

IBM Spectrum Family

IBM Spectrum Control Standard Edition

Karen Orlando

Lloyd Dean

Tiberiu Hajas

Marion Hejny

Johanna Hislop

Ruben Moreno

Hope Rodriguez



 **Analytics**

Storage



International Technical Support Organization

IBM Spectrum Family: IBM Spectrum Control Standard Edition

March 2016

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (March 2016)

This edition applies to IBM Spectrum Control Version 5.2.8.

© Copyright International Business Machines Corporation 2016. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
IBM Redbooks promotions	ix
Preface	xi
Authors	xii
Now you can become a published author, too!	xiv
Comments welcome	xiv
Stay connected to IBM Redbooks	xiv
Chapter 1. Introducing IBM Spectrum Control	1
1.1 IBM Spectrum Family	2
1.2 Software Defined Storage	2
1.3 Overview of IBM Spectrum Storage Family	3
1.3.1 Software Defined Storage control plane	3
1.3.2 Software Defined Storage data plane	4
1.4 Spectrum Control components	4
1.5 Spectrum Control architecture	6
1.6 Spectrum Control products	6
1.6.1 Licensing	7
1.6.2 IBM Virtual Storage Center and Spectrum Control unbundling	7
1.6.3 Spectrum Control Standard Edition features	7
1.6.4 Spectrum Control Advanced Edition comparison to the Standard Edition	11
1.7 IBM Copy Services Manager	13
1.8 Web-based GUI enhancements	14
1.8.1 Identifying hot spots in the web-based GUI	15
1.8.2 Spectrum Control Cinder driver	16
1.8.3 Spectrum Control agile development	17
1.9 VMware	17
1.10 Cognos reporting	17
1.11 Service Management Connect	18
Chapter 2. Installation and upgrade checklist and considerations	19
2.1 Preinstallation planning	20
2.1.1 Planning checklist	20
2.1.2 Installation image considerations	23
2.2 Installation outline	24
2.2.1 Installation checklist	24
2.2.2 Installation	29
2.2.3 Postinstallation tasks	40
2.3 Upgrade summary	41
2.3.1 Planning considerations	41
2.3.2 Pre-upgrade steps	43
2.3.3 Upgrade steps	43
2.3.4 Post-upgrade tasks	44
2.4 Installing the Reporting component later	44
Chapter 3. Setting up authentication	47

3.1 Planning for authentication	48
3.1.1 Authentication and authorization.	49
3.1.2 Assigned roles	49
3.1.3 Single sign-on	50
3.2 Implementing authentication	50
3.2.1 Implementing Local OS authentication and an internal user ID	50
3.2.2 Implementing external authentication (direct LDAP) and a local internal user ID	52
3.2.3 Implementing external authentication (implicit) and a local internal user ID	53
3.2.4 Implementing external authentication and a domain internal user ID	54
3.3 Use case: Configuring Spectrum Control to authenticate with Bluepages for an intranet user	55
3.3.1 Adding a Bluepages SSL certificate	57
3.3.2 Configuring the federated repository.	61
3.3.3 Configuring web security.	66
3.3.4 Performing a test communication with Bluepages	67
3.3.5 Adding a specific Bluegroup to Spectrum Control role-based authorization.	67
Chapter 4. Configuration and administration considerations, tasks, and use cases.	71
4.1 Adding resources to Spectrum Control	72
4.1.1 Adding IBM storage systems	72
4.1.2 Adding storage systems from other vendors.	76
4.1.3 Adding fabrics and switches	76
4.1.4 Next steps	79
4.1.5 Updating and testing a data source	80
4.2 Adding servers	81
4.2.1 Choosing an agentless server versus a Storage Resource agent	81
4.2.2 Adding an agentless server	82
4.2.3 Deploying a Storage Resource agent	83
4.3 Storage Resource agent use cases	84
4.3.1 Updating agent properties.	84
4.3.2 Changing the IP address or DNS name for Storage Resource agents.	85
4.4 Additional configuration tasks	86
4.4.1 Installing additional applications	86
4.4.2 Backing up Spectrum Control	87
4.4.3 User authentication	87
4.4.4 Tivoli Common Reporting/Cognos Reporting single sign-on	88
4.4.5 User management in Spectrum Control	88
4.4.6 Preferred practices configuration tips	89
4.5 Dashboard administration use cases	92
4.5.1 Viewing the quick status on the dashboard	92
4.5.2 Handling alerts: Critical and warning.	93
4.5.3 Handling device condition: Warnings and errors.	95
4.5.4 Clearing unreachable switches	99
4.5.5 Handling performance monitors	100
4.6 Fabric administration use cases	103
4.6.1 Viewing inter-switch connections	103
4.6.2 Viewing NPV switch ports and internal routes	110
Chapter 5. Service-oriented storage management using IBM Spectrum Control	113
5.1 Getting insights into how your storage is consumed	114
5.1.1 How to mitigate the risk of your pool running out of space	115
5.1.2 Monitoring your Easy Tier pools	117
5.1.3 How to mitigate the risk of overprovisioning	121

5.1.4	Creating input for your chargeback report.	126
5.1.5	Rollup reports across different Spectrum Control servers	137
5.2	Organizing capacity data around applications and business units	145
5.3	Getting an overview of the resources that are associated with your applications. . . .	164
5.4	Drilling into resource details for problem determination and performance analysis . .	165
Chapter 6.	Advanced alert management.	177
6.1	Alerting overview.	178
6.2	Alerts.	178
6.2.1	Features	179
6.2.2	Triggering conditions for alerts	179
6.2.3	Reporting	180
6.3	Alerting and event processing.	180
6.3.1	Scenario 1: Monitoring and being notified about storage utilization	182
6.3.2	Scenario 2: Performance troubleshooting with advanced alert management. . .	187
6.3.3	Scenario 3: Closely monitoring new application usage with advanced alert management.	191
Chapter 7.	Managing IBM Spectrum Scale with IBM Spectrum Control	195
7.1	Historical overview of Spectrum Scale support in Spectrum Control	196
7.2	Troubleshooting scenarios	198
7.2.1	Adding Spectrum Scale files and objects to Spectrum Control.	198
7.2.2	Usage scenarios for troubleshooting Spectrum Scale	201
7.3	Advantages of using the Spectrum Control web-based GUI to manage Spectrum Scale clusters.	226
7.4	Advantages of using the Spectrum Scale GUI to manage Spectrum Scale	227
Chapter 8.	IBM Spectrum Control Advanced Edition overview	229
8.1	Introduction	230
8.2	Cloud configuration	231
8.2.1	Tiering.	231
8.2.2	Service classes.	232
8.2.3	Capacity pools	233
8.3	Storage provisioning	234
8.4	Optimization	235
8.4.1	Balancing pools.	235
8.4.2	Analyzing tiering	235
8.4.3	Volume transformation	236
8.5	IBM Spectrum Protect Snapshot.	237
Appendix A.	Managing security vulnerabilities	239
	IBM Product Security Response Team overview.	240
	Communicating about security vulnerabilities	240
	Finding the security bulletins	240
	IBM Product Security Response Team blog.	240
	IBM Support portal	241
	My Notifications.	241
	Resolving security vulnerabilities for Spectrum Control.	241
	Vulnerabilities for other products	241
	Vulnerabilities with the operating systems.	242
	Reporting a security vulnerability	242
Appendix B.	Reporting scripts	243
	Overview	244

Sample script with native SQL statement	244
Sample script to run as a batch report	248
runshowbackexport.bat	249
showbackexport.sql	249
Sample script for Cognos	253
Related publications	265
IBM Redbooks	265
Other publications	265
Online resources	265
Help from IBM	266

Notices

This information was developed for products and services offered in the US. This material might 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, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM®	Linear Tape File System™
Cognos®	IBM FlashSystem®	PartnerWorld®
DB2®	IBM SmartCloud®	Passport Advantage®
developerWorks®	IBM Spectrum™	Redbooks®
DS4000®	IBM Spectrum Accelerate™	Redbooks (logo)  ®
DS5000™	IBM Spectrum Archive™	Storwize®
DS6000™	IBM Spectrum Control™	System Storage®
DS8000®	IBM Spectrum Protect™	Tivoli®
Easy Tier®	IBM Spectrum Scale™	WebSphere®
FlashCopy®	IBM Spectrum Storage™	XIV®
FlashSystem™	IBM Spectrum Virtualize™	
GPFS™	Jazz™	

The following terms are trademarks of other companies:

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

LTO, Ultrium, the LTO Logo and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get personalized notifications of new content
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



Download
Now

Android



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks

About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

IBM® Spectrum Control (Spectrum Control), a member of the IBM Spectrum™ Family of products, is the next-generation data management solution for software-defined environments (SDEs). With support for block, file, object workloads, and software-defined storage and predictive analytics, and automated and advanced monitoring to identify proactively storage performance problems, Spectrum Control enables administrators to provide efficient management for heterogeneous storage environments.

IBM Spectrum Control™ (formerly IBM Tivoli® Storage Productivity Center) delivers a complete set of functions to manage IBM Spectrum Virtualize™, IBM Spectrum Accelerate™, and IBM Spectrum Scale™ storage infrastructures, and traditional IBM and select third-party storage hardware systems.

This IBM Redbooks® publication provides practical examples and use cases that can be deployed with *IBM Spectrum Control Standard Edition*, with an overview of *IBM Spectrum Control Advanced Edition*. This book complements the Spectrum Control IBM Knowledge Center, which is referenced for product details, and for installation and implementation details throughout this book. You can find this resource as the following website:

<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>

Also provided are descriptions and an architectural overview of the IBM Spectrum Family, highlighting Spectrum Control, as integrated into software-defined storage environments.


The following features and topics are demonstrated and described:

- ▶ Web-based GUI enhancements: The Spectrum Control next-generation, management interface offers quick access to status and performance information, and intuitive dashboards that offer at-a-glance visibility into critical infrastructure assets.
- ▶ Installation and upgrade guidelines and tips: Help guide you through the process of installing and upgrading Spectrum Control.
- ▶ Storage management and monitoring: Practical troubleshooting scenarios that use the Spectrum Control web-based GUI. Included are IBM traditional storage systems, all of the latest support for Spectrum Accelerate and Spectrum Scale storage infrastructures, and ports, switches and performance metrics, SAN fabric, and other devices.
- ▶ Performance troubleshooting: Provides a convenient performance monitor and health view, log of alerts of violated thresholds with charts, and simplified views of performance information for all resources.
- ▶ Storage automation: Includes provisioning and practical scenarios that use the extensive set of new alert conditions to better monitor and automate your environment based on identification of critical situations.
- ▶ Advanced analytics overview: Describes cloud configuration, tiered storage optimization, and provisioning.
- ▶ Volume performance reports: Provides server, unmanaged server, hypervisor, and performance data exports to multi-tab Microsoft Excel spreadsheets, including IBM Cognos® based reporting examples and scripts that help create and integrate custom reports on capacity, performance, and utilization.
- ▶ External authentication services: Scenarios illustrate how to use Spectrum Control with external authentication services with Lightweight Directory Access Protocol (LDAP) or Active Directory (AD).

This publication is intended for storage administrators, clients who are responsible for maintaining IT and business infrastructures, and anyone who wants to learn more about employing Spectrum Control and Spectrum Control Standard Edition.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.



Karen Orlando is a Project Leader at the International Technical Support Organization, Tucson Arizona Center. Karen has over 25 years in the IT industry with extensive experience in open systems management, and testing and development of IBM hardware and software for storage. She holds a degree in business information systems from the University of Phoenix and is a Project Management Professional (PMP), certified since 2005.



Lloyd Dean is an IBM Senior Certified Executive IT Architect in Infrastructure Architecture. During Lloyd's 14 plus years at IBM, he has held numerous leadership roles. Within the Sales & Distribution division, Lloyd leads virtual storage solution efforts across the US. In this role, Lloyd leads Storage Solution Architect efforts in storage optimization, business continuance, and cloud enablement. Lloyd maintains Subject Matter Expert (SME) status on the following IBM products: IBM System Storage® SAN Volume Controller, IBM System Storage Storwize® Controllers, Spectrum Control, Tivoli Storage Productivity Storage Center, and IBM Virtual Storage Center. Additionally, Lloyd leads a team by hosting and leading multiple IBM communities focused on knowledge transfer and skill enhancement for IBM storage client technical specialists in IBM virtual storage products.



Tiberiu Hajas is an Advisory Software Engineer from the Almaden Research Center, California, who has been with IBM working in storage technology for 10 years, and in Spectrum Control, Tivoli Storage Productivity Center, and Virtual Storage Center for the past 8 years. He holds a degree in computer science from West University of Timisoara, Romania. Tiberiu has worked and is certified with Tivoli Storage Productivity Center since Version 3.3, but his areas of expertise also include SAN Volume Controller, Storwize V7000, IBM DS8000®, IBM DS6000™, IBM XIV®, N series, EMC, Hitachi, QLogic, LSI, Brocade, and Cisco products.



Marion Hejny joined IBM in 1986 and has worked in many different areas. She provides services around Spectrum Control, including installation, upgrading, and configuration, and creates customized reports for customers in Europe, the Middle East, and Africa. She participates in the Spectrum Control Beta Program and in the Early Adopt Program (EAP) calls. Marion holds a degree in chemical engineering from the University of Applied Science in Darmstadt, Germany, and is a Project Management Professional (PMP), certified since 2008.



Johanna Hislop joined IBM in 1997 and is the team lead for the Spectrum Control Level 3 support team. Her areas of expertise include Spectrum Control, Fibre Channel switches, and working with customers and vendors for both. Johanna has written extensively on the Tivoli Storage Productivity support website, and has several articles and a blog on the IBM developerWorks® website. She has a BS in Computer Science degree from the University of Minnesota, Duluth.



Ruben Moreno joined IBM in 2000 and he has worked on many different software products across IBM Systems Storage. He works on Spectrum Control and has worked on all previous versions of the product, such as Virtual Storage Center and Tivoli Storage Productivity Center. Ruben holds a degree in management information systems from the University of Arizona.



Hope Rodriguez joined IBM in 2008 after studying at New Mexico State University. She holds a degree in information systems. Hope is a Software Test Engineer for Spectrum Control, Tivoli Storage Productivity Center / System Storage Productivity Center in Tucson, AZ. She is also an IBM Certified Deployment Professional on Spectrum Control and Tivoli Storage Productivity Center and has experience with a wide array of IBM storage products.

Thanks to the following people for their contributions to this project:

Larry Coyne

International Technical Support Organization

Daniel Bajema, Christian Bolik, Tom Clark, Peter Flaemig, Stefan Jaquet, Thuan Nguyen, Irina Panaitescu, Dominic Pruitt, Ben Randall, Louise Schillig, Markus Standau, Falk Schneider, Anu Talasu, Scott Trevett, David Warnke, Kevin Webster

IBM Systems Storage

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- Send your comments in an email to:

redbooks@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>



Introducing IBM Spectrum Control

IBM Spectrum Control (Spectrum Control) is integrated data and storage management software that provides monitoring, automation, and analytics for organizations with multiple storage systems. It can reduce storage costs by improving consistency and control for tasks such as storage provisioning, performance management, and copy services. Spectrum Control solutions can be deployed on premises or from the IBM cloud, with a range of license plan options.

This solution delivers a complete set of advanced functions to manage IBM Spectrum Virtualize, IBM Spectrum Accelerate, and IBM Spectrum Scale storage infrastructures, and traditional IBM and select third-party storage hardware systems.

IBM Spectrum Control V5.2.8 is the first version of this product and is based on IBM Tivoli Storage Productivity Center V5.2.7. Spectrum Control, as documented in *IBM Software-Defined Storage Guide*, REDP-5121, is a member of the IBM Spectrum Storage™ Family and Spectrum Storage Family of products. IBM Spectrum Storage is a component of the IBM Software Defined Storage (SDS) solution.

This chapter provides a high-level description of the IBM Spectrum Storage Family products, and how they are integrated into software-defined storage in the control and data plane. Also described are IBM Spectrum Control Architecture, licensing, and IBM Spectrum Control Standard and Advanced Edition functions with comparisons of the products.

Additionally, you can use this chapter to understand the enhancements in the Spectrum Control web-based graphical user interface (GUI) compared to the two GUIs that are available in the prior product offering.

Also included in this chapter is a brief update about the enhancements to IBM Copy Services Manager (formally Tivoli Storage Productivity Center for Replication), VMware support, IBM Cognos reporting, and Service Management Connect.

1.1 IBM Spectrum Family

This section provides a description of the IBM Spectrum Family within the SDS architecture *control plane* and *data plane*. This section emphasizes how Spectrum Control integrates with and enhances this technology.

1.2 Software Defined Storage

SDS is one the three main components of the new software-defined infrastructure (SDI), and the focus of this chapter. The other key components that support the SDI framework are Software Defined Compute and Software Defined Networking (SDN). For more information about these components, see *IBM Software-Defined Storage Guide*, REDP-5121.

The main characteristic of the SDS reference architecture, similar to the SDN environment, is the separation of the storage functions into two main layers:

- ▶ **SDS control plane**

The SDS control plane, where Spectrum Control is, is a software layer that manages the virtualized storage resources. It provides all of the high-level functions that are needed by the customer to run the business workload and enable optimized, flexible, scalable, and rapid provisioning storage infrastructure capacity. These capabilities span functions such as storage virtualization, policies automation, analytics and optimization, backup and copy management, security, and integration with the API services, including other cloud provider services.

- ▶ **SDS data plane**


The data plane encompasses the infrastructure where data is processed. It consists of all basic storage management functions, such as virtualization, RAID protection, tiering, copy services (remote, local, synchronous, asynchronous, and point-in-time), encryption, and data deduplication that can be started and managed by the control plane. The data plane is the interface to the hardware infrastructure where the data is stored. It provides a complete range of data access possibilities, spanning traditional access methods such as block I/O (for example, iSCSI) or File I/O (POSIX compliant), to object-storage or Hadoop Distributed File System (HDFS).

Figure 1-1 on page 3 shows the SDS integration with a mapping of the Spectrum Storage Family of products across the SDS control plane and data plane. The components that comprise Spectrum Control are highlighted.

Storage Control Integration

Storage and Data Control

Control Plane



Data Plane

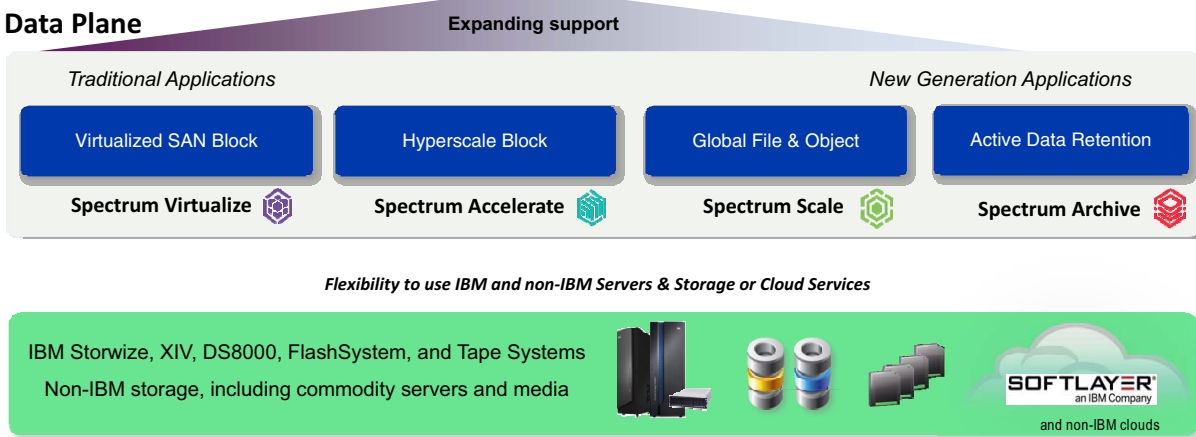


Figure 1-1 Spectrum Control Family mapped to Software Defined Storage control plane and data plane

1.3 Overview of IBM Spectrum Storage Family

This section provides a high-level overview of the IBM Spectrum Storage Family. Descriptions of the products and the functions are provided and identify the separation of the storage functions across the control and data planes.

1.3.1 Software Defined Storage control plane

This section provides an overview of the IBM software product offerings that provide the building blocks for the SDS control plane:

- ▶ **Spectrum Control:** Provides analytics-driven storage management and efficient infrastructure management for virtualized, cloud, and software-defined storage to simplify and automate storage provisioning, capacity management, availability monitoring, and reporting.
- ▶ **IBM Spectrum Protect™:** Enables reliable, efficient data protection and resiliency for software-defined, virtual, physical, and cloud environments, and advanced integrated application snapshot backup and restore capabilities.

1.3.2 Software Defined Storage data plane

This section overviews the IBM software product offerings that provide the building blocks for the SDS data plane:

- ▶ IBM Spectrum Accelerate: Offers grid-scale block storage with rapid deployment that expedites delivery of data across an enterprise and adds flexibility to cloud deployments.
- ▶ IBM Spectrum Virtualize: Software is at the heart of IBM SAN Volume Controller and the IBM Storwize family. It enables these systems to deliver industry-leading virtualization that enhances storage to improve resource utilization and productivity, and streamlines deployment for simpler, more responsive, scalable, and cost-efficient IT infrastructures.
- ▶ IBM Spectrum Scale: Flash-accelerated, industrial strength, highly scalable software-defined storage that enables global shared access to data with scalability and agility for cloud and analytics.
- ▶ IBM Spectrum Archive™: Enables the use of IBM Linear Tape File System™ (LTFS) enabled IBM LTO Ultrium and Enterprise tape drives to consolidate tape storage into low-cost storage pools. These LTFS storage pools, when presented to the Spectrum Scale solution, can then be managed by Spectrum Scale to move automatically infrequently accessed data from disk to tape to lower costs while retaining ease of use without the need for any proprietary tape applications.

1.4 Spectrum Control components

Here are the components of Spectrum Control:

- ▶ Spectrum Control Standard Edition: With this edition, you can manage storage capacity, availability, events, performance, and resources. Spectrum Control Standard, which is the key focus in subsequent chapters, reduces the complexity of managing a storage environment by centralizing, simplifying, and optimizing storage tasks that are associated with storage systems, storage networks, replication services, and capacity management.

For more information, see the following website:

<http://www.ibm.com/software/tivoli/csi/cloud-storage/>

For more information about how to get started quickly with a typical installation of Spectrum Control Standard Edition, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/1K8TUS0>

- ▶ Spectrum Control Advanced Edition: With this edition, you can employ all of the functions of Spectrum Control Standard Edition. In addition, it addresses new workloads that require massive scale and rapid pace, and accelerates business insight by adding advanced analytics functions, such as storage optimization, provisioning, and transformation.

For more information, see the following website:

<http://www.ibm.com/software/products/en/spectrum-control-advanced>

For more information about how to get started quickly with a typical installation of Spectrum Control Advanced Edition, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/1S4S5Ne>

- ▶ IBM Spectrum Control Storage Insights: This component optimizes storage environments with analytics-driven insights that are delivered in the cloud. With the cloud-based delivery model, Spectrum Control Storage Insights can be installed in less than 30 minutes, which avoids the upfront cost and complexity of on-premises solutions. You quickly benefit from the following functions:

- Accurately identify and categorize storage assets
- Monitor capacity and performance in context, including application and department-level views
- Reclaim unused storage
- Optimize data placement based on historical usage patterns

For more information about Spectrum Control Storage Insights, see the following websites:

- <http://www.ibm.com/systems/storage/spectrum/insights>
- <http://www.ibm.com/marketplace/cloud/analytics-driven-data-management/us/en-us>

- ▶ IBM Spectrum Control Base Edition: This component 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.

It provides a single-server back-end location and enables centralized management of IBM storage resources for different virtualization and cloud platforms:

- VMware vCenter Server
- VMware vSphere Web Client (vWC)
- VMware vSphere APIs for Storage Awareness (VASA)
- VMware vCenter Operations Manager (vCOPs)
- VMware vCenter Orchestrator (vCO)

For more information, see to the Spectrum Control Base Edition IBM Knowledge Center at:

<http://ibm.co/10mhKfy>

Note: For more information about the full Spectrum Storage Family, and details regarding Spectrum Control Storage Insights, which is the cloud *software as a service* (SaaS) offering for Spectrum Control, see the following publications:

- ▶ *IBM Software-Defined Storage Guide*, REDP-5121:
<http://www.redbooks.ibm.com/redpieces/abstracts/redp5121.html>
- ▶ *Regain control of your environment with IBM Storage Insights*, REDP-5231:
<http://www.redbooks.ibm.com/abstracts/redp5231.html>

1.5 Spectrum Control architecture

The Spectrum Control architecture is enhanced and includes support for new storage devices from IBM.

The Spectrum Control architecture, which is shown in Figure 1-2, illustrates Spectrum Control and its integration with the Spectrum Storage Family of products within the SDS landscape. Spectrum Control now supports the Spectrum Accelerate and Spectrum Scale solutions, along with all of the storage, SAN, and servers that are supported in Tivoli Storage Productivity Center V5.2.7.

For more information about specific supported products, versions, and detailed guidance regarding how to add, configure, and manage the products that are supported by Spectrum Control, see the Spectrum Control IBM Knowledge Center:

<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>

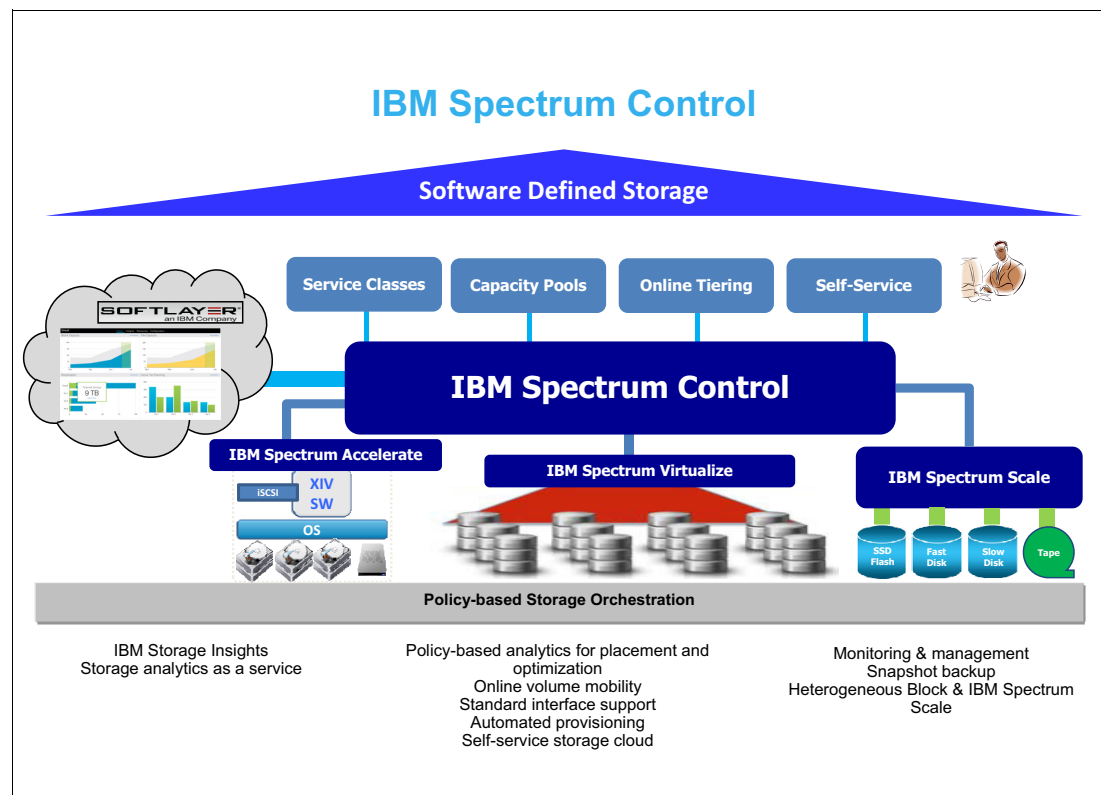


Figure 1-2 Spectrum Control architecture

1.6 Spectrum Control products

This section describes the Spectrum Control product licensing, Spectrum Control Standard Edition features, Spectrum Control Advanced Edition features, and a comparison between the two editions. The purpose of this section is to provide insight into the planning for selecting the correct products to meet your business requirements.

1.6.1 Licensing

Spectrum Control is made available in Standard and Advanced editions, which are priced by the capacity of the storage systems that are managed, and Standard and Advanced Select editions, which are priced by the number of storage enclosures that are managed. To obtain a complete breakdown of roles and features by license, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/1P33zf0>

Note: The Spectrum Control IBM Knowledge Center is enhanced with additional details and depth to provide the Spectrum Control administrator or client with a single source for details regarding installation, upgrades, *how to* examples, and use case scenarios to maintain and use Spectrum Control.

1.6.2 IBM Virtual Storage Center and Spectrum Control unbundling

With Spectrum Control, a change in licensing and bundling occurred regarding the advanced features. In prior releases of Tivoli Storage Productivity Center, to get the required advanced features, you needed to obtain the IBM SmartCloud® Virtual Storage Center (VSC) solution. This solution included the advanced version of Tivoli Storage Productivity Center and was the only way to obtain those advanced functions, such as analytics and fully based storage service class and storage provisioning.

With the announcement of Spectrum Control, IBM also released the Advanced and Advanced Select editions, modified the bundling, and renamed IBM SmartCloud VSC to IBM Virtual Storage Center. Spectrum Control Advanced features can still be obtained in VSC and VSC for Storwize solution packs, but Spectrum Control Advanced Edition features are no longer tied to that solution package.

In addition, Spectrum Control Advanced editions include IBM Spectrum Protect Snapshot (formerly Tivoli FlashCopy® Manager), and Copy Services Manager (formerly Tivoli Storage Productivity Center for Replication). For more information about IBM Spectrum Protect Snapshot, see 8.5, “IBM Spectrum Protect Snapshot” on page 237. For more information about IBM Copy Services Manager, see 1.7, “IBM Copy Services Manager” on page 13.

The key value differentiators between Spectrum Control Advanced editions and VSC is the inclusion of Spectrum Virtualization in to the VSC product, which makes Spectrum Control Advanced Edition a feature-rich solution for many storage shops seeking a full management platform in their storage environment, but not by using external storage virtualization.

1.6.3 Spectrum Control Standard Edition features

Spectrum Control Standard editions and Standard Select editions have numerous features supporting administration. These editions include reporting and monitoring across your block, file, and object storage devices. For specific details by role or by Spectrum Control version type, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/1mQaj9Q>

Note: In the rest of this book, the term *Spectrum Control Standard* refers to both Standard and Standard Select editions of Spectrum Control, and *Spectrum Control Advanced* refers to both Advanced and Advanced Select editions.

A key value of Spectrum Control, available in the Standard and Advanced versions, is the rich reporting capabilities. Reporting for capacity and assets is enabled through the web-based GUI and Cognos external interfaces. This enablement supports reporting by using all data that is available in the common repository. This support includes the most recently collected detail by incorporating all times that are defined by using the *user-configurable historical retention setting*. This user-configurable retention setting supports block, file, and object storage, and SAN fabric and switch details. The specific length of time that the data that is maintained is customer configurable.

For more information about configuring this element, see the Spectrum Control IBM Knowledge Center:

ibm.co/1RwPwFh

Tip: In Spectrum Control Standard, the history retention setting configuration is included in the web-based GUI. This configuration was moved from the stand-alone GUI in Tivoli Storage Productivity Center to the web-based GUI in Spectrum Control, and can be accessed by clicking **Main Menu** → **Settings**(1) → **History Retention** to edit (2), as shown in Figure 1-3.

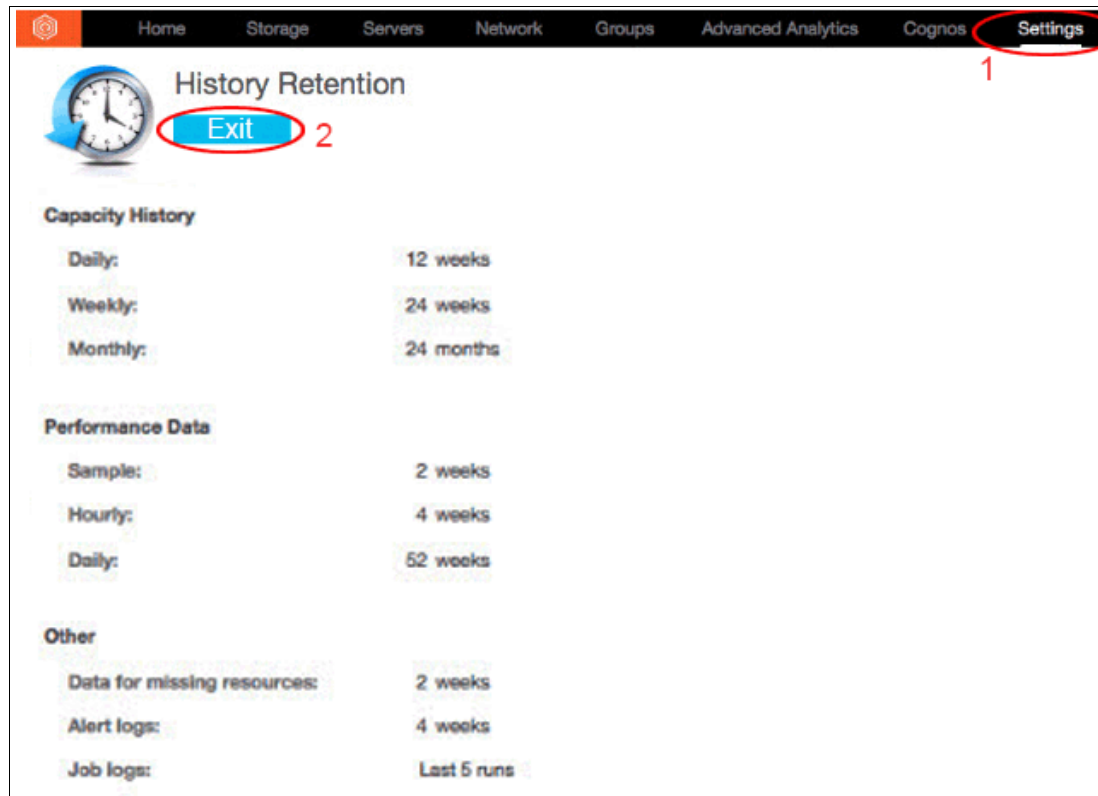


Figure 1-3 Spectrum Control History Retention window

Using the edit option that is shown in Figure 1-3, you can update the default times to support the data availability, and to support reporting or performance debug efforts that are being satisfied through Spectrum Control Standard.

Tip: A preferred practice for the *History Retention* setting for storage performance monitor data is 30 days for Sample level, 180 days for hourly level, and 365 days for daily level. This practice provides the data that supports performance monitoring debug activity.

Storage performance management is enabled for all block storage devices, and for file storage in the IBM Spectrum Scale solution. Capacity and asset reporting covers all three types of storage, making Spectrum Control Standard a critical management tool for heterogeneous block, file, and object storage environments. Storage reporting is enabled through the web-based GUI, command-line interface (CLI), or Cognos.

Reporting through all of these interfaces supports activities such as *screen print* in HTML, PDF, or CSV format, batch reporting by using the predefined capacity, asset, and performance reports in Cognos, or by using the user-defined reporting that is available in Cognos. For usage scenarios and more information about the rich reporting that is available in Spectrum Control Standard, see Chapter 5, “Service-oriented storage management using IBM Spectrum Control” on page 113.

Another key value that is provided in Spectrum Control Standard is alert management. As described in Chapter 6, “Advanced alert management” on page 177, alert management is a major enhancement focus in Spectrum Control Standard. Many elements of the new web-based GUI can now have alerting enabled with conditional capabilities. These conditional capabilities allow alerts to be generated when multiple conditions are identified, based upon your business rules.

This alert management enhancement in Spectrum Control Standard fulfils the user requests for these capabilities. Through the agile development approach, clients can see more enhancements in this area each quarter. For specific enhancements in each release, see the *Spectrum Control Standard Release Notes*, found at:

<http://ibm.co/22Vhz4J>

Spectrum Control Standard provides the means for centralizing access to device-specific management GUIs. This is a time saver for many administrators because it has access to the many storage and SAN devices through a simple click, as shown in Figure 1-4. Clicking the **Open Storage System** web-based GUI control starts the *native element manager* for the device of interest for your review. Look for this option in all storage and SAN switch device windows.

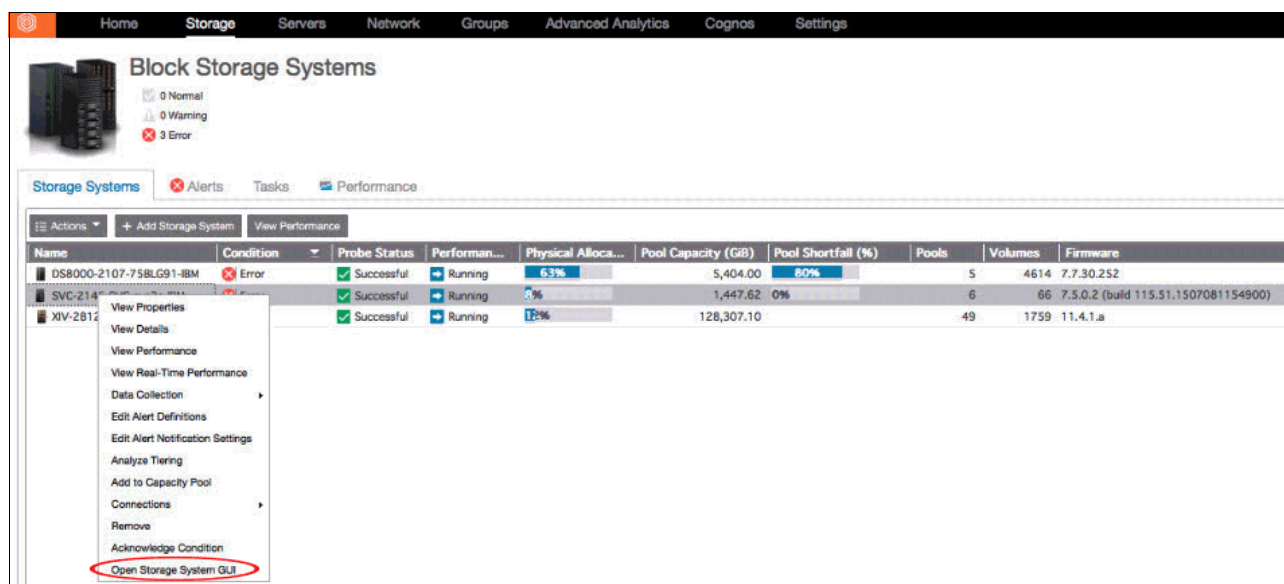


Figure 1-4 Spectrum Control Open Storage System web-based GUI

Combining Spectrum Control Standard's wealth of configuration and available performance details, then using the native element manager to perform actions to make changes, is an acceptable solution for many customers. Clients seeking tight integration with VMWare or OpenStack or seeking to deploy a self-service storage offering might find that Spectrum Control Advanced is the preferred choice.

With Spectrum Control, by using the *roles* feature, you can provide both administrators and users the ability to monitor and to use the rich reporting interfaces. Users that are defined with the *monitor* role have full access to all of the standard reporting available through all three interfaces (the web-based GUI, CLI, or Cognos). This role provides a common reporting solution for all users regarding storage. The roles feature also allows a system or database administrator to have access to reporting for specific applications and departments that are made available by the storage team. This capability can remove much of the demands from a storage team for this type of request.

Two options for reporting, data gathering, and monitoring of your environment are part of Spectrum Control. These options are:

- ▶ Spectrum Control, which is deployed on premises.
- ▶ Spectrum Control Storage Insights has many of the same capabilities of Spectrum Control Standard Edition but is obtained as a cloud-enabled service.

The details of Storage Insights regarding full features, pricing, and terms are available at:

<https://www.ibm.com/marketplace/cloud/analytics-driven-data-management/us/en-us>

The key advantage or decision as to which product or offering to use is directly related to whether your corporation or company can use and deploy cloud-based solutions. From a feature and function capability, Spectrum Control Standard and Storage Insights both offer rich reporting across all of the same storage devices.

Notes:

A key difference between Spectrum Control and Storage Insights is that Storage Insights, as of this writing, does not offer SAN switch device management. If that is critical for your business needs, then Spectrum Control is the clear choice.

Another key differentiator is that Storage Insights does not provide functions to make any updates in your storage system's configuration. It simply collects information and delivers multiple insights to the user that are based on IBM proprietary analytics.

If cloud tools are an option for your environment, and storage device and performance reporting is the business need, then Storage Insights can be the preferred choice. This option enables the same storage monitoring, alerting, and reporting capabilities while removing the management of the tool from the customers because it is part of the service. All updates to the core Spectrum Control solution that are delivered in Storage Insights is performed by the cloud provider. Your only management responsibility is in registering end devices and enabling the one-way and direction secure data collector application in your data center.

For more information about Spectrum Control Storage Insights, see *Regain control of your environment with IBM Storage Insights*, REDP-5231.

1.6.4 Spectrum Control Advanced Edition comparison to the Standard Edition

Section 1.6.3, “Spectrum Control Standard Edition features” on page 7 shows examples of features that are available in Spectrum Control Standard editions. This section provides a high-level view of the Advanced features, which are available only in the Advanced Edition and Advanced Select editions of Spectrum Control. It also provides a comparison of Spectrum Control Advanced versus Standard.

For more information about features that are available in the Advanced version of Spectrum Control, see the Spectrum Control Standard and Advanced features that are described and detailed by feature and user role in the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1mQaj9Q>

The key to understanding which version to use relates directly to what feature and capabilities are required for your corporate enterprise, data center, or remote office needs.

For the smaller remote office or project deployments, an instance of Spectrum Control Standard can provide the remote and local monitoring of storage, servers and if deployed, SAN fabrics by local and corporate storage teams.

For a data center or full corporate enterprise deployments, the key function to review is whether cloud provisioning is being performed, or whether integration with external cloud enabling technologies such as VMWare or OpenStack are deployed or soon to be deployed. In either option, Spectrum Control Advanced might be the clear choice because it supports tight integration with these external hypervisors or full software-defined environment (SDE) management solutions such as OpenStack, which typically feature storage, SAN, and compute provisioning.

Storage provisioning that uses storage service class definitions is key to successful storage operations in these environments. For an overview of these advanced features, such as service classes, see Chapter 8, “IBM Spectrum Control Advanced Edition overview” on page 229. For more information about usage scenarios and explanations, see the following publications:

- ▶ *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204
- ▶ *IBM SmartCloud Virtual Storage Center*, SG24-8239

In reviewing Spectrum Control Advanced features, the key to reaching maximum value is based upon using all of the storage analytics that are enabled in Spectrum Control Advanced. Storage analytics from a Spectrum Control Advanced focus is realized through the awareness of the storage configurations being managed, the knowledge of the workloads being delivered in that configuration, and the wanted action you are attempting to perform. If you are performing any or all of the following actions, you find Spectrum Control Advanced the solution you need to focus on:

- ▶ Balancing storage to mitigate a poor performing storage pool
- ▶ Balancing storage across multiple tiers
- ▶ Defining a service class for storage provisioning with SAN provisioning across an entire capacity pool of storage within a full data center

In the past, Tivoli Storage Productivity Center provided the ability to perform individual volume level storage provisioning. Although this task functioned, it did not integrate well in standardized storage provisioning environments because it was founded on a single storage device focus. As most customers came to this realization, these same customers then used Tivoli Storage Productivity Center in a limited fashion for asset and performance reporting. Although these functions were valuable, they were not achieving the value that was possible.

With Spectrum Control Advanced, you no longer manage storage provisioning at a device volume level. Now, you manage at the pool level and define the attributes that are required for the volumes being provisioned based upon those attributes. To support this design, use storage service classes, and either storage or capacity pools.

Spectrum Control defines a storage pool as a typical client defines it. A pool is defined based upon the storage device capabilities. With Spectrum Virtualize solutions, such as IBM SAN Volume Controller, IBM Storwize V7000, V5000, and V3000, or IBM FlashSystem™ V9000, a pool in Spectrum Virtualize is also equivalent to a pool in Spectrum Control. With IBM XIV, a Spectrum Control pool equals an XIV pool or child pool.

Spectrum Control defined capacity pools include storage pools defined to it. Capacity Pools are an optional configuration in Spectrum Control, and can allow for consistent usage of storage with similar capabilities in your data center.

Using the defined storage or capacity pools, Spectrum Control then analyzes the details that are known about all of the assets that are connected to the request. Then, Spectrum Control runs the request based upon the provisioning request, and the attributes that are defined in the storage service class. This approach provides a scalable and optimal solution that reduces risk because all clients use the same storage service class or set of classes in provisioning storage across the enterprise. Whether the storage requester uses the Spectrum Control web-based GUI, CLI, or an external API to the VMWare vCenter or OpenStack, the method that is employed is based on the storage service class and storage or capacity pool. Finally, the storage service class also includes capability beyond defining the storage attributes. It can also perform the SAN zoning.

In addition to storage provisioning, managing or optimizing storage is an ongoing task in most storage shops. A key storage optimization task that many storage environments typically do not perform is a primary feature of Spectrum Control Advanced Edition. This storage optimization feature supports *horizontal* storage auto-tiering. *Vertical* storage auto-tiering is now a standard element in many enterprise storage products, whether obtained in a storage frame or a storage virtualization solution. Having the ability to move entire storage volumes, not volume extents, between pools (horizontally) of different capabilities (for example, tier 0, all flash memory to hybrid storage or archive storage) and doing so with knowledge of the workload being delivered, is a key advanced function that Spectrum Control Advanced can introduce to your environment.

