bottom of the page.

By adding a comment, you accept our [IBM Knowledge Center Terms of Use](#). Your comments entered on this IBM Knowledge Center site do not represent the views or opinions of IBM. IBM, in its sole discretion, reserves the right to remove any comments from this site. IBM is not responsible for, and does not validate or confirm, the correctness or accuracy of any comments you post. IBM does not endorse any of your comments. All IBM comments are provided "AS IS" and are not warranted by IBM in any way.



Comments (0) [Add Comment](#)

No Comments

Contact Privacy Terms of use Accessibility **Feedback**

Feedback

The feedback form is displayed and you can use it to enter and submit your comments privately.

- You can post a public comment on the Knowledge Center page that you are viewing, by clicking **Add Comment**. For this option, you must first log in to IBM Knowledge Center with your IBM ID.
- You can send your comments by email to starpubs@us.ibm.com. Be sure to include the following information:
 - Exact publication title and product version
 - Publication form number (for example: SC01-0001-01)
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

Note: When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Chapter 1. IBM Spectrum Connect suite

IBM Spectrum Connect is a centralized cloud integration system that consolidates a range of IBM storage provisioning, virtualization, cloud, automation, and monitoring solutions through a unified server platform.

See the following topics for detailed description of the IBM Spectrum Connect suite:

- “Introduction”
- “Installation” on page 7
- “Operation and management” on page 18
- “Installing Spectrum Connect in the shared environment” on page 11
- “Upgrading an existing installation” on page 12

Introduction

IBM Spectrum Connect provides a single-server backend location and enables centralized management of IBM storage resources for different virtualization, cloud and container platforms.

By implementing advanced storage provisioning techniques, IBM Spectrum Connect allows storage and hybrid cloud administrators to supply and integrate IBM storage with a range of VMware solutions, Kubernetes container clusters, and Microsoft PowerShell automation methods. IBM Spectrum Connect provides a web-based UI that makes this entire administration easy and straightforward, and saves a significant amount of time in setting up, connecting, and integrating the required storage resources into your cloud environment. To enable this wide range of functional integration possibilities, the IBM Spectrum Connect suite includes the following individual cloud interfaces

- VMware vCenter Server
- VMware vSphere Web Client (vWC)
- VMware vSphere APIs for Storage Awareness (VASA)
- VMware vRealize Operations Manager (vROps)
- VMware vRealize Orchestrator (vRO)
- Kubernetes
- Microsoft PowerShell

Through its user credential, storage system, storage space and service management options, Spectrum Connect facilitates the integration of IBM storage system resources with the supported virtualization, cloud and container platforms, while providing the foundation for integration with future IBM systems and Independent Software Vendor (ISV) solutions, as illustrated below. Storage profiles (services), defined in Spectrum Connect, are delegated for use in VMware or Kubernetes environment for simplified profile-based volume provisioning.

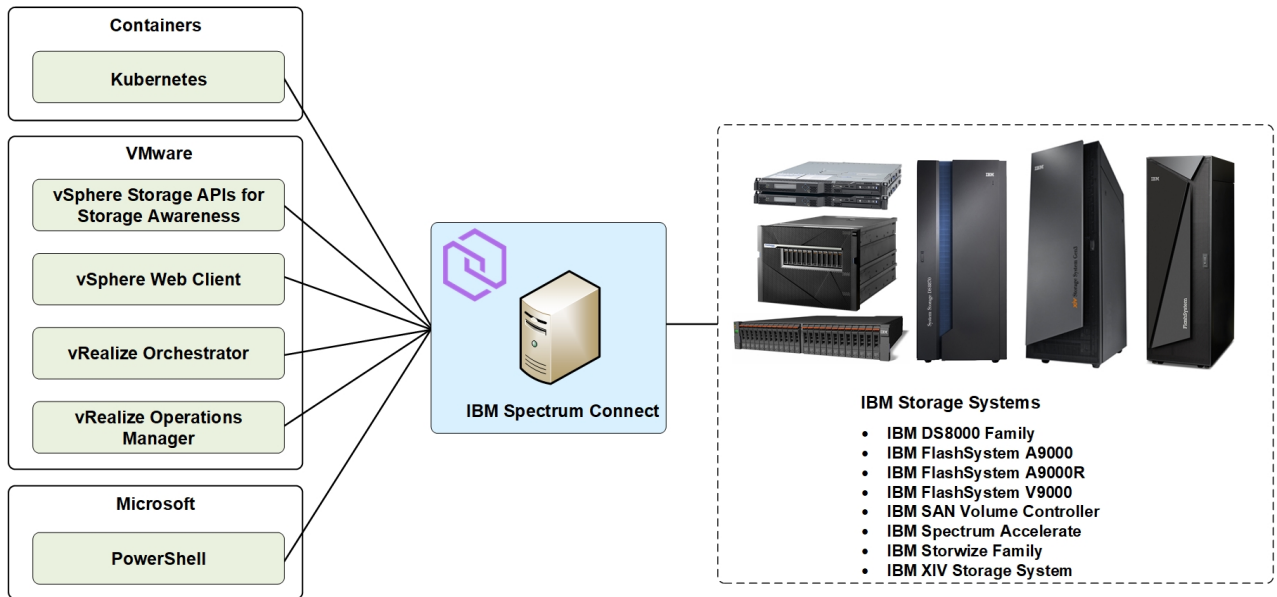


Figure 1. Integration of IBM storage systems with different virtualization and automation environments

Included cloud interfaces

The following cloud interfaces are included in the IBM Spectrum Connect software package:

- IBM Storage Provider for VMware VASA
- IBM Storage Enhancements for VMware vSphere Web Client
- IBM Storage Plug-in for VMware vRealize Orchestrator
- IBM Storage Management Pack for VMware vRealize Operations Manager
- IBM Storage Enabler for Containers
- IBM Storage Automation Plug-in for PowerShell

In addition, the IBM Spectrum Connect lifecycle and compatibility matrix on IBM Knowledge Center details the Spectrum Connect lifecycle with compatible storage system microcodes and supported cloud interfaces.

VMware virtual volumes

IBM Spectrum Connect delivers comprehensive storage virtualization support that use VMware vSphere Virtual Volume (VVOL) technology.

Note: The virtual volume technology is supported only by the IBM XIV (11.5.1 or later) and storage systems that run IBM Spectrum Virtualize™ (7.6 or later).

The VVol architecture, introduced in VMware VASA 2.0, preserves the concept of a traditional datastore, maintaining familiarity and compatibility with previous data storage implementations. With VVol, the IBM storage systems become aware of individual VMs, allowing data operations, such as snapshot and replication, to be performed directly by the storage system at the VM level.

The storage system uses VASA provider to present VVols to the ESXi host and inform the vCenter of the availability of VVol-aware storage. Storage services are

configured on VASA provider by the storage administrator and are used to manage storage resources (pools) and VVols. The services represent storage capacity with a set of attributes, such as encryption or provisioning type.

VVol usage improves system scalability, ensures granular management, leverages hardware features and performance of storage systems at the VM level, providing complete end-to-end cloud solution. An additional entity, a storage space, includes one or more services, and can be assigned to different storage customers.

As illustrated below, the IBM Storage Provider for VMware VASA implements the VMware Virtual Volume API, providing an out-of-band management bridge between vSphere and the storage system. Out-of-band link separates the management path from the data path, which connects the ESXi servers to the virtual disks in a VVol datastore through a Protocol Endpoint (PE). Instead of presenting a LUN to the hypervisor and allowing an ESXi host to perform data operations, a storage system takes on itself a bulk of storage-related functions.

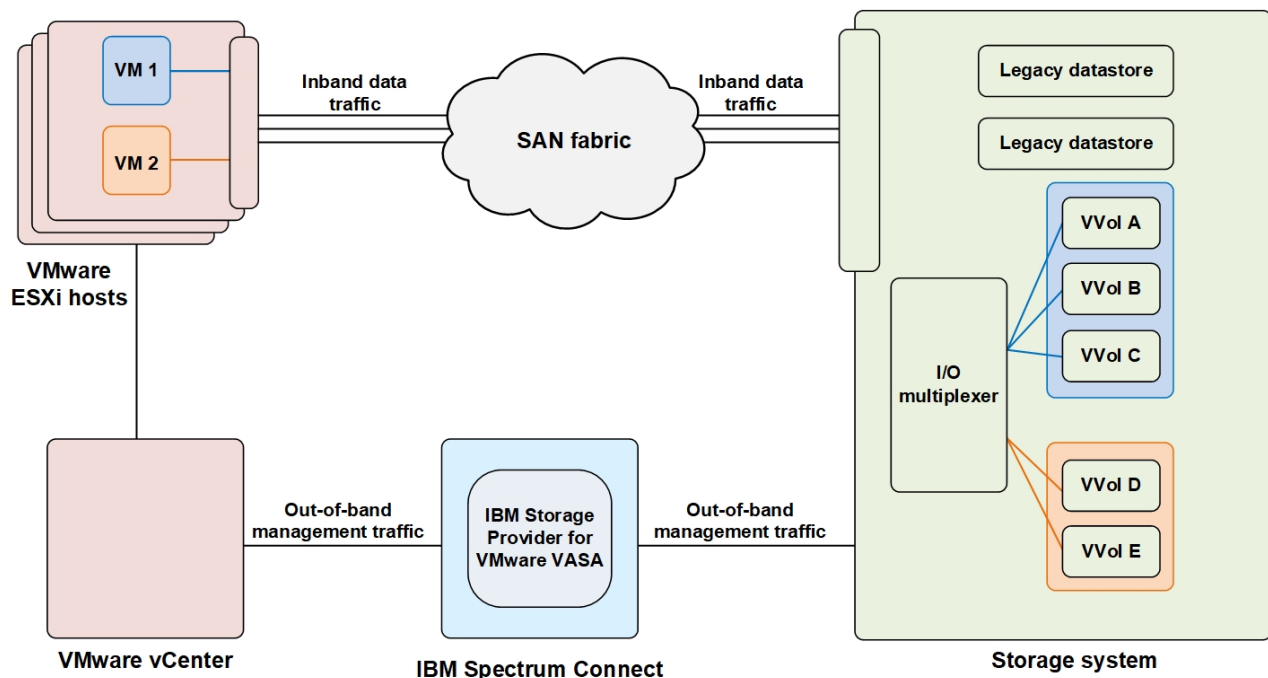


Figure 2. Using virtual volumes with IBM Spectrum Connect

For instructions about how to configure a VVol-enabled storage service, see “Creating a VVol-enabled service” on page 225.

VMware Storage Policy Based Management (SPBM)

IBM Spectrum Connect uses VMware vSphere Storage Policy Based Management (SPBM) technology for optimizing the virtual machine provisioning process.

Delivering only one service level, the traditional storage provisioning models fail to match storage consumer needs with storage provider capabilities. This results in misalignment between the system capacities and application requirements, leading to over-provisioning and waste of IT resources.

The SPBM approach allows dynamic definition of storage policies with their subsequent delivery on a per-VM basis. In this case, the storage consumer can pair

an application with existing storage policies and provision storage resources exactly according to application requirements. Storage policies, referred to as storage services in this user guide, combine storage capacity with a set of attributes (encryption, provisioning type, etc.) to define storage spaces, which are used as virtual datastores to suit requirements of a specific VM, as illustrated below.

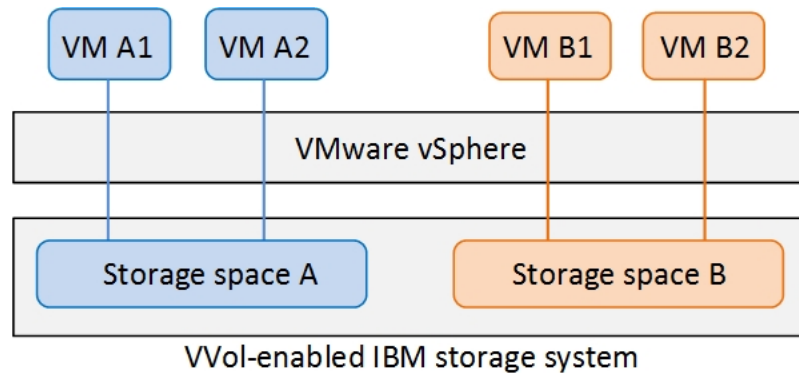


Figure 3. Storage Policy Based Management (SPBM) concept

Note: The virtual volume functionality is supported only by the IBM XIV (11.5.1 or later) and storage systems that run IBM Spectrum Virtualize (7.6 or later).

Storage space and service management

After deployment and storage system attachment, the IBM Spectrum Connect administrators must define the new virtual entities, resulting in simpler and more flexible storage management.

Note: The virtual volume functionality is supported by the IBM XIV (11.5.1 or later) and storage systems that run IBM Spectrum Virtualize (7.6 or later).

The virtual storage entities include:

- **Storage service** – A combination of assigned storage resources (pools) and user-defined policies (capabilities). The storage resources which are assigned to the service may reside on any storage system, as illustrated below. The policies are additional capabilities, or storage requirements for the service. They are compression, encryption, etc.
- **Storage space** – A logical grouping of several storage services. Usually, a single space is assigned to a specific organization (storage tenant).




Figure 4. Storage elements without VVol utilization

When the use of VMware virtual volumes (VVols) is enabled for a service, its storage resources contain XIV group pools or child pools on storage systems that run IBM Spectrum Virtualize. For XIV storage systems, a storage resource consists of the following pools:

- Thin pool for thin provisioning.
- Thick pool for thick provisioning.
- Meta pool for holding VM-related management metadata.

Figures below illustrate different deployment scenarios for VVol-enabled services.




Figure 5. Single service per space and storage system




Figure 6. Multiple services per storage space and single service per storage system




Figure 7. Multiple services per storage space and storage systems

The combination of storage space and service is created by storage administrators to include the required resource capacity and storage capabilities on a matching storage system. Then the space/service combination is used by VMware administrators for volume provisioning instead of physical objects.

Management options

IBM Spectrum Connect can be managed using the following methods:

- Graphical user interface (GUI).
- Command line interface (CLI).

Web-based user interface (UI)

IBM Spectrum Connect includes a simple user-friendly web-based graphical user interface (UI) for storage management.

The Spectrum Connect web UI simplifies storage provisioning, delivering a single control instance for all available resources. The UI has the following management capabilities:

- Addition and management of the physical storage resources (storage systems and pools)
- Creation and configuration of virtual storage entities (spaces and services)
- Integration with cloud interfaces (VMware VASA, vWC, vRO, vROps, PowerShell, Kubernetes)
- User administration
- Certificate management

Command-line interface (CLI)

IBM Spectrum Connect can be managed via a command-line interface (CLI).

The Spectrum Connect CLI is used for user and storage system management, as well as for integration of the cloud interfaces. However, the CLI application scope is limited, and the use of the web-based user interface (UI) is advised.

The CLI tool is supplied as a part of the Spectrum Connect package. It can be run locally from the Linux command prompt environment, or from a remote terminal connection.

Installation

Download and install IBM Spectrum Connect software package as described in the following sections.

- “Compatibility and requirements”
- “Downloading IBM Spectrum Connect software”
- “Performing first-time installation of Spectrum Connect” on page 8
- “Installing Spectrum Connect in the shared environment” on page 11
- “Upgrading an existing installation” on page 12

For information about uninstallation, see “Uninstalling the Spectrum Connect software” on page 17.

For information about installation of IBM Storage Enabler for Containers, see “Installing IBM Storage Enabler for Containers” on page 179.

Compatibility and requirements

For the complete and up-to-date information about the compatibility and requirements of IBM Spectrum Connect, refer to its latest release notes.

You can find the latest Spectrum Connect release notes on IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SS6JWS) or on IBM Fix Central (www.ibm.com/support/fixcentral).

Note: Refer to the relevant VMware documentation for information about how to install the compatible versions of vSphere Web Client Server. You should also refer to the latest installation and configuration instructions for ESXi and vCenter servers.

Downloading IBM Spectrum Connect software

IBM Spectrum Connect is available as a free software solution for IBM storage system customers.

About this task

You can download the latest version of the Spectrum Connect at any time from the IBM Fix Central (www.ibm.com/support/fixcentral). Fix Central provides fixes and updates for your systems software, hardware, and operating system. This procedure describes how to locate the Spectrum Connect package on the website.

Procedure

To download the Spectrum Connect:

1. Go to the Spectrum Connect welcome page on IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SS6JWS).
2. Click **Download the latest Spectrum Connect package from IBM Fix Central**. The IBM Fix Central page for Spectrum Connect is displayed.
3. Download the required software version.

Performing first-time installation of Spectrum Connect

You can install the IBM Spectrum Connect software on a compatible version of Red Hat Enterprise Linux (RHEL) or CentOS. For more information, refer to the release notes.

Before you begin

- Verify that the following TCP ports are open to ensure network connectivity between VMware resources, Spectrum Connect and IBM storage systems:
 - 8440 (vCenter, vROps and vRO servers). Refer to the installation procedure below for details on opening the 8440 port for the 'iptables' (RHEL 6.x) or 'firewall-cmd' (RHEL 7.x, CentOS 7.x). In addition, Spectrum Connect must be allowed to bind to port 8440, if Security-Enhanced Linux (SELinux) is enabled.
 - 7778 (XIV, Spectrum Accelerate, FlashSystem A9000/A9000R).
 - 22 (storage systems that run IBM Spectrum Virtualize).
 - 8452 (DS8000).

Note: You can change the default TCP port (8440) at any time by running a script, as explained in “Changing the Spectrum Connect communication port” on page 77.

- Verify that the TCP ports 5672 and 4369 are open. These ports are used by the 'rabbitmq' and 'ampq' internal processes, respectively. Refer to the installation procedure below for the port opening procedures.
- Make sure that the 'ibmsc' user can access the /opt/ibm and /var/log/sc folders.
- Check that the 'policycoreutils-python' package has been installed for RHEL 6.x.
- Check that the 'zlib' library has been installed for RHEL 7.x and CentOS 7.x.
- Check that the 'bzip2' program has been installed for all operating systems.
- Verify that the 'postgresql' package is not installed on your host. Your server may have a package of 'postgresql' version 8 installed, as a part of operation system distribution. This may result in a conflict with version of the package, installed during the Spectrum Connect deployment. Use the `>> rpm -qa | grep postgres` command to search for the 'postgresql' package.
- A new Linux username – **ibmsc** – is created during installation to be used for the Spectrum Connect management operations. You can customize the user ID for

ibmsc by adding a Linux user (**useradd** command in RHEL) prior to the package installation. In this case, create the `/home/ibmsc` directory before starting the installation process.

Procedure

Follow these steps to install Spectrum Connect:

1. Open the 8440 TCP port:
 - RHEL 6.x:
 - **iptables -I OUTPUT -p tcp --dport 8440 -j ACCEPT**
 - **iptables -I INPUT -p tcp --dport 8440 -j ACCEPT**
 - **service iptables save**
 - **service iptables restart**
 - RHEL 7.x, CentOS 7.x:
 - **firewall-cmd --permanent --add-port=8440/tcp**
 - **firewall-cmd --reload**
2. If you are using SELinux, allow Spectrum Connect to bind to the 8440 TCP port:
 - RHEL 6.x: **semanage port -a -t http_port_t -p tcp 8440**
 - RHEL 7.x, CentOS 7.x: **setsebool -P nis_enabled 1**
3. Open the 5672 and 4369 TCP ports under RHEL 7.x and CentOS 7.x:
 - **firewall-cmd --permanent --zone=trusted --add-interface=lo**
 - **firewall-cmd --permanent --zone=trusted --add-port=5672/tcp**
 - **firewall-cmd --permanent --zone=trusted --add-port=4369/tcp**
4. Download the installation package and the `IBM_Spectrum_Connect_Signing_Key_Pub.key` file, used for the package validation. See “Downloading IBM Spectrum Connect software” on page 7).
5. Copy the installation package and the public key files to a local folder on the Linux host that will be used as Spectrum Connect.
6. Go to the local folder and then use the **gpg --import IBM_Spectrum_Connect_Signing_Key_Pub.key** to import the IBM GNU Privacy Guard (GPG) public key to validate the installation files. This ensures that the files were received from IBM and were not manipulated in any way by a third party.

Note: Downloading the install package from a trusted, SSL-protected resource, such as Fix Central, ensures its authenticity and integrity. However, you can mark the key as trusted by entering **gpg --edit-key "IBM Spectrum Connect Signing Key"**, typing the **trust** command and selecting option 5.

7. Extract the installation package file ('*' represents the build number) :

```
# tar -xzvf IBM_Spectrum_Connect-3.4.0-*-x86_64.tar.gz
```

Depending on the operating system type, the following files are extracted:

- RHEL 6.x:
 - `erlang-18.1-1.el6.x86_64.rpm`
 - `nginx-1.8.0-1.el6ngx.x86_64.rpm`
 - `postgresql92-9.2.14-1PGDG.rhel6.x86_64.rpm`

- postgresql92-contrib-9.2.14-1PGDG.rhel6.x86_64.rpm
 - postgresql92-libs-9.2.14-1PGDG.rhel6.x86_64.rpm
 - postgresql92-server-9.2.14-1PGDG.rhel6.x86_64.rpm
 - rabbitmq-server-3.6.0-1.noarch.rpm
 - uuid-1.6.1-10.el6.x86_64.rpm
 - RHEL 7.x:
 - erlang-18.1-1.el7.centos.x86_64.rpm
 - nginx-1.8.0-1.el7ngx.x86_64.rpm
 - postgresql92-9.2.14-1PGDG.rhel7.x86_64.rpm
 - postgresql92-contrib-9.2.14-1PGDG.rhel7.x86_64.rpm
 - postgresql92-libs-9.2.14-1PGDG.rhel7.x86_64.rpm
 - postgresql92-server-9.2.14-1PGDG.rhel7.x86_64.rpm
 - rabbitmq-server-3.6.0-1.noarch.rpm
 - uuid-1.6.2-26.el7.x86_64.rpm
 - CentOS 7.x:
 - erlang-19.3.0-1.el7.centos.x86_64.rpm
 - jemalloc-3.6.0-1.el7.x86_64.rpm
 - libxslt-1.1.28-5.el7.x86_64.rpm
 - nginx-1.10.3-1.el7ngx.x86_64.rpm
 - postgresql92-9.2.24-5PGDG.rhel7.x86_64.rpm
 - postgresql92-contrib-9.2.24-5PGDG.rhel7.x86_64.rpm
 - postgresql92-docs-9.2.24-5PGDG.rhel7.x86_64.rpm
 - postgresql92-libs-9.2.24-5PGDG.rhel7.x86_64.rpm
 - postgresql92-server-9.2.24-5PGDG.rhel7.x86_64.rpm
 - rabbitmq-server-3.6.0-1.noarch.rpm
 - redis-3.2.5-1.el7.x86_64.rpm
 - uuid-1.6.2-26.el7.x86_64.rpm
 - `ibm_spectrum_connect-3.4.0-*.bin` – product BIN file.
 - `ibm_spectrum_connect-3.4.0-xxxx-x86_64.bin.asc`– digital signature file for the BIN file verification.
8. Enter `# gpg --verify ibm_spectrum_connect-3.4.0-xxxx-x86_64.bin.asc ibm_spectrum_connect-3.4.0-xxxx-x86_64.bin` to verify the digital signature of the installation files.
 9. Go to the extracted directory and then use the `rpm -iv *.rpm` command to run and install all the complementary RPM files. The IBM Storage Provider service starts automatically after the installation (for more information, see “Checking and controlling the Spectrum Connect Linux service” on page 74) and a new Linux username – `ibmsc` – is created so that you can use it for the Spectrum Connect management operations.
 10. Enter `chmod +x ibm_spectrum_connect-3.4.0-*.bin` to authorize the installation of the product BIN file.
 11. Enter `./ibm_spectrum_connect-3.4.0-*.bin` to start the installation.
 12. Review the license agreement which is displayed after you run the installation file.
 13. Enter 1 to accept the license agreement and complete the installation:

Press Enter to continue viewing the license agreement, or enter "1" to accept the agreement, "2" to decline it, "3" to print it, "4" to read non-IBM terms, or "99" to go back.

1

```

Preparing for new installation...
Creating system user ibmsc...
Configuring rsyslog...
Setting up nginx...
Generating SSL certificate...
Configuring postgresql database...
Creating IBM Spectrum Connect user...
Configuring IBM Spectrum Connect...
Adding iptables rules...
Starting Celery services... [OK]
Starting Django service... [OK]
NOTE: An initial username 'admin' with an initial password 'admin!'
has been defined for the initial access (via the CLI or GUI) to the IBM Spectrum Connect.
IMPORTANT: To avoid unauthorized access to the IBM Spectrum Connect,
the password for this username should be changed as soon as possible.
You can control IBM Spectrum Connect services using the
'service ibm_spectrum_connect {start|stop|status}' command.

Installation completed successfully.

```

14. To avoid unauthorized access to Spectrum Connect via GUI, it is strongly recommended to change the default password for the 'admin' user as soon as possible, as described in “Changing the password of a Spectrum Connect user” on page 42.

Note: During installation, IBM Spectrum Connect overwrites several configuration files. The previous versions of these files are renamed and kept in their original locations. A backup version is renamed as `file_name.pre_previous_SCB_version_number`. For example, `sc_nginx.conf` from Spectrum Control Base (3.0.3) installation is renamed as `sc_nginx.conf.pre_3.0.3`. Table below specifies the files renamed during installation.

Table 1. Configuration files renamed during Spectrum Connect installation

File name	Location
40-ibmsyslog.conf	/etc/rsyslog.d/
500.html	/etc/nginx/conf.d/
sc_fastcgi_params	/etc/nginx/conf.d/
sc_nginx.conf	/etc/nginx/conf.d/

Installing Spectrum Connect in the shared environment

You can install the Spectrum Connect software as a part of Hyper-Scale Manager, a standalone application for storage infrastructure management.

Before you begin

- Verify that the following communication ports are open to ensure network connectivity between VMware resources, Spectrum Connect and IBM storage systems:
 - 8440 (vCenter, vROps and vRO servers). Refer to the installation procedure below for details on opening the 8440 port for the 'iptables'(RHEL 6.x) or

'firewall-cmd' (RHEL 7.x, CentOS 7.x). In addition, Spectrum Connect must be allowed to bind to port 8440, if Security-Enhanced Linux (SELinux) is enabled.

- 7778 (XIV, Spectrum Accelerate, FlashSystem A9000/A9000R).
- 22 (storage systems that run IBM Spectrum Virtualize).
- 8452 (DS8000).

Note: You can change the default TCP port (8440) at any time by running a script, as explained in “Changing the Spectrum Connect communication port” on page 77.

- Verify that the TCP ports 5672 and 4369 are open. These ports are used by the 'rabbitmq' and 'ampq' internal processes, respectively. Refer to the installation procedure below for the port opening procedures.
- Make sure that the 'ibmsc' user can access the /opt/ibm and /var/log/sc folders.
- Check that the 'policycoreutils-python' package has been installed for RHEL 6.x.
- Check that the 'zlib' library has been installed for RHEL 7.x and CentOS 7.x.
- Check that the 'bzip2' program has been installed for all operating systems.
- Verify that the 'postgresql' package is not installed on your host. Your server may have a package of 'postgresql' version 8 installed, as a part of operation system distribution. This may result in a conflict with version of the package, installed during the Spectrum Connect deployment. Use the `>> rpm -qa | grep postgres` command to search for the 'postgresql' package.
- A new Linux username – **ibmsc** – is created during installation to be used for the Spectrum Connect management operations. You can customize the user ID for **ibmsc** by adding a Linux user (**useradd** command in RHEL) prior to the package installation. In this case, create the /home/ibmsc directory before starting the installation process.

Procedure

Download and install Hyper-Scale Manager, according to instructions in the product documentation. Spectrum Connect is installed automatically together with the other elements of the Hyper-Scale Manager suite.

Important: Adding storage systems function is disabled for Spectrum Connect instance, when it is installed using Hyper-Scale Manager. Storage systems can be added only using the Hyper-Scale Manager UI.

Upgrading an existing installation

If you are already using earlier releases of IBM Spectrum Connect, you can upgrade to the newer version without having to uninstall the previous one.

Before you begin

- If needed, back up the current Integration Server, Spectrum Control Base or Spectrum Connect database, by entering one of the following commands. Backup and restore procedure is allowed within the same version only. This means that it is not possible to back up files from Spectrum Control Base version 2.2.1 and restore them in version 3.0.3.
 - **isis_configuration backup -f /var/tmp/backup -k <key value>** for IBM Storage Integration Server.

- **sc_configuration backup -f /var/tmp/backup -k <key value>** for IBM Spectrum Connect.
- Verify that the following communication ports are open to ensure network connectivity between VMware resources, Spectrum Connect and IBM storage systems:
 - 8440 (vCenter, vROps and vRO servers)
 - 7778 (XIV, Spectrum Accelerate, FlashSystem A9000/A9000R)
 - 22 (storage systems that run IBM Spectrum Virtualize)
 - 8452 (DS8000)

Note: The 8443 port remains open after the Spectrum Connect upgrade, if the previous version was configured to use this port. If you install the software package anew, the 8440 port is used. You can change the default TCP port at any time by running a script, as explained in “Changing the Spectrum Connect communication port” on page 77. Moreover, certain upgrade scenarios, such as using SELinux, may require additional configuration steps, as detailed in “Performing first-time installation of Spectrum Connect” on page 8.

- Make sure that the 'ibmsc' user can access the /opt/ibm and /var/log/sc folders.
- Verify version of your current Spectrum Connect installation. Spectrum Control Base released before version 3.0.0 cannot be upgraded to version 3.1.0 or later. This includes the IBM Storage Integration Server releases. You must upgrade the earlier release to version 3.0.0, restart vSphere Web Client, then upgrade it to version 3.1.0 or later.

Procedure

Perform the following procedure to upgrade Spectrum Connect:

1. Log out of the Spectrum Connect GUI and close the browser.
2. On the Spectrum Connect side: download the newer installation package and the IBM_Spectrum_Connect_Signing_Key_Pub.key file, used for the package validation. See “Downloading IBM Spectrum Connect software” on page 7).
3. Copy the installation package and the public key files to a local folder on a current Spectrum Connect server.
4. Go to the local folder and then use the **gpg --import IBM_Spectrum_Connect_Signing_Key_Pub.key** to import the IBM GNU Privacy Guard (GPG) public key to validate the installation files. This ensures that the files were received from IBM and were not manipulated in any way by a third party.

Note: Downloading the install package from a trusted, SSL-protected resource, such as Fix Central, ensures its authenticity and integrity. However, you can mark the key as trusted by entering **gpg --edit-key "IBM Spectrum Connect Signing Key"**, typing the **trust** command and selecting option 5.

5. Extract the installation package file ('*' represents the build number) :

```
# tar -xzvf IBM_Spectrum_Connect-3.4.0-*x86_64.tar.gz
```

Depending on the operating system type, the following files are extracted:

- RHEL 6.x:
 - erlang-18.1-1.el6.x86_64.rpm

- nginx-1.8.0-1.el6.ngx.x86_64.rpm
 - postgresql92-9.2.14-1PGDG.rhel6.x86_64.rpm
 - postgresql92-contrib-9.2.14-1PGDG.rhel6.x86_64.rpm
 - postgresql92-libs-9.2.14-1PGDG.rhel6.x86_64.rpm
 - postgresql92-server-9.2.14-1PGDG.rhel6.x86_64.rpm
 - rabbitmq-server-3.6.0-1.noarch.rpm
 - uuid-1.6.1-10.el6.x86_64.rpm
 - RHEL 7.x:
 - erlang-18.1-1.el7.centos.x86_64.rpm
 - nginx-1.8.0-1.el7.ngx.x86_64.rpm
 - postgresql92-9.2.14-1PGDG.rhel7.x86_64.rpm
 - postgresql92-contrib-9.2.14-1PGDG.rhel7.x86_64.rpm
 - postgresql92-libs-9.2.14-1PGDG.rhel7.x86_64.rpm
 - postgresql92-server-9.2.14-1PGDG.rhel7.x86_64.rpm
 - rabbitmq-server-3.6.0-1.noarch.rpm
 - uuid-1.6.2-26.el7.x86_64.rpm
 - CentOS 7.x:
 - erlang-19.3.0-1.el7.centos.x86_64.rpm
 - jemalloc-3.6.0-1.el7.x86_64.rpm
 - libxslt-1.1.28-5.el7.x86_64.rpm
 - nginx-1.10.3-1.el7.ngx.x86_64.rpm
 - postgresql92-9.2.24-5PGDG.rhel7.x86_64.rpm
 - postgresql92-contrib-9.2.24-5PGDG.rhel7.x86_64.rpm
 - postgresql92-docs-9.2.24-5PGDG.rhel7.x86_64.rpm
 - postgresql92-libs-9.2.24-5PGDG.rhel7.x86_64.rpm
 - postgresql92-server-9.2.24-5PGDG.rhel7.x86_64.rpm
 - rabbitmq-server-3.6.0-1.noarch.rpm
 - redis-3.2.5-1.el7.x86_64.rpm
 - uuid-1.6.2-26.el7.x86_64.rpm
 - `ibm_spectrum_connect-3.4.0-*.bin` – product BIN file.
 - `ibm_spectrum_connect-3.4.0-xxxx-x86_64.bin.asc`– digital signature file for the BIN file verification.
6. Enter `# gpg --verify ibm_spectrum_connect-3.4.0-xxxx-x86_64.bin.asc ibm_spectrum_connect-3.4.0-xxxx-x86_64.bin` to verify the digital signature of the installation files.
 7. Go to the extracted directory and then use the `rpm -U *.rpm` command to run and install all the complementary RPM files.
 8. Enter `chmod +x ibm_spectrum_connect-3.4.0-*.bin` to authorize the installation of the product BIN file.
 9. Enter `./ibm_spectrum_connect-3.4.0-*.bin` to start the upgrade.

Note: During the upgrade several configuration files are overwritten by newer versions (as illustrated by the screen output below). If these files were changed in the previous versions, you must apply the same changes to the new files, if you want to preserve the settings.

10. Review the license agreement which is displayed after you run the installation file.
11. Enter 1 to accept the license agreement and complete the installation:

```
Press Enter to continue viewing the license agreement, or enter "1" to accept
the agreement, "2" to decline it, "3" to print it, "4" to read non-IBM terms,
or "99" to go back.

1

Preparing to upgrade [ibm_storage_integration_server] to the new [ibm_spectrum_connect]
Stopping service ibm_storage_integration_server...
Renaming old user isis to ibmsc
Renaming old group isis to ibmsc
Upgrading [ibm_storage_integration_server] to the new [ibm_spectrum_connect]
Moving rpm configuration files to IBM Spectrum Connect
  3 configuration files replaced by newer versions. The original files moved to:
    /opt/ibm/ibm_spectrum_connect/conf.d/ibmsyslog.conf.saverpm
    /opt/ibm/ibm_spectrum_connect/conf.d/nginx/sc_nginx.conf.saverpm
    /opt/ibm/ibm_spectrum_connect/conf.d/vasal/vasa_config.ini.saverpm
Moving Django key to IBM Spectrum Connect
Moving SSL Certificate to IBM Spectrum Connect
Configuring rsyslog
Setting up nginx
Migrating database to IBM Spectrum Connect
Configuring new service [ibm_spectrum_connect]
Update SC_UUID
Upgrading the extension of all registered vCenter servers...
Adding a default vCO server instance...
Starting Celery services ... [OK]
Starting Django service ... [OK]
Unconfiguring old service [ibm_storage_integration_server]

Installation completed successfully.
```

12. Log in (see “Logging in” on page 25) and click **Settings > About** to verify that the Spectrum Connect version number has been updated.

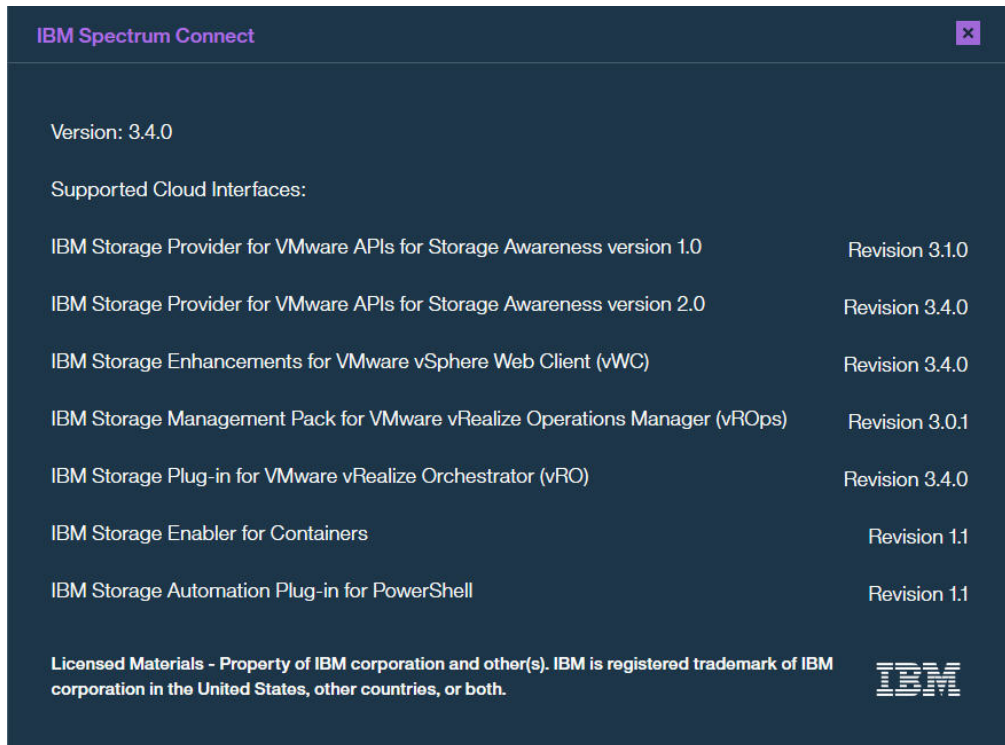


Figure 8. Spectrum Connect version number

13. To avoid unauthorized access to Spectrum Connect via GUI, it is strongly recommended to change the default password for the 'admin' user as soon as possible, as described in “Changing the password of a Spectrum Connect user” on page 42. After the upgrade, Spectrum Connect:
 - Adds a service for each existing storage resource (pool) that was connected to a vCenter or a vRO server. The new services are placed under default storage space. See “Managing storage spaces and services” on page 54 for details on how to configure storage spaces and services.
 - With the Spectrum Connect update, the IBM Storage Enhancements for VMware vSphere Web Client is upgraded automatically for all connected vCenter servers. However, the following components must be updated manually:
 - IBM Storage Plug-in for VMware vRealize Orchestrator. See “Managing integration with vRealize Orchestrator” on page 141.
 - IBM Storage Management Pack for VMware vRealize Operations Manager. See “Managing integration with vRealize Operations Manager” on page 151.
 - Registration of Spectrum Connect as a VASA storage provider on the vCenter server. Re-register Spectrum Connect, as detailed in “Registering Spectrum Connect as a storage provider on vCenter server” on page 103.

Note: During upgrade, Spectrum Connect overwrites several configuration files. The previous versions of these files are renamed and kept in their original locations. A backup version is renamed as file_name.pre_previous_SCB_version_number. For example, sc_nginx.conf from Spectrum Control Base (3.0.3) installation is renamed as sc_nginx.conf.pre_3.0.3. Table below specifies the files renamed during upgrade.

Table 2. Configuration files renamed during Spectrum Connect upgrade

File name	Location
40-ibmsyslog.conf	/etc/rsyslog.d/
500.html	/etc/nginx/conf.d/
sc_fastcgi_params	/etc/nginx/conf.d/
sc_nginx.conf	/etc/nginx/conf.d/

What to do next

To ensure proper transition from the physical to abstracted storage provisioning, perform the following:

- When upgrading from version 2.x to version 3.x, restart the vSphere Web Client via the vCenter administration panel, after the upgrade is completed.
- For VMware vRealize Orchestrator, run configuration workflow, as detailed in “Downloading and installing the plug-in package for vRO” on page 142. Then use relevant workflows for new storage objects (space, service, volume), as explained in Chapter 5, “IBM Storage Plug-in for VMware vRealize Orchestrator,” on page 141.
- For VASA 1.0 and VASA 2.0 (VVols), update VM storage policy and select storage services instead of capabilities, when setting requirements for the storage resources.

Uninstalling the Spectrum Connect software

If you want to completely remove the IBM Spectrum Connect software from the Linux host upon which it is installed, follow the steps in the following procedure.

Before you begin

Important:

- Before removing Spectrum Connect software, remove all vCenter servers that were registered for vSphere Web Client, as explained in “Removing a vCenter server” on page 110. If any vCenter server is not removed prior to the uninstallation, the IBM Storage Enhancements will remain visible but not functional for that vCenter server.
 - To avoid loss of user accounts, credentials, storage system configurations, storage pool attachments, and vCenter server associations – always back up the Spectrum Connect configuration before any uninstallation.
 - Uninstalling the software on the Spectrum Connect side causes the following features to cease functioning:
 - All CLI and GUI management options on the Spectrum Connect side.
 - IBM Storage Enhancements on the vSphere Web Client side.
 - VASA-related operations on the vCenter server side.
-

Procedure

To uninstall the Spectrum Connect software from the Linux host:

1. Log on to the Linux command prompt environment as a root user.
2. Stop the Spectrum Connect service, as explained in “Checking and controlling the Spectrum Connect Linux service” on page 74.

- Run the standard Linux uninstallation command for each installed package (as detailed in “Performing first-time installation of Spectrum Connect” on page 8).

Operation and management

This chapter describes the initial operation tasks, as well as the full range of management options that are available on IBM Spectrum Connect.

- “Required and optional initial tasks”
- “Configuring LDAP-based directory user access” on page 21
- “Logging in” on page 25
- “Web-based UI” on page 27
- “Running initial setup” on page 30
- “Managing high-availability groups” on page 34
- “Managing server certificates” on page 37
- “Managing Spectrum Connect users” on page 40
- “Managing storage systems” on page 43
- “Managing storage spaces and services” on page 54
- “Administration” on page 73
- “Management from the command-line interface” on page 77
- “Troubleshooting IBM Spectrum Connect” on page 92

Required and optional initial tasks

After IBM Spectrum Connect is installed, different tasks are required before the server can become fully operational.

Refer to the following tables for information about the required and optional management tasks.

Note: Unless specified otherwise in the 'Management method' column, you can initiate tasks from either the web user interface (UI) or command-line interface (CLI) .

Table 3. Required tasks in sequential order

Step	Required task	Management method	Relevant cloud interface	Refer to
1.	Set a password for the 'ibmsc' user if necessary (a Linux root user action), and then switch to the 'ibmsc' user.	CLI	All	<ul style="list-style-type: none"> • “CLI – Switching to 'IBMSC' user mode” on page 78
2.	Run initial setup to define high-availability groups, add SSL certificate and set storage system credentials. If you complete the initial setup, you can skip steps 3–6 below.	UI	All	<ul style="list-style-type: none"> • “Running initial setup” on page 30

Table 3. Required tasks in sequential order (continued)

Step	Required task	Management method	Relevant cloud interface	Refer to
3.	Log in to the UI and then change the password of the initial admin user.	UI	All	<ul style="list-style-type: none"> • “Logging in” on page 25 • “Changing the password of a Spectrum Connect user” on page 42
4.	Define a high-availability group	UI	IBM Storage Provider for VMware VASA	<ul style="list-style-type: none"> • “Managing high-availability groups” on page 34
5.	Add SSL certificate for the Spectrum Connect server	CLI or UI	All	<ul style="list-style-type: none"> • “CLI – Managing server certificates” on page 80 • “Managing server certificates” on page 37
6.	Add the storage system access credentials	CLI or UI	All	<ul style="list-style-type: none"> • “CLI – Adding or removing storage system credentials” on page 81 • “Entering the storage system credentials” on page 44
7.	Add the storage systems to be used	CLI or UI	All	<ul style="list-style-type: none"> • “CLI – Managing storage systems” on page 84 • “Adding a storage system” on page 46

Table 4. Optional tasks

Step	Optional task	Management method	Relevant cloud interface	Refer to
1.	Set the VASA access credentials (“VASA Secret”) to allow connection of vCenter servers that require VASA functions.	CLI or UI	IBM Storage Provider for VMware VASA	<ul style="list-style-type: none"> • “CLI – Setting the VASA credentials” on page 87 • “Setting the VASA credentials” on page 101
2.	Define storage spaces and services.	UI	<ul style="list-style-type: none"> • IBM Storage Enhancements for VMware vSphere Web Client • IBM Storage Plug-in for VMware vRealize Orchestrator • IBM Storage Provider for VMware VASA 	<ul style="list-style-type: none"> • “Adding a storage space” on page 55 • “Adding a storage service” on page 56 • “Defining and attaching storage resources” on page 60
3.	Add vCenter servers.	CLI or UI	IBM Storage Enhancements for VMware vSphere Web Client	<ul style="list-style-type: none"> • “Adding a vCenter server” on page 107

Table 4. Optional tasks (continued)

Step	Optional task	Management method	Relevant cloud interface	Refer to
4.	If you want to use vWC plug-in for managing volumes created on storage resources attached to specific services, delegate the services to the previously added vCenter servers.	CLI or UI	IBM Storage Enhancements for VMware vSphere Web Client	<ul style="list-style-type: none"> • “Delegating storage services to a vCenter server” on page 110
5.	Establish integration with vRealize Orchestrator (vRO). Note: Not applicable for the DS8000 family storage systems.	UI	IBM Storage Plug-in for VMware vRealize Orchestrator	<ul style="list-style-type: none"> • “Managing integration with vRealize Orchestrator” on page 141
6.	Delegate storage services to the previously added vRO server. Note: Not applicable for the DS8000 family storage systems.	UI	IBM Storage Plug-in for VMware vRealize Orchestrator	<ul style="list-style-type: none"> • “Managing integration with vRealize Orchestrator” on page 141
7.	Establish integration with vRealize Operations Manager (vROps). Note: Not applicable for the DS8000 family storage systems.	CLI or UI	IBM Storage Management Pack for VMware vRealize Operations Manager	<ul style="list-style-type: none"> • “CLI – Managing integration with vRealize Operations Manager” on page 87 • “Managing integration with vRealize Operations Manager” on page 151
8.	Configure LDAP-based directory user access to Spectrum Connect.	CLI	-	<ul style="list-style-type: none"> • “Configuring LDAP-based directory user access” on page 21
9.	Manage the Spectrum Connect users.	CLI	-	<ul style="list-style-type: none"> • “CLI – Managing Spectrum Connect users” on page 78 • “Managing Spectrum Connect users” on page 40
10.	Run stateful containers with a storage volume provisioned from an external IBM storage system.	UI	IBM Storage Enabler for Containers	<ul style="list-style-type: none"> • “Managing integration with IBM Storage Enabler for Containers” on page 189

Table 4. Optional tasks (continued)

Step	Optional task	Management method	Relevant cloud interface	Refer to
11.	Automate storage provisioning, host mapping, volume expansion and other storage-related tasks by using Microsoft PowerShell cmdlets.	UI	IBM Storage Automation Plug-in for PowerShell	<ul style="list-style-type: none"> “Managing integration with Microsoft PowerShell” on page 208
12.	Back up or restore a Spectrum Connect configuration, including the data of all existing user accounts, credentials, storage systems, and storage resources.	CLI	-	<ul style="list-style-type: none"> “CLI – Backing up or restoring Spectrum Connect configuration” on page 91

Configuring LDAP-based directory user access

You can allow external directory users to connect to Spectrum Connect and manage it without having a locally-defined user account.

The connection to the directory server is established through Lightweight Directory Access Protocol (LDAP) authentication. When directory server access is enabled, any login attempt (attempt to log in to Spectrum Connect) is authenticated against the defined directory server.

Use the **sc_ldap** CLI command to configure LDAP-based directory user access to Spectrum Connect. Use the required argument after the command, as specified in the following table.

Note:

- When directory user access is enabled and configured through **sc_ldap**, the directory users can access and manage only Spectrum Connect. A separate and unrelated authentication system may be used on the storage system side for directory-based management of the storage system. For more information, refer to “CLI – Adding or removing storage system credentials” on page 81 and to your storage system documentation.
 - Once the connection is established, all users that are defined on the directory server can access and manage Spectrum Connect.
 - Once the connection is established, the users that came up in the configured user search DN results can access and manage Spectrum Connect.
-

Table 5. Arguments for sc_ldap

Argument	Use after sc_ldap to:
<pre> configure -e -a -s <server URI> -t <directory server type> -r <user search DN> -k <user search key> -g <user group DN> -o <user group object class or configure --enable --anonymous --server_uri <server URI> --server_type <directory server type> --user_search_dn <user search DN> --user_search_key <user search key> --group_search_dn <user group DN> --group_object_class <user group object class> </pre>	<p>Enable directory access and establish a connection to a directory server as an anonymous user with the following parameters specified after the -a argument on the command line:</p> <ul style="list-style-type: none"> • Server URI (-s; --server_uri) – Uniform resource identifier (URI) of the directory server. This parameter determines which directory server should be accessed and used for directory user management of Spectrum Connect. • Server type (-t; --server_type) – Type of the directory server. One of the following types can be specified: <ul style="list-style-type: none"> – Active Directory (ACTIVE_DIRECTORY) – Open LDAP (OPEN_LDAP) – Custom (CUSTOM) • User search DN (-r; --user_search_dn) – Distinguished name (DN) to be used for the user search. • User search key (-k; --user_search_key) – Search key of the directory user. Valid only if the specified server type (-t; --server_type) is CUSTOM. • Group search DN (-g; --group_search_dn) – Distinguished name (DN) of the user group for search purposes. • Group object class (-o; --group_object_class) – Object class of the user group. Valid only if the specified server type (-t; --server_type) is CUSTOM. <p>For example:</p> <pre> sc_ldap configure -e -a -s ldap://ad1.ibm.com -t ACTIVE_DIRECTORY -r "CN=Users,DC=mydomain,DC=test,DC=com" -g "CN=sc_TestGrp,CN=Users,DC=mydomain,DC=test,DC=com" </pre> <p>When prompted to enter a password, press Enter without entering any password:</p> <pre> Please enter the BIND_DN password (password not shown): The following changes were applied to the LDAP configuration: ENABLED Please restart the IBM Spectrum Connect to apply the new configuration. </pre> <p>After enabling the directory access, test the directory connection by using the test option (see below). After testing, restart the Spectrum Connect service as explained in “Checking and controlling the Spectrum Connect Linux service” on page 74. Then, use the sc_users command to add LDAP administrators to enable them to access the Spectrum Connect GUI, see “CLI – Managing Spectrum Connect users” on page 78.</p>

Table 5. Arguments for `sc_ldap` (continued)

Argument	Use after <code>sc_ldap</code> to:
<p><code>configure -e -u <Bind DN username></code> <code>-p <Bind DN password></code></p>	<p>Enable directory access and establish a connection to a directory server by using the Bind DN user account that was predefined on the directory server (predefined by the directory server administrator). For this command, specify these two parameters in addition to the entries listed for the anonymous user:</p> <ul style="list-style-type: none"> • Bind DN username (<code>-u</code>; <code>--bind_dn</code>) – Username of the Bind DN user through which access to the directory server is established. Spectrum Connect uses this username to log in to the directory server and establish the connection with it. • Bind DN password (<code>-p</code>; <code>--bind_password</code>) – Password of the Bind DN username. <p>For example:</p> <pre>sc_ldap configure -e -s ldap://myad1.ibm.com -t ACTIVE_DIRECTORY -r "CN=Users,DC=sc,DC=test,DC=com" -g "CN=Users,DC=sc,DC=test,DC=com" -u mybinduser -p mypasswd</pre> <p>When prompted to enter a password, enter the directory server's Bind DN user password:</p> <pre>Please enter the BIND_DN password (password not shown): ***** The following changes were applied to the LDAP configuration: ENABLED Please restart the IBM Spectrum Connect to apply the new configuration.</pre>
<p><code>configure -d</code> or <code>configure --disable</code></p>	<p>Disable directory user access.</p> <p>After disabling the directory access, restart the Spectrum Connect service as explained in “Checking and controlling the Spectrum Connect Linux service” on page 74.</p>
<p><code>list</code></p>	<p>Display the current directory server configuration status (on Spectrum Connect) and Bind DN username.</p>
<p><code>test -u <directory username> -p <password></code></p>	<p>Test a directory user account by specifying the username and password of that account. You can test any user account that came up as configured user search DN on the directory server (the test is not for the Bind DN user account, but for an actual directory account).</p> <p>For example:</p> <pre>sc_ldap test -u mytestuser -p mytestuserpasswd IBM Spectrum Connect LDAP configuration has been verified successfully.</pre>
<p><code>-h</code> or <code>--help</code></p>	<p>Display help information that is relevant to <code>sc_ldap</code>.</p> <p>You can also display help for the <code>configure</code>, <code>list</code>, or <code>test</code> argument if it is typed on the command line as well.</p>

Adding a directory server certificate

If the directory server uses Transport Layer Security (TLS), you must edit the `ldap.conf` file and specify the trusted certificate directory location and filename on Spectrum Connect. Complete the following steps to update Spectrum Connect:

1. Log in to the directory server and issue the following command: **`certutil -ca.cert client.crt`**. This command generates the server certificate.
2. Go to the `/etc/openldap/` directory and edit the `ldap.conf` file by setting the value for the `TLS_CACERT` parameter. The following example shows the contents of the `ldap.conf` file:

```
#LDAP Defaults
#
#BASE dc=example,dc=com
#URI ldap://ldap.example.com ldap://ldap-master.example.com:port#
#SIZELIMIT 12
#TIMELIMIT 15
#DEREF never
TLS_CACERT /etc/openldap/certs/trusted_ldap.pem
```

Make sure that the `TLS_CACERT` parameter has the directory and file name of the new certificate that you generated. After editing the `ldap.conf` file, the `ldap.ini` file is automatically updated.

Editing the ldap.ini configuration file

In addition to using the **`sc_ldap`** CLI command, you can edit the `ldap.ini` configuration file to manually change the directory user access settings.

Attention:

- Do not edit the `ldap.ini` file if you are not familiar with directory setting conventions.
 - Restart Spectrum Connect after editing the `ldap.ini` file to apply the changes.
-

The following example shows the editable parameters and their values specified after the '=' sign:

```
enable_ldap = True
server_uri = ldap://servername.domainname:389
server_type = OPEN_LDAP
user_search_dn = ou=users,dc=dcname,dc=com
user_search_key =
group_search_dn = dc=dcname,dc=com
group_object_class =
bind_dn =
bind_password = <encrypted password>=
bind_pwd_verification = <encrypted key>=
```

The following table summarizes the parameters and their indication.

Table 6. ldap.ini configuration parameters

Parameter	Indication
enable_ldap	True or False. When True and enabled, the login attempt is authenticated against the directory server.

Table 6. *ldap.ini* configuration parameters (continued)

Parameter	Indication
server_uri	Uniform resource identifier (URI) of the directory server.
server_type	Type of the directory server: <ul style="list-style-type: none"> • Active Directory • Open LDAP • Custom
user_search_dn	Distinguished name (DN) to be used for user search.
user_search_key	Search tag for obtaining a unique relative distinguished name (RDN). Commonly used values: uid, preferredId
group_search_dn	Distinguished name (DN) to be used for user group search.
group_object_class	Type of the user group. Commonly used values: GroupOfNames, NestedGroupOfNames, GroupOfUniqueNames, NestedGroupOfUniqueNames, ActiveDirectoryGroup, NestedActiveDirectoryGroup
bind_dn	Username of the Bind DN user through which access to the directory server is established.
bind_password	Password of the Bind DN username. The password is displayed in its encrypted form.
bind_pwd_verification	Verification string for the Bind DN password. The string is displayed in its encrypted form.

Note:

- **user_search_key** and **user_search_dn** return unique results. For example:

```
user_search_key=uid
user_search_dn=ou=users,dc=dcname,dc=com
```

In this case, if the user name is "John", the resulting DN for matching the user over LDAP would be: uid=John,ou=users,dc=dcname,dc=com

- When **server_type** type is Active Directory, the following parameters are used by default:

```
user search key = sAMAccountName
user group object class = NestedActiveDirectoryGroup
```
- When **server_type** type is Open LDAP, the following parameters are used by default:

```
user search key = uid
user group object class = GroupOfUniqueNames
```
- You can use a valid username and password (defined on the LDAP server) with the **sc_ldap test** command to test your LDAP configuration.

Logging in

To log in to Spectrum Connect from a browser, you need to enter the web address (URL) of the Linux host upon which Spectrum Connect is installed.

Procedure

Use one of the following login method to start working with IBM Spectrum Connect:

1. Enter the web address (URL) of the Linux host upon which IBM Spectrum Connect is installed. Use the following format `https://[Spectrum Connect IP address]:8440`.

You can change the default TCP port (8440) at any time by running a script, as explained in “Changing the Spectrum Connect communication port” on page 77.

2. To log in to a Spectrum Connect installed in a shared environment, click the **cloud integration** link on the Hyper-Scale Manager welcome page.

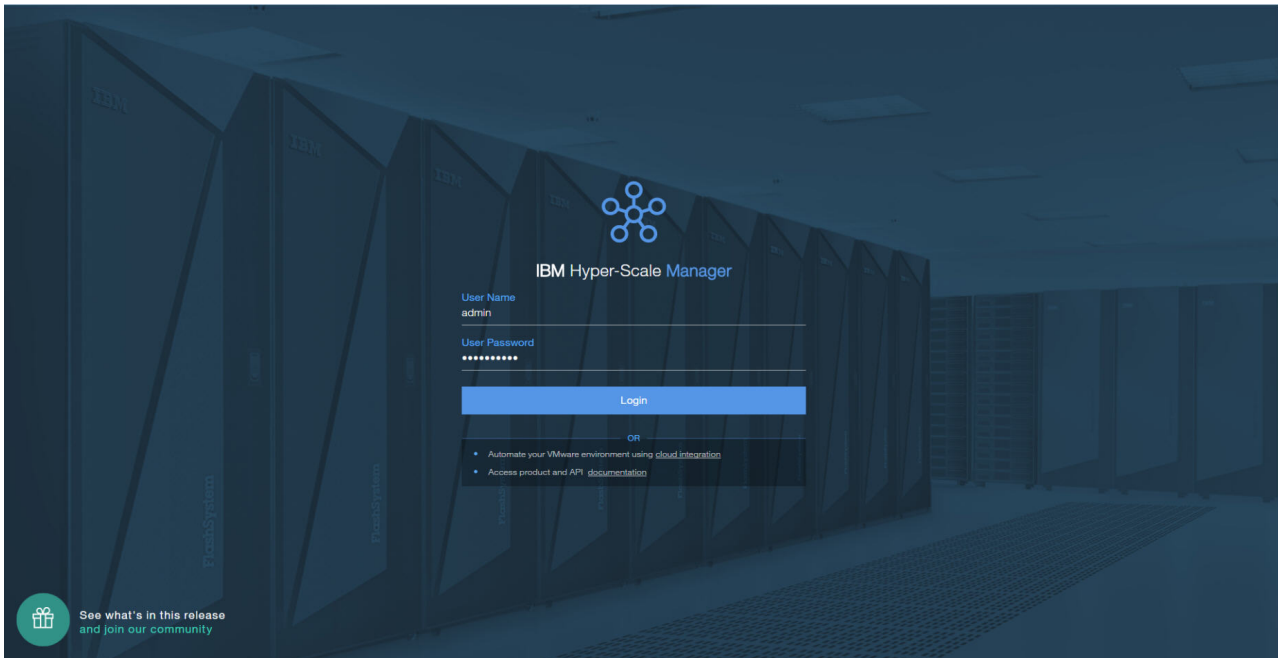


Figure 9. Hyper-Scale Manager welcome page

Important: Adding storage systems function is disabled for Spectrum Connect instance, when it is installed using Hyper-Scale Manager. Storage systems can be added only using the Hyper-Scale Manager UI.

If your web browser displays a connection security message after entering the web address, see “Managing server certificates” on page 37.

The Spectrum Connect login page is loaded .

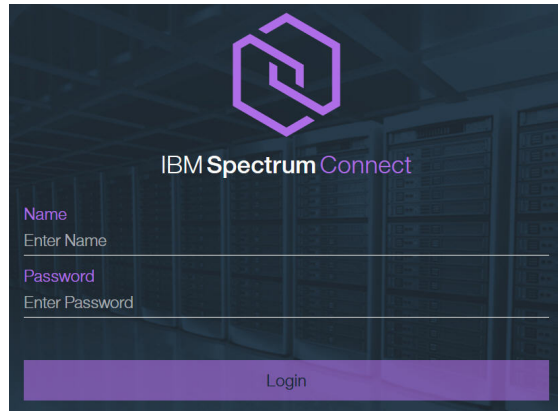


Figure 10. Spectrum Connect login screen in a standard web browser

3. Log in by entering your Spectrum Connect username and password.

Attention:

- After the installation, the initial username is **admin** and the initial password is **admin1!**. To avoid unauthorized access to Spectrum Connect, the default username and password are changed during the initial setup procedure, as explained in “Running initial setup” on page 30. You can also create a new user account and then delete the **admin** account (see “Managing Spectrum Connect users” on page 40).
 - If the Spectrum Connect service is stopped on the Linux host (see “Checking and controlling the Spectrum Connect Linux service” on page 74), it is not possible to log in or perform any GUI operation.
-

What to do next

After successful login, complete the initial setup wizard, which includes the following mandatory configuration steps:

- Setting up a high-availability group.
- Providing SSL certificate.
- Defining storage system credentials.
- Changing default Spectrum Connect credentials.

Web-based UI

The IBM Spectrum Connect UI provides an intuitive easy-to-use browser-based interface for managing IBM storage resources.

The Spectrum Connect web UI consists of the four panes:

- **Interfaces** – integration with vCenter and vRO servers; PowerShell and IBM Enabler for Containers.
- **Spaces/Storage Services** – handling storage spaces and services. These entities are described in the “Storage space and service management” on page 4 section.
- **Storage Systems** – management of storage systems and storage resources.
- **Monitoring** integration with vROps server.

After successful login, the Spaces/Storage Services and Storage Systems panes are displayed.




Figure 11. Spaces/Storage Services and Storage Systems panes

Click **Interfaces** on the left of the screen to go to the **Interfaces** pane, or click **Monitoring** on the right to display the Monitoring pane. When the Interfaces pane is displayed, click **Allocation** to return to the initial view (Spaces/Storage Services and Storage Systems panes).

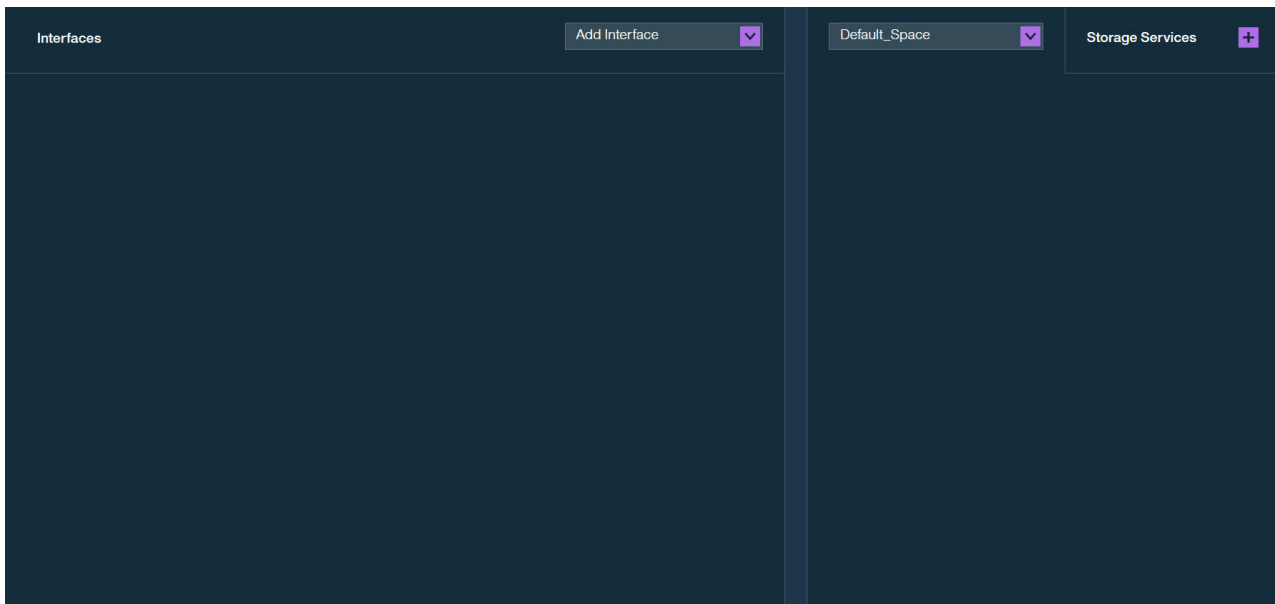


Figure 12. Interfaces and Spaces/Services panes

To return to the initial view from the Monitoring pane, click **Storage Space**.

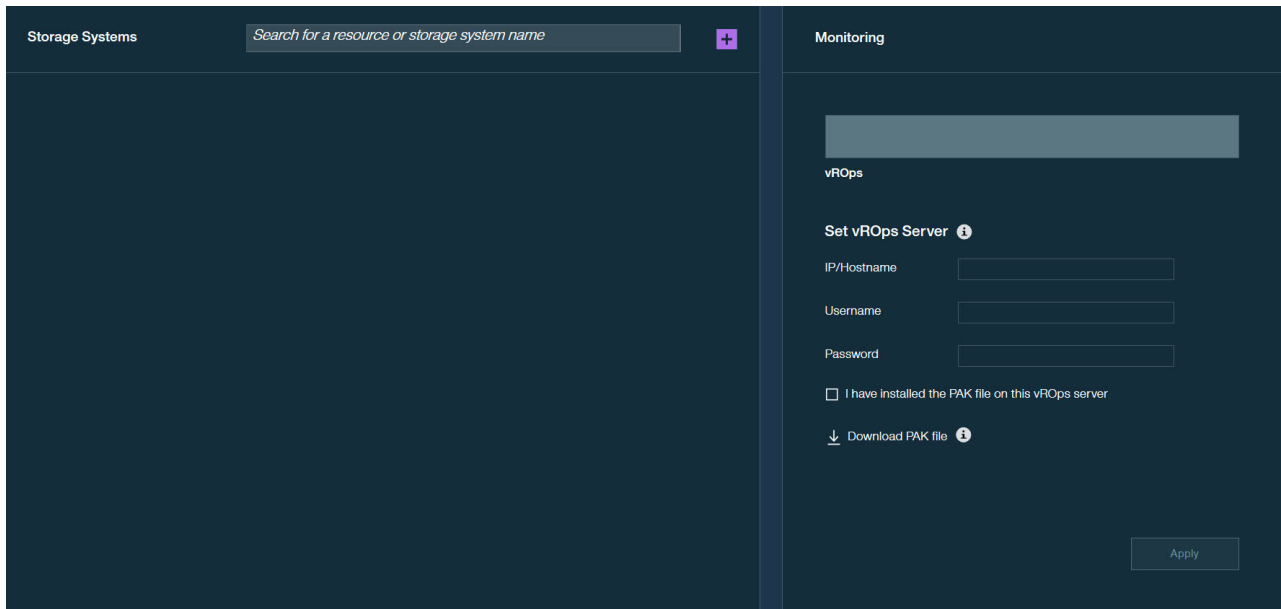








Figure 13. Storage Systems and Monitoring panes

The table below summarizes functionality of the UI elements.

Table 7. Spectrum Connect UI elements

UI element	Description
	Add button. Click this button to add a new object (server, storage service, system, etc).
	Edit button, which is displayed a row in a table is selected. Click this button to configure the object in the table row (user, storage space, etc).
	Home button. Click this button to display the home screen (Spaces/Services and Storage Systems panes).
	Settings button. Use this button to access the Settings menu to: <ul style="list-style-type: none"> • Define storage credentials • Define VASA Provider credentials • Manage certificates • Add and change the Spectrum Connect users • Add and edit storage spaces • Collect logs, display additional help options • Display version number of Spectrum Connect and included applications • Access product documentation
	Remove button. Click this button to remove a storage element or delete a user from the User List.
	Help button. Click this button to display the guided tour, which is also available after the initial login.
	Right and left navigation pointers. They indicate if additional panes are available to the right or to the left of the current display.

Table 7. Spectrum Connect UI elements (continued)

UI element	Description
	Attach/Delegate button. It is available, when an unattached storage object is selected. Click this button to: <ul style="list-style-type: none"> • Delegate a storage service to a previously selected interface (vCenter, vRO, PowerShell, Enabler for Containers) • Attach a storage resource to a previously selected storage service
	Detach/Cancel Delegation button. It is available, when an attached storage resource is selected. Click this button to: <ul style="list-style-type: none"> • Cancel delegation of a storage service to a previously selected interface • Detach a storage resource from a previously selected storage service
	Bar View button. Click this button to display the available storage systems and resources as vertical bars.
	Table View button. Click this button to display the available storage systems and resources as a table.
	Start Monitoring button. It is available, when the mouse pointer is moved over a storage system, which is not monitored via vROps. Click this button to start the monitoring process.
	Stop Monitoring button. It is available, when the mouse pointer is moved over a storage system, which is monitored via vROps. Click this button to stop the monitoring process.

Running initial setup

After you log in to Spectrum Connect, the program invokes a setup wizard to facilitate with the initial configuration.

Procedure

Follow the procedure below to complete the initial setup process, which is run automatically after the first successful login. This procedure is mandatory for subsequent storage provisioning, using Spectrum Connect.

Note: This operation is not available from the CLI .

1. Enter the name of the high-availability (HA) group and the FQDN of the Spectrum Connect server. For detailed description of high-availability groups, see “Managing high-availability groups” on page 34.

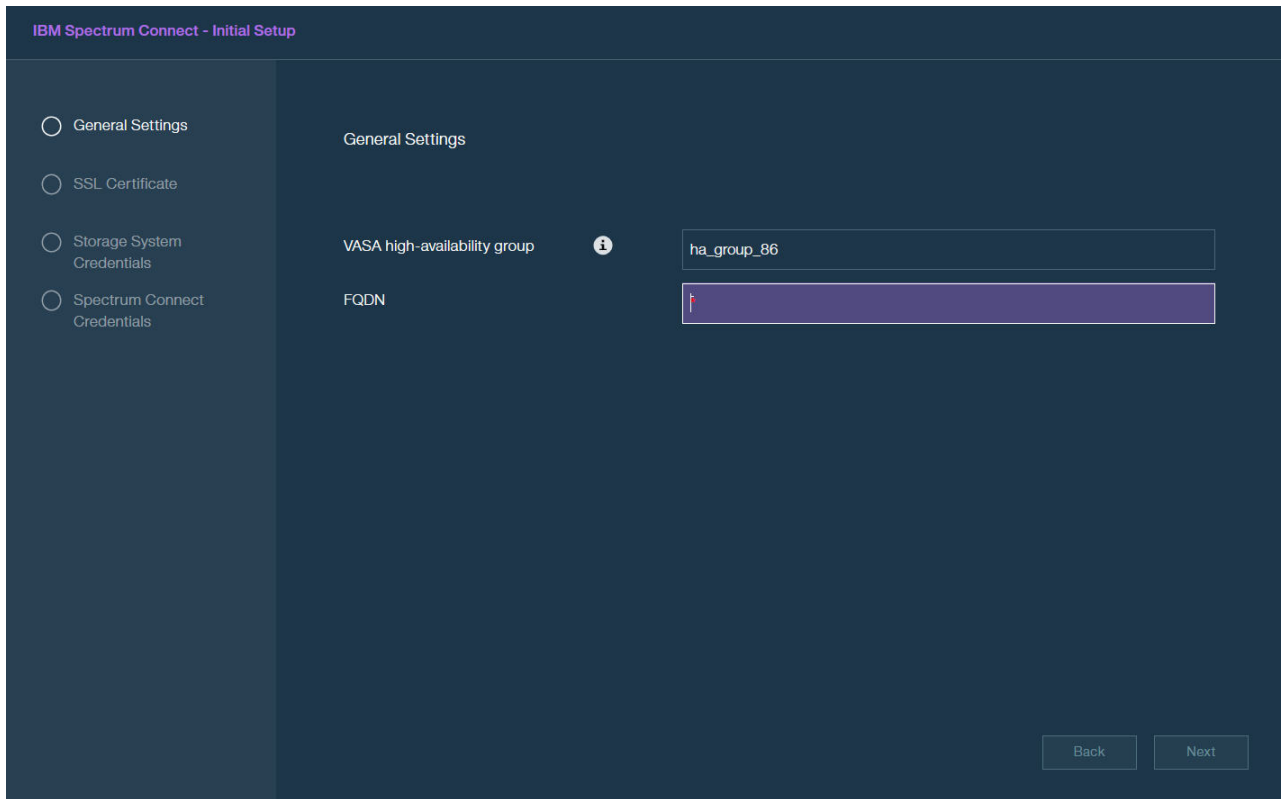


Figure 14. Initial setup wizard, defining HA group

2. Generate or upload SSL certificate, required for establishing a secure communication link with the Spectrum Connect server. For detailed description of SSL certificates, see “Managing server certificates” on page 37.

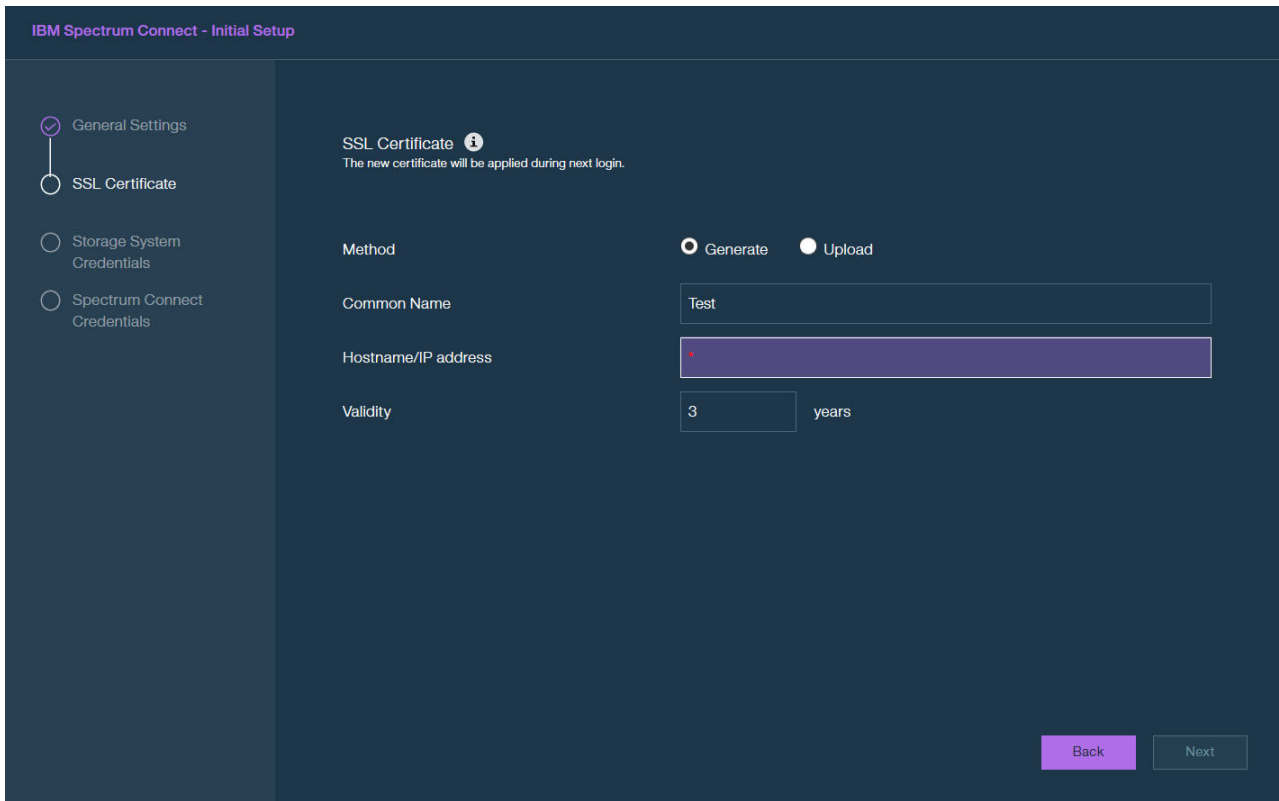


Figure 15. Initial setup wizard, defining SSL certificate

3. Define credentials to be used to connect to the IBM storage systems. For detailed description of storage system credentials, see “Entering the storage system credentials” on page 44.

Important: An identical storage admin user account with identical credentials (the same username and password) must already be predefined on all the IBM storage systems that you intend to use. Spectrum Connect can use only **a single system management account** for accessing all the different storage systems that you use. For storage systems, running Spectrum Virtualize, ensure that the credentials belong to a user account with *VASAProvider* role, if you intend to use VVols. For non-VVol applications, you can use the *Administrator* role.




Figure 16. Initial setup wizard, defining storage system credentials

4. Change default credentials for accessing the Spectrum Connect server.

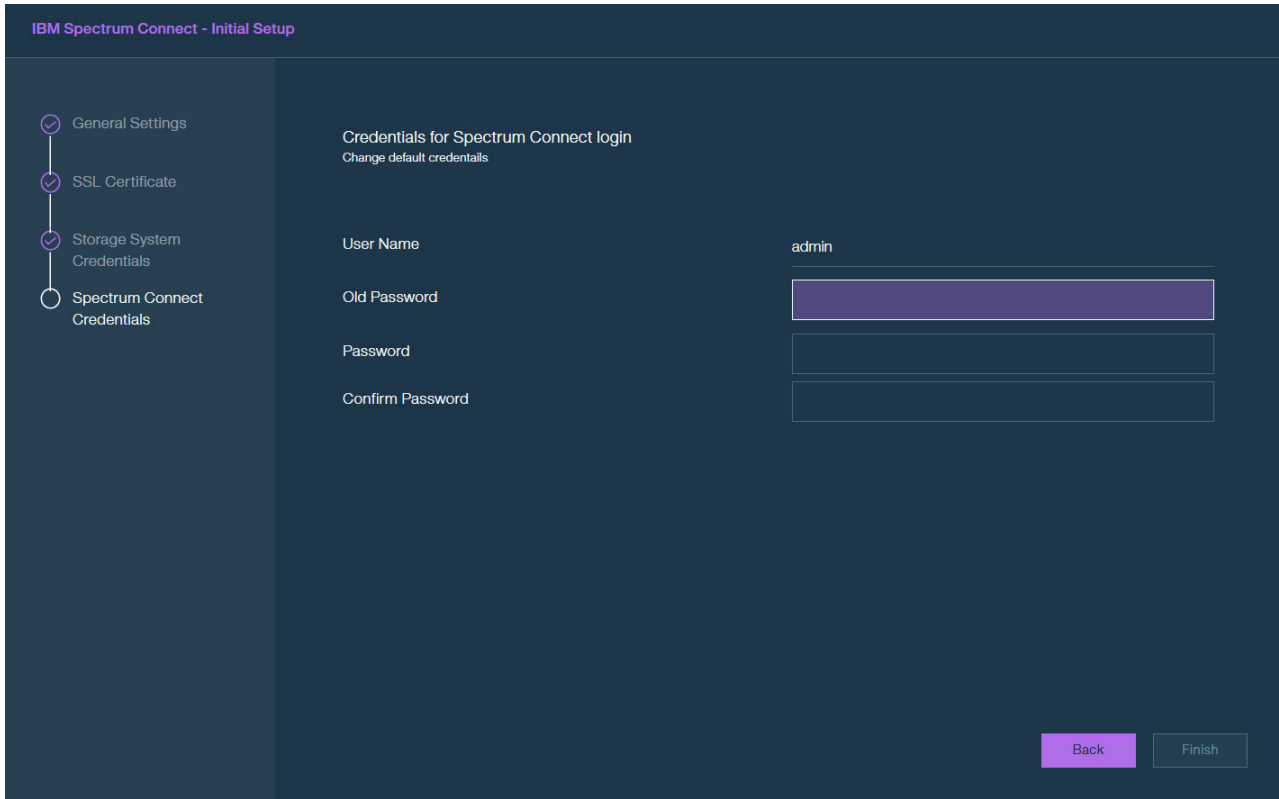


Figure 17. Initial setup wizard, defining Spectrum Connect credentials

What to do next

After successful login, take a guided tour, which is invoked automatically, to familiarize yourself with the main interface elements and principles of operation.

Managing high-availability groups

IBM Spectrum Connect implements VMware's high-availability architecture for multiple VASA providers.

Multiple Spectrum Connect instances that manage the same storage system can be combined into high-availability (HA) group for VASA provider redundancy. The HA technology is implemented in the active/passive mode. In this mode, one Spectrum Connect acts as active and the other acts as standby for transfer of responsibility if the active Spectrum Connect fails or becomes unreachable, as illustrated below.

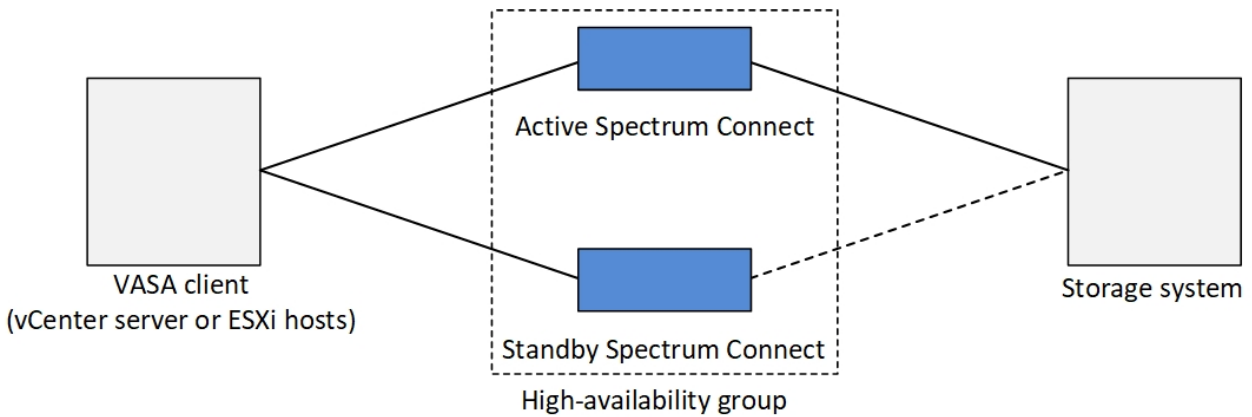


Figure 18. High-availability group concept

To configure an HA group, open the **Settings** menu. For configuration instructions, see “Defining a high-availability group.”




Figure 19. General Settings option on the Settings menu

Defining a high-availability group

To ensure continuous storage management, multiple Spectrum Connect instances must be activated and combined into a high-availability (HA) group.

Before you begin

You must define an HA group before adding storage systems to Spectrum Connect. In addition, verify that both Spectrum Connect servers have a network connection and can communicate with each other.

About this task

An HA group is a combination of two Spectrum Connect instances, operating in the active/passive mode. Active Spectrum Connect instances are backed up by the standby ones to provide uninterrupted service if one of them becomes unavailable. A default HA group (*ha_group_<1-100>*) is created automatically during Spectrum Connect installation.

Procedure

1. Click **General settings** on the Settings menu. The General Settings dialog box is displayed.

2. Enter the FQDN of the first Spectrum Connect and the name of the HA group that you want to create, and then click **Apply**.

Note:

- The FQDN, which you enter, is automatically copied to the Hostname field of the Server Certificate dialog box during the generation of the Spectrum Connect server certificate, as explained in “Managing server certificates” on page 37.
 - A default HA group name is assigned automatically by Spectrum Connect in the *ha_group_number* format. The *number* is randomly selected from 1 to 100. If you intend to change the default HA group name, any XIV storage system with microcode 11.5, connected to Spectrum Connect, which is the default HA group member, must be removed from Spectrum Connect before the group name change.
-

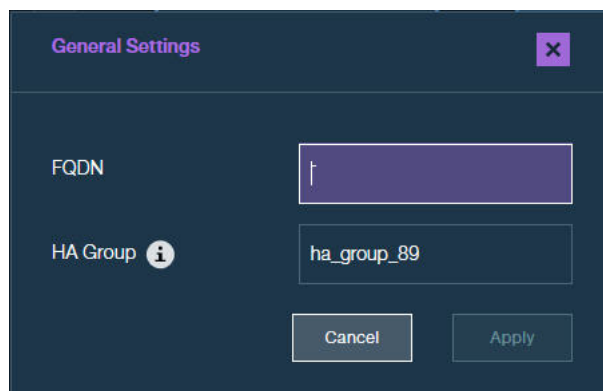


Figure 20. General Settings dialog box

3. Log in to the standby instance of Spectrum Connect and add its FQDN, as explained above. The HA group name must be identical to the value that was provided when registering the active Spectrum Connect instance.

Important:

- The Spectrum Connect software must be installed and configured separately for each HA group member. This is required for obtaining a unique VASA key which is generated during installation.
 - **Do not clone an existing instance of an active Spectrum Connect for use as a standby member.** However, cloning the VM which will host Spectrum Connect prior to installation is acceptable.
 - When deploying a Spectrum Connect virtual machine in a VMware HA cluster, use the Fault Tolerance feature to ensure its continuous availability. If the fault tolerance cannot be achieved, use the VMware High Availability, as the second best choice.
 - Back up your Spectrum Connect configuration regularly, using the **sc_configuration** CLI command. See “CLI – Backing up or restoring Spectrum Connect configuration” on page 91.
-

4. Define storage credentials for all Spectrum Connect instances, as explained in “Entering the storage system credentials” on page 44.
5. Add the same storage system for all Spectrum Connect instances. Spectrum Connect redundancy is available only if active and standby members manage the same storage system. See “Adding a storage system” on page 46.

6. Configure the server certificates. See “Managing server certificates.”
7. Register all HA group members as storage providers in the vCenter server. See “Registering Spectrum Connect as a storage provider on vCenter server” on page 103.
8. Add storage spaces and services to the active Spectrum Connect. There is no need to configure these settings for a standby HA member. See “Managing storage spaces and services” on page 54. Storage spaces and services defined on an active instance of Spectrum Connect do not appear on the standby instance immediately. Any VVol-enabled storage services defined on the active Spectrum Connect are populated on the standby server's GUI during system failover.
A failover to a standby HA group member occurs when the active Spectrum Connect service is stopped or reset. See “Checking and controlling the Spectrum Connect Linux service” on page 74.

Managing server certificates

During the installation, a self-signed Secure Sockets Layer (SSL) certificate is generated to create a secure communication channel for servers and clients. If you already have a trusted certificate that you want to use, you can replace the self-signed certificate with an existing trusted certificate or generate a new certificate.

About this task

A self-signed certificate file, **vp.crt**, and a certificate key file, **vp.key**, are stored in the following directory:

```
/opt/ibm/ibm_spectrum_connect/settings/ssl_cert.
```

Because the self-signed certificate is not automatically recognized by the web browser that you use to log in to Spectrum Connect, you might encounter a connection security warning before you can access the Spectrum Connect login page (see “Logging in” on page 25).




Figure 21. Connection security warning in the Mozilla FireFox web browser

To avoid such warning messages, you need to upload a server certificate which is signed by a public certificate authority (CA), such as VeriSign, or by a CA whose root certificate was imported to your web browser. In addition, you can generate an SSL certificate.

Procedure

1. Click **Server certificate** in the Settings menu. The Server Certificate dialog box is displayed.

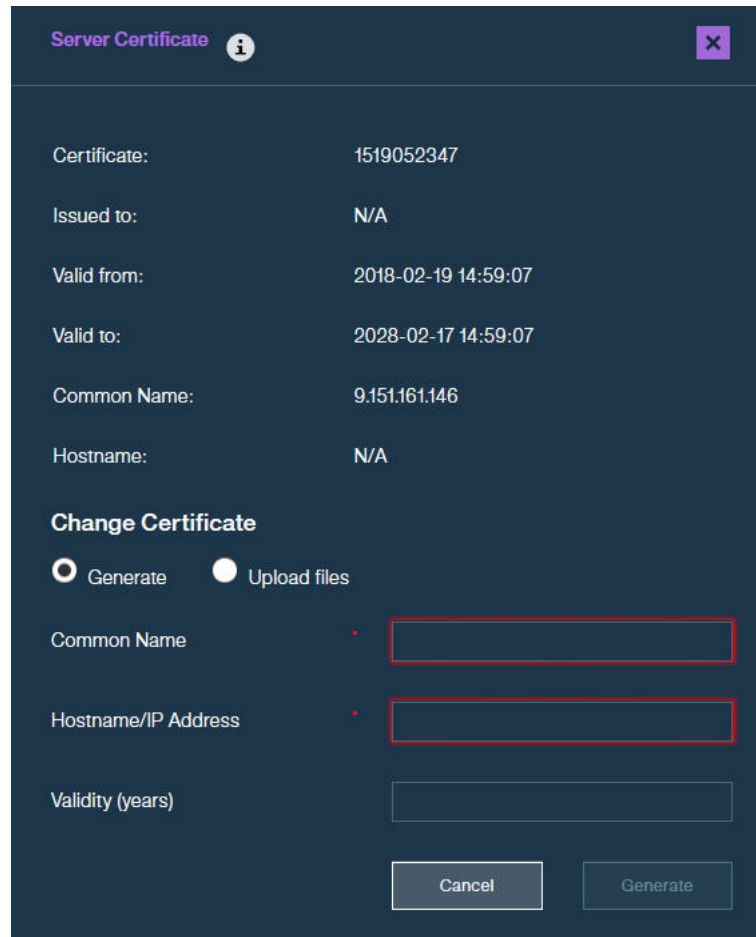


Figure 22. Generate option on Server Certificate dialog box

2. Enter the common name, hostname/IP address of the Spectrum Connect server and certificate validity period, and then click **Generate**.


Note:

- The Spectrum Connect hostname is automatically copied from the FQDN field of the Settings menu. The value is entered during high-availability group definition, as explained in “Defining a high-availability group” on page 35.
- The common name of the Spectrum Connect server must match the hostname/IP address in the URL, which is used, when registering Spectrum Connect as a storage provider on a vCenter server. See “Registering Spectrum Connect as a storage provider on vCenter server” on page 103.
- To support a VMCA subordinate CA, the **ssl_verify_depth** parameter in the `/etc/nginx/conf.d/sc_nginx.conf` file is set to 2 by default. This is sufficient for a single customer root CA and a VMCA, as subordinate CA. If you require more than one CA in the trust chain, increase the value accordingly, and restart the Nginx service. See example below.

```
ssl_verify_client optional;
ssl_verify_depth 3;
ssl_client_certificate ssl_cert/trusted_clients.pem;
```

Spectrum Connect generates the SSL certificate and key files, restarts the Nginx process and refreshes the web browser.

3. Log out and log into Spectrum Connect to complete the certificate generation.
4. To upload a certificate and a certificate key files, select **Upload files** on the Server Certificate dialog box.



The screenshot shows a dark-themed dialog box titled "Server Certificate" with a close button (X) in the top right corner. The dialog displays the following information:

Certificate:	1519052347
Issued to:	N/A
Valid from:	2018-02-19 14:59:07
Valid to:	2028-02-17 14:59:07
Common Name:	9.151.161.146
Hostname:	N/A

Below the information is a section titled "Change Certificate" with two radio buttons: "Generate" (unselected) and "Upload files" (selected). Underneath are two file selection fields:

- "Attach Customer Certificate File:" with a text box containing "No file chosen" and a purple "Browse" button.
- "Attach Key File:" with a text box containing "No file chosen" and a purple "Browse" button.

At the bottom of the dialog are two buttons: "Cancel" and "Upload".

Figure 23. Upload files option on Server Certificate dialog box

5. Click **Browse** and attach your certificate **vp.crt**, and a certificate key files, **vp.key**, and then click **Upload**. Spectrum Connect overwrites the existing SSL certificate and key files, restarts the Nginx process and refreshes the web browser.
6. Log out and log into Spectrum Connect to complete the procedure.

Managing Spectrum Connect users

At any time, you can add new Spectrum Connect user accounts, change user account passwords, or delete existing user accounts.

To access the user management options, click the **Settings** button and select **Users** in the Settings menu. You can then view, add, and manage the Spectrum Connect users as explained in the following sections.




Figure 24. Users option in the Setting menu

Note: The same operations are available from the CLI as well, as explained in “CLI – Managing Spectrum Connect users” on page 78.

- “Adding a new user”
- “Changing the password of a Spectrum Connect user” on page 42
- “Deleting a user” on page 43

Adding a new user

You can add a user account for any authorized user that requires access to Spectrum Connect.

About this task

In addition to storage system credentials, you can define a user for logging into Spectrum Connect and performing GUI or CLI management actions.

Note: All Spectrum Connect users have the same permission level, and can undo any setting or action made by another user.

Procedure

1. Click **Users** in the Settings menu. The User List is displayed, which details all existing users.
2. Click **Add**.




Figure 25. Add option in the Users dialog box

The New User dialog box is displayed.

3. Enter the username and password for the account that you want to create, and then click **Apply**. The minimum password length is seven characters and it must include at least one letter and one digit. The username of the created account is added to the Users list.

Note: The Spectrum Connect GUI is used for managing only local users. To add or delete LDAP admin users, use the CLI `sc_users` command, as explained in “CLI – Managing Spectrum Connect users” on page 78. When directory access is enabled, you can still manage local admin users listed in the **Users** dialog box, but these users will not be able to log into Spectrum Connect.

Changing the password of a Spectrum Connect user

At any time, you can change the password of a Spectrum Connect user account.

About this task

You can modify a user account only by changing its password. Periodic password change is recommended on a regular basis.

Note:

- Spectrum Connect user names cannot be changed. Instead, you can delete a user account (see “Deleting a user” on page 43) and then create a new one (see “Adding a new user” on page 41).
 - This operation is available from the CLI as well, as explained in “CLI – Managing Spectrum Connect users” on page 78.
-

Procedure

1. Click **Users** on the Settings menu. The **User List** is displayed, which details all existing users.
2. Click the **Edit** button in the row of the user account that you want to update.

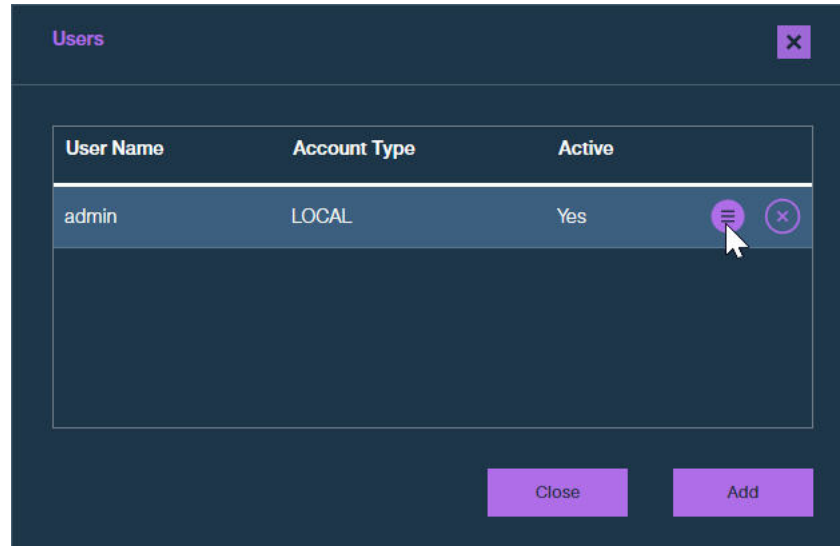


Figure 26. Edit button in user account row

The **Change Password** dialog box is displayed.

3. Enter the current password of the user account, and then the new password as required. The minimum password length is seven characters and it must include at least one letter and one digit. Then, click **Apply**.

Deleting a user

When necessary, you can delete any user account.

About this task

After you delete a *Spectrum Connect admin* user account from the GUI, the specific user can no longer log in to the IBM Spectrum Connect server from either the GUI or the CLI. Although the user deletion is a non-reversible operation, you can redefine the same username as explained in “Adding a new user” on page 41.

Note: Deleting a IBM Spectrum Connect user does not affect the storage credentials.

Procedure

1. Click **Users** in the Settings menu. The **User List** is displayed, which details all existing users.
2. Click the **Remove** button on the row of the user account that you want to remove.
A confirmation message appears.
3. Click **Yes** to delete the user, or **No** to cancel the operation.

Managing storage systems

All IBM storage systems that provide storage resources to your VMware platforms must be defined as storage systems on IBM Spectrum Connect.

To access the storage system management options, go to the **Storage Systems** pane of the Spectrum Connect GUI, illustrated below. You can then view, add, manage

About this task

From the Spectrum Connect GUI you can set or change the current storage system access credentials that are used for accessing all the IBM storage systems.

Important:

- An identical storage admin user account with identical credentials (the same username and password) must already be predefined on all the IBM storage systems that you intend to use. Spectrum Connect can use only **a single system management account** for accessing all the different storage systems that you use. For storage systems, running Spectrum Virtualize, ensure that the credentials belong to a user account with *VASAProvider* role.
For more information about how to define a storage admin account on your IBM storage systems, refer to the relevant storage system management tools documentation.
 - Setting the storage credentials on Spectrum Connect allows you to add the IBM storage systems on the next step.
 - If the system management account is defined on a directory server, see “Checking the format of directory-based storage system credentials” on page 95.
-

Note: The same operations are available from the CLI as well, as explained in “CLI – Adding or removing storage system credentials” on page 81.

Procedure

1. Click the **Settings** button and select **Storage credentials** from the Settings menu.



Figure 28. Settings button

The **Storage Credentials** dialog box is displayed. The dialog box presents the currently defined storage system username in the **User name** configuration box.

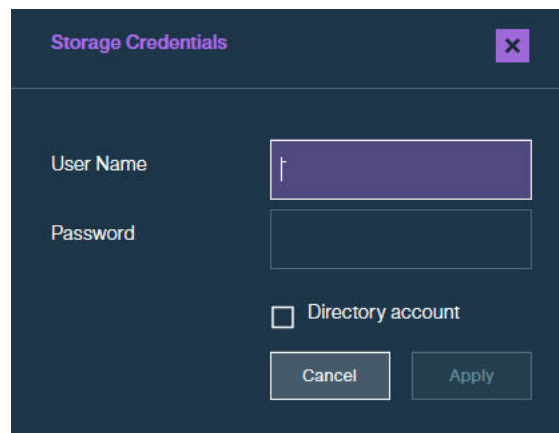


Figure 29. Current storage system username (for all storage systems)

2. Enter the username and password of the storage admin user that was defined on all your IBM storage systems.
3. If the storage admin user account is defined on a directory server, select the **Directory account** check box. If the storage admin user account is locally-defined on the storage system, clear the check box.
4. Click **Apply**.

What to do next

You can now start adding storage systems. Spectrum Connect will use the credentials that you have set in order to connect to the storage systems that you add.

Attention: During regular operation, whenever a directory-based storage admin fails to log in (from the Spectrum Connect side) to any storage system that is in use, Spectrum Connect immediately locks the storage admin user account and all storage systems become inaccessible on the Spectrum Connect side. This is to prevent repeated login attempt failures after which the directory server blocks that user account. In such a case, set the storage system credentials again to unlock the storage admin account on the Spectrum Connect side, with either the same credentials or with updated credentials. The equivalent action in the Spectrum Connect CLI is to use the force credentials options, as described in “CLI – Adding or removing storage system credentials” on page 81.

Adding a storage system

After the storage system credentials are set, you can start adding storage systems to Spectrum Connect.

About this task

The storage systems that you add can be used by the solution components that are included in the Spectrum Connect package (see “Included cloud interfaces” on page 2). You can add each individual storage system separately, as described in the following procedure

Important: Adding storage systems function is disabled for Spectrum Connect instance, when it is installed using Hyper-Scale Manager. In this case, storage systems can be added only using the Hyper-Scale Manager UI.

Procedure

1. Click the **Add** button on the Storage Systems pane.




Figure 30. Add button

The **Add New IBM Storage System** dialog box is displayed.




Figure 31. Add New IBM Storage System dialog box

2. Enter the management IP address or hostname of the array.
3. Click **Add**. If the credentials are correct (as previously defined; see “Entering the storage system credentials” on page 44) and the IP connection is established, the storage system is added to the **Storage Systems** pane.

If the storage system includes previously defined storage pools, you can view their names and sizes on the system. If no pools exist on the system, you can define them, as explained in “Defining and attaching storage resources” on page 60.

Spectrum Connect fetches information about storage resources on a system every ten minutes by default. You can refresh the storage resource information immediately by right-clicking a system that you want to refresh, and then selecting **Refresh**.

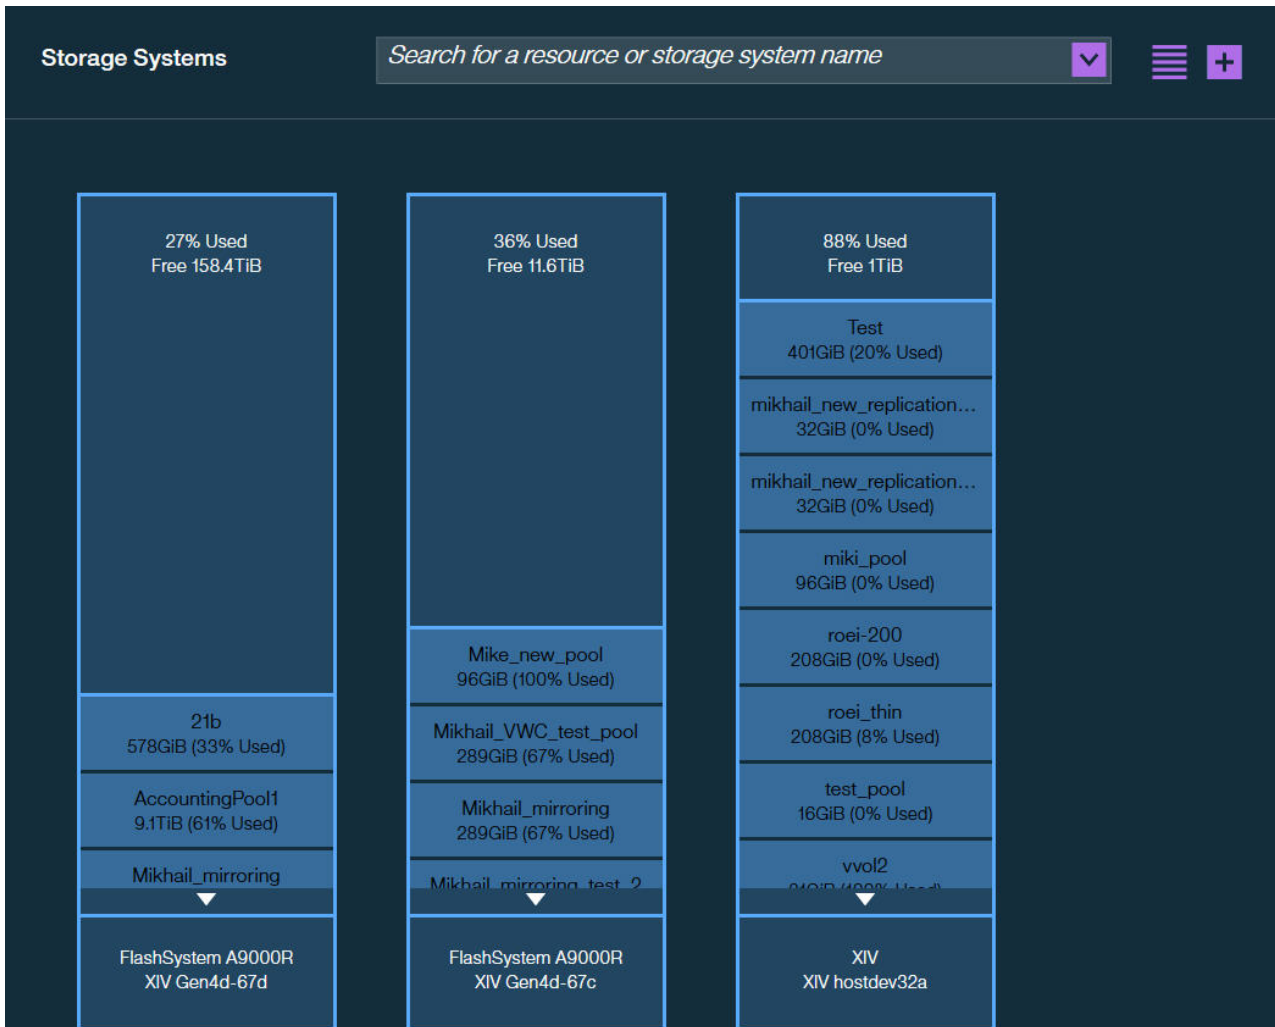


Figure 32. Storage Systems pane, bar view

If a storage system connectivity problem occurs, Spectrum Connect displays a red frame around the system. Move the mouse pointer over the system, and click the red triangle to display the error message.

4. Click the **Table View** button to display the existing storage systems as a table.

Storage Resources











RESOURCE ▲	DOMAIN	STORAGE AR...	USAGE	SIZE(...)	SERVICE	TYPE
18	dana-domain	XIV Gen3P1-22b	100% 	0		XIV
21b	dana-not-managed	XIV Gen4d-67d	33% 	577.59		FlashSystem ..
31d-ayab-pool		XIV hostdev31d	0% 	112.18		XIV
51-ayab-Pool-dev31d		XIV hostdev31d	0% 	112.18		XIV
A	dana-domain	XIV Gen3P1-22b	100% 	0		XIV
AccountingPool1	sra_demo	XIV Gen4d-67d	60% 	9337.76		FlashSystem ..
Mike_new_pool	mikhail_domain	XIV Gen4d-67c	100% 	96.26		FlashSystem ..
Mikhail_VWC_test_p..	mikhail_domain	XIV Gen4d-67c	66% 	288.79		FlashSystem ..
Mikhail_mirroring	mikhail_domain	XIV Gen4d-67c	66% 	288.79		FlashSystem ..
Mikhail_mirroring	mikhail_domain	XIV Gen4d-67d	66% 	288.79		FlashSystem ..

Figure 33. Storage Systems pane, table view

What to do next

See “Working with storage system views” to learn how arrange storage systems in the table view or search for a specific storage system or resource.

Working with storage system views

Spectrum Connect offers powerful search tools to help you locate a specific storage system or resource, according to its name or other attributes.

About this task

After the storage systems become visible on the Storage System pane of Spectrum Connect, you can arrange the way they appear in the table view. In addition, you can use the search box on the top of the Storage System pane to locate a specific storage system or resource, as described below.

Procedure

1. In the table view mode, click on an attribute column header to arrange the existing storage systems and resources, according to one on the following criteria:

- Resource name
- Parent pool name
- Domain name
- Storage array name
- Storage space usage in %
- Size in GiB
- Storage service it is attached to, if any
- Storage system type

You can change the item order by clicking the arrow in the column header.

RESOURCE	DOMAIN	STORAGE AR...	USAGE	SIZE(...)	SERVICE	TYPE
vvol2	igor_domain	XIV hostdev32a	100%	31.31		XIV
vvol3	igor_domain	XIV hostdev32a	100%	31.31		XIV
liort_ubiq_pool_2	liort_domain	XIV Gen4d-67d	80%	481.32		FlashSystem
pool12	olgas-domainc	XIV Gen4d-67c	80%	481.32		FlashSystem
pool12	olgas-domainc	XIV Gen4d-67d	80%	481.32		FlashSystem
roei_replication	roei_domain	XIV Gen4d-67c	72%	1058.92		FlashSystem
vvol_scale	vvol_scale_env	XIV Gen3P1-22b	69%	10079.84		XIV
Mikhail_VWC_test_p..	mikhail_domain	XIV Gen4d-67c	66%	288.79		FlashSystem
Mikhail_mirroring	mikhail_domain	XIV Gen4d-67c	66%	288.79		FlashSystem
Mikhail_mirroring	mikhail_domain	XIV Gen4d-67d	66%	288.79		FlashSystem

Figure 34. Storage Systems pane with storage elements arranged according to descending space usage

2. Enter a full or partial name of a storage system or resource into the search box. Spectrum Connect displays search results, highlighting the names that match the search string.

The screenshot shows a table titled 'Storage Resources' with a search bar containing 'dev'. The table has columns for RESOURCE, STORAGE ARRAY, USAGE, SIZE(...), SERVICE, and TYPE. Two rows are visible, both with 0% usage and a size of 112.18.

RESOURCE	STORAGE ARRAY	USAGE	SIZE(...)	SERVICE	TYPE
xavi-vrops-dev-env	XIV hostdev31d	0%	112.18		XIV
51-ayab-Pool-dev31d	XIV hostdev31d	0%	112.18		XIV

Figure 35. Storage Systems pane with highlighted search string results

3. Click the arrow in the search box to display the advanced filtering tool.

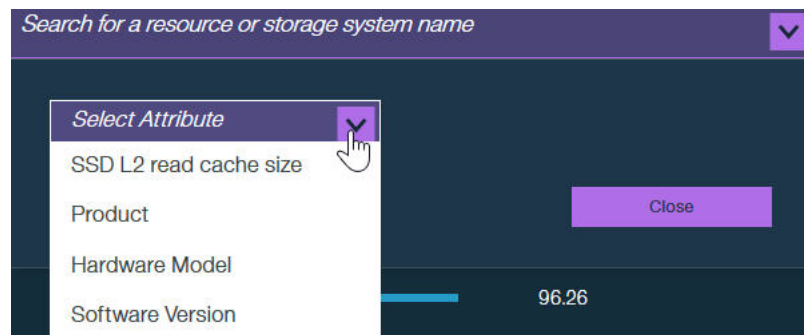


Figure 36. Advanced filtering tool

4. In the advanced filtering tool, select one or more of the following filtering attributes and their values from the drop-down lists. The available attributes may vary, depending on a storage system in use.
 - Software version
 - Disk size. This attribute is not available for storage systems, running IBM Spectrum Virtualize.
 - Hardware model
 - Product type

Spectrum Connect arranges storage elements according to the selected filtering attributes and their values. The best matching results are grouped at the left in the bar view mode and at the top in the table view.

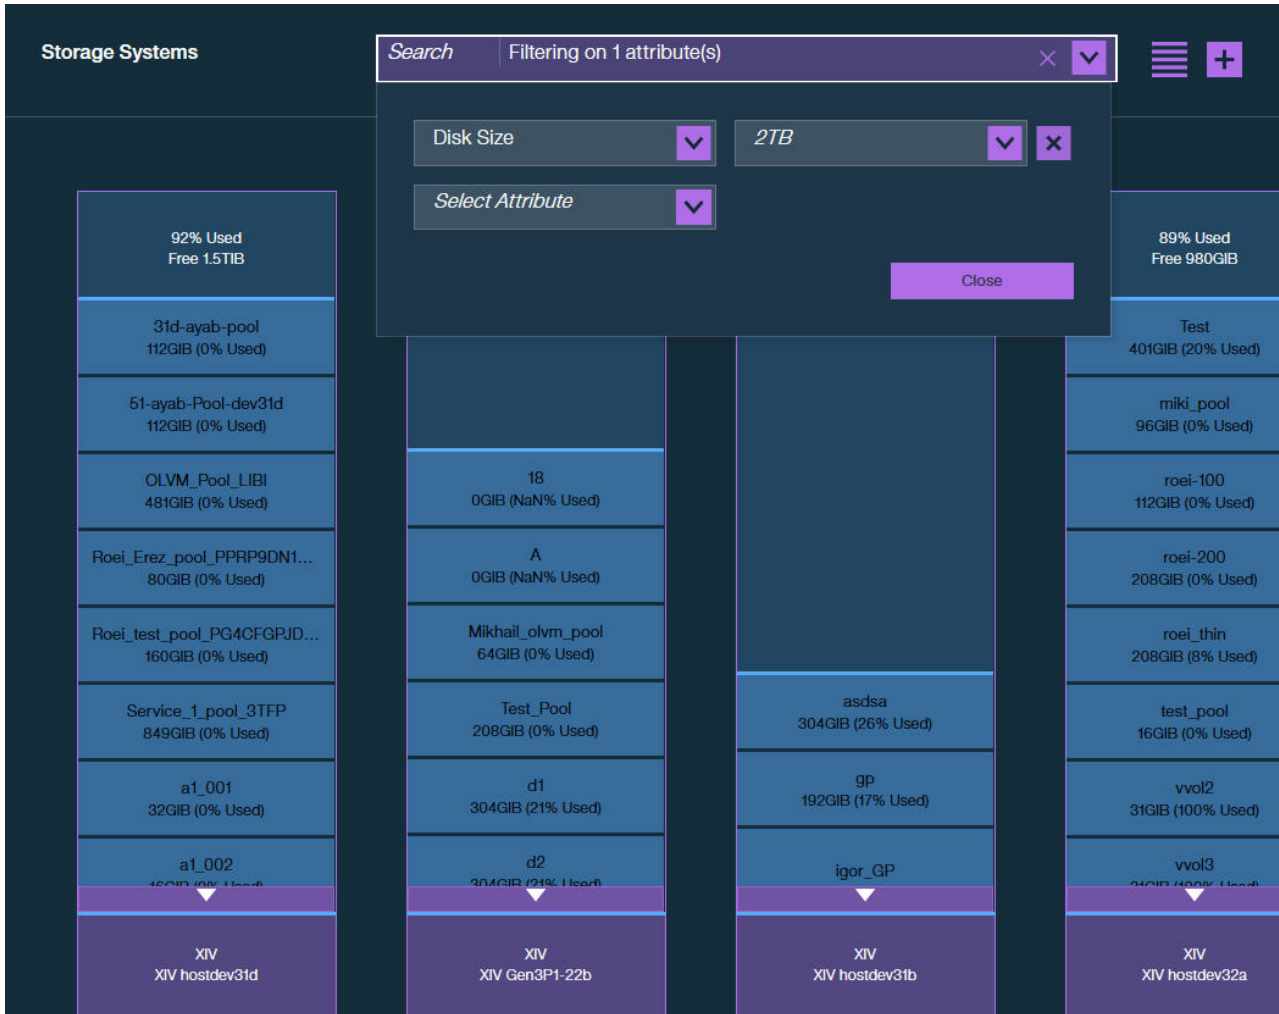


Figure 37. Storage Systems pane with storage elements filtered according to disk size

Modifying the IP address or hostname of a storage system

At any time, you can modify the IP address or hostname of an added storage system.

About this task

If the management IP address or hostname of a storage system changes, you can update this change on Spectrum Connect without having to remove or re-add the storage system.

Procedure

1. In the Storage Systems pane, right-click a storage system that you want to update, and then select **Modify**. The **Array Settings** dialog box is displayed.

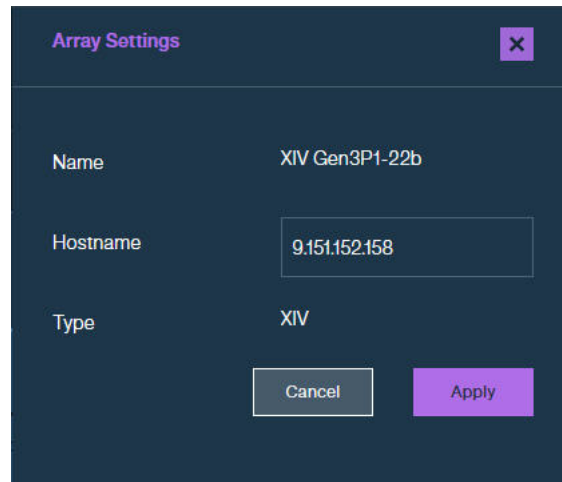


Figure 38. Array Settings dialog box

Note: If a storage system has multiple management IP addresses, you can display all of them by moving the mouse pointer over the **Hostname** field in the **Array Settings** dialog box.

2. Enter the new management IP address or hostname of the storage system, and then click **Apply**.

Removing a storage system

If a storage system is no longer needed, you can remove and disconnect it from Spectrum Connect.

About this task

- A removed storage system, along with its storage pools and volumes, can no longer be managed by the included solution components (see “Included cloud interfaces” on page 2).
- If the removed storage system contains working storage resources and volumes, the information of these storage pools and volumes is no longer displayed in vSphere Web Client. However, **vSphere data access and service level for these storage pools and volumes is not affected**. In addition, the removed system and its storage resources and volumes can be managed from the standard IBM storage system management tools.
- After the removal, you can add the storage system back again to fully restore its management.

Procedure

1. In the **Storage Systems** pane, right-click a storage system that you want to remove, and then select **Delete**. The confirmation message is displayed.
2. Click **OK** to confirm your action. The storage system is removed.

Important: Do not change VVol-related configuration on Spectrum Connect, after the storage system holding the VVol metadata information was detached from Spectrum Connect, and then connected again.

- If Spectrum Connect is not registered as a storage provider on vCenter, perform the registration (this recovers the metadata) and then connect the storage system.
 - If Spectrum Connect is registered as a storage provider on vCenter, wait until vCenter finishes VASA Provider activation for this storage system (it happens automatically). The activation is completed after the storage system appears under Spectrum Connect in the vCenter. To verify that activation is successful, refresh the Spectrum Connect GUI. The storage resource that had been attached to a VVol-enabled service before the storage system was removed, appears as attached to this service.
-

Managing storage spaces and services

After defining physical storage resources on Spectrum Connect you must add virtual storage elements: spaces and services.

Spectrum Connect administers storage at the virtual level, using spaces and services. This simplifies storage provisioning, and allows users to allocate their own storage resources to suit requirements of a specific VM.

Storage spaces and services are described in the “Storage space and service management” on page 4 section of Chapter 1.

Storage spaces are added and managed, using the **Manage Spaces** option in the Settings menu. You can also add a new space via the drop-down menu of the Default Space tab on the **Spaces/Storage Services** pane.

Storage services are added and configured via the Spaces/Storage Services pane of the Spectrum Connect UI.




Figure 39. Spaces/Storage Services pane

- “Adding a storage space” on page 55
- “Removing a storage space” on page 56
- “Adding a storage service” on page 56
- “Removing a storage service” on page 59
- “Defining and attaching storage resources” on page 60
- “Defining and attaching replicated storage resources” on page 65
- “Modifying storage resources” on page 71
- “Modifying storage resource attachments” on page 73

Adding a storage space

Once physical storage elements are defined in the IBM Spectrum Connect, you can continue by adding the first virtual entity – a storage space.


About this task

A storage space is a logical grouping of storage services and underlying physical storage resources. This combination determines the storage that is available when a user requests a storage service provisioning.

A storage space is added by providing it with a meaningful name and description.

Procedure

1. Click the **Settings** button and select **Manage storage spaces** from the Settings menu. The list of storage spaces is displayed.
2. Click **Add** . The **New Storage Space** dialog box is displayed.



The image shows a dark-themed dialog box titled "New Storage Space". At the top left of the dialog is the title "New Storage Space" followed by an information icon (i) and a close button (x). Below the title bar, there are two input fields. The first is labeled "Storage Space Name" and is a single-line text box. The second is labeled "Description" and is a multi-line text area. At the bottom of the dialog, there are two buttons: "Cancel" and "Apply".

Figure 40. New Storage Space dialog box

3. Enter the name and description of the storage space that you want to create, and then click **Apply**. The name and description of the created storage space is added to the Spaces list.

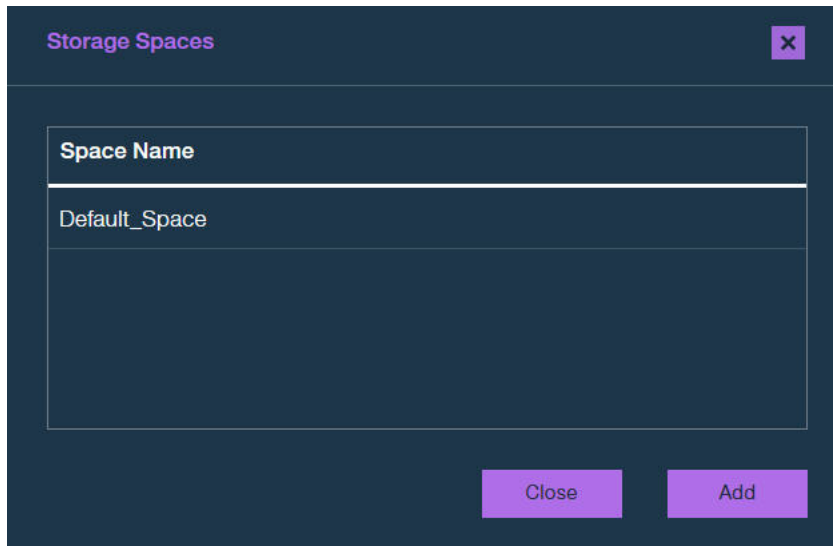



Figure 41. List of storage spaces

Note: You can also add a new space by clicking **Add New Space** in the drop-down menu of the **Default Space** tab on the **Spaces/Storage Services** pane.



Removing a storage space

If a storage space is no longer needed, you can remove it from Spectrum Connect.

Before you begin

Before removing a storage space, delete all services that are defined on this space. See “Removing a storage service” on page 59

About this task

A removed storage system, along with its services and pools, can no longer be available for storage service provisioning.

Procedure

1. Click the **Settings** button and select **Manage spaces** from the Settings menu. The Spectrum Connect space list is displayed, detailing all existing storage spaces.
2. Click the **Remove** button on the row of the storage space that you want to remove. A confirmation message appears.
3. Click **OK** to remove the space, or **Cancel** to cancel the operation.

Adding a storage service

After a storage space is defined, you can start adding storage services to the space.

About this task

Storage services contain one or more physical storage pools. In addition to storage capacity, a service has a set of capabilities, defining the storage quality, such as thin/thick provisioning, compression, encryption, etc.

The services that you add become available for the solution components included in the Spectrum Connect package (see “Included cloud interfaces” on page 2). When provisioning storage, the end users consume it from the spaces and services without dealing underlying physical storage infrastructure.

You can add each individual storage service separately, as described in the following procedure.

Procedure

1. In the Spaces tab of the **Spaces/Storage Services** pane, select a space on which you want to create a new service.
2. Click **Add** button on the Spaces/Storage Services pane. The **New Storage Service** dialog box is displayed.




Figure 42. New Storage Service dialog box

3. Use the **Select Capability** drop-down box to define the new service attributes and their values. See table below for details.

Important: Spectrum Connect uses dynamic filtering to simplify selection of a storage system and a resource, which provides the best match for selected service attributes. When a service is created, Spectrum Connect groups the best matching storage systems in the left section of the Storage Systems pane. Also, the matching storage resources are grouped at the top of the respective storage system. In addition, storage system and resources that do not support the selected service attributes are grayed out in the bar view or removed from the list in the table view.

Table 8. Service parameters

Parameter	Description and values
Name	Alphanumeric string for service identification. This is a mandatory field. The service name must not contain spaces.
Description	Alphanumeric string for service description.
QoS	<p>Enables the use of the Quality of Service (QoS) feature for the service. QoS is applicable to volumes (Max Independent Performance) or storage resources (Max Shared Performance), setting the IOPS and bandwidth limits within the following ranges:</p> <ul style="list-style-type: none"> • IOPS: 0-100000 • BW (bandwidth): 0-10000 MB/s <p>Currently, the QoS capability is not available for IBM SAN Volume Controller (SVC) storage systems or Spectrum Accelerate Family products, using the domain administrator storage credentials.</p>
Encryption	Enables encryption for the service. If enabled, you can attach only encrypted storage resource to the service.
Flash	Enables utilization of a storage resource, located on a flash-based storage resource. This can be one of the following storage systems: FlashSystem 900, FlashSystem V9000, Storwize Family.
Replication	Enables synchronous mirroring for the service provisioned on Spectrum Accelerate Family products. This capability is not available for VVol-enabled services. Additional service capabilities, such as encryption or compression, will be applied to both primary and secondary storage pools, which will be attached to the replicated service. See “Defining and attaching replicated storage resources” on page 65.
Availability	<p>Enables the use of IBM high-availability technologies for HA storage deployments of IBM SAN Volume Controller (SVC) and Storwize storage systems. Select one of the following HA methods:</p> <ul style="list-style-type: none"> • Stretched – SVC stretched cluster topology. The data is mirrored across two different locations, allowing for fast non-disruptive failover with no interruption of service and automatic re-synchronization in the event of small-scale outages. • HyperSwap – SVC/Storwize local HyperSwap topology. The local HyperSwap is similar to the stretched cluster method, spreading the nodes of the system across two sites, with storage at a third site acting as a tie-breaking quorum device. However, it does not support VVols. The HyperSwap functionality is available for storage systems that support more than one I/O group. • Mirrored – Data mirroring. Two mirrored copies of the volume are maintained on the same storage system, while each copy of the volume is kept in a separate storage pool. <p>With the default, Regular, option, a storage system uses a single volume.</p>

Table 8. Service parameters (continued)

Parameter	Description and values
Space Efficiency	<p>Enables storage space efficiency features for the service. When selected, you can configure the service to be attached to a thick- or thin-provisioned storage resource.</p> <p>Configuration considerations</p> <p>When adding a VVol-enabled service, define its space efficiency during VM provisioning via the vSphere web client. To allow this, disable space efficiency in Spectrum Connect.</p>
Data Reduction	<p>Enables the use IBM Real-time Compression™ with or without data deduplication.</p> <p>Configuration considerations</p> <p>A service with enabled IBM Real-time Compression will be able to support the compression-compatible (thin-provisioned) storage resources on XIV and Spectrum Accelerate systems. For the FlashSystem V9000, Storwize Family storage systems, a storage resource must have the data compression enabled prior to service attachment (via product CLI or UI).</p> <p>Currently, DS8000 storage systems do not support IBM Real-time Compression.</p>
VVol Service	<p>Enables virtual volume functionality for the service.</p> <p>The virtual volume functionality is supported by the IBM XIV (11.5.1 or later) and storage systems that run IBM Spectrum Virtualize (7.6 or later).</p> <p>An XIV VVol-enabled service does not support IBM Real-time Compression.</p>

4. Click **Create** to finish the procedure. A new service is added to the current Spaces tab.
5. You can edit the service properties by right-clicking a service which you want modify, and then selecting **Modify Properties**.

What to do next

Define and attach storage resources (pools) to the service, as explained in “Defining and attaching storage resources” on page 60. For replicated serviced, define and attach paired pools, which contain primary and secondary pools, as detailed in “Defining and attaching replicated storage resources” on page 65.

Removing a storage service

If a storage service is no longer needed, you can remove it.

About this task

- A removed storage service, along with its pools, can no longer be managed by the included solution components.
- If the removed storage service contains working storage resources, these resources become available for allocation by other existing services.

Procedure

1. On the **Spaces/Storage Services** pane, right-click a service that you want to remove, and then select **Remove**. A confirmation message is displayed.
2. Click **OK** to confirm the removal.

Defining and attaching storage resources

After the storage systems and services are added, you can start defining resources on the systems.

Before you begin

Verify that the storage resource, which you intend to define and attach, will be compatible with the service in terms of space efficiency, compression mode or any other attribute defined for the service.

About this task

If no resources (pools) exist on the storage systems, you can define the resources and attach them to a service to make the physical storage available for the Spectrum Connect components and end users. You can attach several resources that belong to the same or different storage systems to a single service. However, a specific resource can be added to one service only.

Note: For VVol-enabled services, the minimum recommended storage resource size is 500 GiB.

Procedure

1. On the **Spaces/Storage Services** pane, select the storage space from which you want to choose storage services. The available services that reside on the selected storage space are immediately displayed.
2. Click on a service to which you want to attach a resource. The service color changes to blue to indicate selection.
3. On the Storage Systems pane, right-click a storage system on which you want to define a new resource, and then select **Add Resource**. The Add New Resource dialog box is displayed.

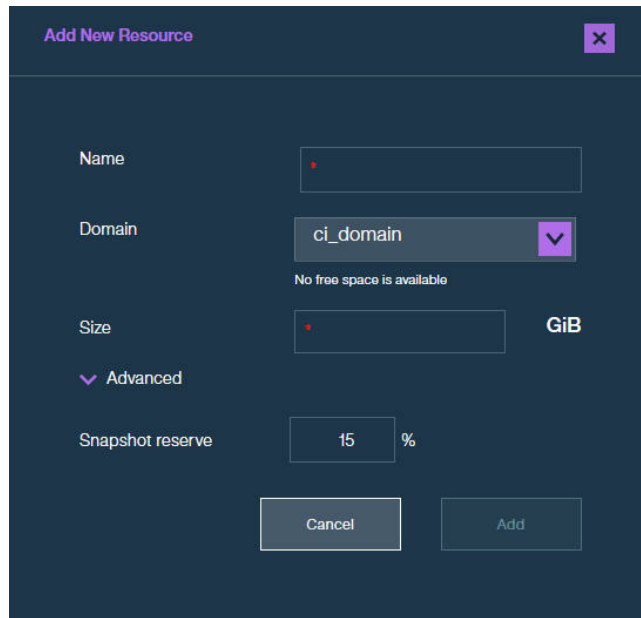


Figure 43. Add New Resource dialog box

4. In the **Add New Resource** dialog box, enter configuration parameters for a new resource. See table below for details.

Table 9. Storage resource parameters

Parameter	Description and values
Name	Alphanumeric string for storage resource identification. The name must not contain spaces. This is a mandatory field.
Domain	<p>Management domain of the pool IBM Spectrum Accelerate systems, or a parent pool for storage systems that run IBM Spectrum Virtualize (7.6 or later).</p> <p>Configuration considerations</p> <p>To use management domains, verify that:</p> <ul style="list-style-type: none"> • Domains are already defined on your storage system. • Default credentials of a storage admin user are assigned to the domain.
Size	Storage resource size in GiB.

Table 9. Storage resource parameters (continued)

Parameter	Description and values
Over-provisioning	<p>Percentage of over-provisioned storage space on the service, defining the ratio between logical and physical storage capacity. For example, when configured to 100, it sets a 1:1 ratio between the two values, while a value of 200 defines the logical capacity (soft size) to be twice the physical capacity (hard size).</p> <p>Configuration considerations</p> <p>The recommended over-provisioning value for a VVol-enabled service with XIV storage systems is 400[®]%.</p> <p>The recommended over-provisioning value for a service with IBM Real-time Compression is 200%.</p> <p>The XIV utilizes thin provisioning for all VM configuration volumes. When creating a regular VM on an XIV disk, you can choose the provisioning mode, according to your application requirements.</p> <p>For other cases, you can choose any value:</p> <ul style="list-style-type: none"> • 100%, if you do not want to utilize thin provisioning. • Above 100%, if you want to take a risk of allocating resources that the XIV storage system may not have. <p>In addition, when utilizing the over-provisioning, XIV storage systems must be also configured to allow support for this feature. When used with the XIV domains, the domain must be also configured with matching soft and hard capacity settings. To enable the VVol functionality at the XIV side, see “Creating a VVol-enabled service” on page 225.</p> <p>Important</p> <p>The managed domain that you created cannot be used for traditional volumes without virtualization. You must create a separate regular domain for them. This domain must have the same user and the ESXi hosts that you intend to manage. However, you need to create a separate storage resource and a new service on the regular domain via Spectrum Connect for subsequent use by the vWC.</p> <p>The over-provisioning is relevant for XIV storage systems only.</p> <p>Configuration examples</p> <ul style="list-style-type: none"> • Storage pools in a regular service with over-provisioning at 100% are created as thick pools. • Storage pools in a regular service with over-provisioning at 150% are created as thin pools. If a hard pool size is 500 GiB, its soft size is 750 GiB. • Group storage pools in a VVol-enabled service with over-provisioning at 400% are created as thin pools with their soft sizes to be four time bigger than their hard sizes.
Snapshot reserve	<p>Percentage of storage space on the service reserved for snapshots.</p>

Note: EasyTier-enabled parent pools defined on storage systems, running Spectrum Virtualize software, can be attached to storage services to create a compatible service. IBM EasyTier is a pool-level feature and it is transparent to Spectrum Connect. As a result, all volumes in the EasyTier-enabled service retain this capability.

5. Click **Add**. The storage resource is created on the storage system and attached to the service. The resource color changes to blue to indicate the successful attachment.

The resources defined on the system before the system was attached to Spectrum Connect, can be allocated to a service by selecting the service and clicking an unattached resource.

A storage service provides indication for the allocated and used storage space.

- Allocated – amount of storage space available on all pools connected to the service.
- Used – amount of storage space used by all storage elements connected to the service (datastores, servers, etc).