Beyond the use of storage tiering in Spectrum Control Advanced for performance improvements, it also supports cost management. Each storage tier that is deployed in an environment has financial costs tied to it. Whether that cost is based upon tier of storage (tier 0 versus tier 1, 2, or each tier), or a software feature that has dollars that are assigned to it, using Spectrum Control Advanced provides cost savings to your business enterprise. Employing its horizontal storage tiering, you can increase the utilization of storage from both a performance and a cost perspective. The cost of the tier can be managed by moving block volumes between tiers to balance both performance and cost.

Another valuable aspect of Spectrum Control Advanced is its ability to reduce the time and keystrokes to perform a task. If your goal is to provision multiple storage volumes for a new server, both Spectrum Control Advanced and native element managers provide this multiple volume creation capability. If your goal is to alter the configuration of existing volumes, that is, *transform volumes* in Spectrum Control Advanced, Spectrum Control Advanced is the clear leader. The storage administrator can use Spectrum Control Advanced to select all of the volumes in an application server and transform them, from thick provisioned to thin, thick provisioned to compressed, or migrate the volumes from one pool to another. These are steps that in native element managers take time, and allow risks that are not present when using Spectrum Control Advanced.

These advanced topics typically take considerable time to understand, research, and perform without injecting risk into applications and business units. With Spectrum Control Advanced, customers discover that this optimization and infrastructure cost management capability is based on workload analytics. The risk of implementing these horizontal changes or volume transformations is low with a high return on the result. Using Spectrum Control Advanced, a storage administrator can easily analyze and identify workloads that have cyclic patterns that can benefit by being in a faster tier at peak usage times, and a lowered cost to the business by being in another tier at slower usage times.

Cyclic workloads cross all industries today, whether the industry is in retail, healthcare, and so on. After you identify a cyclic workload, then you must decide whether the storage will be optimized best by implementing hybrid vertical tiering or horizontal full volume tiering. If the latter is required, then that is the scenario that Spectrum Control supports.

Using Spectrum Control, you can define application groups. These application groups of assets can consist of servers, SAN ports, and storage volumes. Then, by using the scheduling feature in Spectrum Control Advanced, you can then move the volumes that are associated with this application group from their storage pools to the flash pool a day before the workload launch. You can then automatically move that same application set back to the original pool when the event window completes.

All of these advanced features are available today, and through your own research and testing, you can build a plan to introduce these functions. Also, through an IBM Lab Services engagement, you can bring in IBM skilled subject matter experts (SMEs) to guide you through the entire process.

1.7 IBM Copy Services Manager

In Spectrum Control Standard and Advanced versions, unlike the Tivoli Storage Productivity Center releases since Version 5.2, which included Tivoli Storage Productivity Center for Replication, Spectrum Control new installations do not include this component. In place of installing Tivoli Storage Productivity Center for Replication as part of the Spectrum Control installation, a license for the Copy Services Manager product is included. This license directory must be copied to the installation directory for IBM Copy Services Manager before installation.

Copy Services Manager is based on Tivoli Storage Productivity Center for Replication with all of the same features and functions, and more. Copy Services Manager is no longer coupled with Spectrum Control, which provides options regarding deployment options and configurations that are needed across the enterprise. Copy Services Manager can be installed on the Spectrum Control application server, but also can be installed on separate physical or virtual machines (VMs).

Per the release notes of Copy Services Manager, several new features are incorporated that directly enhance the Spectrum Virtualize support. The first enhancement supports the capability to restart a stopped Global Mirror or Metro Mirror Replication Consistency Group. In the past, customers relied upon manual restart actions, or developed custom scripts to auto-restart a replication session. Copy Services Manager now includes an option that allows Copy Services Manager to automatically attempt a restart for an error code 1920. This error code indicates that a Global Mirror relationship stopped because of a persistent I/O error.

Tip: For customers planning on moving from Tivoli Storage Productivity Center for Replication to Copy Services Manager, a CLI option supports creating a backup of the operational Tivoli Storage Productivity Center for Replication configuration in a compressed file. This file can then be used by an import process in Copy Services Manager to restore that working configuration on the new Copy Services Manager instance.

For more information about how to obtain the software, and install, configure, and manage your storage replication solutions, see the Copy Services Manager IBM Knowledge Center, found at:

<http://ibm.co/1Sdq2wr>

Note: When performing an upgrade from Tivoli Storage Productivity Center V5.2 to Spectrum Control, the Tivoli Storage Productivity Center for Replication software is maintained. There will be no further enhancements to Tivoli Storage Productivity Center that is included in the Spectrum Control upgrade, and the upgrade does not remove this software component. It is maintained to give you time to review, test, and prepare to install the replacement IBM Copy Services Manager solution.

A final thought regarding upgrades: Instead of removing Tivoli Storage Productivity Center for Replication, you can simply stop the service.

1.8 Web-based GUI enhancements

Spectrum Control is based upon a new and changed dashboard, and is the successor to the web-based GUI in Tivoli Storage Productivity Center V5.2.7. This GUI provides many existing features from the prior product, new features, and support for SDS products, such as the Spectrum Accelerate and Spectrum Scale solutions.

Figure 1-5 on page 15 shows the enhanced web-based GUI in Spectrum Control. All of the resources (Storage, SAN Fabrics, SAN Switches, Hypervisors, Servers, Departments, and Applications) that were available in Tivoli Storage Productivity Center V5.2.7 are present, with Spectrum Scale being the new device, but in a slightly different arrangement.

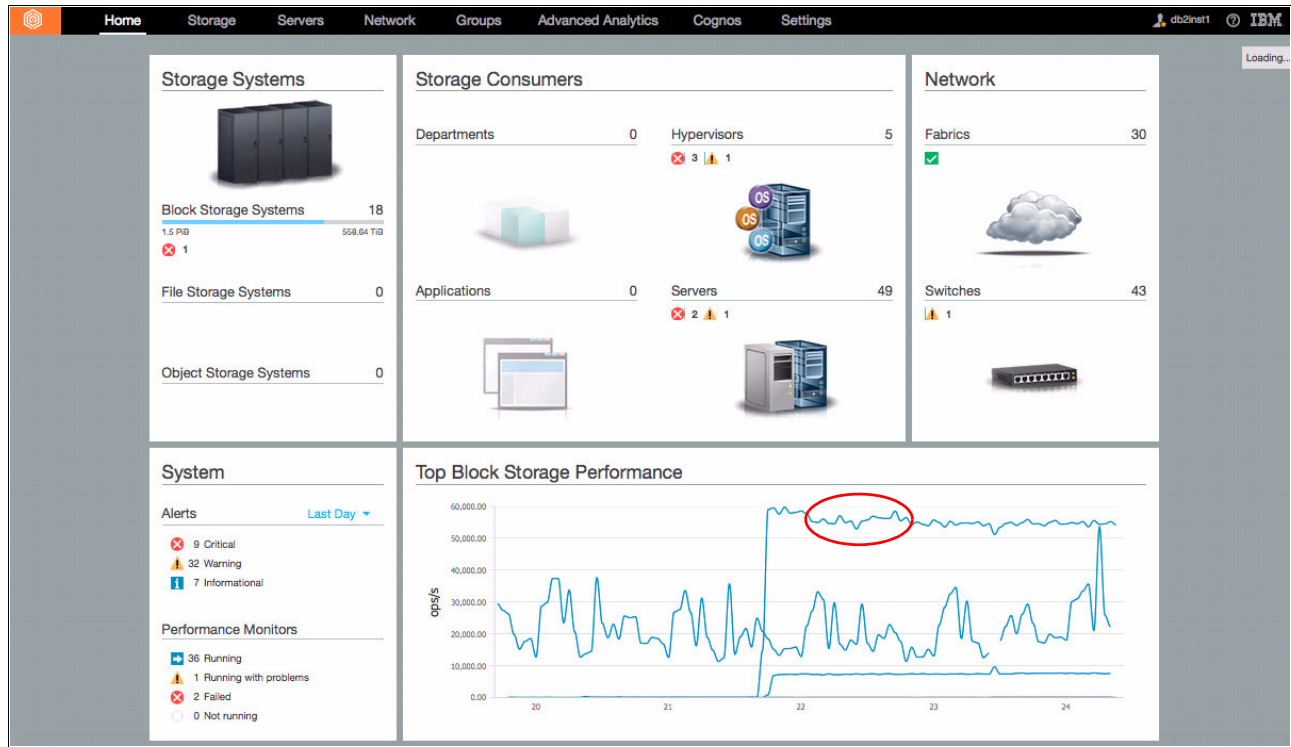


Figure 1-5 New Spectrum Control Dashboard

The dashboard that is shown in Figure 1-5 groups the resources in logical sections for the viewer. Through the new performance pane, the top five most heavily used block storage devices are displayed. Hovering your cursor over the graph items highlights and displays the storage device with specific performance details, as highlighted in the red oval in Figure 1-5.

1.8.1 Identifying hot spots in the web-based GUI

The utilization value for storage pools that is described in *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204 was replaced in Spectrum Control and in later versions of Tivoli Storage Productivity Center when the solution for the Storage Pool Utilization metric was found to be invalid. In Spectrum Control, a performance metric was added to provide a method of comparing similar storage pool configurations, which means that the user can determine whether one pool compared to another is more heavily used and would have more performance potential remaining.

This feature, which shows the Pool Activity and Maximum Pool Activity Score in the tabular section of the performance pane, is highlighted in Figure 1-6. As documented in the IBM Knowledge Center, either the Pool Activity Score or the Max Pool Activity Score reflects the ratio of IOPS / GiB. The Activity Scores are calculated based on the performance data that was collected for the pool on the previous day.

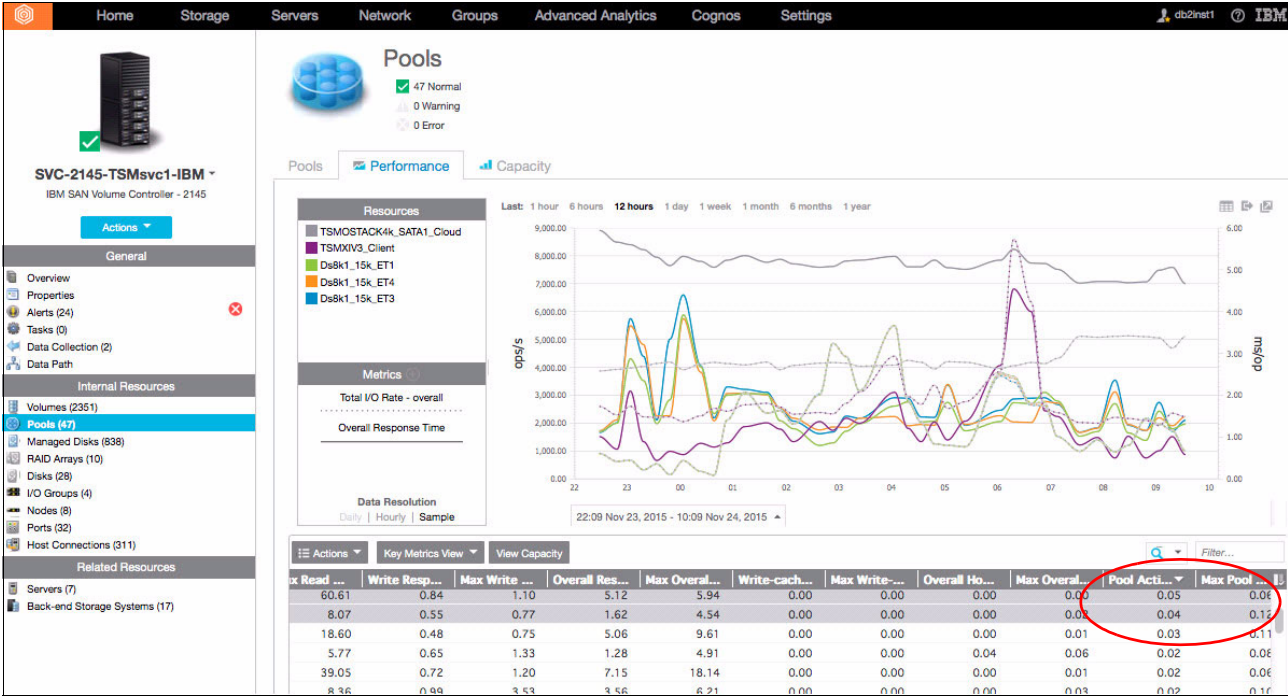


Figure 1-6 Example of Pool Activity and Maximum Pool Activity Score

By scanning the utilization values for pools on a storage virtualizer, such as SAN Volume Controller, you can identify the hot spots or the pools that have high utilization values on the storage virtualizer. You also can determine the course of action to take. For example, if some of the pools on a tier have high utilization values and some have low utilization values, you can use the *Balance Pools* wizard to redistribute volumes across other pools that have the same storage structure. However, if all of the pools on a tier have high utilization values, you can use the *Analyze Tiering* wizard to balance the pools by retiering the most active volumes.

1.8.2 Spectrum Control Cinder driver

If you run Spectrum Control Advanced Edition or Spectrum Control Advanced Select Edition, you can use the Spectrum Control Cinder driver to allow OpenStack environments to manage block storage through the storage service classes that you can define and manage. Each storage service class can have specific users that are authorized to use it so that the storage administrator can develop storage service classes for specific groups of administrators, for example, OpenStack administrators. So, they can safely share your storage environment by using standard tools and storage classes, and either shared pools that are dedicated for them or mixed pools that are available to all teams by using storage service classes that you define and manage. Using these classes limits the risk of this self-service feature.

For more information about the Cinder driver, see the following website:

<http://ibm.co/10XaFrC>

1.8.3 Spectrum Control agile development

Spectrum Control is developed by using an agile development focus. This development approach has been used for nearly two years, and altered the direction for the enhancement of many IBM products. With Spectrum Control, product updates potentially can be delivered as early as every three months, which include all enhancement and defects that were resolved since the last product update.

Sections 1.8.1, “Identifying hot spots in the web-based GUI” on page 15 and 1.8.2, “Spectrum Control Cinder driver” on page 16 reviewed enhancements to Spectrum Control, which highlight a realized potential for this agile development process. To highlight this improvement, Tivoli Storage Productivity Center, for example, did not support SNMP V3 with full SNMP Community user ID and password support. In the matter of two updates, by using agile development, this important function was added, and is now included in Spectrum Control.

The agile approach led to other enhancements, such as the Cinder Driver enhancement that is described in 1.8.2, “Spectrum Control Cinder driver” on page 16, and the fix to the pool utilization, which is described in 1.8.1, “Identifying hot spots in the web-based GUI” on page 15, which was developed and upgraded with the Pool Activity and Maximum Pool Activity Scores.

1.9 VMware

Spectrum Control has the same capabilities as Tivoli Storage Productivity Center V5.2 for supporting VMWare vCenter, and is fully documented in the Spectrum Control IBM Knowledge Center. For more information about configuring the vSphere Web Client extension or the VASA provider, see the IBM Knowledge Center at the following websites:

- ▶ Configuring the VASA provider:
<http://ibm.co/1VZuZYK>
- ▶ Configuring the vSphere Web Client extension:
<http://ibm.co/1UJ5Vnq>

For more information, see *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

1.10 Cognos reporting

Spectrum Control maintains the same capabilities as Tivoli Storage Productivity Center V5.2 with Cognos reporting and is fully documented in the Spectrum Control IBM Knowledge Center. For more information about Cognos reporting, see the Scenarios section of the IBM Knowledge Center, which includes preferred practices regarding use cases with reporting:

<http://ibm.co/1Sdqw5z>

This publication later reviews service management and how Spectrum Control can provide predefined reports to satisfy many service management needs, along with methods to develop reports. For more information, see Chapter 5, “Service-oriented storage management using IBM Spectrum Control” on page 113, and for scripting examples, see Appendix B, “Reporting scripts” on page 243.

These new reports use the same data elements that are available in Spectrum Control, but satisfy your reporting needs. These reports can be made available for *ad hoc* or batch use cases, and then distributed as required. Similar to versions of Tivoli Storage Productivity Center V5.2, Spectrum Control continues to include Cognos. Cognos includes a set of rich features for report creation, including the capability to combine multiple Spectrum Control data stores to support rollup reporting. This report option is not limited to asset-only details, such as in Tivoli Storage Productivity Center Rollup reports. For more information about the rollup report scenario, see 5.1.5, “Rollup reports across different Spectrum Control servers” on page 137.

1.11 Service Management Connect

Connect, learn, and share with Service Management professionals, which are product support technical experts who provide their perspectives and expertise, by accessing Service Management Connect at the following website:

<https://www.ibm.com/developerworks/servicemanagement>

You can use Service Management Connect in the following ways:

- ▶ Become involved with transparent development, in an ongoing, open engagement between other users and IBM developers of Tivoli products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- ▶ Read blogs to benefit from the expertise and experience of others.
- ▶ Use wikis and forums to collaborate with the broader user community.



Installation and upgrade checklist and considerations

This chapter looks at a new IBM Spectrum Control (Spectrum Control) installation and at upgrading an existing installation. The focus of this chapter is to provide guidelines to help you through the installation or upgrade. This chapter refers to the Spectrum Control documentation, providing additional hints and tips that we discovered while creating our example environment.

Note: For our installation and upgrade, we installed Spectrum Control V5.2.8 and IBM DB2® 10.5 Fix Pack 5.

2.1 Preinstallation planning

This section provides steps to help you obtain and prepare the necessary files before starting the installation.

2.1.1 Planning checklist

This section documents the tasks that are part of planning a Spectrum Control installation. It takes you through selecting your product and components to downloading and extracting the code.

Selecting the Spectrum Control product license

The following products are available within the Spectrum Control Family:

- ▶ Spectrum Control Standard Edition
- ▶ Spectrum Control Standard Select Edition
- ▶ Spectrum Control Advanced Edition
- ▶ Spectrum Control Advanced Select Edition

The Spectrum Control product licenses and related functions are summarized in the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1P33zf0>

Selecting the components to install

When running the Spectrum Control installation program, you are presented with the option of installing several components to get started. There are more packages available with the physical media and electronic images, but they are optional and might not be used in all environments. Any optional components can be installed after the initial installation. Knowing what you want to install helps you prepare only the files that you need.

Installation files for the following components are provided with Spectrum Control:

- ▶ DB2
 - DB2 for Windows
 - DB2 for Linux
 - DB2 for IBM AIX®
- ▶ Spectrum Control
 - IBM Spectrum Control for AIX
 - IBM Spectrum Control for Linux
 - IBM Spectrum Control for Windows

In addition, installation files for the following *optional* agents are available with Spectrum Control:

- IBM Spectrum Control Storage Resource agent for Windows
- IBM Spectrum Control Storage Resource agent for UNIX
- IBM Spectrum Control Monitoring Agent for IBM Tivoli Monitoring on Windows
- IBM Spectrum Control Monitoring Agent for IBM Tivoli Monitoring on UNIX

- IBM Jazz™ for Service Management and Tivoli Common Reporting (*optional*)

Install IBM Jazz for Service Management and IBM Tivoli Common Reporting if you want to use Cognos Business Intelligence (BI) reports with Spectrum Control. Cognos BI is offered as an optional tool for producing scheduled reports from the data that is gathered by Spectrum Control. You can view information that is related to your storage resources directly in the Spectrum Control web-based GUI without installing Cognos BI.

If you want to use Cognos BI reports, you need the following components:

- Jazz for Service Management Version 1.1.0.3
- Tivoli Common Reporting Version 3.1.0.1
- Tivoli Common Reporting Version 3.1.0.2
- IBM WebSphere® Application Server Version 8.5.0.1

Note: You can install Spectrum Control without reporting and then later install Jazz for Service Management and the Spectrum Control reports. Choosing to install reporting later allows you to get Spectrum Control running faster.

The optional agent components are used after the initial Spectrum Control installation for certain environments.

After identifying which components you plan to install, obtain the Spectrum Control installation images and download only the ones that you need for your environment.

Obtaining the Spectrum Control V5.2 installation images

Whether you want to do a fresh installation or upgrade your existing installation, you must download the necessary installation packages unless you have them on physical media. In the past, this publication included specific part numbers and file names to help you locate, download, extract, and install the correct packages from IBM Passport Advantage® Online. That was helpful, but it had a downside. The part numbers and file names changed with each release so that this reference was quickly out of date when compared to the latest release. To avoid that problem, this section directs you to content on the IBM Support Portal that can help you get the part numbers and file names for the latest release, even if they are not available at the time of writing.

There are two separate websites for obtaining the Spectrum Control installation packages that are referenced in this section:

- Passport Advantage Online. Use for new installations.
http://www.ibm.com/software/passportadvantage/pao_customer.html
- Fix Central. Use for upgrades.
<http://www.ibm.com/support/fixcentral>

Additional version-specific information for downloading from the locations above is available at the following websites:

- IBM Spectrum Control Support Portal:
http://www.ibm.com/support/entry/portal/product/tivoli/tivoli_storage_productivity_center
- *Latest Downloads for IBM Spectrum Control*, found at:
<http://www.ibm.com/support/docview.wss?uid=swg21320822>

Finding the file names and part numbers by release

The *Latest Downloads for IBM Spectrum Control* web doc links to version-specific documents for all available releases and is frequently updated. Whether you are downloading the electronic images from Passport Advantage Online or Fix Central, use this web doc to identify the latest releases and prepare for installation. The information helps you find the part numbers and file names for what you need to download.

You can find the same download information by going directly to the Spectrum Control support portal. Select **Downloads (view all)** and you are presented with a listing of available downloads for Spectrum Control. You can filter this list by version or operating system. If you already know what you want, there are links at the top of this listing and on the Spectrum Control Support Portal to take you to the fixes on Fix Central or Passport Advantage.

Figure 2-1 show the download options on the IBM Support Portal.



Figure 2-1 Download options on the IBM Support Portal

Tip: In the listing of all downloads, you should find a document for the Spectrum Control release that you want. Start there, and it provides additional information for installing or upgrading to that release. This is the same document that you link to from *Latest Downloads for IBM Spectrum Control*. Each release has its own document.

Downloading for new installations

If you are a client or an IBM Business Partner, the Passport Advantage Online and IBM PartnerWorld® websites provide access to entitled software product installation images in the form of eAssembly bundles. The eAssembly bundles contain installation images that have a product license. They also provide all required and optional parts to get you started. You can access Passport Advantage Online at the following website:

http://www.ibm.com/software/passportadvantage/pao_customer.html

If you have difficulty signing into Passport Advantage Online, follow the links for requesting assistance from the Passport Advantage Online website.

Part numbers and explanations of the files that are included in the eAssembly bundles are provided by the documentation that you can reach from the *Latest Downloads for IBM Spectrum Control* web doc. This information helps you determine which files you need to get started. At a minimum, you need the Spectrum Control server installation images and the DB2 installation image for your operating system. The rest of the images belong to optional components and products that can be installed later.

Downloading for upgrades

The Spectrum Control installation images are available as unlicensed upgrade packages (refresh packs) on the Fix Central support website:

<http://www.ibm.com/support/fixcentral>

These installation images can be downloaded and used for upgrading to a new Spectrum Control version. The links to the version-specific information from the *Latest Downloads for IBM Spectrum Control* can help you find the right set of fixes on Fix Central by providing direct links to the refresh pack version you want.

Note: Use the installation images from IBM Passport Advantage Online if you need the Spectrum Control product licenses or any other bundled products. These licensed installation images can also be used for upgrade.

Extracting the source files

After you have downloaded the installation images, you must extract the compressed files before beginning the installation. It is important to extract the images in dedicated directories as specified for each component.

Spectrum Control

In our scenario, we install Spectrum Control for Windows. The installation images for the server component should be extracted into a single dedicated directory. In our environment, we extracted the files into the `c:\install\spctl_5.2_windows_disk1` directory. If the installation images are downloaded to a Windows directory that has spaces or unusual special characters in the name, Spectrum Control does not install correctly.

Note: Many of the part numbers and file names change between releases. To find the file names that are associated with each product you are installing, check the IBM Support Portal for a web doc on *Downloading IBM Spectrum Control using Passport Advantage Online* for your release.

If you are prompted to merge (or replace) folders or files with the same name while you are extracting the images, select **Merge/Replace**.

IBM DB2 Enterprise Server Edition

Run the self-extracting DB2 executable files. By default, the files are extracted to a temporary location deep in the hierarchy. Therefore, in our example, we specified the `c:\install\db2` directory.

Spectrum Control reports (optional)

Extract Jazz for Service Management, IBM WebSphere Application Server for Jazz for Service Management, and IBM Tivoli Common Reporting into one directory. In our environment, we used the `c:\install\jazzsm` directory. The extractions create separate subdirectories in that location.

2.1.2 Installation image considerations

Here are the special considerations and reminders for handling the installation images:

- ▶ You can download and extract the compressed installation files on any system. However, it is a requirement to have the installation files locally on the server on which you are going to install or upgrade to Spectrum Control V5.2.x. If you mount a network drive with the Spectrum Control installation files and try to run the setup from there, it will not work.
- ▶ Extract the files into separate dedicated directories for each component (for example, DB2, Spectrum Control, and JazzSM).
- ▶ Do not use spaces or special characters in the directory names.

2.2 Installation outline

This section provides a checklist for a fresh Spectrum Control V5.2 installation to help you through all the installation phases and steps.

2.2.1 Installation checklist

Here are the items that you should verify before starting your installation:

1. Choose the Spectrum Control installation environment: single-server or multiple-server.

A multiple-server environment might be suitable for large storage environments, where one server is not sufficient to manage the Spectrum Control components. Customer policies on server location and database placement are also a factor.

Tip: Performance testing has indicated that a local Spectrum Control database repository works best for reporting with Cognos.

Figure 2-2 shows two different scenarios that use the single-server installation type for Spectrum Control.

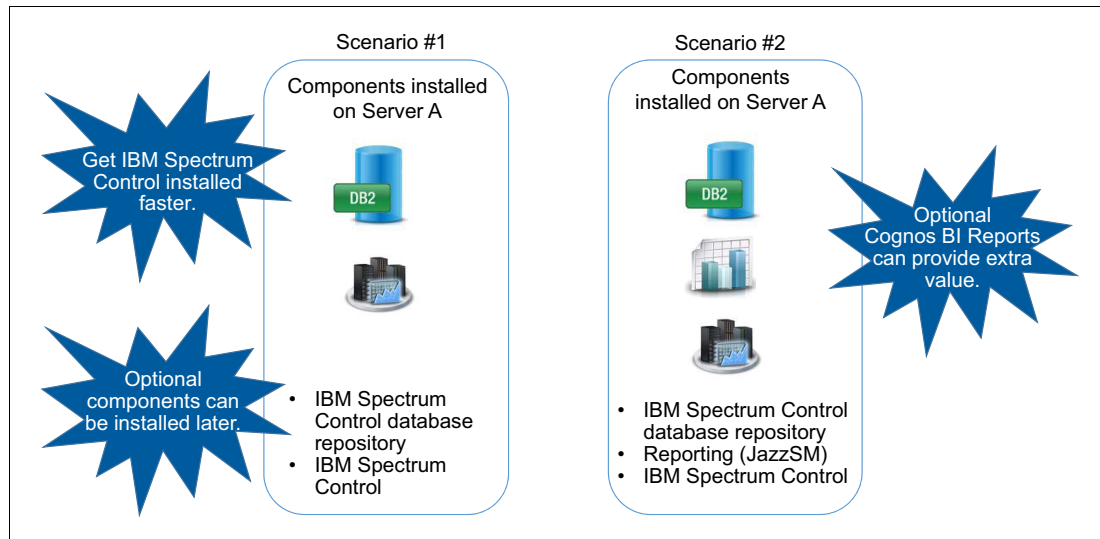


Figure 2-2 Single-server installation scenarios

Figure 2-3 on page 25 shows three different multiple-server installation scenarios where Spectrum Control, DB2, and the Reporting component can be spread across different servers.

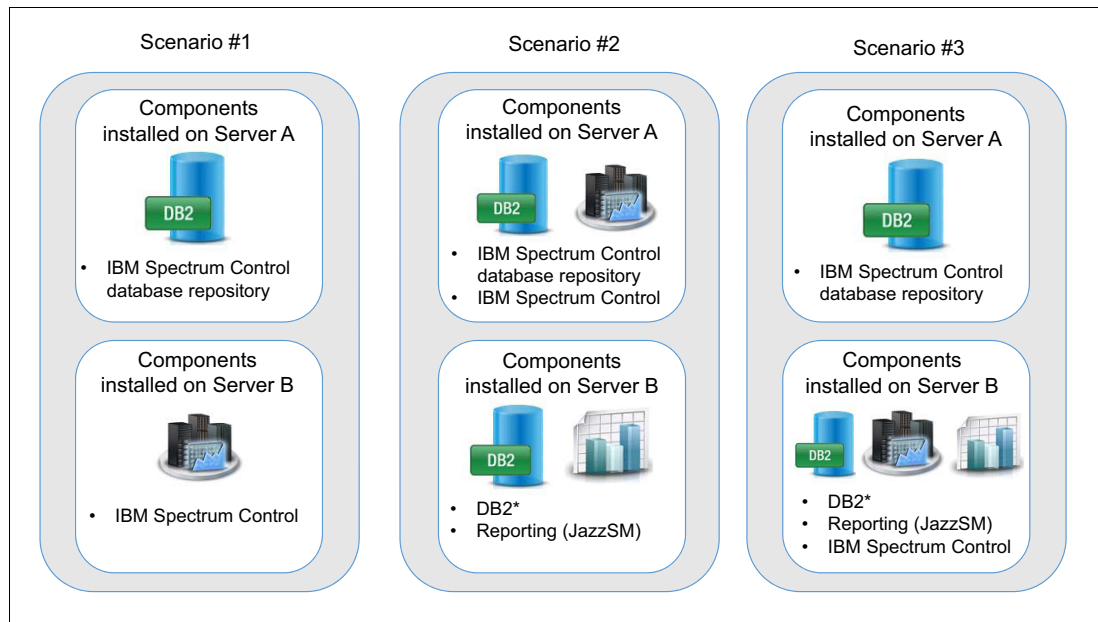


Figure 2-3 Multiple-server installation scenario

In some multi-server installations, two DB2 databases are required because JazzSM and Tivoli Common Reporting require a local DB2 installation. We have indicated that with (*) in Figure 2-3.

Note: Use the single-server installation where possible for several reasons:

- ▶ Less complexity for installation and future upgrades
- ▶ Easier to collect all of the log files for troubleshooting by using the `service.bat/sh` tool that runs on the server where Spectrum Control is installed
- ▶ Simpler backups with only one Spectrum Control server to back up

2. Verify system requirements.

For more information about the hardware requirements for Spectrum Control V5.2, see the hardware support matrix on the IBM Support website, found at:

<http://www.ibm.com/support/docview.wss?uid=swg27047048>

– Operating system requirements:

Spectrum Control is a native, 64-bit application for all server operating systems. You cannot use the Spectrum Control server on 32-bit operating systems. Every application that is included with the server components is a 64-bit application.

– Memory requirements and installation notes:

The preferred memory requirement for Spectrum Control V5.2 is 16 GB. In an evaluation environment with a limited number of volumes, subsystems, fabrics, switches, ports, and agents or agentless servers being monitored, you might not require 16 GB of memory, but at least 8 GB is required. For optimal performance from your Spectrum Control server, you may want to have additional RAM and CPUs.

– Installation notes

In our Windows environment, we successfully installed Spectrum Control on the C: drive of the server. Spectrum Control, DB2, and JazzSM can be installed on drives other than the default C drive for Windows (or / “root” for IBM AIX). Here are the benefits and considerations for not installing on the C drive:

- Space that is allocated for the operating system is not consumed by the Spectrum Control, DB2, and JazzSM applications.
- Performance degradation because a lack of swap space can occur when running out of space on the C drive. This can lead to an operating system crash if Windows cannot write to log and temporary files.

Recommendation: Either have sufficient space on the C drive, or install on another drive, for example, E. If installed on a drive with insufficient space for DB2, see *IBM Tivoli Storage Productivity Center Beyond the Basics*, SG24-8236.

– JazzSM

If you choose to install JazzSM, there are two choices you can make:

- Install JazzSM on your C drive. If you choose to do this, you need an additional 11 GB. This requirement is in addition to the 12 GB that is required for Spectrum Control.
- Install JazzSM on a drive other than C, for example, E.

Note: If you choose to install JazzSM on a drive other than your C drive, in addition to the 11 GB of available space on your other drive, you still need to have 5 GB of free space available on your C drive for temporary files.

The 5 GB of temporary space can be pointed to another location by changing the TMP and TEMP location for \$USER\Application\data\ from your Windows environment, as shown in the following steps:

- i. Click **Control Panel** → **System and Security** → **System** and click **Advanced system settings**, as shown in Figure 2-4.

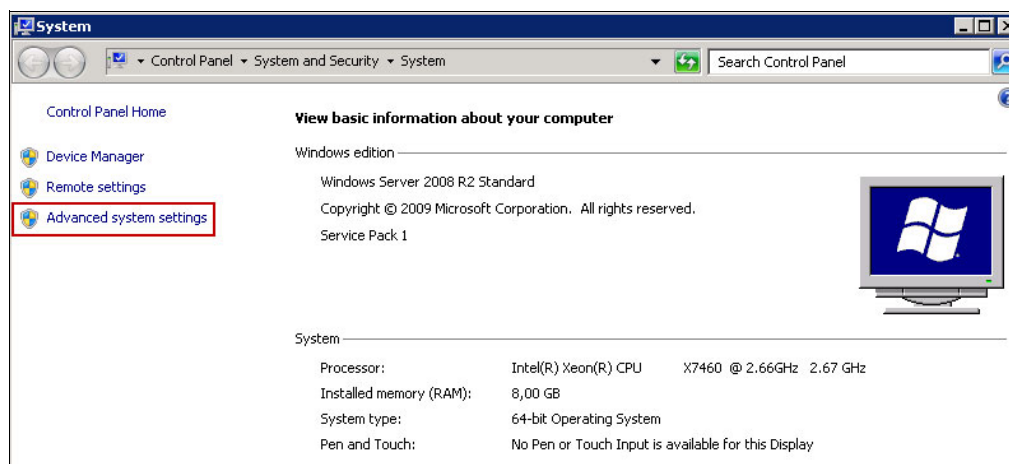


Figure 2-4 System properties

- ii. Click **Environment Variables**, as shown in Figure 2-5 on page 27.

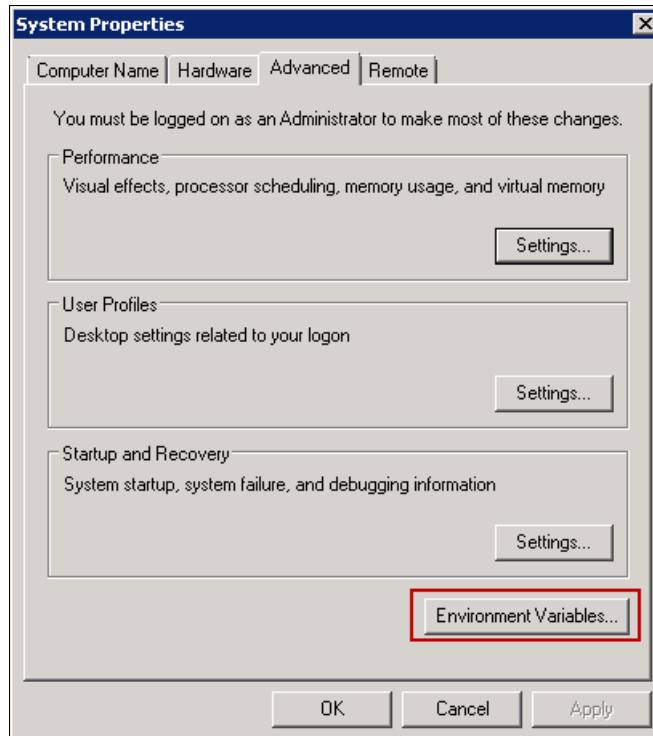


Figure 2-5 Environment variables

- iii. You can now change the TMP and TEMP location for the user that is installing Spectrum Control. In this example, the user is dk021887 and TEMP and TMP need to be changed to E:\temp. Click the line with TMP marked “1”, and then click **Edit**, marked “2”, as shown in Figure 2-6.

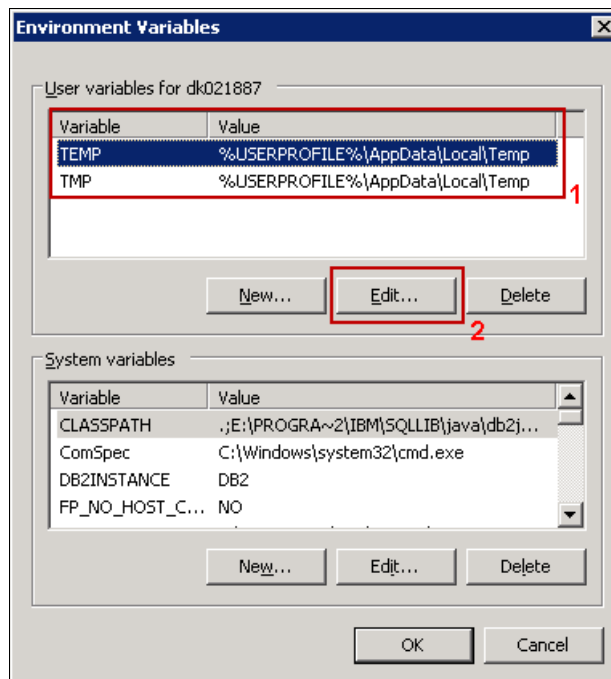


Figure 2-6 TEMP environment variables

3. Plan the installation method in a Windows domain.

Choose between using local user accounts or domain user accounts for the Spectrum Control common user and the Spectrum Control database user.

Note: Do not use a mix of local and domain IDs for a single Spectrum Control server. Use local accounts where possible because they have no effect on the type of user IDs that can log in to Spectrum Control. Using domain IDs for users of Spectrum Control is still possible when you install Spectrum Control with local IDs for the technical user IDs.

For more information about using local or domain user IDs for the installation or on using an external authentication server, see Chapter 4, “Configuration and administration considerations, tasks, and use cases” on page 71. For more information about authentication, see Chapter 3, “Setting up authentication” on page 47.

4. Verify the user name and password requirements and define users.

A useful worksheet to document the user names and passwords is provided in the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/10molXq>

Note: Here are the special characters that are allowed. Because you can use separate users for WebSphere and DB2, we show the allowed characters by application:

WebSphere password: [,] , ? , ` , ~ , ! , (,) , - , _ , .

DB2 password: \$, ~ , @ , # , (,) , - , _ , { , } , .

If you are using the same password for DB2 and WebSphere, any special characters need to be supported by both. Therefore, the Spectrum Control common user password should contain only these special characters:

Spectrum Control common user: ~ , (,) , - , _ , .

The comma (,) is not an allowed special character. It is used here as a separator between the allowed characters.

5. Plan the role definition for Spectrum Control users.

The following roles are available with Spectrum Control:

- Administrator
- Monitor
- External Application

For more information about the roles and role to group mappings, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1Ri0b6c>

6. Plan for managing and monitoring resources, including storage systems, hypervisors, file systems, and switches with Spectrum Control.

Assess the scope of your environment, based on the managing and monitoring needs, to accomplish the following tasks:

- Identify the communication methods (native API, SNMP, Storage Resource agent, or SMI-S agent).
- Verify whether the interoperability requirements are met. You can find the Supported Hardware, Products, and Platforms Interoperability Matrix Links web doc at the following website:
<http://www.ibm.com/support/docview.wss?uid=swg21386446>
- Plan and implement all necessary actions to meet the interoperability requirements (for example, SAN switch or storage device firmware updates, or device driver updates).

7. Plan for Data Path Explorer.

Note: The data path visualization requires Adobe Flash Player to be installed in your web browser.

2.2.2 Installation

This section gives a suggested set of actions for a Spectrum Control V5.2.x fresh installation, pointing out important notes.

We refer to the installation scenario that we implemented in our test environment, which consists of a Spectrum Control Advanced Edition installation on a Windows system in a single-server installation type.

For more information about how to download the installation images for Spectrum Control V5.2 by using the Passport Advantage Online website, see “Obtaining the Spectrum Control V5.2 installation images” on page 21.

For a detailed description of each installation step, review the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/10KAb19>

Preinstallation steps

Before starting the installation of Spectrum Control, complete the following steps:

1. Verify that Spectrum Control server is defined in the domain name system (DNS) Server.

Important: Use DNS. Spectrum Control always uses a fully qualified name, thus if Spectrum Control is not defined to a DNS server, you will have problems when navigating through the Spectrum Control panes and windows.

2. Verify or edit the `etc/hosts` file.

JazzSM requires the entry `localhost` in the `etc/hosts` file. Typically in Windows, the file has the entry `127.0.0.1 localhost` commented out.

Open the C:\Windows\System32\drivers\etc\hosts file by using Notepad and verify that the line localhost is uncommented, as shown in Example 2-1.

Example 2-1 The hosts file

```
# localhost name resolution is handled within DNS itself.  
127.0.0.1      localhost
```

Note: Make sure that you start the Notepad application as an administrator. We have seen situations where you can edit and save the file, but changes were not written into the same location, so the changes did not have any effect. Open the file after your changes to verify that the changes are saved.

3. Verify the language settings.

Often, your operating system is installed and the interface is using one language, but when you install additional software, it uses a different language because there are several ways to detect the language of the system, and if you have not set all settings to the same language, you can end up having mixed languages.

The Spectrum Control web-based GUI uses the language setting of your browser. Set the rest of Spectrum Control to use the English language because that also determines how logs are written. This setting affects how a problem can be handled by IBM Support.

For a white paper about how to set the language for all components on a Windows server, see the following website:

<http://www.ibm.com/support/docview.wss?uid=swg27038231>

Note: Even though this white paper was written for Tivoli Storage Productivity Center V5.1, it remains valid for Spectrum Control V5.2.x.

4. Install DB2.

Install DB2 by following the detailed instructions that are provided in the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1IYVpbd>

Note: The DB2 Control Center is no longer available as of DB2 10. To administer your database, you can use IBM Data Studio. The IBM Data Studio installation is included with DB2 10.1 or the 10.5 product installation. The DB2 Setup wizard provides an option to install Data Studio components.

Tip: If you do not want to create the Spectrum Control database and the Tivoli Common Reporting database on the Windows C: drive, you should change the DB2 default database location setting. To do this, run **update dbm config using dftdbpath <drive>**.

5. Extract the Spectrum Control installation packages (for more information, see 2.2, “Installation outline” on page 24).

In our installation scenario, we chose to install JazzSM at the same time as Spectrum Control installation to use the Spectrum Control reporting capabilities.

Note: If you do not need to implement Spectrum Control reports, you can choose to not install JazzSM and install only Spectrum Control.

You can decide to install JazzSM at any time after the Spectrum Control installation.

- a. Extract all of the Spectrum Control server parts into a single directory. In our example, we used the `C:\install\spctl_5.2_windows_disk1` folder.

Important: If the installation images are downloaded to a Windows directory that has spaces or unusual special characters in the name, Spectrum Control does not install correctly.

Note: Future Spectrum Control program temporary fix (PTF) levels might present a different number of parts for the `disk1` file. Make sure that all the parts are downloaded and extracted into the same folder.

- b. Extract the following files into a new folder. In our example, we used the `C:\install\jazzsm` folder.

- JazzSM file:
Jazz for Service Management 1.1.0.3
- WebSphere Application Server file:
WebSphere Application Server 8.5.0.1
- Common Reporting files:
Tivoli Common Reporting 3.1.0.1
Tivoli Common Reporting 3.1.0.2

Note: The versions that are listed here apply to Spectrum Control V5.2.8. Future releases of Spectrum Control might include updated versions of these applications.

Installation steps

To install Spectrum Control V5.2.x, complete the following steps:

1. Start the Spectrum Control installation program. In our scenario, it is in `C:\install\spctl_5.2_windows_disk1\TPC\setup`.

Note: Run the installation program as an administrator by right-clicking `setup` and selecting **Run as administrator**.

2. Select a language and click **OK**.
3. Read the program license agreement and select **I accept both the IBM and the non-IBM terms**.

4. In the Before You Begin: Prerequisites window, in our example, we chose to install the optional reporting features, JazzSM, and Tivoli Common Reporting. Click **Install now** (Figure 2-7).

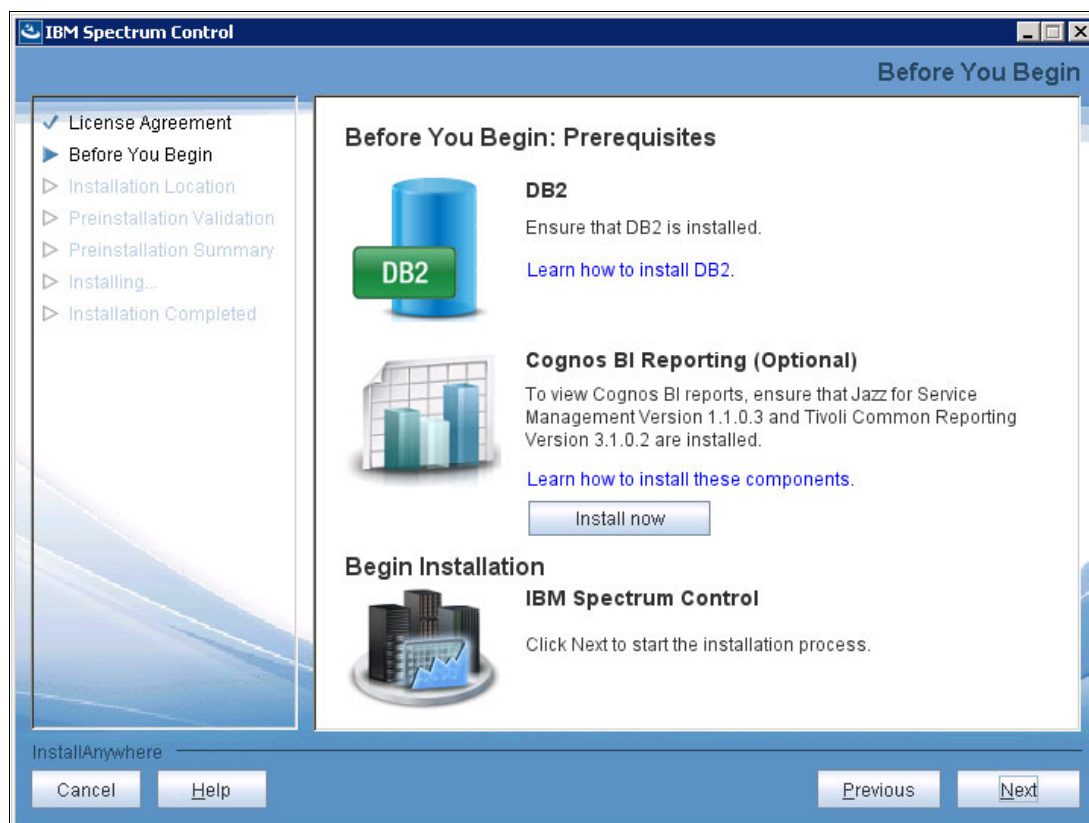


Figure 2-7 Before You Begin: Prerequisites window

Note: If you click the **Learn how to install these components** link, you have access to the JazzSM installation by using the Launchpad on Windows document.

Important: Install JazzSM by using the Spectrum Control installer because this is the only way the JazzSM Windows services are created. The service is not created if the JazzSM installer is started in a stand-alone way.

If you decide to install the reporting function later, be sure to use the Spectrum Control installer.

The following steps are specifically for the JazzSM installation:

- In the Install Jazz for Service Management window, specify the location of the JazzSM installation file. In our installation scenario, it is C:\install\jazzsm.
- In the Welcome to Jazz for Service Management 1.1.0.3 window, select **Custom Installation** (Figure 2-8 on page 33).

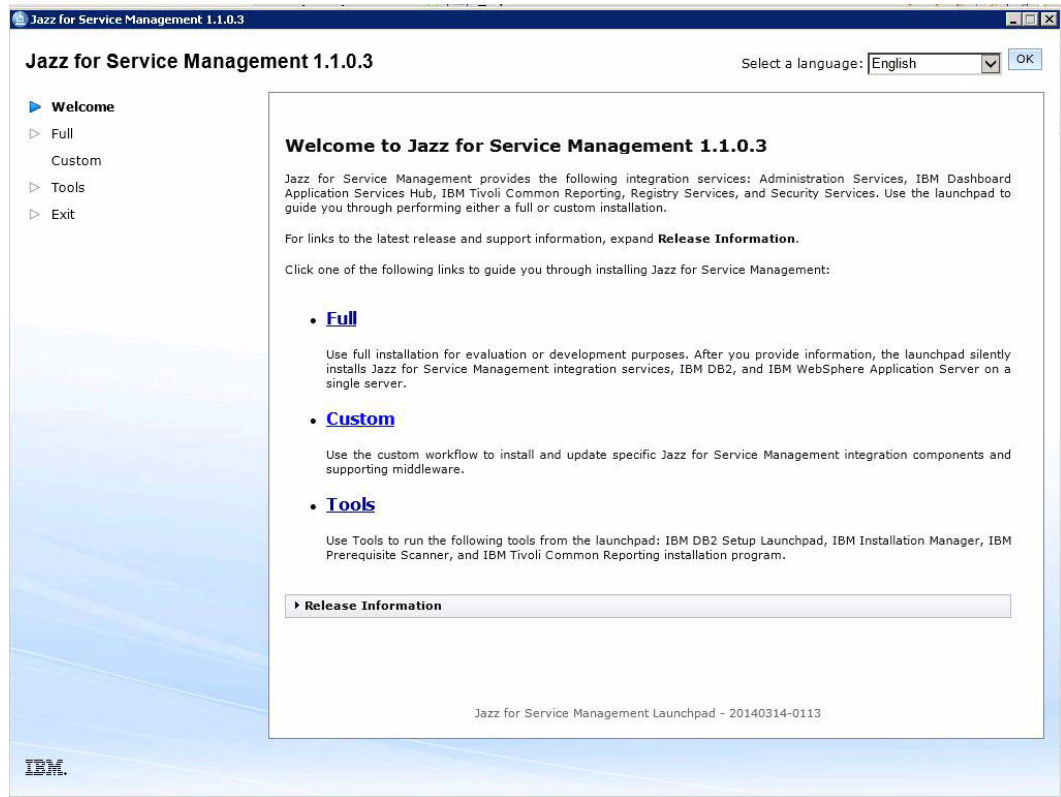


Figure 2-8 Welcome to Jazz for Service Management 1.1.0.3

- c. The Custom Workflow window shows a summary of the available JazzSM services and supporting middleware that can be installed. Click **Next**.
- d. In the Specify Jazz for Service Management Home Location window, click **Next**. The empty Existing environment field means that you do not have an existing Jazz for Service Management installation (Figure 2-9).

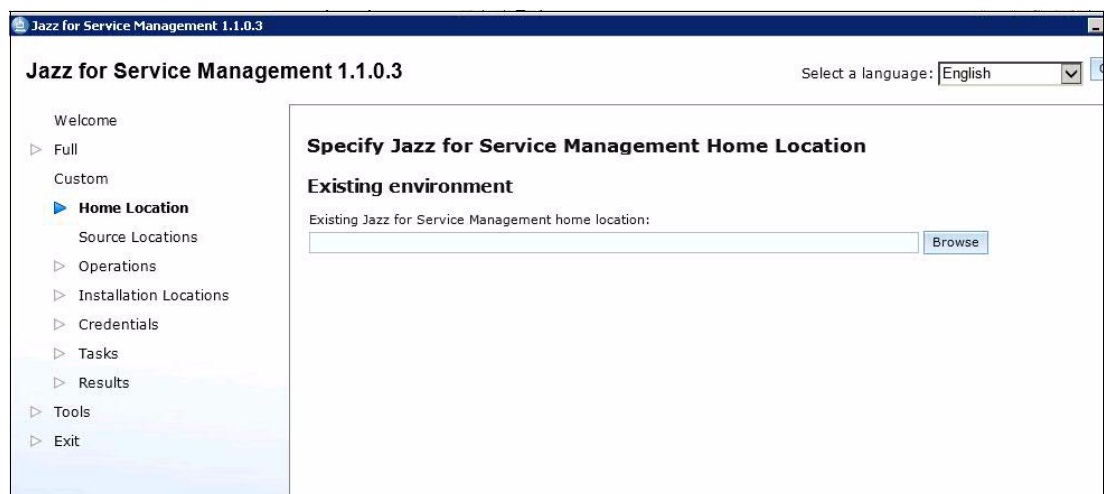


Figure 2-9 Specify Jazz for Service Management Home Location window

- e. In the Specify Source Location window, leave the prefilled values as the defaults and click **Next**.

- f. In the Select Operations window, choose to install only **Reporting Services** and **IBM WebSphere Application Server**, as shown in Figure 2-10. As you can see, the **IBM DB2 Enterprise Server Edition** entry is disabled because we already installed DB2. The DB2 installation files have not been added to the JazzSM installation folder.

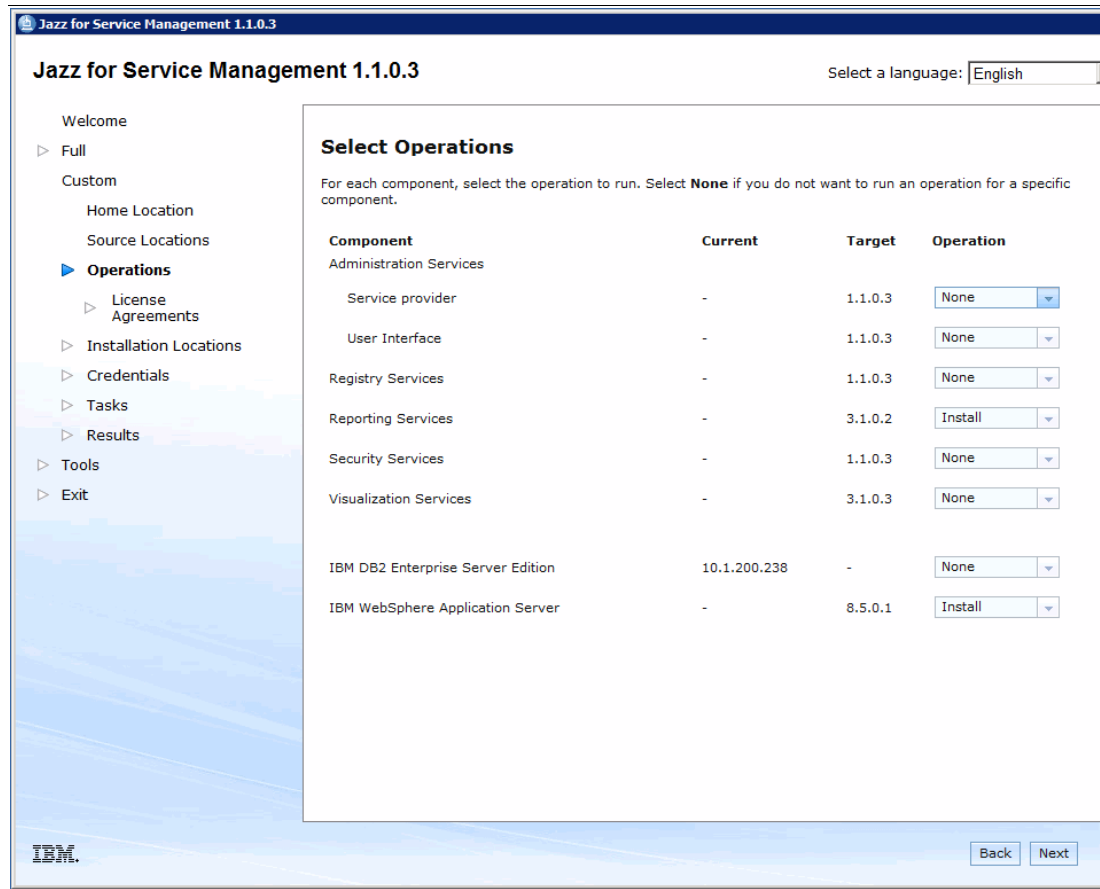


Figure 2-10 Select Operations window

Important: If you plan to configure Lightweight Directory Access Protocol (LDAP) authentication by using IBM System Storage DS8000 Storage Manager or IBM System Storage SAN Volume Controller, you must install Security Services.

- g. In the license agreement window, select **I accept the terms in the license agreement** and click **Next**.

- h. In the Specify Installation Locations window, you can specify an installation location or leave the predefined paths. As a preferred practice, use the predefined paths, as shown in Figure 2-11. The **DB2 installation location** entry is disabled because we already installed DB2.

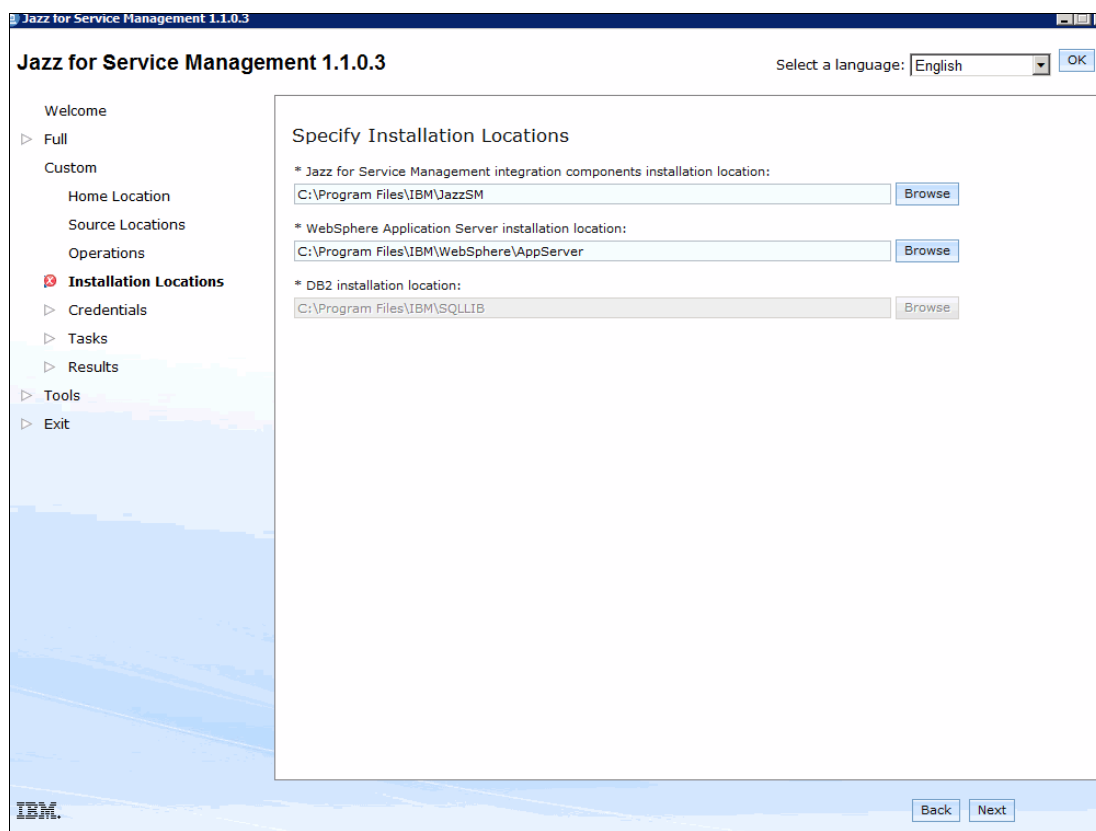


Figure 2-11 Specify Installation Locations

5. In the Specify Credentials window, complete the credential entries for WebSphere Application Server administrator and DB2, as shown in Figure 2-12.

Jazz for Service Management 1.1.0.3

Select a language: English OK

Specify Credentials

WebSphere Application Server Profile

Administrator name: smadmin

Administrator password:

Confirm administrator password:

DB2

Instance user name: db2admin

Instance user password:

Back Next

Figure 2-12 Specify Credentials window

WebSphere Application Server user: The WebSphere Application Server administrator is a new user, and you can change the predefined name and set the password. This new user is not a user in the local operating system; it is created in the WebSphere file repository, and it is the only user that can access the reporting part of Spectrum Control after the installation. You can change the WebSphere configuration later to use the operating system users. You can also enable LDAP or Active Directory servers for reporting. For more information, see 3.2, “Implementing authentication” on page 50.

Connecting to DB2: To allow Spectrum Control access to the DB2 database, you must specify a user with administrative role to DB2 and to the operating system.

For example, change the smadmin user, which is predefined and is the default value, to db2admin to match the DB2 administrator credentials that are created during the DB2 installation.

Important: Be sure to assign different user IDs for the WebSphere Application Server administrator and to the DB2 administrator, or Spectrum Control will have problems later.

- a. Figure 2-13 on page 37 shows the Run Tasks window. Review the listed tasks sequence to be run (marked with “1”). If all wanted components are listed, click **Run** (marked with “2” in the lower right).

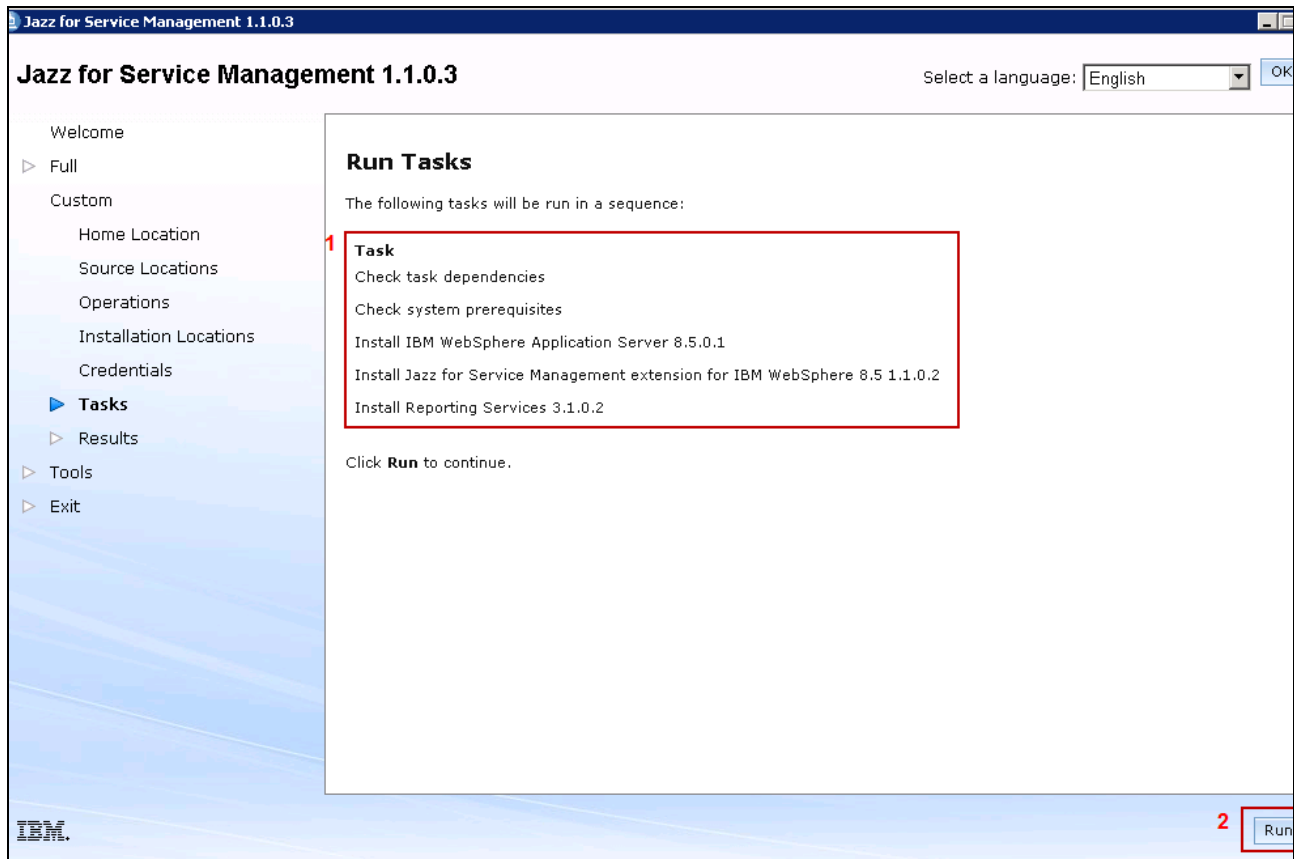


Figure 2-13 Review tasks and start the installation

- You can view the status and details, marked with “1”, of each task as it completes, as shown in Figure 2-14. By clicking the corresponding **View Details** button in the Review Results window, you can get more details about that task. When all tasks complete, click **Exit** in the left navigation tree.

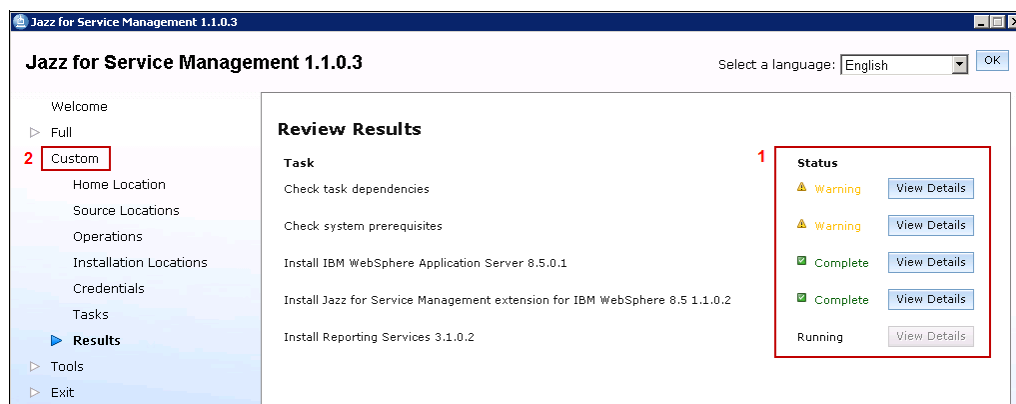


Figure 2-14 Installation progress

- c. If any task fails, click **View Details** button and search for FAIL in the text file. An example of an error during installation is shown in Example 2-2. When the error is corrected, click **Custom** (marked with a “2” in Figure 2-14 on page 37) to restart the installation. All entered and selected options are retained, making a retry simple.

Example 2-2 Installation error

Property	Result	Found	Expected
=====	=====	=====	=====
os.localhostInHostsFile	FAIL	False	True

- d. Return to the Install Jazz for Service Management window either by clicking **Done** in the lower right corner or **Exit** in the menu to the left (shown in Figure 2-14 on page 37).
The Install now button is disabled. The JazzSM installation is complete. You can proceed with Spectrum Control installation by clicking **OK**.
6. In the Choose Installation Location and Type window, select the installation location or accept the defaults and choose **Single server** as Installation Type, as shown in Figure 2-15.

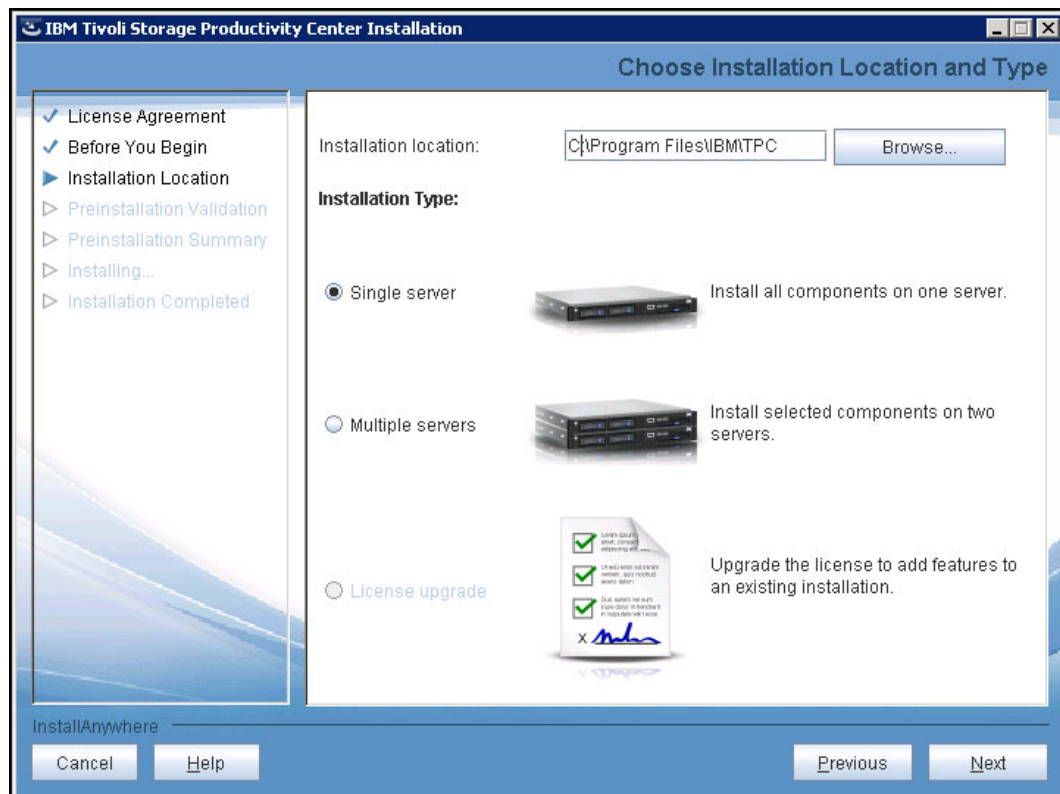


Figure 2-15 Installation location

7. In the Single Server Installation Information window, enter the DB2 administrator credentials that you defined during the DB2 installation, as shown in Figure 2-16 on page 39.

IBM Tivoli Storage Productivity Center Installation

Single Server Installation Information

- ✓ License Agreement
- ✓ Before You Begin
- ✓ Installation Location
- ▶ **Preinstallation Validation**
- ▷ Preinstallation Summary
- ▷ Installing...
- ▷ Installation Completed

Server Information

Host name: ?

Ports: to 9569

[Verify port availability](#)

Common User Name and Password

Specify a user name and password to configure all components.

User name: ?

Password:

Advanced Customization

[Configure Database repository](#)

[Configure Jazz for Service Management and Tivoli Common Reporting](#)

InstallAnywhere

Figure 2-16 Enter DB2 credentials

Note: If you plan to install Spectrum Control in a large environment, you should consider changing the default location of the Spectrum Control database and store it on fast disks, for example, IBM FlashSystem 900 or V9000. In addition, consider using separate disks for the DB2 redo logs, as for any other large database installation.

8. In the Configure Jazz for Service Management and Tivoli Common Reporting window, enter the JazzSM credentials that you defined during the JazzSM installation (Figure 2-12 on page 36), as shown in Figure 2-17.

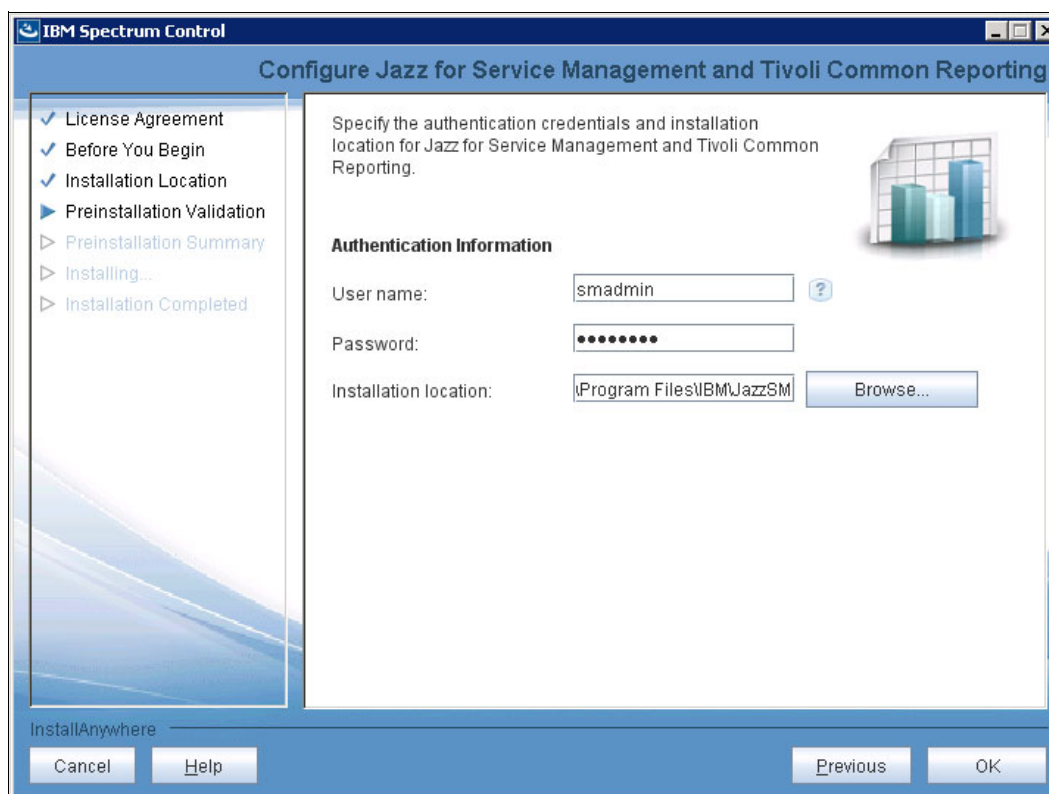


Figure 2-17 JazzSM credentials

9. In the Preinstallation Summary window, review the installation information and click **Install**. The installation might take a couple hours, depending on your server hardware configuration.
10. When the installation is complete, in the Installation Completed window, the links to connect to the Spectrum Control web user interface are shown. Click **Done** to close the wizard.

2.2.3 Postinstallation tasks

After the Spectrum Control installation, you should perform all or some of the following postinstallation tasks, depending on your environment:

- ▶ Verify the Spectrum Control installation, as detailed in the Spectrum Control IBM Knowledge Center, found at:
<http://ibm.co/22VkJ8L>
- ▶ Configure Spectrum Control to send email notifications by clicking **Settings** → **Alert Notifications**.
- ▶ Configure Spectrum Control to send SNMP traps to an SNMP Manager by clicking **Settings** → **Alert Notifications**.
- ▶ Configure the LDAP environment.

- ▶ Implement Role to Group Mapping for Spectrum Control users by clicking **Settings User** → **Management**.
- ▶ To monitor and administer a DB2 environment, download and install IBM Data Studio V3.2 or later. The downloadable file and installation instructions are available at the following website:
<http://www.ibm.com/support/docview.wss?uid=swg24033663>
- ▶ Set up your browser to use the Spectrum Control web-based GUI:
 - Check that your browser is supported.
 - Adjust the firewall so that the Spectrum Control web-based GUI can be reached from other systems.
 - Adjust your browser's software to block pop-up ad windows.
 - For the data path viewer, you need Adobe Flash Player.

Proceed with the Spectrum Control configuration tasks, as described in Chapter 4, “Configuration and administration considerations, tasks, and use cases” on page 71.

2.3 Upgrade summary

You can upgrade from Tivoli Storage Productivity Center Version 5.1 or Version 5.2 to Spectrum Control Version 5.2. You cannot upgrade from any release before Tivoli Storage Productivity Center Version 5.1. In our example, we upgrade to Version 5.2.8.

Spectrum Control uses the same installation directory and database name as Tivoli Storage Productivity Center.

2.3.1 Planning considerations

The following sections describe the planning considerations for Spectrum Control upgrades.

Considerations for upgrades from Tivoli Storage Productivity Center V5.2.x

The following list is a summary of the considerations for upgrading your Tivoli Storage Productivity Center V5.1.x to Spectrum Control V5.2:

- ▶ Spectrum Control is a native, 64-bit application for all server operating systems. You cannot use the Spectrum Control server on 32-bit operating systems. Every application that is included with the server components is a 64-bit application.
- ▶ DB2 9.7 is not supported by Spectrum Control V5.2. DB2 must be upgraded to a supported 64-bit version first. For more information about the DB2 upgrade, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1PnDcAP>

For more information about the supported DB2 levels, see *Supported Platforms*, found at:

<http://www.ibm.com/support/docview.wss?uid=swg27047043>

- ▶ When upgrading from Tivoli Storage Productivity Center V5.2.x to Spectrum Control V5.2.x, maintenance updates are applied to Tivoli Storage Productivity Center for Replication. It remains installed as Tivoli Storage Productivity Center for Replication. A license for Copy Services Manager is included with Spectrum Control to allow migration for maintenance and new features.

Considerations for upgrades from Tivoli Storage Productivity Center V5.1.x

The following list is a summary of the considerations for upgrading your Tivoli Storage Productivity Center V5.1.x to Spectrum Control V5.2:

- ▶ Windows 2003 and Windows 2008 32-bit environments are not supported by Spectrum Control. To upgrade older Tivoli Storage Productivity Center installations, the operating system must be migrated to a Windows 64-bit environment.
- ▶ DB2 9.7 (32-bit or 64-bit) is not supported by Spectrum Control V5.2. DB2 must be upgraded to a supported 64-bit version. For more information about the DB2 upgrade, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1PnDcAP>

For more information about DB2 levels that are supported by Spectrum Control, see *Supported Platforms*, found at:

<http://www.ibm.com/support/docview.wss?uid=swg27047043>

- ▶ The Spectrum Control V5.2 preferred memory requirement is 16 GB. Verify whether additional memory is required for your Spectrum Control server.
- ▶ Tivoli Integrated Portal is not used in Spectrum Control. During the Spectrum Control upgrade process, you can choose to preserve the old Tivoli Integrated Portal instance. As a preferred practice, uninstall Tivoli Integrated Portal unless it is being shared with another product. Tivoli Integrated Portal data is not migrated.
 - LDAP configuration settings are not migrated from Tivoli Integrated Portal to JazzSM. You must configure manually the JazzSM with LDAP repository by using the WebSphere Application Server administrative console. For more information, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1ZWxy03>
 - If you used any custom Tivoli Common Reporting reports or configured DB2 as an external content store with Tivoli Common Reporting V2.1.1, see *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

Important: If your DS8000 or SAN Volume Controller storage system is configured to use Single Sign-On with Tivoli Storage Productivity Center V5.1, the DS8000 or SAN Volume Controller has stored the Tivoli Integrated Portal information in its configuration. This configuration must be switched back to local authentication during the upgrade of Spectrum Control because Tivoli Integrated Portal is no longer used and should be uninstalled.

To be on the safe side, consider reverting to OS authentication for Tivoli Storage Productivity Center V5.1.x during the upgrade. In addition, check your documentation regarding the values that you need for reenabling LDAP and Active Directory later.

- ▶ The roles that were defined in Tivoli Storage Productivity Center V5.1 were consolidated. For more information, see Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1ZWxy03>

- ▶ When upgrading from Tivoli Storage Productivity Center 5V.1.x to Spectrum Control V5.2.x, maintenance updates are applied to Tivoli Storage Productivity Center for Replication. It is upgraded and shows as Tivoli Storage Productivity Center for Replication, not as Spectrum Control. A license for Copy Services Manager is included with Spectrum Control to allow migration for maintenance fixes and new features.
- ▶ If you have CAS-based Data or Fabric Agents from earlier releases, migrate these agents to Storage Resource agents before upgrading.

Considerations for upgrades from Tivoli Storage Productivity Center V4.2.x

You cannot upgrade from Tivoli Storage Productivity Center V4.2.2 or any earlier releases to Spectrum Control V5.2. Upgrade from Tivoli Storage Productivity Center V4.2.x to Tivoli Storage Productivity Center V5.2.7 and then to Spectrum Control. This action avoids installing Tivoli Integrated Portal with Tivoli Storage Productivity Center V5.1.x. For more information about that upgrade, see *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

2.3.2 Pre-upgrade steps

Here are the pre-upgrade steps:

1. Back up the existing Tivoli Storage Productivity Center or Spectrum Control database and server machine to provide a point to restore if the upgrade fails. Instructions for backing up the database by using DB2 backup CLI commands are available in the Spectrum Control IBM Knowledge Center, found at:
<http://ibm.co/1ZgcG1y>
2. If you have a standby server that is configured for the Tivoli Storage Productivity Center for Replication server, make the Tivoli Storage Productivity Center for Replication instance on the stand-by server active. Make the Tivoli Storage Productivity Center for Replication server that you are upgrading the standby server.
3. Follow any necessary actions from the upgrade considerations that are presented in 2.3.1, “Planning considerations” on page 41.

2.3.3 Upgrade steps

The upgrade steps are similar to the installation steps that are described in “Preinstallation steps” on page 29, except that the DB2 installation step is skipped because you have installed DB2.

Complete the following steps:

1. Verify or edit the `etc/hosts` file (see “Preinstallation steps” on page 29).
2. Extract the installation packages (see the extraction steps in “Preinstallation steps” on page 29).
3. Start the Spectrum Control installation program (see “Installation steps” on page 31) and follow the same steps as the fresh installation. During Spectrum Control installation steps, the Uninstall Tivoli Integrated Portal window might open. This window opens when you upgrade from Tivoli Storage Productivity Center V5.1.x. You should uninstall it to reclaim system resources unless it is shared with another product. Spectrum Control V5.2 does not use Tivoli Integrated Portal.

2.3.4 Post-upgrade tasks

After the Spectrum Control migration is complete, verify the correct upgrade of your Spectrum Control environment by implementing verification steps by using the checklist that is provided in 2.2.3, “Postinstallation tasks” on page 40 as a guideline.

In addition, you must verify that all the tasks and schedules that are defined in the former Tivoli Storage Productivity Center environment (performance monitoring jobs and reporting jobs) are correctly redefined and active in your new Spectrum Control V5.2 environment.

2.4 Installing the Reporting component later

If you did not install the Reporting component, you see only a link to learn how to install Cognos when you click the Cognos menu item in the GUI, as shown in Figure 2-18.

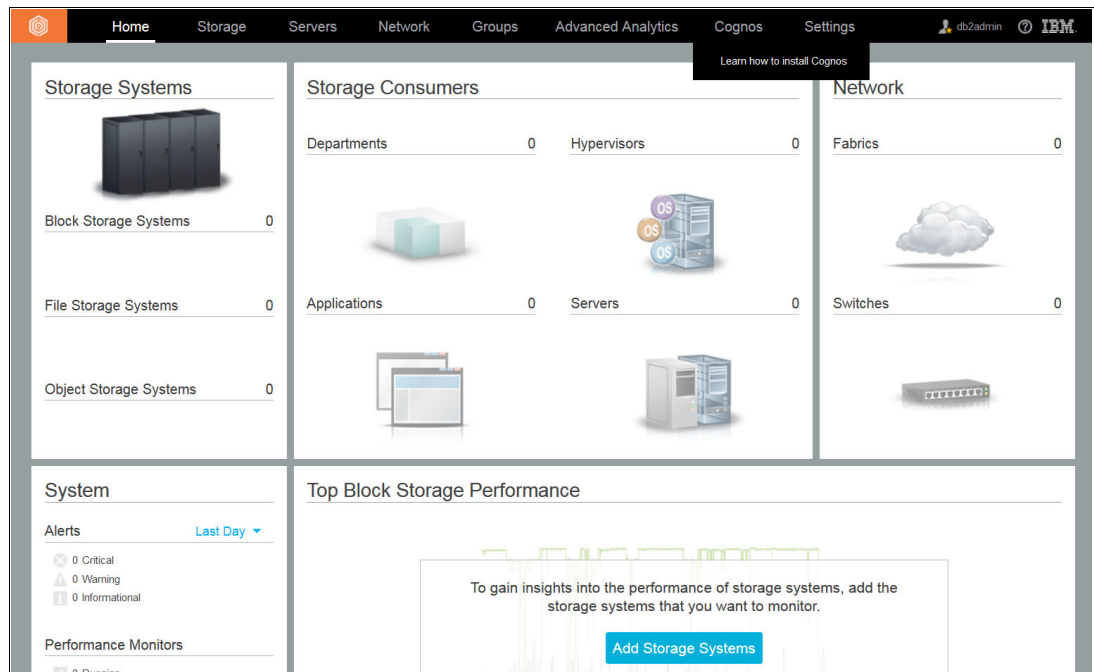


Figure 2-18 Spectrum Control web-based GUI when Cognos Reporting is not installed

Note: Although the Learn how documentation can help you with installing JazzSM, use with Spectrum Control installation program to install JazzSM for a simpler installation workflow.

To install JazzSM by using the Spectrum Control installation program, complete the following steps:

1. Download and extract the images (see 2.1, “Preinstallation planning” on page 20).
2. Verify or edit the etc/hosts file (see “Preinstallation steps” on page 29).
3. Start the Tivoli Storage Productivity Center installer.

In the Before You Begin: Prerequisites window, choose to install the optional Reporting features, JazzSM, and Tivoli Common Reporting by clicking **Install now**, as shown in Figure 2-19 on page 45.

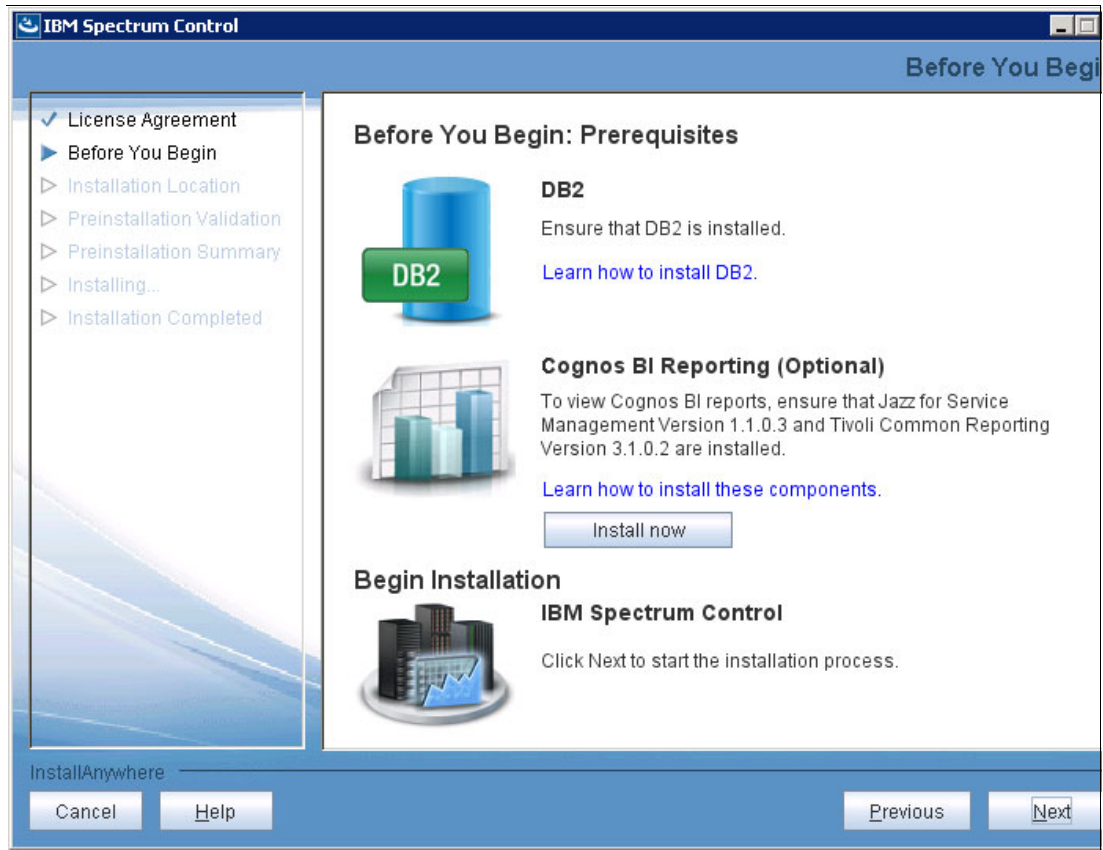


Figure 2-19 Before You Begin: Prerequisites window

Note: If you click the **Learn how to install these components** link, you have access to the JazzSM installation by using the Launchpad on Windows guide.

4. Continue with the installation steps of the JazzSM component, as described in step a on page 32 through step c on page 38.
5. Continue the installation steps in the Spectrum Control installation program to install the Spectrum Control specific reports in the Reporting environment.
6. Restart the Spectrum Control web-based GUI by using the IBM WAS80Service - SpectrumControl service on Windows.



Setting up authentication

Since the introduction of the web-based GUI, the use of external authentication servers continues to grow. IBM Spectrum Control (Spectrum Control) supports using Lightweight Directory Access Protocol (LDAP) or Active Directory (AD) as authentication methods.

This chapter explains how you can use Spectrum Control with external authentication services and what you must understand to configure this setup correctly.

This chapter uses the term *external authentication server* and *directory server* interchangeably. Although both terms have the same meaning, there are times when one or the other characteristic should be emphasized, so both terms are used. Specifically, this chapter describes Spectrum Control user roles and Cognos authentication.

This chapter briefly describes how to implement the different cases of authentication:

- ▶ Implementing Local OS authentication and an internal user ID
- ▶ Implementing external authentication (direct LDAP) and a local internal user ID
- ▶ Implementing external authentication (implicit) and a local internal user ID

For more information about configuring these implementations, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/10UfPmd>

This chapter also demonstrates a use case about how to implement authentication by using the IBM Bluepages (intranet), as described in 3.3, “Use case: Configuring Spectrum Control to authenticate with Bluepages for an intranet user” on page 55.

A complete guide to setting up authentication can be found in *IBM Tivoli Storage Productivity Center Beyond the Basics*, SG24-8236.

3.1 Planning for authentication

Authentication in Spectrum Control can be complex, and this is especially true for large enterprises that have sophisticated security compliance and mechanisms in place.

Planning for authentication should begin even before you decide on the layout to install Spectrum Control. Without fully understanding how authentication works, the person (storage administrator) who is about to install Spectrum Control can easily choose a complex path.

A preferred practice is to install Spectrum Control by using the Local OS administrative account. However, if your security policy does not allow it, then this is the first question that must be answered: Do you need to install Spectrum Control by using a domain account, and what security restrictions are in place?

Typically, applications are installed by using a domain account because of security considerations, but the domain grants other advantages:

- ▶ Centralized user provisioning
- ▶ Centralized user de-provisioning
- ▶ Ability to enforce a policy on the accounts
- ▶ Centralized auditing of authentication attempts

Spectrum Control supports LDAP or AD as authentication methods. However, if you must install Spectrum Control by using an external user, that narrows down your possibilities to AD because LDAP can be used only for external authentication, *not* for installing the product.

You must understand that Spectrum Control has internal user IDs that are used by Spectrum Control components to communicate and external user IDs (users of Spectrum Control) that are used by administrators or users of the product.

Figure 3-1 shows the user ID and repository combinations schematic from *IBM Tivoli Storage Productivity Center Beyond the Basics*, SG24-8236. This chapter describes each use case. In addition, this chapter describes a use case about how to implement authentication by using the IBM Bluepages (intranet), which is described in 3.3, “Use case: Configuring Spectrum Control to authenticate with Bluepages for an intranet user” on page 55.