Important:

- When attaching a resource to an over-provisioned service, verify that a storage system which hosts the pool has enough capacity to accommodate the service space requirements.
 - Datastores created on VVol-enabled services always display the meta and thick pools of its group pool to be 100% full. You can disregard this alert.
-

6. You can display detailed information for all storage resources defined on a storage system by switching to the table view. The **System Storage Resources** table is displayed.











Storage Resources							
Search for a resource or storage system name							
RESOURCE	DOMAIN	STORAGE AR...	USAGE	SIZE(...)	SERVICE	TYPE	
18	dana-domain	XIV Gen3P1-22b	100% 	0		XIV	
21b	dana-not-managed	XIV Gen4d-67d	33% 	577.59		FlashSystem ..	
31d-ayab-pool		XIV hostdev31d	0% 	112.18		XIV	
51-ayab-Pool-dev31d		XIV hostdev31d	0% 	112.18		XIV	
A	dana-domain	XIV Gen3P1-22b	100% 	0		XIV	
AccountingPool1	sra_demo	XIV Gen4d-67d	60% 	9337.76		FlashSystem ..	
Mike_new_pool	mikhail_domain	XIV Gen4d-67c	100% 	96.26		FlashSystem ..	
Mikhail_VWC_test_p..	mikhail_domain	XIV Gen4d-67c	66% 	288.79		FlashSystem ..	
Mikhail_mirroring	mikhail_domain	XIV Gen4d-67c	66% 	288.79		FlashSystem ..	
Mikhail_mirroring	mikhail_domain	XIV Gen4d-67d	66% 	288.79		FlashSystem ..	

Figure 44. Storage Resources table

The table lists all storage resources defined on the storage system, including their sizes, free and used storage space, parent pool for storage systems that run IBM Spectrum Virtualize, and a service that the resource is attached to. When you click on a storage resource row to select it, you can complete the following:

- Click the **Attach/Delegate** button to attach the storage resource to a previously selected storage service.
 - Click the **Detach/Cancel Delegation** button to detach the storage resource from a previously selected storage service.
 - Right-click and select **Modify** to resize the resource or remove it.
 - Right-click and select **Delete** button to delete child pools or Spectrum Accelerate pools. This option does not exist for parent pools or DS8000 pools.
7. You can display detailed information for all storage resources attached to a service on a storage system by right-clicking the service on which you want zoom, and then selecting **View Resources**. The **Service Storage Resources** table is displayed.

Name	Size (GiB)	Free (GiB)	Used (GiB)	Storage system
ilyak_paired_pool-ilyak_paired_pool	288.79	96.26	192.53	XIV Gen4d-67d
ilyak_paired_pool-ilyak_paired_pool	288.79	96.26	192.53	XIV Gen4d-67d

Figure 45. Service Storage Resources table

The table lists all storage resources attached to the storage service, including their sizes, free and used storage space and a system that the resource was created on. When you click on a storage resource row to select it, you can complete the following:

- Click the **Resource Detach** button to detach the storage resource from a storage service.
- Click the **Resource Settings** button to resize the resource or remove it.
- Click the **Remove** button to delete child pools or Spectrum Accelerate pools. This option is not functional for parent pools or DS8000 pools.

Defining and attaching replicated storage resources

For mirrored storage systems and services, you can define primary and secondary storage resources to be used for data replication.

Before you begin

Verify that:

- Primary and secondary storage systems were added to Spectrum Connect.
- A storage service to use the replicated storage resources has the replication capability enabled.
- A target has been defined on a storage system for the management domain to be used by the paired pool. The storage system credentials must be valid for both primary and secondary pool domains.

About this task

You can define paired storage resource (pool), which contains a primary (source) pool and secondary (target) pool. Then, the paired pool is attached to a replicated service. This process will create a storage volume and its mirror on the target pool. You can use the IBM Storage Enhancements for VMware vSphere Web Client to provision the volume as a part of a consistency group. The primary and secondary

pools reside on the source and target storage systems, which can be located at different sites.

Procedure

1. On the **Spaces/Storage Services** pane, select the storage space from which you want to choose storage services. The available services that reside on the selected storage space are immediately displayed.
2. Click on a service with replication capabilities to which you want to attach the resources. The service color changes to blue to indicate selection.
3. On the Storage Systems pane, right-click a compatible storage system on which you want to define a new resource, and then select **Add Resource**. The Add New Resource dialog box is displayed.

You can choose to create a new paired resource, or select existing ones.

The screenshot shows the 'Add New Resource' dialog box with the following fields and options:

- Resource Setting:**
 - Name: [Text Input]
 - Resource Pool: Create new, Choose existing (with info icon)
 - Size: [Text Input] GiB
 - Snapshot reserve: [Text Input] 15 %
 - Type: Paired Pool
- Primary Resource:**
 - Array: [Dropdown Menu]
 - Domain: [Dropdown Menu]
- Secondary Resource:**
 - Array: [Dropdown Menu]
 - Domain: [Dropdown Menu]

Buttons: Cancel, Add

Figure 46. Add New Resource for replicated service dialog box, Create New option




Figure 47. Add New Resource for replicated service dialog box, Choose Existing option

4. In the **Add New Resource** dialog box, enter configuration parameters for a new resource (paired pool). See table below for details.

Table 10. Paired pool parameters

Parameter	Description and values
Name	Alphanumeric string for identification of the replication relationship. The name must not contain spaces.
Create new	New primary and secondary storage resources to be created for the paired pool.
Choose existing	Existing primary and secondary storage resources to be selected for the paired pool.
Array	A storage system to include the primary or secondary storage resource.
Domain	Management domain of the pool on the storage systems, Configuration considerations To use management domains, verify that: <ul style="list-style-type: none"> • Domains are already defined on your storage system. • Default credentials of a storage admin user are assigned to the domain.
Size	Storage resource size in GiB. The size of the replicated pool cannot be larger than the size of the smallest pool of the paired resources (primary or secondary). When you select existing pools for replication, the pool size cannot be changed.

Table 10. Paired pool parameters (continued)

Parameter	Description and values
Over-provisioning	<p>Percentage of over-provisioned storage space on the service, defining the ratio between logical and physical storage capacity on XIV storage systems. For example, when configured to 100, it sets a 1:1 ratio between the two values, while a value of 200 defines the logical capacity (soft size) to be twice the physical capacity (hard size).</p> <p>Configuration considerations</p> <p>The recommended over-provisioning value for a service with IBM Real-time Compression is 200%.</p> <p>The XIV utilizes thin provisioning for all VM configuration volumes. When creating a regular VM on an XIV disk, you can choose the provisioning mode, according to your application requirements.</p> <p>For other cases, you can choose any value:</p> <ul style="list-style-type: none"> • 100%, if you do not want to utilize thin provisioning. • Above 100%, if you want to take a risk of allocating resources that the XIV storage system may not have. <p>In addition, when utilizing the over-provisioning, XIV storage systems must be also configured to allow support for this feature. When used with the XIV domains, the domain must be also configured with matching soft and hard capacity settings.</p> <p>Important</p> <p>The managed domain that you created cannot be used for traditional volumes without virtualization. You must create a separate regular domain for them. This domain must have the same user and the ESXi hosts that you intend to manage. However, you need to create a separate storage resource and a new service on the regular domain via Spectrum Connect for subsequent use by the vWC.</p> <p>The over-provisioning is relevant for XIV storage systems only.</p> <p>Configuration examples</p> <ul style="list-style-type: none"> • Storage pools in a regular service with over-provisioning at 100% are created as thick pools. • Storage pools in a regular service with over-provisioning at 150% are created as thin pools. If a hard pool size is 500 GiB, its soft size is 750 GiB.
Snapshot reserve	Percentage of storage space on the service reserved for snapshots.

5. Click **Add**. The storage resource is created on the storage system and attached to the service. The resource color changes to blue to indicate the successful attachment.

You can check pool pairing by hovering your mouse pointer over a primary or secondary storage resource. The paired pool colors change to purple and a coupling line is displayed to indicate the connection.

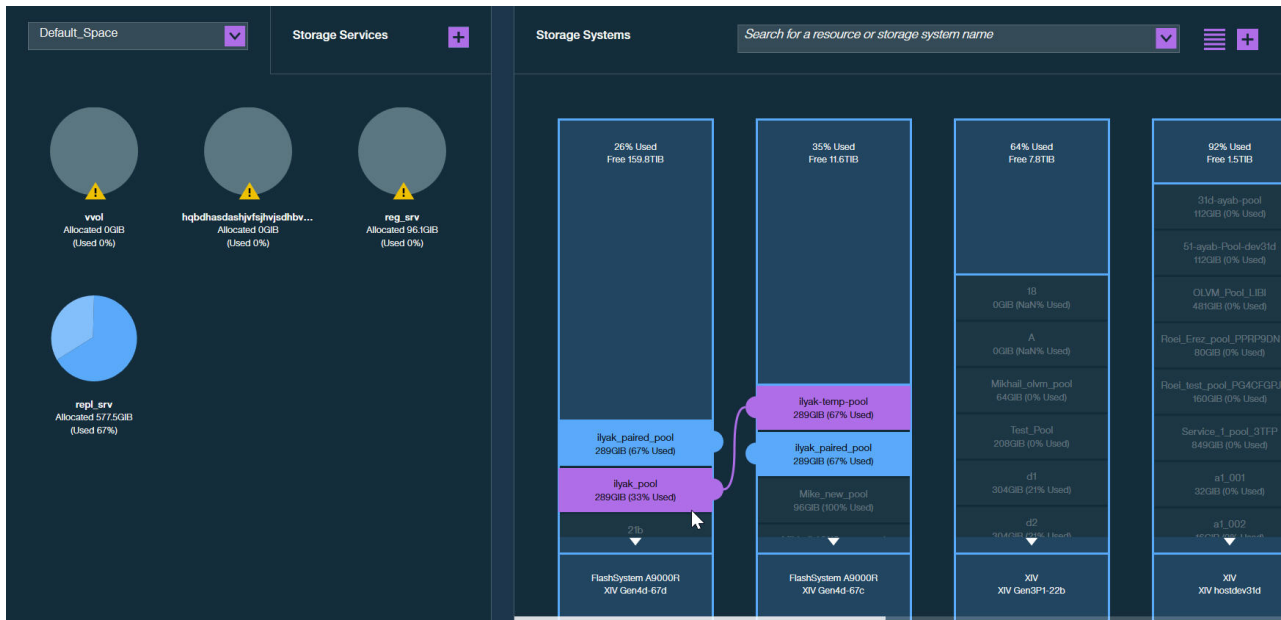


Figure 48. Paired pools

The resources defined on the system before the system was attached to Spectrum Connect, can be allocated to a service by selecting the service and clicking an unattached resource.

A storage service provides indication for the allocated and used storage space.

- Allocated – amount of storage space available on all pools connected to the service.
- Used – amount of storage space used by all storage elements connected to the service (datastores, servers, etc).

Important: When attaching a resource to an over-provisioned service, verify that a storage system which hosts the pool has enough capacity to accommodate the service space requirements.

6. You can display detailed information for all storage resources defined on a storage system by switching to the table view. The **System Storage Resources** table is displayed.











Storage Resources							
Search for a resource or storage system name							
RESOURCE ▲	DOMAIN	STORAGE AR...	USAGE	SIZE(...)	SERVICE	TYPE	
18	dana-domain	XIV Gen3P1-22b	100% 	0		XIV	
21b	dana-not-managed	XIV Gen4d-67d	33% 	577.59		FlashSystem ..	
31d-ayab-pool		XIV hostdev31d	0% 	112.18		XIV	
51-ayab-Pool-dev31d		XIV hostdev31d	0% 	112.18		XIV	
A	dana-domain	XIV Gen3P1-22b	100% 	0		XIV	
AccountingPool1	sra_demo	XIV Gen4d-67d	60% 	9337.76		FlashSystem ..	
Mike_new_pool	mikhail_domain	XIV Gen4d-67c	100% 	96.26		FlashSystem ..	
Mikhail_VWC_test_p..	mikhail_domain	XIV Gen4d-67c	66% 	288.79		FlashSystem ..	
Mikhail_mirroring	mikhail_domain	XIV Gen4d-67c	66% 	288.79		FlashSystem ..	
Mikhail_mirroring	mikhail_domain	XIV Gen4d-67d	66% 	288.79		FlashSystem ..	

Figure 49. Storage Resources table

The table lists all storage resources defined on the storage system, including their sizes, free and used storage space, and a service that the resource is attached to. When you click on a storage resource row to select it, you can complete the following:

- Click the **Attach/Delegate** button to attach the storage resource to a previously selected storage service.
 - Click the **Detach/Cancel Delegation** button to detach the storage resource from a previously selected storage service.
 - Right-click and select **Modify** to resize the resource or remove it.
 - Right-click and select **Unlink** to break the replication connection between the storage resources of the paired pool. The primary and secondary pools themselves remain intact.
7. You can display detailed information for all storage resources attached to a service on a storage system by right-clicking the service on which you want zoom, and then selecting **View Resources**. The **Service Storage Resources** table is displayed.

Name	Size (GiB)	Free (GiB)	Used (GiB)	Storage system
ilyak_paired_pool-ilyak_paired_pool	288.79	96.26	192.53	XIV Gen4d-67d
ilyak_paired_pool-ilyak_paired_pool	288.79	96.26	192.53	XIV Gen4d-67d

Figure 50. Service Storage Resources table

The table lists all storage resources attached to the storage service, including their sizes, free and used storage space and a system that the resource was created on. When you click on a storage resource row to select it, you can complete the following:

- Click the **Resource Detach** button to detach the storage resource from a storage service.
- Click the **Resource Settings** button to resize the resource or remove it.
- Click the **Remove** button to delete a pool.

Modifying storage resources

If needed, you can modify any resource (pool) defined on a storage system.

About this task

Storage resources can be renamed and resized to match the requirements of the connected services.

Note: Pool resizing does not function for parent or DS8000 pools.

Procedure

1. On the **Storage Systems** pane, right-click a pool which you intend to edit and select **Modify** to edit the pool settings or **Unlink** to unlink replicated pools. The **Edit Resource** dialog box is displayed.

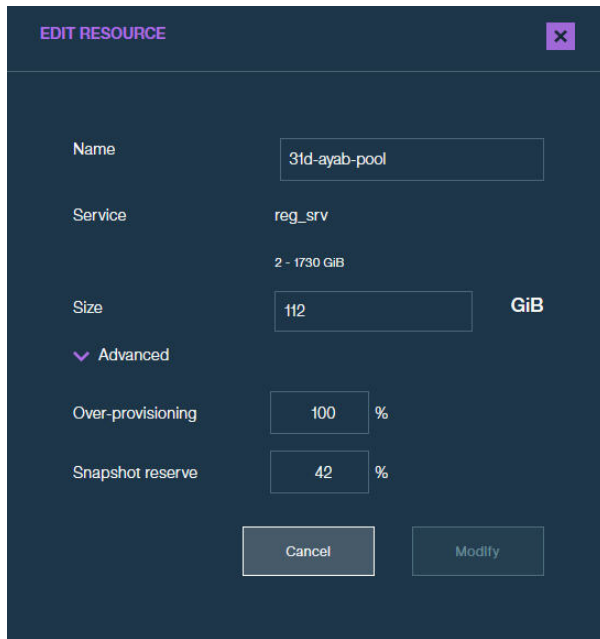


Figure 51. Edit Resource dialog box

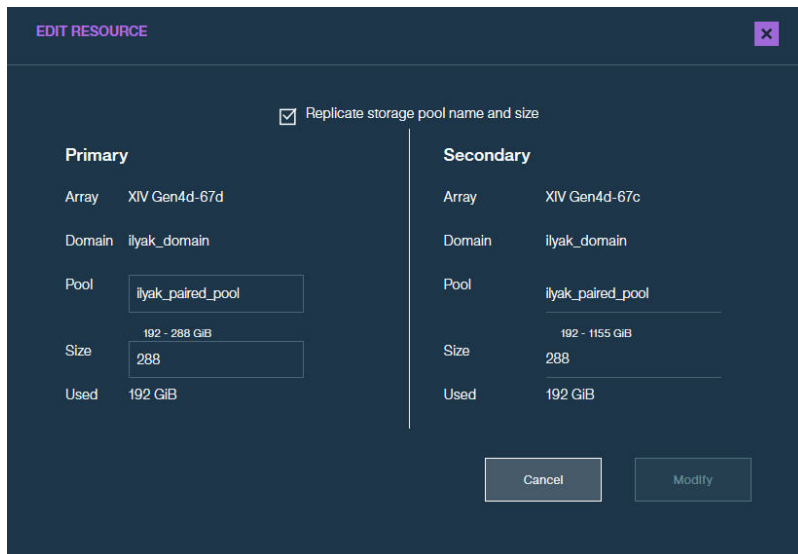


Figure 52. Edit Settings for replicated pools dialog box

2. Configure the following:

- In the **Name** text box, enter the new name for the pool.
- In the **Size** text box, enter the new size for the pool.
- For replicated pools, select **Replicate storage pool and size** to copy the primary pool name and size to the secondary pool in the mirroring relationship.

3. Click **Modify** to save the change.

Note: You can also access the **Resource Settings** dialog box by selecting a storage resource row in the **System Storage Resource** or **Service Storage Resource** table. Refer to “Defining and attaching storage resources” on page 60 for details.

Modifying storage resource attachments

When required, you can modify a storage resource (pool) service attachments.

About this task

Storage resources can be detached from the service or attached to it.

- A detached storage pool can no longer be managed by the included solution components.
- A detached storage pool becomes available for allocation by other existing services.
- After the detachment, you can attach the pool again to fully restore its management.

Procedure

1. On the **Spaces/Storage Services** pane, right-click a storage service and select **Manage Resources**. The pools that are currently attached to the service and those that can be attached to it are highlighted on the **Storage Systems** pane.

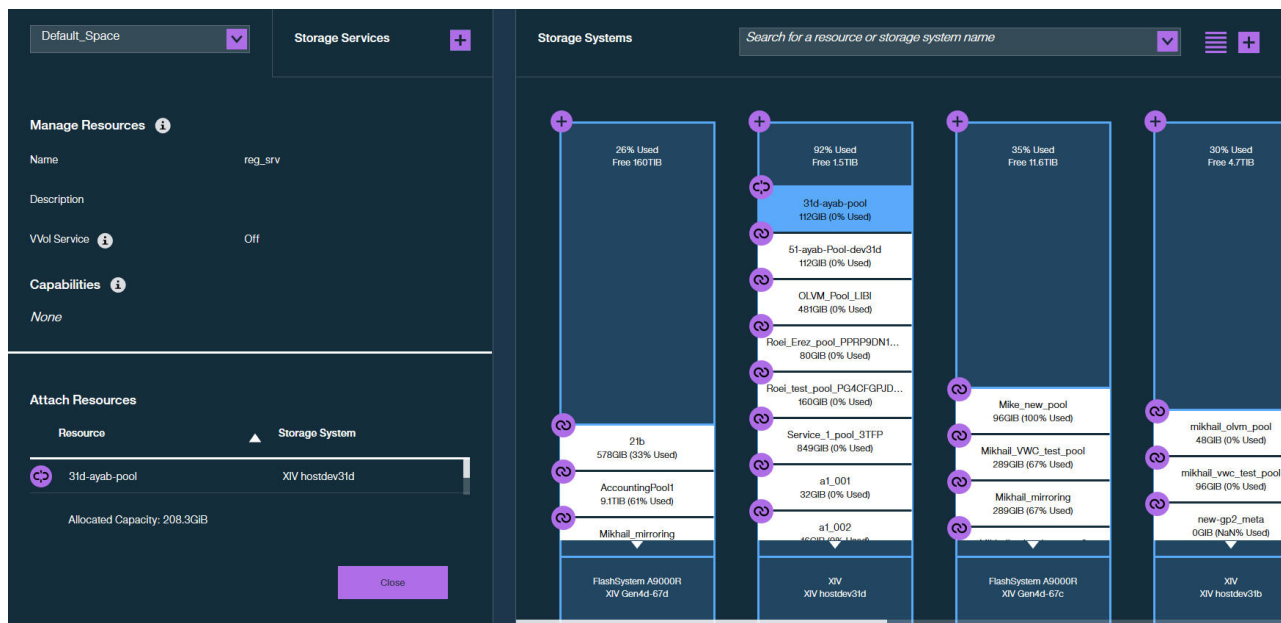


Figure 53. Manage Resources dialog box

2. From the **Storage Systems** pane:
 - Click the **Attach/Delegate** button to attach the storage resource to the storage service. The pool color changes to indicate the successful attachment.
 - Click the **Detach/Cancel Delegation** button to detach the storage resource from the storage service. The pool color changes to indicate the successful detachment.

Administration

This chapter details common administrative tasks that can be performed when using the IBM Spectrum Connect.

See the following sections for more information:

- “Checking and controlling the Spectrum Connect Linux service” on page 74

- “Checking and modifying the configuration files”
- “Changing the Spectrum Connect communication port” on page 77

Checking and controlling the Spectrum Connect Linux service

At any time, you can check whether the Spectrum Connect service runs properly on the Linux host. You can also stop and then start the service if needed.

Procedure

1. Log on to the Linux command prompt environment as a root user.

Important: Only root users can complete service operations.

2. Enter the following command: **service ibm_spectrum_connect status**. The status of the Celery and Django services is displayed.

```
[root@ibmsc]# service ibm_spectrum_connect status
Celery services are running...
Django service is running...
vWC refresh service is running...
```

What to do next

If you want to stop the Spectrum Connect service, use the **stop** command:

```
[root@ibmsc]# service ibm_spectrum_connect stop
Stopping ibm_spectrum_connect (via systemctl): [ OK ]
```

To start the service again, use the **start** command:

```
[root@ibmsc]# service ibm_spectrum_connect start
Starting ibm_spectrum_connect (via systemctl): [ OK ]
```

If you want to stop and then start the Spectrum Connect service in one command, use the **restart** command:

```
[root@ibmsc]# service ibm_spectrum_connect restart
Restarting ibm_spectrum_connect (via systemctl): [ OK ]
```

Checking and modifying the configuration files

Spectrum Connect has several configuration files that store configuration settings that you can change manually if needed.

You can view and modify the contents of each file with any standard text editor, according to the purpose and contents of each file.

After modifying a configuration file, you must restart the Spectrum Connect service by running the following CLI command on the Linux host:

```
service ibm_spectrum_connect restart
```

For more information about this service, see “Checking and controlling the Spectrum Connect Linux service.”

Table 11. Configuration files

File name	Directory location	Purpose or relevant parameters
ibmsyslog.conf	/opt/ibm/ibm_spectrum_connect/conf.d/	Defines the logging standard per application, as well as the target of the log messages.
ibmlogs-rotate	/opt/ibm/ibm_spectrum_connect/conf.d/	Controls the archive and renewal timing attributes of the log files. For the list of log files, see “Checking the log files” on page 93.
ldap.ini	/opt/ibm/ibm_spectrum_connect/conf.d/	See “Configuring LDAP-based directory user access” on page 21.
ldap.conf	/etc/openldap/	See “Configuring LDAP-based directory user access” on page 21.
vasa_config.ini	/opt/ibm/ibm_spectrum_connect/conf.d/vasal/	<ul style="list-style-type: none"> • populate_vasa_events_and_alarms – The time interval in minutes between each operation of filtering relevant events for each connected vCenter server. The default value is 2.
hsgsvr_config.ini	/opt/ibm/ibm_spectrum_connect/conf.d/hsgsvr/	<ul style="list-style-type: none"> • populate_arrays_and_events – The time interval in minutes between each update of information (changes and events) received from each monitored storage system. The default value is 10. When using a large number of storage systems, the value might need to be higher than 10. See “Working with multiple storage systems” on page 224 for additional information.
vcops_config.ini	/opt/ibm/ibm_spectrum_connect/conf.d/vcops	<p>See “Adjusting system update interval.”</p> <p>See “Configuring alarm reporting” on page 76.</p> <p>See “Configuring metrics scope” on page 76.</p> <p>See “Enabling SSL verification” on page 76.</p>

Adjusting system update interval

The IBM Storage adapter pushes the storage system information to the vRealize Operations Manager HTTP Post Adapter, using HTTP post requests.

About this task

By default, the update occurs every five minutes. The system information includes the following:

- Resource definition and all its relevant matrices
- Relationship between the storage resources
- Relationship between storage volumes and VMware datastores
- Storage system events

Procedure

To change the system update interval:

In the `vcops_config.ini` file, change the **vcops_push_interval** parameter to a desired value in minutes.

Configuring alarm reporting

System events are relayed to the vRealize Operations Manager via the IBM Storage adapter.

About this task

By default, the IBM Storage adapter reports only immediate and critical events to the vRealize Operations Manager. You can select a lowest severity level, instructing the IBM Storage adapter to deliver events that are equal or above the specified value. In addition, you can disable event reporting altogether.

Procedure

To configure alarm reporting:

In the `vcops_config.ini` file, set the **event_level** parameter to one of the following values:

- **none** – no events are reported
- **info** – all events are reported
- **warning** – warning, immediate and critical events are reported
- **immediate** – immediate and critical events are reported
- **critical** – only critical events are reported

Configuring metrics scope

The IBM Storage adapter relays the storage system metrics data to the vRealize Operations Manager.

About this task

You can adjust the scope of metrics data that is pushed by the IBM Storage adapter. By default, the detailed metrics are reported, but you can change the setting to deliver only summary of the performance counters.

Procedure

To change the scope of statistics data:

In the `vcops_config.ini` file, change the **push_detailed_statistics** parameter to True (detailed performance metrics) or False (performance metrics summary).

Enabling SSL verification

SSL protocol provides an encrypted communication link between the vROps server and the IBM Spectrum Connect.

About this task

To ensure a secure communication channel between the vROps server and IBM Spectrum Connect, you can enable the SSL certification, which is disabled by default. If you enable the SSL verification, make sure to provide a valid certificate via Linux.

Procedure

To enable SSL verification:

In the `vcops_config.ini` file, change the `verify_ssl_certificate` parameter to True (enable).

Changing the Spectrum Connect communication port

The TCP port used by Spectrum Connect to communicate with vCenter and vRO servers can be changed at any time.

About this task

By default, Spectrum Connect uses TCP ports 8443 or 8440 for communication with vCenter and vRO servers to access vWC, register as VASA storage provider and run orchestration workflows. The 8443 port is used if the current Spectrum Connect has been upgraded from a previous version. The 8440 port is used if the software package has been installed anew.

Procedure

To change the current Spectrum Connect communication port:

1. Verify that you have `root` access privileges.
2. Access the `/opt/ibm/ibm_spectrum_connect/bin/` directory.
3. Run the `sc_port_change.sh -p <new_port_number>` script, replacing `new_port_number` with the desired TCP port value. The script changes the port number and restarts the Nginx and Spectrum Connect services for the changes to take effect.

What to do next

After running the script:

- Re-register a previously registered Spectrum Connect as a VASA storage provider on the vCenter server.
- If you are using vRO, re-run the configuration script to instruct Spectrum Connect to use the new port.

Management from the command-line interface

You can access and control Spectrum Connect by using its command-line interface (CLI) functions locally from the Linux command prompt environment, or from a remote terminal connection.

The Spectrum Connect CLI is used for user and storage system management, as well as for integration of the cloud interfaces. However, the CLI application scope is limited. To employ the full functionality range of Spectrum Connect, use the web-based graphical user interface (GUI).

The following sections describe all the CLI configuration and management functions:

- “CLI – Switching to 'IBMSC' user mode” on page 78
- “CLI – Managing Spectrum Connect users” on page 78
- “CLI – Adding or removing storage system credentials” on page 81
- “CLI – Managing storage systems” on page 84
- “CLI – Setting the VASA credentials” on page 87
- “CLI – Managing integration with vRealize Operations Manager” on page 87

CLI – Switching to 'IBMSC' user mode

To start configuring Spectrum Connect, you must be logged in as the **ibmsc** user in the Linux command prompt environment.

About this task

ibmsc is a user account that is automatically created after the installation, allowing you to carry out the Spectrum Connect CLI-based configuration and management operations.

If needed, you can set the password for accessing the **ibmsc** user account externally (for example, from a remote computer over SSH), as described in the following procedure.

Procedure

To set a password for the **ibmsc** user:

1. Log in to the Linux command prompt environment as a root user.
2. Enter **passwd ibmsc** and then enter the password for the user account:

```
[root]# passwd ibmsc
Changing password for user ibmsc.
New password: *****
Retype new password: *****
passwd: all authentication tokens updated successfully.
[root]#
```

What to do next

To switch to the IBMSC user, enter the **su - ibmsc** command:

```
[root]# su - ibmsc
```

CLI – Managing Spectrum Connect users

All user accounts that can be used to access the Spectrum Connect must be individually defined.

You can either define (add) a single Spectrum Connect user account, or define multiple user accounts to be used separately.

Use the **sc_users** CLI command to add, delete, or display user accounts, and also to change the password of any specific account. Use the required argument after the command, as specified in the following table. In addition, you can configure the password reuse policy and security timeout, using the **sc_setting** CLI command.

Note:

- All CLI command arguments are case-sensitive.
 - The same operations are available from the GUI as well, as explained in “Managing Spectrum Connect users” on page 40.
 - The **sc_users** utility cannot define or affect external directory users. For more information about how to configure directory user access, see “Configuring LDAP-based directory user access” on page 21.
-

Table 12. Arguments for `sc_users`

Argument	Use after <code>sc_users</code> to:
<p><code>add -n <username> -p <password></code></p> <p>or</p> <p><code>add --user_name <username> --user_password <password></code></p>	<p>Add the username and password of the user that may access Spectrum Connect. You can add more than one user. The minimum password length is seven characters and it must include at least one letter and one digit.</p> <p>For example:</p> <pre>sc_users add -n johnvc -p *****</pre>
<p><code>change_password -n <username> -p <new password></code></p> <p>or</p> <p><code>change_password --user_name <username> --user_password <password></code></p>	<p>Change the password of a user account that was already added. The password that you type for the specified username is set as the new password. The minimum password length is seven characters and it must include at least one letter and one digit.</p> <p>For example:</p> <pre>sc_users change_password -n johnvc -p *****</pre>
<p><code>delete -n <username></code></p> <p>or</p> <p><code>delete --user_name <username></code></p>	<p>Delete a user account from the server.</p> <p>For example:</p> <pre>sc_users delete -n johnvc</pre>
<p><code>list</code></p>	<p>Display the names of currently defined user accounts.</p> <p>For example:</p> <pre>sc_users list User list: john_vc zivka1_vc lihit_vc</pre>
<p><code>-h</code></p> <p>or <code>--help</code></p>	<p>Display help information that is relevant to <code>sc_users</code>.</p> <p>You can also display help for the <code>add_user</code>, <code>change_password</code>, or <code>delete_user</code> argument if it is typed on the command line as well.</p>

Table 13. User-related arguments for `sc_setting`

Argument	Use after <code>sc_setting</code> to:
<p><code>modify -n USER_PASSWORD_HISTORY_LEN -v <password retention number></code></p>	<p>Prevent the user to submit a new password that is the same as any of the prior passwords for that account.</p> <p>For example, to prevent the user to submit a new password that is the same as the last four prior passwords, enter:</p> <pre>sc_setting modify -n USER_PASSWORD_HISTORY_LEN -v 4</pre>

Table 13. User-related arguments for `sc_setting` (continued)

Argument	Use after <code>sc_setting</code> to:
<code>modify -n TOKEN_INACTIVITY_TIMEOUT -v <inactivity timeout in minutes></code>	<p>Define a time period in minutes after which the Spectrum Connect GUI management session is terminated, if no user input is detected. By default, the timeout is set to 15 minutes.</p> <p>For example, to set the inactivity timeout to 20 minutes, enter:</p> <pre>sc_setting modify -n TOKEN_INACTIVITY_TIMEOUT -v 20</pre>

CLI – Managing server certificates

During the installation, a self-signed Secure Sockets Layer (SSL) certificate is generated to create a secure communication channel for servers and clients. If you already have a trusted certificate that you want to use, you can replace the self-signed certificate with an existing trusted certificate or generate a new certificate.

A self-signed certificate file, `vp.crt`, and a certificate key file, `vp.key`, are stored in the following directory:

```
/opt/ibm/ibm_spectrum_connect/settings/ssl_cert.
```

Because the self-signed certificate is not automatically recognized by the web browser that you use to log in to Spectrum Connect, you might encounter a connection security warning before you can access the Spectrum Connect login page (see “Logging in” on page 25).

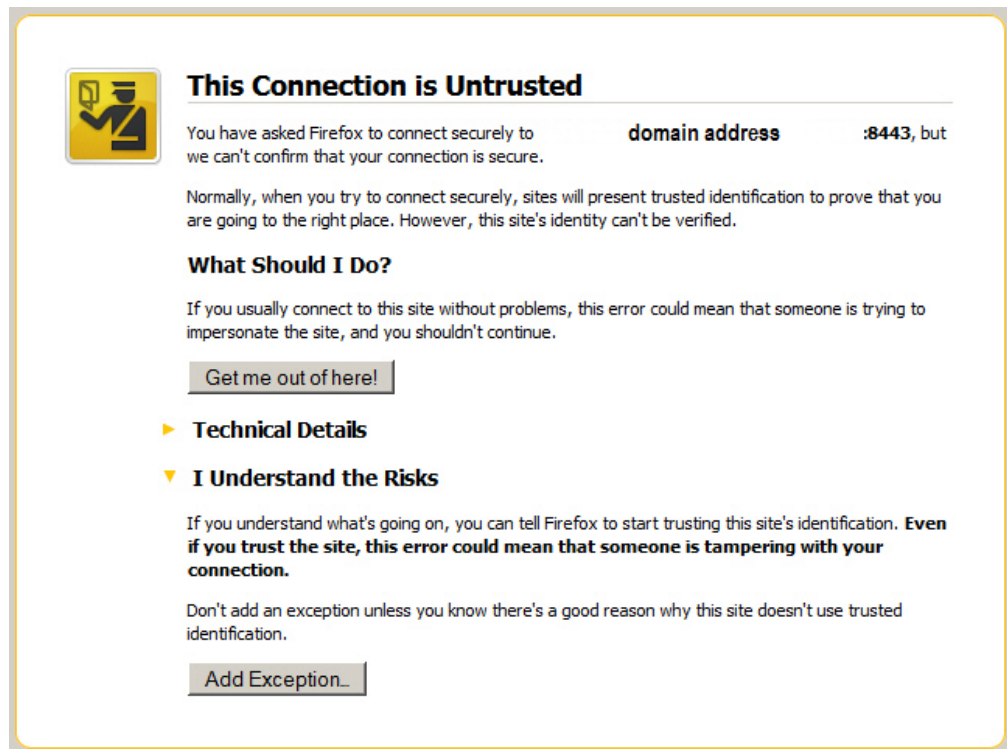


Figure 54. Connection security warning in the Mozilla FireFox web browser

To avoid such warning messages, use the **import** option of the **sc_ssl** command to upload a server certificate which is signed by a public certificate authority (CA), such as VeriSign, or by a CA whose root certificate was imported to your web browser. In addition, you can use the other options of the **sc_ssl** command to generate or to trust an SSL certificate.

Note:

- All CLI command arguments are case-sensitive.
- The same operations are available from the GUI as well, as explained in “Managing server certificates” on page 37.

Table 14. Arguments for **sc_ssl**

Argument	Use after sc_ssl to:
generate -c <common_name> -n <host_name> -i <ip_address> -e <expiration_period>	<p>Enter the hostname, common name, IP address of the Spectrum Connect server and certificate validity period (in days).</p> <p>For example:</p> <pre>sc_ssl generate -c mycommonname -n "sc_serverhostname" -i 1.0.0.200 -e 5000</pre>
trust -c <certificate_path>	<p>Select an SSL certificate to be trusted, by providing a path to its location.</p> <p>For example:</p> <pre>sc_ssl trust -c CA_certificate.crt</pre>
import -c <certificate_path> -k <key_path>	<p>Import a SSL certificate and a key file, by providing paths to their locations.</p> <p>For example:</p> <pre>sc_ssl import -c self_signed_certificate.crt -k private_key.key</pre>
-h	<p>Display help information that is relevant to sc_ssl.</p> <p>You can also display help for the generate, trust, or import argument if it is typed on the command line as well.</p>

CLI – Adding or removing storage system credentials

This section explains how to set the credentials that will be used to connect to the IBM storage system, or systems, that your VMware platforms use for storage provisioning.

Important:

- An identical storage admin user account with identical credentials (the same username and password) must already be predefined on all the IBM storage systems that you intend to use. Spectrum Connect can use only **a single system management account** for accessing all the different storage systems that you use. For more information about how to define a storage admin account on your IBM storage systems, refer to the relevant storage system management tools documentation.
 - Setting the storage credentials on Spectrum Connect allows you to add the IBM storage systems on the next step.
 - If the storage system management account is defined on a directory server, see “Checking the format of directory-based storage system credentials” on page 95.
-

Use the **sc_storage_credentials** CLI command to set (add), remove, or display the current storage system access credentials that Spectrum Connect uses in order to access all the IBM storage systems. Use the required argument after the command, as specified in the following table.

Note:

- All CLI command arguments are case-sensitive.
 - The same operations are available from the GUI as well, as explained in “Entering the storage system credentials” on page 44.
-

Table 15. Arguments for `sc_storage_credentials`

Argument	Use after <code>sc_storage_credentials</code> to:
<p><code>set -u <storage system username></code> <code>-p <storage system password></code> <code>-f</code> <code>-a <user type></code> or <code>set --user <storage system username></code> <code>--password <storage system password></code> <code>--force</code> <code>--user_account <user type></code></p>	<p>Set the username and password for accessing all the relevant IBM storage systems, and specify whether the storage admin user is locally-defined on the storage system or on a directory server. For storage systems running Spectrum Virtualize, ensure that the credentials belong to a user account with <i>VASAProvider</i> role.</p> <p>For example, if the storage admin user is locally defined on the storage system, enter:</p> <pre>sc_storage_credentials set -u john21 -p ***** -a local</pre> <p>And if the storage admin user is defined on a directory server, enter:</p> <pre>sc_storage_credentials set -u john21 -p ***** -a directory</pre> <p>Attention: During regular operation, whenever a directory-based storage admin fails to log in (from the Spectrum Connect side) to any storage system that is in use, Spectrum Connect immediately locks the storage admin user account and all storage systems become inaccessible on the Spectrum Connect side. This is to prevent repeated login attempt failures after which the directory server blocks that user account. In such a case, use the <code>-f</code> or <code>--force</code> argument on the command line to unlock the storage admin account on the Spectrum Connect side, with either the same credentials or with updated credentials. For example:</p> <pre>sc_storage_credentials set -u john21 -p ***** -f -a directory</pre> <p>The equivalent action in the Spectrum Connect GUI is to update the account credentials, as described in “Entering the storage system credentials” on page 44.</p> <p>If the storage credentials are defined on a directory server, see “Checking the format of directory-based storage system credentials” on page 95.</p>
<p><code>remove</code></p>	<p>Delete the existing storage system user account definition from the server.</p> <p>For example:</p> <pre>sc_storage_credentials remove</pre> <p>Attention: If you already added storage systems to the server, deleting the user account disconnects all these storage systems.</p>

Table 15. Arguments for `sc_storage_credentials` (continued)

Argument	Use after <code>sc_storage_credentials</code> to:
<p>list</p>	<p>Display the username of the existing storage system user account definition.</p> <p>The following example shows the command output when the storage admin user account is defined locally on the storage system:</p> <pre data-bbox="662 401 1414 596"> sc_storage_credentials list Username Array Alias User Category Account Source ----- admin XIV hostdev31b storageadmin local admin XIV hostdev32a storageadmin local admin XIV hostdev31a storageadmin local </pre> <p>The following example shows the command output when the storage admin user account is defined on a directory server:</p> <pre data-bbox="662 709 1414 905"> sc_storage_credentials list Username Array Alias User Category Account Source ----- admin XIV hostdev31b storageadmin directory admin XIV hostdev32a storageadmin directory admin XIV hostdev31a storageadmin directory </pre>
<p>-h</p> <p>or --help</p>	<p>Display help information that is relevant to <code>sc_storage_credentials</code>.</p> <p>You can also display help for the <code>set</code> argument if it is typed on the command line as well.</p>

CLI – Managing storage systems

All IBM storage systems that provide storage resources to your VMware platforms must be defined as storage systems on the IBM Spectrum Connect.

Use the `sc_storage_array` CLI command to add, remove, configure or list these IBM storage systems (referred to as *arrays* in the command syntax and output). Use the required argument after the command, as specified in the following table. In addition, you can set a threshold for alerting the user, when a storage resource capacity is running low, using the `sc_setting` CLI command.

Important:

- IBM storage systems can be added only after the storage credentials are set, as explained in “CLI – Adding or removing storage system credentials” on page 81.
 - If you want to remove existing storage systems:
 - A removed storage system, along with its storage pools and volumes, can no longer be managed by the included solution components (see “Included cloud interfaces” on page 2).
 - If the removed storage system contains working storage pools and volumes, the information of these storage pools and volumes is no longer displayed in vSphere Web Client. However, **vSphere data access and service level for these storage pools and volumes is not affected**. In addition, the removed system and its storage pools and volumes can be managed from the standard IBM storage system management tools.
 - After the removal, you can add the storage system back again to fully restore its management.
-

Note:

- All CLI command arguments are case-sensitive.
 - The same operations are available from the GUI as well, as explained in “Managing storage systems” on page 43.
-

Table 16. Arguments for `sc_storage_array`

Argument	Use after <code>sc_storage_array</code> to:
<code>add -i <management IP address> -a <system alias name></code> or <code>add --mgmt_ip <management IP address> --storage_type <storage system type> --alias <system alias name></code>	Add a storage system specified by an IP address or DNS. For example: <pre>sc_storage_array add -i 10.100.155.200</pre> Optional: you can define an alias for the added storage system, by adding <code>-a <alias name></code> or <code>--alias <alias name></code> to the command. For example: <pre>sc_storage_array add -i 10.100.155.200 -a mystorage1</pre> Note: If you choose not to define an alias, the alias that is already defined (if one was defined) on the IBM storage system side is automatically assigned as the alias.
<code>remove -a <system alias name></code> or <code>remove --alias <system alias name></code>	Remove a storage system specified by its alias (alias that was given by you or was automatically assigned). For example: <pre>sc_storage_array remove -a mystorage1</pre>

Table 16. Arguments for `sc_storage_array` (continued)

Argument	Use after <code>sc_storage_array</code> to:
<code>list</code>	<p>List the names and details of all the IBM storage systems that are currently added (and were not removed). The displayed information includes:</p> <ul style="list-style-type: none"> • Array alias • Array identifier • Management IP address • Elapsed time since last update • Connected • Notes <p>See the example after this table.</p>
<code>full_metadata_recovery -a <system alias name></code>	<p>Perform full VVol metadata restoration on a secondary Spectrum Connect in a high-availability (HA) group. The metadata disparity might occur during recovery in the HA group due to timestamp inconsistencies. The timestamp inconsistencies are possible on a storage system, when its internal clock is changed to comply with time zone settings or daylight saving time requirements.</p>
<code>refresh</code>	<p>Refresh information about all managed IBM storage systems.</p>
<code>-h</code> or <code>--help</code>	<p>Display help information that is relevant to <code>sc_storage_array</code>.</p> <p>You can also display help for the <code>add</code> or <code>remove</code> argument if it is typed on the command line as well.</p>

The following example shows the displayed information and format of the `sc_storage_array list` command output.

```
sc_storage_array list
```

Array Alias	Array Identifier	Management IP Addresses	Elapsed time since last update	Connected	Notes
array1000	2810-114-MN65026	9.100.150.155	16 minutes	False	Failed to log in to array 2810 with the provided credentials.
array2000	2810-114-MN65027	9.200.155.155	27 minutes	True	

Table 17. Storage resource-related arguments for `sc_setting`

Argument	Use after <code>sc_setting</code> to:
<code>modify -n RESOURCE_CAPACITY_THRESHOLD -v <capacity threshold in %></code>	<p>Define a threshold for a VVol-based storage resource capacity. When this threshold is reached, an alarm is generated to indicate that the storage capacity is running low. By default, the threshold is set at 70%.</p> <p>For example, to alert the user, when the storage resource has reached 80% of its capacity, enter:</p> <pre>sc_setting modify -n RESOURCE_CAPACITY_THRESHOLD -v 80</pre>

CLI – Setting the VASA credentials

The VASA credentials comprise a user name and a password that VMware vCenter servers can use to connect to the IBM Spectrum Connect and employ VMware vSphere APIs for Storage Awareness (VASA) functions.

Use the `sc_vasa_admin` CLI command to set or display the VASA credentials. vCenter servers can then use these credentials to connect to Spectrum Connect and utilize VASA functions, as explained in “Registering Spectrum Connect as a storage provider on vCenter server” on page 103. Use the required argument after the command, as specified in the following table.

Note:

- Only one set of a username and a password can be used for the VASA credentials, which applies to all vCenter servers that require VASA functions.
 - All CLI command arguments are case-sensitive.
 - The same operations are available from the GUI as well, as explained in “Setting the VASA credentials” on page 101.
-

Table 18. Arguments for `sc_vasa_admin`

Argument	Use after <code>sc_vasa_admin</code> to:
<code>set_secret -n <username> -p <password></code> or <code>set_secret --user_name <username></code> <code>--user_password <password></code>	Set the username and password that the VASA credentials should comprise. For example: <pre>sc_vasa_admin set_secret -n johnvasa -p ***** The secret key for the VASA Provider has been set successfully.</pre>
<code>list_secret</code>	Display the username of the currently defined VASA credentials. For example: <pre>sc_vasa_admin list_secret Secret key username for the VASA Provider: johnvasa</pre>
<code>-h</code> or <code>--help</code>	Display help information that is relevant to <code>sc_vasa_admin</code> . You can also display help for the <code>set_secret</code> argument if it is typed on the command line as well.

CLI – Managing integration with vRealize Operations Manager

Before you can use the IBM Storage Management Pack for VMware vRealize Operations Manager, you must set a connection to at least one vRealize Operations Manager (vROps) server, and then define which storage systems should be monitored in vROps.

After a vROps server connection is defined and storage systems (referred to as *arrays* in the command syntax and output) are associated with the vROps server, detailed monitoring information for these storage systems becomes available in vROps (for more information, see Chapter 6, “IBM Storage Management Pack for VMware vRealize Operations Manager,” on page 151).

Use the **sc_vrops_server** CLI command to add, remove, or disable connections to vROps servers, or to list the current server connections. An HTTP POST adapter is automatically created for each vROps server that you add.

Use the **sc_vrops_adapter** command to attach storage systems to any created HTTP POST adapter. A storage system that is attached to an HTTP POST adapter can be monitored by the vROps server for which the HTTP POST adapter was created. You can use the **sc_vrops_adapter** command for additional options as described below.

Use the required argument after each command, as specified in the following tables.

Important:

- Storage systems can be attached to HTTP POST adapters only after the storage systems have been added to Spectrum Connect, as explained in “CLI – Managing storage systems” on page 84.
 - If the IBM Storage PAK file was not deployed on the vROps server, IBM Storage monitoring information is not displayed with dedicated dashboards, graphic icons, and user-friendly attribute names in vROps. For information about how to deploy the IBM Storage PAK file, see “Downloading the vROps management package” on page 152.
 - If you want to detach storage systems:
 - A detached storage system can no longer be monitored through vRealize Operations Manager.
 - After the detachment, you can reattach the storage system to fully restore its monitoring through vRealize Operations Manager.
-

Note:

- All CLI command arguments are case-sensitive.
 - Apart from the ability to add more than one vROps server and additional HTTP POST adapters from the CLI, the same operations are available from the GUI as well, as explained in “Managing integration with vRealize Operations Manager” on page 151.
-

Table 19. Arguments for **sc_vrops_server**

Argument	Use after sc_vrops_server to:
add -n <hostname> -u <username> -p <password> or add --hostname <hostname> --username <username> --password <password>	Add a vROps server connection by specifying the following parameters on the command line: <ul style="list-style-type: none"> • IP address or hostname of the vROps server that you want to add (connect to). • Username for accessing the vROps server. • Password for accessing the vROps server. For example: <div style="border: 1px solid black; border-radius: 15px; padding: 5px; margin: 10px 0;"> <pre>sc_vrops_server add -n vrops1.domain -u john1 -p *****</pre> </div> <p>Note: The local HTTP POST adapter is added automatically after the vROps server is added.</p>

Table 19. Arguments for `sc_vrops_server` (continued)

Argument	Use after <code>sc_vrops_server</code> to:
<p><code>remove -n <hostname></code></p> <p>or</p> <p><code>remove --hostname <hostname></code></p>	<p>Remove a vROps server connection by specifying its IP address or hostname on the command line.</p> <p>For example:</p> <pre data-bbox="764 369 1446 428" style="border: 1px solid black; border-radius: 10px; padding: 5px;">sc_vrops_server remove -n vrops1.domain</pre> <p>Note: A removed vROps server ceases to receive monitoring information regarding IBM storage resources that are in use. You can add a vROps server back by using the add option (see above).</p>
<p><code>disable -n <hostname></code></p> <p>or</p> <p><code>disable --hostname <hostname></code></p>	<p>Disable reporting to a vROps server (without removing its connection) by specifying its IP address or hostname on the command line.</p> <p>For example:</p> <pre data-bbox="764 751 1446 810" style="border: 1px solid black; border-radius: 10px; padding: 5px;">sc_vrops_server disable -n vrops1.domain</pre> <p>Note: A vROps server for which reporting is disabled ceases to receive monitoring information regarding IBM storage resources that are in use. You can resume reporting to a vROps server by using the enable option (see below).</p>
<p><code>enable -n <hostname></code></p> <p>or</p> <p><code>enable --hostname <hostname></code></p>	<p>Resume reporting to a vROps server by specifying its IP address or hostname on the command line.</p> <p>For example:</p> <pre data-bbox="764 1108 1446 1167" style="border: 1px solid black; border-radius: 10px; padding: 5px;">sc_vrops_server enable -n vrops1.domain</pre>
<p><code>list</code></p>	<p>List all the currently connected (added) vROps servers and their operation status (enabled or disabled).</p> <p>For example:</p> <pre data-bbox="764 1318 1446 1493" style="border: 1px solid black; border-radius: 10px; padding: 5px;">sc_vrops_server list vROps Hostname Status ----- vrops1.domain Enabled vrops2.domain Disabled</pre>
<p><code>-h</code></p> <p>or <code>--help</code></p>	<p>Display help information that is relevant to <code>sc_vrops_server</code> .</p> <p>You can also display help for the add, remove, enable, or disable argument if it is typed on the command line as well.</p>

Table 20. Arguments for `sc_vrops_adapter`

Argument	Use after <code>sc_vrops_adapter</code> to:
<p><code>array_attach -a <alias> -n <hostname></code> or <code>array_attach --alias <alias> --hostname <hostname></code></p>	<p>Attach a storage system to a vROps HTTP POST adapter (of a vROps server) by specifying the following parameters on the command line:</p> <ul style="list-style-type: none"> • Alias name of the storage system that you want to attach. • IP address or hostname of the storage system that you want to attach. <p>For example:</p> <pre>sc_vrops_adapter array_attach -a myXIV -n 9.150.200.100</pre> <p>Note: The storage system must already be added to Spectrum Connect, as explained in “CLI – Managing storage systems” on page 84.</p>
<p><code>array_detach -a <alias> -n <hostname></code> or <code>array_detach --alias <alias> --hostname <hostname></code></p>	<p>Detach a storage system from a vROps HTTP POST adapter (of a vROps server) by specifying the following parameters on the command line:</p> <ul style="list-style-type: none"> • Alias name of the storage system that you want to detach. • IP address or hostname of the storage system that you want to detach. <p>For example:</p> <pre>sc_vrops_adapter array_detach -a myXIV -n 9.150.200.100</pre>
<p><code>add -s <server> -n <hostname></code> or <code>add --server <server> --hostname <hostname></code></p>	<p>Add an HTTP POST adapter for a vROps server by specifying the following parameters on the command line:</p> <ul style="list-style-type: none"> • IP address or hostname of a currently added (connected) vROps server. • IP address or hostname of the HTTP POST adapter that you want to add. <p>For example:</p> <pre>sc_vrops_adapter add -s vrops1.domain -n adapter2.domain</pre>
<p><code>remove -n <hostname></code> or <code>remove --hostname <hostname></code></p>	<p>Remove an HTTP POST adapter by specifying its IP address or hostname on the command line.</p> <p>For example:</p> <pre>sc_vrops_adapter remove -n adapter2.domain</pre>
<p><code>report_thresholds -n <hostname></code> or <code>report_thresholds --hostname <hostname></code></p>	<p>Report the thresholds to a vROps server by specifying its IP address or hostname on the command line.</p> <p>For example:</p> <pre>sc_vrops_adapter report_thresholds -n vrops1.domain</pre>

Table 20. Arguments for `sc_vrops_adapter` (continued)

Argument	Use after <code>sc_vrops_adapter</code> to:															
<code>report_thresholds -a</code> or <code>report_thresholds -all</code>	Report the thresholds to all vROps servers that are currently added (connected). Usually, this command can be omitted, because the thresholds are defined by default, when a vROps server is added. <pre>sc_vrops_adapter report_thresholds -a</pre>															
<code>list</code>	List all the currently defined HTTP POST adapters and display their associated vROps server, alias name of attached storage systems, last report time (report to the vROps server), and indication regarding whether the adapter is remote or locally defined. For example: <pre>sc_vrops_adapter list</pre> <table border="1"> <thead> <tr> <th>vROps Hostname</th> <th>HTTP POST Hostname</th> <th>Remote</th> <th>Last Reported</th> <th>Array alias</th> </tr> </thead> <tbody> <tr> <td>vrops1.domain</td> <td>adapter1.domain</td> <td>No</td> <td>5 minutes ago</td> <td>myXIV</td> </tr> <tr> <td>vrops2.domain</td> <td>adapter2.domain</td> <td>Yes</td> <td>8 minutes ago</td> <td>myXIV</td> </tr> </tbody> </table>	vROps Hostname	HTTP POST Hostname	Remote	Last Reported	Array alias	vrops1.domain	adapter1.domain	No	5 minutes ago	myXIV	vrops2.domain	adapter2.domain	Yes	8 minutes ago	myXIV
vROps Hostname	HTTP POST Hostname	Remote	Last Reported	Array alias												
vrops1.domain	adapter1.domain	No	5 minutes ago	myXIV												
vrops2.domain	adapter2.domain	Yes	8 minutes ago	myXIV												
<code>-h</code> or <code>--help</code>	Display help information that is relevant to <code>sc_vrops_adapter</code> . You can also display help for the <code>array_attach</code> , <code>array_detach</code> , <code>add</code> , <code>remove</code> , or <code>report_thresholds</code> argument if it is typed on the command line as well.															

CLI – Backing up or restoring Spectrum Connect configuration

At any point, you can back up the current Spectrum Connect configuration and save it to a file, or load a previously saved configuration to restore a configuration.

The configuration includes storage credentials, storage systems, vCenter credentials, and storage resource attachments.

Use the `sc_configuration` CLI command to save the existing configuration, or load a saved configuration to replace the existing one. Use the required argument after the command, as specified in the following table.

Note:

- All CLI command arguments are case-sensitive.
 - The backup and restore operations are not available on the Spectrum Connect GUI.
-

Attention:

- Before using the `restore` option, the IBM VASA Provider service must be stopped as explained in “Checking and controlling the Spectrum Connect Linux service” on page 74. Start the service again after the configuration has been loaded.
 - `restore` should be used only with a freshly installed Spectrum Connect that has not yet been configured.
-

Table 21. Arguments for `sc_configuration`

Argument	Use after <code>sc_configuration</code> to:
<code>backup -f <file name> -k <8 characters></code> or <code>backup --file <file name> --key <8 characters></code>	Save the current IBM Storage Provider configuration to the specified file using an AES-256 encryption key that comprises 8 characters. For example: <pre>sc_configuration backup -f confbackup -k abcdefghijklmnop</pre> <p>Important: You will need to provide this key in any restore operation (see below).</p>
<code>restore -f <file name> -k <8 characters></code> or <code>restore --file <file name> --key <8 characters></code>	Load a configuration from a specified file by providing the file name and the encryption key that was used in the creation of the file. <p>Attention: See the notes above this table.</p> For example: <pre>sc_configuration restore -f confgbackup -k abcdefghijklmnop</pre>
<code>restore -f ./vasa115exported.db</code>	Restore an existing Spectrum Connect configuration that was made with IBM Storage Provider for VMware VASA version 1.1.5 (applicable only to XIV systems). For this restore operation, you do not need to provide an encryption key: <pre>sc_configuration restore -f ./vasa115exported.db</pre>
<code>-h</code> or <code>--help</code>	Display help information that is relevant to <code>sc_configuration</code> . You can also display help for the <code>backup</code> or <code>restore</code> argument if it is typed on the command line as well.

Troubleshooting IBM Spectrum Connect

This chapter can help you detect and solve problems that you might encounter when using the IBM Spectrum Connect.

Note:

- For up-to-date information about known issues and possible workarounds, refer to the latest release notes.
- When contacting IBM Support, specify the storage system you are managing, using Spectrum Connect.

See the following sections for more information:

- “Checking the log files” on page 93.
- “Checking the format of directory-based storage system credentials” on page 95.
- “Configuring event forwarding” on page 96.
- “Deleting virtual volumes and group pools via XCLI” on page 97
- “Self-assist options for IBM Spectrum Connect” on page 97.

Checking the log files

The IBM Spectrum Connect maintains log files that record different types of events.

You can find the following log files in the `/var/log/sc/` directory:

- `events.log` – Records Spectrum Connect events according to their type: **Info**, **Error**, or **Warning**. The event logging is compatible with the Rsyslog application, an open source utility for forwarding log messages over IP networks (for more information, see the Rsyslog website). The following example shows different events that might be recorded:

```
IBMSC-0001, INFO, "User {user_name} has logged in."
IBMSC-0002, INFO, "User {user_name} has logged out."
IBMSC-0003, WARNING, "User {user_name} login attempt failed."
IBMSC-0004, INFO, "IBM Spectrum Connect local user account {user_name} was created."
IBMSC-0005, INFO, "IBM Spectrum Connect local user account {user_name} was deleted."
IBMSC-0006, INFO, "IBM Spectrum Connect local user account {user_name} password was reset."
IBMSC-0007, INFO, "Storage credentials were set for user {user_name}."
IBMSC-0008, ERROR, "Storage credentials for user {user_name} were disabled. Reason: {reason}."
IBMSC-0009, INFO, "Storage array identified as {identifier} with IP address {ip_address} has been added."
IBMSC-0010, WARNING, "Storage array {identifier} was removed."
IBMSC-0011, WARNING, "Storage array {identifier} was modified. Its new IP address is {ip_address}."
IBMSC-0014, INFO, "vCenter server with IP address {ip_address} was added by user {user_name}."
IBMSC-0015, WARNING, "vCenter server with IP address {ip_address} was removed."
IBMSC-0016, WARNING, "Credentials for vCenter server with IP address {ip_address} were updated by user {user_name}."
IBMSC-0017, INFO, "Storage pool {pool_name} on storage array {identifier} was attached to vCenter server with IP address {ip_address}."
IBMSC-0018, WARNING, "Storage pool {pool_name} on storage array {identifier} was detached from vCenter server with IP address {ip_address}."
IBMSC-0019, ERROR, "Failed to connect to storage array {identifier}. Reason: {reason}."
IBMSC-0020, INFO, "Information retrieval from storage array {identifier} was completed."
IBMSC-0021, ERROR, "Failed to retrieve information from storage array {identifier}. Reason: {reason}."
IBMSC-0022, INFO, "LDAP authentication was enabled."
IBMSC-0023, INFO, "LDAP authentication was disabled."
IBMSC-0024, INFO, "Completed the vSphere Web Client extension task {task_name} with the following parameters: {parameter_list}."
IBMSC-0025, ERROR, "Failed to complete the vSphere Web Client extension task {task_name} with the following parameters {parameter_list}."
Reason: {reason}."
```

- `hsgsvr.log` – Records events regarding monitoring and operations on storage systems and volumes.
- `vasa1.log` – Records events regarding the communication between Spectrum Connect and the connected vCenter servers that utilize VASA 1.0 functions.
- `vasa2.log` – Records events regarding the communication between Spectrum Connect and the connected vCenter servers that utilize VASA 2.0 functions. In addition, several events related to the VASA 1.0 activity may be recorded in the `vasa2.log` file as well.
- `vwc.log` – Records events regarding the communication between Spectrum Connect and the vSphere Web Client Server on which the IBM Storage Enhancements are installed.
- `celery.log` – Records events regarding the Celery services on the Linux host on which Spectrum Connect is installed.
- `django.log` – Records events regarding the Django service on the Linux host on which Spectrum Connect is installed.
- `vco.log` – Records events regarding operation of the vSphere Orchestrator, complementing information stored in the `hsgsvr.log` file. The `vco.log` file can be accessed via the **Log** tab of the vCO plug-in interface. If the Inventory folder within the 'IBM Storage' context is empty and the following message is stored in the `vco.log`:

```
[SCRepository] com.sun.jersey.api.client.ClientHandlerException:
javax.net.ssl.SSLHandshakeException: java.security.cert.CertificateException:
No name matching sc8.ps.xiv.ibm.com found,
```

replace the default Spectrum Connect SSL certificate and key files as described in “Managing server certificates” on page 37.

- `vcops.log` – Records events regarding the communication between Spectrum Connect and the connected vROps servers.
- `traffic.log` – Records XCLI events. This log file is reserved for debug purposes.

You can retrieve and save the current Spectrum Connect log files in a compressed TAR archive file by using the **Collect Log** option on the Settings menu of the Spectrum Connect GUI. This option allows you to save different instances of the log files at different times.

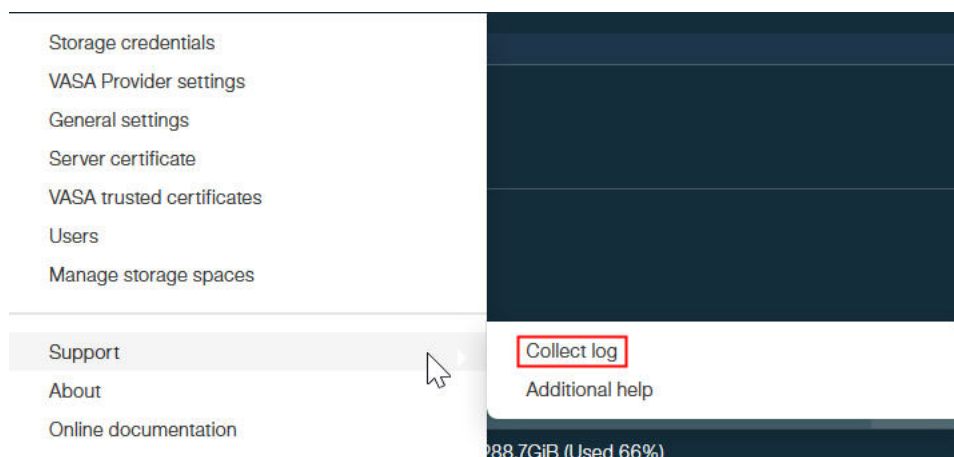


Figure 55. Collect log option

Additional information can be retrieved from other resources. For example, the **Failed to establish connection <hostname>** message stored in the `/var/log/vvold.log` file on an ESXi host indicates a loss of IP connectivity between the Spectrum Connect server and the ESXi host that uses a VVol. To resolve this issue, re-establish the connection with all ESXi host, consuming VVols via the Spectrum Connect web login port.

Also, vSphere components generate assorted logs that contain additional information about activities in vSphere environment. This information might help you resolve technical issues. The log files locations are detailed in the following table.

Note: The file locations refer to external resources and might change without prior notice. See VMware vSphere documentation for the current locations of the relevant system log files.

Table 22. VMware log file locations

VMware component	Operating system		GUI
	Windows	Appliance	
vSphere Web Client	C:\ProgramData\VMware\ vCenterServer\logs\ vsphere-client\logs	/var/log/vmware/vsphere- client/logs	–

Table 22. VMware log file locations (continued)

VMware component	Operating system		GUI
	Windows	Appliance	
vRo	6.x: <ul style="list-style-type: none"> • <instal_dir>\VMware\Orchestrator\app-server\logs, • <instal_dir>\VMware\Orchestrator\apps • <instal_dir>\VMware\Orchestrator\configuration\logs 	6.x, 7.x: /var/log/vco	<ul style="list-style-type: none"> • 6.x: Logs > Generate log report • 7.x: Control center > Export Logs
vROps	–	6.0.1: <ul style="list-style-type: none"> • storage/vcops/log/adapters/VinAdapter/VinAdapter_<id>.log • /storage/vcops/log/collector.log 	<ul style="list-style-type: none"> • 6.0.1: Administration > Support > Logs • 6.5: Administration > Support Bundles > + > Create log bundle > Download the bundle

Checking the format of directory-based storage system credentials

If you are using directory-based storage credentials for adding storage systems to the IBM Spectrum Connect (not for logging in to Spectrum Connect), you must verify that the directory user name is provided in the correct format.

Different formats are possible for a directory-based user name. For example:

- User name without the domain name, for example: john21
- User name with the domain name, for example: john21@domain_name

The format that should be used depends on the **directory user name attribute string** that is defined on the storage system. For example:

- **sAMAccountName** – User name without the domain name (john21).
- **userPrincipalName** – User name with the domain name (john21@domain_name).

Important: Other user name formats that are not specified above may be used. Consult with your directory server administrator about the required user name format, and make sure that the user name format is properly defined on the storage system.

The following figure shows the user name attribute in the LDAP role mapping definitions for XIV (defined via the XIV management GUI).