		Authentication for users of IBM Spectrum Control		
		Local	External	
Internal User IDs	Local	1 Simple Installation	2 Sufficient in most environments	Method 1 (recommended)
	External	3 Possible, but makes no sense	4 Most Complex configuration	Method 2

Figure 3-1 Authentication in Spectrum Control

3.1.1 Authentication and authorization

Authentication is the process in which the provided credentials are compared to the ones that are stored. In Spectrum Control, the provided credentials are the ones that are used in different interfaces, such the web-based GUI, the command-line interface (CLI) based tpctool, or Cognos Reporting. These credentials are compared to the ones that are stored in the File Registry, Local OS, or LDAP repository. A high-level overview is shown in Figure 3-2 on page 51.

Authorization is a function that specifies the access rights a user has that are related to a resource in a computer system. In Spectrum Control, role-based authorization is used, for example, the administrator role is authorized to view and accomplish more actions than the monitor role. For more information about role-based authorization, see the IBM Knowledge Center, found at:

<http://ibm.co/1N2ZB4s>

Section 3.2, “Implementing authentication” on page 50 describes the different authentication possibilities. It shows the Spectrum Control environment as a single-server installation for simplicity. Spectrum Control can also be installed in a multi-server environment. For more information about Spectrum Control in a multiple-server environment, see the IBM Knowledge Center, found at:

<http://ibm.co/10KD2qP>

Different actions must be applied on separate servers, depending on the installation type that you have.

3.1.2 Assigned roles

The user role concept has not changed as of Spectrum Control V5.2.8. It has remained the same since Version 5.2.0. The number of roles were reduced since earlier versions. There are now only three roles available, as listed in Table 3-1.

Table 3-1 Roles in Spectrum Control V5.2

Role	Users that are assigned this role
Administrator	Can use every function in Spectrum Control. By default, the following groups are assigned the Administrator role: <ul style="list-style-type: none">▶ Windows: Administrators.▶ Linux: root.▶ AIX: system.
External Application	Cannot log in to the Tivoli Storage Productivity Center GUI. This role should be used for external applications that use the Tivoli Storage Productivity Center provisioning functions, such as: <ul style="list-style-type: none">▶ vSphere Web Client Extension for Tivoli Storage Productivity Center.▶ IBM SmartCloud Storage Access. IBM SmartCloud Storage Access is no longer sold, and is in maintenance mode only.
Monitor	Can log in to Tivoli Storage Productivity Center, but cannot run any function. A user with this role still can see all the information and open log files, but the only actions that this person can do are the following ones: <ul style="list-style-type: none">▶ Acknowledge alerts▶ Acknowledge a non-normal status▶ Set the tier level of a storage pool

The basic concept is still the same: You assign a role to a group of users so all users in that group can perform certain actions in Spectrum Control.

3.1.3 Single sign-on

Even when Cognos Reporting is set up by using the same authentication server (LDAP repository) that Spectrum Control uses, you still must reauthenticate when you start reporting because the Spectrum Control web-based GUI is running a different WebSphere Application Server than the Cognos component. For this reason, you must authenticate again when you access the reports for the first time during a session.

There is a way to configure the WebSphere Application Server running Cognos Reporting to accept and use single sign-on tokens from the other WebSphere Application Server so that you do not need to enter a user ID and password. This procedure is described in the IBM Knowledge Center, found at:

<http://ibm.co/1RwYBxQ>

Important: This configuration works only if you switch from the default file repository authentication to LDAP or to Local OS in the WebSphere Application Server that Cognos is using.

The single sign-on feature is not supported if you use the Element Managers of various devices that Spectrum Control manages.

3.2 Implementing authentication

This section briefly describes how to implement the different cases of authentication. It does not emphasize the whole installation it has not changed in Spectrum Control since Version 5.2.2.

3.2.1 Implementing Local OS authentication and an internal user ID

This implementation is the easiest and most often used one. Both the Spectrum Control internal user ID and users of Spectrum Control are Local OS users, typically db2admin for Windows OS or db2inst1 for Linux and AIX.

This implementation is illustrated in Figure 3-2 on page 51.

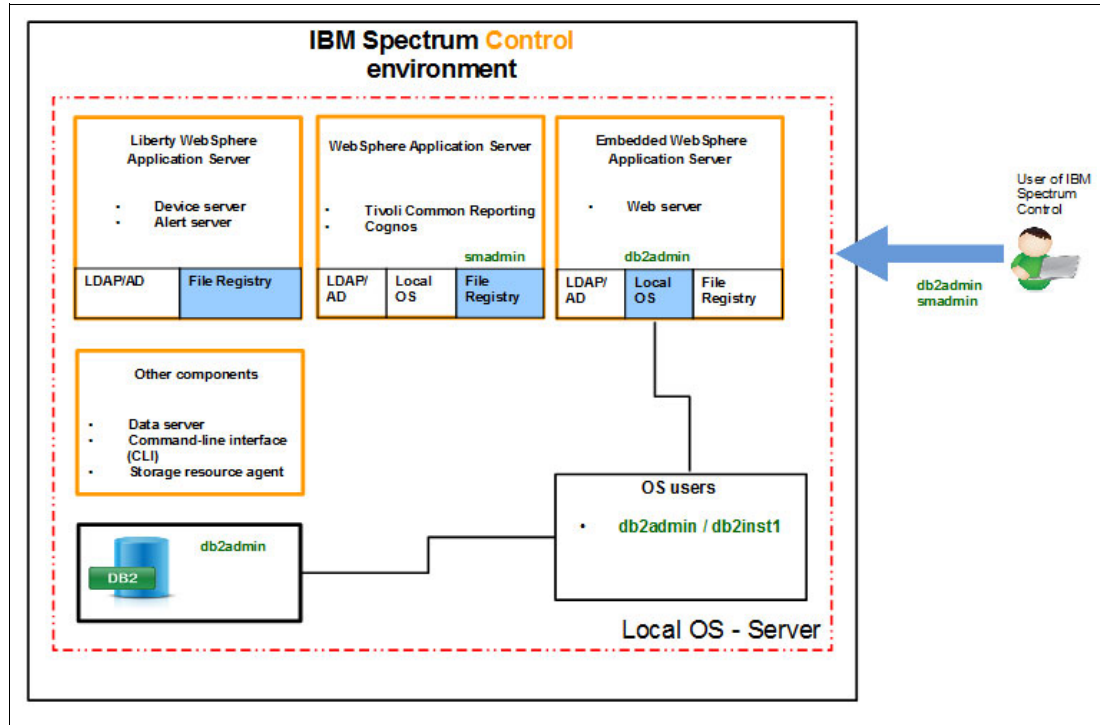


Figure 3-2 Install and perform authentication by using Local OS users

To install Spectrum Control in a single-server environment, see the IBM Knowledge Center, found at:

<http://ibm.co/1mQ1gI6>

As shown in Figure 3-2, Cognos Reporting is running under a different IBM WebSphere Application Server, which is installed by using a File Registry account. Use `smadmin` as the user ID for this File Registry account.

To learn about the installation process for Cognos Reporting, see the IBM Knowledge Center, found at:

<http://ibm.co/1ndk2XD>

Figure 3-2 shows other Spectrum Control components, for example, *Alert Server* and *Device Server*. These components are under a separate *IBM WebSphere Application Server*, which is a lightweight application server called *Liberty*. This server has a file registry as the authentication mechanism, which is not described here because it is used only for internal communication and presents no interest to external users. The same is true for the other component, such as Data Server, Storage Resource Agent, or the CLI.

The preferred practice is to install Spectrum Control and to set up authentication locally because this process enables the setting of external authentication later, as described in 3.2.2, "Implementing external authentication (direct LDAP) and a local internal user ID" on page 52.

3.2.2 Implementing external authentication (direct LDAP) and a local internal user ID

After you install Spectrum Control by using Local OS accounts, the internal user ID stays as-is. However, the users of Spectrum Control want to use their own centralized user account to log in to a web-based GUI, as shown in Figure 3-3.

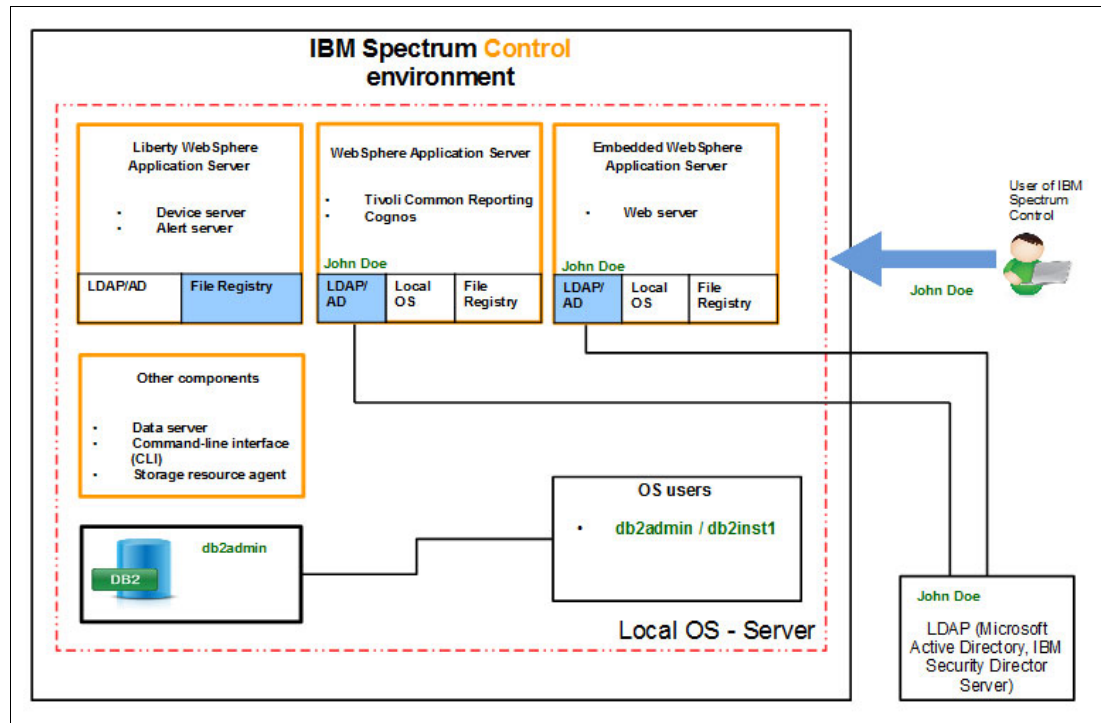


Figure 3-3 Install by using a Local OS user and perform authentication by using external users

This is the preferred method of implementation. It has the greatest flexibility in terms of external repositories by because you can use any LDAP directory server that supports LDAP V3. A comprehensive list can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg27036471>

Using this implementation, the Spectrum Control server can be installed on any supported platform, whether it is Linux, AIX or Windows.

This implementation is referred to as *Method 1* in the repository combination schematic in Figure 3-1 on page 48.

One of the drawbacks of Method 1 is that the administrator of the Spectrum Control must change the authentication of the following items:

- ▶ Web server: You must change WebSphere Application Server from Local OS authentication to the LDAP repository authentication.
- ▶ Cognos Reporting: You must change WebSphere Application Server from File Registry authentication to LDAP repository authentication.

The procedures are similar and are described in the IBM Knowledge Center, found at:

<http://ibm.co/1RwZTs0>

Important: Before applying this method of authentication, you must ensure that there are no duplicated user names or group names in the local file repository, the operating system repository, or the LDAP repository.

In all the cases where external authentication is used, make sure that the firewall is configured so that it can communicate between all the components.

For a complete list of ports that are used by Spectrum Control, see the IBM Knowledge Center:

<http://ibm.co/1K95tcz>

A use case of this implementation is described in 3.3, “Use case: Configuring Spectrum Control to authenticate with Bluepages for an intranet user” on page 55.

3.2.3 Implementing external authentication (implicit) and a local internal user ID

This implementation is referred as *Method 2* in the repository combination schematic in Figure 3-1 on page 48.

Using this implementation, the Spectrum Control server can be installed only on a Windows platform, which limits the platform flexibility. The LDAP repository type can be AD only, so this is not a preferred practice. Use this method only when the Spectrum Control server must be part of the domain.

This implementation is illustrated in Figure 3-4.

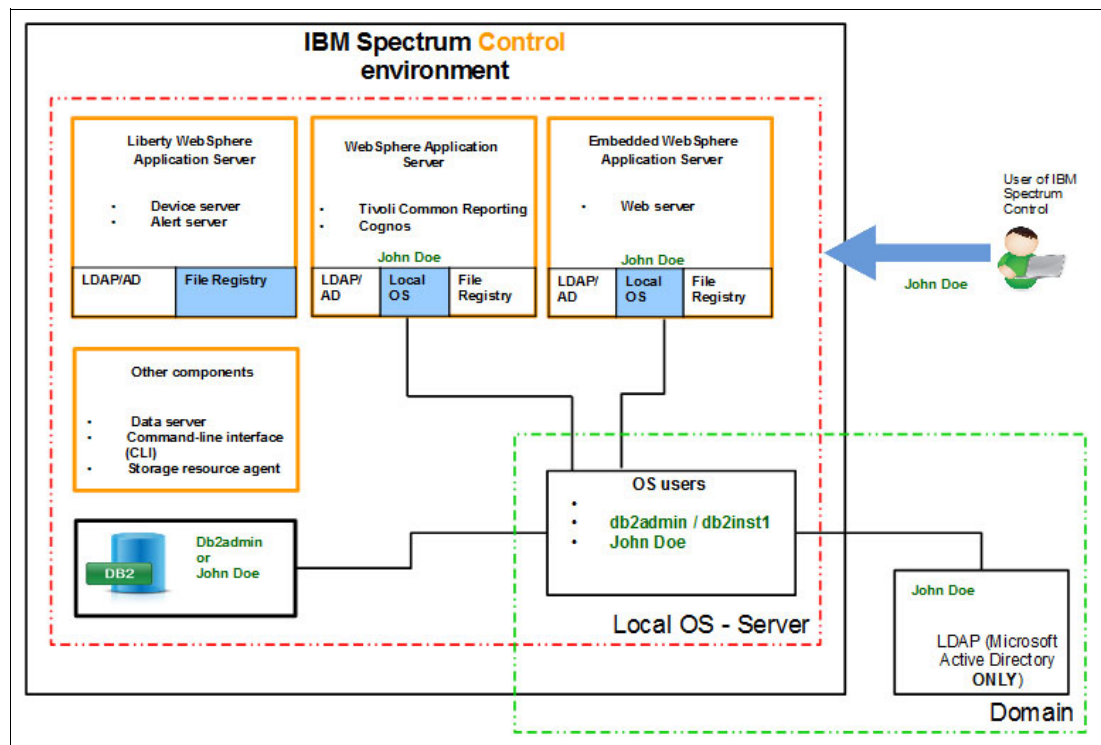


Figure 3-4 Install by using a Local OS user and perform authentication by using external users

Note: Implementing direct communication with LDAP as described in 3.2.2, “Implementing external authentication (direct LDAP) and a local internal user ID” on page 52 can be done even in the cases when the Spectrum Control server is part of a domain. It is a preferred practice rather than using implicit authentication.

The advantage of using this implementation is that there is no need to configure the web server. The authentication can be left on, and as-is, by using Local OS because technically the authentication is handed over to the OS to resolve with AD.

The disadvantage of using this implementation is that for Cognos Reporting, the File Registry must be changed to Local OS. This procedure is described in *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

In addition to the WebSphere Application Server change, the server must be added into the domain. For more information, see the IBM Knowledge Center:

<http://ibm.co/22VndE3>

Also, the computer browser service must be enabled and running on the server. To verify this setting, see the IBM Knowledge Center:

<http://ibm.co/10dPT3p>

Note: The installation of the Spectrum Control server and DB2 must be accomplished by using the local users ID, as described in 3.2.1, “Implementing Local OS authentication and an internal user ID” on page 50.

3.2.4 Implementing external authentication and a domain internal user ID

This implementation is the most complex. It is the closest to the implementation that is described in 3.2.3, “Implementing external authentication (implicit) and a local internal user ID” on page 53. The only major difference is that the internal user ID must be a domain user. This method is supported only on AD, which limits LDAP repository support.

Figure 3-5 on page 55 shows this implementation.

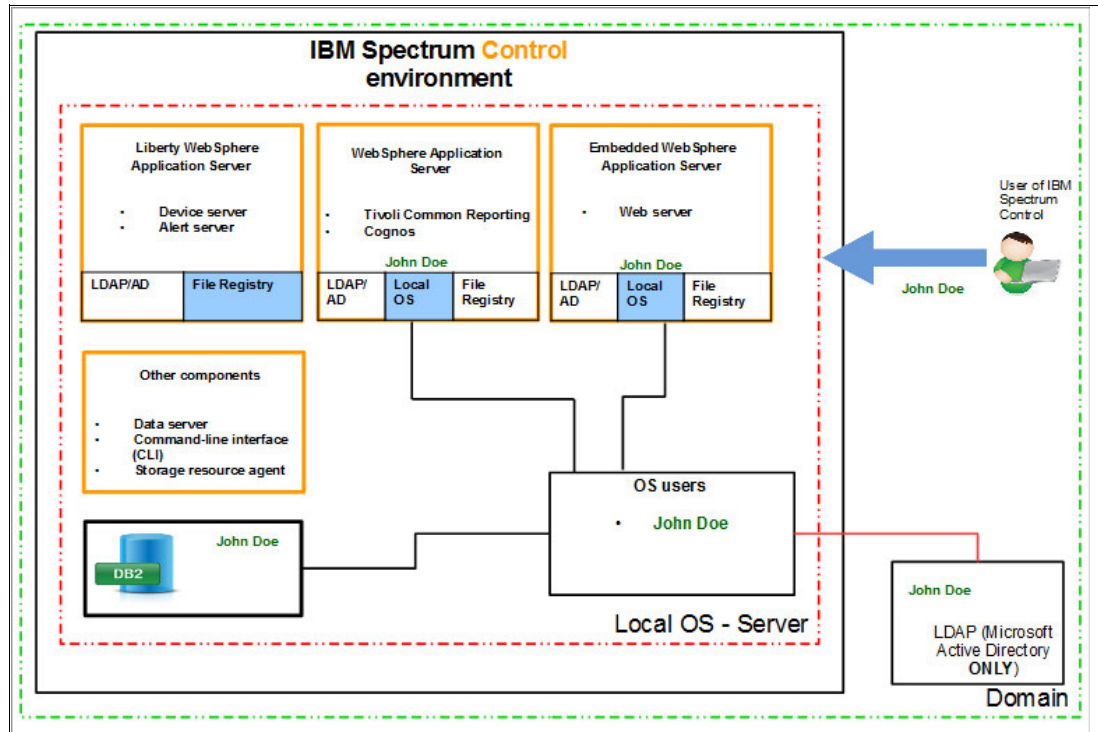


Figure 3-5 Install by using a domain OS user and authentication by using external users

This method also has all the advantages and disadvantages described in 3.2.3, “Implementing external authentication (implicit) and a local internal user ID” on page 53. However, regarding the requirements that are related to the domain, there are a couple of additional items. These items are described in the IBM Knowledge Center, found at:

<http://ibm.co/1mQmjay>

For installation details, see the *Installing DB2 and Spectrum Control by using domain user accounts* topic in the IBM Knowledge Center, found at:

<http://ibm.co/1Znna4x>

3.3 Use case: Configuring Spectrum Control to authenticate with Bluepages for an intranet user

This use case describes how to configure authentication and authorization to allow IBM intranet users (Bluepages) to log in to the Spectrum Control web-based GUI.

The IBM LDAP repository is powered by the IBM Security Directory Server, which provides a trusted identity data infrastructure for authentication, with the following benefits:

- ▶ Standardized architecture and broad platform support for a large range of operating system and heterogeneous environments
- ▶ Scalability, availability, and flexibility to support hundreds of millions of entries
- ▶ Easy user and group management tool
- ▶ Robust auditing and reporting

As method of implementation, we use the preferred practice that is described in 3.2.2, “Implementing external authentication (direct LDAP) and a local internal user ID” on page 52.

This use case does not describe how to set up the IBM Security Directory Server. In this example, we use a test instance of the real IBM intranet (Bluepages). For more information about how to install and configure IBM Security Directory Server as an LDAP repository, see the following resources:

- Installation:
<http://ibm.co/1PSewmE>
- Configuration:
<http://ibm.co/1UJh6MZ>

If your organization uses IBM Security Directory Server, you must contact the administrator to provide you with the configuration information of the server. This information is required for the upcoming steps.

Note: Some of the steps in this section are optional. Skip them if they are not applicable to your environment.

Complete the following steps:

1. Open the admin console for WebSphere Application Server.
2. After you log in to Spectrum Control with Local OS credentials, click **Settings** → **User Management**. Click the **Modify Authentication mechanism** link to start the administration console (Figure 3-6).

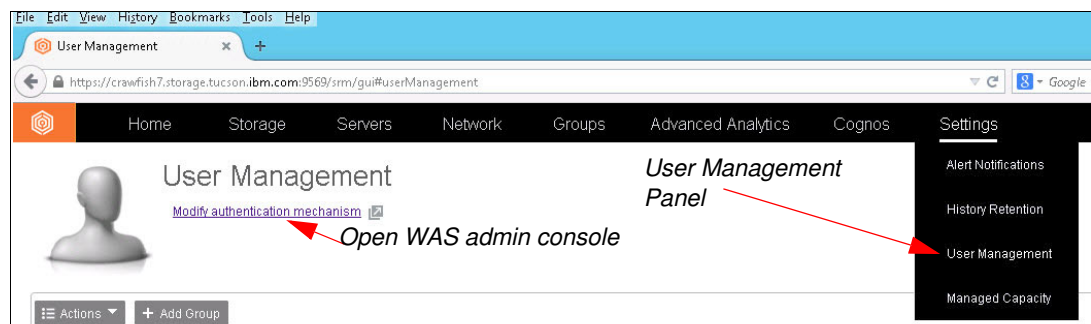


Figure 3-6 User Management window using Local OS authentication

3. The WebSphere Application Server admin console opens. Provide the File Registry / Local OS credentials (Figure 3-7 on page 57).

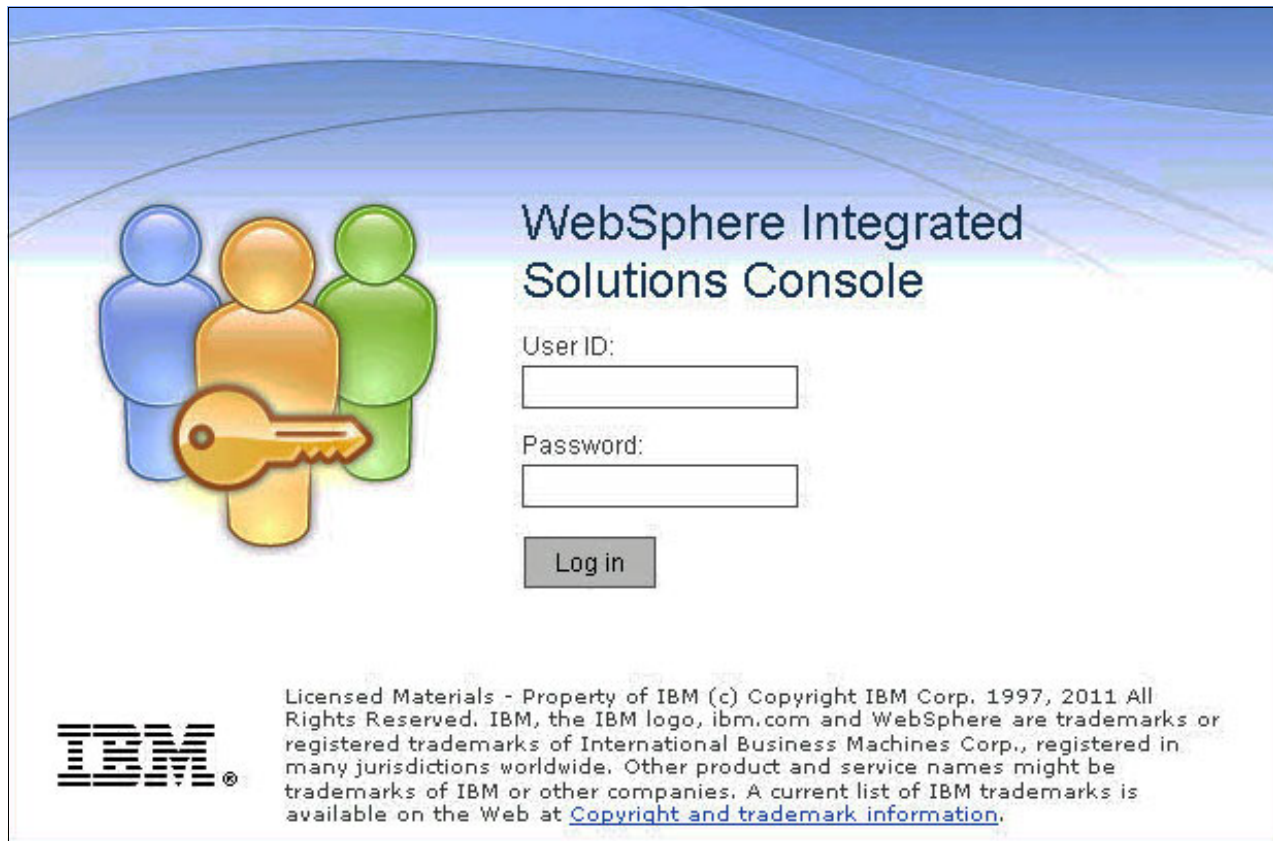


Figure 3-7 WebSphere Application Server admin console login window

3.3.1 Adding a Bluepages SSL certificate

This step is optional and must be done only if your IBM Security Directory Server is configured with SSL. Complete the following steps:

1. After you log in, from the left menu click **Security** → **SSL certificate and key management**.
2. From Related Items links, click **Key stored and certificates** (Figure 3-8).

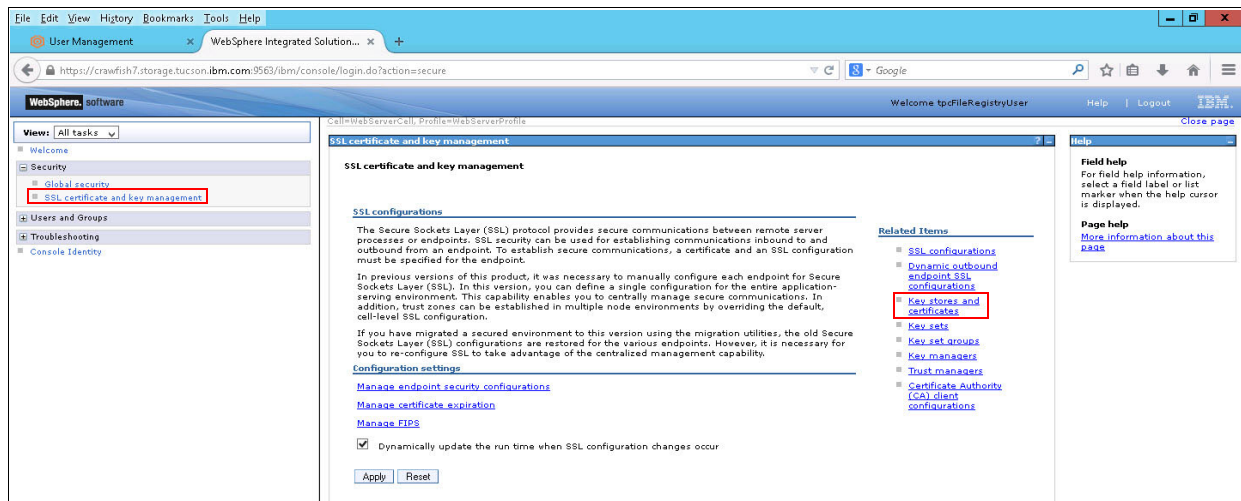


Figure 3-8 SSL certificate and key management - keystores and certificates window

3. From the Keystores and certificates window, click **NodeDefaultTrustStore** (Figure 3-9).

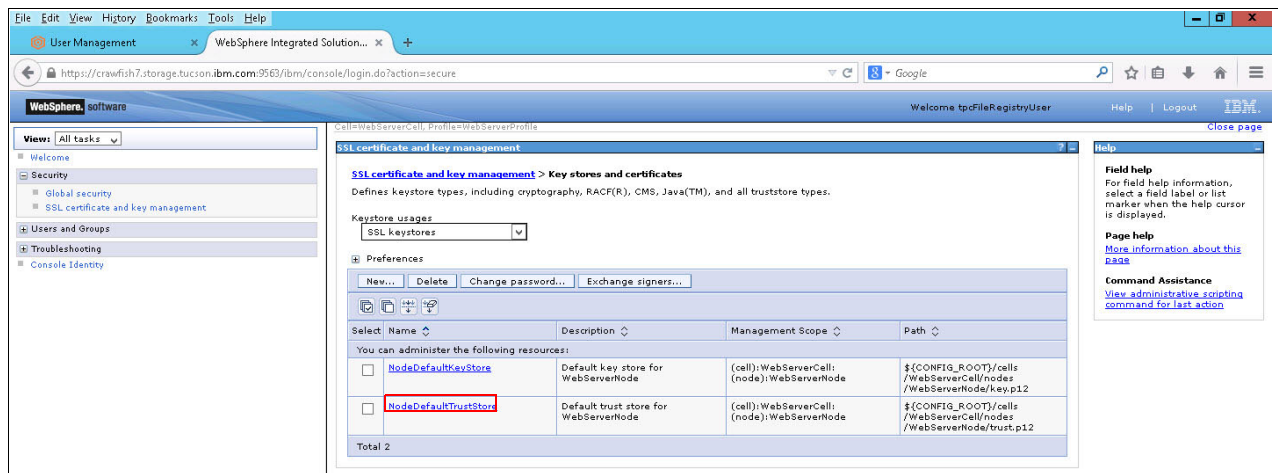


Figure 3-9 NodeDefaultTrustStore

4. In the Additional Properties window, click **Signer certificate** (Figure 3-10).

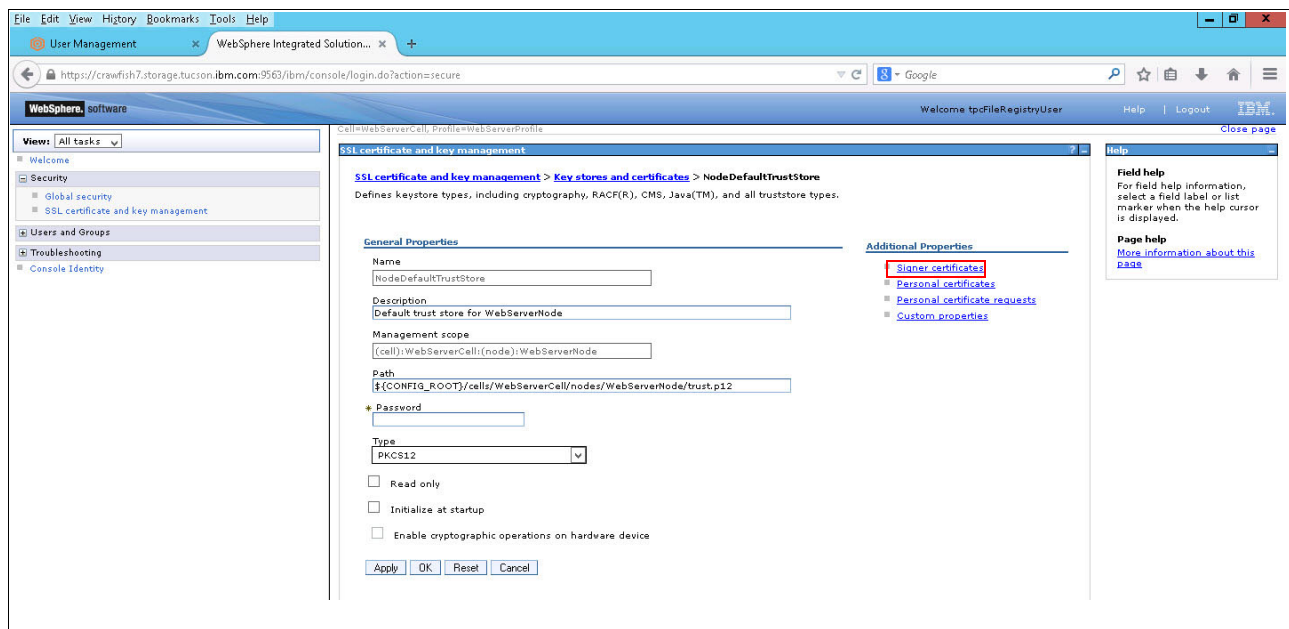


Figure 3-10 Signer certificate

5. In the next window, click **Retrieve from port** (Figure 3-11 on page 59).

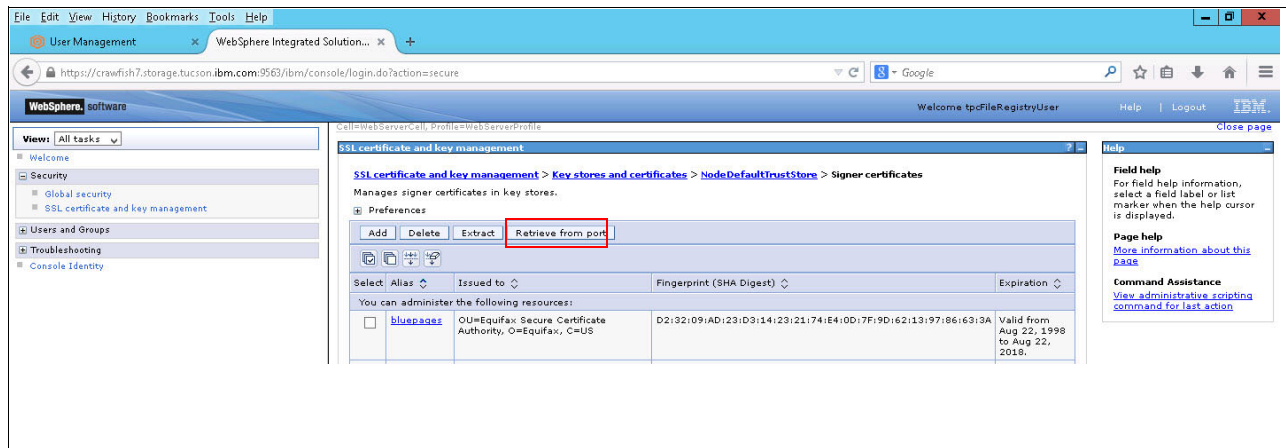


Figure 3-11 Retrieve from port

- Complete the Properties fields with the details of your IBM Security Directory Server. In this use case, we use our test Bluepages directory server. Click **Retrieve signer information** after the properties are entered (Figure 3-12).

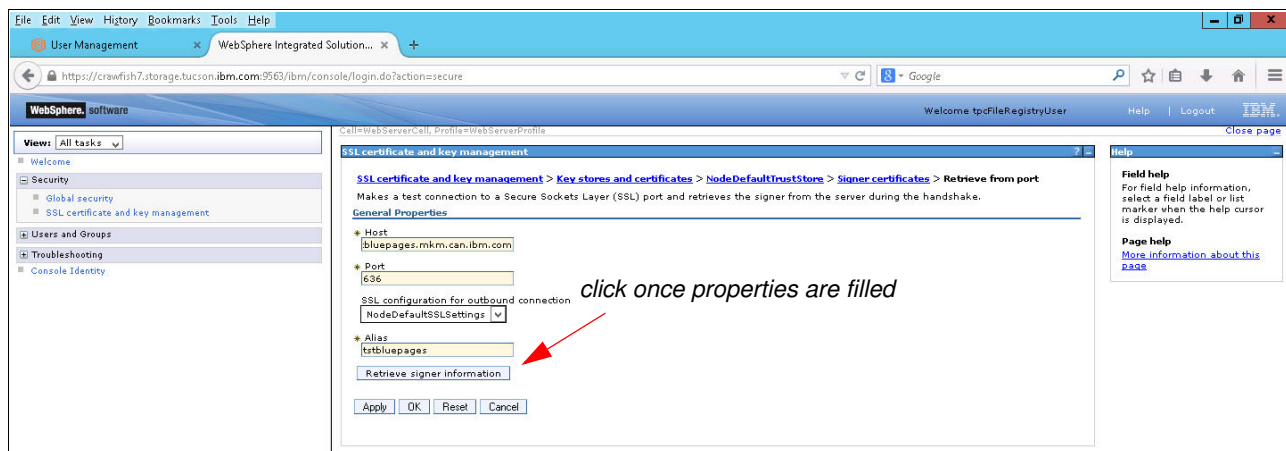


Figure 3-12 Retrieve from port properties window

7. The certificate is retrieved. You are prompted to save the changes directly to the master configuration after you click **Apply** (Figure 3-13).

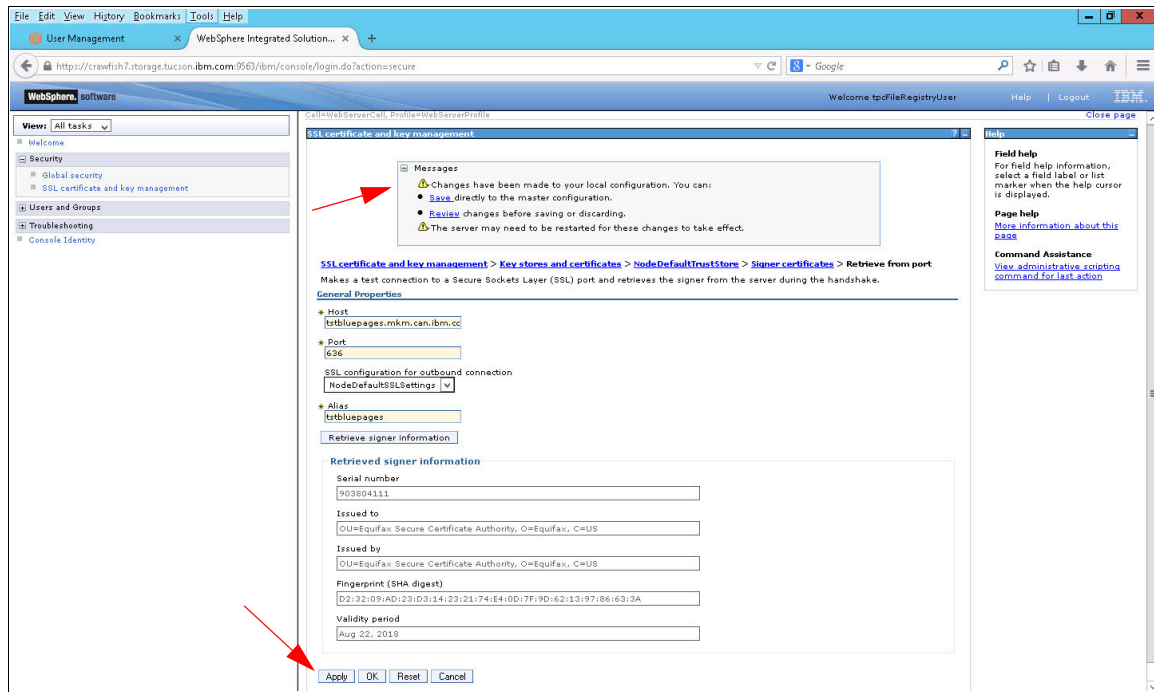


Figure 3-13 Retrieve the certificate and save changes

8. After the certificate is retrieved and the changes are saved, make sure that this certificate shows up in the table, as shown in Figure 3-14.

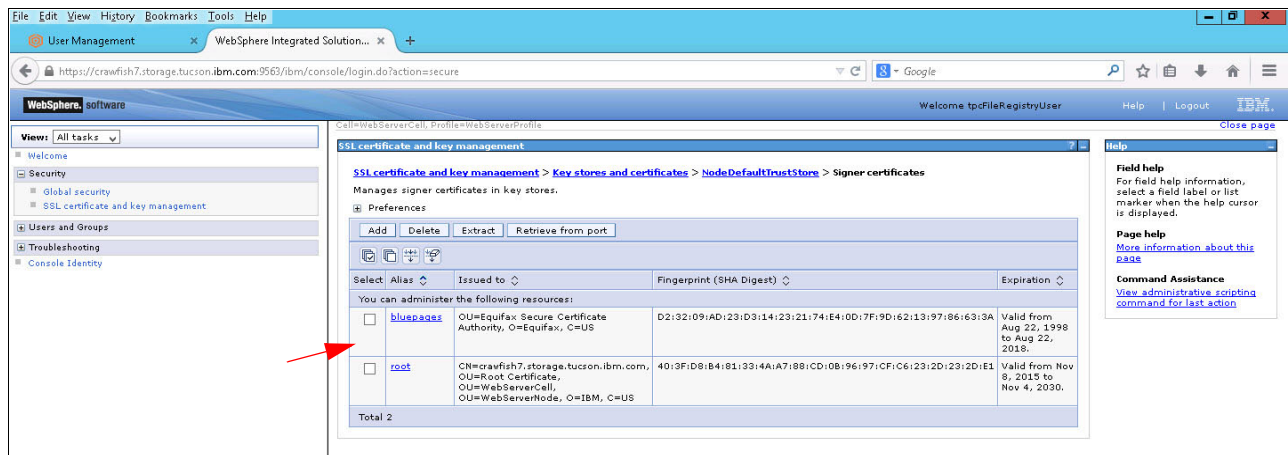


Figure 3-14 Certificate shows up in the Signer certificate table

9. Restart the WebSphere Application Server by using the scripts that are provided in Spectrum Control under <IBM Spectrum Control install location>/scripts:
- stopTPCWeb.bat or .sh
 - startTPCWeb.bat or .sh

3.3.2 Configuring the federated repository

This step is mandatory. Depending on whether IBM Security Directory Server is configured with or without SSL, different ports are used and there is also an additional check box that you must select when using SSL.

The following steps assume that you already logged in to the WebSphere Application Server admin console:

1. Click **Security** → **Global security**.
2. In the **Available realm definitions** section, click **Configure**, as shown in Figure 3-15.

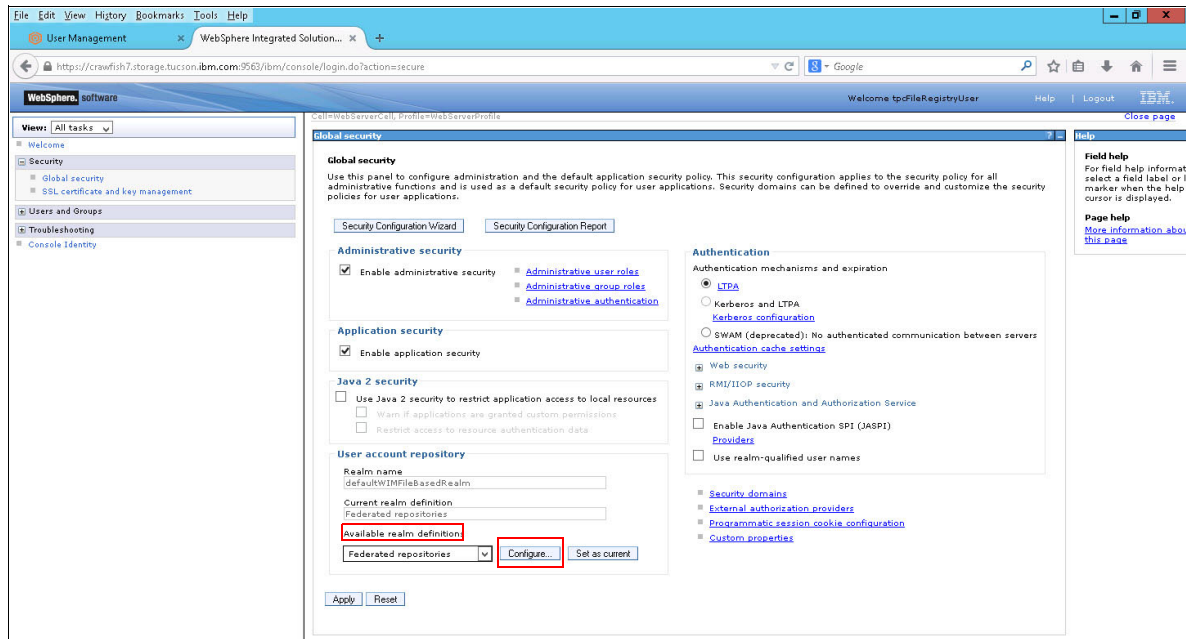


Figure 3-15 Available realm definitions

3. In the **Related Items** section, click **Manage Repositories** (Figure 3-16).

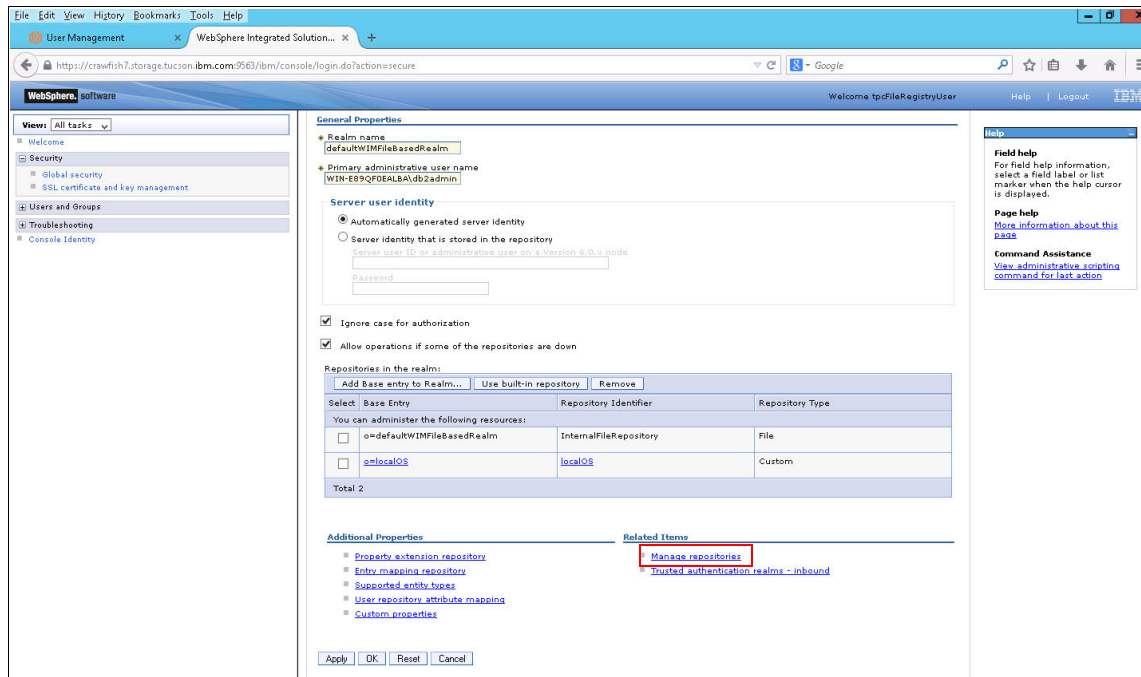


Figure 3-16 Managed repositories

4. From the Managed repositories tables, click **Add** → **LDAP repository** (Figure 3-17).

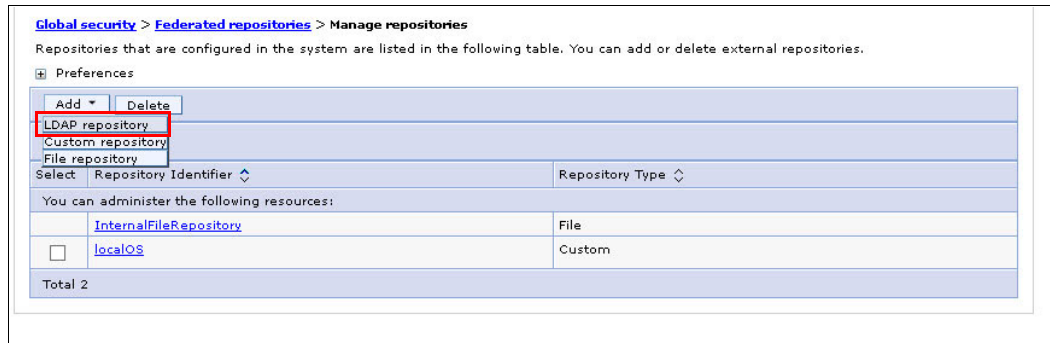


Figure 3-17 Managed repositories - add an LDAP repository

5. You must enter the general properties of the LDAP repository, as shown in Figure 3-18 on page 63. This information typically is provided by the IBM Security Directory Server administrator. In this use case, the primary host name is `tstbluepages.mkm.can.ibm.com` and the port is 636 for SSL. Make sure that the **Require SSL communication** check box is selected if SSL is used; clear this check box for non-SSL communication. The Login properties field also must be changed from `uid` to `mail` if you want the email to be used (instead of the serial number) for authentication.

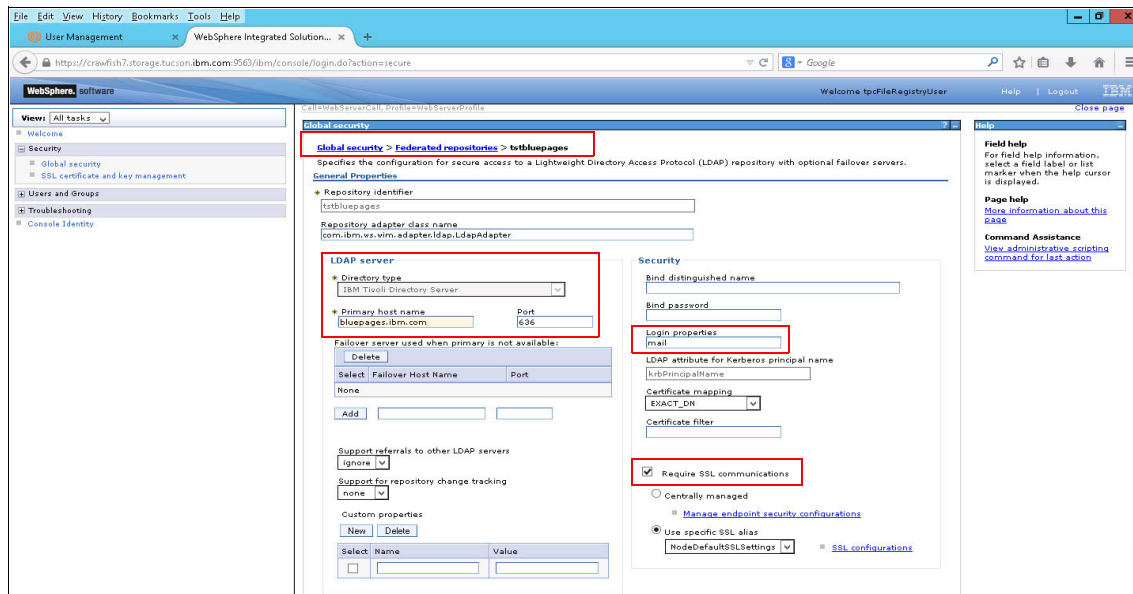


Figure 3-18 Configuring Bluepages as a federated repository

6. At the bottom of the window, click **Apply**. You are prompted to save the changes to the master repository. Click **Save**.
7. While you are in the Manage repositories window, as shown in Figure 3-18, click **tstbluepages**, and in the Additional Properties section, click **LDAP entity types**, as shown in Figure 3-19.

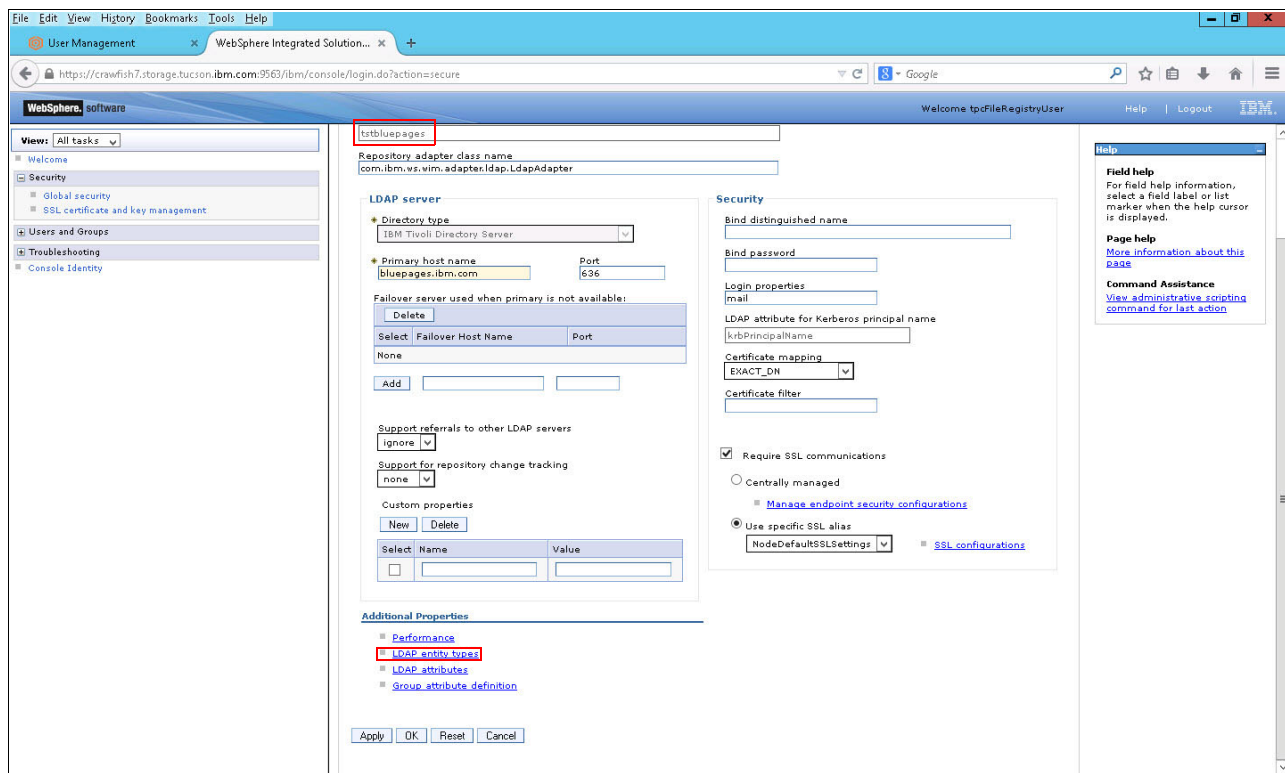


Figure 3-19 LDAP entity types

8. Change the Group entity type to groupOfUniqueNames and the PersonAccount entity type to ibmPerson, as shown in Figure 3-20.

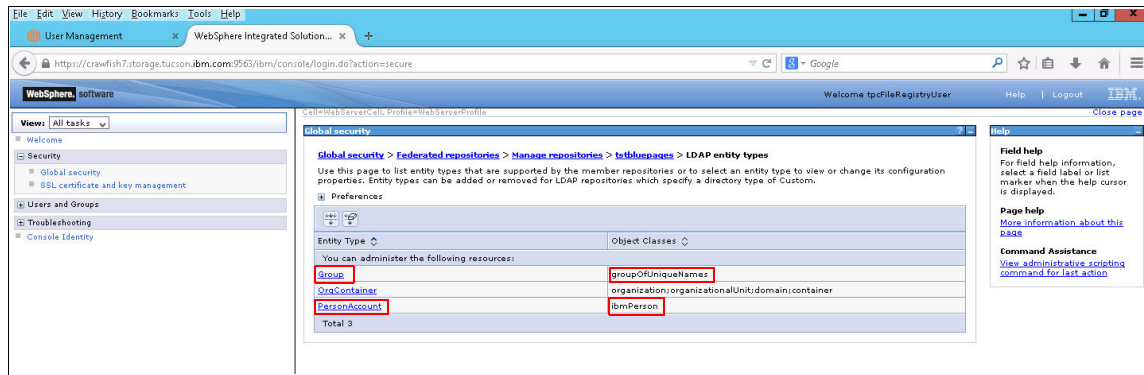


Figure 3-20 Group and PersonAccount entity type changes

9. You are prompted to save the changes to the master configuration. Click **Apply** for all the updates and save these changes.
10. Go to **Federated repositories** and click **Add Base entry to Realm** (Figure 3-21).

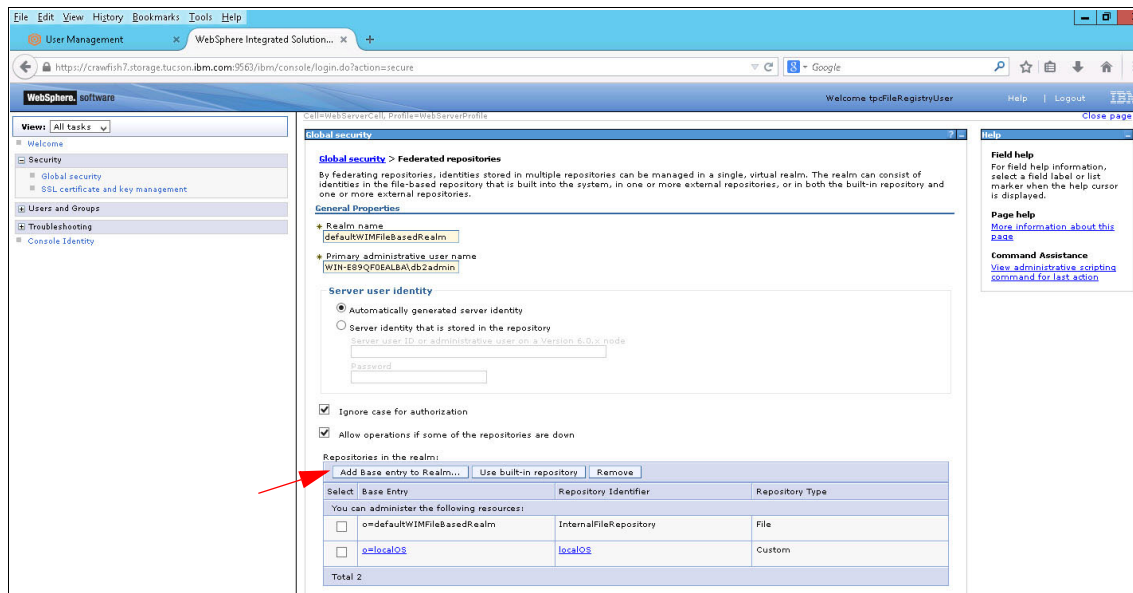


Figure 3-21 Add Base entry to Realm

11. In the Distinguished name of base entry field, enter o=ibm.com (Figure 3-22 on page 65).

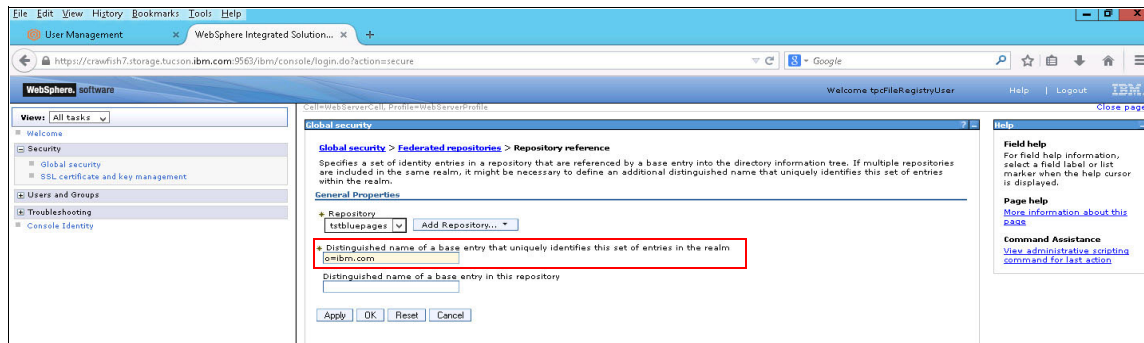


Figure 3-22 Distinguished name of a base entry

12. Go back to **Manage repositories**, and in the **Additional Properties** section, select the **Group attribute** definition for this newly defined repository. Click **Member attributes** (Figure 3-23).

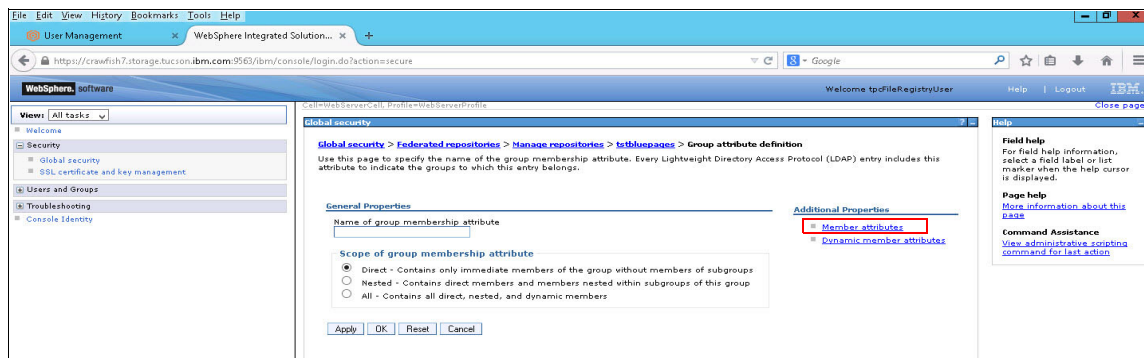


Figure 3-23 Group attribute definition

13. You must update the member to contain the Object Class: groupOfNames and add a property into the table uniquemember, which contains the Object class : groupOfUniqueNames (Figure 3-24). These values are provided by your IBM Security Directory Server administrator.

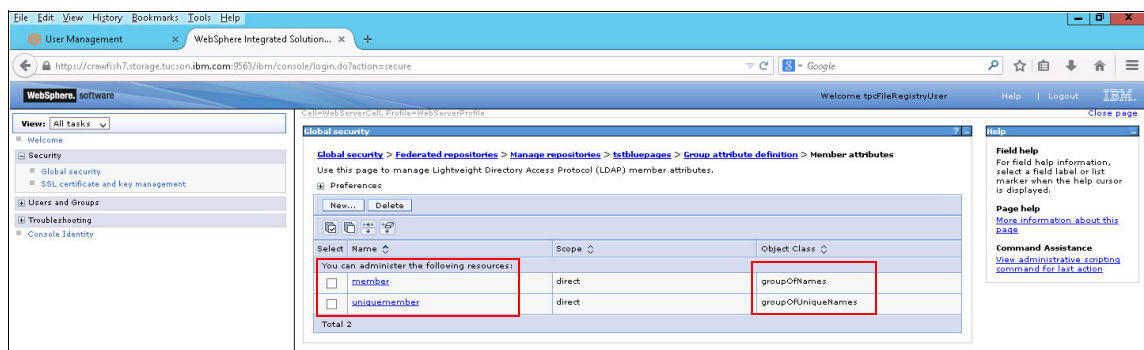


Figure 3-24 Member attributes window

14. Click **Apply** and save the changes to the master configuration.

15. Restart WebSphere Application Server by using the scripts that are provided in Spectrum Control under <IBM Spectrum Control install location>/scripts:

- stopTPCWeb.bat or .sh
- startTPCWeb.bat or .sh

3.3.3 Configuring web security

This step is another optional step. You need to complete it only if your IBM Security Directory Server is configured with SSL. Complete the following steps:

1. Click **Security** → **Global security**.
2. Click **Web security** → **General settings**, as shown in Figure 3-25.

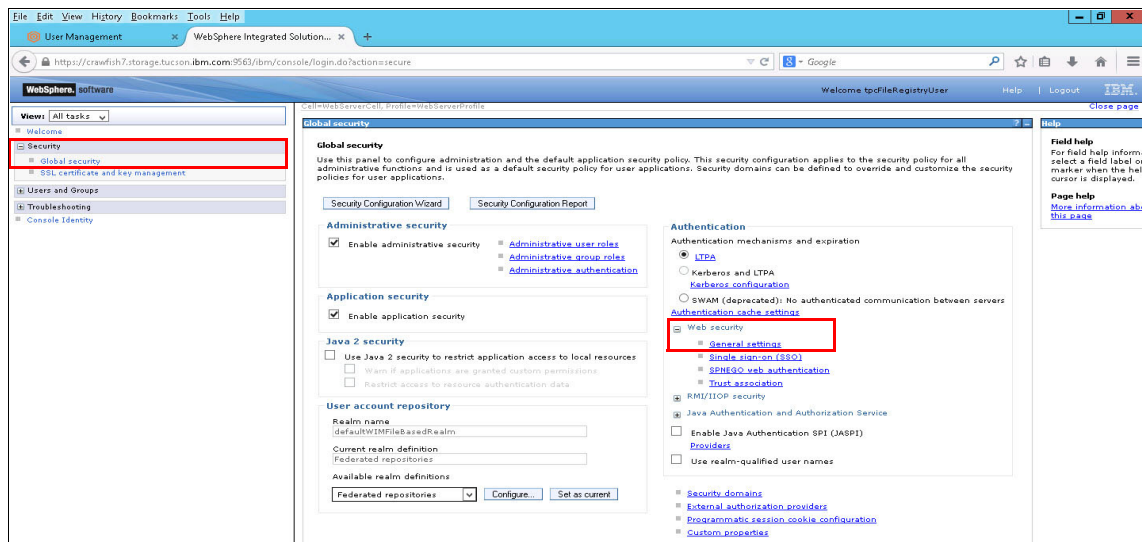


Figure 3-25 Web security - General settings

3. Make sure that authentication is selected only when the URI is protected (Figure 3-26).

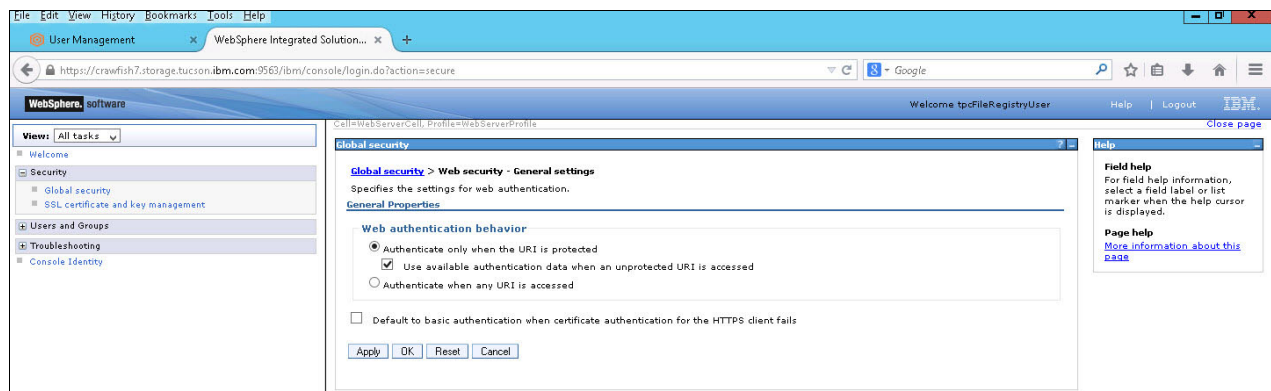


Figure 3-26 General Properties for web security general settings

4. Click **Apply** and save the changes to the master configuration file.

3.3.4 Performing a test communication with Bluepages

After the IBM Security Directory Server (LDAP repository) is added to the WebSphere Application Server federated repositories, you can list the LDAP repository users in the WebSphere Admin Console under **Users and Groups** → **Manage Users**, as shown in Figure 3-27. There is a search function available that you can use to find your user in the LDAP repository.

If WebSphere Application Server could not successfully connect to the LDAP repository, the users are not listed.

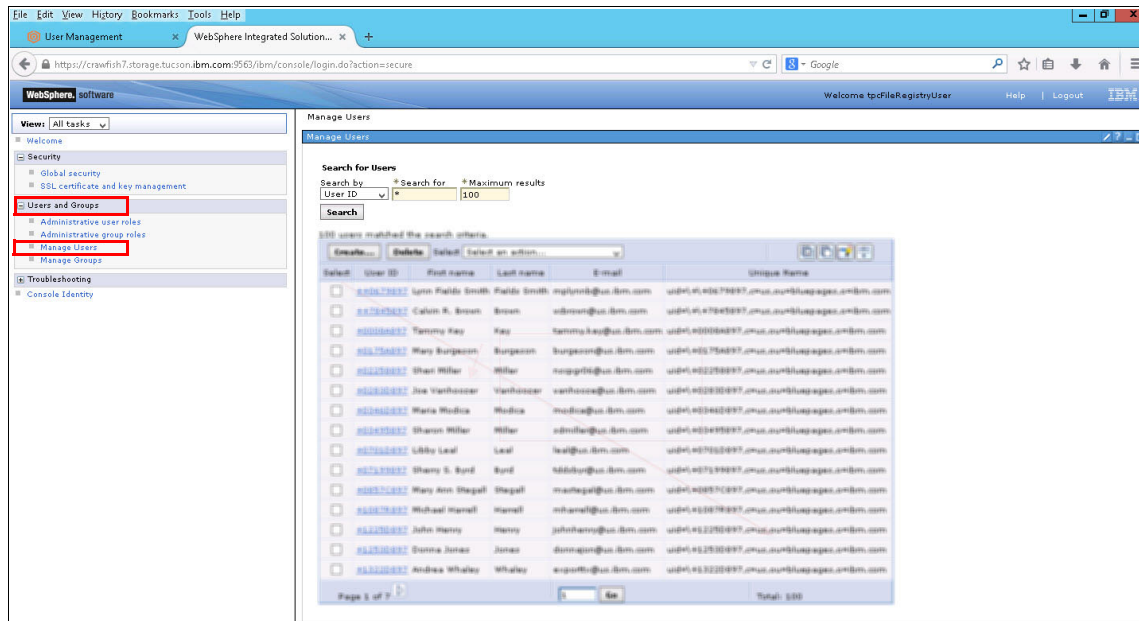


Figure 3-27 LDAP repository users that are listed in the WebSphere Application Server admin console

Now, you can log out of the WebSphere Application Server admin console and use the Bluepages (intranet) user account to log in to the admin console (Figure 3-28).



Figure 3-28 Log in with Bluepages credential

3.3.5 Adding a specific Bluegroup to Spectrum Control role-based authorization

It is not yet possible to log in to Spectrum Control by using the Bluepages (intranet) account. There is one more step that must be completed in the Spectrum Control web-based GUI, which is to map the LDAP group to an authorization level for Spectrum Control.

For more information about role-based authorization, go to the IBM Knowledge Center, found at:

<http://ibm.co/1N31MVM>

Complete the following steps:

1. Using the Bluegroup administration tool, create a group to be used by Spectrum Control. In this use case, we name it SpectrumControl_LDAP.

This administration tool is a web interface for the IBM Security Directory Server. A group can be created on the server itself by using other methods, such as the command line. Make sure that your user is part of this Bluegroup (Figure 3-29).

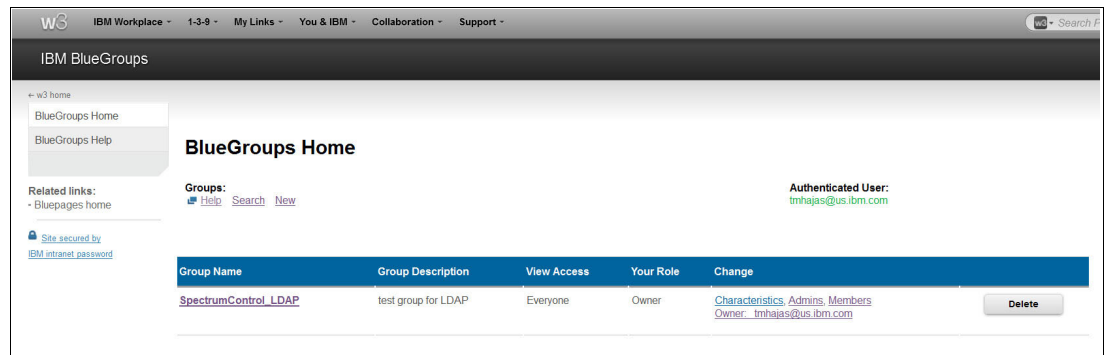


Figure 3-29 BlueGroups administration application

2. After the LDAP repository is added to the federated repository in the WebSphere Application Server admin console, the LDAP users can be listed under Users and Groups and Manage Users, as shown in Figure 3-27 on page 67. The Spectrum Control User Management window should show the Bluegroup that was created earlier.
3. Log in to the Spectrum Control web-based GUI by using a local OS account, click **Settings** → **User Management**, and select **Add Group** (Figure 3-30).

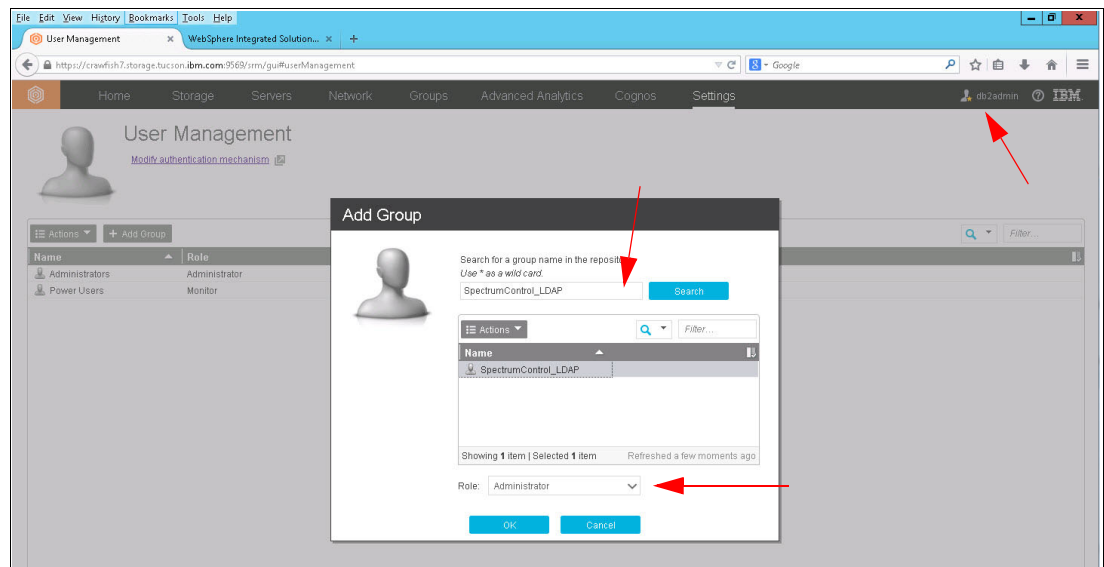


Figure 3-30 Add the Bluegroup to role-based authorization

4. After the Bluegroup is added to the Spectrum Control authorization, log out of the local OS user and log in by using the Bluepages (intranet) user ID (Figure 3-31 on page 69).

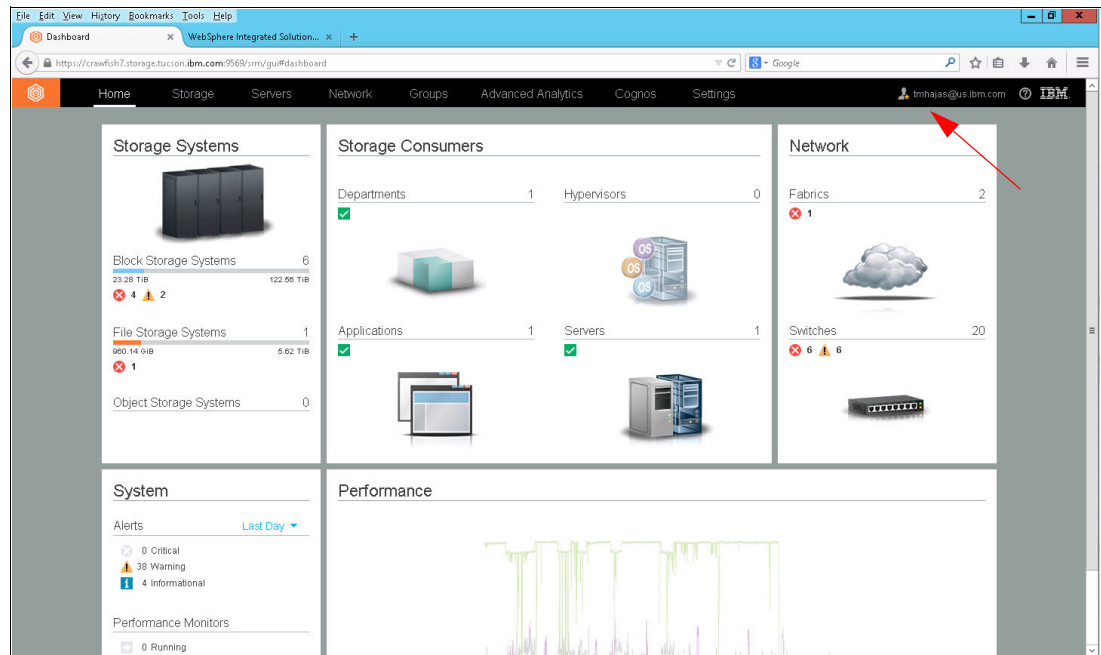



Figure 3-31 Log in to Spectrum Control by using the Bluepages (intranet) account

Now, you can log in to Spectrum Control by using your Bluepages (intranet) credentials, but you still cannot log in to Cognos Reporting with these credentials. As described in 3.2.2, “Implementing external authentication (direct LDAP) and a local internal user ID” on page 52, the WebSphere Application Server of Cognos Reporting must undergo the same procedures as described in this use case.

Important: Using the single sign-on function from Spectrum Control to Cognos Reporting does circumvent the necessity of configuring direct LDAP for Cognos Reporting.

To access the WebSphere Application Server admin console of Cognos Reporting, use the following link, where the default port number is 16316:

`https://host_name:port/ibm/console/logon.jsp`



Configuration and administration considerations, tasks, and use cases

This chapter describes the tasks that you must do to configure your Spectrum Control server to monitor, manage, and report on devices in your environment.

Additionally, this chapter provides administrative use cases for activities such as the following ones:

- ▶ Monitoring the Spectrum Control dashboard
- ▶ Cleaning up alerts, errors, and warnings in the dashboard
- ▶ Viewing inter-switch connections and their performance metrics
- ▶ Viewing connections for N_Port Virtualization (NPV) mode switches

4.1 Adding resources to Spectrum Control

One of your first actions after you complete a fresh installation is to add devices to Spectrum Control. The Add Device wizard in the Spectrum Control V5.2 GUI provides a simplified way for adding, discovering, and configuring devices.

The options to add a device depend on the type:

- ▶ To add storage systems, click **Storage** → **Block Storage Systems, File Storage Systems, or Object Storage Systems**. The types of storage systems that you can choose depends on the menu selection.
- ▶ To add fabrics and switches, click **Network** → **Fabrics** or **Network** → **Switches**.
- ▶ To add VMware Hypervisors or vCenters, click **Servers** → **Hypervisors**.
- ▶ To add Storage Resource agents, click **Servers** → **Servers**.

The following additional resources can be configured and added to Spectrum Control:

- ▶ To add agentless servers, click **Servers** → **Servers**.
- ▶ To add applications, click **Groups** → **Applications**.
- ▶ To add departments, click **Groups** → **Departments**.

Note: The configuration is performed in the web-based GUI for all device types. The stand-alone GUI from previous versions of Tivoli Storage Productivity Center no longer exists. For more information about the stand-alone GUI removal, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1SJHV1H>

The Add Device wizard has three main steps:

1. Select a resource.
2. Add and discover a device.
3. Configure a device.

Examples are provided in the subsequent sections for adding different device types. If you need additional details for configuring and administering the different device types as resources, see the IBM Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1Zghrs0>

4.1.1 Adding IBM storage systems

Most storage systems from IBM are added by using the native API instead of a CIM agent. The type of information that you must provide to add the device to Spectrum Control depends on the device type.

This section describes the tasks to add a SAN Volume Controller to Spectrum Control. The tasks are similar for other IBM storage systems

Complete the following steps:

1. Check the support matrix for the required level of firmware before you begin adding a device at the following website:

<http://www.ibm.com/support/docview.wss?&uid=swg21386446>

From the support matrix page, click **Spectrum Control 5.2.8+** in the column labeled Storage. A window with the list of supported storage systems and the supported firmware levels opens.

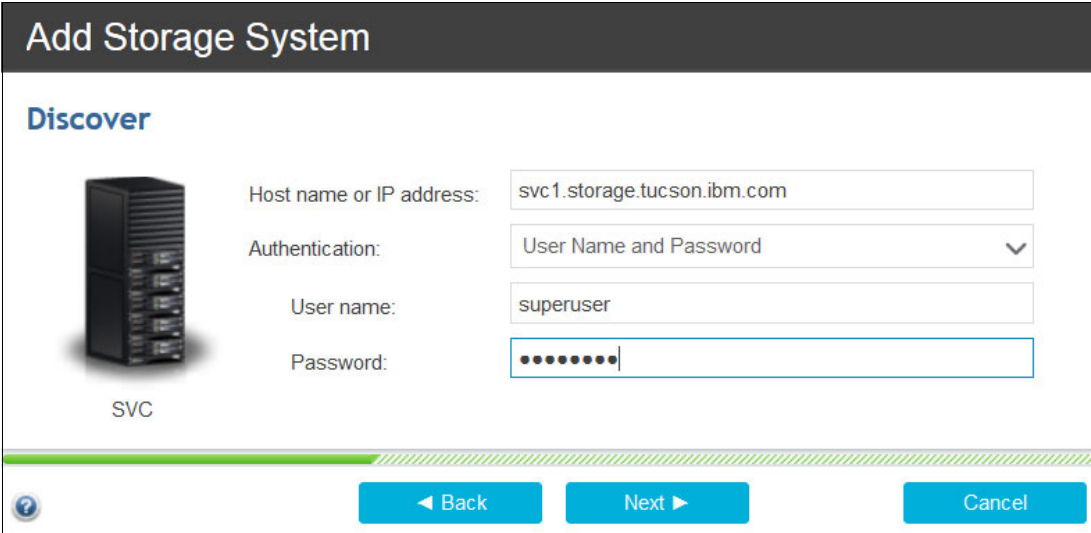
2. After verifying the required firmware levels, go to the Spectrum Control GUI to add the SAN Volume Controller. Click **Storage Resources** → **Block Storage Systems**. From the list of configured Block Storage Systems, click **Add Storage System**.
3. Select the type of storage system to add (see Figure 4-1). IBM devices are listed individually. “All others” should be used for other vendors.



Figure 4-1 Select the type of device for adding block storage systems


Tip: If you do not see the IBM device that you want, make sure that you selected the correct type of storage (block, file, or object) from the menu.

4. Clicking the SAN Volume Controller icon opens the Add Device wizard, as shown in Figure 4-2.



Add Storage System

Discover

 SVC

Host name or IP address:

Authentication:

User name:

Password:

Figure 4-2 Specify connection details

5. Complete the connection information so that Spectrum Control can discover the devices that are available for this data source.

Note: Spectrum Control can add a SAN Volume Controller or Storwize V7000 server by using only a user ID and password instead of Secure Shell (SSH) keys for SAN Volume Controller Version 6.3 or later, which makes adding the devices simple. If a customer has strict rules about how often a password must be changed, using SSH keys might be the better option.

Tip: If you plan to use the provisioning function, or if you have more than one Spectrum Control server in your environment, consider using dedicated user IDs and SSH keys per Spectrum Control server. This way, you can easily perform the following tasks:

- ▶ Delete the ID or change the password to prevent a Spectrum Control server from accessing a device.
- ▶ Identify the commands that are run by a Spectrum Control server by reviewing the audit log files of a device.

6. When you have entered the required information, click **Next**. Spectrum Control attempts to connect to the device.
7. After Spectrum Control successfully connects, you see configuration options similar to what is shown in Figure 4-3 on page 75.

Figure 4-3 Initial device configuration

8. When you have provided all the information and adjusted the scheduling based on your needs, click **Configure**.

Performance data collection begins after a successful initial probe.

Notes:

The 1-minute performance monitoring interval facilitates finer granularity for troubleshooting, debugging, and reporting.

Be sure to consider the performance and space usage when running 1-minute performance monitoring on your server. The space usage is DB2 space usage. The size of the data that is collected is related to the number of objects that the devices contain (for example, LUNs and disks).

Here are some limitations of 1-minute performance monitoring:

- ▶ The 1-minute intervals show only in a chart if the complete time range has 1-minute intervals; otherwise, it “falls back” to 5-minute intervals.
- ▶ In Cognos, the 1-minute intervals are not available.
- ▶ The TPC exposed views do not support 1-minute intervals.

The alerting is evaluated based on 1-minute intervals, so you might need to think about adjusting your alert thresholds (no 5-minute averages, so the threshold might need to be set higher, with longer intervals between repeating alerts).

For the full set of administrative actions for storage systems, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1mQnGX7>

4.1.2 Adding storage systems from other vendors

Non-IBM storage devices are typically connected to Spectrum Control by using a CIM agent. Complete the following steps:

1. From the Add Storage System window that is shown in Figure 4-1 on page 73, click the **All others** device icon.
2. Specify the connection information for your device type.
3. Complete the configuration steps in step 7 on page 74.

For the full set of administrative actions for storage systems, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1mQnGX7>

4.1.3 Adding fabrics and switches

SAN switch and fabric configuration options differ by switch vendor. The Add Device wizard guides you through choosing the correct options:

SNMP Agents	SNMPv1 or SNMPv3 can be used to communicate with all vendor switches.
CIM Agents	Spectrum Control can communicate with CIM Agents for all switches except Cisco. The CIM Agent is the preferred method for Brocade switches.

Tip: The CIM agent is the preferred method for Brocade switches because it collects the richest set of data. If both agent types are configured, Spectrum Control uses the CIM Agent to collect data.

The SNMP agent is the only method available for Cisco switches and provides full functionality.

Note: The Storage Resource agent does not collect fabric information.

The Spectrum Control IBM Knowledge Center explains what information is collected by each agent for each switch vendor. It can be found at the following website:

<http://ibm.co/1K98km1>

Adding switches to Spectrum Control

To add one or more switches to Spectrum Control, complete the following steps:

1. Check the support matrix for supported switches, firmware versions, and SMI Agents before you start adding devices, which is found at the following website:

<http://www.ibm.com/support/docview.wss?&uid=swg21386446>

Click **Spectrum Control 5.2.x** in the column labeled Switches.

Tip: The preceding URL will not change when new Spectrum Control versions are released, so it is a good idea to bookmark the link.

2. Add the switch by clicking **Network** → **Switches**.
3. Click **Add Switch** to add a switch and its data source.

4. A dialog box opens (see Figure 4-4) in which you select the type of device that you want to add. By selecting the type of device, the agent configuration options are customized.



Figure 4-4 Select the type of switch you are adding

5. Configure the connection to the switch based on the vendor and the available agent options:
- For Brocade switches, you should monitor with Network Advisor as the CIM Agent. Change the default communication protocol, port, and namespace if needed for your environment. When you have entered the information, click **Next** to continue, as shown in Figure 4-5.

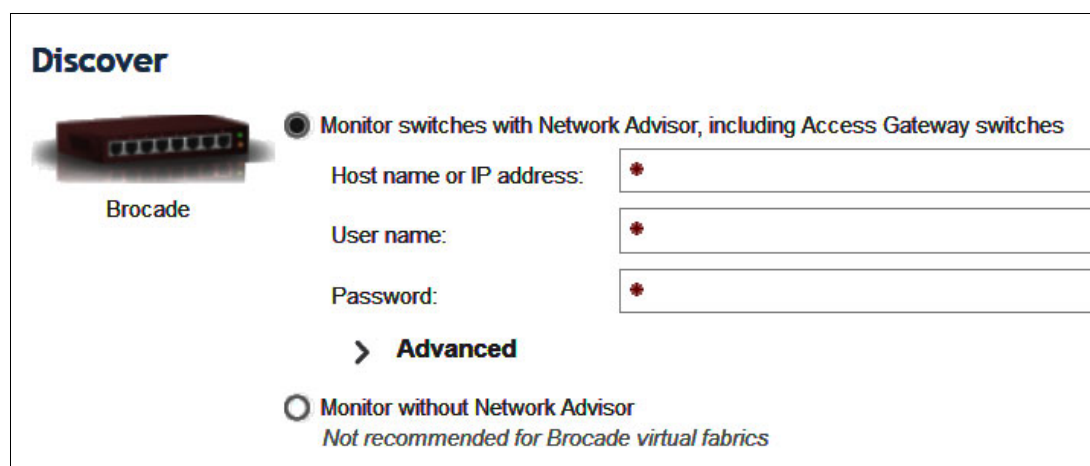


Figure 4-5 Adding Brocade switches with Network Advisor

You can add the Brocade switches by using SNMP agents, but it is not necessary if you are using Network Advisor as a CIM Agent.

Tip: You can download and install the Network Advisor Software for Brocade switches without requiring a license and incurring additional costs. During the Network Advisor installation, select the **SMI-S Agent only** option.

- b. For Cisco switches, the only option is to add them by using SNMPv1 or SNMPv3. Select the protocol and then enter the connection information and the IP addresses of the switches to add. Click **Next** to continue, as shown in Figure 4-6.

Add Switches and Fabrics for Monitoring

Connection Details:

☒ **SNMPv3 (Recommended)**
☐ **SNMPv1**


User name:

Authentication password:

Authentication protocol:
MD5

Encryption key:

Encryption protocol:
None



Cisco

Add Switches:

Add one or more host names or IP addresses:

Host name or IP address	Host name or IP address	Host name or IP address
Host name or IP address	Host name or IP address	Host name or IP address
Host name or IP address	Host name or IP address	Host name or IP address

[Add More](#)

◀ Back
Next ▶
Cancel

Figure 4-6 Add switches by using SNMP

Tip: All of the switches that are added in this window are added by using the same connection details. If your switches have different user names, passwords, or community strings, you should add similar switches together and then click **Add Switch** again for the next set.

6. Spectrum Control verifies that it can communicate with the data source and then discovers which switches and fabrics are available through this data source.
7. Next, probe and performance monitor schedules are configured. When you have provided all the information and adjusted the scheduling based on your needs, click **Configure**, as shown in Figure 4-7 on page 79.

Figure 4-7 Configure probe and performance monitor schedules

Probes

Network resources for Brocade are almost always probed on the fabric level, but are probed at the switch level for Cisco. Brocade NPV (access gateway) switches are the exception. They are probed on a per-switch basis, and not as part of a fabric probe. For all other Brocade switches, if you run a probe for a single Brocade switch in a fabric, it is the same as running a probe for the Brocade fabric. The same is true when you schedule the probes or look at the log files. You can do these activities for both Brocade and Cisco from either the context of a switch or the corresponding fabric.

Although Brocade probes are typically at the fabric level, performance data collection is performed at the individual switch level.

4.1.4 Next steps

The Add Device wizard simplifies the process of adding devices to Spectrum Control by including the scheduling setup and customization. To complete the process of adding devices, here are several tasks that you should complete:

- ▶ Wait for the probe to finish to see whether there are any components with a status other than Normal that you should acknowledge (for example, ports that are not connected to a fabric).
- ▶ Review names, Custom Tags, and Location in the property window of the device, and update them as necessary.
- ▶ If you added a storage system, you might want to continue with setting the tier levels for the storage pools and add the storage pools to capacity pools.
- ▶ Review retention settings, job logs, performance data, and removed resources. Click **Settings** → **History Retention** in the menu bar to review or edit them.
- ▶ Check the default alerts and set alerts for events or thresholds and specify email addresses or other alert notification options. For more information about alerting in Spectrum Control, see Chapter 6, “Advanced alert management” on page 177 or the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1UJizmy>

4.1.5 Updating and testing a data source

A data source is the link between Spectrum Control and a device. Over time, it is not uncommon to expect that the credentials (user ID and passwords) will need to be updated. In some cases, the IP address might need to be changed.