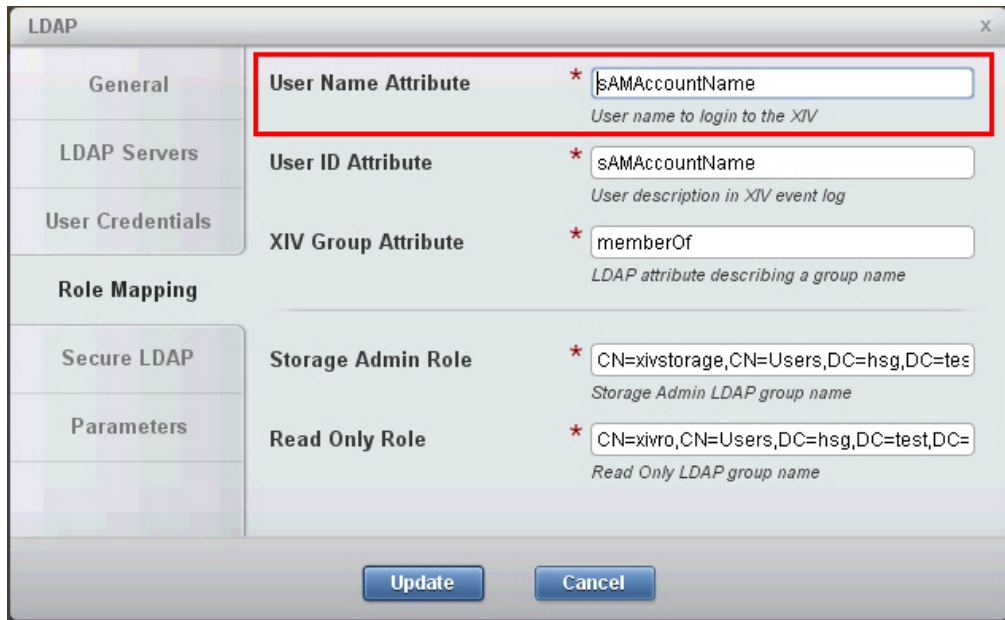


Figure 56. XIV role mapping attributes for directory (LDAP) users

Related tasks:

“Entering the storage system credentials” on page 44

The storage system credentials are used to connect to the IBM storage system or systems, which your cloud interfaces use for storage provisioning.

Configuring event forwarding

Storage system level events generated by a storage system can be forwarded by the IBM Spectrum Connect to the VMware vRealize Log Insight for monitoring and analysis.

Procedure

To configure event forwarding:

1. Go to the /opt/ibm/ibm_spectrum_connect/conf.d/ directory.
2. Open the ibmsyslog.conf file for editing.
3. Look for the following code block:

```
if $programname == 'array_events' then ?SCLogFileName;SCLogFormat
& ~
```

4. Add the IP address of the VMware vRealize Log Insight server and UDP port in the following format: & @@<log_insight_address>:<port>.

Example

To forward the storage system events to the VMware vRealize Log Insight server with IP address 9.151.163.122 via UDP port 514 in addition to writing them to the /var/log/sc/array_events.log file, enter the following:

```
if $programname == 'array_events' then ?ISISLogFileName;ISISLogFormat
& @@9.151.163.122:514
& ~
```


To relay the storage system events to the VMware vRealize Log Insight server with IP address 9.151.163.122 via UDP port 514 without writing them to the local log file, enter the following:

```
if $programname == 'array_events' then @09.151.163.122:514
& ~
```

Deleting virtual volumes and group pools via XCLI

A group pool, used by virtual volumes, can be deleted only after its VVols are removed via XCLI.

About this task

This procedure details how to use XCLI to delete virtual volumes, as they are not visible in the Spectrum Connect GUI. When the virtual volumes are removed, you can delete a group pool, which was used by VVols as well.

Procedure

To delete the virtual volumes and group pools via XCLI:

1. Launch the XIV XCLI tool and log in as a storage integration administrator (*storageintegrationadmin*). Make sure that you are allowed to manage the current domain.
2. List all existing virtual volumes, by using the **vol_list** command.

```
XIV hostdev31c>>vol_list managed=yes domain=test
Name Size (GB) Master Name Consistency Group Pool Creator Compressed Compression Ratio (%) Used Capacity (GB) Compression Saving (GB) Managed
test_1 17 gp_1_meta tester_1 no 0
test_2 18 gp_2_meta tester_1 no 0
```

3. Delete the required virtual volume, by using the **vol_delete** command.

```
XIV hostdev31c>>vol_delete vol=test_1 -y
Warning: This is a managed object. Performing manual operations on it may cause severe problems to the managing software. Are you sure you want to perform the operation on this managed object? y/n: y
```

4. List all group pools in the domain, by using the **gp_list** command.

```
XIV hostdev31c>>gp_list domain=test
Name Meta Pool Name Thin Pool Name Thick Pool Name
gp_1 gp_1_meta gp_1_thin gp_1_thick
gp_2 gp_2_meta gp_2_thin gp_2_thick
```

5. List all pools in the group pool, by using the **pool_list** command.

```
pool_list gp=gp_1 managed=yes
Name Size (GB) Soft Vols (GB) Snap Size (GB) Soft Empty (GB) Hard Size (GB) Hard Vols (GB) Locked Hard Snaps (GB) Hard Empty (GB) Domain Create Compressed Volumes Managed
gp_1_thin 653 0 34 619 309 0 no 0 309 test no yes
gp_1_meta 0 0 0 0 0 0 no 0 0 test no yes
gp_1_thick 34 0 34 0 34 0 no 0 34 test no yes
```

6. Delete the group pool, by using the **gp_delete** command.

```
XIV hostdev31c>>gp_delete gp=gp_1 -y
```

7. Delete the all pools that belonged to the deleted group pool, by using the **pool_delete** command.

```
XIV hostdev31c>>pool_delete pool=gp_1_thin -y
XIV hostdev31c>>pool_delete pool=gp_1_thick -y
XIV hostdev31c>>pool_delete pool=gp_1_meta -y
```

Self-assist options for IBM Spectrum Connect

IBM Support provides several online self-service tools for IBM Spectrum Connect users.

You can try using the following tools to find information and resolve issues without having to contact IBM Support:

- Spectrum Connect-related questions on IBM developerWorks (developer.ibm.com/answers/topics/spectrum%20control/#) – Allows you to ask questions online and get answers from IBM experts or other users. The issue of interest can also be searched for in older discussions.
- IBM Redbooks® (redbooks.ibm.com) – Technical documents where IBM experts share their expertise and best practices for using Spectrum Connect. Among others, you can refer to these publications:
 - Using the IBM Spectrum Accelerate Family in VMware Environments
 - IBM Spectrum Connect and IBM Storage Enabler for Containers -Practical Example with IBM FlashSystem A9000

The above resources are constantly being indexed by web search engines such as Google (google.com).

Chapter 2. Concept diagram for VMware environment

The following concept diagram illustrates how IBM storage systems are accessed and utilized in the VMware environment through IBM Spectrum Connect.

The storage administrator uses Spectrum Connect to select which IBM storage systems (arrays) and what storage resources should be available for use in the VMware environment, and control which specific vCenter servers can utilize the IBM storage resources.

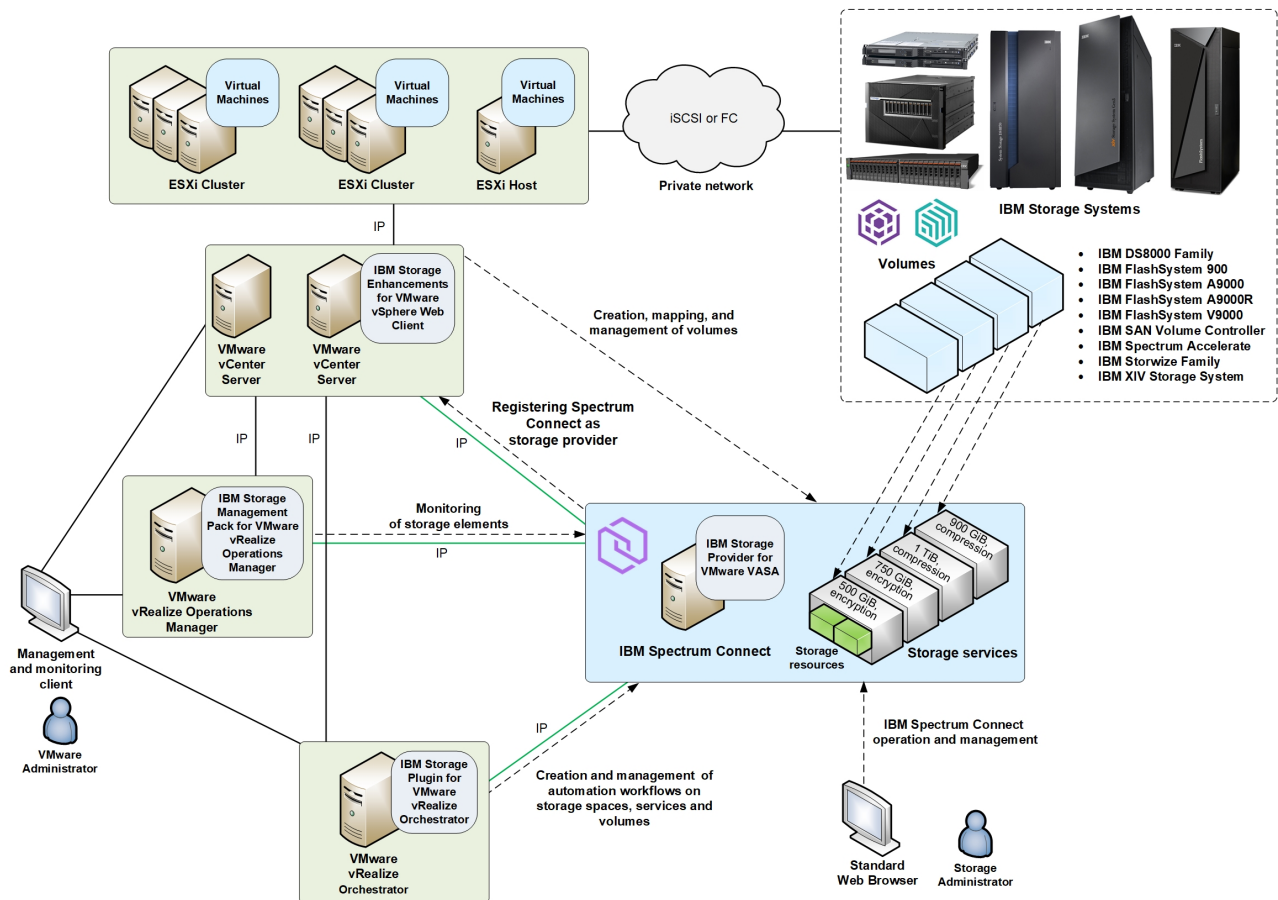


Figure 57. Integration of IBM storage systems with a VMware environment

Spectrum Connect allows registered VMware vCenter servers to utilize its VASA functions, which can be monitored on the vSphere Web Client station.

In parallel, the following operations are enabled on the VMware environment side:

- Through vSphere Web Client, administrators can manually create, map, and fully control storage volumes on the available storage systems and storage resources.
- Through vRealize Orchestrator, administrators can issue workflows for automating the same volumes operations that are available through vSphere Web Client. The automation is run by the VMware vCloud Automation Center (vCAC) platform.

- Through vRealize Operations Manager, administrators can obtain comprehensive monitoring information about the IBM storage resources that are utilized in their virtualized environment.

Note: New storage resources (pools) can be created in advance through the dedicated storage system management tools or from Spectrum Connect. New resources cannot be added from the VMware environment.

Version 3.0.0 introduced a new storage provisioning method, which replaces physical objects (pools and storage systems) with abstracted storage entities (spaces and services), as detailed in “VMware Storage Policy Based Management (SPBM)” on page 3 and “Storage space and service management” on page 4.

Chapter 3. IBM Storage Provider for VMware VASA

The IBM Storage Provider for VMware VASA improves the ability to monitor and automate storage-related operations on VMware platforms.

From its Spectrum Connect host, the IBM Storage Provider for VMware VASA provides a standard interface for any connected VMware vCenter server using the VMware vSphere APIs for Storage Awareness (VASA). It delivers information about IBM storage topology, capabilities, and state, together with storage events and alerts to vCenter server in real time. The storage capabilities are presented as a combination of space and service, facilitating in provisioning volumes based on a predefined set of storage capacities, as detailed in “Storage space and service management” on page 4.

To visualize how this cloud interface is integrated in a VMware environment, see Chapter 2, “Concept diagram for VMware environment,” on page 99.

Managing and monitoring VASA access

You can control and monitor the utilization of the IBM Storage Provider for VMware VASA, which is part of the IBM Spectrum Connect solution.

Refer the following sections on details how to set VASA credentials and manage trusted certificates.

- “Setting the VASA credentials”
- “Managing VASA trusted certificates” on page 103

Setting the VASA credentials

The VASA credentials comprise a username and a password that is set separately from the Spectrum Connect user accounts, and separately from the storage system credentials.

About this task

VMware vCenter servers can use the VASA credentials to connect to Spectrum Connect and utilize VASA functions, as explained in “Registering Spectrum Connect as a storage provider on vCenter server” on page 103.

From the Spectrum Connect UI you can set, change, or display these VASA credentials at any time.

Note:

- Only one set of a username and a password can be used for the VASA credentials, which applies to all vCenter servers that require VASA functions and connect to Spectrum Connect.
 - The same operations are available from the CLI as well, as explained in “CLI – Setting the VASA credentials” on page 87.
-

Procedure

1. Click **VASA Provider settings** on the Settings menu. The **VASA Provider settings** dialog box is displayed.




Figure 58. VASA Provider settings dialog box

2. Enter the username and password that you want to set, confirm the password and then click **Apply**.

Note: The VASA username must be different from any username stored in the Spectrum Connect user database. See “Managing Spectrum Connect users” on page 40.

3. In the **Priority** drop-down box, select a priority (0–255) of Spectrum Connect as a storage provider in vCenter. In its turn, it defines priority of all storage systems added to the current Spectrum Connect instance on the vCenter server.

Note: You can use the Priority setting for resolving an issue, when a Spectrum Connect instance defined for the same storage system is active on one vCenter server and passive on another vCenter. This might occur, when two SC instances with the same priority, managing the same storage system are registered as storage providers on two different vCenters. A loss of connectivity between any of the Spectrum Connect instances and vCenter servers might result in the following configuration:

- Spectrum Connect A is an active storage provider on vCenter 1.
- Spectrum Connect B is an active storage provider on vCenter 2.

This renders Spectrum Connect A to be the active storage provider on vCenter 1 and passive on vCenter 2. In its turn, Spectrum Connect B is active provider on vCenter 2 and passive on vCenter 1. To avoid this situation, use the Priority parameter to define the required Spectrum Connect instance to have a higher priority than another one. This ensures that it remains active after a failover.

Managing VASA trusted certificates

You can view and, if needed, remove all registered vCenter servers as well as upload new certificates.

Before managing VASA certificates, verify that the VASA credentials have been set, as explained in the previous section.

To display a list of all vCenter servers that are currently registered and are utilizing VASA functions, go to the **VASA Trusted Certificates** dialog box (**Settings > VASA trusted certificates**). Instructions on adding Spectrum Connect as a storage provider to the VMware vCenter Server are detailed in the “Registering Spectrum Connect as a storage provider on vCenter server” section. During Spectrum Connect registration, the certificate validity is verified for VMware VASA 1.0 and VASA 2.0. Also, for VMware VASA 2.0, it is verified that the certificate is signed by the VC root CA (when working with VMCA) or it is uploaded as a trusted certificate.

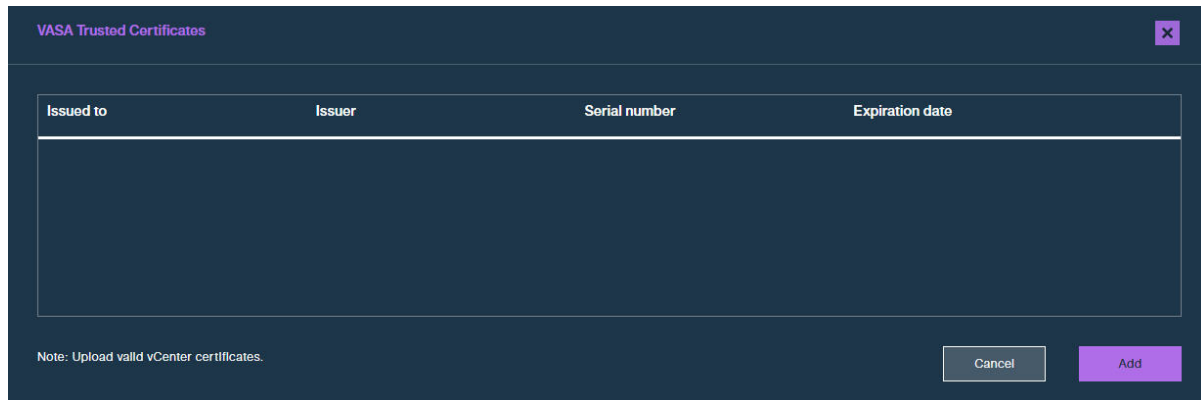


Figure 59. Registered VASA servers (vCenter servers that employ VASA services)

If you want to disconnect a vCenter server, click and highlight the row of the vCenter server that you want to disconnect, and then click the **Remove** button.

If you want to upload a new certificate, click **Add > Upload Certificate > Browse** and attach your certificate file.

Using the IBM Storage Provider for VMware VASA

This chapter focuses on how to use the IBM Storage Provider for VMware VASA after the required configuration on IBM Spectrum Connect has been completed.

After the IBM storage systems have been added to Spectrum Connect and after the VASA access credentials were set (see “Required and optional initial tasks” on page 18), you can start using the IBM Storage Provider for VMware VASA by registering Spectrum Connect on the relevant vCenter servers.

Registering Spectrum Connect as a storage provider on vCenter server

To use the IBM Storage Provider for VMware VASA solution component, you need to register Spectrum Connect as a storage provider on VMware vCenter server.

Before you begin

- When the IBM Storage Provider for VMware VASA and the IBM Spectrum Control storage provider, formerly known as IBM Tivoli® Storage Productivity Center (TPC), are registered on the same VMware vCenter server, while the same storage system is configured for both, vCenter uses IBM Spectrum Control provider **as the only source of information** for that system's storage views on vSphere Web Client.

In such a case, determine whether the IBM Spectrum Control capabilities are sufficient for replacing the IBM Storage Provider for VMware VASA. If the IBM Storage Provider for VMware VASA is still needed, IBM TPC 5.2 (or later) provides a method of excluding storage systems from a specific storage provider, allowing you to remove the system association with the IBM TPC provider.

- The **date and time** that are defined on both the vCenter server and on Spectrum Connect must be identical. To accurately synchronize the date and time between the two servers, you can connect both to a Network Time Protocol (NTP) server.

Note: VMware VASA 2.0 is required for implementing virtual volume functionality.

About this task

The following procedure describes how to register Spectrum Connect on a single vCenter server by using vSphere Web Client.

Procedure

To register Spectrum Connect on VMware vCenter server, complete the following steps.

1. In vSphere Web Client, click **vCenter > vCenter Servers**, and click the vCenter server, on which you want to register the IBM storage provider.
2. On the **Manage** tab, click **Storage Providers**.
3. Click the plus sign to add a new storage provider.

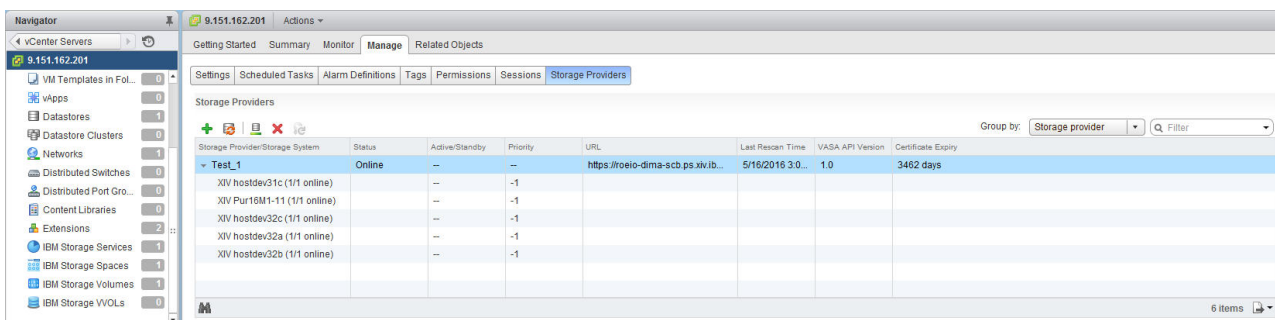


Figure 60. vSphere Web Client – Storage Providers list

The New Storage Provider dialog box is displayed.

4. Enter the name, URL, and pre-configured username and password (VASA Secret) for accessing Spectrum Connect (the VASA Secret is predefined as explained in “Setting the VASA credentials” on page 101). The URL should be entered in the *ip:port* format, specifying the relevant IP address and port number of Spectrum Connect:
 - `https://[Spectrum Connect IP address]:8440/services/vasa1` for VASA 1.0

- [https://\[Spectrum Connect IP address\]:8440/services/vasa](https://[Spectrum Connect IP address]:8440/services/vasa) for VASA 2.0

Note:

- The hostname/IP address in the URL must match the common name, which is used, when generating an SSL certificate for the Spectrum Connect server. See “Managing server certificates” on page 37.
- In this example, "Spectrum Connect IP address" stands for the IP address or domain namespace of Spectrum Connect.
- You can change the default TCP port (8440) at any time by running a script, as explained in “Changing the Spectrum Connect communication port” on page 77.

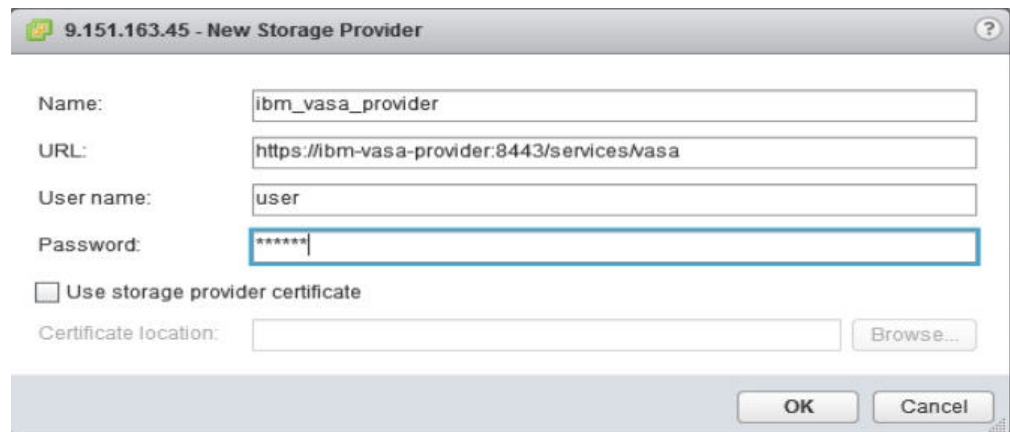


Figure 61. New Storage Provider dialog box for VASA 2.0

5. Click **OK**. A security alert dialog box is displayed.

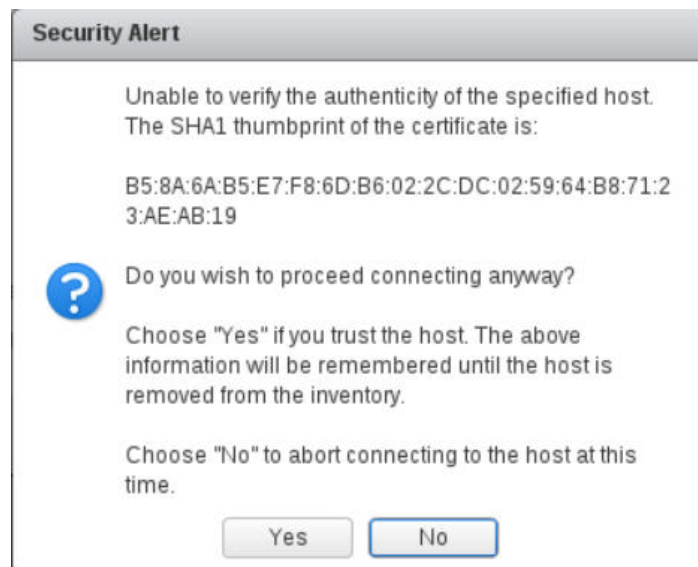


Figure 62. vCenter certificate thumbprint dialog box

6. Click **Yes** to accept the certificate. Spectrum Connect is added to vCenter Server.

Note: The certificate provides improved security by adding server authentication.

If, during the registration process, you have an active Spectrum Connect instance, restart your web browser or refresh the Spectrum Connect GUI window to ensure the successful acquisition of the new server certificate.

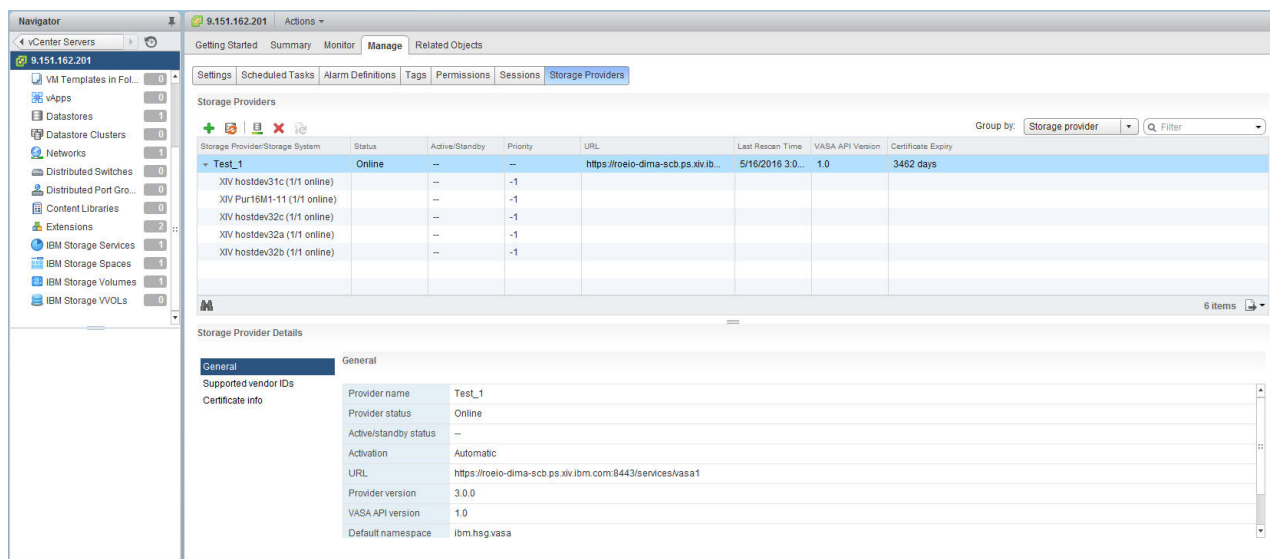


Figure 63. Storage Providers list displaying Spectrum Connect

7. You can ensure continuous storage management by combining multiple Spectrum Connect instances, registered as storage providers, into high-availability groups. This process is described in “Defining a high-availability group” on page 35.

Note: Storage spaces and services defined on an active storage provider do not appear on the standby Spectrum Connect. The spaces and services become visible on the Spaces/Services pane of the Spectrum Connect GUI, when it becomes active after system failover.

What to do next

After Spectrum Connect is registered as a storage provider on vCenter server, you can start managing VMware storage resources on IBM storage systems via vSphere Web Client. Complete the following tasks, referring to the relevant VMware documentation for details.

1. Create a new VM storage policy to set requirements for the storage resources.

Note: During a VM storage policy creation, a VVol-enabled storage service is exposed as a capability value of the IBM Storage Service entry. This information is displayed in the *storage container:storage system* format.

2. Add a new datastore.

Note: A VVol container is created for a single storage system on a storage space. For example, storage space A and storage system A are visible as container 1; storage space A and storage system B are visible as container 2.

3. Create a new virtual machine on the datastore.

Chapter 4. IBM Storage Enhancements for VMware vSphere Web Client

The IBM Storage Enhancements for VMware vSphere Web Client integrate into the VMware vSphere Web Client platform and enable VMware administrators to independently and centrally manage their storage resources on IBM storage systems.

Depending on the IBM storage system in use, VMware administrators can self-provision volumes (LUNs) in selected storage services that were predefined by the storage administrators. The volumes are mapped to ESXi hosts, clusters, or datacenters as logical drives that can be used for storing VMware datastores (virtual machine data containers). See “VMware Storage Policy Based Management (SPBM)” on page 3 and “Storage space and service management” on page 4.

As opposed to the IBM Storage Management Console for VMware vCenter, which is individually installed on each vCenter server, the IBM Storage Enhancements for vSphere Web Client are installed only on the vSphere Web Client Server, allowing multiple vCenter servers to utilize IBM storage resources. In addition, procedures for attaching and detaching storage resources to the services are performed on the Spectrum Connect side, rather than on the vSphere Web Client side.

The IBM Storage Enhancements for VMware vSphere Web Client are automatically deployed and enabled for each and every vCenter server that is registered for vSphere Web Client services on the connected Spectrum Connect.

To visualize how this cloud interface is integrated in a VMware environment, see Chapter 2, “Concept diagram for VMware environment,” on page 99.

Managing integration with vSphere Web Client

Before you can use the IBM Storage Enhancements for VMware vSphere Web Client on the web client side, you need to define on the Spectrum Connect side the vCenter servers for which you want to provide storage resources. Then, you can attach storage services that you want to make available to each vCenter server.

The storage services that you attach on the Spectrum Connect side become visible on vSphere Web Client, and can be used for volume creation by using the IBM Storage Enhancements for vSphere Web Client (for more information, see Chapter 4, “IBM Storage Enhancements for VMware vSphere Web Client”).

- “Adding a vCenter server”
- “Updating the credentials of a vCenter server” on page 109
- “Removing a vCenter server” on page 110
- “Delegating storage services to a vCenter server” on page 110
- “Canceling service delegation to a vCenter server” on page 112

Adding a vCenter server

Add the VMware vCenter servers for which you want to provide storage resources through IBM Spectrum Connect.

Before you begin

Log out of any vSphere Web Client application connected to a vCenter server that you want to add to Spectrum Connect. If you stay logged in, you will be able to use the extension only after you log out and log into vCenter after the connection.

About this task

You need to connect to and add the vCenter servers for which you can later attach storage services that would be visible and accessible on the vSphere Web Client side. You can add a single vCenter server at a time, as described in the following procedure.

Note: For any vCenter server that you add, the IBM Storage Enhancements for VMware vSphere Web Client (see Chapter 4, “IBM Storage Enhancements for VMware vSphere Web Client,” on page 107) are automatically deployed and enabled on the vSphere Web Client Server side.

Procedure

1. Click **Add Interface** on the **Interfaces** pane and then select **Add vCenter**. The **Add New vCenter Server for vWC** dialog box is displayed.




Figure 64. Add vCenter Server for vWC dialog box

2. Enter the IP address or hostname of the vCenter server, as well as the username and password for logging into that vCenter server. If the provided IP address and credentials are accepted by the vCenter server, it is added to the list of servers on the **Interfaces** pane.

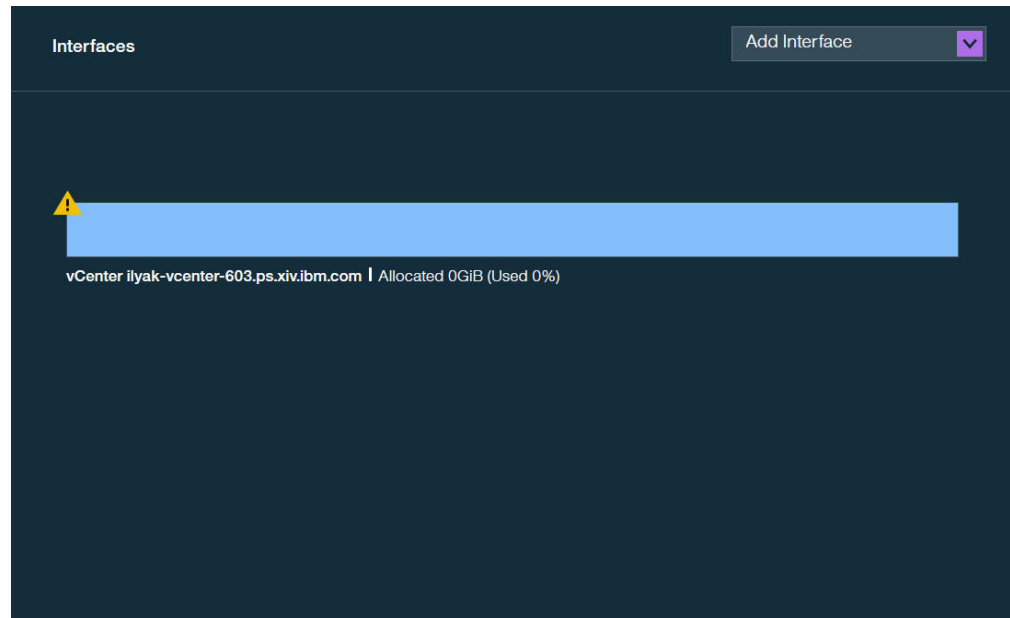


Figure 65. Interfaces pane

Note:

- If you want to use the vSphere Web Client extension on all vCenter servers that operate in linked mode, each server instance must be added to Spectrum Connect. This ensures that the extension is registered on all linked servers properly.
 - The same vCenter server cannot be added to more than one Spectrum Connect instance. Any attempt to add an already registered vCenter server to another Spectrum Connect will override the primary connection.
-

What to do next

Attach storage services to the vCenter server, as explained in “Delegating storage services to a vCenter server” on page 110.

Updating the credentials of a vCenter server

At any time, you can update the credentials that are used by Spectrum Connect to access a vCenter server.

About this task

Whenever the credentials on the vCenter server side change, you can update these credentials on the Spectrum Connect side to allow continued management of storage resources on the vSphere Web Client side.

Note: Prior to changing the vCenter credentials on the Spectrum Connect side, verify that the vCenter user has sufficient access level to complete this procedure.

Procedure

1. On the **Interfaces** pane, right-click the vCenter server for which you want to update the credentials, and then select **Modify**. The **vCenter Server Settings**

dialog box is displayed.




Figure 66. vCenter Server Settings dialog box

2. Enter the new username and password for accessing the vCenter server. Then, click **Apply**.

Removing a vCenter server

When a vCenter server is no longer needed, you can disconnect it from Spectrum Connect by removing it from the Interfaces pane.

About this task

- A removed vCenter server, along with the storage services attached to it, can no longer be managed by the IBM Storage Enhancements for VMware vSphere Web Client. Following the removal, the IBM Storage Enhancements become disabled for that vCenter on the vSphere Web Client Server side.
- If the removed vCenter server is attached to active storage service, the information for these services, as well as their pools is no longer displayed in vSphere Web Client. However, vSphere data access and service level for the services and pools is not affected.
- After the removal, you can add the vCenter server again to fully restore its management.

Procedure

1. On the Interfaces pane, right-click the vCenter server that you want to remove, and then select **Remove**. A confirmation message is displayed.
2. Click **OK** to remove the vCenter server.

Delegating storage services to a vCenter server

You must delegate any storage service that you want to use for volume management operations on the vSphere Web Client side to the vCenter server.

Before you begin

- Storage services can be delegated only with *Spectrum Connect Admin*, *Storage Admin*, *System Admin*, or *Security Admin* storage credentials (see “Entering the storage system credentials” on page 44).

- When working with DS8000, services can be delegated only with *Spectrum Connect Admin* or *Logical Operator* storage credentials.
- Working with VMware VVols requires the *Storage Integration Admin* access level configured at the storage system side.

Any other type of storage credentials (read-only, application admin) cannot perform service delegation. If your credentials are not sufficient, contact your storage administrator for assistance.

About this task

The service delegation is done on the Spectrum Connect side, providing more control to storage administrators as opposed to the VMware administrators on the vSphere environment side. After the delegation, the services and their pools become visible and manageable on the vSphere Web Client side (by using the IBM Storage Enhancements).

Procedure

1. On the **Interfaces** pane, click the vCenter server to which you want to delegate one or more services.
2. On the **Spaces/Storage Services** pane, select the storage space from which you want to choose storage services. The available services that reside on the selected storage space are immediately displayed.
3. Right-click on a service that you want to delegate to the vCenter server, and then select **Delegate to vCenter_IP_address**, or click the **Attach/Delegate** button on the service. The service color changes to indicate the successful delegation, as illustrated below.

You can continue the process by right-clicking available services under the current storage space.

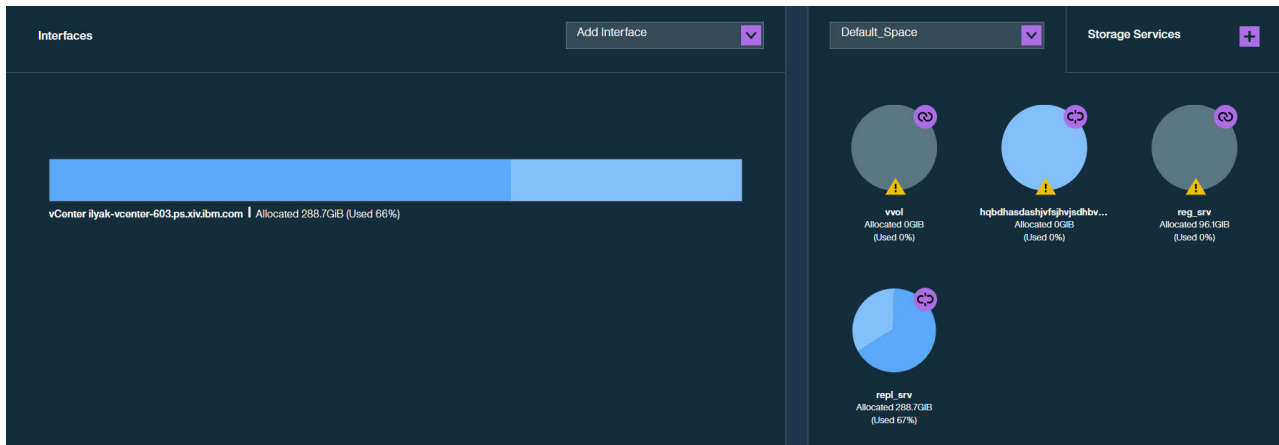


Figure 67. vCenter server with delegated service

The vCenter server provides indication for the allocated and used storage space.

- Allocated – amount of storage space available on all pools connected to the delegated services.
- Used – amount of storage space used by volumes on all pools connected to the delegated services. The volumes are created using IBM Storage Enhancements for VMware vSphere Web Client.

Canceling service delegation to a vCenter server

When required, you can cancel a storage service delegation to the vCenter server.

About this task

- A storage service, which delegation has been canceled, can no longer be managed by the included solution components (see “Included cloud interfaces” on page 2).
- If the storage resources on such service contain working volumes, the information of these volumes is no longer displayed in vSphere Web Client. However, **vSphere data access and service level for these volumes is not affected**. In addition, the removed storage pools (including its volumes) can be managed from the standard IBM storage system management tools.
- The working volumes of such service remain visible in vSphere Web Client, as long as they are mapped to ESXi hosts. However, these volumes cannot be managed via vWC.
- At any time, you can delegate the storage service again to fully restore its management.

Procedure

1. On the **Interfaces** pane, click a vCenter server which service delegation you want to cancel. The services that are currently delegated to the vCenter server are highlighted on the **Spaces/Storage Services** pane.
2. Right-click on a service which delegation you want cancel, and then select **Cancel delegation to vCenter_IP_address**, or click the **Detach/Cancel Delegation** button on the service. The service color changes to indicate the successful cancellation of service delegation.

You can continue the process by right-clicking delegated services under the current space.

Using the IBM Storage Enhancements for VMware vSphere Web Client

Together with supported IBM storage systems that are managed by Spectrum Connect, the deployed IBM Storage Enhancements enable the following management features on vSphere Web Client for registered vCenter servers:

- Full control over storage volumes, including volume creation, resizing, renaming, mapping, unmapping, multipath policy enforcement, and deletion.
- Easy and integrated allocation of volumes to VMware datastores, used by virtual machines that run on ESXi hosts, clusters, or datacenters.

Note:

- The IBM Storage Enhancements are automatically deployed and made available for the vCenter servers that were registered (added) on IBM Spectrum Connect (see “Adding a vCenter server” on page 107).
 - For information about the required vSphere user privileges, see “Required vSphere privileges” on page 113.
-

See the following sections for more information:

- “Viewing the IBM storage object information” on page 114
- “Creating and mapping a new storage volume (LUN)” on page 124
- “Extending a volume” on page 131
- “Renaming a volume” on page 132

- “Setting multipath policy enforcement for a volume” on page 133
- “Unmapping a volume from one or more hosts” on page 135
- “Deleting an unused volume” on page 135
- “Managing consistency group attachment” on page 136
- “Displaying the virtual volume information” on page 138

Required vSphere privileges

To operate the IBM Storage Enhancements for VMware vSphere Web Client, you must have the minimum required privileges defined in your vSphere user role.

Use the **Role Manager** extension in vSphere Web Client to define the required privileges for your user role as detailed in the following table.

Table 23. Required vSphere privileges

Task	Required vSphere user privilege
Adding a vCenter server to the IBM Spectrum Connect (see “Adding a vCenter server” on page 107)	<ul style="list-style-type: none"> • Extension – In this category, select Register extension, Unregister extension, and Update extension. • Global – In this category, select Log event and Cancel task. • Tasks – In this category, select Create task and Update task. • Sessions – In this category, select Impersonate user. • Host – In this category, select Configuration > Storage partition configuration.
Storage provisioning (volume creation and management) from vSphere Web Client (see Chapter 4, “IBM Storage Enhancements for VMware vSphere Web Client,” on page 107)	<ul style="list-style-type: none"> • Tasks – In this category, select Create task and Update task. • Sessions – In this category, select Impersonate user. • Host – In this category, select Configuration > Storage partition configuration.




Figure 68. VMware vSphere Web Client – Create Role dialog box

For more detailed information about how to set the vSphere Web Client privilege types, refer to the VMware vSphere 5.1 Documentation Center (pubs.vmware.com/vsphere-51/index.jsp).

Viewing the IBM storage object information

After the IBM Storage Enhancements for VMware vSphere Web Client are properly installed, the IBM Storage categories are shown under the standard vSphere Web Client categories for each vCenter server, as shown in the following figure.

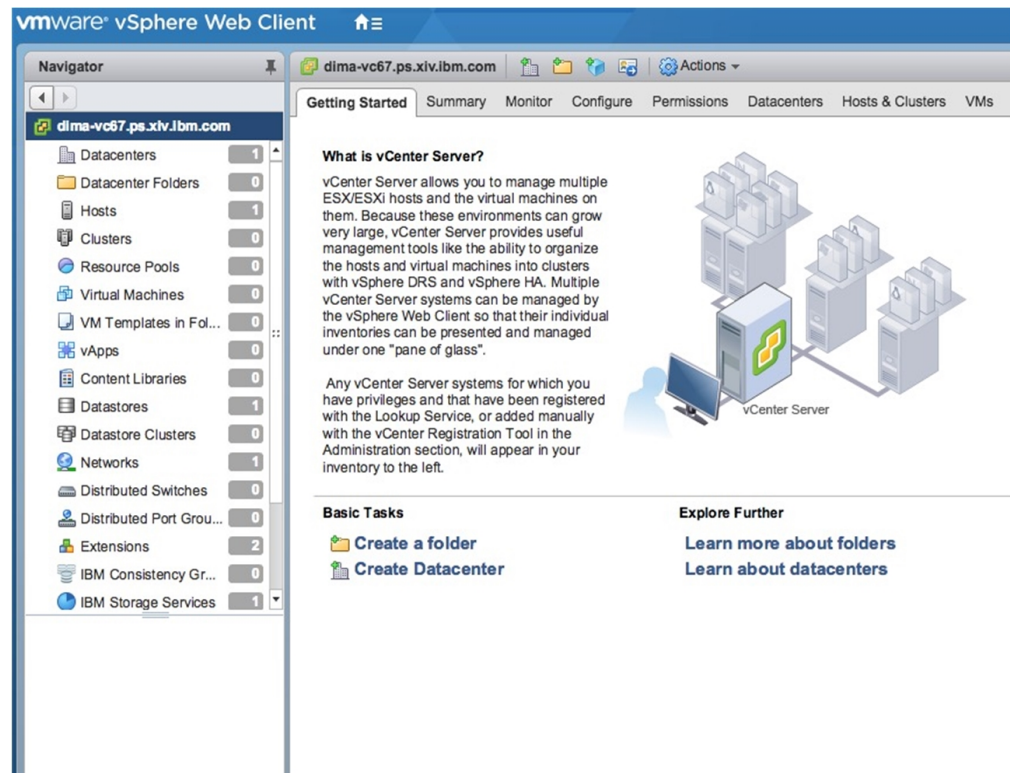


Figure 69. IBM Storage categories in vSphere Web Client

For each vCenter server, the following IBM Storage categories are available for that vCenter server:

- Storage services
- Storage spaces
- Storage volumes
- Consistency groups
- Storage virtual volumes (VVols)

You can click and open an IBM Storage category to view the entities which are currently available for the selected vCenter server.

The screenshot displays the VMware vSphere Web Client interface. The left-hand side features a Navigator pane with a tree view of storage-related objects. The main content area shows the 'Replication_Pool' configuration page, specifically the 'Related Objects' tab. Under the 'IBM Storage Volumes' sub-tab, a table lists the storage volumes. The table has the following columns: Storage Device Name, Volume Identifier, Storage Array, Volume Name, Storage Type, Volume Size (GiB), Usage, Serial, and Path Selection. One volume, 'Demo_Vol', is listed with a size of 1.0 GiB and is currently 'Not in use'.

Storage Device Name	Volume Identifier	Storage Array	Volume Name	Storage Type	Volume Size (GiB)	Usage	Serial	Path Selection
Demo_Vol	naa.6001738cfc9	XIV Gen4d-67c	Demo_Vol	FlashSystem A9K	1.0	Not in use	5EA000000000	Round Robin

Figure 70. IBM Storage Service information

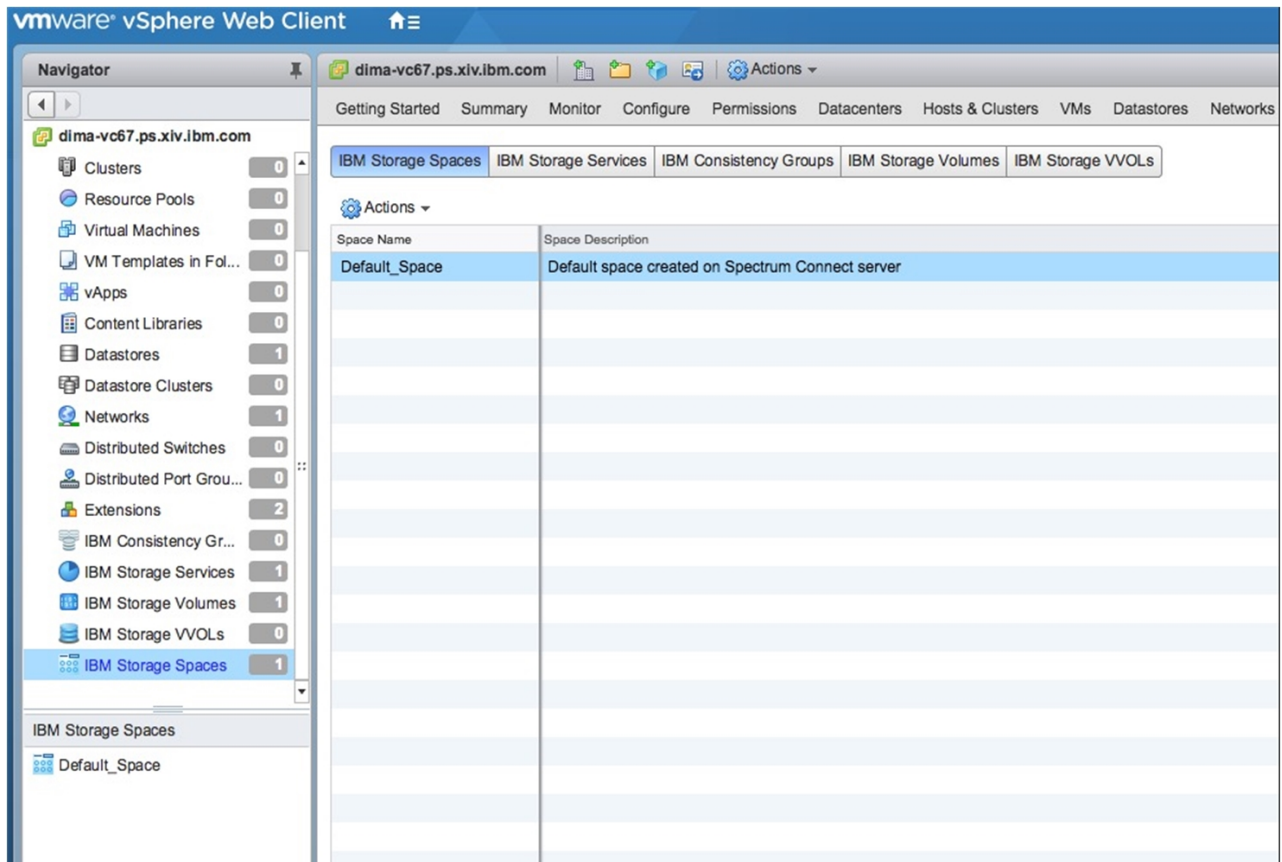


Figure 71. IBM Storage Space information

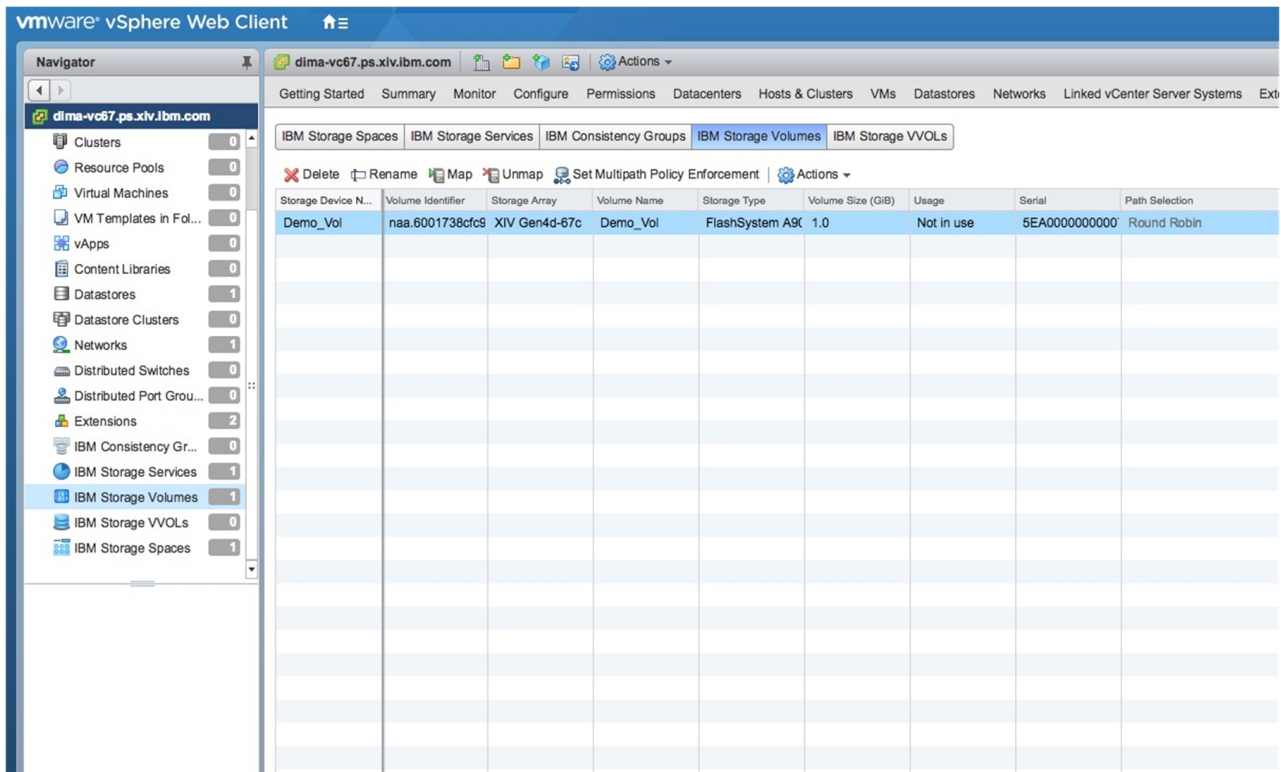


Figure 72. IBM Storage Volume information

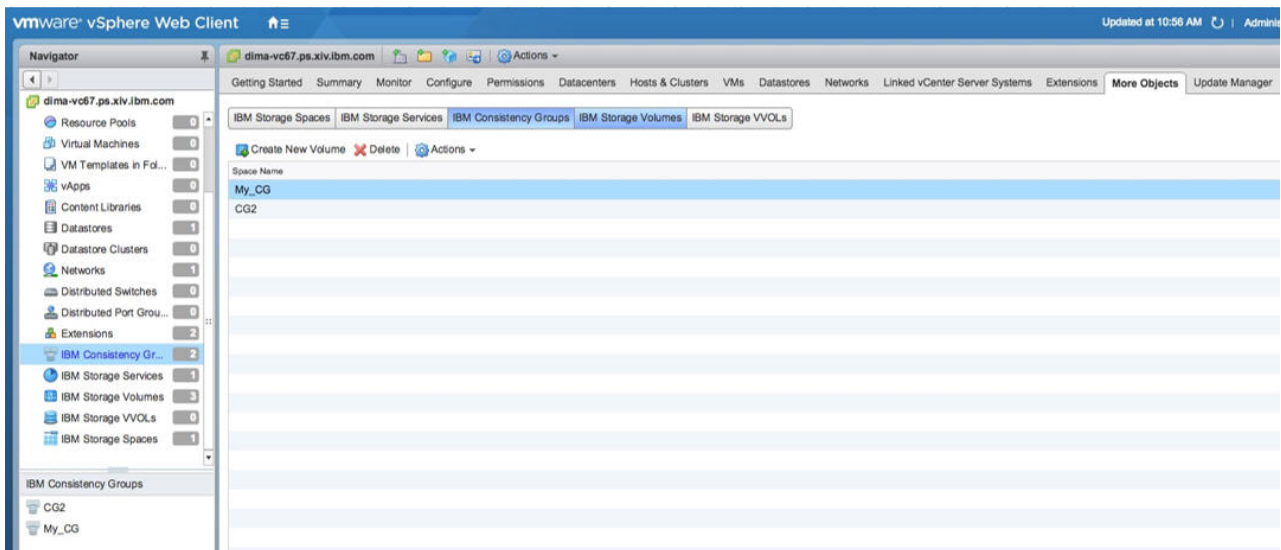


Figure 73. IBM Consistency Group information

Storage Device Name	Volume Identifier	Volume Name	Storage Array	Model	Volume Size (GB)	Compression Rate	Use
rfc4122.4fb8930-53bc-4562-98a0-26b3fb357f7	rfc4122.4fb8930-53bc-4562-98a0-26b3fb357f7	rfc4122.4fb8930-53bc-4562-98a0-26b3fb357f7	2810-214-MN65016		4	N/A	
rfc4122.844bb13f-4c51-48ac-b4ea-cc6485e7f38d	rfc4122.844bb13f-4c51-48ac-b4ea-cc6485e7f38d	rfc4122.844bb13f-4c51-48ac-b4ea-cc6485e7f38d	2810-214-MN65016		42	N/A	

Figure 74. IBM Storage VVol information

In addition, you can display a summary of IBM storage items, as well as information on objects related to them, as shown in the following figures.

The screenshot displays the VMware vSphere Web Client interface for the 'Replication_Pool' service. The left-hand 'Navigator' pane shows a tree view of the environment, with 'IBM Storage Services' and 'Replication_Pool' selected. The main content area shows the 'Summary' tab with the following details:

- Service Name: Replication_Pool
- Service Description: (empty)
- Service Type: regular
- Space: Default_Space
- Max Volume Size For Provisioning: 385.0 GiB

Two summary panels are visible:

- Capabilities:** A table with one entry:

Capability	Value
Replication	Synchronous
- Physical Capacity:** A table showing capacity metrics:

Allocated Capacity	192.5 GiB
Free Capacity	385.1 GiB
Total Capacity	577.6 GiB

Figure 75. IBM Storage Service summary

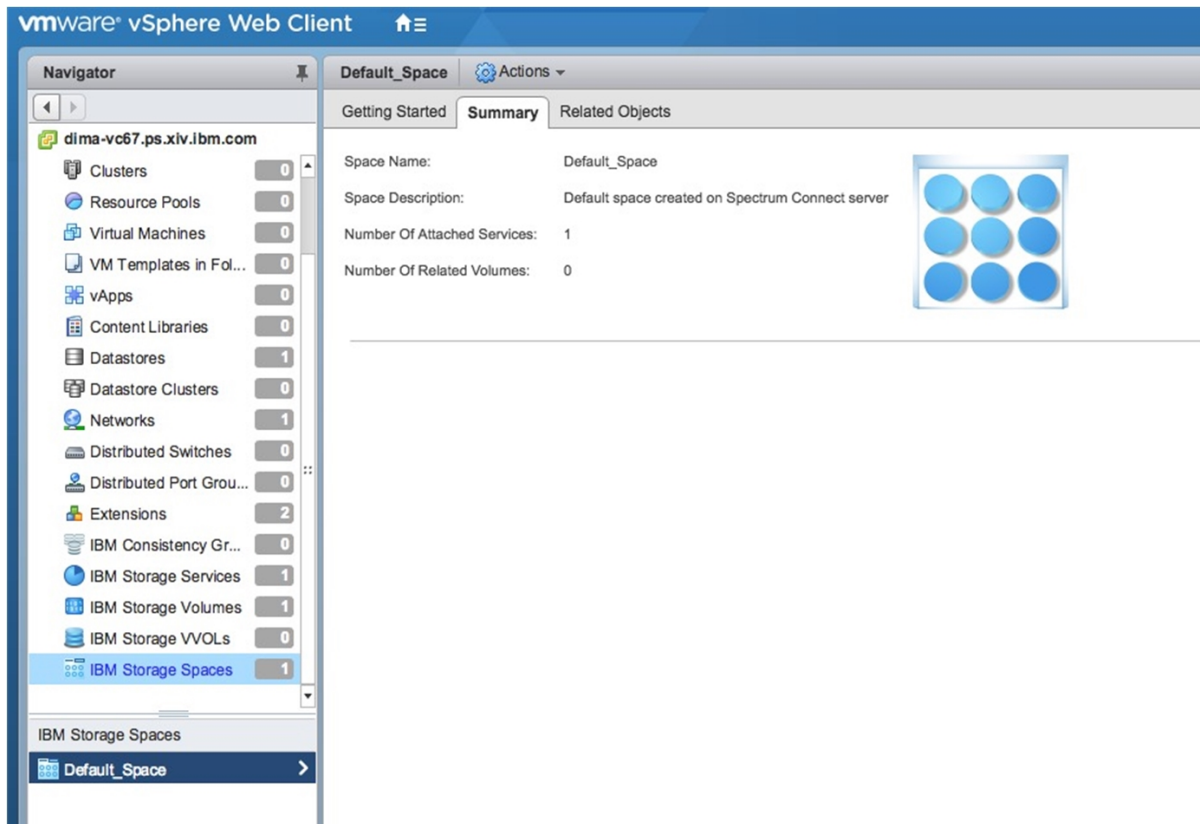


Figure 76. IBM Storage Space summary


vmware vSphere Web Client

Navigator: dlma-vc67.ps.xiv.ibm.com

- Clusters: 0
- Resource Pools: 0
- Virtual Machines: 0
- VM Templates in Fol...: 0
- vApps: 0
- Content Libraries: 0
- Datastores: 1
- Datastore Clusters: 0
- Networks: 1
- Distributed Switches: 0
- Distributed Port Grou...: 0
- Extensions: 2
- IBM Consistency Gr...: 0
- IBM Storage Services: 1
- IBM Storage Volumes: 1
- IBM Storage VVOLS: 0
- IBM Storage Spaces: 1

Demo_Vol

Getting Started | Summary | Monitor | Manage | Related Objects



Storage Device Name: Demo_Vol
 Volume Identifier: naa.6001738cfc9035ea000000000788431
 System Name: XIV Gen4d-67c
 Pool Name: Demo_Replication
 Domain Name: roei_domain
 System Type: FlashSystem A9000R

Details	
Volume Name	Demo_Vol
Usage	Not in use
Serial	5EA0000000...
Number Of Snapshots	0
Last Snapshot	--
Consistency Group	--
Mirroring	Yes
Data Compression	Enabled
LUN	1
Service Compliance	True
Compression Savings (Estimated)	0%

Capacity	
0 GiB 1 GiB	
Used	0 GiB
Free	1 GiB
Total	1 GiB

Path Selection	
Multipath Policy Enforcement	Disabled
Path Selection	Round Robin

Related Objects	
Datastore	--
Storage Space	Default_Space
Storage Service	Replication_Pool

Figure 77. IBM Storage Volume summary

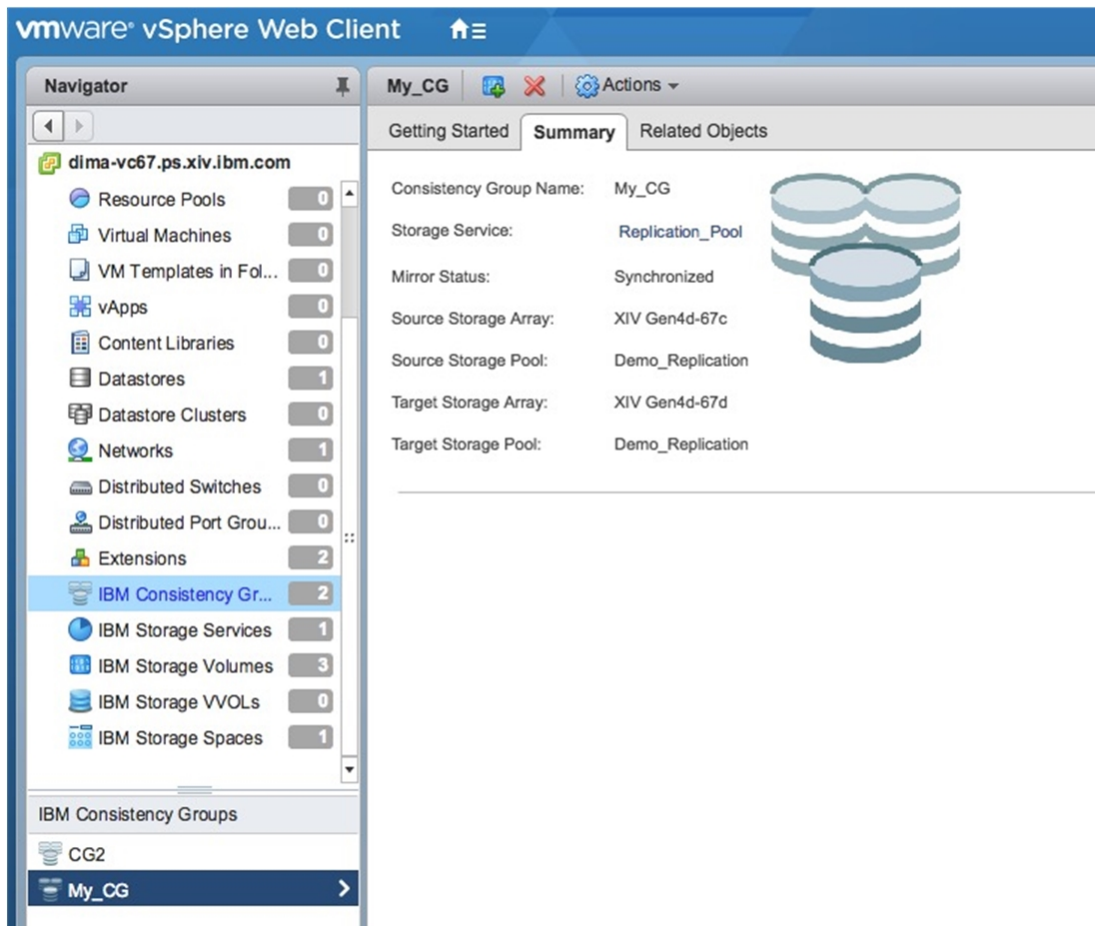


Figure 78. IBM Consistency Group summary




Figure 79. IBM Storage VVol summary

And so on, click the other information categories that are available in vSphere to view the relevant IBM storage information in these categories as well.

Creating and mapping a new storage volume (LUN)

The IBM Storage Enhancements for VMware vSphere Web Client allow you to create new volumes (LUNs) directly from the vSphere Web Client interface. These volumes can be used as storage devices in the vSphere environment.

About this task

Any created volume is mapped to either ESXi hosts, clusters, or datacenters, so that the virtual machines on these hosts, clusters, or datacenters would be able to save datastore information on that volume. The volume is created on a storage service, according to the service definitions from Spectrum Connect.

In addition to single volume creation, you can create multiple volumes simultaneously. If you choose this option, the created volumes are appended with differently numbered suffixes that are automatically generated by the system in consecutive order.

Important:

- You can create volumes only on storage pools that have been attached to the relevant vCenter server on the Spectrum Connect side. For more information, see “Managing integration with vSphere Web Client” on page 107
 - The ESXi hosts and clusters to which you map the created volumes must be predefined on the storage system side. For more information, refer to your IBM storage system documentation.
-

Procedure

1. In vSphere Web Client, navigate to the relevant vCenter server and then to the specific IBM storage object (space, service or volume). The IBM storage object and the relevant storage resources (pools) should already be associated with the vCenter server (see “Viewing the IBM storage object information” on page 114).
2. Go to the **Getting Started** tab of an IBM Storage Services or IBM Storage Volumes entry.
3. Click **Create a new IBM Storage Volume**. Alternatively, select a storage object, and click **Actions > Create New Volume**.