Testing the connection

It is important to know whether Spectrum Control can communicate with a device. To test the connection, complete the following steps:

1. From the list of devices (or from the devices detail window), start the test from the Actions menu by selecting **Connections** → **Test Connection**.

In Figure 4-8, we used the context menu because using the Actions menu would have hidden some more important parts of the window.

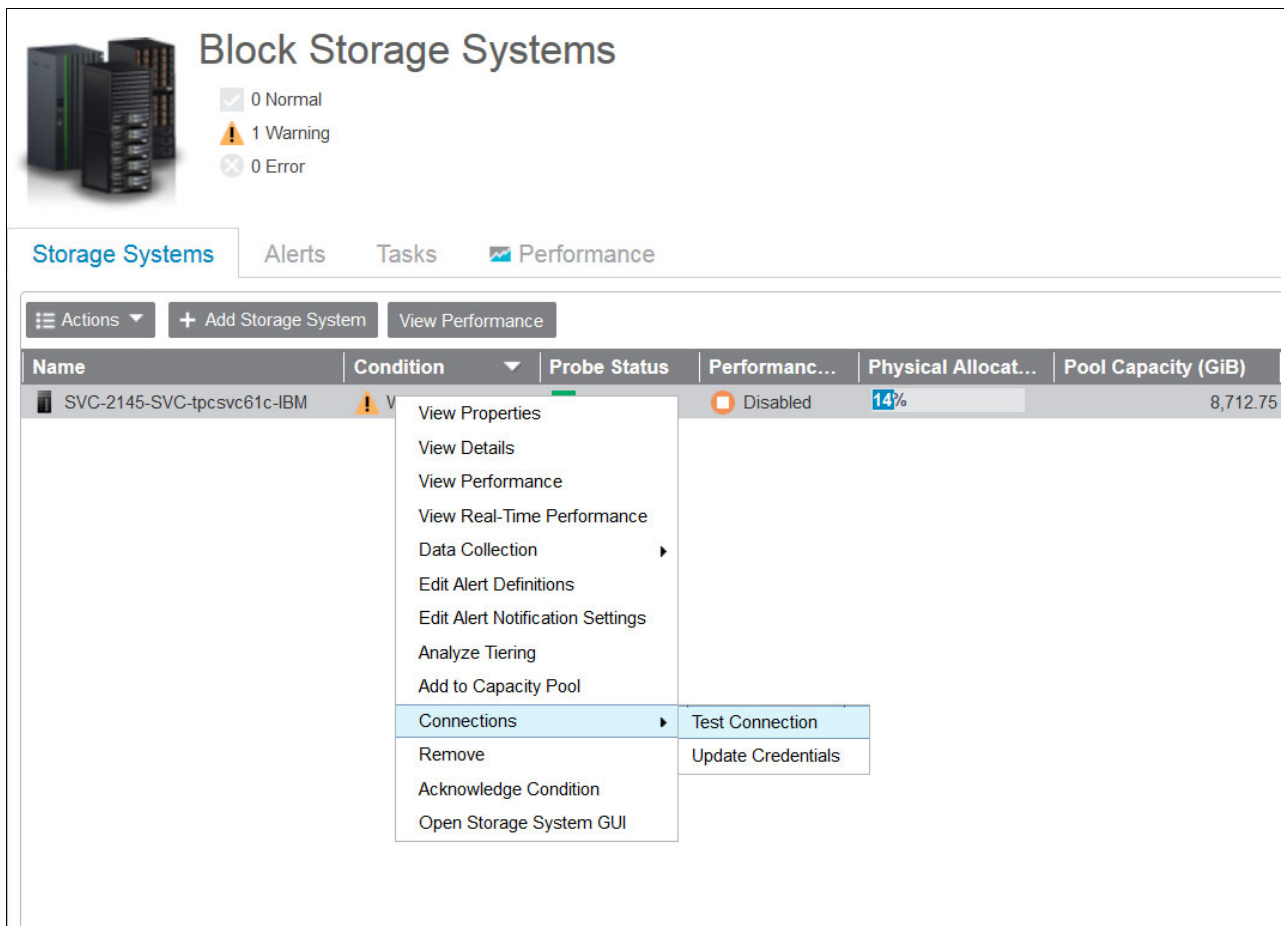


Figure 4-8 Test Connection window

2. Spectrum Control tries to use the data source to communicate with the device. You see a message indicating whether the connection test was successful. Click **Close** to return to the previous window.

If the test was successful, Spectrum Control updates the status. If you are working with a device that has multiple data sources that are configured, this menu and the Update Credentials menu have submenus that list the data sources, as shown in Figure 4-9 on page 81.

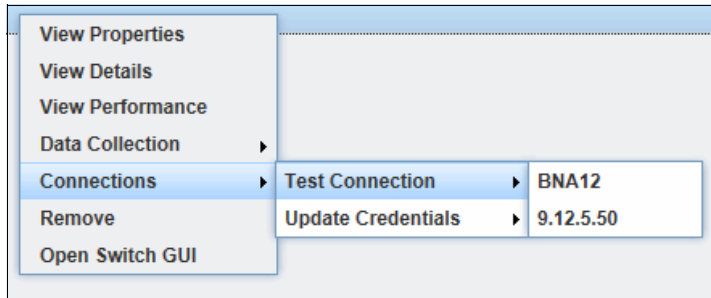


Figure 4-9 Update Credentials window

Updating the credentials

To update the device credentials, complete the following steps:

1. To update the user name and password, from the Actions menu click **Connections** → **Update Credentials** from either the device detail window or from a list of the devices.
2. Update the User name and Password fields as required.

The menu and the window where you enter the information looks slightly different from device to device because the interface depends on the data source type.

Updating the data source IP address

There is no explicit function to update the IP address for a data source. This section describes the way we simulated this scenario:

1. We swapped the cluster IP address and the service IP address of one of our SAN Volume Controllers.
2. We then added a storage system to Spectrum Control by using the new cluster IP address.

What happens in the background is that Spectrum Control runs a discovery and realizes that the device is known and simply updates the IP address. A message displays and indicates that the resource is being managed. When you see this message, click **Cancel**. The IP change is saved to the database. Spectrum Control now uses the new IP address.

4.2 Adding servers

You can add servers for monitoring by deploying a Storage Resource agent or by adding an agentless server to Spectrum Control. An agentless server is a server or virtual machine (VM) that is monitored by Spectrum Control but does not have a Storage Resource agent deployed to it. The agents are used to get a view of the connectivity between servers and other top-level resources, such as storage systems and switches.

4.2.1 Choosing an agentless server versus a Storage Resource agent

In some environments, deploying a Storage Resource agent to the server is not an option, either because of security restrictions in providing administration credentials, or simply to avoid loading a production server with agent code. An agentless server gives you the ability to model a host server, either a physical server or a VM, in Spectrum Control without deploying a Storage Resource agent.

By the definition of agentless servers, it is possible (as with Storage Resource agents) to get a view of the connectivity between servers and other top-level resources. This view is useful in troubleshooting scenarios to view performance data for connected subsystem and switch resources.

To gather server information, Spectrum Control correlates resources that are consumed by the server from storage system host connections and fabric zone aliases. The agentless server shows server Fibre Channel ports, HBAs, SAN-attached disks, and related SAN components. In many cases, you find that the agentless server provides the necessary information.

To determine which type of server to add, follow the guidelines in Table 4-1.

Table 4-1 Choose the type of server to add

Condition	Add an agentless server	Deploy a Storage Resource agent
View monitored resources to which a server is connected.	X	X
View storage assigned to a server from other monitored resources.	X	X
View the performance of the SAN-related resources connected to the server.	X	X
Organization restrictions prohibit deploying agents on servers.	X	
View information about internal server resources, such as controllers and disks.		X

Note: Agentless servers do not provide file system or local disk information. Use Storage Resource agents to access that information.

4.2.2 Adding an agentless server

An agentless server can be any grouping of worldwide port names (WWPNs) on which you want to report, for example, a Cluster or Application.

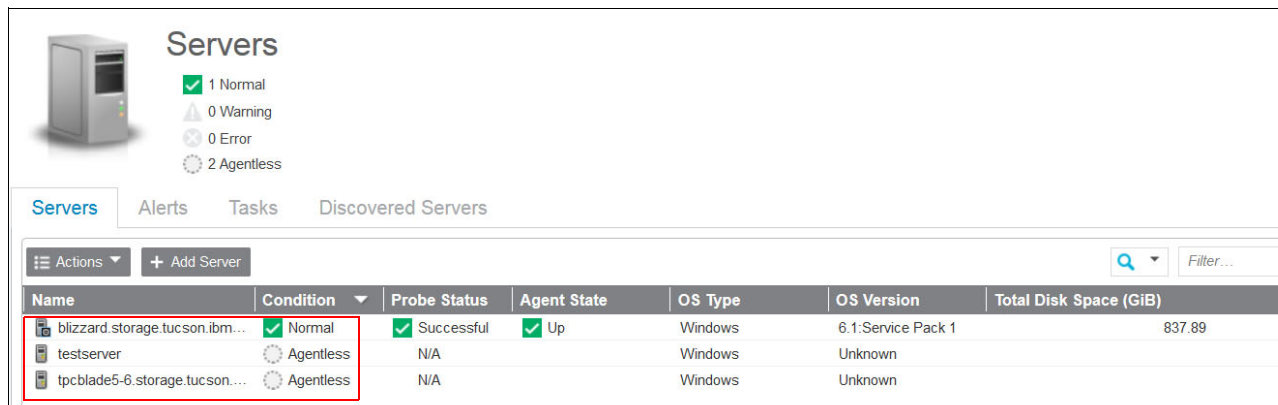
Add an agentless server by providing the IP address or host name of the server. The following optional entries can be provided:

- ▶ OS type
- ▶ Location
- ▶ VM or physical server
- ▶ WWPN

Important: Although adding WWPNs is optional, without WWPNs, Spectrum Control cannot correlate the SAN resources that are consumed by the server.

The WWPNs are used by Spectrum Control to correlate the agentless server with the storage systems, fabrics, and switches that are being monitored. For example, the ports that you enter for an agentless server are compared to the ports on host connections for monitored storage systems. If there is a match between the server and a host connection, you can view the connectivity between the server and storage system and the capacity of volumes that are assigned to the server.

After an agentless server is defined to Spectrum Control, it is added to the Servers window (which can be viewed by clicking **Servers** → **Servers**), as shown in Figure 4-10.



Name	Condition	Probe Status	Agent State	OS Type	OS Version	Total Disk Space (GiB)
blizzard.storage.tucson.ibm...	Normal	Successful	Up	Windows	6.1:Service Pack 1	837.89
testserver	Agentless	N/A		Windows	Unknown	
tpcblade5-6.storage.tucson...	Agentless	N/A		Windows	Unknown	

Figure 4-10 Servers window

Use any of the following methods to add agentless servers for monitoring:

- ▶ Add a server by manually entering its information.
- ▶ Add a server from a host connection on a monitored storage system.
- ▶ Add a server from a discovered port that is connected to a monitored switch.
- ▶ Add a server from a discovered port that is connected to a monitored fabric.

The Add Server button is available in the Servers window, as are the panes for host connections and discovered ports. For more information about adding an agentless server by using these methods, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1P3q9E3>

4.2.3 Deploying a Storage Resource agent

You deploy Storage Resource agents if you want to enable full server monitoring for internal resources, such as controllers, disks, and file systems.

Open the Add Server window by clicking **Servers** → **Servers** → **Add Server**, as shown in Figure 4-11. Select the **Deploy an agent for full server monitoring** check box.

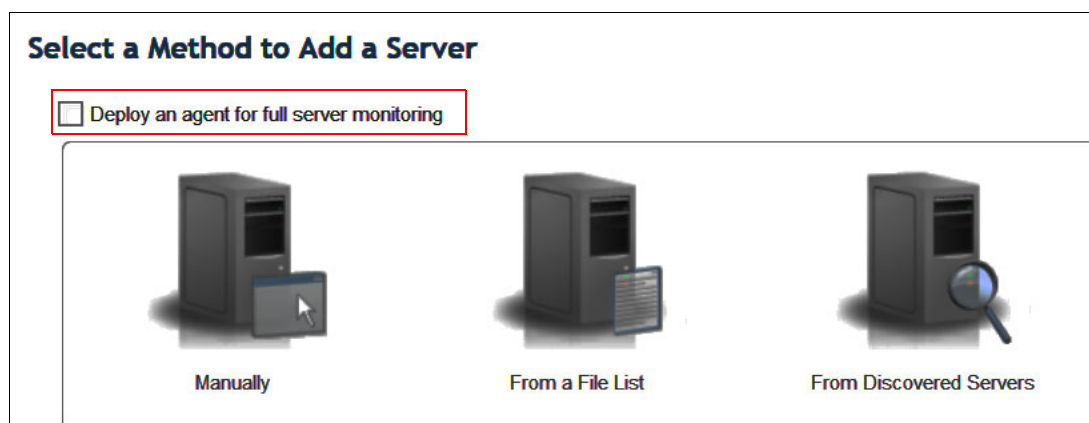


Figure 4-11 Deploy a Storage Resource agent

For more information about adding a Storage Resource agent, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1P3q9E3>

Note: The Storage Resource agent for Spectrum Control does not collect information about fabrics or switches. For the relevant agent options, see 4.1.3, “Adding fabrics and switches” on page 76.

4.3 Storage Resource agent use cases

This section describes administrative use cases that are based on monitoring the Storage Resource agents for Spectrum Control. They involve updating agent properties and handling changes to the DNS name or IP address for a server.

4.3.1 Updating agent properties

This section describes how to edit Storage Resource agent properties. Complete the following steps:

1. Open the Agent tab in the Properties window of the server, as shown in Figure 4-12 on page 85.

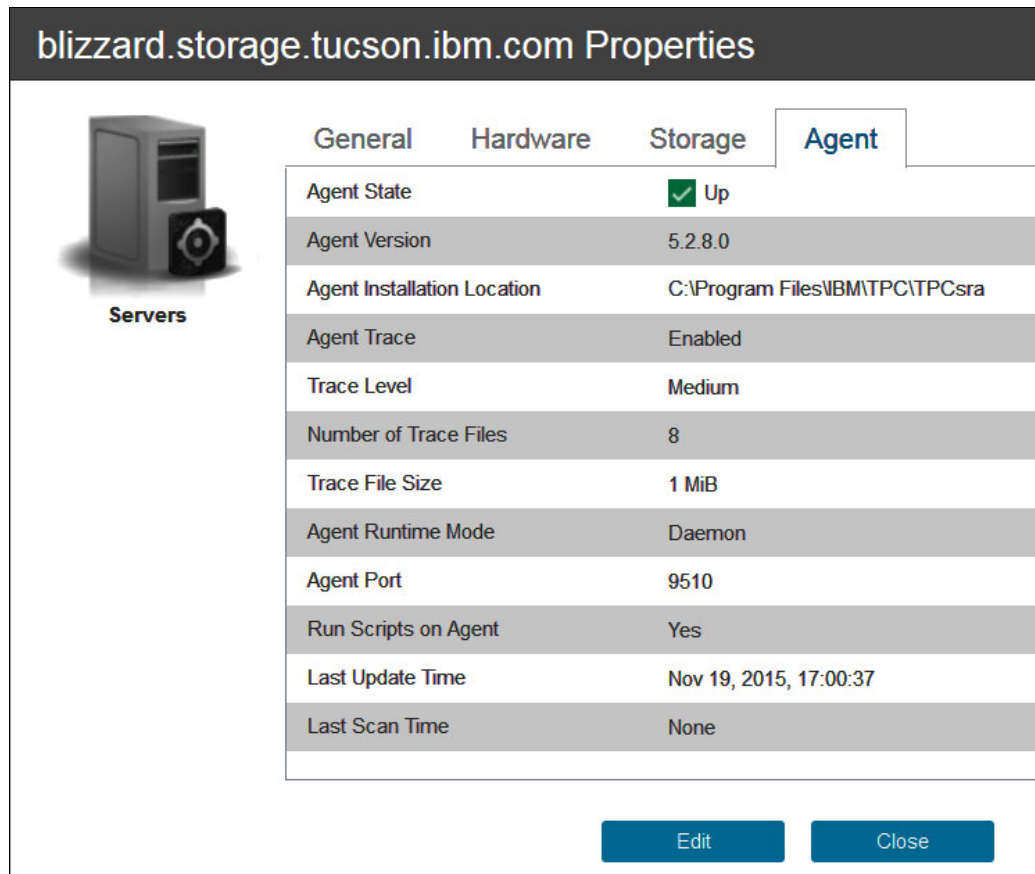


Figure 4-12 Agent properties

2. Click **Edit** to modify the agent properties. Several of the agent properties can be changed from the server Action list or the menu.

4.3.2 Changing the IP address or DNS name for Storage Resource agents

There are situations where you must make changes to the network setup that affect your Storage Resource agents. This section describes common network changes and how to update the Storage Resource agent to comply with the network change.

Changing the IP address of a Storage Resource agent

If you must change the IP address of a computer running the Storage Resource agent, you can do this change without any impact to Spectrum Control because Spectrum Control is always using the DNS name to connect to the Storage Resource agent.

Changing the DNS name of a Storage Resource agent

When you change the DNS name, Spectrum Control can no longer contact the agent. The new DNS name must be communicated to the Spectrum Control server. To update the Spectrum Control server, complete the following steps:

1. Stop the Storage Resource agent.
2. Create a file called REGISTERSRA (uppercase with no extension) in the root directory of the Storage Resource agent.
3. Restart the Storage Resource agent.

When the agent detects the REGISTERSRA file, it connects to the Spectrum Control server and registers itself, or simply updates the DNS name if Spectrum Control already knows that SRA.

This procedure is typically used when Spectrum Control agents are included in a master image that is used to install servers. For more information about the procedures, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/IndsINI>

Note: You can use the same process to register a Storage Resource agent with a new Spectrum Control server when you modify the connection information in the config\Agent.config file.

4.4 Additional configuration tasks

After you have the Spectrum Control software successfully installed with resources added and are successfully collecting data, there are still configuration tasks to help you understand and use Spectrum Control. The configuration can be changed later if you want.

4.4.1 Installing additional applications

Table 4-2 shows additional applications and tools that can be installed to complete the implementation of the Spectrum Control environment.

Table 4-2 Further configuration items

Item	Description
IBM Data Studio	<p>As of DB2 10, the Control Center is no longer available. The replacement for it is IBM Data Studio. The software is available at no charge and can be downloaded from the following website: http://www.ibm.com/developerworks/downloads/im/data/index.html</p> <p>You do not need this software to run Spectrum Control. However, if you used the DB2 Control Center to do some basic administration tasks or to develop SQL queries for custom reporting, you should consider downloading and installing the Data Studio.</p> <p>Data Studio is useful in troubleshooting Spectrum Control application slowdowns resulting from slow-running SQL queries. This tool can also recommend actions to speed up the performance of Spectrum Control (for example, adding indexes and reorgs).</p> <p>The following links provide more information about Data Studio:</p> <ul style="list-style-type: none"> ► http://www.ibm.com/software/products/en/data-studio ► http://ibm.co/1Q01a9T
Mozilla Firefox	<p>You might want to install Firefox on the Spectrum Control server. Normally, you use the browser on your computer to access Spectrum Control, but there might be situations where you directly log on to the server, so having a working browser is helpful. For more information, see the “Supported Platforms Agents, Servers and Browsers” topic, found at: http://www.ibm.com/support/docview.wss?uid=swg21386446</p>
Adobe Flash Player	The View Data Path function requires Flash Player be installed.

4.4.2 Backing up Spectrum Control

It is important to have a backup strategy for your Spectrum Control environment. There can be several items that you need to back up depending on how you use Spectrum Control. Here is a list of items that you should include in your backup strategy:

- ▶ Spectrum Control database
- ▶ Tivoli Common Reporting/Cognos Content Store
- ▶ Reports that were saved to the file system
- ▶ Tivoli Common Reporting/Cognos Report definitions

Because Spectrum Control and the Tivoli Common Reporting/Cognos Content Store are both stored in DB2, the backup can be easily integrated with Spectrum Control or other backup tools. Any reports that are saved in the file system should be part of a regular file system backup.

The backup of report definitions is more complex. Because any authorized user can create a report any time, you should set up a job in Cognos to create an export of the report definitions regularly. The export is a compressed file that is written to the local file systems so that it can be backed up by the regular file system backup.

For more information about the backup of Spectrum Control, see *IBM Tivoli Storage Productivity Center Beyond the Basics*, SG24-8236.

Important: For this approach to work, you should not use the “My Folders” location for your reports. In our testing, we did not find a way to export regularly reports that are stored there. As a preferred practice, create a folder in the Public Folders tab, and save your reports there.

You should not save any of your reports in the folders that Spectrum Control is creating. The reason is that this might make future upgrades more complex than necessary. If you have a folder that is created in the Public Folders tab, you can create links to the Spectrum Control provided reports, which makes it easier to find and go to the reports.

4.4.3 User authentication

Spectrum Control supports a series of federated repositories: Local OS, LDAP, or AD. LDAP and AD are external authentication methods, and Local OS is not.

The web-based GUI allows most actions to occur without requiring you to log in to the server locally. So, the use of external authentication servers is growing.

Depending on your specific environment, there are different methods for authentication in a Spectrum Control environment, as shown in Figure 3-1 on page 48. By default, Spectrum Control uses Local OS authentication for the web-based GUI.

Note: Tivoli Common Reporting/Cognos Reporting uses a default file-based repository with the JazzSM package. The Cognos GUI checks the user ID and password that is entered against a file that contains user IDs and passwords instead of checking with the local OS. Therefore, you must use the same user ID and password that you entered during the JazzSM installation to access Cognos.

4.4.4 Tivoli Common Reporting/Cognos Reporting single sign-on

Even when Tivoli Common Reporting/Cognos Reporting is set up by using the same authentication server as Spectrum Control, you still must authenticate again when you start the reporting.

There is a way to configure the WebSphere Application Server server running Tivoli Common Reporting/Cognos Reporting to accept and use single sign-on (SSO) tokens so that you do not need to enter a user ID and password. For instructions about how to set up this form of authentication, see the following website:

<http://ibm.co/1ZnpSa6>

Important: Single sign-on works only if you have switched from the default file repository authentication to LDAP in the WebSphere Application Server server that Cognos is using.

4.4.5 User management in Spectrum Control

Roles determine the functions that are available to users of Spectrum Control. When a user ID is authenticated to Spectrum Control through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

There are three roles that are available in Spectrum Control, as listed in Table 4-3.

Table 4-3 User roles in Spectrum Control

Role	Users that are assigned this role
Administrator	Can use every function in Spectrum Control. By default, the following groups are assigned the Administrator role: <ul style="list-style-type: none">▶ Windows: Administrators▶ Linux: root▶ AIX: system
External Application	Cannot log in to the Spectrum Control GUI. This role should be used for external applications that use the Spectrum Control provisioning functions, such as: <ul style="list-style-type: none">▶ vSphere Web Client Extension for Spectrum Control▶ IBM SmartCloud Storage Access
Monitor	Can log in to Spectrum Control but cannot run any function. A user with this role still can see all information and open log files, but here are the only actions that the user can do: <ul style="list-style-type: none">▶ Acknowledge alerts▶ Acknowledge a non-normal status▶ Set the tier level of a storage pool

You assign a role to a group of users, so all users in that group can perform certain actions in Spectrum Control. The role to group mappings are available by clicking **Settings** → **User Management** from the menu bar. Figure 4-13 on page 89 shows the User Management window where user groups can be configured.

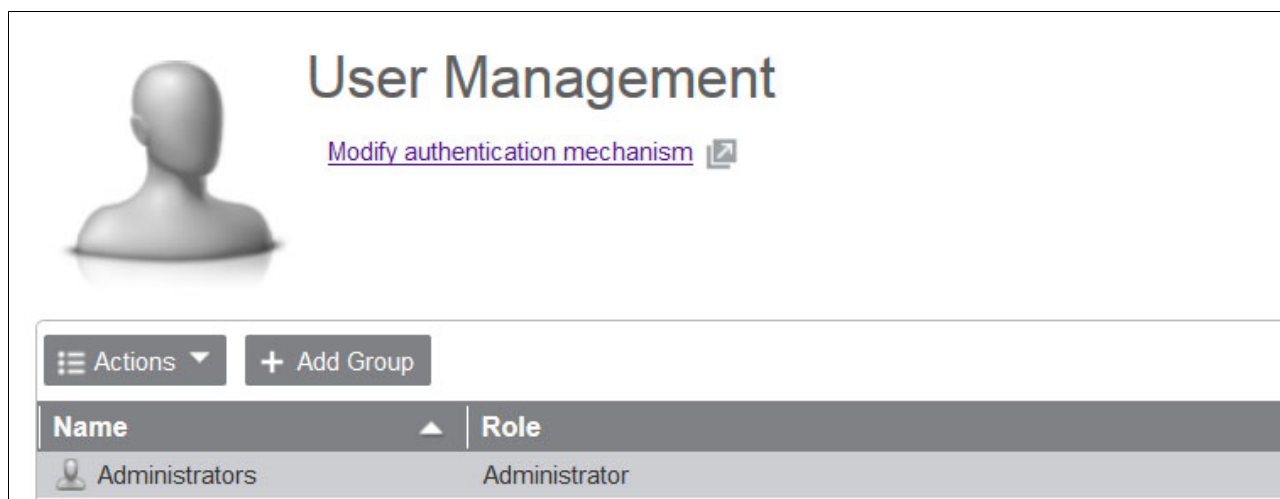


Figure 4-13 Spectrum Control User Management

From this window, you can also directly start the WebSphere Integrated Solutions Console, which is where you configure the authentication mechanism, such as AD or LDAP.

For more information about user management, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1mQpj76>

4.4.6 Preferred practices configuration tips

This section provides tips about how we typically configure a Spectrum Control environment and the benefits of doing it this way.

User IDs for device logins

It is easier to use a superuser ID when you add a device to Spectrum Control than create a user ID that is specific for the storage management task. However, as a preferred practice, create a special user ID for Spectrum Control on the storage device. When you have two Spectrum Control servers (one for test and one for production), consider using separate users IDs for both Spectrum Control servers so that you have more control over which Spectrum Control server can work with a specific device.

In our lab environment, we created a user who is called `tpc_87` for the Spectrum Control server whose IP ends in 87 on the devices that we added to that Spectrum Control instance.

For some data sources, this task requires extra effort because some CIM agents do not require authentication. Because Spectrum Control is often the only application that uses the CIM agent, we did not go through the extra effort of adding that user for the connection between Spectrum Control and the CIM agent.

The IBM System Storage DS4000® and IBM System Storage DS5000™ CIM agent is an example of a data source where you do not configure a special user ID setup. Typically, use the user ID *any* and enter the password *any* when you configure the data source to document that no special user ID is set up between Spectrum Control and the CIM agent.

For the communication between the DS4000 or DS5000 CIM agent and the device, create and use a special user ID.

Tip: The added benefit for doing this setup is that you can look into audit files and understand easily which commands Spectrum Control has run. Some devices log only commands that change the configuration. For those devices, you probably do not see any entries in the audit log unless you use the optimization, provisioning, or transformation functions.

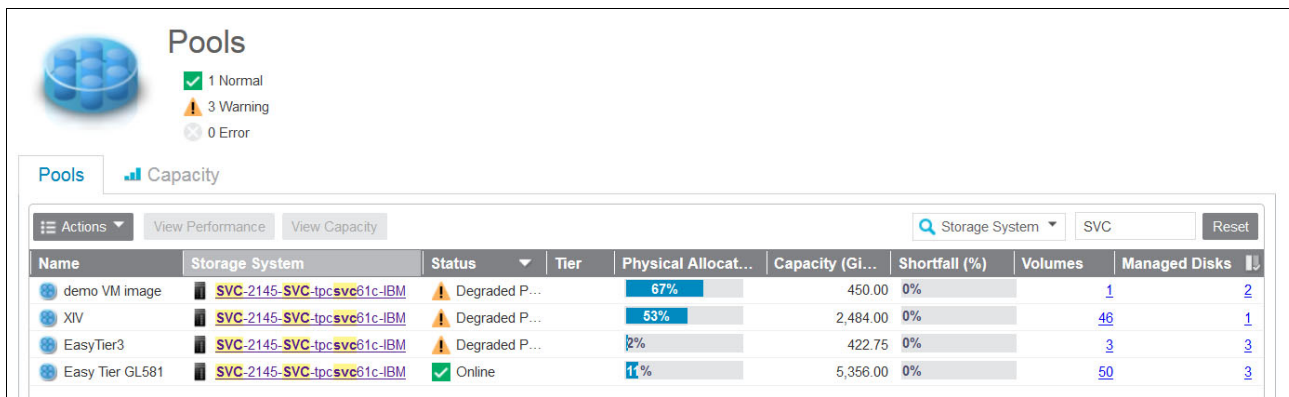
If you are using SAN Volume Controller or Storwize based storage systems and you do not want to or cannot use the user ID and password method, use a public/private key pair for each Spectrum Control server.

Naming conventions

Using a consistent naming scheme is a good idea. Here is one example of why this scheme can make your work so much easier.

The window that show information about pools does not provide information about the storage system type. If you have SAN Volume Controller or Storwize V7000 storage systems, you might want to see only those storage systems in the lists. The SAN Volume Controller and Storwize V7000 storage systems are like the pools that contain the volumes that are assigned to servers. So, the names of a storage system should include some information that you can use as filters in the GUI.

The filtering typically works across all columns in the displayed tables. You can also restrict Spectrum Control to look for the search string in a particular column, as shown in Figure 4-14, for example, for the one that is labeled *Storage System*.



Name	Storage System	Status	Tier	Physical Allocat...	Capacity (Gi...	Shortfall (%)	Volumes	Managed Disks
demo VM image	SVC-2145-SVC-tpcsvc61c-IBM	Degraded P...		67%	450.00	0%	1	2
XIV	SVC-2145-SVC-tpcsvc61c-IBM	Degraded P...		53%	2,484.00	0%	46	1
EasyTier3	SVC-2145-SVC-tpcsvc61c-IBM	Degraded P...		2%	422.75	0%	3	3
Easy Tier GL581	SVC-2145-SVC-tpcsvc61c-IBM	Online		1%	5,356.00	0%	50	3

Figure 4-14 Filter on Storage System

Tip: When the search function is used without restrictions, it also looks for the search string in columns that are not displayed.

By default, the Custom Tag columns are not displayed, but because Spectrum Control looks into those columns anyway, you can use the Custom Tags to customize your environment in a way that lets you use filters to focus on what is important to you.

Because the values in a name depend on your environment and requirements, there is no general schema to use. Also, you might have well-established names that you do not want to change. The next section provides alternative ideas.

Adding information to Spectrum Control

Spectrum Control offers ways that you can customize it. These ways might be more applicable than naming conventions in some cases. Here is a list of items that you might find useful:

- ▶ Custom Tags/User-defined properties

These properties are the most versatile custom attributes that you can use for many different purposes. They are available for the following items:

- Storage Systems
- Storage Pool
- Servers
- Hypervisors
- Switches
- Fabrics

- ▶ Location

The location property can be set for the following items:

- Storage Systems
- Servers
- Hypervisors
- Fabrics
- Switches
- Capacity Pools

For custom reporting, you can find the properties in the data model within the “Component Properties” folder of the corresponding component. An example is the label “Storage Virtualizer Pool Custom Tag 1”.

- ▶ Tiers for storage pools

Because there are many ways that users define what a tier is, Spectrum Control offers a simple and flexible way to define what *you* think a tier is. Basically, you can choose a number 1 - 10 and tag a storage pool with this information.

The Tier property is available for any storage pool no matter what type of storage system to which the pools belong. Because it is a simple tag, Spectrum Control does not check for any inconsistencies. For example, you can define all your pools containing 15 K rpm disks as Tier 5, and all 10k rpm disks as Tier 3 independent from the storage system type.

For custom reporting, you can find the properties in the data model within the “Component Properties” folder of the corresponding component as “Storage Virtualizer Pool Tier” or “Storage Pool Tier”.

- ▶ Capacity pools

Capacity pools are part of the cloud configuration concept in Spectrum Control. They provide a way of grouping of storage pools. Essentially, they are a special “type” of Storage Resource Group, which can contain only storage pools. Another difference is that one storage pool can be assigned only to one capacity pool.

Note: Not all of the preceding properties are available within Tivoli Common Reporting/Cognos Reporting for creating your custom reports through dragging. Nevertheless, some of the information is available through the TPCREPORT views.

4.5 Dashboard administration use cases

Here are some administrative use cases that are based on monitoring the Spectrum Control dashboard that can help you keep a clean and easy-to-monitor dashboard. They involve handling the alert status, performance monitor status, and the condition of your top-level resources.

4.5.1 Viewing the quick status on the dashboard

As a storage administrator that is responsible for SAN switches and storage, you might be contacted by an application team reporting that they saw a loss of paths, added latency, or possibly both. It is your responsibility to find the source of the problem and resolve it.

With Spectrum Control, you can monitor both SAN switches and storage within a common application, which can be a time saver. However, you must keep your monitoring environment in a state such that problems are apparent and not masked by unimportant alerts, which involves making sure Spectrum Control is actively monitoring the devices that you need and keeping your server in a state such that you can identify issues. For example, if everything is always reported as critical or warning, you learn to ignore the alerts.

Figure 4-15 shows a dashboard view with a storage system warning and various fabric and switch issues in the upper half. The lower left corner shows recent alerts and a quick view of how your performance monitors are running.

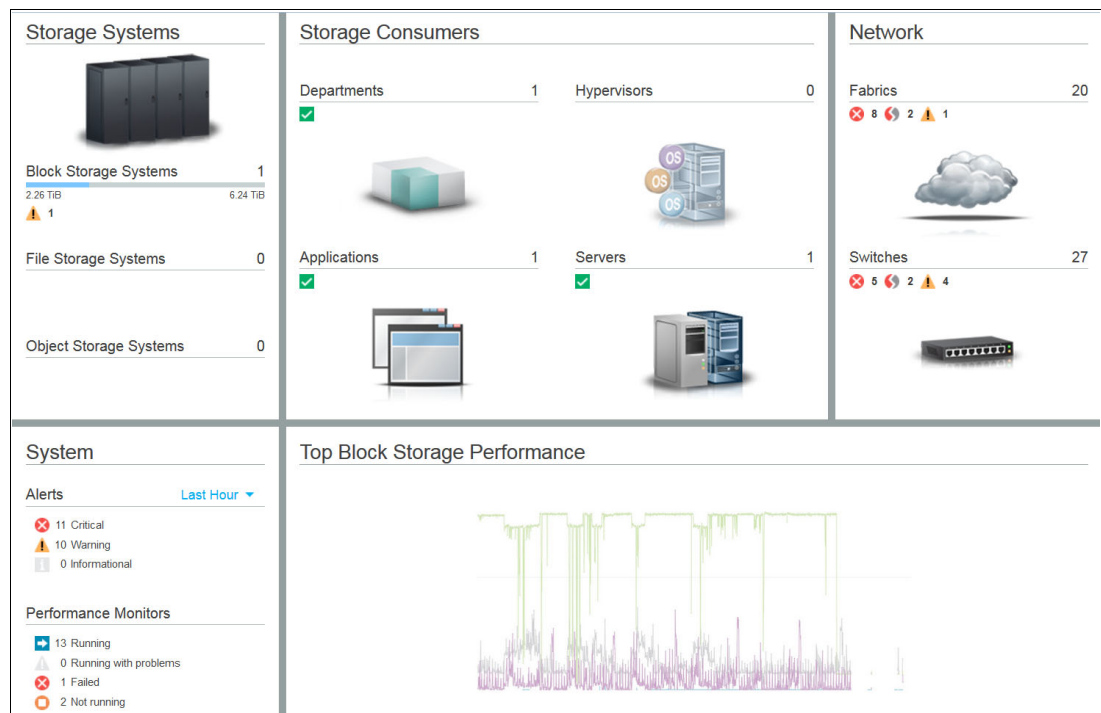


Figure 4-15 Spectrum Control dashboard

Tip: If you are not familiar with the displayed icons, select the ? in the upper right menu bar to open a help topic for the dashboard that explains each icon. You can open help anywhere in Spectrum Control in this manner.

If your switches and storage systems are normally without error or warning, you quickly can see where new issues are occurring by using the dashboard. Handles the alerts and any warning or error conditions on a regular basis for the monitored resources, which helps you keep your dashboard clean.

4.5.2 Handling alerts: Critical and warning

You can find details about the alerts that are shown in Figure 4-15 on page 92 by clicking **Alerts** in the lower left corner to open the Alerts window, as shown in Figure 4-16.

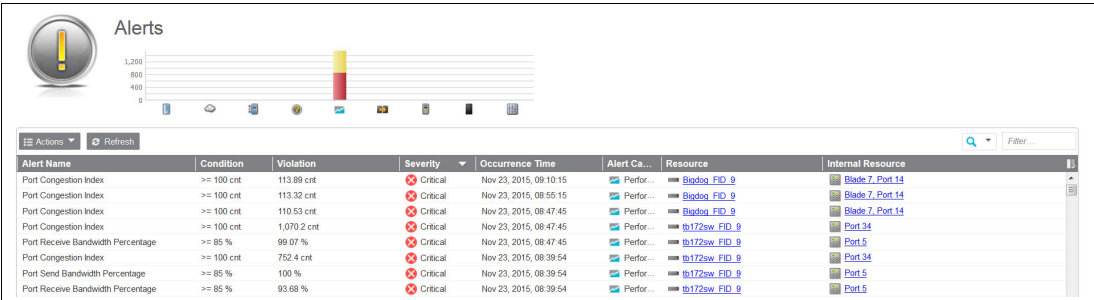


Figure 4-16 Alerts

In this case, the graph at the top of the Alerts window in Figure 4-16 shows that the errors are all related to the performance category. As you look at the most recent ones, you can see that they are related to two switch resources and an internal resource for each occurrence, such as a particular port. To see more details, select an individual alert row and select **View Alert** from the Actions menu. Figure 4-17 shows the details for the Port Congestion Index alert from row one. From this view, you can see that the threshold was crossed multiple times, how often, and to what level.

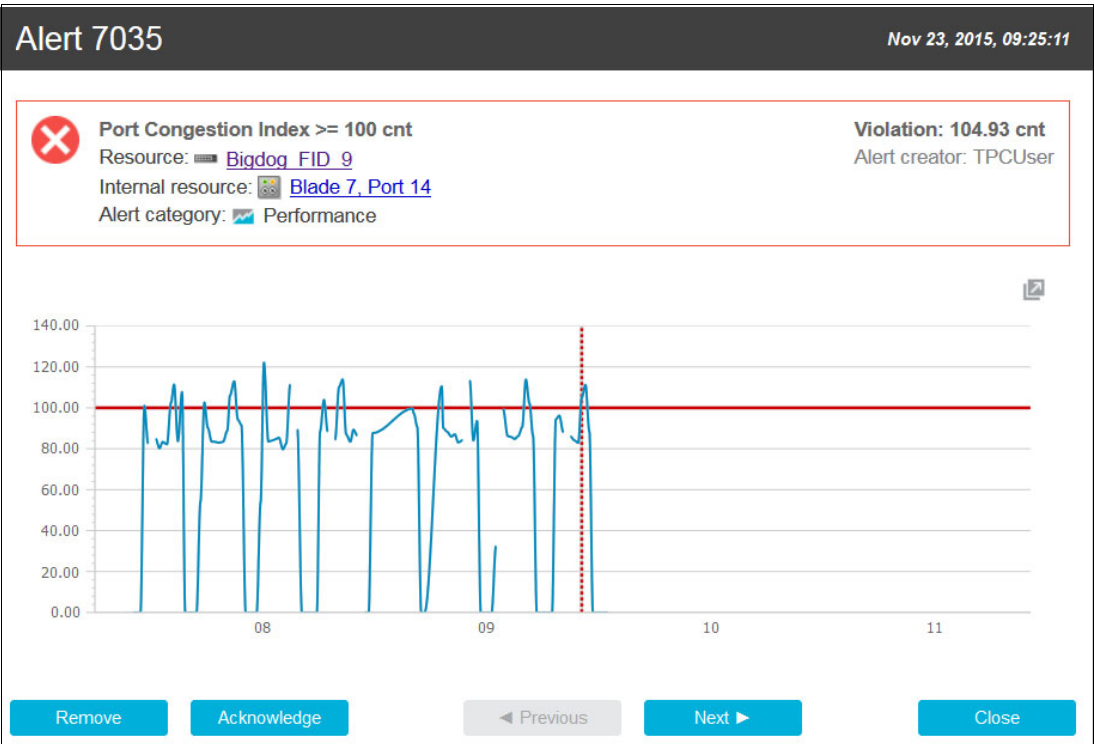


Figure 4-17 Port Congestion Index alert

What do you do in this case? There are a few options:

- Resolve the issue that is causing the alert.

Drill into the linked resource and internal resource to find additional information, such as whether a port is reporting an error status. You may also start an external element manager for the resource to investigate further. After investigating, you might want to do one of the other alert actions in this list.

- Acknowledge the alert.

Some alerts in Spectrum Control are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution.

- Remove the alert.

Remove alerts when you no longer need to view them. By default, alerts are automatically removed based on retention settings that are defined on the History Retention page in Spectrum Control.

- Set a different alert definition.

Determine whether the current alert condition is appropriate. If there is nothing of concern with how the resources are running, you might want to set a different threshold to avoid unnecessary alerts. The Spectrum Control default values are a starting point. The alert definition can be modified by drilling into the affected resource that is listed in the alert and then clicking **Alerts** in the left navigation and then clicking the **Definitions** tab, as shown in Figure 4-18.

The screenshot shows the 'Alerts' page in IBM Spectrum Control. On the left, a sidebar lists navigation options: Overview, Properties, Alerts (1545), Data Collection (2), Internal Resources (Inter-Switch Connections (6), Blades (6), Ports (26)), and Related Resources (Fabric, Switches (4), Discovered Ports (7)). The main content area is titled 'Alerts' and shows a summary of 838 Critical, 707 Warning, and 0 Informational alerts. Below this, there are tabs for 'Alerts', 'Definitions', and 'Notification Settings'. The 'Definitions' tab is active, showing a table of alert definitions for 'Ports' (10/11). The table has columns for 'General' (2/3) and 'Performance' (8/8). Under 'Performance', there are two sections: 'Port: Error Rates' and 'Port: Miscellaneous'. Each section contains several alert definitions with toggle switches, comparison operators, thresholds, units, and icons for alert status (warning, error, info, etc.). At the bottom, there are 'Save' and 'Cancel' buttons.

Figure 4-18 Alerts definitions for a selected resource

For more information about alerting in Spectrum Control, see Chapter 6, “Advanced alert management” on page 177, or the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/110k2sk>

4.5.3 Handling device condition: Warnings and errors

Use the Spectrum Control dashboard to get an at-a-glance snapshot of the health of your resources. After identifying top-level resources that need attention, use Spectrum Control to identify the internal resources that are encountering problems. The upper left corner of Figure 4-15 on page 92 shows a storage system warning condition, which is an example for how you can relate the condition value for a device to the status values of the internal resources that lead to this condition. The relationship between condition and status works similarly for storage systems, servers, hypervisors, fabrics, and switches.

Condition represents the most critical status that was detected on a resource itself and on its internal resources. For example, if an error status is detected on a disk, the overall condition of its storage system is error. Status represents the status of a resource as reported by its own hardware. For example, if a disk on a storage system is disconnected, an error status is reported for that disk by the storage system.

In this case, complete the following steps:

1. Drill into Storage Systems to view the listing of your storage systems and their conditions. Figure 4-19 shows the entry for the storage system with the warning condition.