Figure 80. IBM storage service view – Clicking Create New Volume

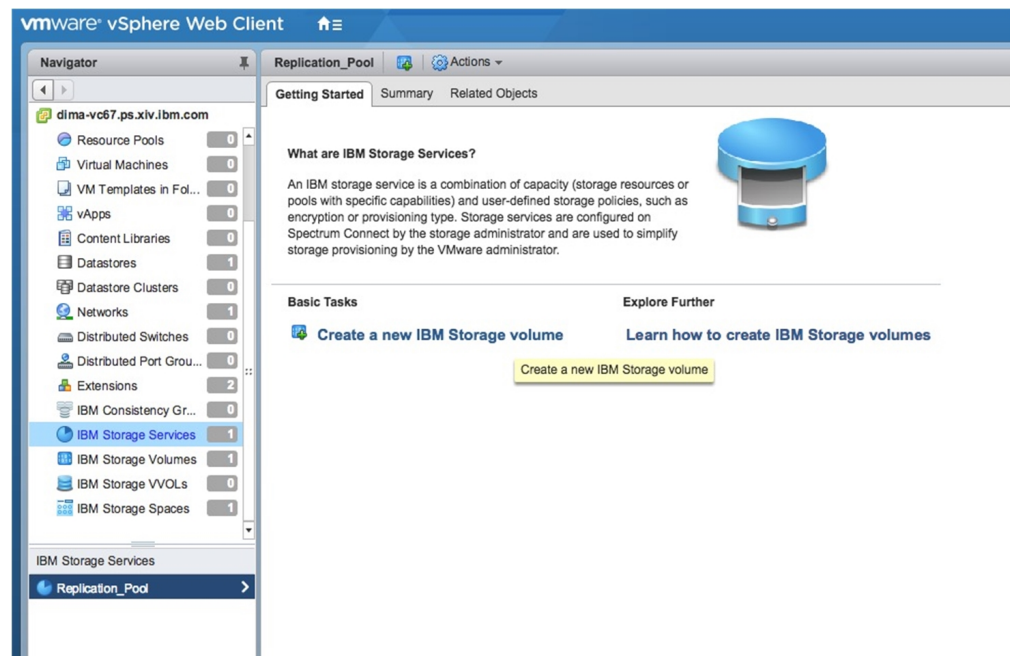


Figure 81. Top Level Objects view – Clicking Create a new IBM Storage volume

The Create New Volume dialog box is displayed.

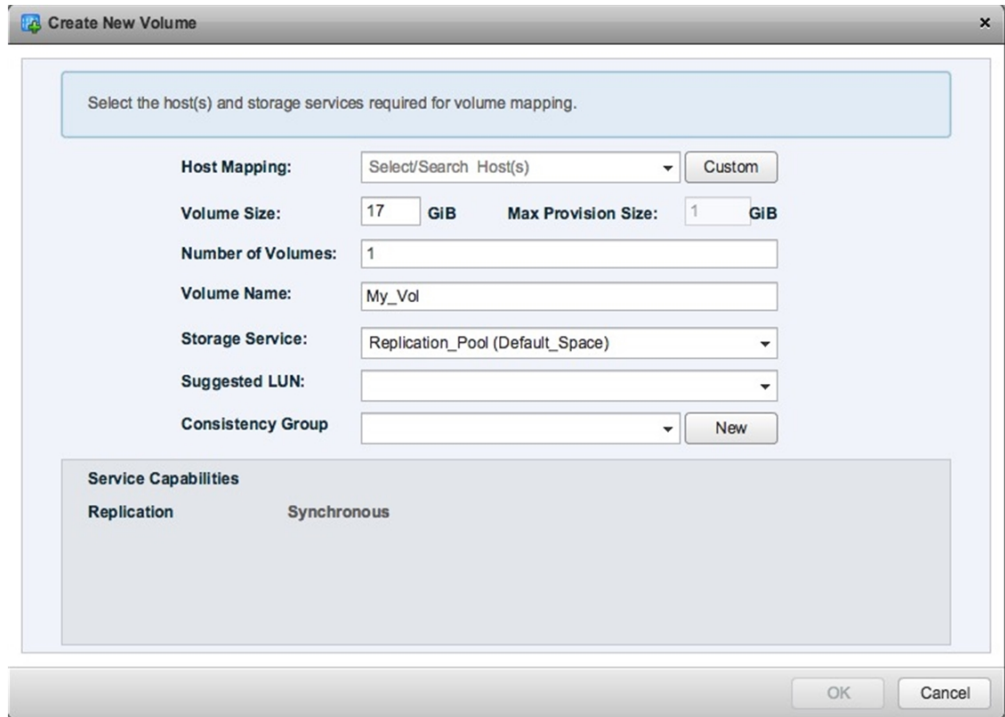


Figure 82. Create New Volume wizard

Note: When you create a single volume, a LUN (logical unit number) is assigned to that volume, and you can later change the LUN assignment. If you create multiple volumes, LUNs (logical unit numbers) are automatically assigned to those volumes and cannot be modified later.

4. In the **Volume Size** text box, enter the size for the new volume.

Note:

- It is recommended to define the size of an XIV volume in a multiple of 17 GB. The Volume Size box appears with a yellow rectangle around it if the size value is not a multiple of 17 GB.
 - A storage resource with largest amount of free space, which is currently attached the service, is automatically selected.
 - The minimum size for compressed XIV volumes is 51 GB.
-

5. In the **Volume Name** text box, enter the name that you want to assign to the new volume.
6. If you want to create multiple volumes simultaneously: In the **Number of Volumes** text box, enter the number of volumes that you want to create simultaneously. The text box next to the **Volume Name** entry displays vol_{1} by default. The {1} represents the suffix value, and it must be kept as part of the volume name. You edit the volume name and also move the suffix value within the name (the {1} suffix does not have to be at the end of the volume name).

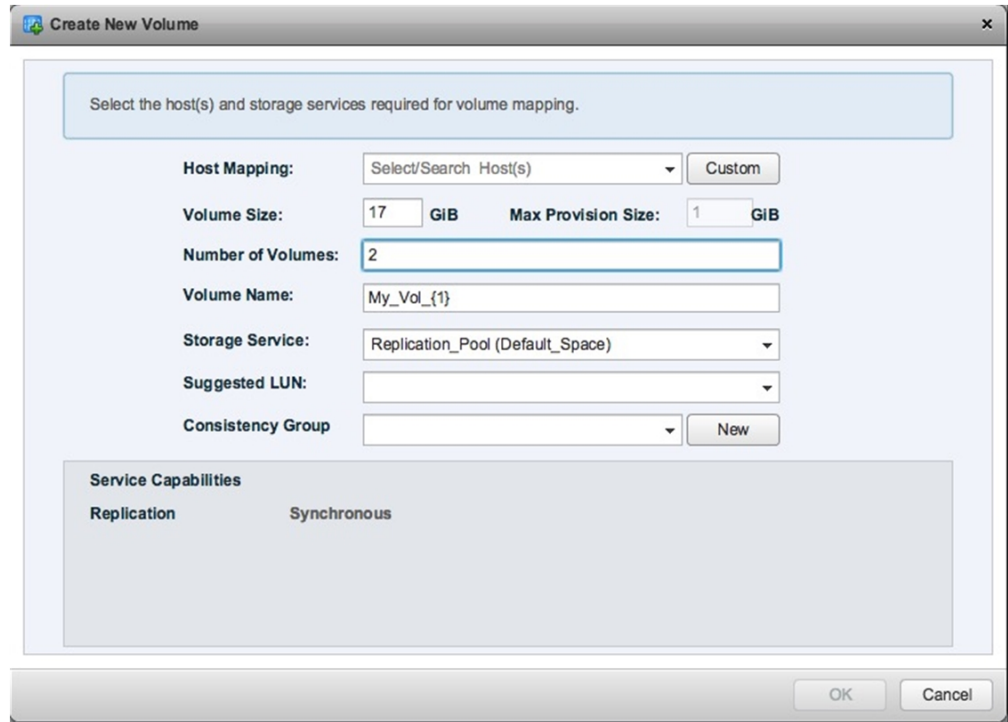


Figure 83. Creating multiple volumes

7. In the **Storage Service** text box, select a storage service, on which the volume will be created.

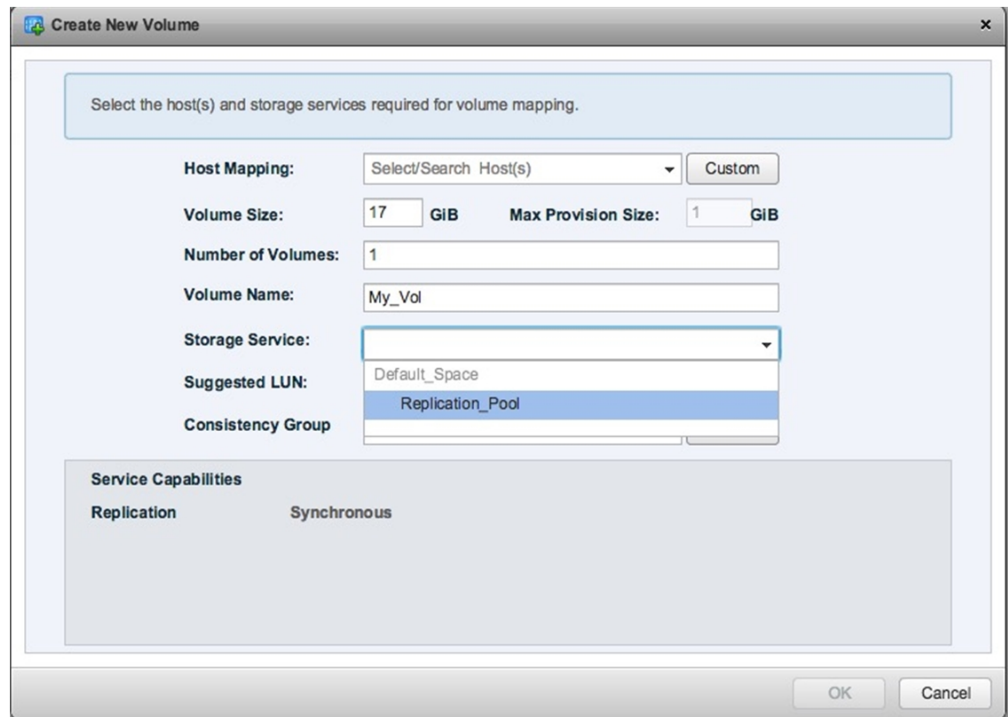


Figure 84. Selecting storage service

8. From **Host Mapping**, select the host(s), cluster(s), or datacenter(s) to which you want to map the new volume. You can click **Custom** to specify a custom

mapping in the **Advanced Host Mapping** dialog box.




Figure 85. Advanced Host Mapping dialog box

Note: Any ESXi host that is connected to the storage system can be selected. Hosts that are not connected to the storage system are marked and a message notifies you about any connectivity problem. If you select a datacenter, its member clusters and hosts are automatically selected under it.

Important: You must map the volume to at least one ESXi host, cluster, or datacenter in order to enable vSphere management of the created volume.

9. If you are creating a single volume, you can select the LUN that should represent the new volume on the storage system, or keep the automatically selected LUN. The LUNs are automatically selected when creating multiple volumes.

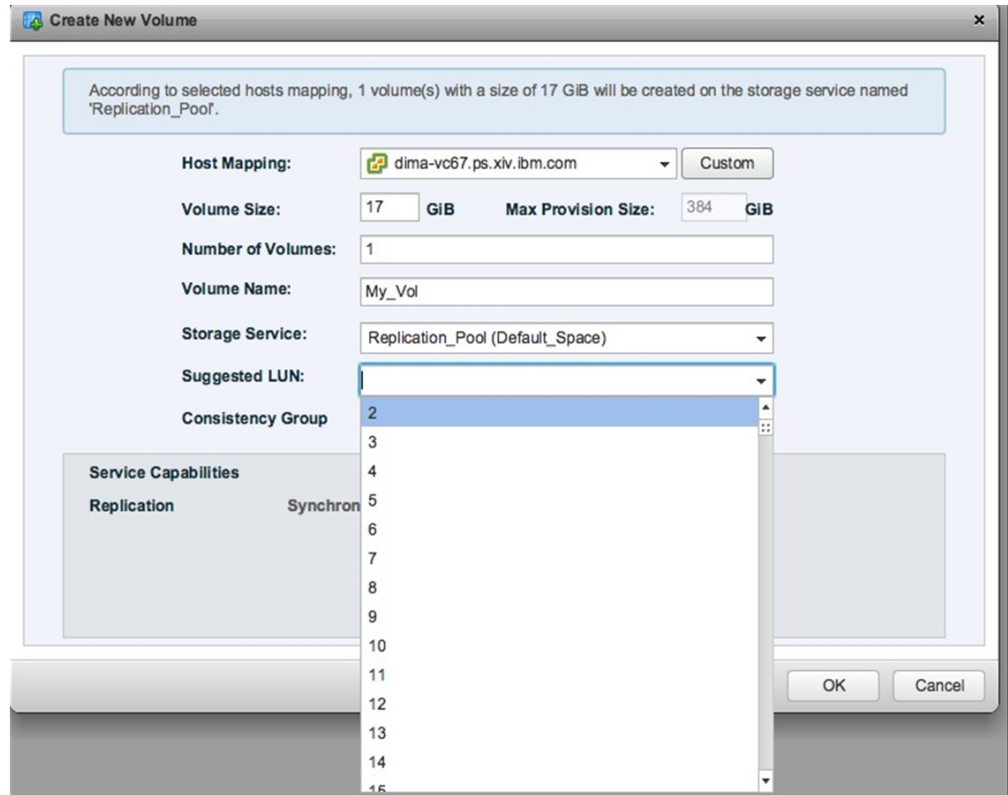


Figure 86. Selecting LUN

Note: After the volume is created, the specific LUN association cannot be changed, and the same number cannot be assigned to a different volume. The specific LUN can become available for reassignment only after its associated volume is deleted.

10. If you want to provision the volume in an existing consistency group, select the required group from the **Consistency Group** drop-down box.

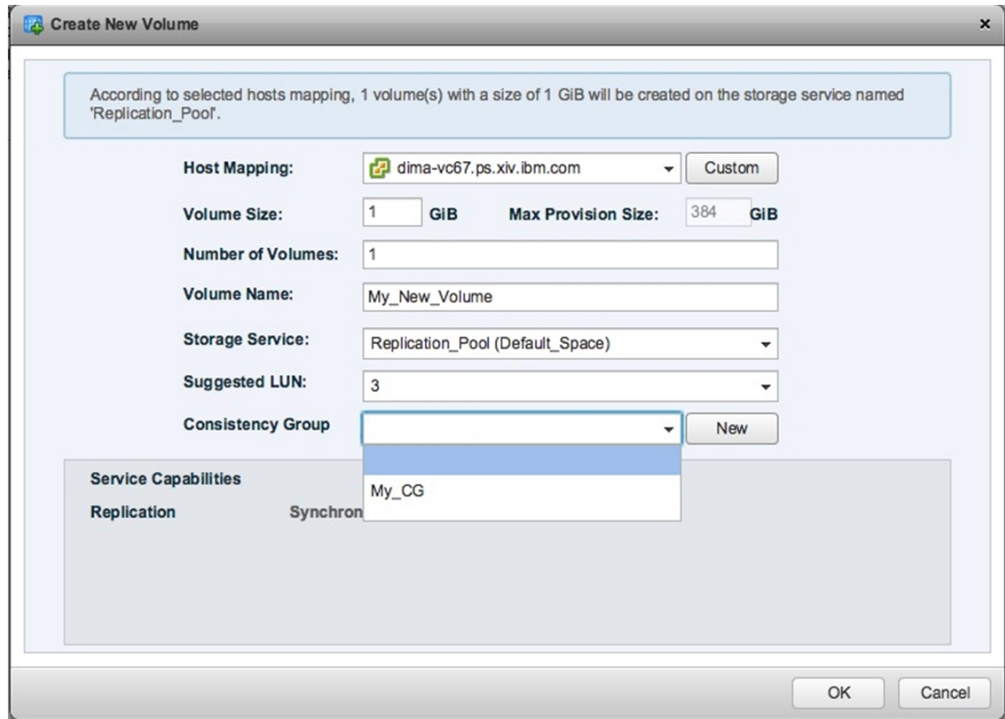


Figure 87. Selecting consistency group

11. If you want to provision the volume in a new consistency group, click the **New** button next to the **Consistency Group** drop-down box. The New Consistency Group dialog box is displayed.




Figure 88. Naming a new consistency group

12. Enter a name for the new consistency group and click **OK**. The dialog box closes, and the consistency group name is displayed in the drop-down box, marked as **new**. Pay attention that the consistency group itself will be activated only after the volume is created.

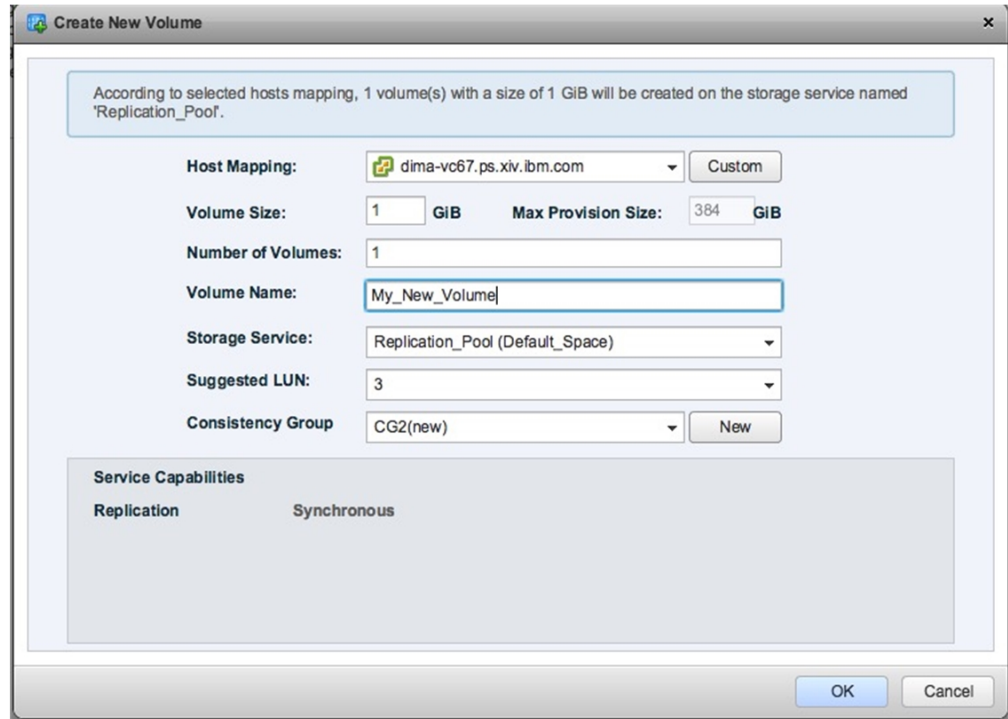


Figure 89. Create New Volume dialog box with all fields filled out

13. Review the details of the new volume that is about to be created, and then click **OK** to confirm its creation as detailed.

Extending a volume

If enough free space is available on the relevant storage pool, you can extend the size of an existing volume.

Procedure

Complete the following procedure to extend the size of a volume.

1. In vSphere Web Client, locate a volume that you want to extend:
 - Go to **IBM Storage Services**, select a service that contains the volume, open the **Related Objects** tab, and select the required volume row, or
 - Go to **IBM Storage Volumes**, select the required volume.
2. Right-click the volume and choose **Extend** or use the **Actions** menu to select **Extend**.

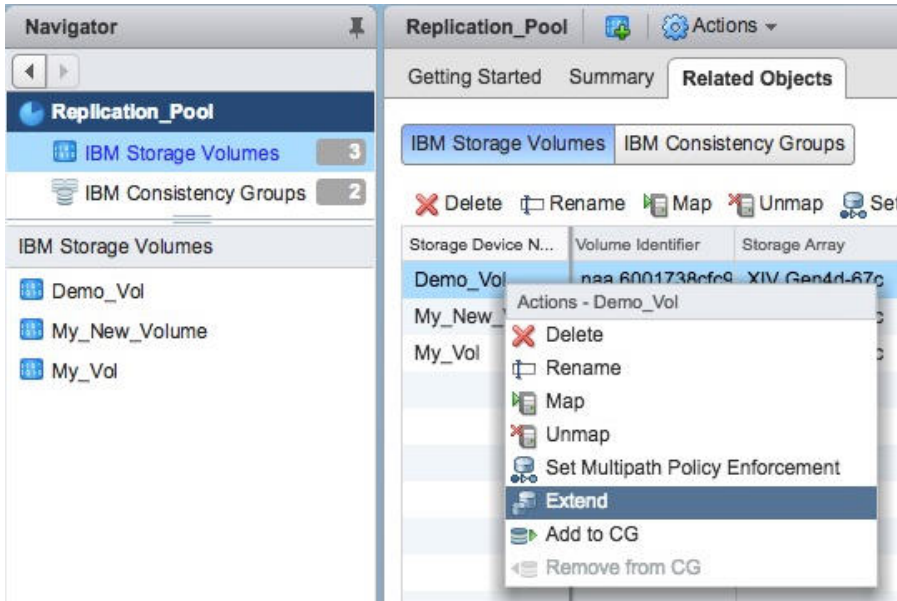


Figure 90. Clicking Extend on the pop-up menu

The Extend Volume dialog box is displayed.

3. In the **Volume Size** text box, enter the new size for the volume. Alternatively, place the mouse pointer on the graphic image of the storage pool, and then click and slide the space marker rightward to set the new volume size. The numerical value in Volume Size is automatically updated accordingly.




Figure 91. Extend Volume dialog box

4. Click **OK**.

Important: Extending the size of a volume does not automatically increase the datastore capacity.

Renaming a volume

Whenever required, you can rename any existing volume by performing the following procedure.

About this task

Renaming a volume is a logical action that does not have any physical effect on the volume or its logical connections. Renaming a volume also changes its displayed name in the vSphere environment.

Procedure

1. In vSphere Web Client, locate a volume that you want to rename:
 - Go to **IBM Storage Services**, select a service that contains the volume, open the **Related Objects** tab, and select the required volume row, or
 - Go to **IBM Storage Volumes**, select the required volume.
2. Right-click the volume and choose **Rename** or use the **Actions** menu to select **Rename**.

The Rename Volume dialog box is displayed.

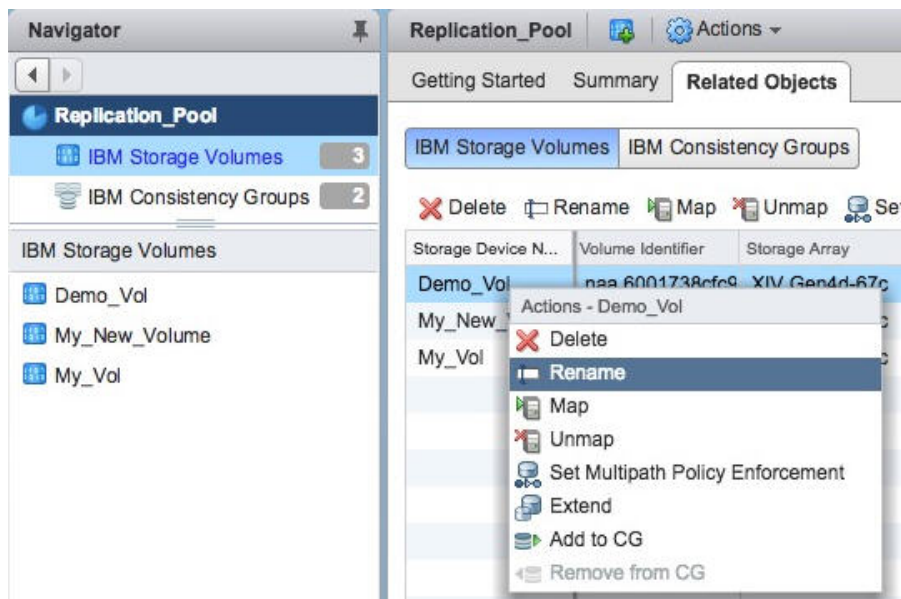


Figure 92. Rename volume option

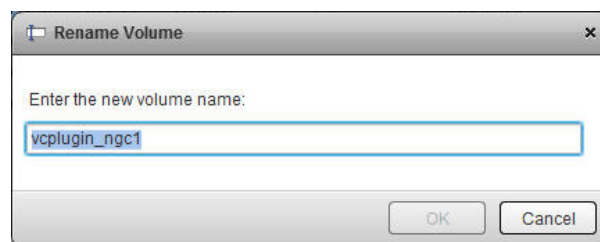


Figure 93. Rename Volume dialog box

3. Enter the new name that you want to assign to the volume, and then click **OK**.

Setting multipath policy enforcement for a volume

You can set the multipath policy enforcement for a single volume.

About this task

By default, the **Round Robin** multipath policy is enforced on volumes. You can disable or change this enforcement for a specific volume if needed.

Note: If you are using ESXi version 5.1 or earlier with DS8000 or Storwize Family systems, see “Setting the multipath policy for DS8000 and Storwize Family systems” on page 224.

Procedure

1. In vSphere Web Client, locate a volume for which you want to change the enforcement:
 - Go to **IBM Storage Services**, select a service that contains the volume, open the **Related Objects** tab, and select the required volume row, or
 - Go to **IBM Storage Volumes**, select the required volume.
2. Right-click the volume and choose **Set Multipath Policy Enforcement** or use the **Actions** menu to select **Set Multipath Policy Enforcement**. The Change Multipath Policy Enforcement dialog box is displayed.

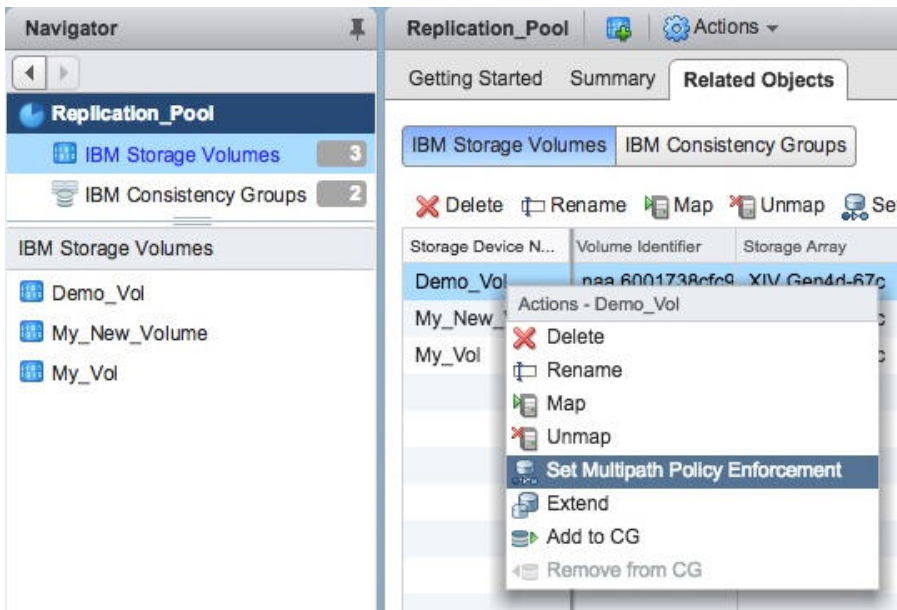


Figure 94. Set Multipath Policy Enforcement option

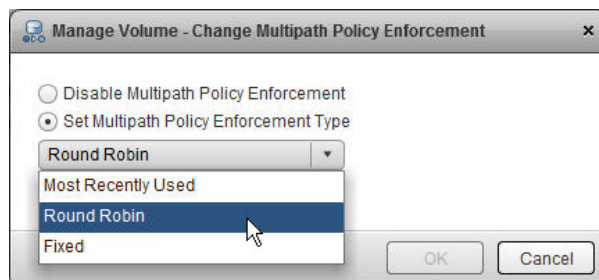


Figure 95. Change Multipath Policy Enforcement dialog box

3. Select the required option and then click **OK**.

Important: After the policy is set, it is enforced by overriding any existing policy for this volume.

Unmapping a volume from one or more hosts

When volumes or ESXi hosts are no longer needed, or if new ones are to replace the current ones, you can unmap volumes from the hosts.

About this task

Important: A volume (LUN) must remain mapped to at least one host. Otherwise, you cannot view the volume or perform any actions on it from vSphere Web Client.

Procedure

1. In vSphere Web Client, locate a volume that you want to unmap:
 - Go to **IBM Storage Services**, select a service that contains the volume, open the **Related Objects** tab, and select the required volume row, or
 - Go to **IBM Storage Volumes**, select the required volume.
2. Right-click the volume and choose **Unmap** or use the **Actions** menu to select **Unmap**.

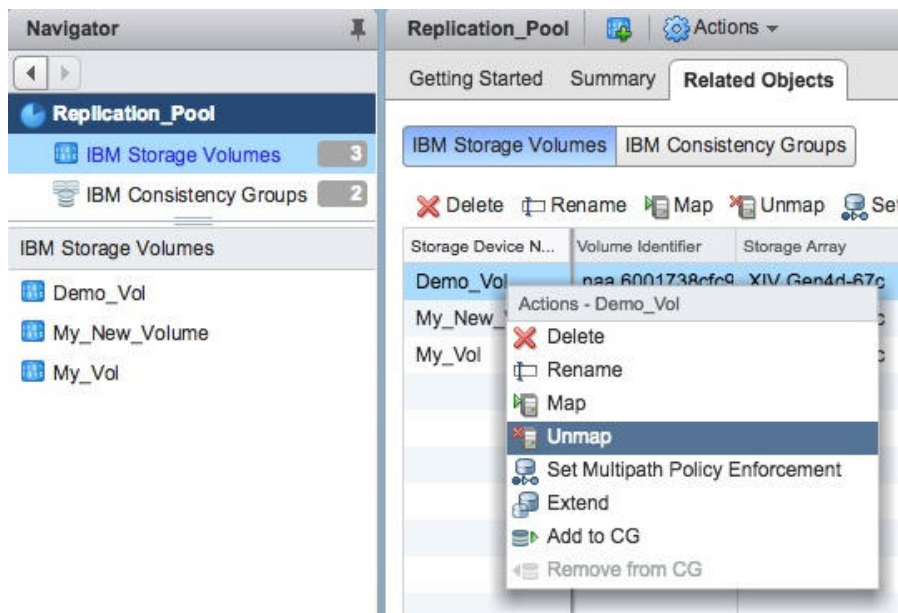


Figure 96. Unmap volume

The Unmap Volume dialog box is displayed.

3. Select the hosts or clusters from which you want to unmap the volume, and then click **OK**.

Deleting an unused volume

When a storage volume is unused and no longer required, you can delete it.

Before you begin

Important: You cannot delete volumes that are currently used by datastores or as a raw-mapped LUN.

Procedure

1. In vSphere Web Client, locate a volume that you want to delete:
 - Go to **IBM Storage Services**, select a service that contains the volume, open the **Related Objects** tab, and select the required volume row, or
 - Go to **IBM Storage Volumes**, select the required volume.
2. Right-click the volume and choose **Delete** or use the **Actions** menu to select **Delete**.

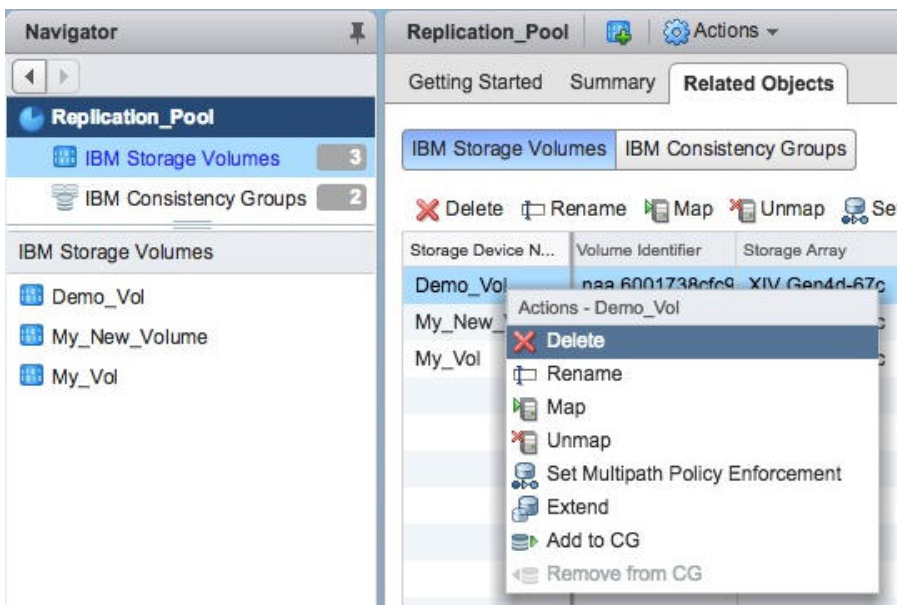


Figure 97. Delete volume

The Delete Volume confirmation message is displayed.

3. Click **OK** to confirm the deletion, or **Cancel** to exit without deleting the volume.

Note: A volume, whose deletion fails, disappears from the volume list. The volume reappears in the list after the next population.

Managing consistency group attachment

A storage volume can be added to a consistency group or removed from it, if required.

Procedure

1. In vSphere Web Client, locate a volume that you want to add to a consistency group, or remove it from the group:
 - Go to **IBM Storage Services**, select a service that contains the volume, open the **Related Objects** tab, and select the required volume row, or
 - Go to **IBM Storage Volumes**, select the required volume.

2. Right-click the volume and choose **Add to CG** or use the **Actions** menu to select **Add to CG**.

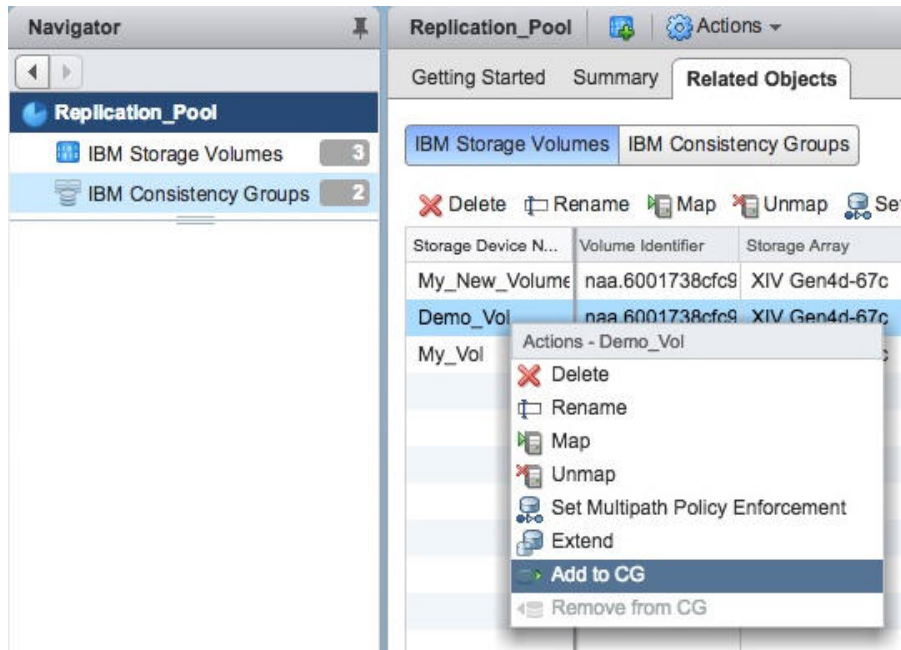


Figure 98. Adding volume to consistency group

The **Add to CG** dialog box is displayed.




Figure 99. Selecting a consistency group

3. Select the required consistency group and click **OK** to finish.
4. You can remove a volume from the existing consistency, by selecting the volume, as explained above, right-clicking the volume and choosing **Remove from CG**.

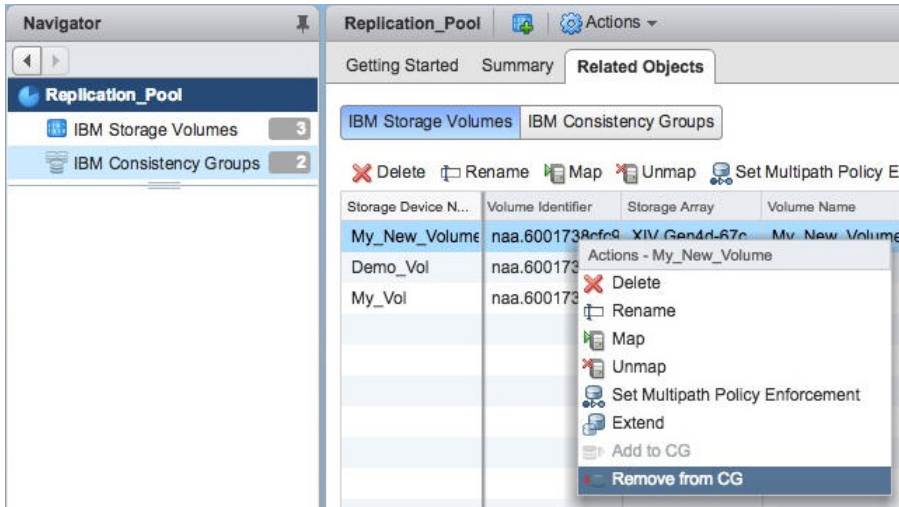


Figure 100. Removing volume from consistency group

A confirmation message is displayed.

5. Click **OK** to confirm the removal.

Displaying the virtual volume information

If virtual volumes are used as storage resources for virtual machines, you can view the VVol summary, snapshot information and details of VMs which are using them.

Before you begin

Verify that a virtual machine has been created on a VVol datastore.

Procedure

1. In vSphere Web Client Navigator, go to **IBM Storage VVols** and select a virtual volume that you intend to view.
2. In the right-hand pane, select one of the following tabs:
 - **Getting Started** to review general information on IBM's implementation of VMware virtual volume technology.
 - **Summary** to display the general information about the virtual volume and its characteristics.
 - **Monitor** to view the snapshot status, if this functionality is enabled.
 - **Related Objects** to list the VMs which are using the VVol.

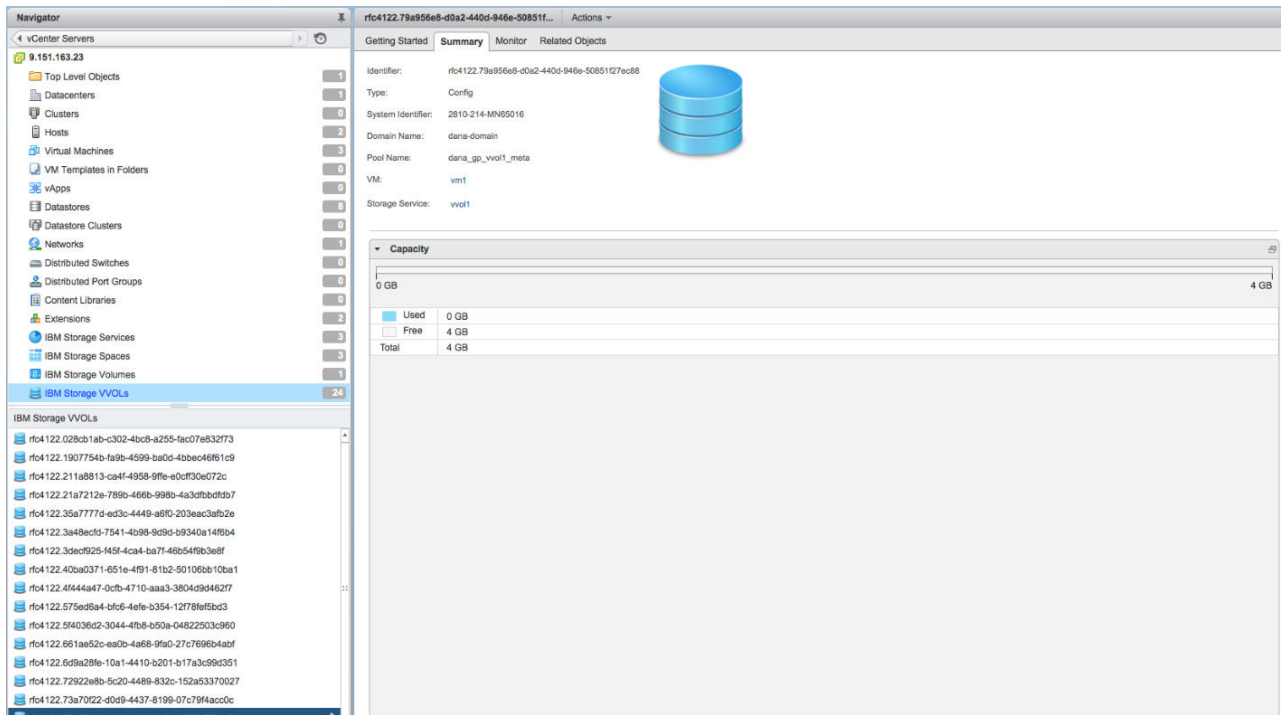


Figure 101. VVol Summary tab

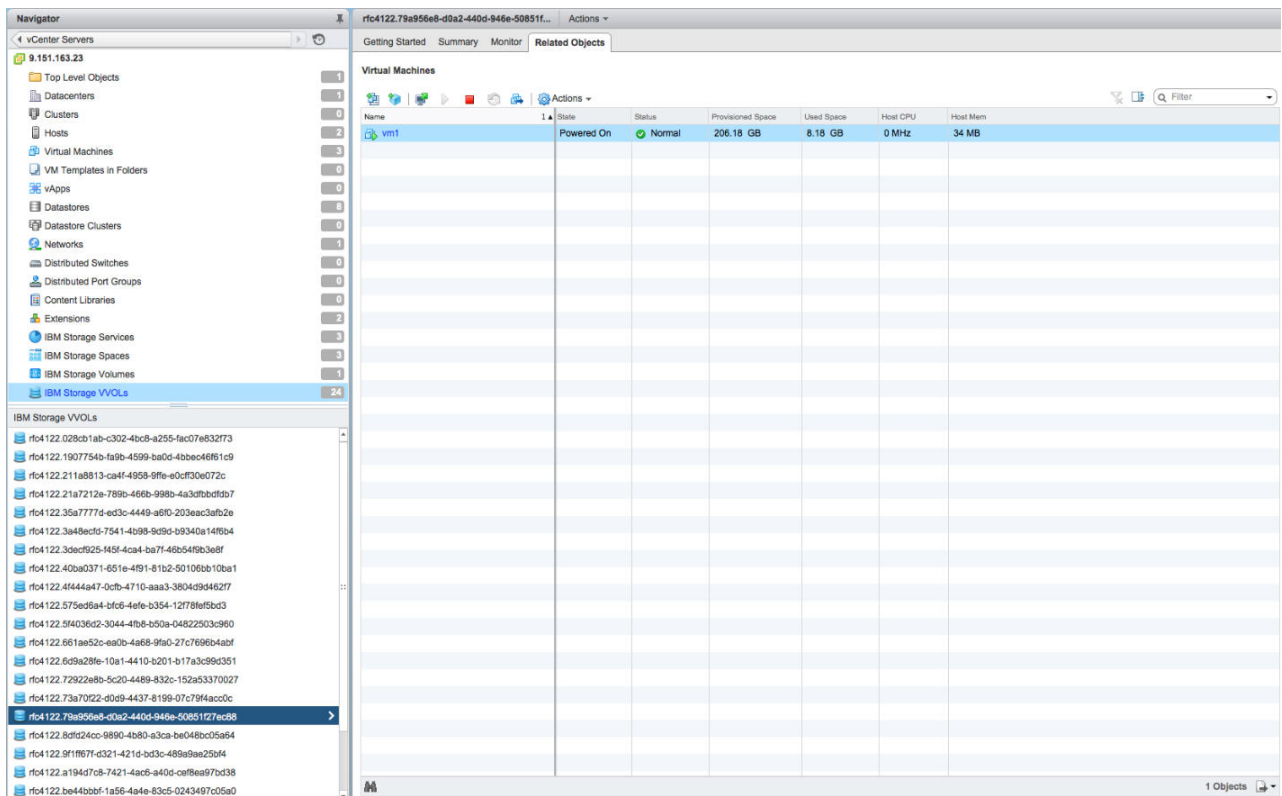


Figure 102. VVol Related Objects tab

Important: Some operations on VVol-based VMs with deployed Spectrum Connect instance, such as hard disk removal, may result in the Invalid Virtual Machine configuration message, displayed by the vWC. This message is not related to the VM functionality and indicates a loss of connectivity between Spectrum Connect and the vCenter server.

Chapter 5. IBM Storage Plug-in for VMware vRealize Orchestrator

The IBM Storage Plug-in for VMware vRealize Orchestrator allows VMware administrators to include IBM storage discovery and provisioning in their vRealize Orchestrator (vRO) automation workflows.

The plug-in package can be downloaded from Spectrum Connect and can then be deployed on the vRealize Orchestrator server. The deployment includes the matching of a unique token key that is set on both servers.

Through vRealize Orchestrator Client, dedicated IBM Storage control elements become available, allowing the issuing of workflows with storage volumes that are attached to the vRealize Orchestrator server.

Rather than issuing volume operations manually and being limited to one manual operation at a time, VMware administrators can preplan and automate storage operations in their virtualized cloud environments, either directly from vRealize Orchestrator or through the VMware vCloud Automation Center (vCAC) platform.

To visualize how this cloud interface is integrated in a virtualized environment, see Chapter 2, “Concept diagram for VMware environment,” on page 99.

Managing integration with vRealize Orchestrator

The IBM Storage Plug-in for the VMware vRealize Orchestrator is used for discovery of the IBM storage resources and provisioning automation workflows in the vRealize Orchestrator (vRO).

Note: In version 3.4.x, the IBM Storage Plug-in for VMware vRealize Orchestrator does not support the DS8000 family storage systems.

To access the vRO management options, add the vRO server (interface) on the **Interfaces** pane. You can then manage the integration with vRO as explained in the following sections:

- “Downloading and installing the plug-in package for vRO” on page 142.
- “Delegating storage services to the vRO server” on page 146.
- “Canceling service delegation to a vRO server” on page 147
- “Regenerating the vRO token” on page 147.




Figure 103. vRO server on the Interfaces pane

Downloading and installing the plug-in package for vRO

To enable the IBM Storage workflows in vRealize Orchestrator, you must first download the IBM Storage plug-in package from Spectrum Connect and install it on the vRealize Orchestrator server.

Before you begin

To allow the IBM Storage Plug-in to securely identify Spectrum Connect and work properly, the SSL certificate that is automatically created on Spectrum Connect upon installation must be replaced with a new one. For more information about how to replace the certificate, see “Managing server certificates” on page 37.

Important: After the IBM Storage Plug-in is installed on vRealize Orchestrator, the Java™ security APIs validate that the hostname received from Spectrum Connect is identical to the common name (CN) value that is in the server certificate.

About this task

The following procedure details how to download, install, and properly configure the IBM Storage Plug-in for VMware vRealize Orchestrator.

Procedure

1. On the **Interfaces** pane, right-click the vRO server, and then select **Modify**. The **vRO Settings** dialog box is displayed.
2. On the bottom of the dialog box, click **Download plug-in package**.

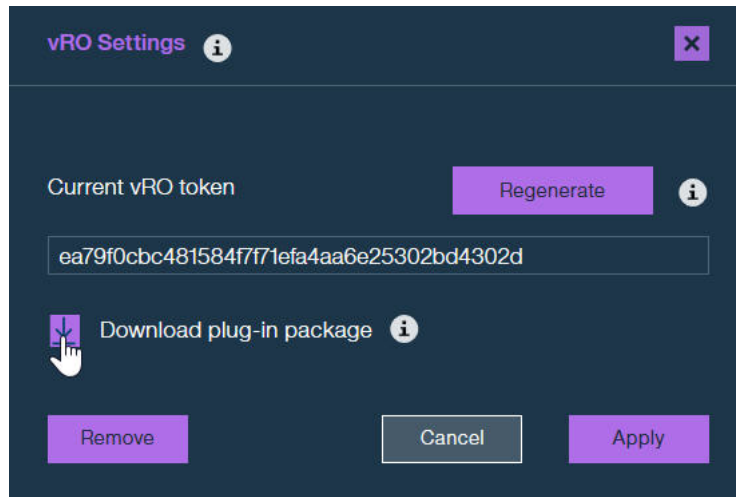


Figure 104. Download plug-in package button

Alternatively, you can download the package from the following directory on Spectrum Connect:

```
/opt/ibm/ibm_spectrum_connect/downloads/static/o11nplugin-ibm-storage.vmoapp
```

3. Copy the current vRO token key from the **Current vRO Token** box.




Figure 105. Current vRO Token

4. Launch the VMware vRealize Orchestrator configuration interface.
5. For vRO 6.x, go to the **Network** tab, click **SSL Trust Manager**; for vRO 7.x, click **Certificates** in the **Manage** category. Then import the Spectrum Connect certificate. The certificate URL format should be: `https://IP_address: TCP port in use`.
6. For vRO 6.x, select **Plug-ins**; for vRO 7.x, click **Manage Plug-Ins** in the **Plug-Ins** category. Then click **Upload and Install** (vRO 6.x), or **Browse > Install** (vRO 7.x). Continue by locating and choosing the downloaded plug-in file. Accept the license agreement. The message 'IBM Storage (3.x.x.x) New plug-in installed.' is displayed. Installation is completed and the IBM Storage plug-in is displayed in the list of vRO plug-ins.




Figure 106. Successful installation of IBM Storage plug-in for vRO 6.x

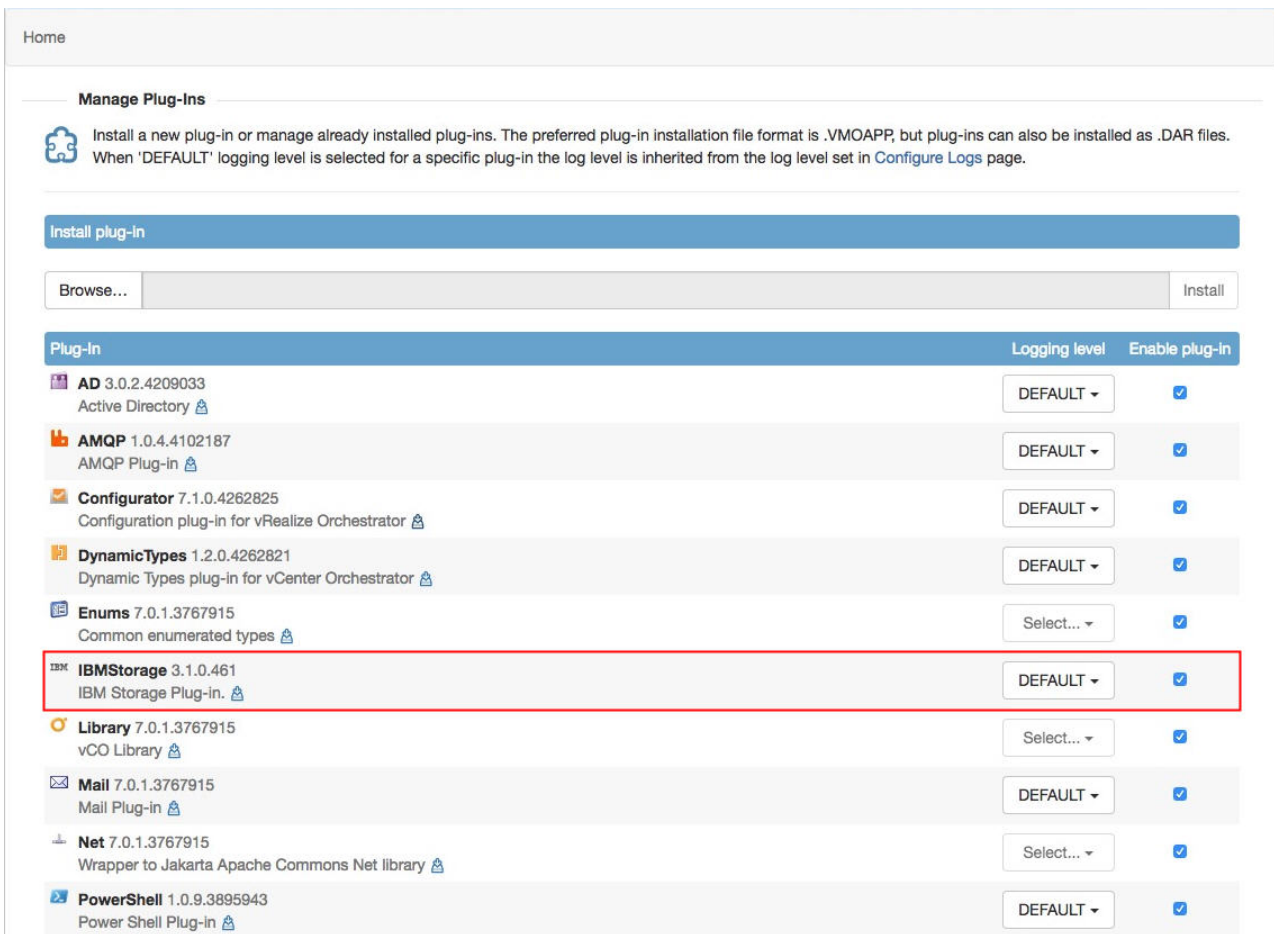


Figure 107. Successful installation of IBM Storage plug-in for vRO 7.x

7. Go to the **Startup Options**, and click **Restart service** to restart the vRO Server service.

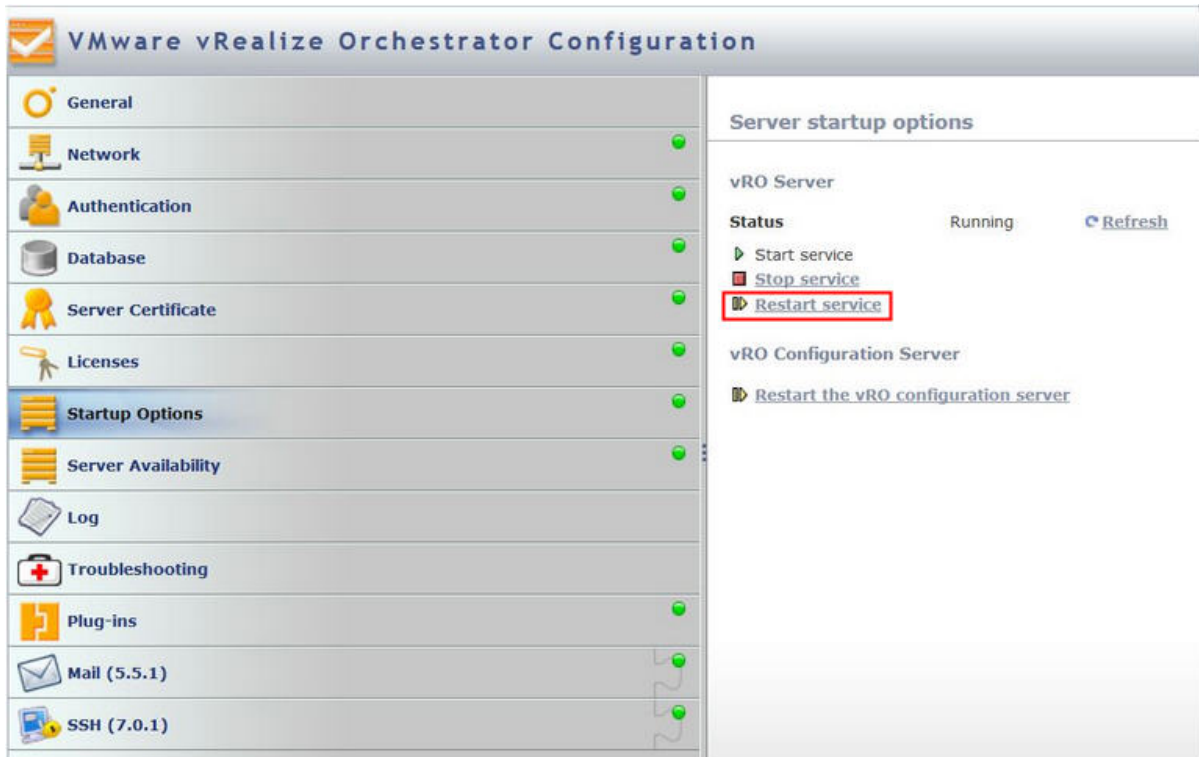


Figure 108. Restarting vRO Server service (vRO 6.x)

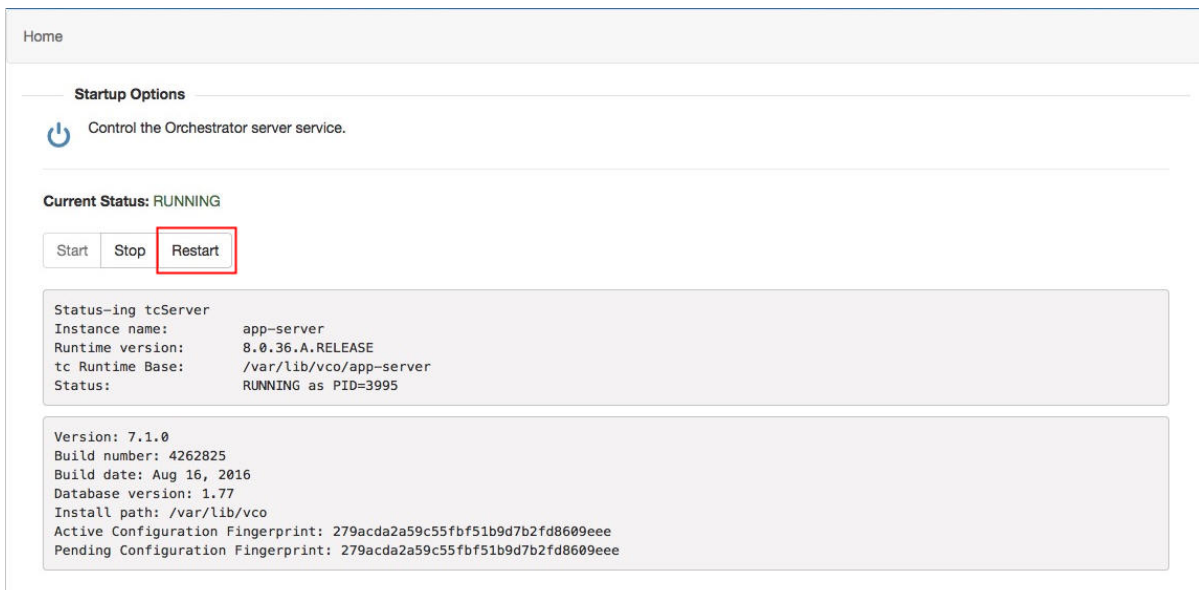


Figure 109. Restarting vRO Server service (vRO 7.x)

8. Start the VMware vRealize Orchestrator client and go to the **Workflows** tab.
9. On the **Workflows** tab, go to **Library > IBM > Storage > Configuration**.

10. Select and run the **Set Server and Token** workflow. The **Start Workflow: Set Server and Token** dialog box is displayed.




Figure 110. Start Workflow: Set Server and Token dialog box

11. Enter the Spectrum Connect IP address or hostname, TCP port in use and the token that you obtained earlier.
12. Click **Submit** to finish the configuration workflow.

Delegating storage services to the vRO server

Before you can use the IBM Storage Plug-in for VMware vRealize Orchestrator on the vRealize Orchestrator (vRO) server side, you must delegate the storage services that you want to make available for vRealize Orchestrator.

About this task

The services and their storage resources that you delegate on Spectrum Connect can be used for issuing volume workflows through vRealize Orchestrator (for more information, see Chapter 5, “IBM Storage Plug-in for VMware vRealize Orchestrator,” on page 141).

Procedure

To delegate storage services to the vRO server:

1. On the **Interfaces** pane, click the vRO server to select it.
2. On the **Spaces/Storage Services** pane, select the storage space from which you want to choose storage services. The available services that reside on the selected storage space are immediately displayed.
3. Right-click on a service that you want to delegate to the vRO server, and then select **Delegate to vRO**, or click the **Attach/Delegate** button on the service. The service color changes to indicate the successful delegation.

You can continue the process by right-clicking available services under the current space.

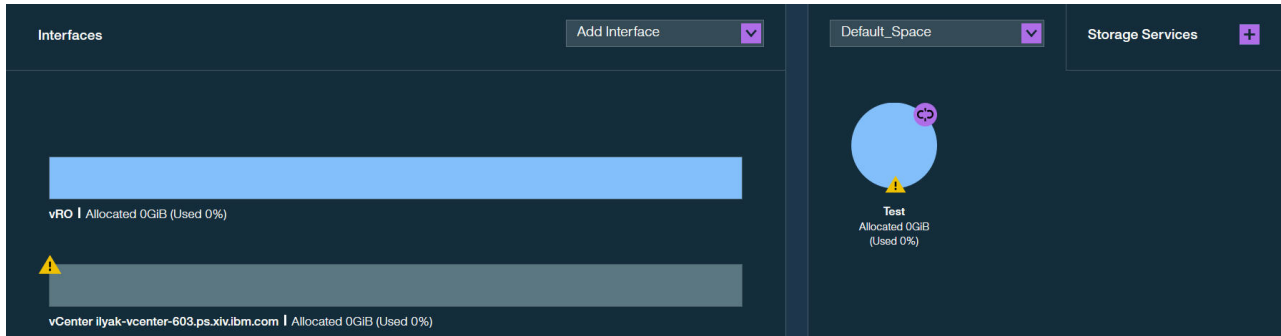


Figure 111. vRO server with delegated services

The vRO server provides indication for the allocated and used storage space.

- Allocated – amount of storage space available on all pools connected to the delegated services.
- Used – amount of storage space used by workflows on all pools connected to the delegated services.

Canceling service delegation to a vRO server

When required, you can cancel a storage service delegation to the vRO server.

About this task

Storage services, which delegation has been canceled, and their resources (pools) are removed from the vRO inventory and cannot be used for workflows. Any scheduled workflows, involving these elements, will fail to run.

Procedure

1. On the **Interfaces** pane, click the vRO server. The services that are currently delegated to the vRO server are highlighted on the **Spaces/Storage Services** pane.
2. Right-click on a service which delegation you want cancel, and then select **Cancel delegation to vRO** , or click the **Detach/Cancel Delegation** button on the service. The service color changes to indicate the successful detachment. You can continue the process by right-clicking delegated services under the current space.

Regenerating the vRO token

The vRO token used during installation of the IBM Storage Plug-in for vRO can be regenerated.

About this task

If the communication link between the vRealize Orchestrator server and Spectrum Connect has been compromised, you can regenerate the vRO token and reconfigure the IBM Storage Plug-in for vRO. For the plug-in installation instructions, see “Downloading and installing the plug-in package for vRO” on page 142.

Procedure

To regenerate the vRO token:

1. On the **Interfaces** pane, right-click the vRO server, and then select **Modify**. The **vRO Settings** dialog box is displayed.
2. Click the **Regenerate** button. A new vRO token is regenerated.

Using the IBM Storage Plug-in for VMware vRealize Orchestrator

Use the IBM Storage Plug-in for VMware vRealize Orchestrator to include IBM Storage discovery and provisioning in your vRealize Orchestrator (vRO) automation workflows.

After the IBM Storage Plug-in is deployed (see “Downloading and installing the plug-in package for vRO” on page 142), the IBM storage objects become available in the vRO, as detailed in the table below.

Table 24. IBM storage objects and events in vRO

Object	Attribute	Type
StorageSpace	name	String
	description	String
	numServices	long
	numVolumes	long
	physicalSize	long
	physicalFree	long
StorageService	name	String
	description	String
	physicalSize	long
	physicalFree	long
	capabilityValues	String
	maxResourceLogicalFree	long
StorageVolume	name	String
	spaceName	String
	serviceName	String
	serviceCompliance	String
	poolName	String
	domainName	String
	array	String
	storageModel	String
	logicalCapacity	long
	usedCapacity	long

The Orchestrator workflows, supported by the IBM storage objects are as follows:

- Create and map a volume
- Delete a volume
- Extend a volume
- Map a volume
- Unmap a volume

Follow these guidelines for vRO workflows:

- Volume creation: the allowed volume sizes are set in whole numbers. Any number after a decimal point is ignored by the vRO. For example, when the volume size is set to 1.6 GB, the 1 GB volume is created.
- Volume extension: the allowed volume sizes are set in whole numbers. Any number after a decimal point is ignored by the vRO. For example, when the volume size is extended to 3.6 GB, the volume size is set to 3 GB.
- Volume mapping is performed by passing initiators to the workflow. If a host definition at the storage system includes two initiators, only one of them is passed by the workflow, when it is run. Also, a volume must be mapped to all hosts using the same LUN.

The following figures display the IBM Storage workflows and elements in vRO.