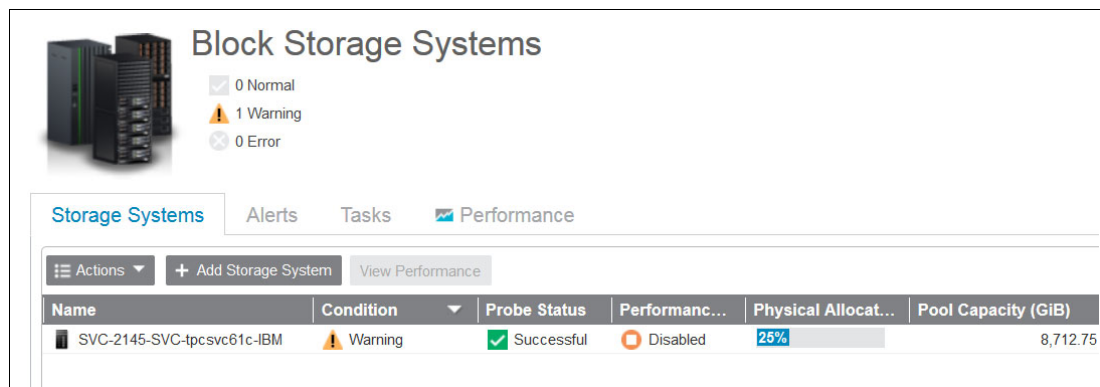
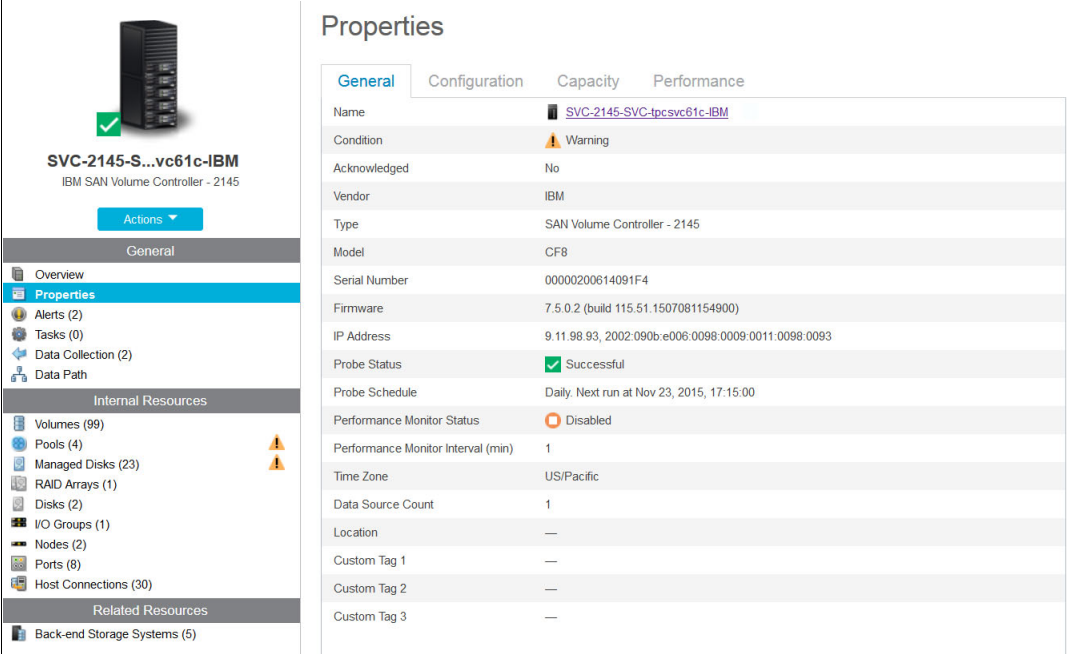


Figure 4-19 Block Storage Systems listing

2. Double-click the storage system row to open the detail window for it. Figure 4-20 shows that the Pools and Managed Disks under the Internal Resources for this storage system are listed with a warning status.



SVC-2145-S...vc61c-IBM
IBM SAN Volume Controller - 2145

Properties

General Configuration Capacity Performance

Name	SVC-2145-SVC-ipcvc61c-IBM
Condition	Warning
Acknowledged	No
Vendor	IBM
Type	SAN Volume Controller - 2145
Model	CF8
Serial Number	00000200614091F4
Firmware	7.5.0.2 (build 115.51.1507081154900)
IP Address	9.11.98.93, 2002:090b:e006:0098:0009:0011:0098:0093
Probe Status	Successful
Probe Schedule	Daily, Next run at Nov 23, 2015, 17:15:00
Performance Monitor Status	Disabled
Performance Monitor Interval (min)	1
Time Zone	US/Pacific
Data Source Count	1
Location	—
Custom Tag 1	—
Custom Tag 2	—
Custom Tag 3	—

Internal Resources


- Volumes (99)
- Pools (4)
- Managed Disks (23)
- RAID Arrays (1)
- Disks (2)
- I/O Groups (1)
- Nodes (2)
- Ports (8)
- Host Connections (30)

Related Resources

- Back-end Storage Systems (5)

Figure 4-20 Storage system detail window with internal resources with a warning status

3. The pools and managed disks are the resources contributing to this condition and should be reviewed. Drill into each of them to get more information. Figure 4-21 shows the pools with a warning status, and Figure 4-22 on page 97 shows managed disks with a warning status.



SVC-2145-S...vc61c-IBM
IBM SAN Volume Controller - 2145

Pools

1 Normal
3 Warning
0 Error

Pools Performance Capacity

Name	Status	Tier	Physical Allocat...	Capacity (GI...	Shortfall (%)	Volumes	Managed Disks	Compressio...	Tier Distribution (%)
demo VM image	Degraded Ports	67%	450.00	0%	1	2			
XIV	Degraded Ports	53%	2,484.00	0%	45	1			
EasyTier3	Degraded Ports	8%	422.75	0%	3	3			
Easy Tier GL581	Online	10%	5,356.00	0%	50	3			

Figure 4-21 Pools with a warning status

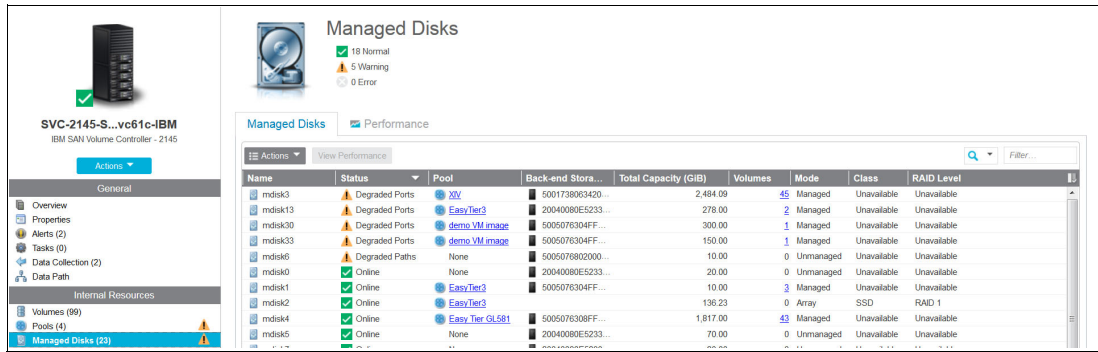


Figure 4-22 Managed disks with a warning status

Both the managed disks and the pools listings indicate a warning status because of degraded ports. You can use this information to investigate further and resolve your issue. However, you might not want to resolve it now, but do not need the condition for that resource reflected in the dashboard view, does not continue to influence the top-level condition. This situation can make it easier to clear out warnings and errors that do not matter to you.

4. Right-click and select **Acknowledge Condition** for a device, as shown in Figure 4-23.

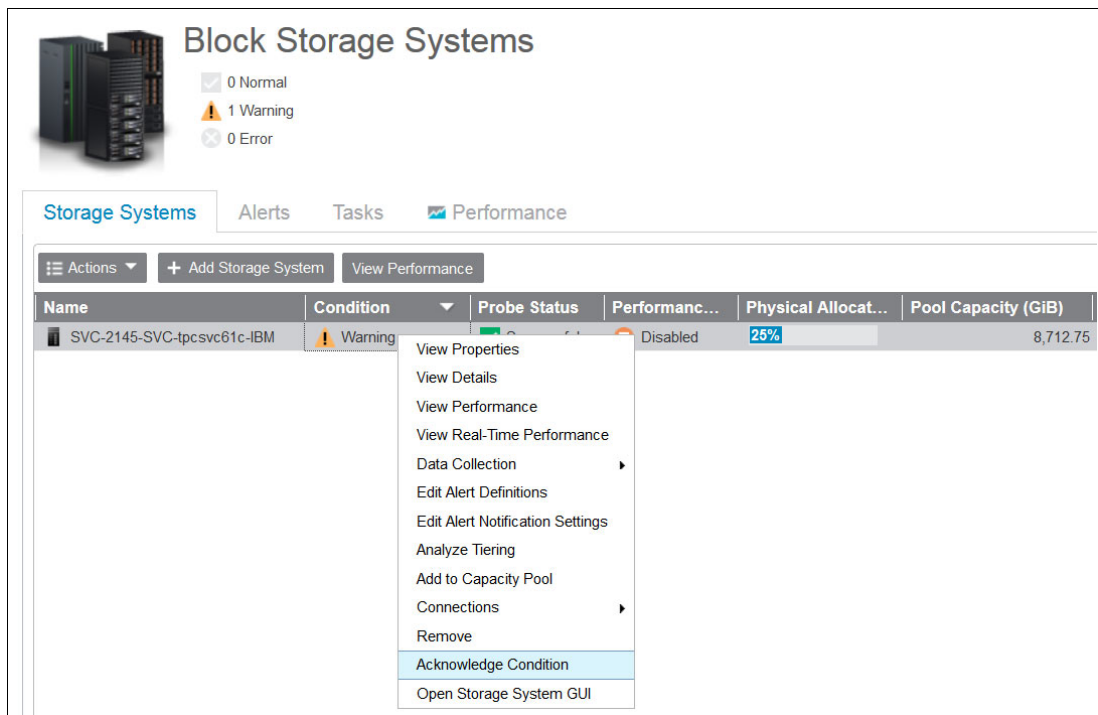


Figure 4-23 Acknowledge a condition for a top-level device

- You also can select individual resources, such as managed disks. Right-click and select **Acknowledge Status** to be more selective about what you choose to acknowledge, as shown in Figure 4-24.

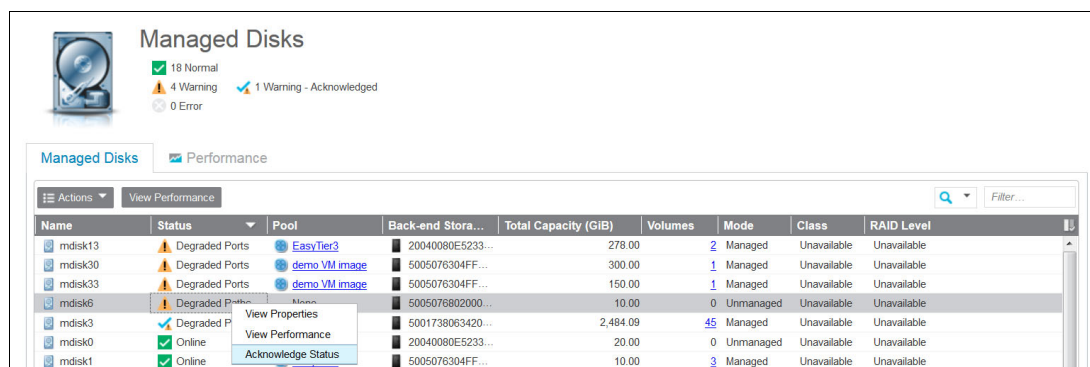


Figure 4-24 Acknowledge a status for an internal resource

After you acknowledge the storage system warning condition that is shown in Figure 4-24, the dashboard no longer reports a warning condition. The acknowledged condition does not contribute to the calculation like the warning condition did. Figure 4-25 shows the cleaned-up storage system section on the dashboard. The overall condition is now shown as normal.

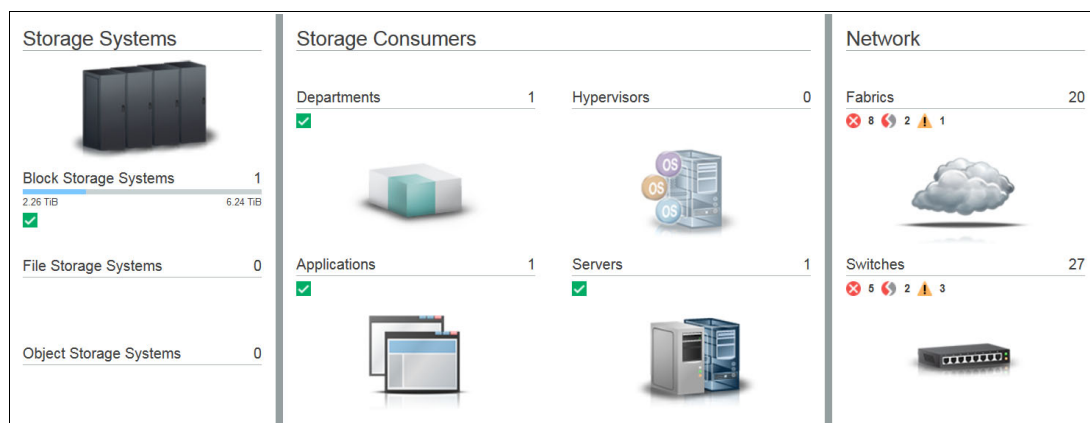


Figure 4-25 Storage system condition acknowledged

Tip: If you acknowledge a condition, status, or alert, you have a similar option to unacknowledge it.

Resolving the device condition warnings and errors for switches and fabrics is a similar process where you open the listing for Fabrics or Switches. Find the resource with the warning or error condition and then drill into the internal resources for it to find their status values. As shown in Figure 4-26 on page 99, if you click **Ports** from the left navigation, you can see that the ones marked as warning are because of a degraded status. The affected switches and their port numbers are listed. It also indicates that the ports are disabled. This information can help you tackle this issue. Acknowledging the status values for these ports so that it does not contribute to the overall condition is an option too. If internal resources are listed with an error or warning condition, they contribute to the condition of the top-level resource. Work your way into each resource category to understand and resolve what you see.

Ports

102 Normal
0 Warning
0 Error

Switch	Blade Slot	Port Number	Status	Port Type	State	Speed	WWPN	Connected Resource
Tetrisblast		3	Degraded	G_Port	Disabled	8	20030005336C2B38	
Tetrisblast		9	Degraded	G_Port	Disabled	8	20090005336C2B38	
tpcblade0c1		15	Degraded	G_Port	Disabled	8	200F0005336E574C	
tpcblade0c1		16	Degraded	G_Port	Disabled	8	20100005336E574C	
tpcblade0c1		17	Degraded	G_Port	Disabled	8	20110005336E574C	
tpcblade0c1		18	Degraded	G_Port	Disabled	8	20120005336E574C	
Bigdog_FID_9		288	Operational	E_Port	Online	5000533803F7B120		tp172new_FID_9
Bigdog_FID_9	1	0	Operational	G_Port	Disabled	8	2000000533803F01	
Bigdog_FID_9	1	19	Operational	F_Port	Online	8	2013000533803F01	tpcblade0c1
Bigdog_FID_9	2	0	Operational	F_Port	Online	8	2040000533803F01	Tetrisblast
Bigdog_FID_9	2	1	Operational	F_Port	Online	8	2041000533803F01	Tetrisblast
Bigdog_FID_9	2	13	Operational	F_Port	Online	4	204D000533803F01	Unavailable
Bigdog_FID_9	7	13	Operational	F_Port	Online	4	208D000533803F01	Unavailable
Bigdog_FID_9	7	14	Operational	E_Port	Online	8	208E000533803F01	tp172new_FID_9
Bigdog_FID_9	7	15	Operational	E_Port	Online	8	208F000533803F01	tp172new_FID_9

Figure 4-26 Fabric ports with warning status

4.5.4 Clearing unreachable switches

Now that the Storage Systems section is cleaned up, as shown in Figure 4-25 on page 98, you can turn your attention to the switches and fabrics. Looking at the dashboard, there are two unreachable switches and fabrics, which are shown on the right. An unreachable fabric is typically caused by an unreachable switch or data source, such as a CIM agent or SNMP agent. The device might be unexpectedly unavailable or unavailable as the result of environmental changes or a plan to decommission hardware. In any of these cases, Spectrum Control should be cleaned up to remove the old data sources or correctly configure them.

Complete the following steps:

1. Start by drilling into the switches or fabrics listing. Figure 4-27 shows switches with several different conditions. The two switches with unreachable condition are used in this example. You might notice that the two unreachable switches have the same name, tpcf-b155. This switch is a Cisco physical switch (black switch icon) and its logical switch (blue logical switch icon). Both use the same data source connection to the SNMP agent and either one can be used for the next actions.

Switches

17 Normal
3 Warning
2 Unreachable
5 Error

Switches Alerts Performance

Name	Condition	Probe Status	Performance M...	Fabric	Principal Switch of Fabric	Ports	Connected Ports	Links
director	Error	Successful	Running	san_director		80	12	4
director_vsan3	Error	Successful	Disabled	VSAN_SNMPTE...	director_vsan3	9	2	
director_vsan1	Error	Successful	Disabled	VSAN1_director	director_vsan1	1	1	director-Port 106 ISL 1
director_vsan2	Error	Successful	Disabled	VSAN0002	director_vsan2	1	1	director-Port 106 ISL 1
director_vsan77	Error	Successful	Disabled	VSAN_77	director_vsan77	88	9	2
tpcf-b155	Unreachable	Failed	Disabled	san_tpcf-b155		16	1	port-channel 1
tpcf-b155	Unreachable	Failed	Disabled	VSAN0001	tpcf-b155	16	1	port-channel 1
Bigdog_FID109...	Warning	Successful	Running	jH_VFSquash-1...	Bigdog_FID109_DefaultSwitchT	70	5	2
Tetrisblast	Warning	Successful	Running			24	4	Bigdog_FID_9 Trunk 1
tpcblade0c1	Warning	Successful	Running			20	14	2

Figure 4-27 Switches listing that shows conditions

2. Right-click one of the unreachable entries and select **Connections** → **Test Connection** to check it. If it can be reached, running a probe is suggested, as shown in Figure 4-28. If you select to start the probe, the GUI is updated sooner.

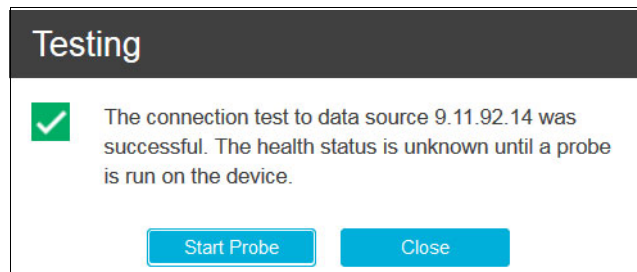


Figure 4-28 Successful test connection

3. If the test connection does not pass, the data source might not be available. You can check and update the credentials. Right-click the switch and select **Connection** → **Update Credentials**. You can make any adjustments that are needed to reestablish a connection.
4. If you know that the switch is decommissioned or you no longer want to monitor it with Spectrum Control, you can right-click the switch and select **Remove**. This action removes the resource and all associated data, including probe and performance data.

Whatever method you use to clean up the unreachable switches, it is preferable to do so because unreachable switches are not being monitored.

4.5.5 Handling performance monitors

Performance monitors are an integral feature of Spectrum Control and its function. As an administrator, you should know whether your resources are being monitored so that the performance data that you need is available when you need it. The lower left corner of the dashboard, as shown in Figure 4-29 on page 101, provides a quick view to check the status for the defined performance monitors.

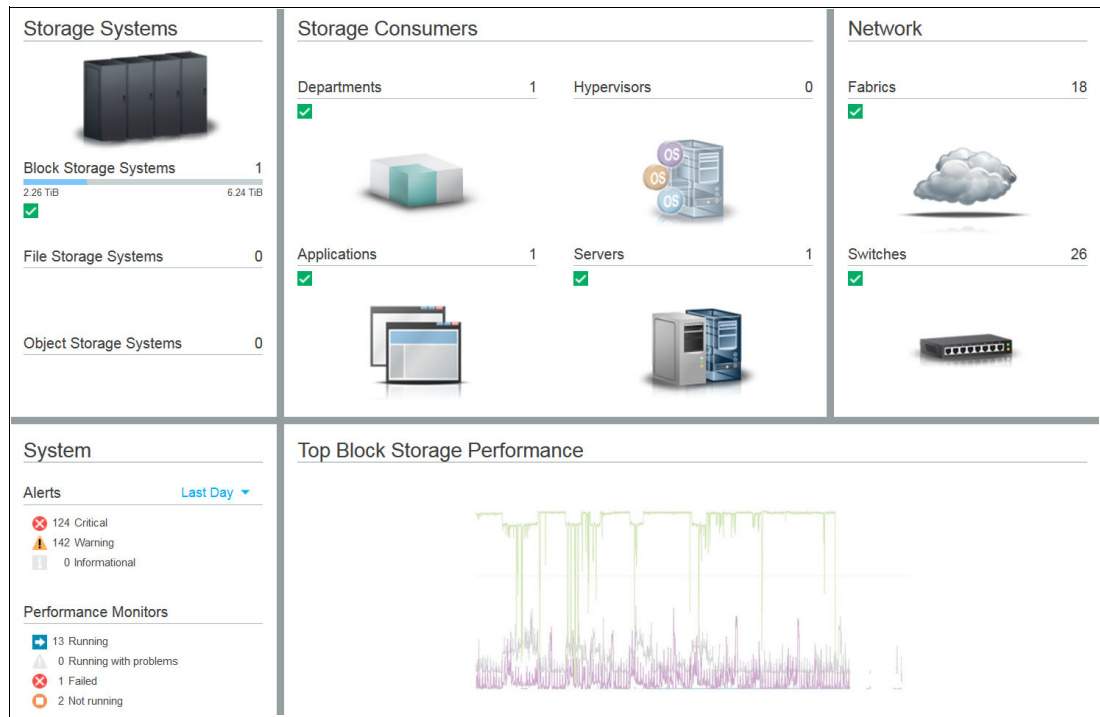


Figure 4-29 Dashboard view with Performance Monitors in the lower left

Ideally, the performance monitors are running without problems. If they are not, you should evaluate them further. In the case that is shown in Figure 4-29, most are running well, but there is one failed monitor and two that are not running. They are not actively collecting performance data for their resources.

Complete the following steps:

1. Select the Performance Monitors heading to see a listing of all performance monitors that are defined in Spectrum Control, as shown in Figure 4-30.

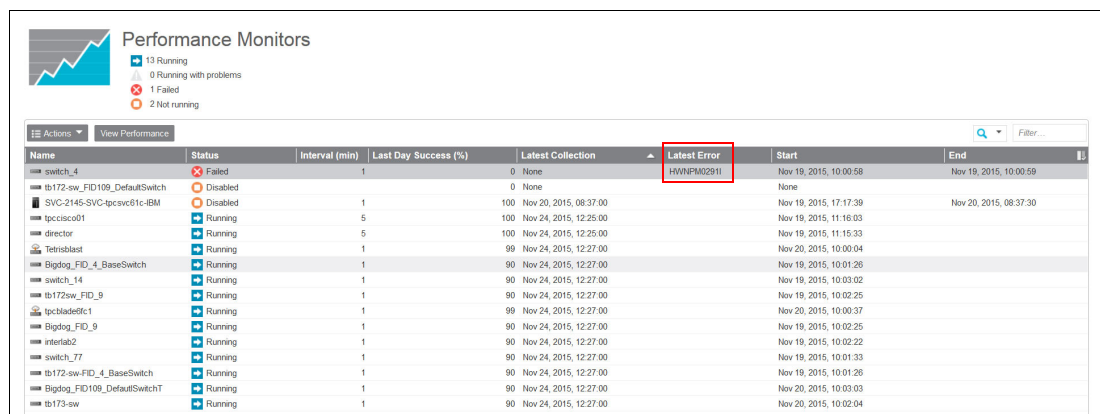


Figure 4-30 Performance Monitors status

- The column that is labeled *Latest Error* in the performance monitors listing in Figure 4-30 on page 101 shows the error code HWNPM0291 for the resource in the first row. It also shows that the start and stop times are close together for the one with a failed status. It did not run for long. To find out why a performance monitor failed, right-click its row and select **Open Logs**.
- Scroll to the bottom of the log to find the reason that the monitor failed. Figure 4-31 shows that in this case that the monitor failed because the switch has no ports. It might be that a performance monitor should not have been defined for this switch, that the switch is incorrectly configured, or that something unexpected happened to the switch ports. Whatever the reason, you now have information that you can use to resolve this failure.

Information	Nov 19, 2015 10:00:59.099	HWNPM0291I	Performance monitoring is unavailable for switch switch_4 (10000005336E697E) because the switch has no ports.
Error	Nov 19, 2015 10:01:02.796	HWNPM2128E	The performance monitor for device switch_4 (10000005336E697E) is stopping due to an unexpected failure.

Figure 4-31 Performance monitor log failure

- Figure 4-30 on page 101 also shows two performance monitors with a disabled status. You might want to disable these particular performance monitors, but it might also be that they were not scheduled when the resource was initially added to Spectrum Control. To enable a disabled performance monitor, right-click it and select **Schedule**. Set it as **Enabled** and select a data collection frequency, as shown in Figure 4-32. A similar action can be done to disable a performance monitor.

tb172-sw_FID109_DefaultSwitch

Data Collection Schedule

Performance monitor:

Enabled

▼

Every minute

▼

Save

Cancel

Figure 4-32 Enable a performance monitor

From time to time, you might see performance monitors report a status of “Running with problems”. If the performance monitor can recover, it goes back to Running. If it cannot, it reports a status of Failed. Figure 4-33 is an example of what you might see in the performance monitor log for a resource that is running with problems. One collection interval was successful, followed by a failure, and then followed by a successful one.

Information	Nov 24, 2015 13:03:12.732	HWNPM2123I	Performance data for resource timestamp 11/24/15 01:02:07 PM GMT-07:00 was collected and processed successfully. 2 performance data records were inserted into the database repository.
Warning	Nov 24, 2015 13:04:15.611	HWNPM2124W	Performance data continuity has been broken. The device might have been reset or rebooted. 1 performance data records were discarded.
Information	Nov 24, 2015 13:05:16.129	HWNPM2123I	Performance data for resource timestamp 11/24/15 01:04:15 PM GMT-07:00 was collected and processed successfully. 2 performance data records were inserted into the database repository.

Figure 4-33 Log example of running with problems status

Note: You can view performance monitors from the individual resources in other views of Spectrum Control.

4.6 Fabric administration use cases

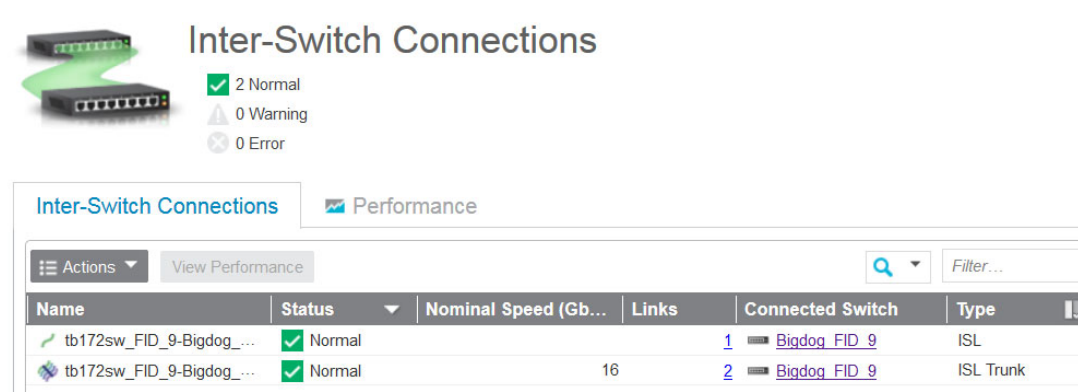
Here are administrative use cases based on monitoring switches and fabrics with Spectrum Control. They include monitoring inter-switch connections, the performance metrics for inter-switch connections, and N_Port Virtualization (NPV) mode switches connected to a fabric.

4.6.1 Viewing inter-switch connections

Inter-switch connections include the following types of connections:

- ▶ Inter-switch link (ISL)
- ▶ NPV Link (a single link between a native fabric switch and an NPV switch)
- ▶ F_Port trunk (also known as an F_port channel)
- ▶ ISL Trunk (also known as an E_port channel)
- ▶ Port channel
- ▶ Inter-chassis link (ICL) trunk

Inter-switch connections are viewable from both the switch and fabric details views. These connections may be between fabric switches, NPV mode switches, or access gateway mode switches. Two types of inter-switch connections, ISL and ISL Trunk, are displayed in Figure 4-34 as an example. You can drill into the links or the connected switch. The number under the *Links* column provides a reference to how many links make up that connection. An ISL is a single link, and an ISL trunk can contain multiple links between the switches.



Name	Status	Nominal Speed (Gb...	Links	Connected Switch	Type
tb172sw_FID_9-Bigdog_...	Normal		1	Bigdog_FID_9	ISL
tb172sw_FID_9-Bigdog_...	Normal	16	2	Bigdog_FID_9	ISL Trunk

Figure 4-34 Inter-switch connections listing

NPV link versus ISL link

An ISL is a single link between two native mode switches, and a single link between a native mode switch and an NPV mode switch is a specific type of inter-switch connection called *NPV Link*. An NPV Link may be a link between a Brocade native mode switch and a Brocade NPV mode switch or a link between a Cisco native mode switch and a Cisco NPV mode switch. Additionally, you may have a Cisco NPV mode switch connected to a Brocade fabric or a Brocade Access Gateway connected to a Cisco fabric. Both are supported options for NPV links. These single links are considered ISL or NPV Links only when they are not part of a Brocade trunk or Cisco port channel.

Figure 4-35 shows the properties for an NPV Link from the inter-switch connections listing for a native mode fabric switch. It provides details for the connected NPV switch and switch port. It was opened by selecting the number in the *Links* column for the NPV Link entry.

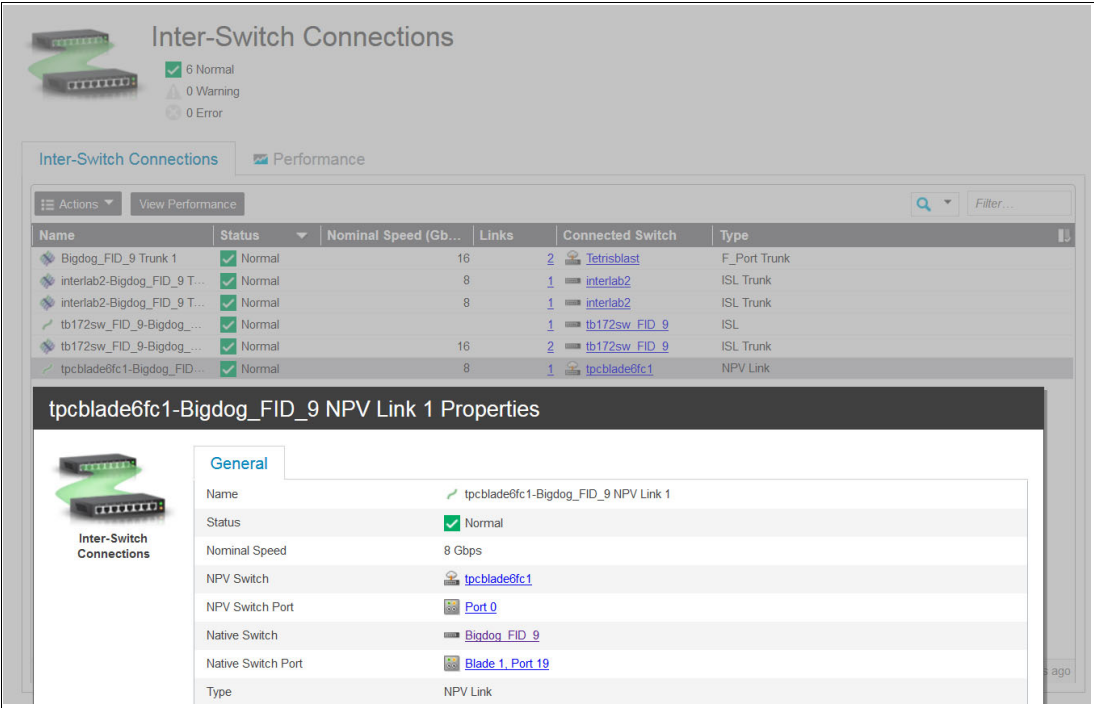


Figure 4-35 NPV Link shown from a native mode fabric switch

F-port trunk versus F-port channel

Whether Spectrum Control displays an inter-switch connection type of F-port trunk versus an F-port channel depends on the switch vendor, but the concept is the same for both. They are each a logical aggregation of links between a native mode fabric switch and an NPV switch. Brocade calls it an F-port trunk and refers to the NPV switch as an access gateway. Cisco calls it an F-port channel. Spectrum Control shows the specific type for the respective switch vendor.

When viewing a list of switches, the links are shown in a column on the right by default. The Properties window that is shown in Figure 4-36 on page 105 was generated by selecting the entry in the *Links* column for the switch named Tetrisblast (the highlighted row in the background). Tetrisblast is an NPV switch and the link that is shown is an F-port trunk.

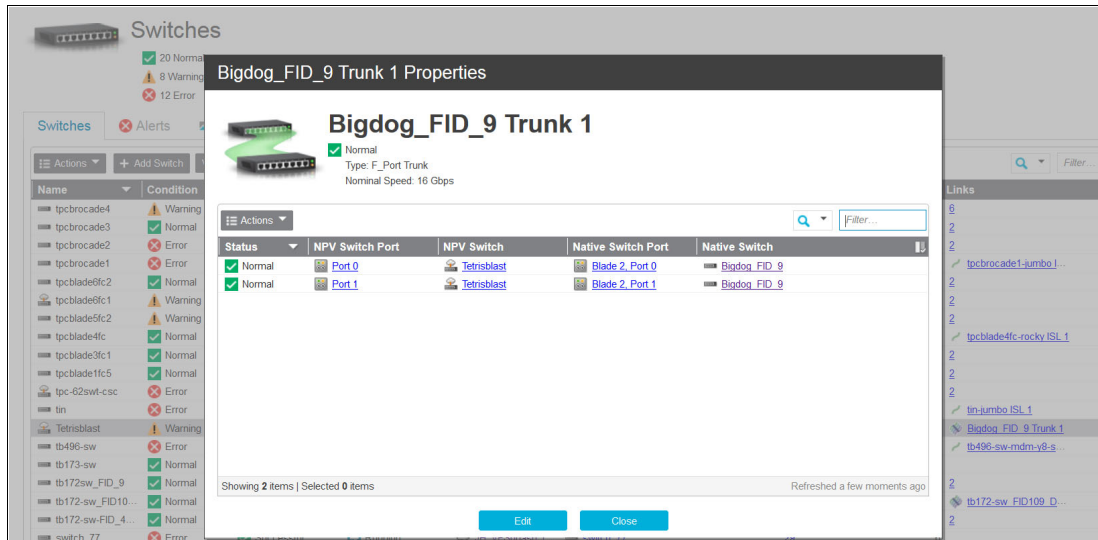


Figure 4-36 F-port trunk properties from an NPV (access gateway) switch

For the F-port trunk in Figure 4-36, the NPV switch and the NPV switch port are shown along with the native switch and native switch port that make up the connection. Each row represents an individual link within the link aggregation for the trunk.

ISL trunk versus E-port channel and ICL trunk

Like the F-port trunks and F-port channels, Spectrum Control displays an inter-switch connection type of ISL trunk versus an E-port channel based on the switch vendor. Both of these types are a logical aggregation of ISLs between two native mode fabric switches, typically between two or more E-port links. Brocade calls the aggregation of these links an ISL trunk and Cisco calls it an E-port channel. Spectrum Control shows the specific type for the respective switch vendor.

Figure 4-37 shows the inter-switch connections for a Brocade native mode switch. From this view, the switches that are listed under the *Switch 1* and *Switch 2* columns indicate which switches are connected. You can select either of those switch names to drill into the switch itself for details, but what you can notice from this window is that the ISL trunk has two links while the ISLs each have one. The ISLs represent a single connection between two switches, and the Brocade ISL trunk or Cisco E-port channel is the aggregation of those ISLs.

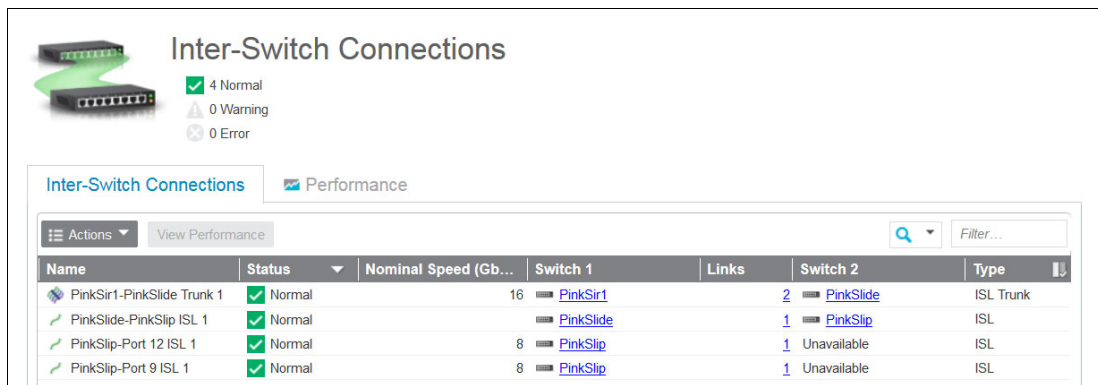


Figure 4-37 Inter-Switch Connections showing ISL Trunk and ISL

For more information about this Brocade ISL trunk and its links, select the number 2 under the *Links* column in Figure 4-37 on page 105 to open the Properties window for the ISL trunk, as shown in Figure 4-38. From here, you can see all of the links making up the Brocade ISL trunk and the specific switch ports where they connect between each switch. From here, you can drill into any of the switch or switch port links for more information.

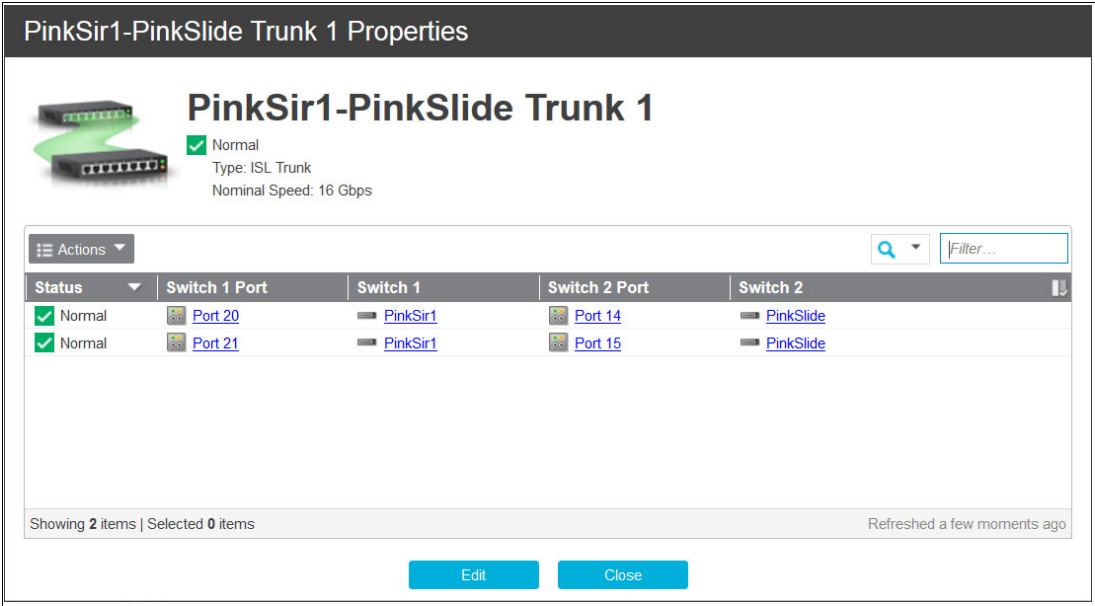


Figure 4-38 Properties for an ISL trunk

Additionally, there is another inter-switch connection type of ICL trunk that you might see in Spectrum Control. This is a logical aggregation of inter-chassis links (ICLs), which are the connections between core routing blades of two Brocade director switches. The core blades have special ICL E_ports that are used only for connections to core blades on other director switches.

Problematic links for inter-switch connections

Figure 4-39 shows an example of two different port channels that are configured on a switch that are recognized as being problematic. No connected switch is listed.

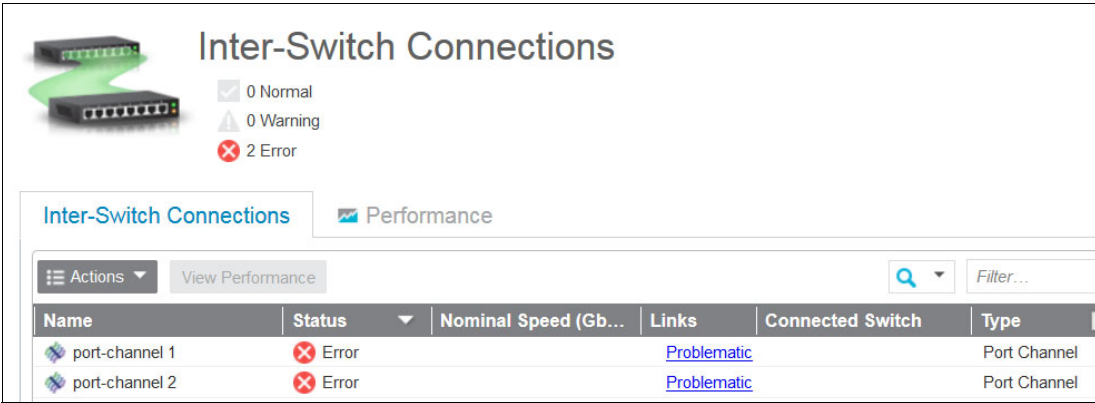


Figure 4-39 Inter-switch connections with problematic links

To understand why a connection is reported as problematic, click **Problematic** under the *Links* column for a particular connection, which opens a Properties window providing additional information about the port channel. Figure 4-40 shows that in this case the issue is because there are no ports in the port channel. You can drill into a list of associated ports to view more details or go to the management interface for the switch to configure the ports and resolve the problem.

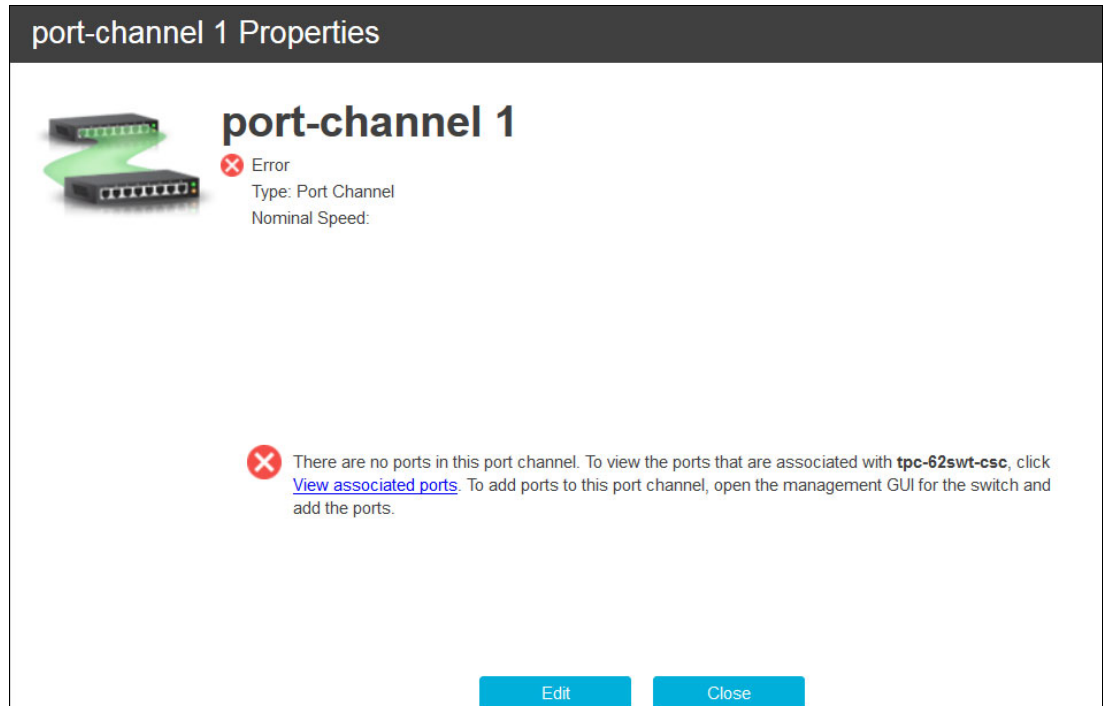


Figure 4-40 Problematic port channel properties

Performance metrics for inter-switch connections

Performance metrics are available for inter-switch connections and can be displayed by clicking the **Performance** tab to display the information, as shown in Figure 4-41.

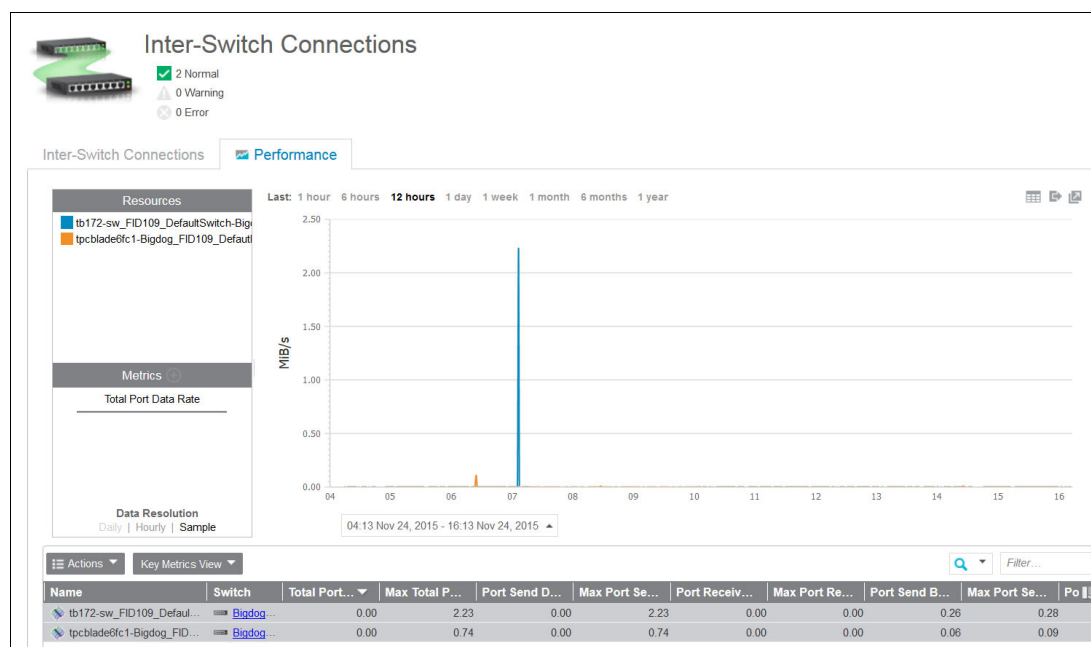


Figure 4-41 Performance tab for inter-switch connections

The performance tab for inter-switch connections works similarly to the performance tab for other resource types. The ISL or NPV Link data is the same as the port data because ISLs and NPV Links represent a port connection. A Brocade trunk or Cisco port channel is an aggregation of the data for the ports that are included in that connection. The performance metrics that are shown for a Brocade trunk or Cisco port channel are the same set of metrics as what is shown for ports.