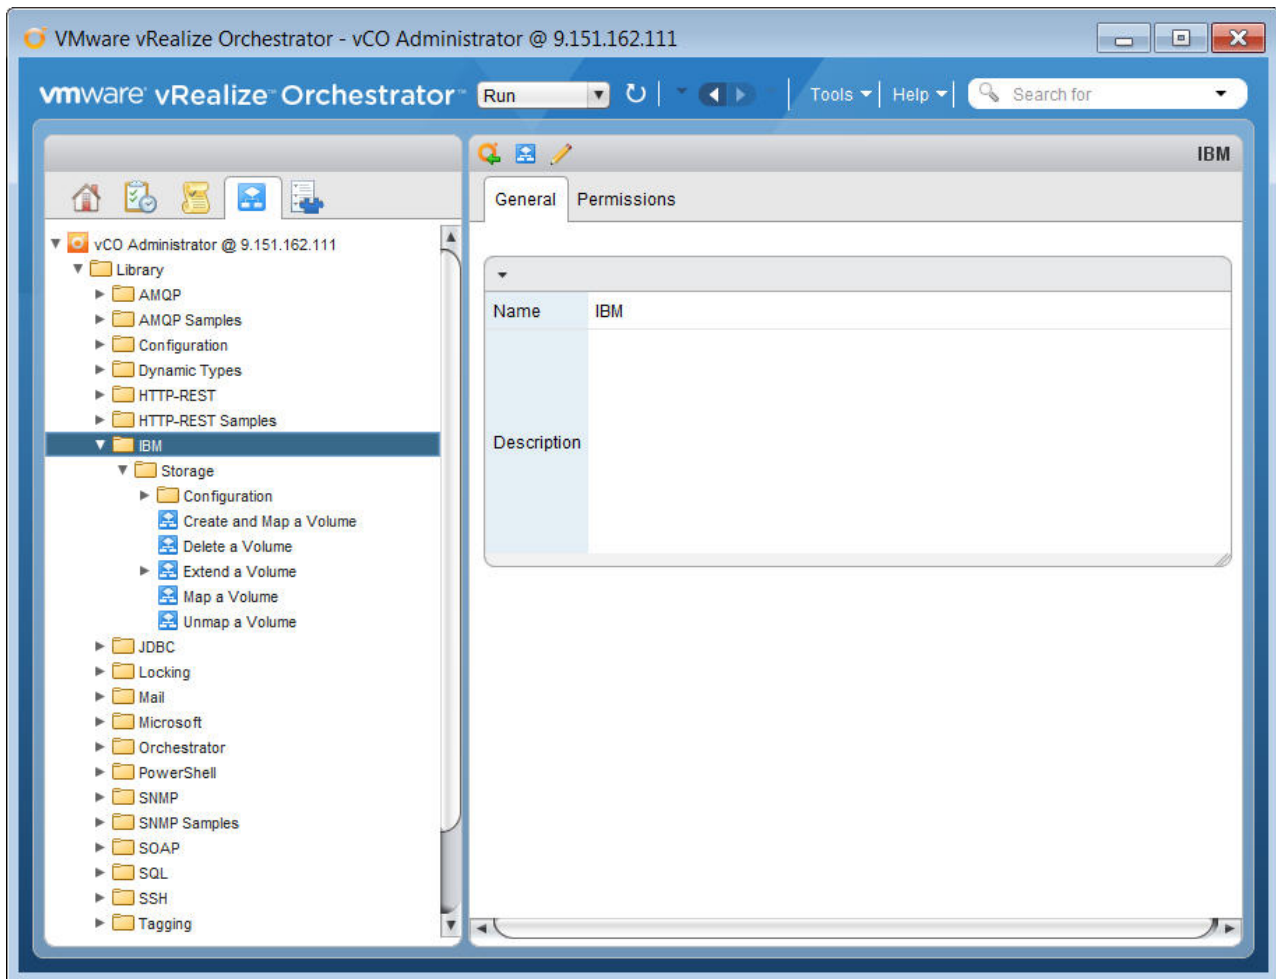


Figure 112. vRealize Orchestrator – available workflows – General tab

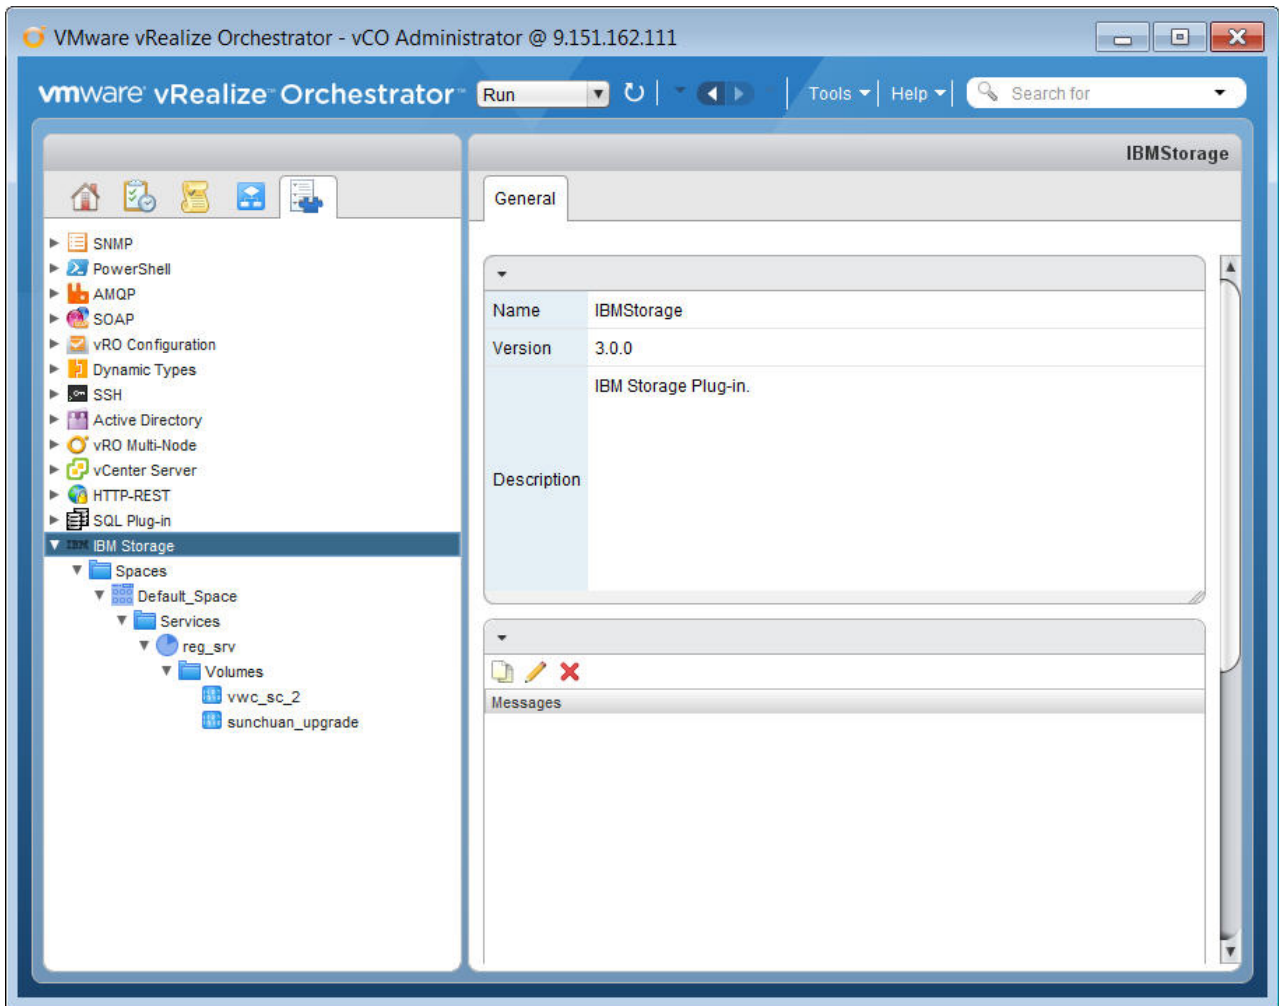


Figure 113. vRealize Orchestrator – available objects

For information about how to initiate workflows in vRealize Orchestrator, refer to the relevant VMware documentation.

Chapter 6. IBM Storage Management Pack for VMware vRealize Operations Manager

The IBM Storage Management Pack for VMware vRealize Operations Manager allows Operations Manager users to obtain comprehensive monitoring information about the IBM storage resources that are utilized in their virtualized environment.

The management pack can be downloaded from Spectrum Connect and can then be deployed on the vRealize Operations Manager server.

After a VMware vRealize Operations Manager server is registered on Spectrum Connect that is configured with storage systems, storage spaces, services, and vRealize servers, storage-related data is pushed to the Operations Manager server every five minutes by default.

The dedicated IBM storage system adapter deployed on the vRealize Operations Manager server enables monitoring the supported IBM storage system via the vROps Manager. This adapter reports the storage-related information, such as monitoring data of all logical and physical elements, covering storage systems, storage domains, storage pools, volumes, hosts, modules, target ports, disks, health status, events, thresholds, and performance. It also provides the dashboards that display detailed status, statistics, metrics, and analytics data alongside hierarchical flowcharts with graphic representation of IBM storage system elements.

Relationships between the IBM storage elements (storage systems, ports, storage pools, volumes, host, host initiator, modules, domain) and datastores, virtual machines, and hosts are displayed graphically in a drill-down style, providing VMware administrators a complete and up-to-date picture of their utilized storage resources.

To visualize how this cloud interface is integrated in a virtualized environment, see Chapter 2, “Concept diagram for VMware environment,” on page 99.

Managing integration with vRealize Operations Manager

Before you can use the IBM Storage Management Pack for VMware vRealize Operations Manager, you must set a connection to at least one vRealize Operations Manager (vROps) server, and then define which storage systems should be monitored in vROps.

After a vROps server connection is defined and storage systems are associated with the vROps server, detailed monitoring information for these storage systems becomes available in vROps (for more information, see Chapter 6, “IBM Storage Management Pack for VMware vRealize Operations Manager”).

Note: In version 3.4.x, the IBM Storage Management Pack for VMware vRealize Operations Manager does not support the DS8000 family storage systems.

To access these options, go to the **Monitoring** pane of Spectrum Connect GUI. You can then manage the integration with vROps as explained in the following sections:

- “Downloading the vROps management package” on page 152.

- “Deploying the management package on vROps” on page 153.
- “Connecting the vROps server to Spectrum Connect” on page 154.
- “Controlling storage system monitoring on the vROps server” on page 155.

Figure 114. Monitoring pane

The vROps integration, except for the management package operations, can be performed from the CLI as well, as explained in “CLI – Managing integration with vRealize Operations Manager” on page 87.

Downloading the vROps management package

IBM Spectrum Connect provides management package in the form of a PAK file which can be deployed on the vRealize Operations Manager.

About this task

Although vROps can display IBM Storage information even without the management package (storage adapter) installation, the IBM Storage adapter is required for displaying the dedicated dashboards, graphic icons, and user-friendly attribute names for the IBM storage elements. The adapter is provided through a PAK file that you need to download, as described in the following procedure.

Procedure

To download the PAK file from Spectrum Connect:

1. Go to **Monitoring** pane of the Spectrum Connect GUI. The **Set vROps Server** dialog box is displayed.

2. On the bottom of the dialog box, click **Download PAK file**.



Figure 115. Download PAK File button

Alternatively, you can download the package from the following directory on Spectrum Connect:

```
/opt/ibm/ibm_spectrum_connect/downloads/static/IBM_Storage_Adapter-3.0.0-x.pak, where x designates the current build number.
```

3. Save the file to your computer to later upload it to the vRealize Operations Manager.

What to do next

See “Deploying the management package on vROps.”

Deploying the management package on vROps

After the management package is downloaded to the computer, it must be deployed on the vRealize Operations Manager.

About this task

The management package must be deployed on the vROps, as described below.

Procedure

To deploy the management package on the vROps:

1. After the management package is downloaded to the computer, access the vRealize Operations Manager administrative web console using `https://hostname` or IP address of the vROps UI.
2. Go to **Administration > Solutions**.
3. In the **Solutions** pane, click the plus sign on the top toolbar to add a new management package. The **Add Solution** dialog box is displayed.
4. In the **Add Solution** dialog box, click **Browse** and select the management package downloaded from Spectrum Connect. After that, click **Upload** to start deployment.




Figure 116. Deploying the management package on the vROps

The IBM license agreement is displayed.

5. Accept the IBM license agreement, click **OK** to continue. The confirmation message is displayed.
6. Click **OK** to confirm the update. The vROps Manager displays a confirmation message after the management package is deployed successfully.

No additional configuration of the management package is required on the vROps. Under certain conditions, its status may appear as **Not configured** on the vROps. You can disregard this information.

Connecting the vROps server to Spectrum Connect

After the management package is successfully deployed and described, you must add the vROps Manager server to Spectrum Connect.

About this task

The vRealize Operations Manager server must be connected to Spectrum Connect, as explained below.

Procedure

1. Go to the **Monitoring** pane.
2. Enter IP address or FQDN of the vRealize Operations Manager server, user name, password and select the check box to confirm you have installed the PAK file on the vROps Manager server; click **Apply** to save the settings.

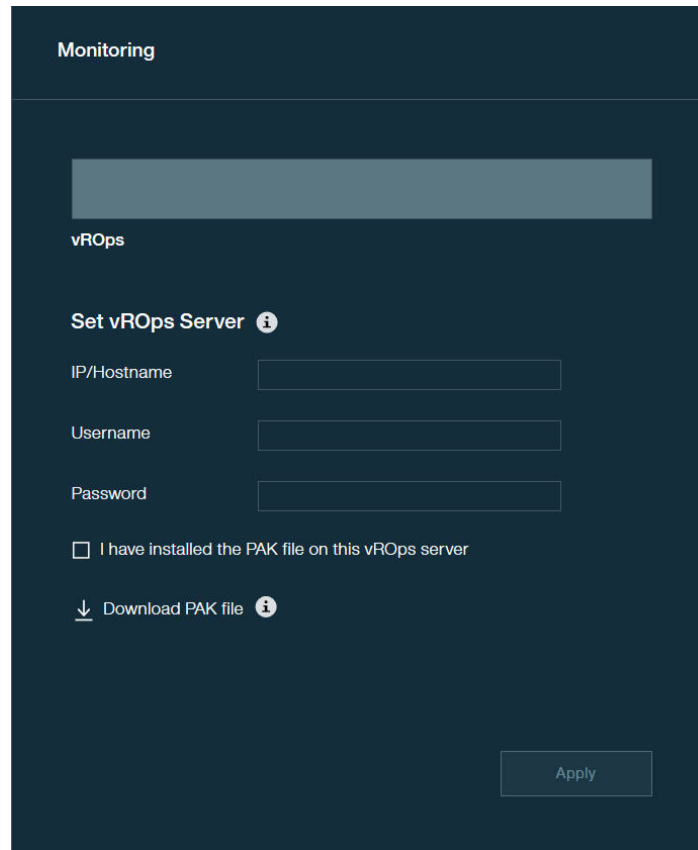


Figure 117. Adding the vROps server to Spectrum Connect

If the vROps server connection is successful, its color changes from gray to blue.

What to do next

See “Controlling storage system monitoring on the vROps server.”

Controlling storage system monitoring on the vROps server

The IBM storage systems connected to Spectrum Connect must be added to the vROps Manager as well to enable their monitoring.

About this task

To enable monitoring of the IBM storage systems, they must be added to the vROps server. You can also remove the systems that do not require monitoring by the vROps.

Procedure

To monitor the storage systems on the vROps server:

1. On the **Monitoring** pane, click the vROps server which you want to use for monitoring.
2. In the **Storage Systems** pane, right-click a storage system that you intend to monitor, and select **Start vROps monitor...**, or click the **Start Monitoring** button on the storage system. The monitored system color changes to indicate the

connection to the vROps server. Spectrum Connect starts pushing the information to vRealize Operations Manager, using HTTP post requests.

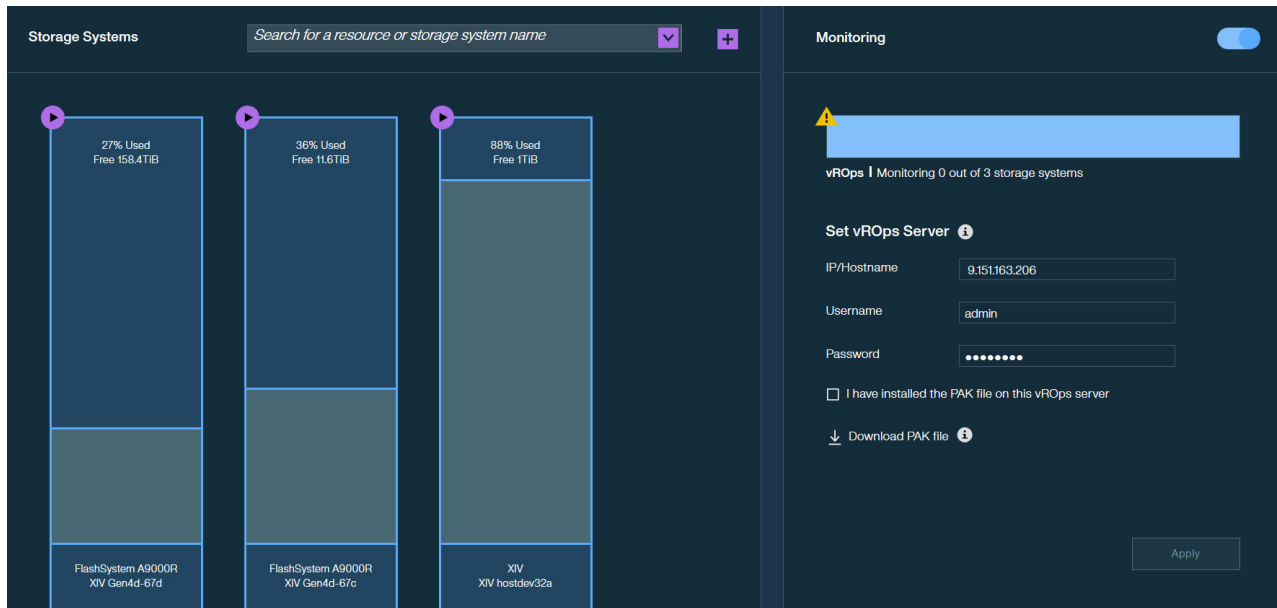


Figure 118. Storage system monitored by the vROps server

You can stop monitoring a storage system by clicking the **Stop Monitoring** button on the monitored system.

You can prevent the vROps server from collecting performance data from the monitored storage systems by toggling the **Activate** switch (). The legend below the vROps server specifies how many storage systems are being monitored by the server out of total number of systems defined on Spectrum Connect. The color of the vROps server changes from dark to bright blue in accordance with the number of monitored systems.

Using the IBM Storage Management Pack for VMware vRealize Operations Manager

Use the IBM Storage Management Pack for VMware vRealize Operations Manager to obtain comprehensive monitoring information about the IBM storage resources that are utilized in your virtualized environment.

About this task

After successfully configuring Spectrum Connect with vRealize Operations Manager, it periodically starts sending the storage system information to vRealize Operations Manager. You can view the detailed IBM storage dashboards, together with the graphical relationships between the storage elements (storage systems, ports, storage pools, volumes) and virtual elements (datastores, virtual machines, hosts) in a drill-down interactive style. Three main dashboards are available for IBM storage systems:

- “Overview dashboard” on page 159
- “Performance dashboard” on page 161
- “Top 10 dashboard” on page 169

In addition, you can:

- Monitor storage system resources. See “Displaying the overview of IBM storage objects” on page 170.
- Define thresholds and alerts for storage resources and metrics in the Overview and Performance dashboards. See “Defining thresholds and alerts for storage objects” on page 171.

Note: Not all informational-level severity events are reported by Spectrum Connect to vRealize Operations Manager. All other severity events (Warning, Minor, Major, Critical) are reported.

Procedure

To view the IBM storage dashboards, complete the following steps.

1. Log in to vRealize Operations Manager user interface.
2. Click **Dashboard List > IBM Storage** and select the required dashboard. The selected dashboard is displayed. The following icons are used to represent

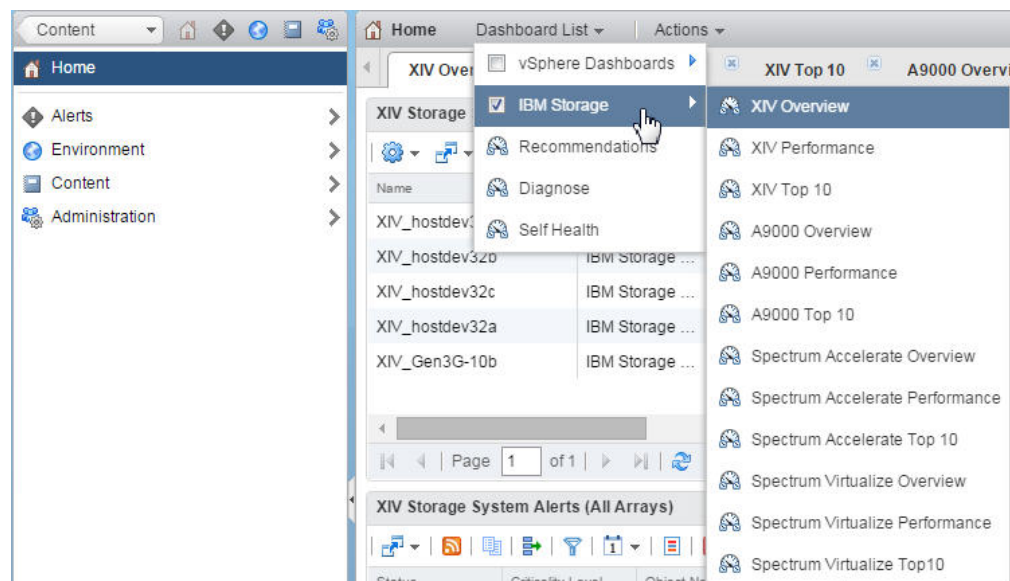


Figure 119. vROps GUI – IBM STORAGE option

the IBM storage elements.

Table 25. IBM Storage Icons in vROps.




Icon	Description
	IBM storage adapter
	IBM storage system
	IBM storage system, running Spectrum Virtualize

Table 25. IBM Storage Icons in vROps (continued).

















Icon	Description
	IBM disk
	IBM MDisk
	IBM domain
	IBM host
	IBM Host Fibre Channel (FC) initiator
	IBM host iSCSI initiator
	IBM module
	Flash canister
	Flash enclosure
	Flash card
	IBM module FC port
	IBM module iSCSI port
	IBM pool
	IBM child pool
	IBM volume

Table 25. IBM Storage Icons in vROps (continued).

Icon	Description
	IBM I/O group

Overview dashboard

The Overview dashboard presents relationships between all virtual elements and storage elements that are in use.

Procedure

To display IBM storage resource overview:

1. In the Overview dashboard, click a storage system in the **Storage Systems** pane. The list of all the resources related to the selected storage system are displayed in the right pane of the Overview dashboard. The resources include virtual machines, host systems, datastores, volumes, pools, storage systems (arrays), hosts, FC initiators, FC ports, iSCSI initiators, iSCSI ports, modules, flash canisters, flash enclosures, flash cards and disks.

The alert widget is detailed in “Using the alert widget” on page 160.

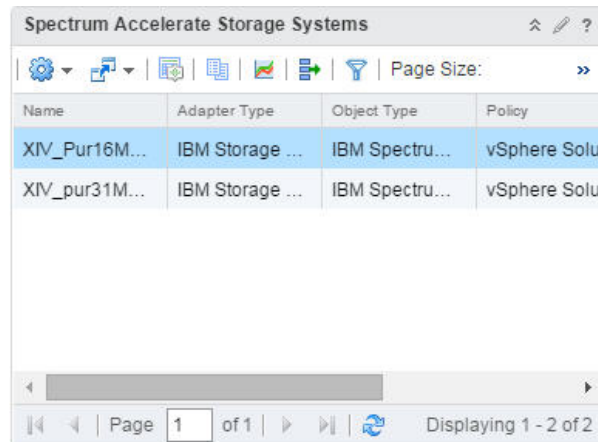


Figure 120. IBM Spectrum Accelerate Storage Systems pane

2. Move the mouse pointer over a resource element or click on it to select it. A tooltip is displayed, detailing the element name and its health score. The health score is calculated automatically by the vROps Manager, according to the number of alarms and statistic information. However, disk health score is based on the *health.requires service* metric.

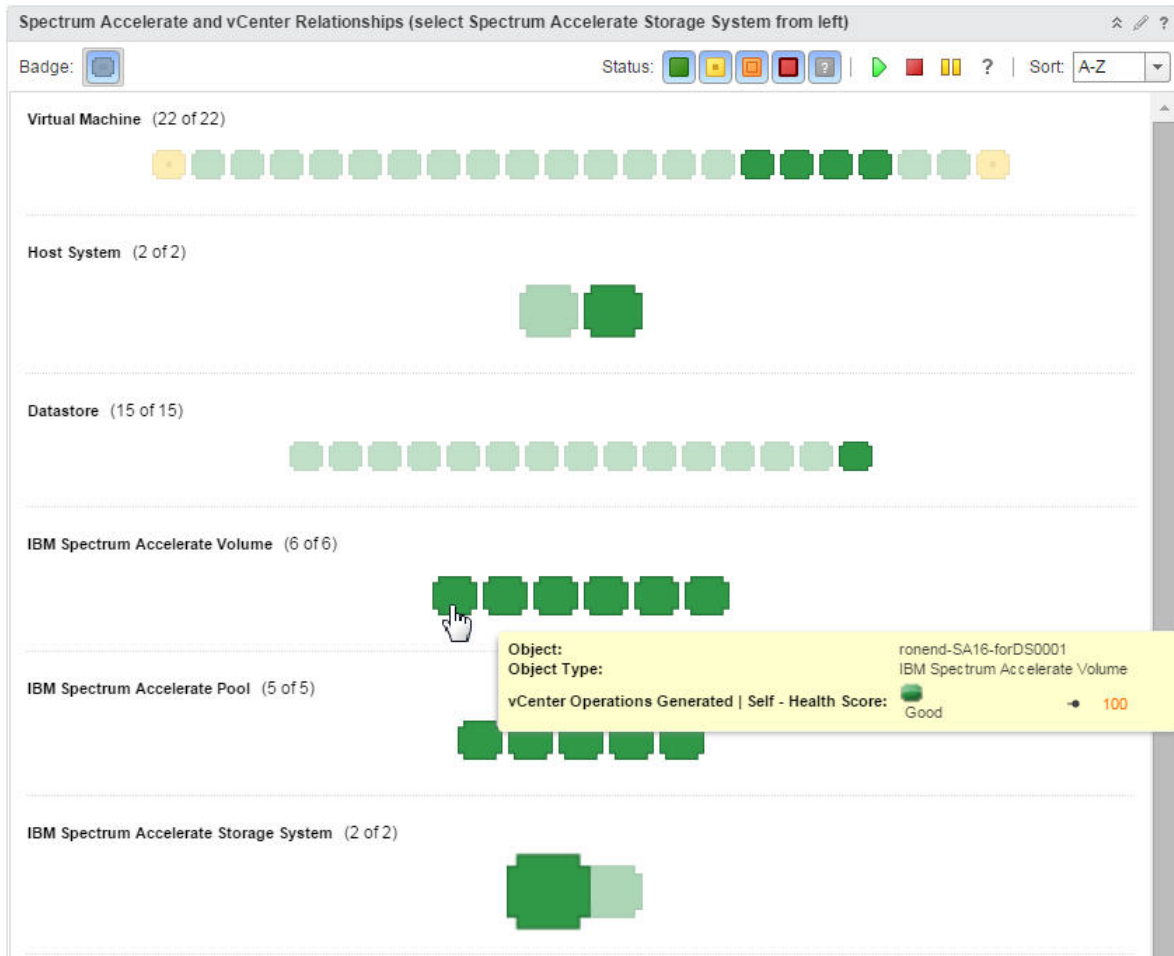


Figure 121. IBM Spectrum Accelerate volume health status

3. Double-click on a selected resource element to display the resource details (health tree, metrics, etc). These are detailed in the Performance dashboard below.

Using the alert widget

The Storage System Alerts (All Arrays) widget is located at the lower-left pane of the Overview dashboard. This widget displays alerts generated by all storage systems, monitored by the vROps server.

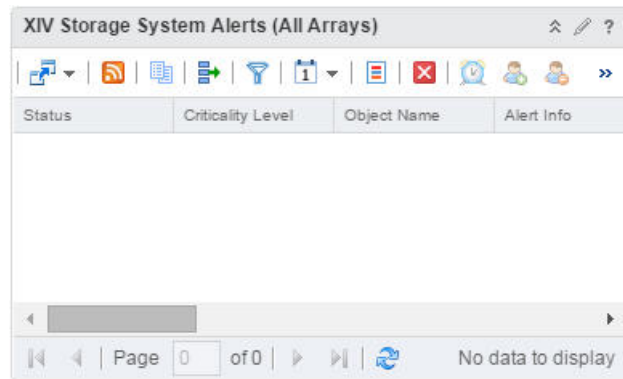


Figure 122. Alert widget

By default, the events are pushed by the IBM storage adapter to the vRealize Operations Manager every 10 minutes.

Performance dashboard

Performance dashboard provides health and performance information for the IBM storage resources.

About this task

Performance information presented in the dashboard is collected for a time period defined by the `vcops_push_interval` parameter in the `vcops_config.ini` file (the default time period is 5 minutes). This period can be changed, as explained in “Adjusting system update interval” on page 75.

Procedure

To display performance information:

1. In the **Virtual Machines** pane of the Performance dashboard, locate a relevant virtual machine and select it.
2. In the **Storage System and vCenter Relationships** pane, select an object, which performance you intend to monitor.

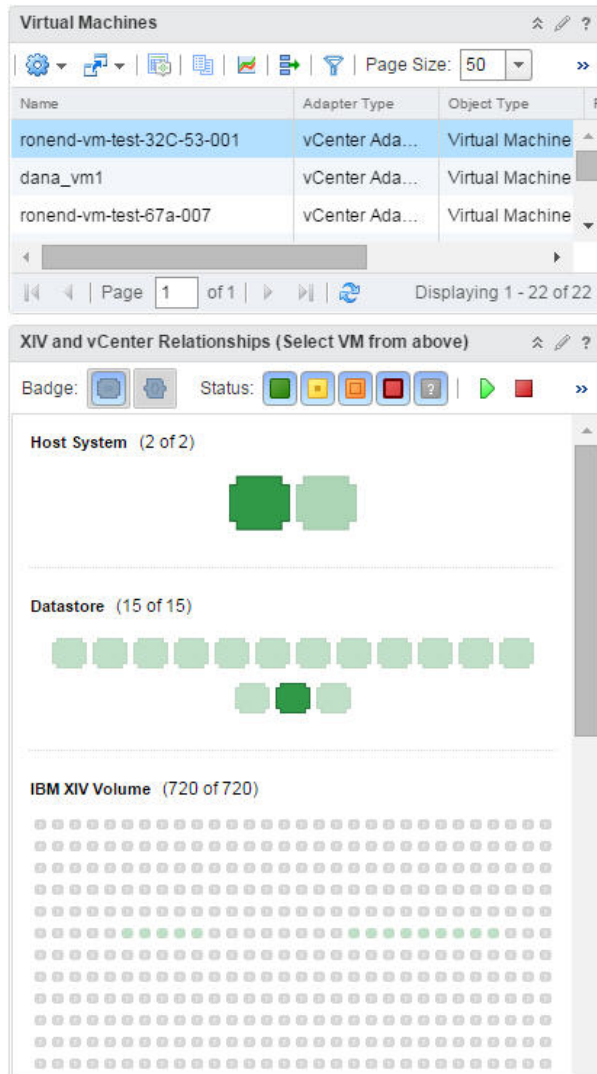


Figure 123. Virtual Machines and XIV and vCenter Relations panes

The **HEALTH TREE** pane displays the selected element and its relation to other storage resources in an hierarchical manner.




Figure 124. HEALTH TREE pane

3. In the **HEALTH TREE** pane, select a resource element to display all relevant performance metrics in the **METRIC SELECTOR** pane. Different metric types are available for different resource elements, as detailed in “Performance metrics.”
4. Select a metrics of an element to display its metric graph in the **METRIC GRAPH** pane.

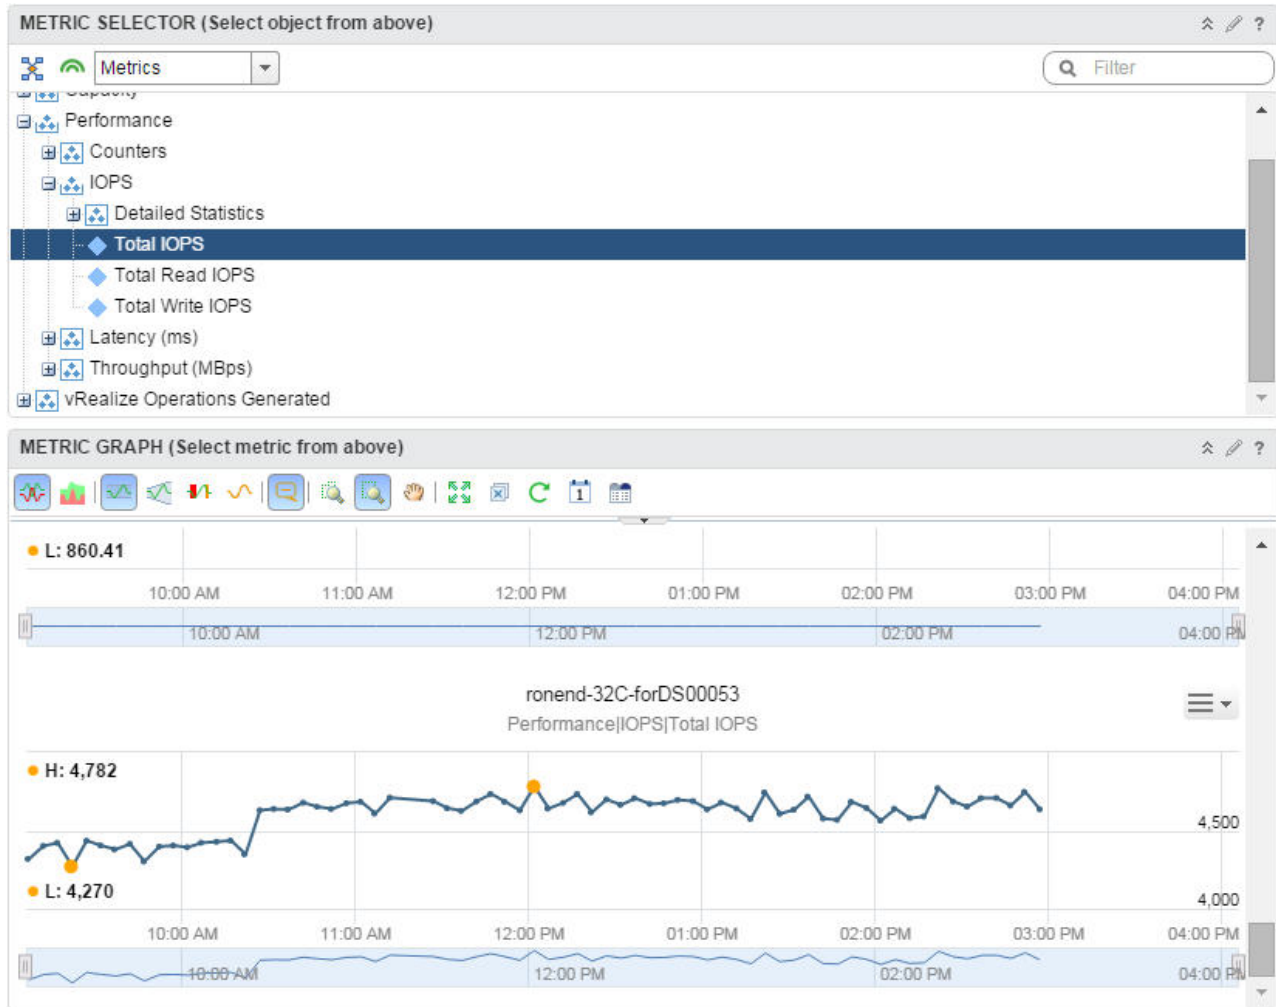


Figure 125. METRIC SELECTOR and METRIC GRAPH panes of the Performance dashboard

Performance metrics

Metrics data available in the Performance dashboard provides health and performance information for the IBM storage objects.

Different metric types are available for different storage elements, as detailed in tables below.

Table 26. Capacity metrics

Metrics	Description	Relevant storage object
Hard (GB)	Hard size (actual physical capacity) of the storage object	Volume, Pool, Array
Soft (GB)	Soft size (maximum capacity) of the storage object, as seen by the hosts	
Utilization Hard (%)	Utilization ratio of the hard capacity	
Utilization Soft (%)	Utilization ratio of the soft capacity	
Free Hard (GB)	Free hard size of the storage system (array)	Array
Free Soft (GB)	Free soft size of the storage system	
Over Provisioning Ratio (%)	The ratio between virtual capacity and real capacity of the pool	Pool
Compression Ratio (%)	Volume compression ratio, calculated as 1-(compressed volume size/uncompressed volume size)	Volume
Compression Saving (GB)	Volume compression savings, calculated as (number of 17GB uncompressed partitions in use – number of 17GB partitions in use)/1000	
Effective (GB)	Potential capacity of the storage object	IBM FlashSystem A9000/A9000R Array or Volume
Physical (GB)	Actual physical capacity of the storage object	
Utilization Effective (%)	Utilization ratio of the effective capacity	
Utilization Physical (%)	Utilization ratio of the physical capacity	
Thin Provisioning Savings (%)	Total savings achieved due to thin provisioning	
Free Effective (GB)	Free effective size of the storage system	IBM FlashSystem A9000/A9000R Array
Free Physical (GB)	Free physical size of the storage system	
Compression Factor	Ratio between unique data (after deduplication) and actual stored data	
Data Reduction Savings (%)	Total savings achieved due to data reduction	
Total Stored (GB)	Actual size of data currently stored on the system	
Total Written (GB)	Total amount of data written by hosts before compression or deduplication	
Estimated Compression Factor	Ratio between unique data and physical data	
Estimated Deduplication Factor	Ratio between written and unique data	
Estimated Minimum Delete Size (GB)	Amount of unique written data divided by estimated compression factor	
Estimated Data Reduction Savings (%)	Estimated savings due to data reduction	

Table 27. Health metrics

Metrics	Description	Relevant storage object
Connected	Connection status of the object	Host, Host FC Initiator, Host iSCSI Initiator
Status	Health status of the module	Module, IBM Spectrum Virtualize MDisk

Table 27. Health metrics (continued)

Metrics	Description	Relevant storage object
Requires Service	Health status of the disk. The REPLACE value indicates disk failure.	Disk
Online	Connection status of the object. For IBM Spectrum Virtualize Host, the status can be as follows: <ul style="list-style-type: none"> • 0 – Offline • 5 – Excluded • 10 – Degraded • 20 – Is Online 	Module FC Port, Module iSCSI Port, IBM Spectrum Virtualize Node FC Port, IBM Spectrum Virtualize Node iSCSI Port, IBM Spectrum Virtualize Host

Table 28. Counter metrics

Metrics	Description	Relevant storage object
Hosts Count	Total number of hosts connected to the storage system	Array, IBM Spectrum Virtualize I/O Group
Volumes Count	Total number of volumes existing in the storage system	
Mirror Relations Count	Total number of mirror relations (master or slave) existing in the storage system	Array
Pools Count	Total number of pools existing in the storage system	
Snapshots Count	Total number of snapshots existing in the storage system	
Volumes and Snapshots Count	Total number of volumes and snapshots existing in the storage system	
Nodes Count	Total number of nodes in the I/O group	IBM Spectrum Virtualize I/O Group

Table 29. Performance metrics

Metrics	Description	Relevant storage object
Total IOPS	Total number of IOPS performed by the object	Array, Host FC Initiator, Host iSCSI Initiator, Module FC Port, Module iSCSI Port
Total Read IOPS	Total number of read IOPS performed by the object	
Total Write IOPS	Total number of write IOPS performed by the object	
Read Hit Large IOPS	Number of IOPS for 64–512 KB packets read from cache	
Read Hit Medium IOPS	Number of IOPS for 8–64 KB packets read from cache	
Read Hit Small IOPS	Number of IOPS for 0–8 KB packets read from cache	
Read Hit Very Large IOPS	Number of IOPS for over 512 KB packets read from cache	
Read Miss Large IOPS	Number of IOPS for 64–512 KB packets read from disk	
Read Miss Medium IOPS	Number of IOPS for 8–64 KB packets read from disk	
Read Miss Small IOPS	Number of IOPS for 0–8 KB packets read from disk	
Read Miss Very Large IOPS	Number of IOPS for over 512 KB packets read from disk	
Write Hit Large IOPS	Number of IOPS for 64–512 KB packets written to cache	
Write Hit Medium IOPS	Number of IOPS for 8–64 KB packets written to cache	
Write Hit Small IOPS	Number of IOPS for 0–8 KB packets written to cache	
Write Hit Very Large IOPS	Number of IOPS for over 512 KB packets written to cache	
Write Miss Large IOPS	Number of IOPS for 64–512 KB packets written to disk	
Write Miss Medium IOPS	Number of IOPS for 8–64 KB packets written to disk	
Write Miss Small IOPS	Number of IOPS for 0–8 KB packets written to disk	
Write Miss Very Large IOPS	Number of IOPS for over 512 KB packets written to disk	
Average Latency	Average response time	
Average Write Latency	Average response time of a write operation	
Average Read Latency	Average response time of a read operation	
Read Hit Large Latency	Response time of cache read operations for 64–512 KB packets	
Read Hit Medium Latency	Response time of cache read operations for 8–64 KB packets	
Read Hit Small Latency	Response time of cache read operations for 0–8 KB packets	
Read Hit Very Large Latency	Response time of cache read operations for over 512 KB packets	

Table 29. Performance metrics (continued)

Metrics	Description	Relevant storage object
Read Memory Hit Large Latency	Response time of DRAM cache read operations for 64–512 KB packets	Array
Read Memory Hit Medium Latency	Response time of DRAM cache read operations for 8–64 KB packets	
Read Memory Hit Small Latency	Response time of DRAM cache read operations for 0–8 KB packets	
Read Memory Hit Very Large Latency	Response time of DRAM cache read operations for over 512 KB packets	
Read Miss Large Latency	Response time of disk read operations for 64–512 KB packets	Array, Host FC Initiator, Host iSCSI Initiator, Module FC Port, Module iSCSI Port
Read Miss Medium Latency	Response time of disk read operations for 8–64 KB packets	
Read Miss Small Latency	Response time of disk read operations for 0–8 KB packets	
Read Miss Very Large Latency	Response time of disk read operations for over 512 KB packets	
Write Hit Large Latency	Response time of cache write operations for 64–512 KB packets	
Write Hit Medium Latency	Response time of cache write operations for 8–64 KB packets	
Write Hit Small Latency	Response time of cache write operations for 0–8 KB packets	
Write Hit Very Large Latency	Response time of cache write operations for over 512 KB packets	
Write Miss Large Latency	Response time of disk write operations for 64–512 KB packets	
Write Miss Medium Latency	Response time of disk write operations for 8–64 KB packets	
Write Miss Small Latency	Response time of disk write operations for 0–8 KB packets	
Write Miss Very Large Latency	Response time of disk write operations for over 512 KB packets	
Total Throughput	Total bandwidth	
Total Read Throughput	Total bandwidth used by read operations	
Total Write Throughput	Total bandwidth used by write operations	
Read Hit Large Throughput	Bandwidth used by cache read operations for 64–512 KB packets	
Read Hit Medium Throughput	Bandwidth used by cache read operations for 8–64 KB packets	
Read Hit Small Throughput	Bandwidth used by cache read operations for 0–8 KB packets	
Read Hit Very Large Throughput	Bandwidth used by cache read operations for over 512 KB packets	

Table 29. Performance metrics (continued)

Metrics	Description	Relevant storage object
Read Memory Hit Large Throughput	Bandwidth used by DRAM cache read operations for 64–512 KB packets	Array
Read Memory Hit Medium Throughput	Bandwidth used by DRAM cache read operations for 8–64 KB packets	
Read Memory Hit Small Throughput	Bandwidth used by DRAM cache read operations for 0–8 KB packets	
Read Memory Hit Very Large Throughput	Bandwidth used by DRAM cache read operations for over 512 KB packets	
Read Miss Large Throughput	Bandwidth used by disk read operations for 64–512 KB packets	Array, Host FC Initiator, Host iSCSI Initiator, Module FC Port, Module iSCSI Port
Read Miss Medium Throughput	Bandwidth used by disk read operations for 8–64 KB packets	
Read Miss Small Throughput	Bandwidth used by disk read operations for 0–8 KB packets	
Read Miss Very Large Throughput	Bandwidth used by disk read operations for over 512 KB packets	
Write Hit Large Throughput	Bandwidth used by cache write operations for 64–512 KB packets	
Write Hit Medium Throughput	Bandwidth used by cache write operations for 8–64 KB packets	
Write Hit Small Throughput	Bandwidth used by cache write operations for 0–8 KB packets	
Write Hit Very Large Throughput	Bandwidth used by cache write operations for over 512 KB packets	
Write Miss Large Throughput	Bandwidth used by disk write operations for 64–512 KB packets	
Write Miss Medium Throughput	Bandwidth used by disk write operations for 8–64 KB packets	
Write Miss Small Throughput	Bandwidth used by disk write operations for 0–8 KB packets	
Write Miss Very Large Throughput	Bandwidth used by disk write operations for over 512 KB packets	
Aborts	Total number of I/Os aborted during the current mirroring operation	
Failures	Total number of I/Os failed during the current mirroring operation	

Table 29. Performance metrics (continued)

Metrics	Description	Relevant storage object
Read Operations	Number of read operations per second	IBM Spectrum Virtualize Vdisk
Write Operations	Number of write operations per second	
Read Blocks	Number of blocks read per second (in MB)	
Write Blocks	Number of blocks written per second (in MB)	
Average Read Response Time	Average response time for disk reads per read operation	
Average Write Response Time	Average response time for disk writes per write operation	
Worst Read Response Time Since Last Statistics Collection	Worst response time for disk reads within the last statistics collection sample period	
Worst Write Response Time Since Last Statistics Collection	Worst response time for disk writes within the last statistics collection sample period	
Average Transfer Response Time	Average transfer latency for both read and write operations	
System CPU Utilization	CPU busy percentage for I/O process cores	
Compression CPU Utilization	CPU busy percentage for compression process cores	

Top 10 dashboard

The IBM Top 10 dashboard represents top ten IBM volumes and hosts in all storage systems monitored by the vROps Manager.

The IBM XIV, IBM Spectrum Accelerate and IBM FlashSystem A9000 Top 10 dashboards include the following information:

- Top 10 volumes by IOPS (last hour)
- Top 10 volumes by IOPS (last 24 hours)
- Top 10 volumes by throughput (last hour)
- Top 10 volumes by throughput (last 24 hours)
- Top 10 hosts by IOPS (last hour)
- Top 10 hosts by IOPS (last 24 hours)
- Top 10 hosts by throughput (last hour)
- Top 10 hosts by throughput (last 24 hours)

The Top 10 dashboards for the storage systems that run Spectrum Virtualize include the following information:

- Top 10 VDisks by read operations per second (last hour)
- Top 10 VDisks by write operations per second (last hour)
- Top 10 VDisks by read blocks per second (last hour)
- Top 10 VDisks by write blocks per second (last hour)
- Top 10 VDisks by worst read response in us since last statistics collection
- Top 10 VDisks by worst write response in us since last statistics collection
- Top 10 VDisks by average transfer response time in us (last hour)

You can double-click on any resource element (disk, host or VDisk) to display its details.

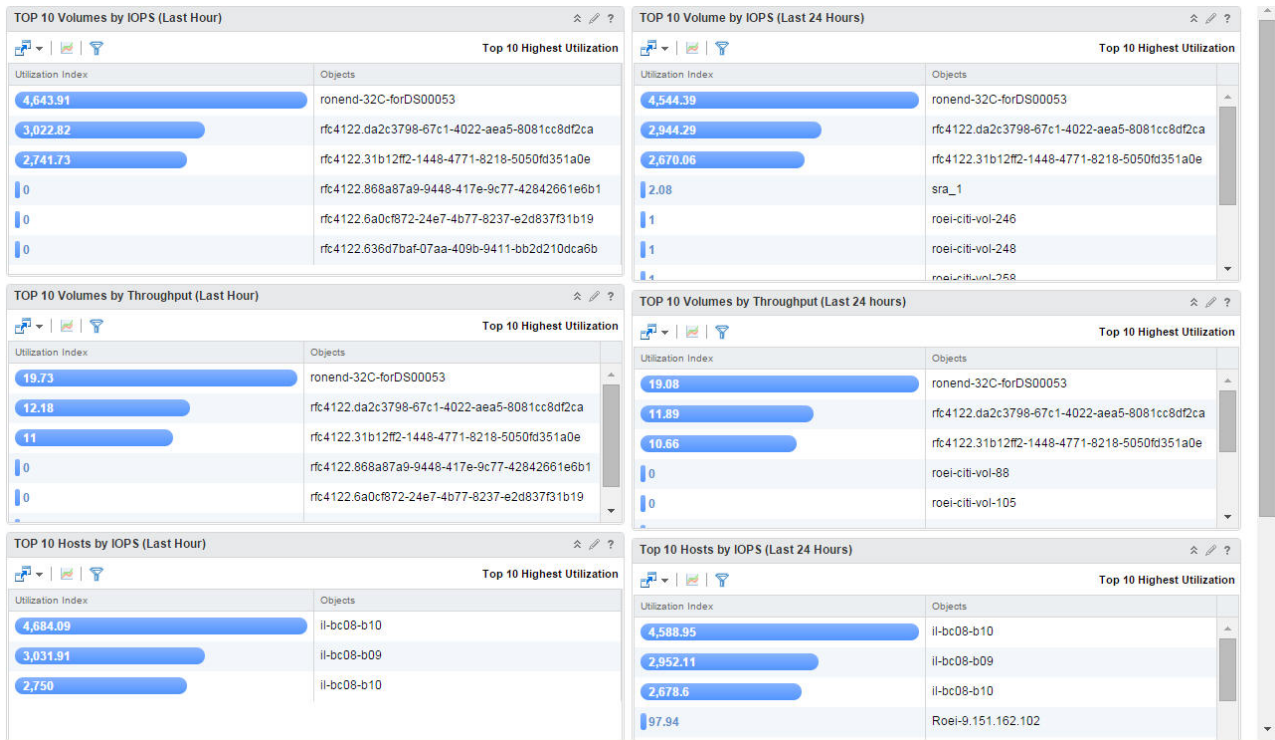


Figure 126. IBM XIV Top 10 dashboards

Displaying the overview of IBM storage objects

You can display the overview of IBM storage objects, including its health, risk and efficiency status, and other attributes in a centralized manner, using the vROps Environment display.

Procedure

To display the overview of IBM storage objects:

1. Click **Environment** > **All Objects** > **IBM Storage Adapter**. The list of IBM storage objects is displayed in the left-hand pane of the vROps Manager.
2. In the IBM storage object list, select an object that you intend to monitor. The overview of the selected object status is displayed in the right-hand pane of the vROps Manager.

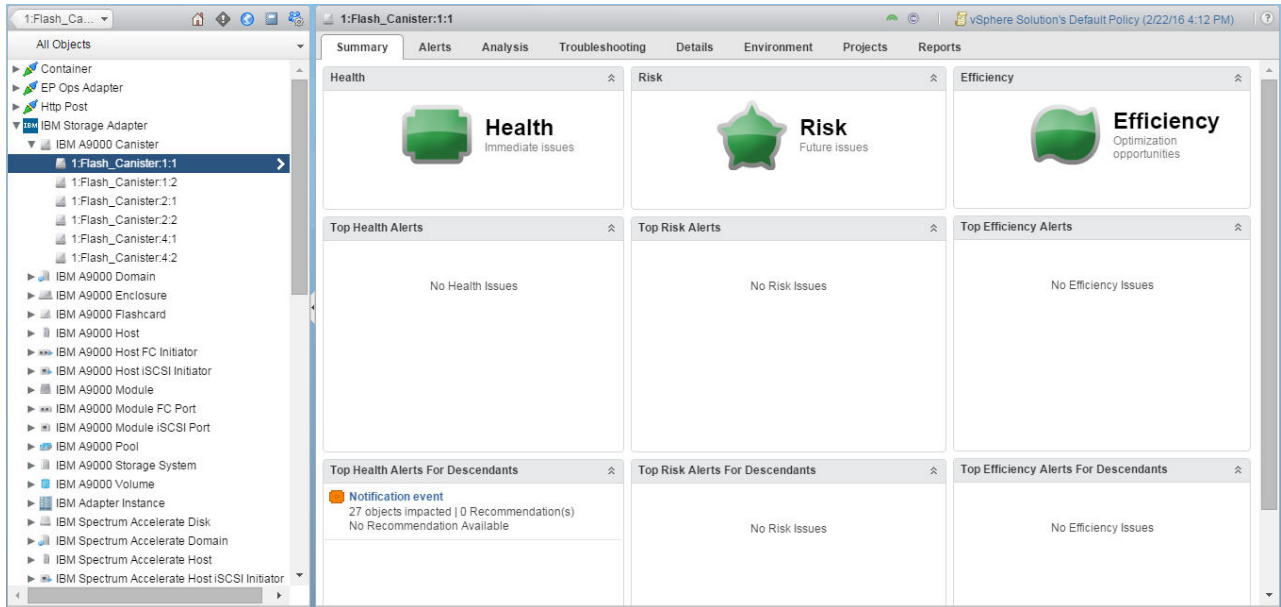


Figure 127. Overview of an IBM FlashSystem A9000 flash canister

Defining thresholds and alerts for storage objects

The vROps Manager maintains thresholds of normal behavior for each storage element.

About this task

The vROps Manager alerts the user, when a metric violates a threshold. The thresholds can be set to fit the needs of the user. Afterward, you can define an alert to be triggered, when the threshold is crossed, changing the color of the storage object in the Overview and Performance dashboards. In addition, you can use the storage object metrics as Key Performance Indicator (KPI).

Procedure

To set thresholds and define their alerts:

1. Go to **Content > Symptom Definitions > Metric/Property Symptom Definitions**.
2. In the right-hand pane of the **Metric/Property Symptom Definitions** dialog box, click on the '+' sign to add a new symptom definition. The **Add Symptom Definition** dialog box is displayed.
3. In the **Base Object Type** drop-down list, select **IBM Storage Adapter**, and then choose a storage element and a metrics, which thresholds you intend to define.
4. Double-click the metrics related to the selected storage resource, or drag it to the working area of the right-hand pane. The selected object and metrics attributes are displayed.

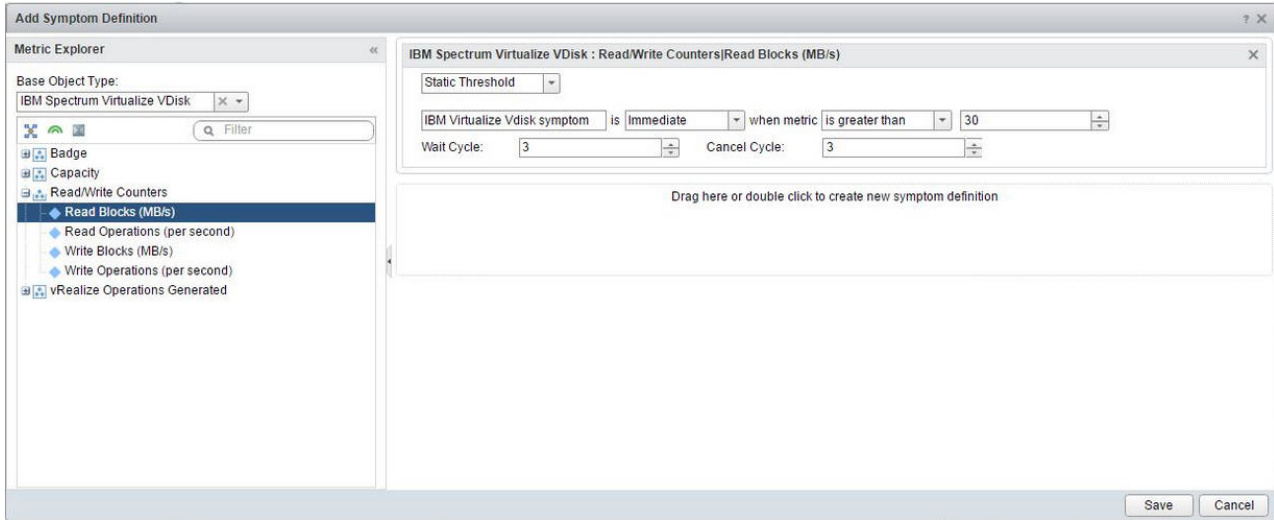


Figure 128. Add Symptom Definition dialog box

5. Enter the symptom definition name, criticality level, threshold triggering condition, wait cycle and cancel cycle values. Then click **Save** to add the symptom definition.
6. Go to **Content > Alert Definitions**, click on the '+' sign to add a new alert definition. The **Alert Definition Workspace** dialog box is displayed.
7. In the left-hand pane of the **Alert Definition Workspace** dialog box, enter alert name and description.

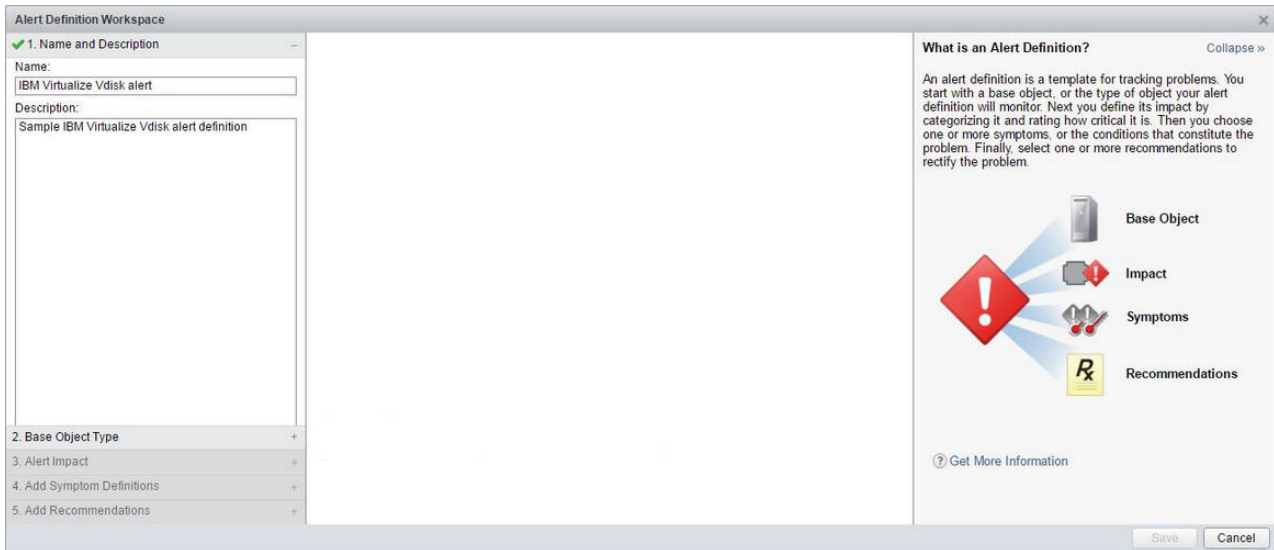


Figure 129. Adding name and description for alert definition

8. Click **Base Object Type** and in the drop-down list, select **IBM Storage Adapter**, and then choose a storage element and a metrics, to which you intend to create an alert.




Figure 130. Selecting base object type for alert definition

9. Click **Alert Impact** and enter impact, criticality, alert type/subtype, wait cycle and cancel cycle values for the alert.




Figure 131. Selecting impact for alert definition

10. Click **Add Symptom Definition**, enter *Self* for the **Defined On** parameter, *Metric/Property* for the **Symptom Definition Type** parameter, locate the previously defined symptom in the **Symptom** list and add it to the alert definition by dragging it to the Symptoms area on the right-hand pane.

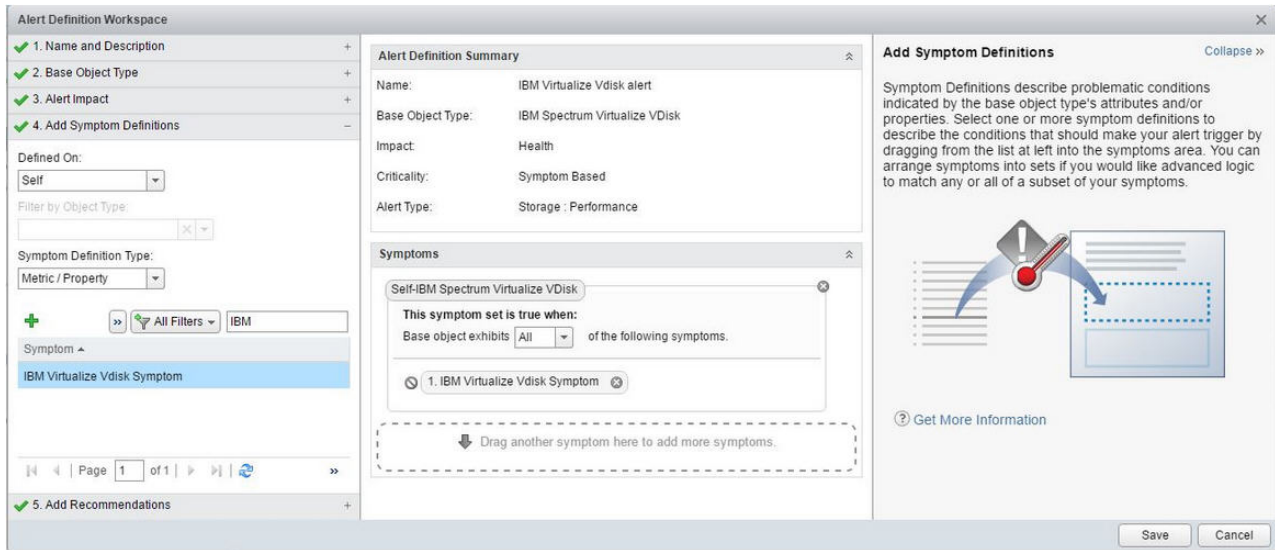






Figure 132. Adding symptom for alert definition

11. Click **Save** to save the changes. After the duration of wait cycle that you defined has elapsed, the storage object color on a dashboard changes, if the metrics threshold was crossed. Storage objects display different colors, according to the criticality of symptoms:
 - Info – 
 - Warning – 
 - Immediate – 
 - Critical – 
12. Navigate to **Administration > Policies**.
13. In the **Policies** pane, click the **Policy Library** tab, select the **Default Policy** entry and click the **Edit** button. The **Edit Monitoring Policy** dialog box is displayed.

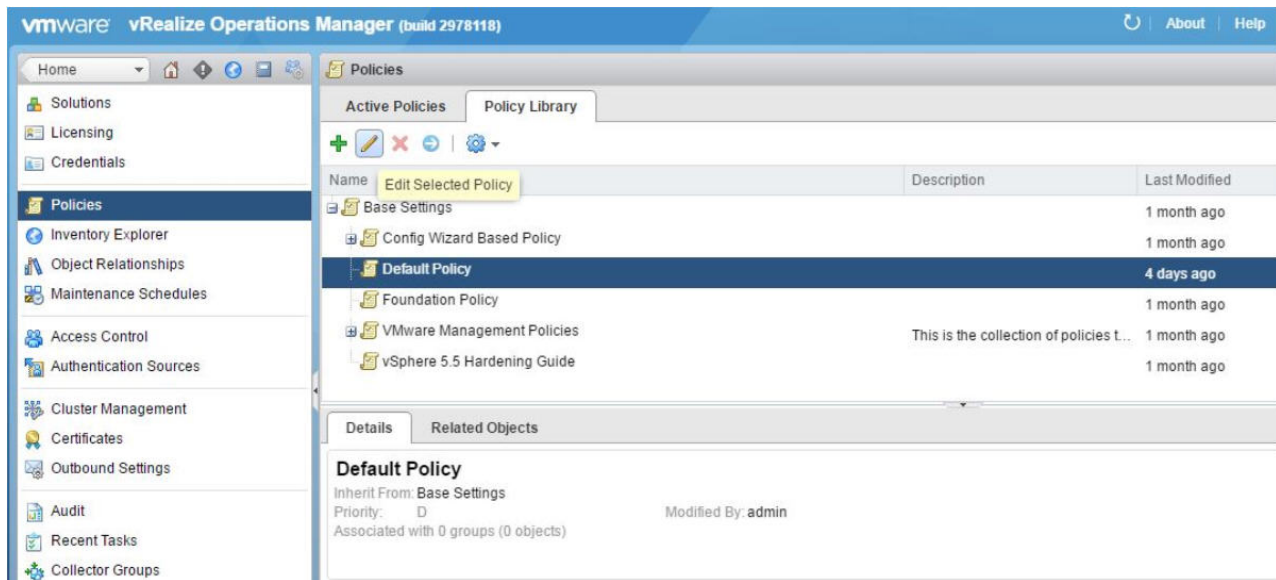


Figure 133. Editing default monitoring policy

14. In the **Edit Monitoring Policy** dialog box, leave the configuration steps 1 to 4 unchanged.
15. In the left-hand pane of the **Edit Monitoring Policy** dialog box, select **Collect Metrics and Properties**. The list of attributes is displayed in the right-hand pane of the dialog box.
16. In the list of attributes, filter by the object type to locate the storage object used for alert definitions, then select the metrics that you want to use as a Key Performance Indicator, set its KPI to **Yes**, and click **Save** to finish.

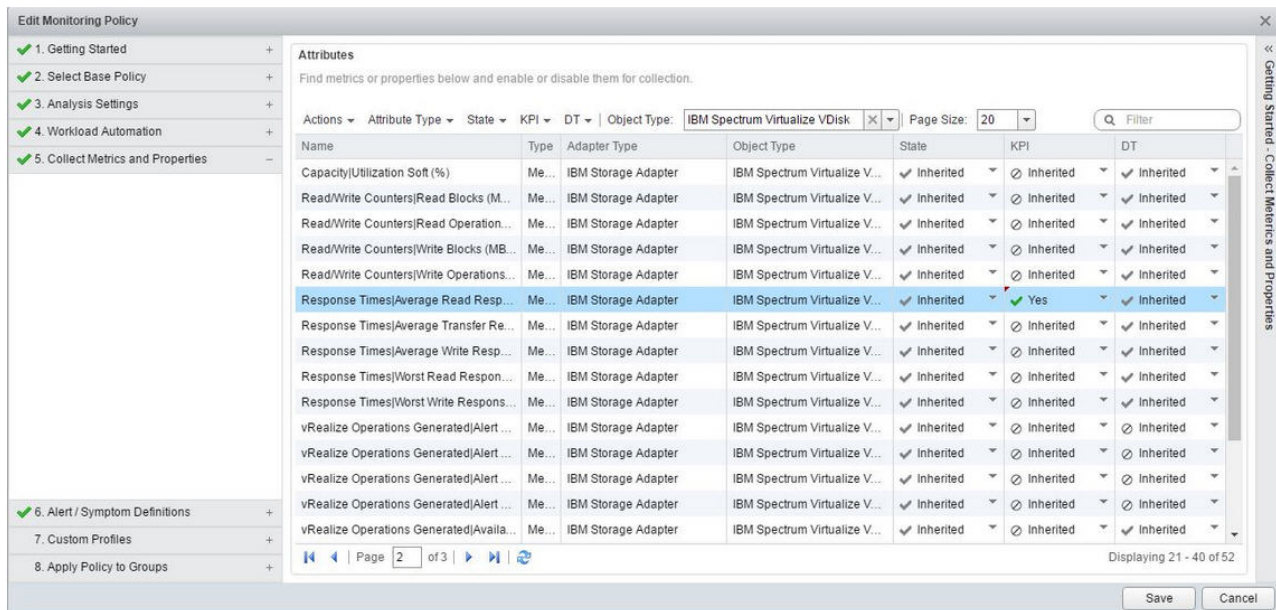


Figure 134. Enabling KPI for a storage object metrics

Chapter 7. IBM Storage Enabler for Containers

IBM Storage Enabler for Containers allows IBM storage systems to be used as persistent volumes for stateful application running in Kubernetes clusters.

IBM Storage Enabler for Containers is based on an open-source IBM project, Ubiquity, integrating it with IBM Spectrum Connect. Through the IBM Storage Enabler for Containers, Kubernetes persistent volumes (PVs) can be provisioned from IBM storage. This is performed by specifying the Spectrum Connect storage service for Kubernetes storage class object. Thus, the containers can be used with stateful microservices, such as database applications (MongoDB, PostgreSQL etc).

IBM Storage Enabler for Containers uses Kubernetes dynamic provisioning for creating and deleting volumes on IBM storage systems. For details about volume provisioning in the Kubernetes environment, refer to Persistent volumes on Kubernetes (kubernetes.io/docs/concepts/storage/volumes). In addition, IBM Storage Enabler for Containers utilizes the full set of Kubernetes FlexVolume APIs for volume operations on a host. The operations include initiation, attachment/detachment, mounting/unmounting etc.

Taking full advantage of the flexible service-based storage provisioning model in Spectrum Connect, the IBM Storage Enabler expands storage profile (service) policy. This allows defining specific capabilities per storage service and creating a storage volume according to these requirements. This policy-driven approach helps the Kubernetes administrators easily define Kubernetes storage classes, such as gold, silver or bronze, by using the Spectrum Connect services. The storage administrators define Spectrum Connect storage services and select their capabilities, sizes and storage types. In their turn, the Kubernetes administrators do not need to be fully aware of the existing storage capabilities; all that they have to do is to select the relevant services to be used for the stateful containers.

Note: For the user convenience, this guide might refer to IBM Storage Enabler for Containers as Enabler for Containers.

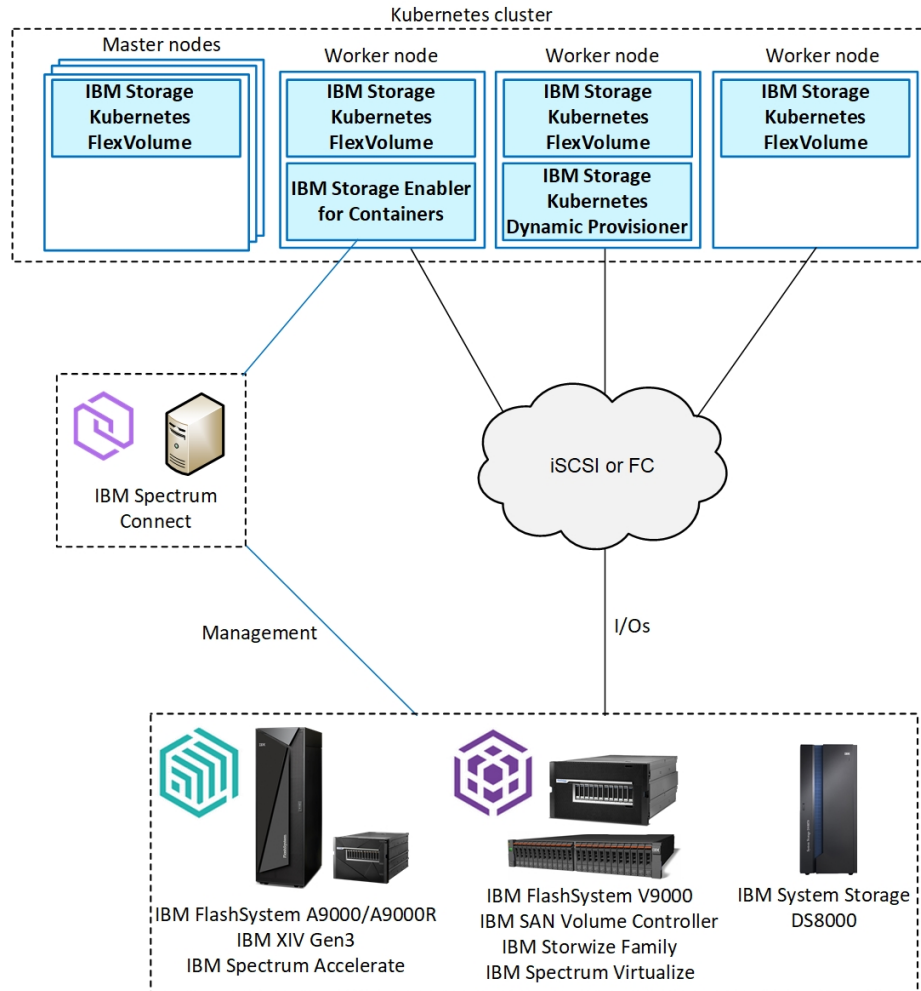


Figure 135. Integration of IBM storage systems in Kubernetes environment

The IBM Storage Enabler for Containers ensures that the data persists (stays intact) even after the container is stopped or removed. The IBM Storage Enabler communicates with the IBM storage systems through IBM Spectrum Connect. Spectrum Connect creates a storage service (for example, gold, silver or bronze) and makes it available for Kubernetes Dynamic Provisioner and FlexVolume, automating IBM storage provisioning for Kubernetes persistent volumes.

- The Dynamic Provisioner allows storage volumes to be created on-demand, using Kubernetes storage classes based on Spectrum Connect storage services. This provides abstraction for the underlying storage platform, eliminating the need for cluster administrators to pre-provision storage.
- The FlexVolume is deployed as a DaemonSet on all nodes of the cluster, enabling the users to attach and mount storage volumes into a pod within a Kubernetes node. The DaemonSet installs the FlexVolume CLI on every node in the cluster in the Kubernetes plug-in directory.

Note: The instances of IBM Storage Enabler for Containers (*ubiquity*), its database (*ubiquity-db*) and IBM Storage Kubernetes Dynamic Provisioner (*ubiquity-k8s-provisioner*) are protected, using standard Kubernetes methods for high-availability. They are deployed as Kubernetes Deployment objects with `replica=1`, so if a node fails, Kubernetes automatically reschedules them to run on another node. IBM Storage Kubernetes FlexVolume (*ubiquity-k8s-flex*) is deployed as a Kubernetes DaemonSet on all the worker and master nodes.

Installing IBM Storage Enabler for Containers

Download and install the IBM Storage Enabler for Containers in Kubernetes cluster as described in the following sections.

- “Compatibility and requirements for IBM Storage Enabler for Containers”
- “Managing SSL certificates with IBM Storage Enabler for Containers” on page 182
- “Downloading IBM Storage Enabler for Containers software” on page 183
- “Performing installation of IBM Storage Enabler for Containers” on page 183

For information about uninstallation, see “Uninstalling the IBM Storage Enabler for Containers software” on page 189.

Compatibility and requirements for IBM Storage Enabler for Containers

For the complete and up-to-date information about the compatibility and requirements for using IBM Storage Enabler for Containers with Kubernetes, refer to the IBM Spectrum Connect latest release notes. The release notes detail supported operating system and Kubernetes versions, as well as microcode versions of the supported storage systems. You can find the latest Spectrum Connect release notes on IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SS6JWS).

About this task

Follow these steps to prepare your environment for installing the IBM Storage Enabler for Containers in the Kubernetes cluster that requires persistent volumes for stateful containers.

Procedure

1. Contact your storage administrator and make sure IBM Storage Enabler for Containers interface has been added to active Spectrum Connect instance and at least one storage service has been delegated to it. See “Managing integration with IBM Storage Enabler for Containers” on page 189 and “Delegating storage services to the IBM Storage Enabler for Containers interface” on page 191 for details.
2. Verify that there is a proper communication link between Spectrum Connect and Kubernetes cluster.
3. Perform these steps for each worker node in Kubernetes cluster:
 - a. Install the following Linux packages to ensure Fibre Channel and iSCSI connectivity. Skip this step, if the packages are already installed.
 - RHEL:

- sg3_utils.
- iscsi-initiator-utils (if iSCSI connection is required).

```
sudo yum -y install sg3_utils
sudo yum -y install iscsi-initiator-utils
```

- Ubuntu:

- scsitools.
- open-iscsi (if iSCSI connection is required).

```
sudo apt-get install scsitools
sudo apt-get install open-iscsi
```

- Configure Linux multipath devices on the host. Create and set the relevant storage system parameters in the `/etc/multipath.conf` file. You can also use the default `multipath.conf` file located in the `/usr/share/doc/device-mapper-multipath-*` directory.

Verify that the `systemctl status multipathd` output indicates that the multipath status is active and error-free.

- RHEL:

```
yum install device-mapper-multipath
sudo modprobe dm-multipath
systemctl start multipathd
systemctl status multipathd
multipath -ll
```

- Ubuntu:

```
apt-get install multipath-tools
sudo modprobe dm-multipath
systemctl start multipathd
systemctl status multipathd
multipath -ll
```

Important: When configuring Linux multipath devices, verify that the `find_multipaths` parameter in the `multipath.conf` file is disabled.

- RHEL: Remove the `find_multipaths yes` string from the `multipath.conf` file.
 - Ubuntu: Add the `find_multipaths no` string to the `multipath.conf` file.
-

- Configure storage system connectivity.

- Define the hostname of each Kubernetes node on the relevant storage systems with the valid WWPN or IQN of the node. The hostname on the storage system must be the same as the hostname defined in the Kubernetes cluster. Use the `$> kubectl get nodes` command to display hostname, as illustrated below. In this example, the `k8s-worker-node1` and the `k8s-worker-node2` hostnames must be defined on a storage system.

Note: In most cases, the local hostname of the node is the same as the Kubernetes node hostname as displayed in the `kubectl get nodes` command output. However, if the names are different, make sure to use the Kubernetes node name, as it appears in the command output.

```

root@k8s-user-v18-master:~# kubectl get nodes
NAME                STATUS    ROLES    AGE     VERSION
k8s-master          Ready    master   34d    v1.8.4
k8s-worker-node1    Ready    <none>   34d    v1.8.4
k8s-worker-node2    Ready    <none>   34d    v1.8.4

```

- For iSCSI, perform these three steps .
 - Make sure that the login used to log in to the iSCSI targets is permanent and remains available after a reboot of the worker node. To do this, verify that the **node.startup** in the `/etc/iscsi/iscsid.conf` file is set to *automatic*. If not, set it as required and then restart the `iscsid` service (**`$> /etc/init.d/iscsid restart`**).
 - Discover and log into the iSCSI targets of the relevant storage systems.

```

$> iscsiadm -m discoverydb -t st -p ${storage system iSCSI port IP}:3260
--discover
$> iscsiadm -m node -p ${storage system iSCSI port IP/hostname} --login

```

- Verify that the login was successful and display all targets that you logged in. The *portal* value must be the iSCSI target IP address.

```

$> iscsiadm -m session --rescan
Rescanning session [sid: 1, target: {storage system IQN},
portal: {storage system iSCSI port IP},{port number}]

```

- d. Make sure that the node *kubelet* service has the attach/detach capability enabled, **enable-controller-attach-detach=true** (enabled by default). To verify the current status, run the following command and check that the Setting node annotation to enable volume controller attach/detach message is displayed:

```

$> journalctl -u kubelet | grep 'Setting node annotation to .
* volume controller attach/detach' | tail -1
Jan 03 17:55:05 k8s-19-master-shay kubelet[3627]: I0103 17:55:05.437720 3627
kubelet_node_status.go:273] Setting node annotation to enable volume controller
attach/detach

```

If the volume controller attach/detach functionality is disabled, enable it, as detailed in Kubernetes documentation.

4. Perform these steps for every master node in Kubernetes cluster:
 - a. Enable the attach/detach capability for the *kubelet* service (**controller-attach-detach-enabled=true**).
 - b. Configure the controller-manager pod to let it access the Kubernetes plug-in directory (`/usr/libexec/kubernetes/kubelet-plugins/volume/exec`), where the FlexVolume driver is located. This step is required, when the controller-manager on the master nodes is deployed as a static pod under Kubernetes versions 1.6 and 1.7.

Skip this step if you use Kubernetes version 1.8 or if you run the controller-manager as a regular Linux service, which already has access to the required folder.

 - Stop the controller-manager pod by moving the `kube-controller-manager.yml` file to temporary directory: **`mv /etc/kubernetes/manifests/kube-controller-manager.yml /tmp`**.
 - Edit the `kube-controller-manager.yml` file: **`vi /tmp/kube-controller-manager.yml`**.
 - Add the following lines under the **volumes** tag.

```
- hostPath:
  path: /usr/libexec/kubernetes/kubelet-plugins/volume/exec
  type: DirectoryOrCreate
  name: flexvolume-dir
```

- Add the following lines under the **volumeMounts** tag:

```
- mountPath: /usr/libexec/kubernetes/kubelet-plugins/volume/exec
  name: flexvolume-dir
```

- Restart the controller-manager pod by moving the `kube-controller-manager.yml` file to its original location:
mv /tmp/kube-controller-manager.yml /etc/kubernetes/manifests/.
 - Verify that the controller-manager pod is in the Running state: **kubectl get pod -n kube-system | grep controller-manager.**
5. If dedicated SSL certificates are required, see the relevant section of the “Managing SSL certificates with IBM Storage Enabler for Containers” procedure. When no validation is required and you can use the self-signed certificates, generated by default by the IBM Storage Enabler for Containers server, skip this procedure.

Managing SSL certificates with IBM Storage Enabler for Containers

IBM Storage Enabler for Containers uses SSL certificates for maintaining a secure communication link between the IBM Storage Enabler for Containers server, its database, the Dynamic Provisioner, the FlexVolume, and the Spectrum Connect server.

Before you begin

Download and extract the IBM Storage Enabler for Containers installer to gain access to the installation script (`ubiquity_installer.sh`). See steps 1 to 3 of the “Performing installation of IBM Storage Enabler for Containers” on page 183 section.

About this task

IBM Storage Enabler for Containers supports two SSL modes, when communicating with its components:

- *require*, when no validation is required. The IBM Storage Enabler for Containers server generates self-signed certificates on the fly. In this mode, you can skip the procedure detailed below and continue with the installation of the IBM Storage Enabler for Containers without any special SSL configuration.
- *verify-full*, expecting the user to provide relevant certificates. When enabled, this SSL mode requires additional configuration steps as listed below.

Procedure

1. When operating in the *verify-full* mode, you will need to generate the following three pairs of the public-private keys for:
 - Spectrum Connect server. You can upload these certificates to the server, as explained in “Managing server certificates” on page 37.
 - IBM Storage Enabler for Containers (*ubiquity*) service object.
 - IBM Storage Enabler for Containers database (*ubiquity-db*) service object.

2. Verify that:

- The SSL certificates that you have generated are valid and signed by root CA.
- The SSL certificates have valid common and alternative names. The alternative names list must contain valid DNS names and/or IP addresses of the SCBE server, *ubiquity* service object, and *ubiquity-db* service object. Run this command to obtain the required network parameters for the *ubiquity* and *ubiquity-db* services:

```
$> ./ubiquity_installer.sh -s create-services
```

The script generates two Kubernetes services that provide the required DNS/IP address combinations.

- The private certificate and certificate key files have the following names:
 - *ubiquity.crt* and *ubiquity.key* for the *ubiquity* service object.
 - *ubiquity-db.crt* and *ubiquity-db.key* for the *ubiquity-db* service object.
 - The trusted CA files contain the root CA certificate and have the following names:
 - *scbe-trusted-ca.crt* for the Spectrum Connect server.
 - *ubiquity-trusted-ca.crt* for the *ubiquity* service object.
 - *ubiquity-db-trusted-ca.crt* for the *ubiquity-db* service object.
 - Copy all generated *.crt and *.key files to a dedicated directory.
3. Run the `$> ubiquity_installer.sh -s create-secrets-for-certificates -t <certificate directory>` command to create the following ConfigMap and secrets:
- ConfigMap *ubiquity-public-certificates* for all the trusted CA files.
 - The *ubiquity-private-certificate* secret for the private certificates used by the *ubiquity* service object.
 - The *ubiquity-db-private-certificate* secret for the private certificates used by the *ubiquity-db* service object.
4. Proceed with installation of the IBM Storage Enabler for Containers, as detailed in “Performing installation of IBM Storage Enabler for Containers.”

Downloading IBM Storage Enabler for Containers software

IBM Storage Enabler for Containers is available as a free software solution for IBM storage system customers.

About this task

You can download the latest version of the Installer for IBM Storage Enabler for Containers at any time from the IBM Fix Central. Fix Central provides fixes and updates for your systems software, hardware, and operating system.

Performing installation of IBM Storage Enabler for Containers

You can install the IBM Storage Enabler for Containers software on a compatible version of Kubernetes. For more information, refer to the release notes of the IBM Spectrum Connect.

Before you begin

Verify that you have completed the preliminary configuration steps, as detailed in “Compatibility and requirements for IBM Storage Enabler for Containers” on page 179.

Important:

- IBM Storage Dynamic Provisioner for Kubernetes uses the Kubernetes configuration file to access the Kubernetes API server and monitor the Persistent Volume Claims (PVCs). Usually, Kubernetes configuration file is located either in the `~/kube/config` or `/etc/kubernetes` directory. Make sure that this configuration file has access to all the namespaces intended persistent volume provisioning. If a PVC comes from namespace that cannot be accessed, it will not be served.
 - During installation of the IBM Storage Enabler for Containers, the IBM Storage Kubernetes FlexVolume driver is automatically installed on all master and worker nodes in a Kubernetes cluster, using the `ubiquity-k8s-flex` DaemonSet.
 - A single IBM Storage Enabler for Containers instance can be installed per one Kubernetes cluster.
-

Procedure

Follow these steps to install IBM Storage Enabler for Containers:

1. Download the installer. See “Downloading IBM Storage Enabler for Containers software” on page 183).
2. Copy the installer to a local folder on a host that can access the Kubernetes cluster, using the `kubect1` command. Usually, a master node has access to the cluster.
3. Extract the installer file ('*' represents the build number):

```
# tar -xzvf installer-for-ibm-storage-enabler-for-containers-x.y.z-*.tar.gz
```

4. Update the `ubiquity_installer.conf` configuration file, according to your environment requirements. Replace the *VALUE* placeholders in the files with your values.

Important: Any change to be made after running the installation script must be performed manually in the corresponding `yml` files themselves.

Table 30. Configuration parameters in `ubiquity_installer.conf`

Parameter	Description
UBIQUITY_IMAGE	Docker image of the IBM Storage Enabler for Containers to be deployed as Kubernetes deployment/ <code>ubiquity</code> . Set by default to <code>ibmcom/ibm-storage-enabler-for-containers:x.y.z</code> . The image is available in the Docker Hub.
UBIQUITY_DB_IMAGE	Docker image of the IBM Storage Enabler for Containers database to be deployed as Kubernetes deployment/ <code>ubiquity-db</code> . Set by default to <code>ibmcom/ibm-storage-enabler-for-containers-db:x.y.z</code> . The image is available in the Docker Hub.

Table 30. Configuration parameters in `ubiquity_installer.conf` (continued)

Parameter	Description
UBIQUITY_K8S_PROVISIONER_IMAGE	Docker image of the IBM Storage Dynamic Provisioner to be deployed as Kubernetes deployment//ubiquity-k8s-provisioner. Set by default to <code>ibmcom/ibm-storage-dynamic-provisioner-for-kubernetes:x.y.z</code> . The image is available in the Docker Hub.
UBIQUITY_K8S_FLEX_IMAGE	Docker image of the IBM Storage FlexVolume to be deployed as Kubernetes daemonset/ubiquity-k8s-flex. Set by default to <code>ibmcom/ibm-storage-flex-volume-for-kubernetes:x.y.z</code> . The image is available in the Docker Hub.
SCBE_MANAGEMENT_IP_VALUE	IP address or FQDN of the Spectrum Connect server.
SCBE_MANAGEMENT_PORT_VALUE	Communication port of the Spectrum Connect server. Default value is <code>8440</code> .
SCBE_DEFAULT_SERVICE_VALUE	Default Spectrum Connect storage service to be used, if not specified by the storage class.
UBIQUITY_INSTANCE_NAME_VALUE	A prefix for any new volume created on the storage system. For example, <code>u_<instance_name>_<PVC_ID></code> .
DEFAULT_FSTYPE_VALUE	File system type of a new volume, if not specified by the user in the storage class. Allowed values: <code>ext4</code> or <code>xf</code> s. Default value is <code>ext4</code> .
DEFAULT_VOLUME_SIZE_VALUE	Default volume size (in GB), if not specified by the user when creating a new volume. Default value is <code>1</code> .
SKIP_RESCAN_ISCSI_VALUE	Rescanning mode. Allowed values: <code>true</code> or <code>false</code> . Set it to <code>true</code> if the nodes have FC connectivity. Setting to <code>false</code> in iSCSI environment, triggers a rescan. Default value is <code>false</code> .
LOG_LEVEL_VALUE	Log level. Allowed values: <code>debug</code> , <code>info</code> , <code>error</code> . Default value is <code>info</code> .
SSL_MODE_VALUE	SSL verification mode. Allowed values: <code>require</code> (No validation is required, the IBM Storage Enabler for Containers server generates self-signed certificates on the fly.) or <code>verify-full</code> (Certificates are provided by the user). The <code>verify-full</code> mode requires additional configuration steps, as detailed in the "Managing SSL certificates with IBM Storage Enabler for Containers" on page 182 section.
SCBE_USERNAME_VALUE	Username defined for the IBM Storage Enabler for Containers interface in Spectrum Connect. Note: The IBM Storage Enabler for Containers interface appears as Enabler for Containers in Spectrum Connect GUI.
SCBE_PASSWORD_VALUE	Password defined for the IBM Storage Enabler for Containers interface in Spectrum Connect. Note: The IBM Storage Enabler for Containers interface appears as Enabler for Containers in Spectrum Connect GUI.

Table 30. Configuration parameters in `ubiquity_installer.conf` (continued)

Parameter	Description
UBIQUITY_DB_PV_NAME_VALUE	Name of the persistent volume to be used for the <i>ubiquity-db</i> database. For the Spectrum Virtualize and Spectrum Accelerate storage systems, use the default value (<i>ibm-ubiquity-db</i>). For the DS8000 storage system, use a shorter value, such as (<i>ibmdb</i>). This is necessary because the DS8000 volume name length cannot exceed 16 characters. As a result, the UBIQUITY_INSTANCE_NAME and UBIQUITY_DB_PV_NAME combination (<code>u_<instance_name>_<PV_name></code>) cannot exceed 16 characters.
UBIQUITY_DB_USERNAME_VALUE	Username and password for the deployment of <i>ubiquity-db</i> database. Do not use the <i>postgres</i> username, because it already exists.
UBIQUITY_DB_PASSWORD_VALUE	Username and password for the deployment of <i>ubiquity-db</i> database.
STORAGE_CLASS_NAME_VALUE	Storage class name. Note: The storage class parameters are used for creating an initial storage class for the <i>ubiquity-db</i> PVC. You can use this storage class for other applications as well. It is recommended to set the storage class name to be the same as the Spectrum Connect storage service name.
STORAGE_CLASS_PROFILE_VALUE	Storage class profile, directing to the Spectrum Connect storage service name.
STORAGE_CLASS_FSTYPE_VALUE	File system type for the storage class profile. Allowed values: <i>ext4</i> or <i>xf</i> s. Default value is <i>ext4</i> .

- Apply the `ubiquity_installer.conf` settings to the relevant `yml` files of the installer. Run the

```
$> ./ubiquity_installer.sh -s update-ymls -c ubiquity_installer.conf
```

command to do this. The code example below illustrates a successful configuration of the `yml` files.

```
[root@k8s-18-master-ibm installer-for-ibm-storage-enabler-for-containers-x.y.z-*]#
./ubiquity_installer.sh -s update-ymls -c /var/tmp/ubiquity_installer.conf
Executing STEP [update-ymls]...
Updating yml files with placeholders from /var/tmp/ubiquity_installer.conf file. Are you sure (y/n): y
Update placeholder [UBIQUITY_IMAGE                ] in files : ./ymls/ubiquity-deployment.yml
Update placeholder [UBIQUITY_DB_IMAGE             ] in files : ./ymls/ubiquity-db-deployment.yml
Update placeholder [UBIQUITY_K8S_PROVISIONER_IMAGE] in files : ./ymls/ubiquity-k8s-provisioner-deployment.yml
...
...
...
Update placeholder [STORAGE_CLASS_NAME_VALUE     ] in files : ./ymls/storage-class.yml
./ymls/ubiquity-db-pvc.yml ./ymls/sanity_ymls/sanity-pvc.yml
Update placeholder [STORAGE_CLASS_PROFILE_VALUE  ] in files : ./ymls/storage-class.yml
Update placeholder [STORAGE_CLASS_FSTYPE_VALUE  ] in files : ./ymls/storage-class.yml
Finished updating yml files according to /var/tmp/ubiquity_installer.conf
```

- Start the first installation stage, which is performed without installing the database (*ubiquity-db*). Run the `$> ./ubiquity_installer.sh -s install -k <path_to_Kubernetes_configuration_file>` command to do this. Make sure that the name of the target configuration file is `config`. The first stage is complete, illustrated by the following message
"IBM Storage Enabler for Containers" installation finished, but the deployment is not ready yet".

- For Kubernetes version 1.6 or 1.7, you must reload the FlexVolume. To do this, manually restart the *kubelet* on all the Kubernetes worker and master nodes. Skip this step for Kubernetes version 1.8, in which the driver is discovered automatically.
- Start the second stage to install the IBM Storage Enabler for Containers database. Run the `$> ./ubiquity_installer.sh -s create-ubiquity-db` command. The installation is complete, illustrated by the following message "IBM Storage Enabler for Containers" installation finished successfully in the Kubernetes cluster".

What to do next

- Verify the post-installation status of the IBM Storage Enabler for Containers service. Run the following command: `$> ./ubiquity_cli.sh -a status`. Check that the database status is Bound, and all other elements are available and running, as illustrated in the following example.

```

NAME                                TYPE      DATA      AGE
secrets/ubiquity-db-credentials      Opaque    3           2m
secrets/scbe-credentials              Opaque    2           2m
NAME                                DATA     AGE
cm/k8s-config                         1         2m
cm/ubiquity-configmap                 10        2m
NAME                                CAPACITY  ACCESS MODES  RECLAIM POLICY  STATUS  CLAIM                                STORAGECLASS  REASON  AGE
pv/ibm-ubiquity-db                    20Gi      RWO           Delete          Bound   ubiquity/ibm-ubiquity-db            gold        1m
NAME                                STATUS    VOLUME        CAPACITY  ACCESS MODES  STORAGECLASS  AGE
pvc/ibm-ubiquity-db                   Bound    ibm-ubiquity-db  20Gi      RWO           gold          2m
NAME                                TYPE     CLUSTER-IP  EXTERNAL-IP  PORT(S)      AGE
svc/ubiquity                          ClusterIP IP_address  | <none>      9999/TCP     2m
svc/ubiquity-db                       ClusterIP IP_address  | <none>      5432/TCP     2m
NAME                                DESIRED   CURRENT   READY   UP-TO-DATE   AVAILABLE   NODE SELECTOR  AGE
ds/ubiquity-k8s-flex                  3         3         3       3             3           <none>         1m
NAME                                DESIRED   CURRENT   UP-TO-DATE   AVAILABLE   AGE
deploy/ubiquity                       1         1         1             1           2m
deploy/ubiquity-db                     1         1         1             1           1m
deploy/ubiquity-k8s-provisioner        1         1         1             1           2m
kubectl get --namespace ubiquity pod | egrep "^ubiquity|^NAME"
-----
NAME                                READY     STATUS    RESTARTS  AGE
ubiquity-db-6b4cf9fd84-8g4k8        1/1      Running   0          1m
ubiquity-f7b99f75c-5p5db            1/1      Running   0          2m
ubiquity-k8s-flex-724qq              1/1      Running   0          1m
ubiquity-k8s-flex-g7lls              1/1      Running   0          1m
ubiquity-k8s-flex-gd7cs              1/1      Running   0          1m
ubiquity-k8s-provisioner-5b8888fd5-j6w4g 1/1      Running   0          2m

```

Figure 136. IBM Storage Enabler for Containers post-installation status

Note: You can use the `$> ./ubiquity_cli.sh -a status_wide` command to display the full system status. You can use it to verify that there is *ubiquity-k8s-flex* pod on each worker and master node.

- Perform the sanity test to spin up and down the Kubernetes PVC and pod. Run this command: `$> ./ubiquity_cli.sh -a sanity`. The following message must be displayed: Sanity finished successfully (pvcl and podl were successfully created and deleted).
- In addition to the default storage class, define more Kubernetes storage classes, if needed. This default storage class, named as a value of the **STORAGE_CLASS_NAME_VALUE** parameter in the `ubiquity_installer.conf` file, was created during the installation. Template for setting storage classes is included in the `./yaml/templates/storage-class-template.yml` file. See "Configuring storage classes, PVCs and pods" on page 193 for details.
- Use the IBM Storage Enabler for Containers for creating persistent volume claims (PVCs) on IBM storage systems. Template for PVC configuration is

included in the `./ymls/templates/pvc-template.yml` file. See “Configuring storage classes, PVCs and pods” on page 193 for details.

Upgrading IBM Storage Enabler for Containers

If you are already using earlier releases of IBM Storage Enabler for Containers, you can upgrade to the newer version without having to uninstall the previous one or using the installation package (*.tar) file.

Procedure

Perform the following procedure to upgrade IBM Storage Enabler for Containers:

1. Update the current deployment and DaemonSet images from version `<current version number>` to `<new version number>` by running the following commands:
 - `kubectl set image daemonset/ubiquity-k8s-flex ubiquity-k8s-flex=ibmcom/ibm-storage-flex-volume-for-kubernetes:<new version number> -n ubiquity`
 - `kubectl set image deployment/ubiquity-k8s-provisioner ubiquity-k8s-provisioner=ibmcom/ibm-storage-dynamic-provisioner-for-kubernetes:<new version number> -n ubiquity`
 - `kubectl set image deployment/ubiquity ubiquity=ibmcom/ibm-storage-enabler-for-containers:<new version number> -n ubiquity`
 - `kubectl set image deployment/ubiquity-db ubiquity-db=ibmcom/ibm-storage-enabler-for-containers-db:<new version number> -n ubiquity`
2. Verify that the image update was successful by checking the status of the IBM Storage Enabler for Containers pods. Navigate to the directory where the IBM Storage Enabler for Containers was installed and run the following command `ubiquity_cli.sh -a status_wide` and check that all relevant pods are in the *Running* state.

Note: You can also use the `#> kubectl get -n ubiquity pod,daemonset,deployment -o wide` command to verify the update.

The image update is finished.

3. After the update, change the image version from `<current version number>` to `<new version number>` in all relevant yml files. Otherwise, you will not be able to start the images using `ubiquity_cli.sh -a start`. Use the following procedure:
 - a. Navigate to your installer directory and go the `/yaml` folder in it.
 - b. Locate the following yml files:
 - `ubiquity-deployment.yml`
 - `ubiquity-k8s-provisioner-deployment.yml`
 - `ubiquity-k8s-flex-daemonset.yml`
 - `ubiquity-db-deployment.yml`
 - c. Edit the files and manually change the image version from `<current version number>` to `<new version number>`. For example, image:
`ibmcom/ibm-storage-enabler-for-containers-db:<current version number>`
to image: `ibmcom/ibm-storage-enabler-for-containers-db:<new version number>`.

Uninstalling the IBM Storage Enabler for Containers software

If you want to completely remove the IBM Storage Enabler for Containers software, use the following procedure.

Before you begin

Verify that there are no persistent volumes (PVs) that have been created, using IBM Storage Enabler for Containers.

Important: The uninstallation process removes the IBM Storage Enabler for Containers components, its metadata, user credentials, *ubiquity* namespace and other elements.

Procedure

To uninstall the IBM Storage Enabler for Containers software:

1. Log on to the host which contains the Enabler for Containers and navigate to the installation directory.
2. Run this script to completely uninstall IBM Storage Enabler for Containers: **\$> ./ubiquity_uninstall.sh**. The installation is complete, illustrated by the following message
"IBM Storage Enabler for Containers" uninstall finished..
3. As this script does not uninstall the FlexVolume from all Kubernetes nodes, perform these additional steps for the complete FlexVolume removal:
 - a. If you use Kubernetes version 1.6 or 1.7 and have the `/usr/libexec/kubernetes/kubelet-plugins/volume/exec` directory on the controller-manager static pod, remove the Kubernetes FlexVolume from the controller-manager pod and restart the *kubelet* service.
 - b. Delete the `/usr/libexec/kubernetes/kubelet-plugins/volume/exec/ibm^ubiquity-k8s-flex` directory from all Kubernetes nodes.
 - c. Restart the *kubelet* service on all nodes.

Managing integration with IBM Storage Enabler for Containers

The IBM Storage Enabler for Containers is used for provisioning of storage volumes from an external IBM storage system to Kubernetes containers.

About this task

The following procedure details how to add the IBM Storage Enabler for Containers interface to Spectrum Connect. A single IBM Storage Enabler for Containers interface can be operated per one Kubernetes cluster.

Procedure

1. On the **Interfaces** pane, click **Add Interface**, and then select **Enabler for Containers**. The **Add New Enabler for Containers Interface** dialog box is displayed.




Figure 137. Add New Enabler for Containers Interface dialog box

2. Enter credentials for the new IBM Storage Enabler for Containers user, and click **Apply**.

Note: When entering a user name and password for an Enabler for Containers interface on the Spectrum Connect with LDAP authentication, make sure that these credentials are the same as defined for LDAP. In addition, you can choose between using a single user or user group, if LDAP is enabled for this Spectrum Connect instance.

The IBM Storage Enabler for Containers interface is now added to Spectrum Connect.




Figure 138. IBM Storage Enabler for Containers interface on the Interfaces pane

What to do next

You can continue integrating the IBM Storage Enabler for Containers interface, as explained in the following sections:

- “Delegating storage services to the IBM Storage Enabler for Containers interface” on page 191.
- “Canceling service delegation to IBM Storage Enabler for Containers” on page 192.

Delegating storage services to the IBM Storage Enabler for Containers interface

Before you can use the IBM Storage Enabler for Containers to provision storage volumes from an external IBM storage system to Kubernetes containers, you must delegate the storage services that will be used by container plug-ins.

About this task

The services and their storage resources that you delegate on Spectrum Connect can be used in creating storage volumes in Kubernetes. Spectrum Connect storage services are translated into Kubernetes storage classes allowing for dynamic (on-demand) provisioning of storage for containers.

Service delegation is a prerequisite for deploying IBM Storage Enabler for Containers, IBM Storage Kubernetes Dynamic Provisioner and IBM Storage Kubernetes FlexVolume. For more information about deployment requirements, see “Compatibility and requirements for IBM Storage Enabler for Containers” on page 179).

Procedure

To delegate storage services to IBM Storage Enabler for Containers:

1. On the **Interfaces** pane, click the Enabler for Containers interface to select it.
2. On the **Spaces/Storage Services** pane, select the storage space from which you want to choose storage services. The available services that reside on the selected storage space are immediately displayed.
3. Right-click on a service that you want to delegate to the Enabler for Containers interface, and then select **Delegate to <interface_name>**, or click the **Attach/Delegate** button on the service. The service color changes to indicate the successful delegation.

You can continue the process by right-clicking available services under the current space.

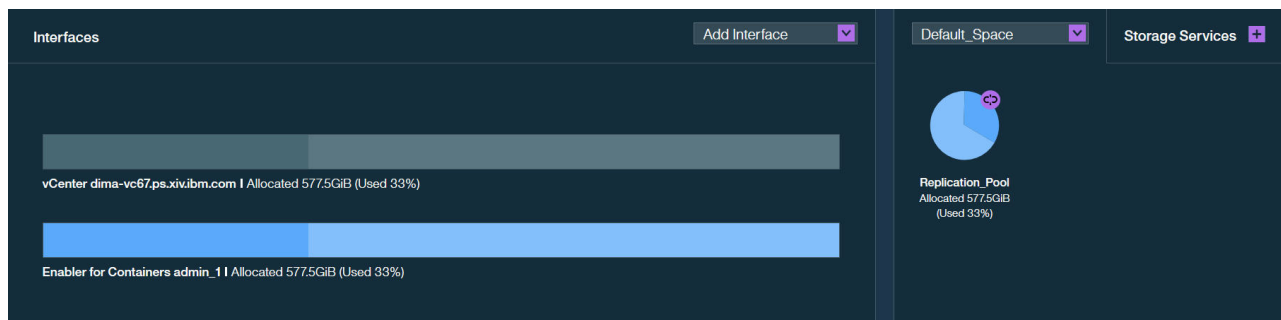


Figure 139. Enabler for Containers interface with a delegated service

The Enabler for Containers interface provides indication for the allocated and used storage space.

- Allocated – total amount of storage space available on all pools connected to the delegated services.
- Used – amount of storage space used by containers and snapshots on all pools connected to the delegated services.

What to do next

After service delegation, you can proceed with installation of the IBM Storage Enabler for Containers for further use of the allocated storage resources as persistent volumes for containers. See “Installing IBM Storage Enabler for Containers” on page 179.

If this is the first service defined before installation of the IBM Storage Enabler for Containers, a default storage class is created automatically during the installation. To link a Spectrum Connect storage service to a Kubernetes storage class, set the value of the `STORAGE_CLASS_PROFILE_VALUE` parameter in the `ubiqity.config` file to be the same as the service name.

If you already installed the IBM Storage Enabler for Containers, add more services and delegate them to the Enabler for Containers interface. Then create Kubernetes storage classes and link them to the services. These storage classes can be used for creating new PVCs based on the Spectrum Connect services.

Canceling service delegation to IBM Storage Enabler for Containers

When required, you can cancel a storage service delegation to a IBM Storage Enabler for Containers interface.

Before you begin

Before canceling service delegation to IBM Storage Enabler for Containers, delete all Kubernetes storage classes linked to the services, which delegations are to be canceled.

About this task

Storage services, which delegation has been canceled, and their resources (pools) cannot be used as external storage for containers.

Procedure

1. On the **Interfaces** pane, click the IBM Storage Enabler for Containers interface. The services that are currently delegated to the interface are highlighted on the **Spaces/Storage Services** pane.
2. Right-click on a service which delegation you want cancel, and then select **Cancel delegation to <Interface_name>** , or click the **Detach/Cancel Delegation** button on the service. The service color changes to indicate the successful detachment.

You can continue the process by right-clicking delegated services under the current space.

Using the IBM Storage Enabler for Containers

This chapter covers the following topics:

- “Configuring storage classes, PVCs and pods” on page 193.
- “Sample configuration for running a stateful container” on page 195.
- “Recovering a crashed Kubernetes node” on page 198.
- “Updating the Enabler for Containers configuration files” on page 200.

Configuring storage classes, PVCs and pods

This section details how to configure Kubernetes storage classes, persistent volume claims and pods.

Defining additional Kubernetes storage classes

Define additional Kubernetes storage classes, if needed. Template for setting storage classes is illustrated below. The template is provided as the `./ymls/templates/storage-class-template.yml` file. As the only storage class, created during installation, is used for the database, you might need additional storage classes for volume provisioning on IBM storage. A separate storage class must be assigned for each storage service delegated to the IBM Storage Enabler for Containers interface in Spectrum Connect.

```
kind: StorageClass
apiVersion: storage.k8s.io/v1beta1
metadata:
  name: "<Storage class name>"
  labels:
    product: ibm-storage-enabler-for-containers
# annotations:
# storageclass.beta.kubernetes.io/is-default-class: "boolean"
provisioner: "ubiquity/flex"
parameters:
  profile: "<SCBE service name>"
  fstype: "<Filesystem type>"
  backend: "scbe"
```

Table 31. Configuration parameters in `storage-class-template.yml`

Parameter	Description
name	Storage class name. It is recommended to use the storage class profile value as the storage class name.
profile	Spectrum Connect storage service name.
fstype	File system type of a new volume Allowed values: <i>ext4</i> or <i>xfs</i> .
is-default-class	Configures the storage class to be the default <i>true</i> or not <i>false</i> . By default, it is set to <i>false</i> .
product	Permanently set to <i>ibm-storage-enabler-for-containers</i>
provisioner	Permanently set to <i>ubiquity/flex</i> .
backend	Permanently set to <i>scbe</i> .

Creating persistent volume claims (PVCs)

Use the IBM Storage Enabler for Containers for creating persistent volume claims (PVCs) on IBM storage systems. Template for PVC configuration is illustrated below. The template is provided as the `./ymls/templates/pvc-template.yml` file.

When a PVC is created, the IBM Storage Dynamic Provisioner generates a persistent volume (PV), according to Spectrum Connect storage service, defined for the PVC storage class, then it binds the PV to the PVC. By default, the PV name will be PVC-ID. The volume name on the storage will be `u_[ubiquity-instance]_[PVC-ID]`. Keep in mind that the `[ubiquity-instance]` value is set in the IBM Storage Enabler for Containers configuration file.

```

kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: "<PVC name>"
  labels:
    product: ibm-storage-enabler-for-containers
  # pv-name: "<PV name>"
spec:
  storageClassName: <Storage Class Name>
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: <Number>Gi

```

Table 32. Configuration parameters in *pvc-template.yml*

Parameter	Description
name	Persistent volume claim name
storageClassName	Storage class name used for the PVC provisioning
pv-name	Persistent volume name. This name is used for creating a PV with a specific name, which is different from the default PV. The default PV name is its PVC ID. However, this dedicated PV name must be unique. No other PV with the same name is allowed within the Kubernetes cluster.
accessModes	Permanently set to <i>ReadWriteOnce</i> . Other access modes, such as <i>ReadWriteMany</i> , are not supported.
storage	Volume size in Gb. Other volume size units are not supported.

Creating a pod to use the PVC for storage

The PVCs can be used by Kubernetes pods for running stateful applications. Below is the example for of the template for using PVC in the pod *yml* file. When a pod is created, The IBM Storage FlexVolume performs the following actions automatically:

- Volume attachment to the host, on which Kubernetes scheduled the pod to run.
- Rescanning and discovering the multipath device of the new volume.
- Creating an XFS or EXT4 filesystem on the device (if filesystem does not exist on the volume).
- Mounting the new multipath device on `/ubiquity/[WWN of the volume]`.
- Creating a symbolic link from `/var/lib/kubelet/pods/[pod ID]/volumes/ibm~ubiquity-k8s-flex/[PVC ID]` to `/ubiquity/[WWN of the volume]`.

As a result, the pod goes up with the mounted PV on the container pod in the `mountPath` defined in the *yml* file.

```

kind: Pod
apiVersion: v1
metadata:
  name: <Pod name>
spec:
  containers:
  - name: <Container name>
    image: <Image name>
    volumeMounts:
    - name: <yaml volume name>
      mountPath: <Mount point>
  volumes:
  - name: <yaml volume name>
    persistentVolumeClaim:
      claimName: <PVC name>

```

Table 33. Configuration parameters in *sanity-pod.yml*

Parameter	Description
name	Pod name
containers.name	Container name
containers.image	Container image
volumeMounts.name	Internal volume name
volumeMounts.mountPath	Mounting point for the PVC in the container
volumes.name	Internal volume name
volumes.persistentVolumeClaim	Name of the persistent volume claim

When a pod is deleted, the following actions are performed automatically:

- Removing a symbolic link from `/var/lib/kubelet/pods/[pod ID]/volumes/ibm~ubiquity-k8s-flex/[PVC ID]` to `/ubiquity/[WWN of the volume]`.
- Unmounting the new multipath device on `/ubiquity/[WWN of the volume]`.
- Removing the multipath device of the volume.
- Detaching (unmapping) the volume from the host.
- Rescanning in the cleanup mode to remove the physical and multipath device files of the detached volume.

Sample configuration for running a stateful container

You can use IBM Storage Enabler for Containers for running stateful containers with a storage volume provisioned from an external IBM storage system.

About this task

This example illustrates a basic configuration required for running a stateful container with volume provisioned on a Spectrum Connect storage service.

- Creating a storage class `gold` that is linked to Spectrum Connect storage service `gold` with XFS file system.
- Creating a PersistentVolumeClaim (PVC) `pvc1` that uses the storage class `gold`.
- Creating a pod `pod1` with container `container1` that uses PVC `pvc1`.
- Starting I/Os into `/data/myDATA` in `pod1\container1`.
- Deleting the `pod1` and then creating a new `pod1` with the same PVC. Verifying that the file `/data/myDATA` still exists.

- Deleting all storage elements (pod, PVC, persistent volume and storage class).

Procedure

1. Open a command-line terminal.
2. Create a storage class, as shown below. The storage class gold is linked to a Spectrum Connect storage service on a pool from IBM FlashSystem A9000R with QoS capability and XFS file system. As a result, any volume with this storage class will be provisioned on the gold service and initialized with XFS file system.

```
$> cat storage_class_gold.yml
kind: StorageClass
apiVersion: storage.k8s.io/v1beta1
metadata:
  name: "gold"                # Storage Class name
  annotations:
    storageclass.beta.kubernetes.io/is-default-class: "true"
provisioner: "ubiquity/flex"
parameters:
  profile: "gold"
  fstype: "xfs"
  backend: "scbe"
```

```
$> kubectl create -f storage_class_gold.yml
storageclass "gold" created
```

3. Display the newly created storage class to verify its successful creation.

```
$> kubectl get storageclass gold
NAME          TYPE
gold (default)  ubiquity/flex
```

4. Create a PVC pvc1 with the size of 1 Gb that uses the storage class gold.

```
$> cat pvc1.yml
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: "pvc1"
spec:
  storageClassName: gold
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
```

```
$> kubectl create -f pvc1.yml
persistentvolumeclaim "pvc1" created
```

The IBM Storage Enabler for Containers creates a persistent volume (PV) and binds it to the PVC. The PV name will be PVC-ID. The volume name on the storage will be u_[ubiquity-instance]_[PVC-ID]. Keep in mind that the [ubiquity-instance] value is set in the IBM Storage Enabler for Containers configuration file.

5. Display the existing PVC and persistent volume.

```
$> kubectl get pvc
NAME      STATUS    VOLUME                                     CAPACITY   ACCESSMODES   AGE
pvc1     Bound    pvc-254e4b5e-805d-11e7-a42b-005056a46c49  1Gi        RWO           1m
```

```
$> kubectl get pv
NAME                                     CAPACITY   ACCESSMODES   RECLAIMPOLICY   STATUS   CLAIM           REASON   AGE
pvc-254e4b5e-805d-11e7-a42b-005056a46c49  1Gi        RWO           Delete          Bound    default/pvc1
```

6. Display the additional persistent volume information, such as its WWN, location on a storage system, etc.

```
$> kubectl get -o json pv pvc-254e4b5e-805d-11e7-a42b-005056a46c49 | grep -A15 flexVolume
  "flexVolume": {
    "driver": "ibm/ubiquity",
    "options": {
      "LogicalCapacity": "1000000000",
      "Name": "u_PROD_pvc-254e4b5e-805d-11e7-a42b-005056a46c49",
      "PhysicalCapacity": "1023410176",
      "PoolName": "gold-pool",
      "Profile": "gold",
      "StorageName": "A9000 system1",
      "StorageType": "2810XIV",
      "UsedCapacity": "0",
      "Wwn": "36001738CFC9035EB0CCCCC5",
      "fstype": "xfs",
      "volumeName": "pvc-254e4b5e-805d-11e7-a42b-005056a46c49"
    }
  },
```

7. Create a pod pod1 with a persistent volume vol1.

```
$> cat pod1.yml
kind: Pod
apiVersion: v1
metadata:
  name: pod1
spec:
  containers:
  - name: container1
    image: alpine:latest
    command: [ "/bin/sh", "-c", "--" ]
    args: [ "while true; do sleep 30; done;" ]
    volumeMounts:
    - name: vol1
      mountPath: "/data"
  restartPolicy: "Never"
  volumes:
  - name: vol1
    persistentVolumeClaim:
      claimName: pvcl

$> kubectl create -f pod1.yml
pod "pod1" created
```

As a result, the IBM Storage Kubernetes FlexVolume performs the following:

- Attaches the volume to the host.

Note: Volume attachment is triggered by the controller-manager which runs on the master node.

- Rescans and discover the multipath device of the new volume.
- Creates XFS or EXT4 filesystem on the device (if filesystem does not exist on the volume).
- Mounts the new multipath device on /ubiquity/[WWN of the volume].
- Creates a symbolic link from /var/lib/kubelet/pods/[pod ID]/volumes/ibm~ubiquity-k8s-flex/[PVC ID] to /ubiquity/[WWN of the volume].

8. Display the newly created pod1 and write data to its persistent volume. Make sure that the pod status is Running.

```
$> kubectl get pod pod1
NAME      READY   STATUS    RESTARTS   AGE
pod1     1/1     Running   0           16m
```

```

$> kubectl exec pod1 -c container1 -- bash -c "df -h /data"
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/mpathi 951M  33M  919M   4% /data

$> kubectl exec pod1 -c container1 -- bash -c "mount | grep /data"
/dev/mapper/mpathi on /data type xfs (rw,relatime,seclabel,attr2,inode64,noquota)

$> kubectl exec pod1 touch /data/FILE
$> kubectl exec pod1 ls /data/FILE
File

$> kubectl describe pod pod1 | grep "^Node:"
Node:  k8s-node1/hostname

```

9. Log in to the worker node that has the running pod and display the newly attached volume on the node.

```

> multipath -ll
mpathi (36001738cfc9035eb0cccc5) dm-12 IBM      ,2810XIV
size=954M features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
  |- 3:0:0:1 sdb 8:16 active ready running
  `- 4:0:0:1 sdc 8:32 active ready running

$> df | egrep "ubiquity|^Filesystem"
Filesystem      1K-blocks    Used Available Use% Mounted on
/dev/mapper/mpathi 973148    32928    940220    4% /ubiquity/6001738CFC9035EB0CCCC5

$> mount |grep ubiquity
/dev/mapper/mpathi on /ubiquity/6001738CFC9035EB0CCCC5 type xfs (rw,relatime,seclabel,attr2,inode64,noquota)

$> ls -l /var/lib/kubelet/pods/*/volumes/ibm~ubiquity-k8s-flex/*
lrwxrwxrwx. 1 root root 42 Aug 13 22:41 pvc-254e4b5e-805d-11e7-a42b-005056a46c49 -> /ubiquity/6001738CFC9035EB0CCCC5

```

10. Delete the pod.

```

$> kubectl delete pod pod1
pod "pod1" deleted

```

As a result, the IBM Storage Kubernetes FlexVolume performs the following:

- Removes symbolic link from `/var/lib/kubelet/pods/[pod ID]/volumes/ibm~ubiquity-k8s-flex/[PVC ID]` to `/ubiquity/[WWN of the volume]`.
- Unmounts the new multipath device on `/ubiquity/[WWN of the volume]`.
- Removes the multipath device of the volume.
- Detaches (unmap) the volume from the host.
- Rescans in cleanup mode to remove the physical device files of the detached volume.

11. Remove the PVC and its PV (volume on the storage system).

```

$> kubectl delete -f pvc1.yml
persistentvolumeclaim "pvc1" deleted

```

12. Remove the storage class. This command removes the Kubernetes storage class only, the Spectrum Connect storage service remains intact.

```

$> kubectl delete -f storage_class_gold.yml
storageclass "gold" deleted

```

Recovering a crashed Kubernetes node

This section details a manual operation required to revive Kubernetes pods that reside on a crashed node.

Identifying a crashed node

When a worker node shuts down or crashes, all stateful pods that reside on it become unavailable, and the node status appears as *NotReady*.

```
# kubectl get nodes
```

NAME	STATUS	AGE	VERSION
kuber-node1	Ready	2h	v1.7.5
kuber-node2	NotReady	2h	v1.7.5
kuber-serv1	Ready	2h	v1.7.5

When this node status persists for more than five minutes (default setting, see note below for instructions on how to change this value), the following occurs:

- Status of a pod scheduled on the pod becomes *Unknown*.
- The new pod is scheduled on another node in the cluster with status *ContainerCreating*, denoting that the pod is scheduled on a crashed node.

As a result, the pod scheduled on a crashed node appears twice on two nodes with two statuses, as illustrated below.

```
# kubectl get pods -o wide
```

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE
sanity-deployment-2414-538d2	1/1	Unknown	0	15m	IP_address	kuber-node2
sanity-deployment-2414-n8cfv	0/1	ContainerCreating	0	34s	<none>	kuber-node1

Note: The time period between the node failure and creation of a new pod on another node is user-configurable. Use the following procedure to change the `pod-eviction-timeout` value:

1. Move the `kube-controller-manager.yml` file to `/tmp` folder (`mv /etc/kubernetes/manifests/kube-controller-manager.yml /tmp`).
 2. Edit the controller-manager file (`vim /tmp/kube-controller-manager.yml`).
 3. Add the `--pod-eviction-timeout=60s` line to the `kube-controller-manager` command.
 4. Move the `kube-controller-manager.yml` file to its original location (`mv /tmp/kube-controller-manager.yml /etc/kubernetes/manifests/kube-controller-manager.yml`).
-

Recovering a crashed node

To allow Kubernetes to recover the stateful pods from a crashed node and schedule them on a functional node in the cluster:

- Remove the crashed node from the cluster to free up all its pods (`kubectl delete node <node_name>`),
or
- Force delete the stateful pods, which are in *Unknown* state (`kubectl delete pods <pod_name> --grace-period=0 --force -n <namespace>`).

After the mandatory five-minute timeout, as set by Kubernetes itself, the pod runs on a scheduled node. The pod status changes from *ContainerCreating* to *Running*. See example below for the `sanity-deployment-2414-n8cfv` pod.

If the crashed node recovers by itself or the user reboots the node, no additional actions are required to release its pods. The pods recover automatically after the node restores itself and joins the cluster. When a crashed node is recovered, the following occurs:

1. The pod with the *Unknown* status is deleted.
2. The volume(s) is detached from the crashed node.
3. The volume(s) is attached to node, on which the new pod is scheduled.
4. After the mandatory five-minute timeout, as set by Kubernetes itself, the pod runs on a scheduled node. The pod status changes from *ContainerCreating* to *Running*. See example below for the `sanity-deployment-2414-n8cfv` pod.

```
# kubectl get pods -o wide
```

```
NAME                                READY STATUS  RESTARTS  AGE  IP            NODE
sanity-deployment-2414-n8cfv        1/1   Running   0         8m   IP_address   kuber-node1
```

Updating the Enabler for Containers configuration files

If required, you can adjust configuration parameters of the IBM Storage Enabler for Containers after its installation, using the `yml` files on the host:

- `./ubiquity-configmap.yml`
- `./scbe-credentials-secret.yml`
- `./ubiquity-db-credentials-secret.yml`
- `./ubiquity-db-credentials-secret.yml`
- `./yamls/storage-class.yml`
- `./yamls/ubiquity-db-pvc.yml`
- `./yamls/sanity_yamls/sanity-pvc.yml`

Table 34. Configuration parameters in `ubiquity-configmap.yml`

Parameter	Description
SCBE-MANAGEMENT-IP	IP address or FQDN of the Spectrum Connect server. To update this parameter after installation: 1. Update the parameter in the <code>ubiquity-configmap</code> configMap (<code>\$> kubectl edit -n ubiquity configmap ubiquity-configmap</code>). 2. Delete the <code>ubiquity</code> pod of the <code>ubiquity</code> deployment. Then, it is automatically recreated by the deployment with the new parameters.
SCBE-MANAGEMENT-PORT	Communication port of the Spectrum Connect server. To update this parameter after installation, use the same procedure as the one for SCBE-MANAGEMENT-IP .
SCBE-DEFAULT-SERVICE	Default Spectrum Connect storage service to be used, if not specified by the storage class. To update this parameter after installation, use the same procedure as the one for SCBE-MANAGEMENT-IP .
UBIQUITY-INSTANCE-NAME	A prefix for any new volume created on the storage system. This parameter cannot be changed.

Table 34. Configuration parameters in *ubiquity-configmap.yml* (continued)

Parameter	Description
UBIQUITY-IP-ADDRESS	<p>IP address of the <i>ubiquity</i> service object. The <code>ubiquity_installer.sh</code> automatically updates this parameter during the initial installation.</p> <p>If the <i>ubiquity</i> service is recreated, it receives a new IP address. As a result, this parameter must be updated to allow the FlexVolume pods to access it.</p> <p>To update the IP address after installation:</p> <ol style="list-style-type: none"> 1. Update the IP address in the <code>ubiquity_configmap</code> configMap (<code>\$> kubectl edit -n ubiquity configmap ubiquity-configmap</code>). 2. Delete all FlexVolume pods (<code>\$> kubectl get pod -n ubiquity grep ubiquity-k8s-flex awk '{print \$1}' xargs kubectl delete pod -n ubiquity</code>).
DEFAULT-FSTYPE	<p>File system type of a new volume, if not specified by the user in the storage class.</p> <p>Allowed values: <i>ext4</i> or <i>xfs</i>.</p> <p>To update this parameter after installation, use the same procedure as the one for SCBE-MANAGEMENT-IP.</p>
DEFAULT-VOLUME-SIZE	<p>Default volume size (in GB), if not specified by the user when creating a new volume.</p> <p>To update this parameter after installation, use the same procedure as the one for SCBE-MANAGEMENT-IP.</p>
SKIP-RESCAN-ISCSI	<p>Rescanning mode.</p> <p>Allowed values: <i>true</i> or <i>false</i>. Set to <i>true</i> if the nodes have FC connectivity.</p> <p>To change the rescanning mode after installation:</p> <ol style="list-style-type: none"> 1. Update the modes in the <code>ubiquity_configmap</code> configMap (<code>\$> kubectl edit -n ubiquity configmap ubiquity-configmap</code>). 2. Delete all FlexVolume pods (<code>\$> kubectl get pod -n ubiquity grep ubiquity-k8s-flex awk '{print \$1}' xargs kubectl delete pod -n ubiquity</code>).
LOG-LEVEL	<p>Log level.</p> <p>Allowed values: <i>debug</i>, <i>info</i>, <i>error</i>.</p> <p>To receive more detailed events about the IBM Storage Enabler for Containers operation, you can adjust the log level after installation. Use this information your troubleshooting and debug processes.</p> <p>To change the log level after installation:</p> <ol style="list-style-type: none"> 1. Change the value in the <code>ubiquity_configmap</code> configMap (<code>\$> kubectl edit -n ubiquity configmap ubiquity-configmap</code>). 2. Stop the IBM Storage Enabler for Containers (<code>\$> ubiquity_cli.sh -a stop</code>). 3. Restart the IBM Storage Enabler for Containers (<code>\$> ubiquity_cli.sh -a start</code>).

Table 34. Configuration parameters in *ubiquity-configmap.yml* (continued)

Parameter	Description
SSL-MODE	<p>SSL verification mode.</p> <p>Allowed values: <i>require</i> (No validation is required, the IBM Storage Enabler for Containers server generates self-signed certificates on the fly.) or <i>verify-full</i> (Certificates are provided by the user.)</p> <p>To change the SSL mode to <i>verify-full</i> after installation:</p> <ol style="list-style-type: none"> 1. Stop the IBM Storage Enabler for Containers (\$> ubiquity_cli.sh -a stop). 2. Edit the <i>ubiquity-configmap.yml</i> file, setting the SSL_MODE to <i>verify-full</i>. 3. Edit the <i>ubiquity_installer.conf</i> file, setting the SSL_MODE to <i>verify-full</i>. 4. Run \$> ubiquity_installer.sh -a update-ymls -c ubiquity_installer.conf. 5. Create SSL certificates, as detailed in the “Managing SSL certificates with IBM Storage Enabler for Containers” on page 182 section. 6. Restart the IBM Storage Enabler for Containers (\$> ubiquity_cli.sh -a start).

Table 35. Configuration parameters in *scbe-credentials-secret.yml*

Parameter	Description
SCBE-USERNAME	<p>Username defined for the IBM Storage Enabler for Containers interface in Spectrum Connect.</p> <p>To change the Enabler for Containers interface credentials after installation:</p> <ol style="list-style-type: none"> 1. Change the credentials for the Enabler for Containers interface, using the Spectrum Connect GUI. 2. Enter the new Base-64-encoded username and password (\$> kubectl edit -n ubiquity scbe-credentials). 3. Delete the <i>ubiquity-k8s-flex</i> pods, <i>ubiquity-k8s-provisioner</i> pod and <i>ubiquity</i> pod. Do not delete the <i>ubiquity-db</i> pod After the deletion, Kubernetes restarts the pods with new credentials.
SCBE-PASSWORD	<p>Password defined for the IBM Storage Enabler for Containers interface in Spectrum Connect. Instructions on how to change the interface credentials after installation are detailed in the SCBE-USERNAME description above.</p>

Table 36. Configuration parameters in *ubiquity-db-credentials-secret.yml*

Parameter	Description
UBIQUITY-DB-USERNAME	<p>Username and password for the deployment of <i>ubiquity-db</i> database.</p> <p>Do not use the <i>postgres</i> username, because it already exists.</p>

Table 36. Configuration parameters in *ubiquity-db-credentials-secret.yml* (continued)

Parameter	Description
UBIQUITY-DB-PASSWORD	Username and password for the deployment of <i>ubiquity-db</i> database.

Table 37. Configuration parameters in *storage-class.yml*, *ubiquity-db-pvc.yml*, *sanity-pvc.yml*

Parameter	Description
STORAGE-CLASS-NAME	Storage class name.
STORAGE-CLASS-PROFILE	Storage class profile, directing to the Spectrum Connect storage service.
STORAGE-CLASS-FSTYPE	File system type for the storage class profile. Allowed values: <i>ext4</i> or <i>xfs</i> .
Note: The storage class parameters are used for creating an initial storage class for the <i>ubiquity-db</i> PVC.	

Troubleshooting the IBM Storage Enabler for Containers

This section can help you detect and solve problems that you might encounter when using the IBM Storage Enabler for Containers.

Checking logs

You can use the IBM Storage Enabler for Containers logs for problem identification. To collect and display logs, related to the different components of IBM Storage Enabler for Containers, use the following Kubernetes commands:

- Log collection – **./ubiquity_cli.sh -a collect_logs**. The logs are kept in a folder, named as *./ubiquity_collect_logs_MM-DD-YYYY-h:m:s*. The folder is placed in the directory, from which the log collection command was run.
- IBM Storage Enabler for Containers – **\$> kubectl logs -n ubiquity deploy/ubiquity**.
- IBM Storage Enabler for Containers database – **\$> kubectl logs -n ubiquity deploy/ubiquity-db**.
- IBM Storage Kubernetes Dynamic Provisioner – **\$> kubectl logs -n ubiquity deploy/ubiquity-k8s-provisioner**.
- IBM Storage Kubernetes FlexVolume for a pod – **\$> kubectl logs -n ubiquity pod ubiquity-k8s-flex<pod_ID>**. In addition, events for all pods on a specific Kubernetes node are recorded in the *ubiquity-k8s-flex.log* file. You can view this file in the following directory: */usr/libexec/kubernetes/kubelet-plugins/volume/exec/ibm~ubiquity-k8s-flex*.
- Controller-manager:
 - Static pod – **kubectl get pods -n kube-system** to display the master pod name. Then, **kubectl logs -n kube-system pod_name** to check the logs.
 - Non-static pod – **journalctl** to display the system journal. Then, search for the lines that have controller-manager entries.

Detecting errors

This is an overview of actions that you can take to pinpoint a potential cause for a stateful pod failure. The table at the end of the procedure describes the problems and provides possible corrective actions.

1. Run the **ubiquity_cli.sh -a status_wide** command to check if:
 - All Kubernetes pods are in Running state.
 - All PVCs are in Bound state.
 - *ubiquity-k8s-flex* pod exists on each master node in the cluster. If you have three master nodes and five worker nodes, you must see a eight *ubiquity-k8s-flex* pods.

Note: The output of the **ubiquity_cli.sh -a status_wide** is similar to the **./ubiquity_cli.sh -a status** output, illustrated in the *What to do next* section of “Performing installation of IBM Storage Enabler for Containers” on page 183.

2. If you find no errors, but still unable to create or delete pods with PVCs, continue to the next step.
3. Display the malfunctioned stateful pod (**\$> kubectl describe pod pod_ID**). Usually, pod description contains information about possible cause of the failure. Then, proceed with reviewing the IBM Storage Enabler for Containers logs.
4. Display the IBM Storage Kubernetes FlexVolume log for the active master node (the node that the controller-manager is running on). Use the **\$> kubectl logs -n ubiquity ubiquity-k8s-flex-<pod_ID_running_on_master_node>** command. As the controller-manager triggers the storage system volume mapping, the log displays details of the FlexVolume attach or detach operations. Additional information can be obtained from the controller-manager log as well.
5. Review the IBM Storage Kubernetes FlexVolume log for the worker node, on which the container pod is scheduled. . Use the **\$> kubectl logs -n ubiquity ubiquity-k8s-flex-<pod_ID_running_on_worker_node>** command. As the *kubelet* service on the worker node triggers the FlexVolume mount and unmount operations, the log is expected to display the complete volume mounting flow. Additional information can be obtained from the *kubelet* service as well, using the **\$> journalctl -u kubelet** command.
6. Display the IBM Storage Enabler for Containers server log (**\$> kubectl logs -n ubiquity deploy/ubiquity** command) or its database log (**\$> kubectl logs -n ubiquity deploy/ubiquity-db** command) to check for possible failures.
7. Display the IBM Storage Dynamic Provisioner log (**\$> kubectl logs -n ubiquity ubiquity-k8s-provisioner**) to identify any problem related to volume provisioning.
8. View the Spectrum Connect log (*hsgsrv.log*) for list of additional events related to the storage system and volume operations.

Table 38. Troubleshooting for IBM Storage Enabler for Containers

Description	Corrective action
IBM Storage Kubernetes FlexVolume log for the active master node has no attach operations	Verify that: <ul style="list-style-type: none"> • Controller-manager pod can access the Kubernetes plug-in directory. See “Compatibility and requirements for IBM Storage Enabler for Containers” on page 179 for instructions on configuring the access. • The correct hostname of the node is defined on the storage systems with the valid WWPN or IQN of the node, as described in “Compatibility and requirements for IBM Storage Enabler for Containers” on page 179. This information appears in the controller-manager log.
IBM Storage Kubernetes FlexVolume log for the worker node that runs the pod has no new entries, except for <i>ubiquitytest</i> (Kubernetes 1.6 or 1.7 only)	Restart the <i>kubelet</i> on Kubernetes worker and master nodes. See “Performing installation of IBM Storage Enabler for Containers” on page 183.
IBM Storage Kubernetes FlexVolume log for the worker node that runs the pod contains errors, related to WWN identification in the <i>multipath -ll</i> output	Check that: <ul style="list-style-type: none"> • Fibre Channel zoning configuration of the host is correct. • The Kubernetes node name is defined properly on the storage system. • Node rescan process was successful.
No connectivity between the FlexVolume pod and the IBM Storage Enabler for Containers server	Log into the node and run the FlexVolume in a test mode (<code>\$> /usr/libexec/kubernetes/kubelet-plugins/volume/exec/ibm~ubiquity-k8s-flex/ubiquity-k8s-flex testubiquity</code>). If there is an error, make sure the IP of <i>ubiquity</i> service is the same as configured in the <i>ubiquity-configmap.yml</i> file. If not, configure the IP properly, then delete the FlexVolume DaemonSet and re-create it to apply the new address value.
Failure to mount a storage volume to a Kubernetes node	If the FlexVolume fails to locate a WWPN within multipath devices, verify your multipathing configuration and connectivity to a storage system. See “Compatibility and requirements for IBM Storage Enabler for Containers” on page 179.
IBM Storage Enabler for Containers database fails to achieve the <i>Running</i> status after the configured timeout expires	<ul style="list-style-type: none"> • Check the <i>kubectl</i> logs for the FlexVolume pod on a node where the database was scheduled to run. Verify that the mount and rescan operations were successful. Another reason might be that the Docker image pulling is taking too much time, preventing the deployment to become active. • Check the <i>kubectl</i> logs for the FlexVolume pod that runs on the master node. Check any error related to attachment of the <i>ibm-ubiquity-db</i> volume. • Check the Kubernetes scheduling. Verify that it does not exceed the timeout configured in the installation script. • After you resolve the issue, verify that the <i>ibm-ubiquity-db</i> status is <i>Running</i>.
IBM Storage Enabler for Containers database persists in the <i>Creating</i> status. In addition, the Volume has not been added to the list of VolumesInUse in the node's volume status message is stored in <i>/var/log/message</i> file on the node, where the database is deployed.	To resolve this, move <i>kube-controller-manager.yaml</i> out and into <i>/etc/kubernetes/manifests/</i> to be recreated the control-manager pod: <pre data-bbox="540 1581 1453 1770"> mv /etc/kubernetes/manifests/kube-controller-manager.yaml /tmp sleep 5 mv /tmp/kube-controller-manager.yaml /etc/kubernetes/manifests/ sleep 15 #check the control-manager pod is running. \$> kubectl get pod -n kube-system grep controller-manager # Verify it is in Running state. </pre>
Persistent volume remains in the Delete state, failing to release	Review the Provisioner log (<code>\$> kubectl logs -n ubiquity deploy/ubiquity-k8s-provisioner</code>) to identify the reason for deletion failure. Use the <code>\$ kubectl delete</code> command to delete the volume. Then, contact the storage administrator to remove the persistent volume on the storage system itself.

Table 38. Troubleshooting for IBM Storage Enabler for Containers (continued)

Description	Corrective action
<p>Communication link between IBM Storage Dynamic Provisioner and other solution elements fails due to Provisioner token expiration</p>	<p>IBM Storage Dynamic Provisioner uses a token that in some environments has an expiration time, for example twelve hours. To keep the link alive for an unlimited time, you can use a <i>service-account</i> token without expiration time. You can replace the current token with the <i>service-account</i> token, as follows:</p> <pre data-bbox="516 380 1403 527"> \$> TOKEN=\$(kubectl get secret --namespace default \$(kubectl get secret --namespace default grep service-account awk '{print \$1}') -o yaml grep token: awk '{print \$2}' base64 -d) \$> kubectl config set-credentials <mycluster.user> --token=\${TOKEN} </pre>
<p>A pod creation fails and the following error is stored in the FlexVolume log of the node intended for the pod: DEBUG 4908 executor.go:63 utils::Execute Command executed with args and error and output. [[{command=iscsiadm}{args=[-m session --rescan]} {error=iscsiadm: No session found.} {output=}]]"</p>	<p>Verify that the node has iSCSI connectivity to the storage system. If the node has none, see the “Compatibility and requirements for IBM Storage Enabler for Containers” on page 179 section for instructions on how to discover and log into iSCSI targets on the storage system.</p>
<p>Status of a stateful pod on a malfunctioned (crashed) node is <i>Unknown</i></p>	<p>Manually recover the crashed node, as described in the “Recovering a crashed Kubernetes node” on page 198 section.</p>
<p>A pod becomes unresponsive, persisting in the <i>ContainerCreating</i> status. The "error=command [mount] execution failure [exit status 32]" error is stored in the FlexVolume log of the node, where the pod was scheduled.</p> <p>The failure occurs because the mountPoint already exists on this node. This might happen due to earlier invalid pod deletion.</p>	<p>Manually recover the pod, using the following procedure:</p> <ol data-bbox="505 1094 1403 1402" style="list-style-type: none"> 1. Check if there is a symbolic link to the mountPoint by running <code>\$> ls -l /var/lib/kubelet/pods/<pod_ID>/volumes/ibm~ubiquity-k8s-flex/<PVC_ID></code>. 2. If the file exists and there is a symbolic link to the /ubiquity/<PVC_WWN>, remove it by running <code>rm /var/lib/kubelet/pods/<pod_ID>/volumes/ibm~ubiquity-k8s-flex/<PVC_ID></code>. 3. Unmount the PV by running <code>umount /ubiquity/<PVC_WWN></code>. 4. Wait for several minutes for Kubernetes to rerun the mountFlow. Then, at the end of the process, display the FlexVolume log by running <code>kubectl logs -n ubiquity ubiquity-k8s-flex-<pod_ID_on_the_node></code> to verify the <i>Running</i> status of the pod.
<p>A pod becomes unresponsive, persisting in the <i>ContainerCreating</i> status. An error indicating a failure to discover a new volume WWN, while running the <code>multipath -ll</code> command, is stored in the FlexVolume log. This log belongs to the node, where the pod was scheduled.</p>	<p>Restart the <i>multipathd</i> service by running the <code>service multipathd restart</code> command on the worker node, where the pod was scheduled.</p>

Chapter 8. IBM Storage Automation Plug-in for PowerShell

The IBM Storage Automation Plug-in for PowerShell offers lightweight commands in the Microsoft PowerShell environment (cmdlets) for provisioning and management of the IBM storage systems.

The plug-in package can be downloaded from Microsoft PowerShell Gallery (<https://www.powershellgallery.com/packages/SpectrumConnect-Client>) and can then be deployed on a PowerShell host. It can also be used together with PowerCLI, Windows PowerShell interface for automating storage-related tasks in VMware vSphere environment.

After the plug-in is installed, storage-related cmdlets can be used in the PowerShell environment. The PowerShell runtime invokes these cmdlets for automated end-to-end storage provisioning, host mapping, volume expansion and other storage-related tasks. .

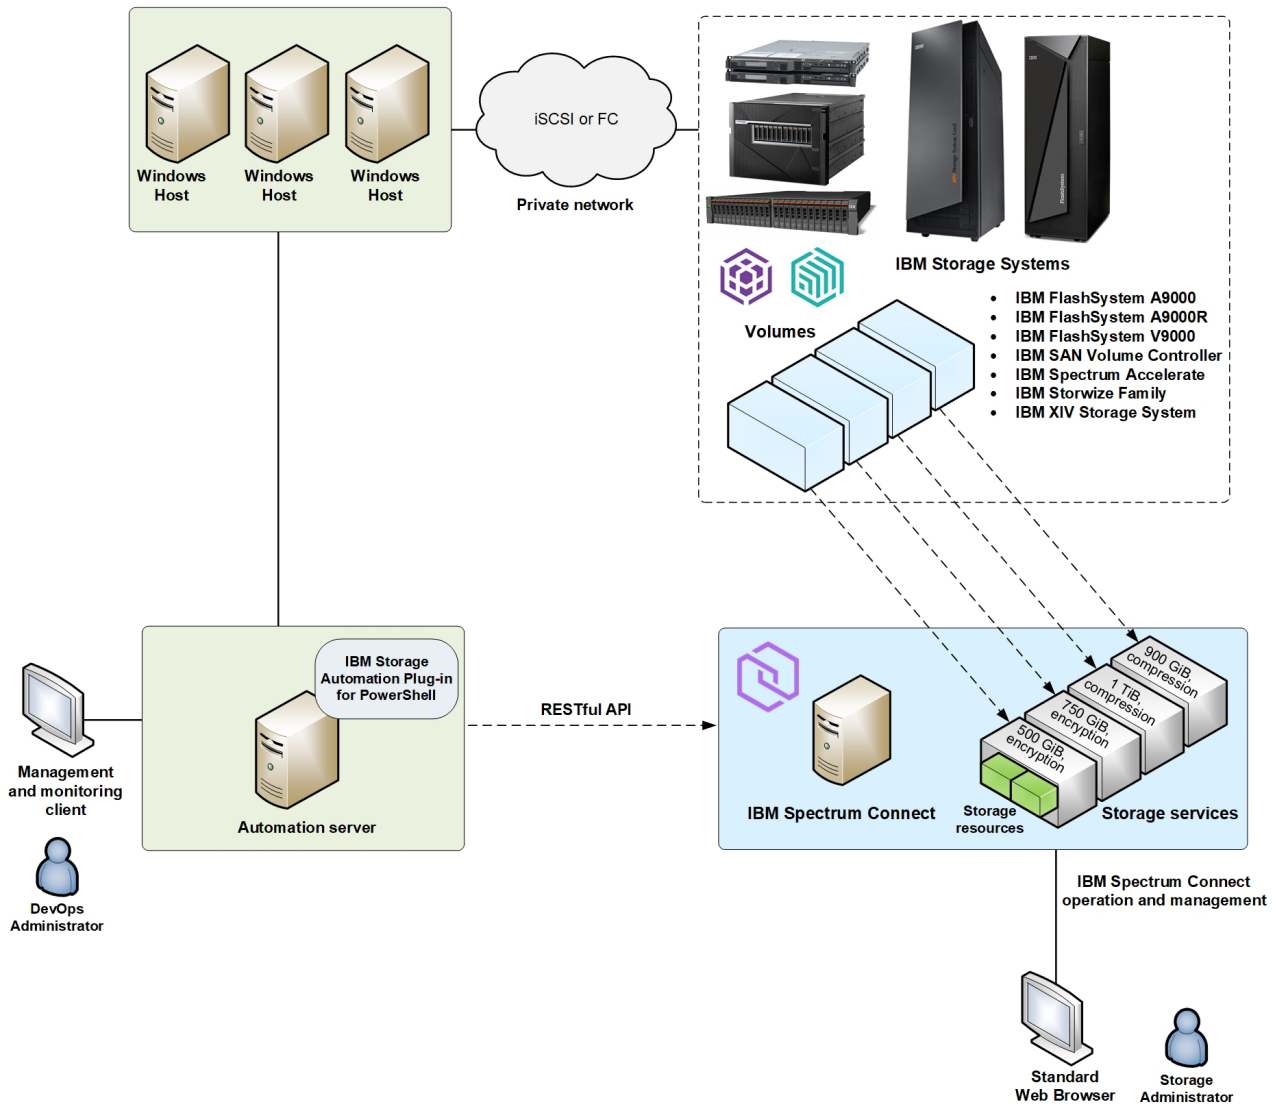


Figure 140. Integration of IBM storage systems in PowerShell environment

Managing integration with Microsoft PowerShell

The IBM Storage Automation Plug-in for PowerShell is used for automated provisioning of storage volumes from an external IBM storage system via PowerShell cmdlets.

About this task

The following procedure details how to add the PowerShell client interface to Spectrum Connect.

Procedure

1. On the **Interfaces** pane, click **Add Interface**, and then select **Add PowerShell**. The **Add New PowerShell Interface** dialog box is displayed.




Figure 141. Add New PowerShell Interface dialog box

2. Enter credentials for the new PowerShell user, and click **Apply**.

Note: When entering a user name and password for a PowerShell interface on the Spectrum Connect with LDAP authentication, make sure that these credentials are the same as defined for LDAP. In addition, you can choose between using a single user or user group, if LDAP is enabled for this Spectrum Connect instance.

The PowerShell interface is now added to Spectrum Connect.




Figure 142. PowerShell interface on the Interfaces pane

What to do next

You can continue integrating the plug-in with Microsoft PowerShell, as explained in the following sections:

- “Downloading and installing the plug-in package for PowerShell” on page 210.
- “Delegating storage services to the PowerShell interface” on page 210.
- “Canceling service delegation to PowerShell” on page 211.

Downloading and installing the plug-in package for PowerShell

To start using PowerShell cmdlets for storage provisioning, you must first download the IBM Storage Automation Plug-in for PowerShell from Microsoft PowerShell Gallery and install it on a Microsoft PowerShell host.

Before you begin

Verify that:

- Windows Management Framework 5.0 has been installed.
- Microsoft PowerShell 5.0 has been installed.
- PowerShell user credentials have been defined in Spectrum Connect.

About this task

The following procedure details how to download, install, and properly configure the IBM Storage Automation Plug-in for PowerShell.

Procedure

1. Go to Microsoft PowerShell Gallery (<https://www.powershellgallery.com/packages/SpectrumConnect-Client>) and download the plug-in to every Microsoft PowerShell host. You can also navigate to the PowerShell Gallery from within the Spectrum Connect GUI by clicking **Download plug-in package** on the bottom of the **PowerShell Settings** dialog box (right-click the PowerShell interface on the **Interfaces** pane, and select **Modify**).
2. Initiate Microsoft PowerShell.
3. In the PowerShell environment, use the **Find-Module** and **Install-Module** commands to locate and install the SpectrumConnect-Client package.
4. Create connection between the IBM Storage Automation Plug-in for PowerShell and Spectrum Connect, using credentials previously defined in Spectrum Connect. Use the following procedure:

```
$client=New-SCConnection -ConnectionUri https://9.115.247.148:8440  
-UserName powershell -Password Passw0rd
```

The session is active now.
5. If you need to uninstall the plug-in, use the **Uninstall-Module** command to uninstall the SpectrumConnect-Client package.

Delegating storage services to the PowerShell interface

Before you can use the IBM Storage Automation Plug-in for PowerShell, you must delegate the storage services that you want to make them available for PowerShell scripts.

About this task

The services and their storage resources that you delegate on Spectrum Connect can be used for issuing PowerShell cmdlets (for more information, see “Using the IBM Storage Automation Plug-in for PowerShell” on page 212).

Procedure

To delegate storage services to the PowerShell:

1. On the **Interfaces** pane, click the PowerShell interface to select it.
2. On the **Spaces/Storage Services** pane, select the storage space from which you want to choose storage services. The available services that reside on the selected storage space are immediately displayed.
3. Right-click on a service that you want to delegate to the PowerShell interface, and then select **Delegate to <PowerShell_interface_name>**, or click the **Attach/Delegate** button on the service. The service color changes to indicate the successful delegation.

You can continue the process by right-clicking available services under the current space.

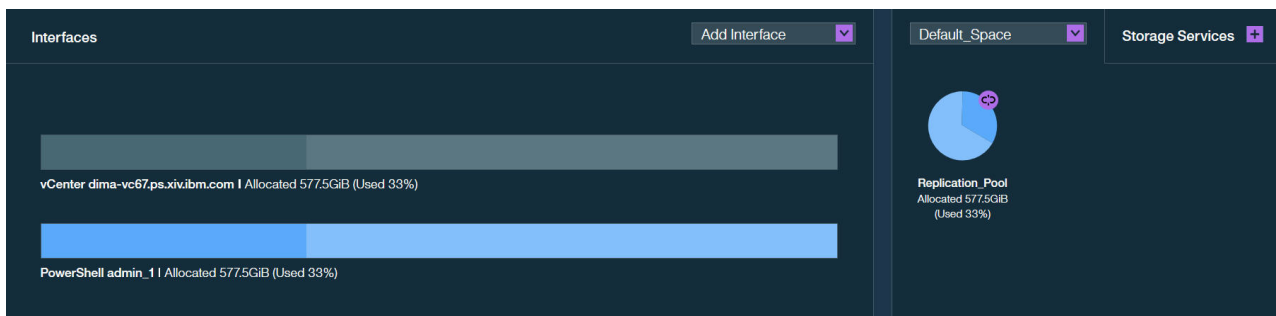


Figure 143. PowerShell interface with a delegated service

The PowerShell interface provides indication for the allocated and used storage space.

- Allocated – amount of storage space available on all pools connected to the delegated services.
- Used – amount of storage space used by PowerShell scripts on all pools connected to the delegated services.

Canceling service delegation to PowerShell

When required, you can cancel a storage service delegation to a PowerShell interface.

About this task

Storage services, which delegation has been canceled, and their resources (pools) cannot be used for running PowerShell cmdlets. Any scripts, involving these elements, will fail to run.

Procedure

1. On the **Interfaces** pane, click the PowerShell interface. The services that are currently delegated to PowerShell interface are highlighted on the **Spaces/Storage Services** pane.
2. Right-click on a service which delegation you want cancel, and then select **Cancel delegation to <interface_name>**, or click the **Detach/Cancel Delegation** button on the service. The service color changes to indicate the successful detachment.

You can continue the process by right-clicking delegated services under the current space.

Using the IBM Storage Automation Plug-in for PowerShell

Use the IBM Storage Automation Plug-in for PowerShell to run storage provisioning cmdlets in the PowerShell environment.

After the IBM Storage Plug-in is deployed (see “Downloading and installing the plug-in package for PowerShell” on page 210), the storage-related cmdlets become available for PowerShell users, as detailed in the table below.

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell

Object/URI	Description	Syntax and parameters
SCConnection	Set up the connection between the Spectrum Connect server and the IBM Storage Automation Plug-in for PowerShell. Before using the PowerShell client, create a PowerShell interface user.	<p>New-SCConnection -ConnectionUri <String> [-Credential] <PSCredential> [<CommonParameters>]</p> <p>New-SCConnection -ConnectionUri <String> [-UserName] <String> [-Password] <String> [<CommonParameters>]</p> <ul style="list-style-type: none"> • ConnectionUri <String> Alias L. The connection URL for Spectrum Connect server in the https://SCIPAddress:ServicePort(8440) format. • Credentials <PSCredential>. User credentials for connection to the Spectrum Connect server. • UserName <String> Alias U. User name and password for the Spectrum Connect server connection must be specified. The user can be created on the Spectrum Connect GUI, when creating a PowerShell interface. Other non-PowerShell-interface users are limited to access some Spectrum Connect functions. • Password <String> Alias P. The corresponding password. <p>For example, connect to the Spectrum Connect server at 9.115.250.45 with user name powershell and password admin2:</p> <pre>PS C:\>\$client=New-SCConnection -ConnectionUri https://9.115.250.45:8440 -UserName powershell -Password admin2</pre> <p>Note: If the token is expired, re-issue New-SCClient to establish a new connection.</p>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
SCSpace	Retrieve information about all storage spaces or show the detailed information according to space ID or name.	<p>Get-SCSpace [[-SpaceID] <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCSpace [-SpaceName] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> • SpaceID <String>. Get the space details via the space ID. • SpaceName <String>. Get the space details via the space name. • SCConnection <SCConnection>The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, display the detailed information, using space ID:</p> <pre>PS C:\>Get-SCSpace -SCConnection \$client -SpaceID efd54ac-70c8-4693-ad7f-3df1cf56c187 id : efd54ac-70c8-4693-ad7f-3df1cf56c187 num_services : 0 unique_identifier : efd54ac-70c8-4693-ad7f-3df1cf56c187 name : Default_Space description : Space used as the default. storage_array_metadata : {}</pre>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
SCService	Retrieve information about storage services delegated to the existing PowerShell interfaces.	<p>Get-SCService [[-ServiceID] <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCService [-ServiceName <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCService [-SpaceID <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCService [-SpaceName <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> • ServiceID <String>. The service ID or its unique identifier. • ServiceName <String>. The service name. • SpaceID <String>. Retrieve service information via its storage space ID. • SpaceName <String>. Retrieve service information via its storage space name. • SCConnection <SCConnection>. The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, display the detailed information, using service ID:</p> <pre>PS C:\>Get-SCService f69be4d0-b56c-449b-a3ae-a727a6ba972d -SCConnection \$client id : f69be4d0-b56c-449b-a3ae-a727a6ba972d unique_identifier : f69be4d0-b56c-449b-a3ae-a727a6ba972d name : service_vvol description : container : 7af7432d-172d-495b-a287-a1df2d2f4b77 capability_values : type : vvol physical_size : 214748364800 logical_size : 214748364800 physical_free : 214748364800 logical_free : 214748364800 total_capacity : 214748364800 used_capacity : 0 max_resource_logical_free : 214748364800 max_resource_free_size_for_provisioning : 214748364800 num_volumes : 0 has_admin : True</pre>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
SCHost	Retrieve information about hosts.	<p>Get-SCHost [[-HostID] <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCHost [-ArrayID] <String> [-HostName] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCHost [-ArrayID] <String> [-Initiator] <Object> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCHost [-ArrayID] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> • HostID <String>. The host ID unique to the Spectrum Connect server. • ArrayID <String>. The ID of the managed storage system. • HostName <String>. The host name. This string is case-sensitive. • Initiators <Object>. iSCSI or FC ports of the host. • SCConnection <SCConnection>The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, display the detailed information about a single host:</p> <pre>PS C:\>Get-SCHost -ID 6 id : 6 array_type : 2145 array : 00000200AE62F4E4 name : fakeFCHost port_count : 2 iogroup_count : 4 status : offline host_type : generic storage_cluster : physical_host : iogroups : {2, 3, 4, 1} initiators : {21000024FF2DAFF6, 21000024FF2DAFF5}</pre>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
SCVolume	<p>Perform volume operations:</p> <ul style="list-style-type: none"> Retrieve information about volumes managed by the PowerShell interfaces. Create a volume via the services delegated to the PowerShell interface. Resize a volume. Remove a volume. 	<p>Get</p> <p>Get-SCVolume [[-VolumeID] <String>] [[-SCConnection] <SCConnection>] [<CommonParameters>]</p> <p>Get-SCVolume [-VolumeName] <String> [[-ArrayID] <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCVolume [-ArrayID] <String> [<CommonParameters>] Get-SCVolume [-ServiceID] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCVolume [-ScsiID] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> VolumeID <String>. The volume ID, its unique identifier or alias ID. VolumeName <String>. The volume name or alias name. The VolumeName is not unique, especially when multiple storage systems are managed by same Spectrum Connect instance. To identify the volume on certain storage system, the combination VolumeName and ArrayID can be used. ArrayID <String>. The ID of the managed storage system. This parameter can be used separately, or together with VolumeName. ServiceID <String>. The ID of the service managing the volume. SCConnection <SCConnection>. The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, retrieve detailed information of a volume on a managed storage system:</p> <pre>PS C:\>Get-SCVolume -VolumeName SCVOL304 -ArrayID 00000200AE62F4E4 -SCConnection \$client scsi_identifier : 6005076802B98BD3900000000000000AD array_type : 2145 array : 00000200AE62F4E4 array_name : V7K_71 id : 6005076802B98BD3900000000000000AD pool_name : SC61_THIN pool_id : 3 max_extendable_size : 6596353938773 service_compliance : True domain_name : service_name : ThinProvisionService container_name : DemoSpace service_id : 8d4b6c91-670f-40bd-8d69-901e4a1219f9 container_id : 4cfc8298-965f-4f90-b23f-4d04fe70d9ee storage_model : Storwize V7000 volume_id : 24 name : SCVOL304 logical_capacity : 10737418240</pre>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
		<p>New</p> <p>New-SCVolume [-VolumeName] <String> [-Size] <Int64> [[-SizeUnit] <String>] [-ServiceID] <String> [[-Initiator] <Object>] [[-SCConnection] <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> • VolumeName <String>. The name of a new volume. • Size <Int64>. The size of a new volume. This parameter is together with SizeUnit. • SizeUnit <String>. The unit of the new volume size (Byte, GB, GiB, TiB). Default setting is GB. • ServiceID <String>. The ID of the service managing the volume. • Initiators <Object>. iSCSI or FC ports of the host to be mapped to the new volume. System.Collections.ArrayList object or just a Object[] can be used for multiple initiators. • SCConnection <SCConnection>The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, create a new volume for a specific service:</p> <pre>PS C:\>New-SCVolume -Name DemoThin04 -Size 1 -SizeUnit GiB -ServiceID \$service.id ft array_name,id,name,pool_name,service_name, logical_capacity array_name : cim75 id : 6005076801A707416800000000000B76 name : DemoThin04 pool_name : SCReserved service_name : SVCThinService container_name : DemoSpace logical_capacity : 1073741824</pre>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
		<p>Resize</p> <p>Resize-SCVolume [-VolumeID] <String> [-NewSize] <Int64> [[-SizeUnit] <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> • VolumeID <String>. ID of the volume to be resized. • NewSize <Int64>. The new size of the volume. This parameter is together with SizeUnit. • SizeUnit <String>. The unit of the new volume size (Byte, GB, GiB, TiB). Default setting is GB. • SCConnection <SCConnection>The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, create a new volume for a specific service:</p> <pre>Resize-SCVolume -ID 6005076801A7074168000000000000B76 -NewSize 2 -SizeUnit GiB ft name,id,logical_capacity name id logical_capacity ----- DemoThin03 6005076801A7074168000000000000B74 2147483648</pre>
		<p>Remove</p> <p>Remove-SCVolume [-VolumeID] <Object> [[-Force]] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> • VolumeID <String>. ID of the volume to be removed. • Force <SwitchParameter>. Deletes mapping relations of the volume to be removed. • SCConnection <SCConnection>The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, remove a volume:</p> <pre>PS C:\>Remove-SCVolume -VolumeID 6005076801A7074168000000000000B85 -SCConnection \$client True</pre>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
SCHostVol Mapping	<p>Perform volume mapping operations:</p> <ul style="list-style-type: none"> Retrieve information about host-volume mapping. Map the volume to the host. Unmap the volume from the host. 	<p>Get</p> <p>Get-SCHostVolMapping [[-HostVolMappingID] <String>] [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCHostVolMapping [-HostID] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Get-SCHostVolMapping [-VolumeID] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> HostVolMappingID <String>. The mapping relation ID, its unique identifier or alias ID. HostID <String>. Displays mapping relations of the host with the specified ID. VolumeID <String>. Displays mapping relations of the volume with the specified ID. SCConnection <SCConnection>The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, display mapping relations for the host with ID 5:</p> <pre>PS C:\>Get-SCHostVolMapping -HostID 5 -SCConnection \$sc ft id volume host lun_number ----- 113 6005076802B98BD3900000000000000AC 5 2 3 6005076802B98BD390000000000000018 5 1 4 6005076802B98BD390000000000000017 5 0 111 6005076802B98BD3900000000000000AB 5 7 107 6005076802B98BD3900000000000000A8 5 5</pre>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
		<p>New</p> <p>New-SCHostVolMapping [-VolumeID] <String> [-HostID] <Object> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>New-SCHostVolMapping [-VolumeID] <String> [-HostName] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>New-SCHostVolMapping [-VolumeID] <String> [-Initiators] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> • VolumeID <String>. ID of the volume to be mapped. • HostID <String>. ID of the host that the specified volume will be mapped to. • Initiators <Object>. The host endpoint (iSCSI or FC). • SCConnection <SCConnection>The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, map a single volume to multiple hosts:</p> <pre>PS C:\>New-SCHostVolMapping -VolumeID 6005076801A707416800000000000B76 -HostIDs 22,21,1</pre> <pre>{@{id=21; volume=6005076801A707416800000000000B76; host=22; lun_number=1}}{@{id=21; volume=6005076801A707416800000000000B76; host=22; lun_number=4}, @{}; volume=6005076801A707416800000000000B76; host=21; lun_number=4}, @{};</pre>

Table 39. Cmdlets available via IBM Storage Automation Plug-in for PowerShell (continued)

Object/URI	Description	Syntax and parameters
		<p>Remove</p> <p>Remove-SCHostVolMapping [-VolumeID] <String> [-HostID] <Object> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Remove-SCHostVolMapping [-VolumeID] <String> [-HostName] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <p>Remove-SCHostVolMapping [-VolumeID] <String> [-Initiators] <String> [-SCConnection <SCConnection>] [<CommonParameters>]</p> <ul style="list-style-type: none"> • VolumeID <String>. ID of the volume to be unmapped. • HostID <String>. ID of the host that the specified volume will be unmapped from. • Initiators <Object>. The host endpoint (iSCSI or FC). • SCConnection <SCConnection>The client handle for Spectrum Connect server. <p>It can be created via New-SCConnection. If the SCConnection is not specified, a most recently used SCConnection is utilized by default. The connection can be found via the global variable \$Global:DefaultConnection. If multiple Spectrum Connect servers are connected, you must explicitly specify the SCConnection parameter.</p> <p>For example, unmap a volume from two hosts:</p> <pre style="border: 1px solid black; border-radius: 10px; padding: 5px;">PS C:\>Remove-SCHostVolMapping -VolumeID 6005076801A707416800000000B76 -HostIDs 22,21 -SCConnection \$client</pre>

Chapter 9. Best practices and use cases

Refer to the general guidance and best practices that are described in the following sections.

- “Handling datastores”
- “Handling ESXi hosts that use storage volumes”
- “Distributing volumes evenly on DS8000 systems”
- “Setting the multipath policy for DS8000 and Storwize Family systems” on page 224
- “Working with multiple storage systems” on page 224
- “Upgrading or installing Spectrum Connect with vSphere failover” on page 225
- “Creating a VVol-enabled service” on page 225
 - “Creating a VVol-enabled service on XIV storage systems” on page 226
 - “Creating a VVol-enabled service on storage systems that run IBM Spectrum Virtualize” on page 227
- “Additional reference information” on page 236

Handling datastores

For best performance of VMware datastores:

- Create each datastore on a separate storage volume.
- If you use snapshots/mirroring for volumes, place all Datastore Extents volumes (the building block LUNs of a datastore) in a consistency group (defined by using the storage system GUI or CLI).
- Avoid storing static files in a VVol-enabled datastore. When a new folder is created in the root path of a VVol-based datastore, a 4-GB virtual volume is defined, which cannot be extended.

Handling ESXi hosts that use storage volumes

For the best performance of ESXi hosts that use volumes provisioned on a Spectrum Accelerate or Spectrum Virtualize storage system, define all ESXi hosts within a cluster as cluster hosts on the storage system as well.

Following this practice prevents complications associated with the mapping of a storage volume to different ESXi hosts in a cluster using different LUN numbers.

Distributing volumes evenly on DS8000 systems

DS8000 storage systems have two rank groups, 0 and 1, each managed by a single server. In addition, each DS8000 extent pool is based on one rank group.

Accordingly, it is recommended to spread volumes evenly across the DS8000 systems. Spreading the volumes equally on the extent pools of rank groups 0 and 1 balances the workload across the DS8000 system.

Setting the multipath policy for DS8000 and Storwize Family systems

When using the IBM Storage Enhancements for VMware vSphere Web Client, the recommended multipath policy for DS8000 and Storwize Family (including SAN Volume Controller) storage systems is **Round Robin**.

If you are using VMware ESXi servers of version 5.5 or later, the **Round Robin** multipath policy is enforced by default. However, if you are using earlier ESX or ESXi versions, the **Fixed** policy is chosen by default, and it is recommended to change the multipath policy on those servers to **Round Robin**.

For information about how to change the default multipath policy enforcement for earlier ESX or ESXi versions, refer to article 1017760 on the VMware Knowledge Base website (kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1017760).

Working with multiple storage systems

IBM Spectrum Connect, running on RHEL 6.3–6.6 64-bit operating systems, utilizes CPU and memory resources in accordance with the amount of objects monitored by the vROps server.

Before you begin

For best results, deploy Spectrum Connect on a dedicated server. The minimum hardware requirements are detailed in the latest release notes, available on the IBM Knowledge Center or on the IBM Fix Central (www.ibm.com/support/fixcentral).

About this task

Any increase in the amount of the monitored objects requires additional hardware resources and population interval adjustment. The recommended hardware and software requirements for different amounts of monitored objects are as follows:

- Up to 1000 objects per storage system:
 - Up to 20 storage systems – 64-bit dual-core CPU, 4 GB of RAM
 - Up to 40 storage systems – 64-bit quad-core CPU, 6 GB of RAM
 - Up to 100 storage systems – 64-bit six-core CPU, 6 GB of RAM
- Up to 5000 objects per storage system:
 - Up to 20 storage systems – 64-bit quad-core CPU, 4 GB of RAM
 - Up to 40 storage systems – 64-bit six-core CPU, 6 GB of RAM, 15 min. population interval
 - Up to 100 storage systems – 64-bit six-core CPU, 6 GB of RAM, 25 min. population interval

The procedure for adjusting population interval is detailed below.

Procedure

To adjust a population interval:

1. Go to the `/opt/ibm/ibm_spectrum_control/conf.d/hsgsvr` directory.
2. Edit the `hsgsvr_config.ini` file and change the population interval in the **populate_arrays_and_events** parameter to a desired value in minutes.
3. Save the changes.

- Restart the Spectrum Connect service (service ibm_spectrum_connect restart).

Upgrading or installing Spectrum Connect with vSphere failover

Spectrum Connect can be upgraded or installed together with vSphere failover procedure.

Before you begin

- Verify that at least two Spectrum Connect instances are registered as storage providers on a vCenter server.
- Both Spectrum Connect instances must be online, one of them must be active and the other one – standby.

About this task

To minimize system downtime, you can upgrade the existing Spectrum Connect or install its new release along with completing failover between active and standby instances in vSphere environment.

Procedure

- Copy the upgrade or installation package files to a temporary folder on a virtual machine that is hosting Spectrum Connect.
- Upgrade or install the standby Spectrum Connect instance. See “Upgrading an existing installation” on page 12 or “Performing first-time installation of Spectrum Connect” on page 8.
- Verify the standby instance connectivity to a storage system, by using the **sc_storage_array list** command. A proper connectivity to a storage system is indicated by *Yes* in the *Connected* field, as illustrated in the following example.

Array Alias	Array Identifier	Management IP Addresses	Elapsed time since last update	Connected	Notes
SVC232	0000020062A1D16C	9.115.246.232	7 minutes	Yes	

- Complete the active Spectrum Connect failover to the standby instance, by restarting the guest operation system of the virtual machine that is hosting the active Spectrum Connect. The failover process can take up to 8 minutes to complete. As a result, the standby instance becomes active, running the newly upgraded or installed Spectrum Connect.
- Repeat the upgrade/install and failover processes for the remaining Spectrum Connect instances.

Creating a VVol-enabled service

The IBM Spectrum Connect introduces a comprehensive storage virtualization support, using VMware virtual volume (VVol) technology.

Before you begin

Note: The virtual volume functionality is supported by the IBM XIV (11.5.1 or later) and storage systems that run IBM Spectrum Virtualize (7.6 or later).

- Verify that all required ESXi hosts are connected and defined at the storage system side.
- Verify that the Spectrum Connect time is synchronized with the time used by the vCenter server

About this task

This section details how to create a VVol-enabled storage service on XIV or storage systems that run IBM Spectrum Virtualize. The service or a group of services can be used to define storage spaces, serving as virtual datastores for VM deployment. See the following sections, depending on the storage system in use:

- “Creating a VVol-enabled service on XIV storage systems”
- “Creating a VVol-enabled service on storage systems that run IBM Spectrum Virtualize” on page 227

Creating a VVol-enabled service on XIV storage systems

The IBM Spectrum Connect features a comprehensive storage virtualization support, using VMware Virtual Volume (VVol) technology.

About this task

This section details how to create a VVol-enabled storage service on XIV storage systems.

To create a VVol-enabled storage service:

- Enable VVol utilization at the XIV side.
- Create the service using Spectrum Connect.

Procedure

1. Launch the XIV management GUI and log in as a storage administrator.
2. Create a domain with required soft and hard capacity. Make sure that the soft capacity is four times larger than the hard capacity.
3. For XIV and Spectrum Accelerate storage systems, define a user with category *storageintegrationadmin*.
4. Associate the *storageintegrationadmin* user with the domain.
5. Associate all ESXi hosts with the domain.

Important: The managed domain that you created cannot be used for traditional volumes without virtualization. You must create a separate regular domain for them. This domain must have the same user and the ESXi hosts that you intend to manage. However, you need to create a separate storage resource and a new service on the regular domain via Spectrum Connect for subsequent use by the VMware vWC.

6. Launch the XIV XCLI tool.
7. Enable the metadata service for the XIV, using the `metadata_service_enable` command.
8. Close the XIV XCLI tool and return to the XIV management GUI.
9. Change the storage administrator user to *storageintegrationadmin* user and re-launch the XIV XCLI tool
10. Create a new Administrative Logical Unit (ALU) per each ESXi host. Use the following XCLI format:
`alu_create alu=<alu-name> host=<host-name> lun=logical-unit-number.`
Make sure that the LUN is in the 512–755 range.
11. Launch Spectrum Connect.

12. Configure a fully qualified domain name for the Spectrum Connect server and define a high-availability group. See “Defining a high-availability group” on page 35.
13. Generate a self-signed Spectrum Connect server certificate. See “Managing server certificates” on page 37.
14. Set up VASA credentials. See “Setting the VASA credentials” on page 101.
15. Enable the *storageintegrationadmin* to access the XIV storage resources. See “Entering the storage system credentials” on page 44.
16. Add the XIV storage system to the Spectrum Connect. See “Adding a storage system” on page 46.
17. Add a new storage space. See “Adding a storage space” on page 55.
18. Add a VVol-enabled service to the storage space. See “Adding a storage service” on page 56.
19. Define a storage resource and attach it to the VVol-enabled service. See “Defining and attaching storage resources” on page 60.
20. Register Spectrum Connect as a storage provider on VMware vCenter server. See “Registering Spectrum Connect as a storage provider on vCenter server” on page 103.
21. Launch the vWC and create a VVol-enabled datastore. Select the storage space that you defined as an underlying storage resource for the datastore.

Note: The storage space is presented as a combination of a VVol container and a storage system. For example, storage space A and storage system A are visible as container 1; storage space A and storage system B are visible as container 2. This information is displayed in the *storage container:storage system* format.

Creating a VVol-enabled service on storage systems that run IBM Spectrum Virtualize

The IBM Spectrum Connect features a comprehensive storage virtualization support, using VMware Virtual Volume (VVol) technology.

About this task

This section details how to create a VVol-enabled storage service on storage systems that run IBM Spectrum Virtualize (7.6 or later).

To create a VVol-enabled storage service:

- Enable VVol utilization at the side of a storage system that runs IBM Spectrum Virtualize.
- Create the service, using Spectrum Connect.

Procedure

1. Activate the storage system CLI utility and log in as *SecurityAdmin*, and then create a user group with role *VASAProvider* (**mkusergrp -name <vasa_group_name> -role VASAProvider**).
2. Create a *VASAProvider* user in the user group (**mkuser -name <user_name> -usergrp <vasa_group_name>**).
3. Create a metadata volume (**mkmetadatavdisk -mdiskgrp <pool_name>**). For the storage systems that run microcode 7.6.1.0 or later, this operation requires the

SecurityAdmin access level. Also, on these storage systems, you can enable the VVol feature by toggling the Enable VVOL switch under Settings -> System -> VVOL, using the web GUI.

4. For each ESXi host, define the host on the storage system as *adminlun* type (`svctask mkhost -name <host_name> -fcwwpn <fibre_channel_wwpn> -iscsiname <iscsi_wwn> -type adminlun`).
5. Launch Spectrum Connect.
6. Define a high-availability group. See “Defining a high-availability group” on page 35.
7. Generate a self-signed Spectrum Connect server certificate. See “Managing server certificates” on page 37.
8. Set up VASA credentials. See “Setting the VASA credentials” on page 101.
9. Enable the *VASAProvider* user to access the storage resources. See “Entering the storage system credentials” on page 44.
10. Add the storage system to the Spectrum Connect. See “Adding a storage system” on page 46.
11. Add a new storage space. See “Adding a storage space” on page 55.
12. Add a VVol-enabled service to the storage space. See “Adding a storage service” on page 56.
13. Define a storage resource and attach it to the VVol-enabled service. See “Defining and attaching storage resources” on page 60.
14. Register Spectrum Connect as a storage provider on VMware vCenter server. See “Registering Spectrum Connect as a storage provider on vCenter server” on page 103.
15. Launch the vWC and create a VVol-enabled datastore. Select the storage space that you defined as an underlying storage resource for the datastore.

Note: The storage space is presented as a combination of a VVol container and a storage system. For example, storage space A and storage system A are visible as container 1; storage space A and storage system B are visible as container 2. This information is displayed in the *storage container:storage system* format.

Configuring an LDAP user for a managed domain

To ensure that proper LDAP authentication is used in a storage system managed domain, you must create a user with the storage integration administrator role on the Active Directory server.

Before you begin

Verify that you configured the following entities:

- A managed domain on your storage system. In the example of the configuration procedure, the *dana-domain* domain is used as a managed domain on XIV storage system.
- Operational Microsoft Active Directory service with an active group and a user attached to the group. In the example of the configuration procedure, they are illustrated as the *xivstorage* group and the *xivuser* user.

About this task

The following procedure details how to configure an LDAP user for a managed domain.

Procedure

1. Start the XIV management GUI and log in as a storage administrator.
2. Go to **Systems > System Settings > LDAP**. The **LDAP** dialog box is displayed.

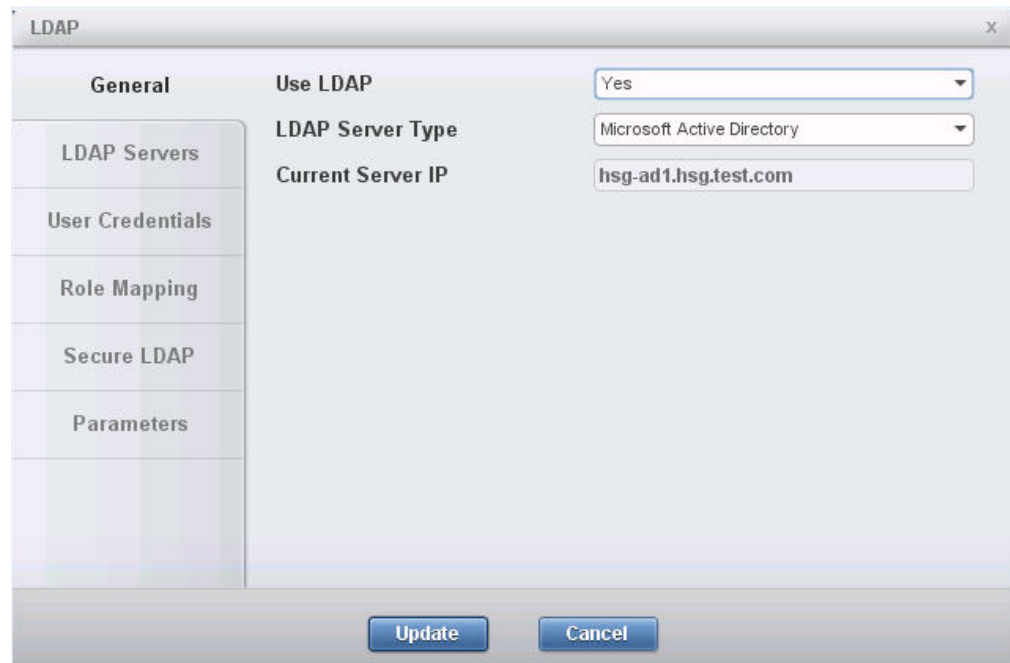


Figure 144. General tab, LDAP dialog box

3. On the **General** tab, enable the LDAP use and select the Microsoft Active Directory, as a directory service. Then, click **Update**.
4. Go to the **LDAP Servers** tab, and verify that the FQDN and IP address of the Active Directory server are correct.
5. Go to the **User Credentials** tab, and define the service user (xivuser in the example) and its password. This user is bound to the Active Directory service. It retrieves credentials data, which is stored in the LDAP directory. Then, click **Update**.




Figure 145. User Credentials tab, LDAP dialog box

6. Go to the **Role Mapping** tab, and set the necessary values for the user attributes, group attributes, and roles. Pay attention to the Storage Integration Admin Role setting (xivstorageintegrationadmin in the example). This parameter, along with the managed domain name, is used as a group name on the Active Directory server. Then, click **Update**.




Figure 146. Role Mapping tab, LDAP dialog box

Note: On the **Role Mapping** tab, the values of the **Storage Admin Role** and **Storage Integration Admin Role** parameters appear truncated. The full value designations are as follows:

- **Storage Admin Role:** *CN=xivstorage,CN=Users,DC=hsg,DC=test,DC=com.*
 - **Storage Integration Admin Role:**
CN=xivstorageintegrationadmin,CN=Users,DC=hsg,DC=test,DC=com.
-

7. Start your Active Directory management software and go to the group configuration section.
8. In the group configuration section, add a new group with the following attributes:
 - Group name: **xivstorageintegrationadmin@dana-domain**. The group name must be the same as the Storage Integration Admin Role setting on XIV (**xivstorageintegrationadmin**) and it must include the name of the XIV managed domain (**dana-domain**).
 - Description: **StorageIntegrationAdmin**
 - Group type: **Security**
 - Group scope: **Global**




Figure 147. Group configuration on Active Directory server

9. Go to the user configuration section, create a new user and add it to the xivstorageintegrationadmin@dana-domain group. The user should have the following attributes:
 - Full name: **danasia**
 - User UPN logon: **danasia@hsg.test.com**
 - User SamAccountName logon: **hsg* danasia**
 - Description: **StorageIntegrationAdmin**
 - Member of: **xivstorageintegrationadmin@dana-domain**

The screenshot shows the user configuration interface for 'danasia' on an Active Directory server. The interface is organized into three main sections:

- Account:**
 - First name: danasia
 - Middle initials: (empty)
 - Last name: (empty)
 - Full name: danasia
 - User UPN logon: danasia@hsg.test.com
 - User SamAccountName: hsg\ danasia
 - Account expires: Never, End of []
 - Password options:
 - User must change password at next log on
 - Other password options
 - Smart card is required for interactive log on
 - Password never expires
 - User cannot change password
 - Encryption options: (dropdown)
 - Other options: (dropdown)
 - Protect from accidental deletion
 - Log on hours... Log on to...
- Organization:**
 - Display name: danasia
 - Job title: (empty)
 - Office: (empty)
 - Department: (empty)
 - E-mail: (empty)
 - Company: (empty)
 - Web page: (empty)
 - Manager: (empty) [Edit... Clear]
 - Phone numbers:
 - Main: (empty)
 - Home: (empty)
 - Mobile: (empty)
 - Fax: (empty)
 - Pager: (empty)
 - IP Phone: (empty)
 - Direct reports: (empty) [Add... Remove]
 - Address:
 - Street: (empty)
 - City: (empty) State/Province: (empty) Zip/Postal code: (empty)
 - Country/Region: (dropdown)
 - Description: StorageIntegrationAdmin
- Member Of:**
 - Filter: (empty)
 - Table:

Name	Active Directory Domain	Services Folder	Primary
Domain Users	hsg.test.com/Users		
xivstorageintegrationadmin@dana-domain	hsg.test.com/Users		✓
 - [Add... Remove]

Figure 148. User configuration on Active Directory server

10. Use the following XCLI commands to verify the LDAP configuration:

- Run the **ldap_mode_get** command to make sure that the LDAP authentication is active:

```
>>ldap_mode_get
Mode
-----
Active
```

- Run the **ldap_test** command to verify that the LDAP user xivuser has been configured correctly:

```
>>>> ldap_test fqdn=hsg-ad1.hsg.test.com user=xivuser password=<password>
command 0:
administrator:
  command:
    code = "SUCCESS"
    status = "0"
    status_str = "Command completed successfully"
aserver = "DELIVERY_SUCCESSFUL"
```

- Run the **ldap_test** command again to verify that the LDAP storage integration admin user **danasia** has been configured correctly:

```
>>>> ldap_test fqdn=hsg-ad1.hsg.test.com user=danasia password=<password>
command 0:
administrator:
  command:
    code = "SUCCESS"
    status = "0"
    status_str = "Command completed successfully"
aserver = "DELIVERY_SUCCESSFUL"
```

11. Start IBM Spectrum Connect and go to **Setting > Storage Credentials**. The Storage Credentials dialog box is displayed.




Figure 149. Storage Credentials dialog box

12. In the Storage Credentials dialog box, enter the user name defined on the Active Directory server (**danasia**), define a password, and select the **Directory account** check box to specify that the credentials are stored on the Active Directory server.
13. Click **Apply** to finish.

Restoring VVol-based virtual machines

You can restore virtual machines deployed on virtual volumes in case of Spectrum Connect failure.

Before you begin

Verify that:

- VVol metadata database is available on a storage system.
- Spectrum Connect installation package is ready for deployment.

About this task

This section details how to restore VMs deployed on virtual volumes, if all Spectrum Connect instances have been permanently destroyed and no backups are available. This scenario may occur, when VMs, hosting Spectrum Connect, were migrated to VVols.

To restore VVol-based VMs:

- Obtain a name of the VVol metadata database located a storage system. This example illustrates procedure for storage systems that run IBM Spectrum Virtualize.
- Install and configure a new copy of Spectrum Connect.

Procedure

1. Activate the storage system CLI utility and log in as *VASAProvider*. The code example below uses the *vasa-admin* ID in the *VASAProvider* role.
2. Display the name of the VVol metadata database, by running the **svctask metadata_db_list** command.

```
IBM_Storwize:vasa-admin> svctask metadata_db_list
DB_name
vvol-db
```

Note: If you have several matching databases on your storage system, use the **svctask metadata_entry_list -db <db_name>** command to display the database contents. A database with multiple entries is the most likely candidate to store the VVol metadata information.

3. Install Spectrum Connect, as explained in the previous chapters.
4. In the **General Settings** dialog box, use the VVol metadata database name (*vvol-db* in example below), as the HA group designation.

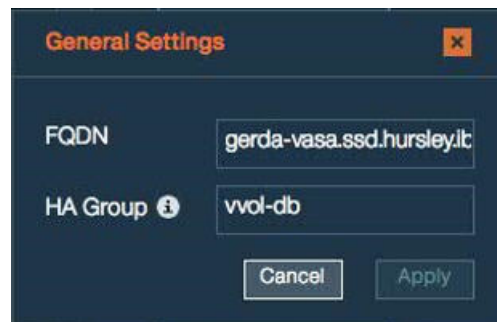


Figure 150. Defining *vvol-db*, as HA group name

5. In the **Storage Credentials** dialog box, define the user name, according to the ID name, defined in step 1 of this procedure. This allows the new Spectrum Connect instance to communicate with the storage system.

Spectrum Connect begins populating relevant storage space with the information recovered from the database on the storage system. Depending on the amount of information, this can take some time to complete. You may need to refresh your browser to see the updated information in the Spectrum Connect GUI.

6. Complete the Spectrum Connect configuration as usual, registering it as a storage provider on vCenter server.

7. Launch the vSphere Web Client for registered vCenter servers and power on your VVol-based VMs as normal.

Additional reference information

IBM provides additional reference documentation to help users integrate, implement and operate realistic scenarios that include Spectrum Connect and various IBM storage systems.

IBM Redbooks are technical documents where IBM experts share their expertise and best practices for using Spectrum Connect. Among others, you can refer to these publications:

- Using the IBM Spectrum Accelerate Family in VMware Environments
- IBM Spectrum Connect and IBM Storage Enabler for Containers - Practical Example with IBM FlashSystem A9000

Chapter 10. RESTful API for Spectrum Accelerate software-defined storage

RESTful API for IBM Spectrum Connect provides an application programming interface (API) for retrieving status information and performing basic maintenance of IBM Spectrum Accelerate™, deployed as a software storage solution.

Note: All other storage systems, except the deployable IBM Spectrum Accelerate, are not supported by the RESTful APIs.

RESTful API for Spectrum Connect relies on a transport protocol to process the following requests.

- Query – Returns object entities managed by Spectrum Connect.
- Create – Creates a new entity.
- Delete – Deletes an existing entity.
- Update – Performs a partial update of an existing entity.
- Action – Performs a complimentary action (for example, phases out a disk on a specific storage system).

RESTful API protocol

The API uses HTTP as the transport protocol and relies on HTTP for some features (for example, security).

The protocol relies on HTTP methods to support CRUD operations (Create/Read/Update/Delete).

- POST – Create and action operations.
- GET – Query (read) operations.
- DELETE – Delete operations.
- PATCH – Partial update operations.

The API URL is built from scope specifiers and resource identifiers.

- **api** – Specifies the main API scope.
- **v1** – Specifies the protocol version.
- **resource type** – Specifies type of the system resource.
- **id** – Identifies the resource.

For example, GET /api/v1/disks/7.

RESTful API for Spectrum Connect returns the following codes for request status:

- **200** – The request has succeeded.
- **401** – The request requires user authentication.
- **405** – The request method is not allowed for the specified resource.
- **500** – The request cannot be implemented due to an unexpected condition.

Note: Additional codes may be used according to the requirements of the REST standards.

Query request and response

RESTful API for IBM Spectrum Connect uses query requests to return the state of a single resource or a list of resources.

You can run a query by issuing the **GET** command (on a URI). The following list shows the query format types:

- A single resource – Single resource properties are returned.
- A collection of resources – A list of resources and their properties is returned.

For example, GET /api/v1/interfaces/22 lists the properties of the interface with identification number 22, as illustrated below.

```
{
  "name": "aaab",
  "array": "pu16",
  "module": "1:Module:1",
  "type": "iSCSI",
  "address": "3.3.3.5",
  "netmask": "255.255.255.0",
  "gateway": "3.3.3.254",
  "address6": "",
  "gateway6": "",
  "mtu": "1500",
  "ports": "1",
  "id": 22
}
```

Note: Query requests can contain additional URI arguments, such as filtering parameters. When used, it returns only resources that match your filtering criteria. For example, GET /api/v1/interfaces?array=pu21 captures interfaces that exist on array *pu21*.

The **GET** command can be used to monitor a running task. For example, GET /api/v1/tasks/04fc6120-60ae-4182-baa9-687d6ae96ffe returns the current task status:

```
{
  "reason": null,
  "task_id": "04fc6120-60ae-4182-baa9-687d6ae96ffe",
  "start_time": "2015-03-03T06:57:48.573",
  "task_state": "Running",
  "array_id": "2810-999-PR16118",
  "name": "disk phase-in"
}
```

A task ID is generated, when the task is initiated by the action request. See “Action request and response” on page 240.

Create request and response

RESTful API for IBM Spectrum Connect uses create requests for object creation.

You can create an object by issuing the **POST** command followed by request object. The request must contain a set of parameters required for object creation.

For example, POST /api/v1/interfaces entered with the parameters detailed below, creates an interface on the specified storage system module.

```
name aaab
address 3.3.3.3
netmask 255.255.255.0
gateway 3.3.3.254
array pu16
module 1:Module:1
ports 1
```

The response to the POST request is as follows:

- OK, followed by the new object, if the task is completed successfully.
- An error message, detailing a reason for the failed request, as illustrated below.

```
{
  "detail": "One of the physical ports specified is already assigned to an IP Interface"
}
```

Delete request and response

RESTful API for IBM Spectrum Connect uses delete requests for object deletion.

You can delete an object by issuing the **DELETE** command on a full URL. The request must contain a single object, which is a target of the delete request.

For example, `DELETE /api/v1/interfaces/22` deletes the interface 22.

The response to the delete request is OK if the task is completed successfully, or an error message with a status code, detailing a reason for the failed request, as illustrated below.

```
{
  "detail": "Not found"
}
```

Update request and response

RESTful API for IBM Spectrum Connect uses update requests for partial object update.

You can update an object by issuing the **HTTP PATCH** command on a full URL. The request must contain a single object, which is the target of the update request.

For example, `PATCH /api/v1/interfaces/26` entered with the parameters detailed below, changes the name and IP address of the interface 26.

```
name aaac
address 3.3.3.5
```

The response to the update request is as follows:

- OK, followed by the updated object, if the task is completed successfully
- An error message with a status code, detailing a reason for the failed request, as illustrated below.

```
{
  "detail": "IP address specified for the default gateway is not in the subnet of the IP Interface"
}
```

Action request and response

RESTful API for IBM Spectrum Connect uses action requests to perform complimentary actions on the requested objects, which are out of scope of other request types.

You can initiate an action by issuing the **POST** command (on a URI). For example, `POST /api/v1/disks/82/phasein` starts a phase-in procedure for disk 82.

The response to the request is OK, followed by the task ID, as illustrated below. The task ID can be used in a query request to monitor the task progress, see "Query request and response" on page 238.

```
OK
{
  "task id": "04fc6120-60ae-4182-baa9-687d6ae96ffe"
}
```

If the request fails, an error message is displayed, carrying a status code with a reason for the failed request.

Storage system operations

RESTful API for IBM Spectrum Connect uses the **GET** and **POST** commands for storage system (array) operations.

You can run a storage system query by issuing the **GET** command to retrieve a list of storage systems and their properties, or display the properties of a single storage system.

For example, `GET /api/v1/arrays/2810-999-dc21011` lists the properties of the storage with identification number 2810-999-dc21011, as illustrated below.

```
OK 200
{
  "id": 2810-999-dc21011,
  "alias": "pu21",
  "mgmt_addresses": [
    "9.151.153.87",
    "9.151.153.86",
    "9.151.153.39"
  ],
  "last_updated": "2015-04-15T17:55:02.431",
  "error_message": "",
  "connected": true,
  "name": "XIV pur21m10m11m15",
  "firmware_version": "11.5.0",
  "scsi_model_identifier": "2810XIV",
  "array_type": "2810XIV",
  "storage_model": "XIV",
  "physical_capacity": 12011310153728,
  "serial": "21011",
  "capacity_max_pool_size": 12011,
  "capacity_soft_mib": 11454878,
  "capacity_hard_mib": 11454878,
  "capacity_free_soft_mib": 32822,
  "capacity_free_hard_mib": 6859798,
  "capacity_spare_disks": 3,
  "capacity_spare_modules": 1,
  "capacity_target_spare_disks": 3,
  "capacity_target_spare_modules": 1,
  "capacity_limit_percentage": 100
}
```


The upgrade procedure is implemented by issuing the **POST** command with this syntax: `POST /api/v1/arrays/<array-id>/upgrade`. The input parameters include the following mandatory entries: *username*, *password* and *pkg_name*.

Note: Before running the upgrade, verify that:

- The correct installation file is stored in the `/opt/ibm/ibm_spectrum_control/downloads/` directory. The package must be compatible with the valid upgrade path from the current microcode.
 - The credentials to be supplied have sufficient access level to complete the upgrade procedure.
-

For example, to upgrade the storage system 2810-999-dc21011 microcode to version 11.5.0.c, enter `POST /api/v1/arrays/2810-999-dc21011/upgrade` with the required parameters.

```
POST /api/v1/interfaces
{
  "username": "opsadmin",
  "password": "opspasswd",
  "pkg_name": "xiv_ver_11.5.0.c.tgz"
}
```

If the response to the upgrade request is OK, followed by the task ID, as illustrated below. The task ID can be used in a query request to monitor the upgrade progress, see “Query request and response” on page 238.

```
OK
{
  "task id": "04fc6120-60ae-4182-baa9-687d6ae96ffe"
}
```

In addition, you can use the **GET** to retrieve the current state of the upgrade procedure.

For example, to get the current upgrade status for storage system 2810-999-dc21011, enter `GET /api/v1/arrays/2810-999-dc21011/upgrade`. The output is illustrated below.

```
OK 200
{
  "array": "2810-999-dc21011",
  "state": "Upgrade Not Underway",
  "consequence": "New version has not been downloaded yet",
  "package_target_version": ""
}
```

Module operations

RESTful API for IBM Spectrum Connect uses the **GET** and **POST** commands for module operations.

You can run a module query by issuing the **GET** command to get a list of modules that belong to a storage system and their properties, or display the properties of a single module.

For example, `GET /api/v1/modules/1` lists the properties of the module with identification number 1, as illustrated below.

```
OK 200
{
  "id": 1,
```

```

"component_id": "1:Module:7",
"status": "OK",
"type": "g3.0_interface",
"requires_service": "REPLACE",
"service_reason": "HARDWARE_ERROR",
"disk_bay_count": 12,
"fc_port_count": 4,
"ethernet_port_count": 4,
"memory_gb": 15
}

```

You can initiate phase-in or phase-out action on a module by issuing the **POST** command with the following syntax:

- POST /api/v1/modules/<module-id>/phasein
- POST /api/v1/modules/<module-id>/phaseout

For example, POST /api/v1/modules/3/phasein starts a phase-in procedure for the module with identification number 3. The response to the request is OK, followed by the task ID, as illustrated below. The task ID can be used in a query request to monitor the task progress, see “Query request and response” on page 238.

```

OK
{
  "task id": "04fc6120-60ae-4182-baa9-687d6ae96ffe"
}

```

Disk operations

RESTful API for IBM Spectrum Connect uses the **GET** and **POST** commands for disk operations.

You can run a disk query by issuing the **GET** command to get a list of disks that belong to a storage system and their properties, or display the properties of a single disk.

For example, GET /api/v1/disks/1 lists the properties of the disk with identification number 1, as illustrated below.

```

OK 200
{
  "id": "1",
  "array": "pur15m1",
  "name": "1:Disk:1:2",
  "status": "OK",
  "capacity": "2TB",
  "vendor": "IBM",
  "model": "ST32000444SS",
  "size": "1878633",
  "serial": "9WM1YM3M",
  "requires_service": "REPLACE",
  "service_reason": "HARDWARE_ERROR",
  "temperature": 19,
  "encryption": "Not supported",
  "controller": "SAS"
}

```

You can initiate phase-in or phase-out action on a disk by issuing the **POST** command with the following syntax:

- POST /api/v1/disks/<disk-id>/phasein
- POST /api/v1/disks/<disk-id>/phaseout

For example, `POST /api/v1/disks/82/phasein` starts a phase-in procedure for the disk with identification number 82. The response to the request is OK, followed by the task ID, as illustrated below. The task ID can be used in a query request to monitor the task progress, see “Query request and response” on page 238.

```
OK
{
  "task id": "04fc6120-60ae-4182-baa9-687d6ae96ffe"
}
```

Interface operations

RESTful API for IBM Spectrum Connect uses the **GET**, **PATCH**, **POST** and **DELETE** commands for interface operations.

You can run an interface query by issuing the **GET** command to get a list of IP interfaces that belong to a storage system and their properties, or display the properties of a single interface. The command output can be filtered by module ID or interface type (management, iSCSI, etc.).

For example, `GET /api/v1/interfaces/11` lists the properties of the interface with identification number 11, as illustrated below.

```
OK 200
{
  "name": "management",
  "array": "pu16",
  "module": "1:Module:1",
  "type": "Management",
  "address": "9.151.156.3",
  "netmask": "255.255.248.0",
  "gateway": "9.151.159.254",
  "address6": "",
  "gateway6": "",
  "mtu": "1500",
  "ports": "",
  "id": 11
}
```

You can initiate a partial update action on an interface by issuing the **PATCH** command with the following syntax: `PATCH /api/v1/interfaces/<interface-id>`. The input parameters include the following entries: **name**, **mtu**, **netmask**, **address**, **gateway**.

For example, the following request updates the required parameters for interface 25.

```
PATCH /api/v1/interfaces/25
{
  "name": "interface_name",
  "address": "9.151.151.3"
}
```

The response to the request is as follows:

```
OK 200
{
  "name": "interface_name",
  "array": "pu16",
  "module": "1:Module:1",
  "type": "iSCSI",
  "address": "9.151.151.3",
  "netmask": "255.255.255.248",
  "gateway": "9.151.151.222",
}
```

```

"address6": "",
"gateway6": "",
"mtu": "1500",
"ports": "1",
"id": 25
}

```

You can create a new iSCSI interface entry by issuing the **POST** command with the following syntax: `POST /api/v1/interfaces`. The input parameters include the following entries:

- Mandatory – *name, address, array, netmask, module, ports*.
- Optional – *mtu, gateway*.

For example, the following request creates a new interface with the required parameters.

```

POST /api/v1/interfaces
{
  "name": "iSCSI_1_1",
  "array": "pu16",
  "module": "1:Module:1",
  "address": "9.151.151.151",
  "netmask": "255.255.255.248",
  "gateway": "9.151.151.222",
  "mtu": "1500",
  "ports": "1"
}

```

The response to the request is as follows:

```

OK 200
{
  "name": "iSCSI_1_1",
  "array": "pu16",
  "module": "1:Module:1",
  "type": "iSCSI",
  "address": "9.151.151.151",
  "netmask": "255.255.255.248",
  "gateway": "9.151.151.222",
  "address6": "",
  "gateway6": "",
  "mtu": "1500",
  "ports": "1",
  "id": 25
}

```

You can delete an iSCSI interface by issuing the **DELETE** command with the following syntax: `DELETE /api/v1/interfaces/<interface-id>`.

For example, `DELETE /api/v1/interfaces/1` deletes the interface with identification number 1.

Port operations

RESTful API for IBM Spectrum Connect uses the **GET** command for port operations.

You can run an interface query by issuing the **GET** command to get a list of IP ports that belong to a storage system and their properties, or display the properties of a single port. The command output can be filtered by module ID or interface type (management, iSCSI, etc.).

For example, GET /api/v1/ports/1 lists the properties of the port with identification number 1, as illustrated below.

```
OK 200
{
  "index": 1,
  "array": "pur15m1",
  "role": "iSCSI",
  "ip_interface": "iSCSI_1_1",
  "module" : "1:Module:1"
}
```

Emergency shutdown

RESTful API for IBM Spectrum Connect uses the **POST** command for initiating an emergency storage system shutdown.

You can shut down a storage system in emergency mode by issuing the **POST** command, followed by the arrays argument, the storage system ID and the shutdown argument.

For example, POST /api/v1/arrays/1/shutdown shuts down the storage system with identification number 1, issuing the OK response.

Notices

These legal notices pertain to the information in this IBM Storage product documentation.

This information was developed for products and services offered in the US. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
USA*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
USA*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of the International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Copyright and trademark information website (www.ibm.com/legal/us/en/copytrade.shtml).

VMware, the VMware logo, ESX, ESXi, vSphere, vCenter, and vCloud are trademarks or registered trademarks of VMware Corporation in the United States, other countries, or both.

Microsoft, Windows Server, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.



Printed in USA

SC27-5999-24