As you view the information in the performance tab, you can go both up and down from the inter-switch connections. Figure 4-42 shows the options that are available as you right-click an inter-switch connection in the lower table. From here, you can choose to go to the port performance data for the entire switch or drill down into the performance data for individual switch ports.

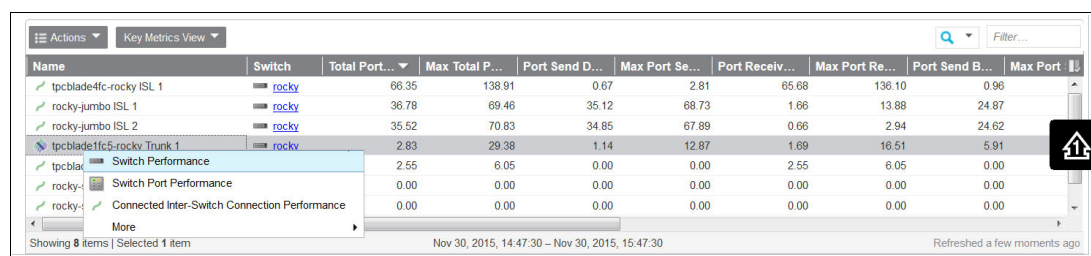


Figure 4-42 Go to and drill down options for inter-switch connection performance data

If you choose the last option in the list from Figure 4-42, **Connected Inter-Switch Connection Performance**, you can go to the performance data from the perspective of the other switch participating in the inter-switch connection. If everything is working well on the two switches, the performance data for the inter-switch connections on each switch should be symmetrical.

For example, the send rates on one switch should match the receive rates on the other, and vice versa. However, you might see more errors being logged on one switch than the other, and that might help you identify a problem. Figure 4-43 shows an example view from the perspective of both switches that are involved in a Brocade trunk after selecting this option.

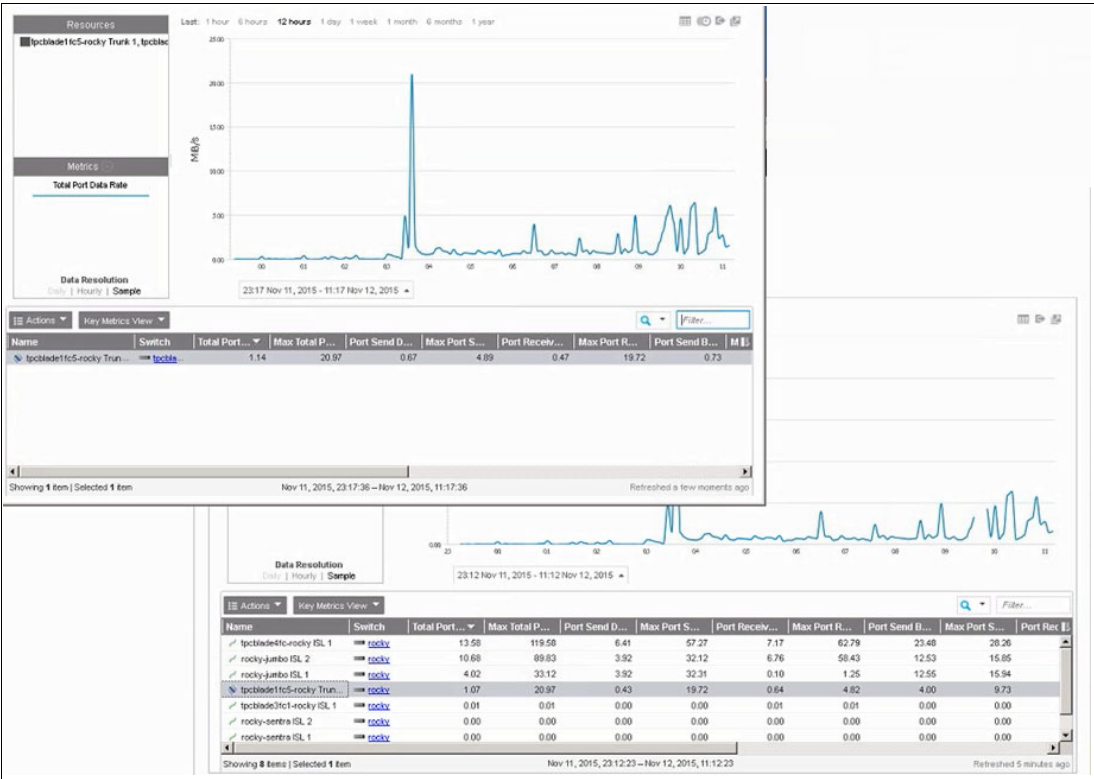


Figure 4-43 Viewing inter-switch connection performance data from both switches

Note: Performance data for a Brocade trunk or Cisco port channel is the aggregation of the performance data for the ports that are included in that trunk or port channel. If you choose to drill down to the switch port performance data for a simple ISL or NPV Link, the same information is shown because the ISL or NPV Link represents a single port to port connection.

Although performance data is not available for the fabric specifically, it is available for the switches within the fabric and you can view the data for the inter-switch connections or ports all at once.

Figure 4-44 shows the inter-switch connection performance data for a single fabric across multiple switches. The table below the chart lists the resources for which performance data is available in the selected time range. Two rows are displayed for each inter-switch connection, where each row represents the performance from the perspective of one of the two switches participating in the connection. (This situation assumes that both switches are part of the fabric and that both had performance monitors running. NPV mode switches are not considered part of the fabric.) The *Switch Name* column in each table row differentiates the switch perspective for the inter-switch connection that is listed in the *Name* column. The performance data is not identical for each side of the connection, but it is symmetrical. For a given trunk, the individual links that are members of the trunk should be load-balanced by the switches such that similar performance and throughput characteristics are seen.

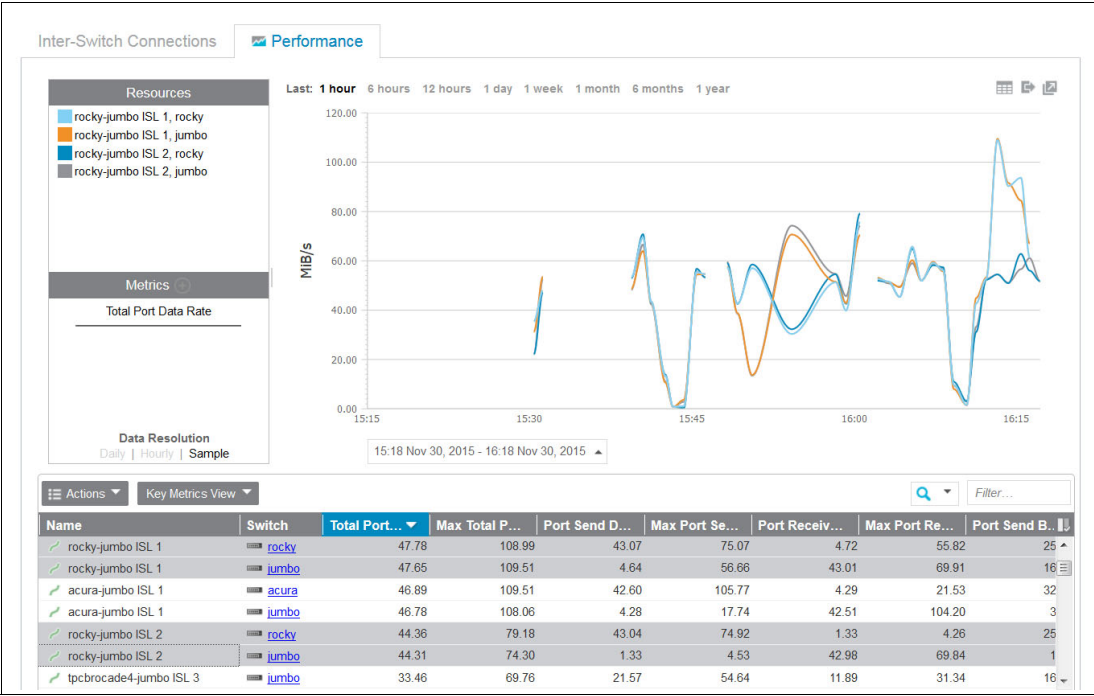


Figure 4-44 Inter-switch connection information for a single fabric

Note: If the inter-switch connection is between a fabric switch and an NPV or access gateway mode switch, you do not see both sides of the inter-switch connection. The NPV mode switch is not part of the fabric and is not shown.

4.6.2 Viewing NPV switch ports and internal routes

In this case, an administrator wants to see where a server is connected to a fabric, particularly the connections for an NPV (N_Port Virtualization) switch. Figure 4-45 on page 111 shows the port listing for an NPV switch. Port 0 and port 1 are members of an F-port trunk, and information about that connection can be found in the *Link* column.

Ports
 19 Normal
 3 Warning
 0 Error
 2 Warning - Acknowledged

Ports Internal Routes Performance

Port Number	Status	Port Type	State	Speed (...)	WWPN	Connected Resource	Connected Port	Link
0	Operational	NP_Port	Online	8	20000005336C2B38	Bigdog_FIO_9	Blade 2_Port 0	Bigdog_FIO_9 Trunk 1
1	Operational	NP_Port	Online	8	20010005336C2B38	Bigdog_FIO_9	Blade 2_Port 1	Bigdog_FIO_9 Trunk 1
2	Operational	G_Port	Enabled but offline	8	20020005336C2B38			
3	Degraded	G_Port	Disabled	8	20030005336C2B38			
4	Operational	G_Port	Enabled but offline	8	20040005336C2B38			
5	Operational	F_Port	Online	8	20050005336C2B38	Unavailable	Unavailable	
6	Operational	G_Port	Enabled but offline	8	20060005336C2B38			
7	Operational	G_Port	Enabled but offline	8	20070005336C2B38			
8	Operational	F_Port	Online	4	20080005336C2B38	Unavailable	Unavailable	
9	Degraded	G_Port	Disabled	8	20090005336C2B38			
10	Operational	G_Port	Enabled but offline	8	200A0005336C2B38			
11	Operational	G_Port	Enabled but offline	8	200B0005336C2B38			
12	Operational	G_Port	Disabled	8	200C0005336C2B38			
13	Operational	G_Port	Disabled	8	200D0005336C2B38			
14	Operational	G_Port	Enabled but offline	8	200E0005336C2B38			
15	Degraded	G_Port	Disabled	8	200F0005336C2B38			
16	Operational	G_Port	Enabled but offline	8	20100005336C2B38			
17	Operational	G_Port	Enabled but offline	8	20110005336C2B38			
18	Operational	G_Port	Enabled but offline	8	20120005336C2B38			

Showing 24 items | Selected 1 item Refreshed a few moments ago

Figure 4-45 Ports listing for an NPV switch

The Internal Routes tab is available only for NPV switches and is shown in Figure 4-46. Click the **Internal Routes** tab to view the routes and connections from a node to a fabric switch as it goes through an NPV switch. Figure 4-46 shows the node ports that are connected to the NPV switch, which ports on the NPV switch to which they are connected, which ports to which traffic is routed, and the connections to the fabric switch ports. These internal routes show you the data path through the NPV switch.

Ports
 20 Normal
 4 Warning
 0 Error

Ports Internal Routes Performance

Node	Node Port	Port Connected to Node	Port Connected to Fabric	Link	Fabric Switch Port	Switch	Fabric
3862812-7825410-IBM	1.FC_Port 4.2	Port 5	Port 0	Bigdog_FIO_9 Trunk 1	Blade 2_Port 0	Bigdog_FIO_9	1000000533802F01
3862812-7825410-IBM	1.FC_Port 4.2	Port 5	Port 1	Bigdog_FIO_9 Trunk 1	Blade 2_Port 1	Bigdog_FIO_9	1000000533802F01
myServer	21000018321459E0	Port 8	Port 1	Bigdog_FIO_9 Trunk 1	Blade 2_Port 1	Bigdog_FIO_9	1000000533802F01
myServer	21000018321459E0	Port 8	Port 0	Bigdog_FIO_9 Trunk 1	Blade 2_Port 0	Bigdog_FIO_9	1000000533802F01

Figure 4-46 Internal routes and data path from a server through an NPV switch to a native fabric switch

Gather additional information about any of the listed resources by drilling into the links.

Note: Performance metrics can be collected for F-port trunks and all other inter-switch connection types, as described in “Performance metrics for inter-switch connections” on page 108.



Service-oriented storage management using IBM Spectrum Control

As a storage provider, you are confronted with managing the constantly changing requirements regarding capacity growth and performance optimization, which are necessary to provide appropriate resources for the business needs of your customers at minimum cost.

This chapter describes how IBM Spectrum Control (Spectrum Control) helps you to accomplish the following tasks:

- ▶ Getting insights into how your storage is consumed.
- ▶ Organizing capacity data around applications and business units.
- ▶ Getting an overview of the resources that are associated with your applications
- ▶ Drilling into resource details for problem determination and performance analysis.

5.1 Getting insights into how your storage is consumed

The new consumer-oriented dashboard design of Spectrum Control that is shown in Figure 5-1, which shows the Dashboard of Spectrum Control, with the Storage Systems pane highlighted, provides a first glance at your storage systems. Underneath each type of storage (Block, File, and Object Storage) is a horizontal bar chart that provides details about the used and free space and the condition of the monitored storage systems in your environment.

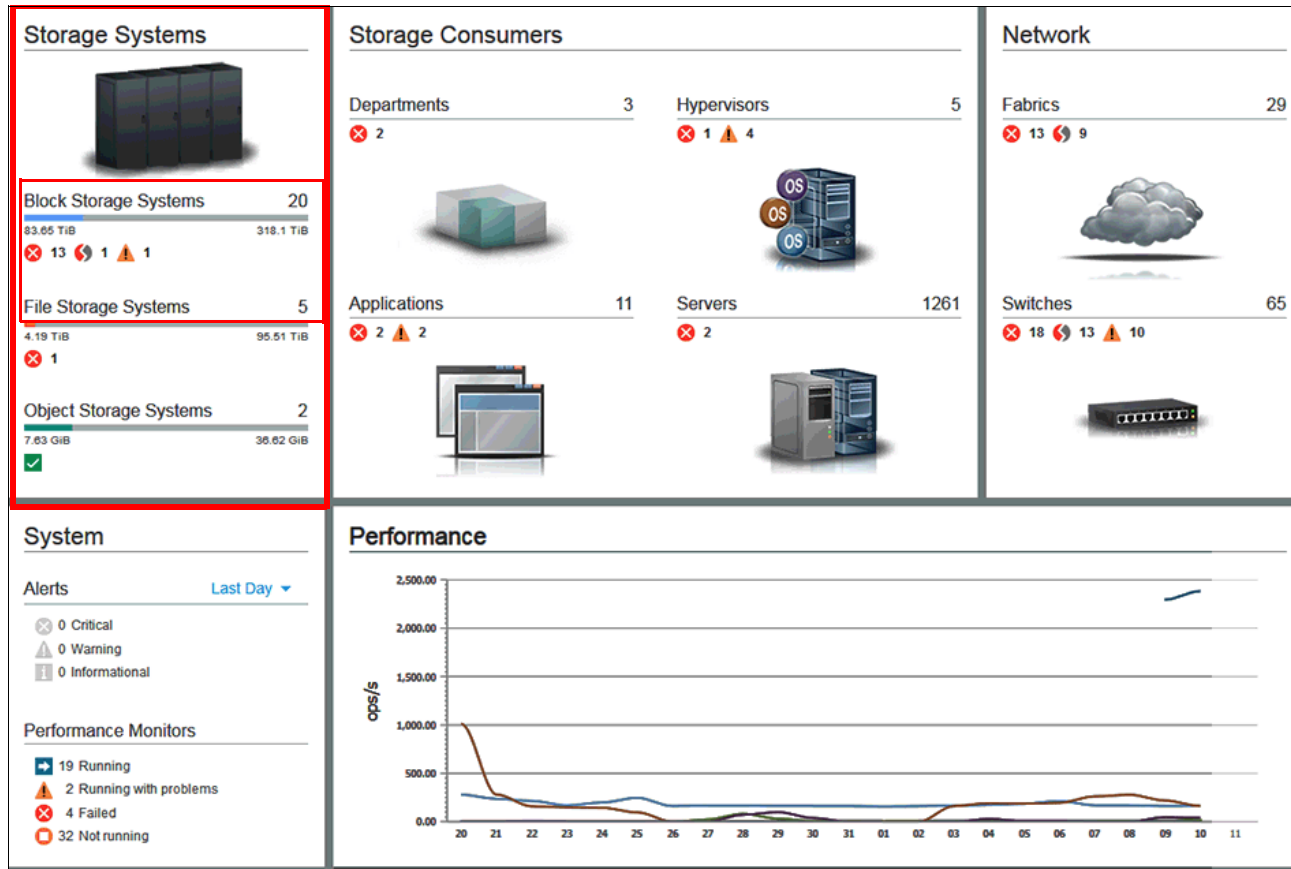


Figure 5-1 Dashboard Design of Spectrum Control

In the upper left pane of the Dashboard, you can see the number of monitored Storage systems. Below that, you can see the number of storage systems that are not in a normal state, which are grouped by status. For example, in Figure 5-1, 13 block storage systems have the status “error”, one is “unreachable,” and one has the “warning” status.

Reference: For more information about the statuses of monitored resources, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1UJk17k>

Review the Spectrum Control IBM Knowledge Center in the “What’s new” section of future versions for further improvements of the Dashboard at the following website:

<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>

5.1.1 How to mitigate the risk of your pool running out of space

In addition to the performance charts, Spectrum Control offers capacity charts that show the growth of pools over time. Figure 5-2 shows an example of how you can initiate viewing historical capacity charts for pools.

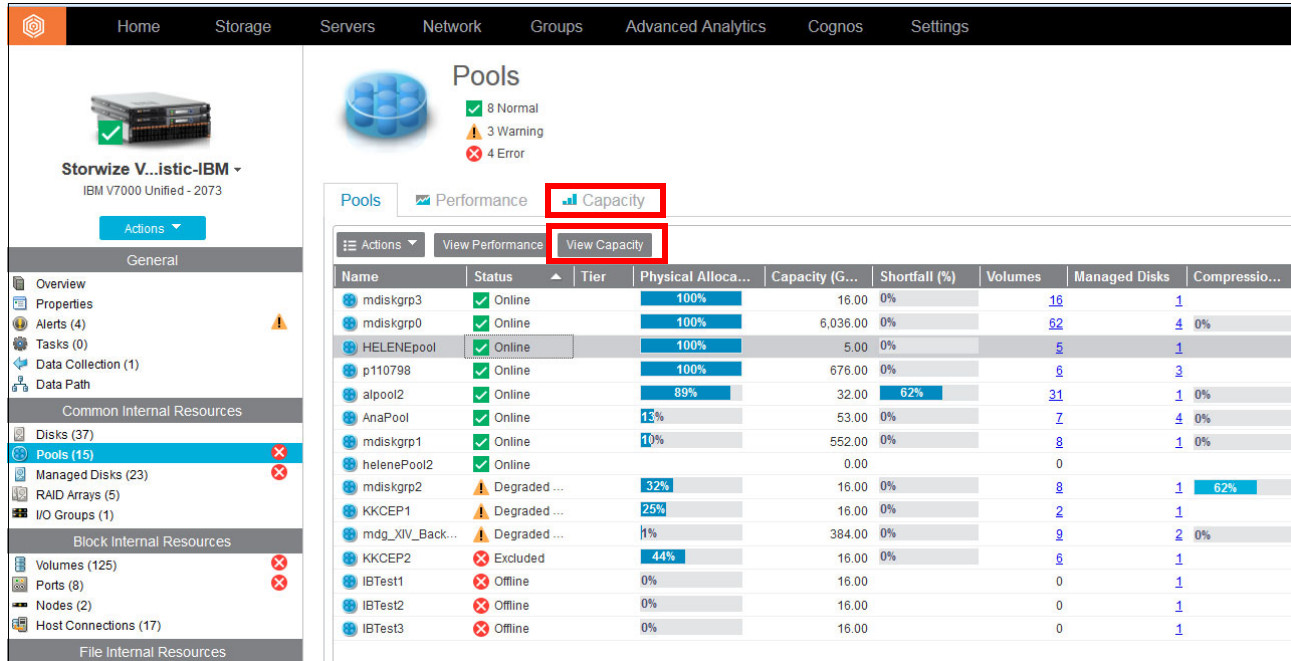


Figure 5-2 Start the historical capacity trends for pools

If you select **Capacity** in Figure 5-2, the pools with the highest Available Pool Space are shown automatically in the chart. If you want to see specific pools, select the pools and click **View Capacity** in Figure 5-2. The capacity chart opens in a new browser window.

Figure 5-3 shows the historical capacity of the fb_flash_1 pool.

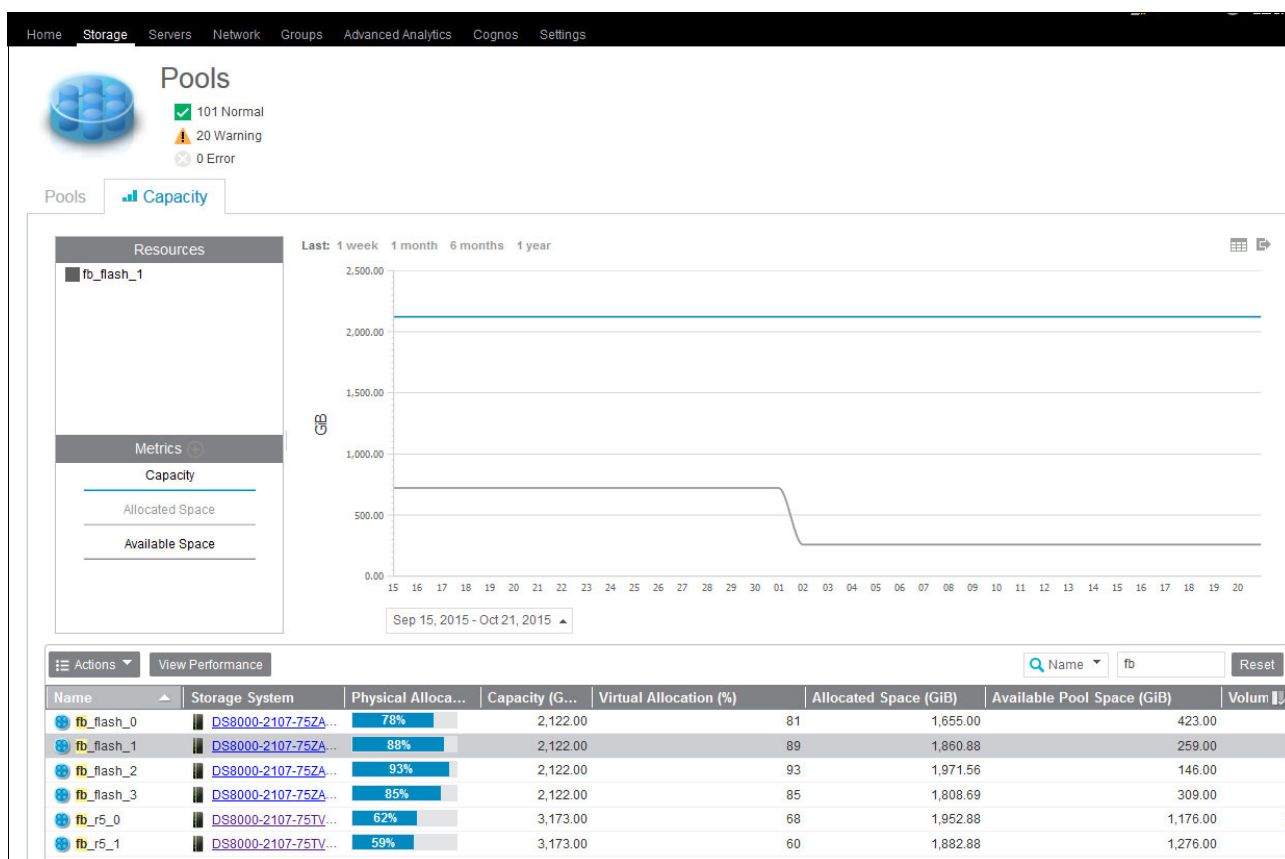


Figure 5-3 Example of a historical capacity chart

Spectrum Control gives you a projection of when your pools run out of space. Based on linear regression of the last 180 days, Spectrum Control calculates when a pool runs out of space (assuming that the capacity behavior does not change drastically). See the Zero Capacity metric, which is shown in Figure 5-4.

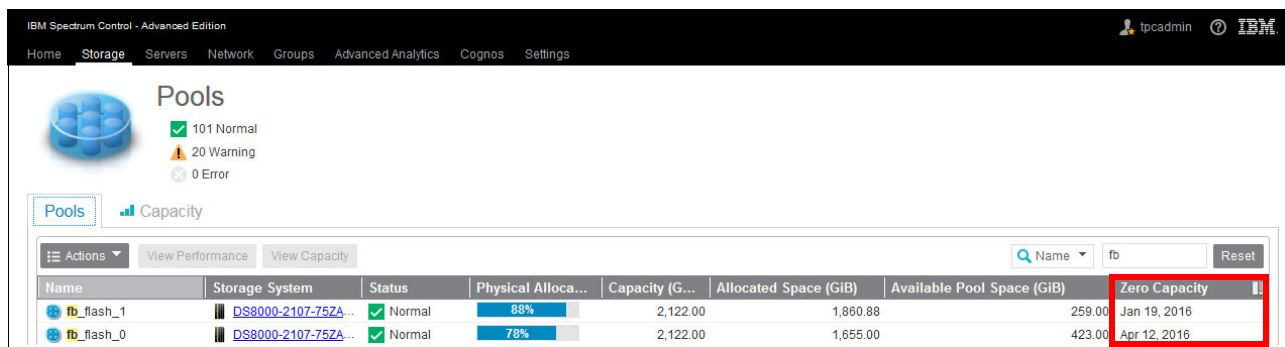


Figure 5-4 Example of the Zero Capacity metric in Spectrum Control

Spectrum Control calculates a *trend* line for the *existing* data by using the least square method to find a line that fits the existing data. This trend line's equation is then used to calculate the Zero Capacity Metric as a *forecast*.

The Zero Capacity metric is available on all pool window, but it is not shown by default. Right-click the table header row and select the **Zero Capacity** metric. You can sort the Zero Capacity Column in an ascending order so that it shows you the pools that will run out of space at the top of the column.

5.1.2 Monitoring your Easy Tier pools

Spectrum Control shows you at a glance what kind of storage of which your pool is composed. With this information, you can easily determine how much space is left in each of the tiers of your IBM Easy Tier® pools or you can watch how Easy Tier is performing.

Easy Tier in IBM SAN Volume Controller and IBM Storwize family

Implementing the IBM Storwize V5000, SG24-8162 describes the Easy Tier operating modes and the resulting Easy Tier statuses that are available in IBM SAN Volume Controller and IBM Storwize family.

Easy Tier operating modes

Here is a list of Easy Tier operating modes in IBM SAN Volume Controller and the IBM Storwize family for MDisk groups (pools):

- ▶ Off: No statistics are recorded and no extents are moved.
- ▶ On: When the Easy Tier function is turned on, Easy Tier measures the I/O activity for all extents. When you have a multi-tiered pool, the extents are migrated dynamically by the Easy Tier processes to achieve the preferred performance.
- ▶ Measured: Easy Tier measures the I/O activity for all extents, but does not move any extents in the storage pool.
- ▶ Auto: Enables Easy Tier functions when the storage pool contains MDisk from more than one tier, and enables automatic rebalancing when the storage pool contains MDisk from only one tier.

Easy Tier status in IBM SAN Volume Controller and IBM Storwize family

Depending on the Easy Tier mode attributes, a storage pool can have one of the following Easy Tier statuses:

- ▶ Active: This status indicates that Easy Tier is actively managing the extents of the storage pool.
- ▶ Balanced: This status applies to homogeneous storage pools and indicates that Easy Tier is actively managing the extents to provide enhanced performance by rebalancing the extents among the MDisk within the tier.
- ▶ Measured: This status means that Easy Tier is constantly measuring the I/O activity for all extents to generate an I/O statistics report.
- ▶ Inactive: When the Easy Tier status is inactive, no extents are being monitored and no statistics are being recorded.

Spectrum Control shows the operation mode and status for Easy Tier pools in IBM SAN Volume Controller and IBM Storwize family, as shown in Figure 5-5.

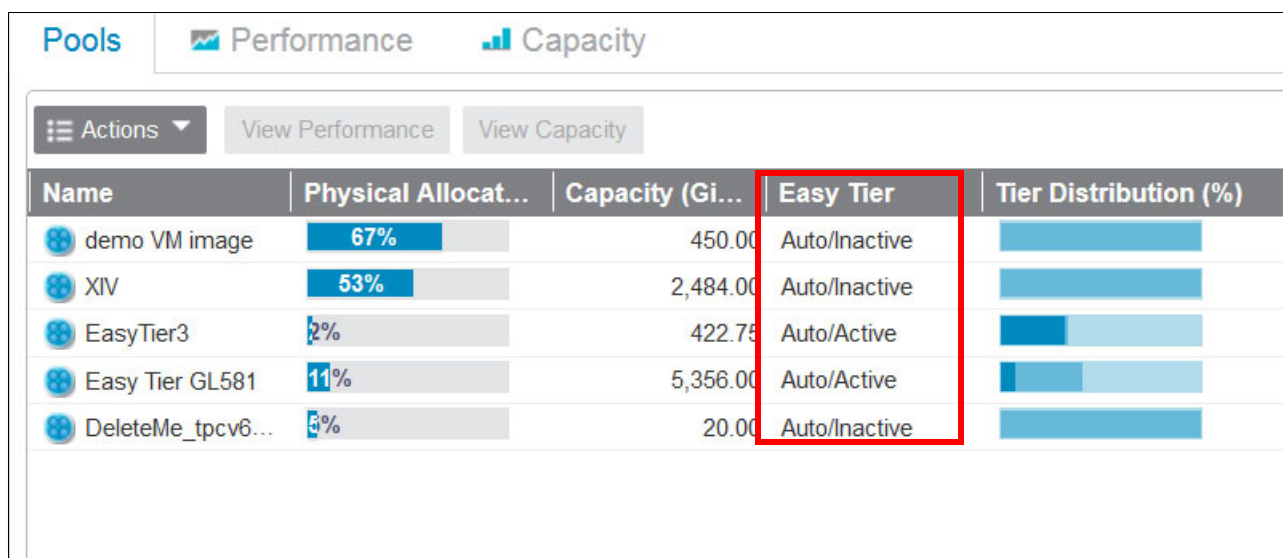


Figure 5-5 Easy Tier mode and status

Table 5-1 shows the mapping of SAN Volume Controller mode and status to the values that are shown in Spectrum Control. Check the Spectrum Control IBM Knowledge Center for changes in future versions.

Table 5-1 Mapping of SAN Volume Controller Easy Tier mode and status to the values that are shown in Spectrum Control

SAN Volume Controller Easy Tier Mode	SAN Volume Controller Easy Tier Status for a single tier	SAN Volume Controller Easy Tier Status for multitier	Spectrum Control value for single tier	Spectrum Control value for multitier
Auto	Balanced	Active	Auto/Active	Auto/Active
On	Balanced	Active	Enabled/Inactive	Enabled/Active
Off	Inactive	Inactive	Disabled	Disabled

Easy Tier control in IBM System Storage DS8000

This section describes how to use the Easy Tier *monitor control* function in a DS8000 storage system.

Easy Tier Monitor Mode control in a DS8000 storage system

The Easy Tier monitor control setting can be set to *automode*, *all*, or *none*, which refers to the volumes to be monitored. The default is *automode*.

The following settings are available:

- **Automode:** Extent statistics for logical volumes in managed extent pools under control of Easy Tier Automatic Mode are updated to reflect the current workload activity.
- **All:** Extent statistics for all logical volumes in managed and non-managed extent pools are updated to reflect the current workload activity.
- **None:** Extent statistics collection is suspended.

IBM DS8000 *Easy Tier*, REDP-4667 describes the DS8000 Easy Tier concepts, design, and implementations.

Easy Tier Automatic Mode control in a DS8000 storage system

As a prerequisite, monitor mode is set to *all*, which is preferable.

This control can be set to *tiered*, *all*, or *none*, and defaults to *tiered*.

The following settings are available:

- Tiered: Automatic extent migrations that are performed by Easy Tier Automatic Mode are enabled only for multitier or hybrid pools.
- All: Automatic extent migrations that are performed by Easy Tier Automatic Mode are enabled for hybrid (multitier) and homogeneous (single-tier) pools.
- None: Easy Tier Automatic Mode is disabled and all automatic extent migrations are deferred for all the extent pools.

Table 5-2 shows the mapping of DS8000 *Monitor Mode* and DS8000 *Automatic Mode* to the values that are shown in Spectrum Control.

Table 5-2 Mapping of DS8000 Easy Tier Monitor and Automatic Mode to Spectrum Control values

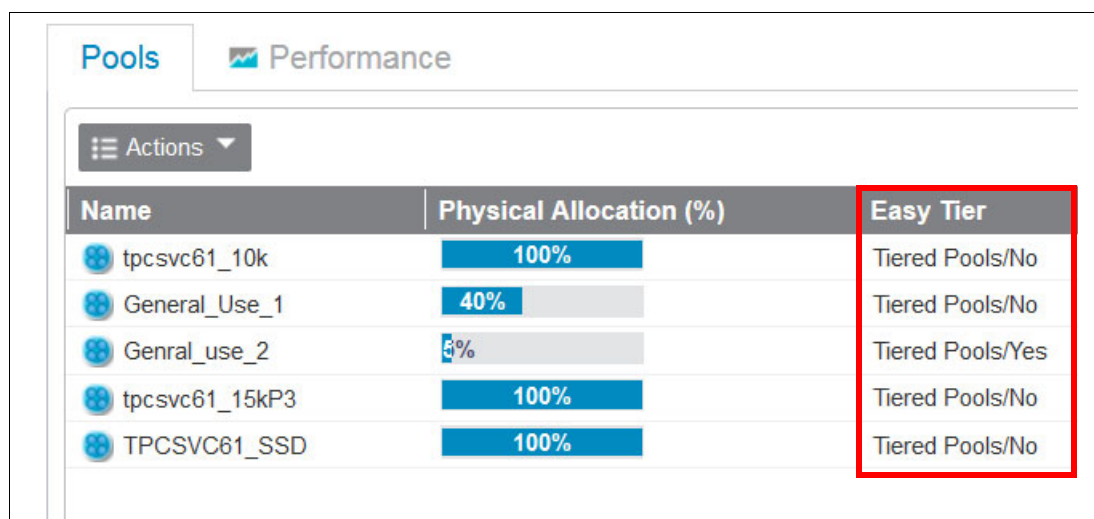
DS8000 Monitor Mode	DS8000 Automatic Mode	Spectrum Control value for single tier	Spectrum Control value for multiple tiers
Automatic Mode	Tiered	Tiered Pools/No	Tiered Pools/Yes
Automatic Mode	All	All Pools/Yes	All Pools/Yes
Automatic Mode	None	Disabled	Disabled
All	All	All Pools/Yes	All Pools/Yes
All	Tiered	Tiered Pools/No	Tiered Pools/Yes
All	None	Disabled	Disabled
None	All	Disabled	Disabled
None	Tiered	Disabled	Disabled
None	None	Disabled	Disabled

Table 5-3 shows a summary of possible Easy Tier values in Spectrum Control for DS8000 Easy Tier pools.

Table 5-3 Spectrum Control Easy tier values for DS8000 pools

Spectrum Control Easy Tier values for DS8000 pools	Explanation
Tiered Pools / Yes	Easy Tier is enabled for multitiered pools only and there is more than one tier in the pool.
Tiered Pools / No	Easy Tier is enabled for multitiered pools only and there is only one tier in the pool.
All Pools / Yes	Easy Tier is enabled for multitier and single-tier pools.
Disabled	Either Easy Tier mode or Monitor Mode is disabled.

Figure 5-6 shows an example of Easy Tier values of DS8000 pools in Spectrum Control.



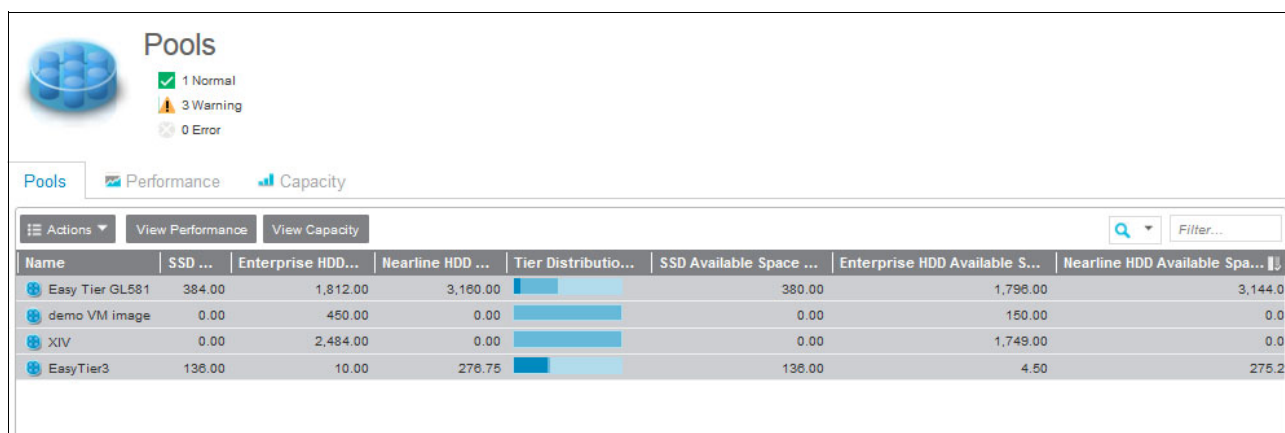
Name	Physical Allocation (%)	Easy Tier
tpcsvc61_10k	100%	Tiered Pools/No
General_Use_1	40%	Tiered Pools/No
Genral_use_2	5%	Tiered Pools/Yes
tpcsvc61_15kP3	100%	Tiered Pools/No
TPCSVC61_SSD	100%	Tiered Pools/No

Figure 5-6 Easy Tier values of DS8000 pools in Spectrum Control

Easy Tier distribution shown in Spectrum Control

Spectrum Control shows the distribution and the capacity values of the volume extents across each Easy Tier pool for volume extents on solid-state drives (SSDs) or flash modules, on enterprise hard disk drives (HDDs), and on nearline HDDs.

In Figure 5-7, Spectrum Control shows the capacity values for up to three different types of extents in a pool.



Name	SSD ...	Enterprise HDD...	Nearline HDD ...	Tier Distributio...	SSD Available Space ...	Enterprise HDD Available S...	Nearline HDD Available Spa...
Easy Tier GL581	384.00	1,812.00	3,160.00		380.00	1,796.00	3,144.0
demo VM image	0.00	450.00	0.00		0.00	150.00	0.0
XIV	0.00	2,484.00	0.00		0.00	1,749.00	0.0
EasyTier3	136.00	10.00	276.75		136.00	4.50	275.2

Figure 5-7 Example of the tier distribution of a pool

These metrics might not be shown by default. If not, right-click the table header row and select the metrics.

The second, third, and fourth columns show the *total amount* of storage space for each different type of extent that can be used by Easy Tier for retiering the volume extents in the pool.

The fifth column shows the tier distribution of the total amount of storage space in percentage (%).

The sixth, seventh, and eight columns show the amount of storage space that is *available* for each different type of extent that can be used by Easy Tier for retiring the volume extents in the pool.

The volumes table (Figure 5-8) shows the tier distribution. If you want to know how tiers are distributed in the volumes of an Easy Tier pool, open the pool's Properties window and click the **Volumes** tab.

Figure 5-8 shows the volume's Easy Tier distribution.

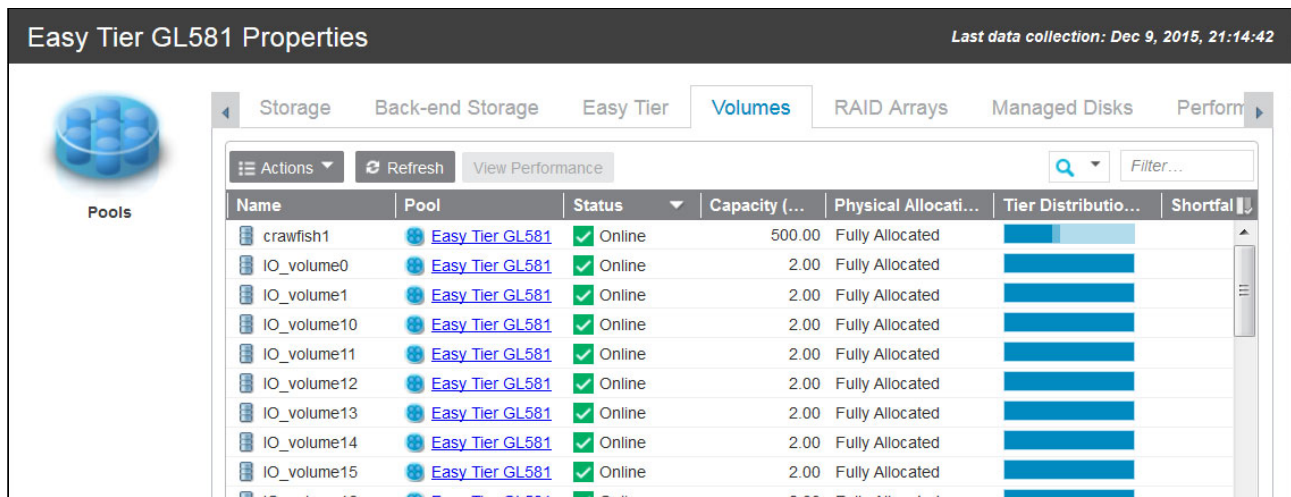


Figure 5-8 Volume's Easy Tier distribution

5.1.3 How to mitigate the risk of overprovisioning

If you are concerned that you will run out of space because of overprovisioning, there are some ways that Spectrum Control can help you to prevent this situation:

- ▶ Computed metrics: Virtual Allocation Percentage and Shortfall Percentage
- ▶ Alerts and thresholds
- ▶ Restriction settings in Service Class Definitions

Virtual Allocation Percentage and Shortfall Percentage

This section describes how the metrics *Shortfall Percentage* and *Virtual Allocation Percentage* are calculated:

The Virtual Allocation Percentage is defined as follows:

The percentage of physical space in storage virtualizer pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

$$\text{Virtual Allocation Percentage} = (\text{Total Volume Capacity} \div \text{Pool Capacity}) * 100$$

Shortfall Percentage is defined as follows:

The percentage of the remaining unallocated volume space in storage virtualizer pools that is not available to be allocated (= "Unallocatable Volume Space").

For more information about these definitions, see the *Storage Systems Capacity Report* topic in the Spectrum Control V5.2.8 IBM Knowledge Center, found at:

http://www-01.ibm.com/support/knowledgecenter/SS5R93_5.2.8/com.ibm.spectrum.sc.doc/fqz0_r_cr_storage_systems_space_rpt.html?lang=en

Figure 5-9 shows the Unallocatable Volume Space.




Figure 5-9 Unallocatable Volume Space, Unallocated Volume Space, and Available Pool Space

Shortfall Percentage = (Unallocatable Space ÷ Unallocated Volume Space) * 100, and
Unallocatable Volume Space = Total Volume Capacity - Allocated Space - Available Space

The Shortfall Percentage cannot exceed 100%.

In summary, Unallocatable Volume Space is the space that you must add to your pool so that all thin-provisioned volumes can grow to their maximum size, and Shortfall Percentage is the relationship of the Unallocatable Space to the Unallocated Volume Space.

Figure 5-10 on page 123 explains the pool metrics that are needed to calculate Shortfall Percentage.




Figure 5-10 Pool metrics

The pool metrics that are shown in Figure 5-10 are displayed in Spectrum Control.

Figure 5-11 shows an example of a pool with thin-provisioned volumes.

Pools Capacity							
Actions View Performance View Capacity General_use_1 Reset							
Name	Capacity (GiB)	Total Volume C...	Allocated Sp...	Available Poo...	Unallocated Vo...	Unallocatable Vol...	Shortfall
General_Use_1	1,480.00	5,215.50	982.50	467.00	4,233.00	3,766.00	89%

Figure 5-11 Example for the Shortfall Percentage metric - thin-provisioned volumes

The following formulas show how the Unallocated Volume Space, Shortfall Percentage, and Virtual Allocation Percentage metrics from Figure 5-11 are calculated in Spectrum Control:

- ▶ Unallocatable Volume Space = 5215.50 GiB - 982.5 GiB - 467 GiB = 3766 GiB
- ▶ Shortfall Percentage = 3766 GiB ÷ 4233 GiB * 100 = 89%
- ▶ Virtual Allocation Percentage = (5215.50 GiB ÷ 1480 GiB) * 100 = 352%

Note: If Virtual Allocation Percentage is less than 100% Shortfall Percentage will be 0%, and there is no risk. When Virtual Allocation Percentage is greater than 100% Shortfall Percentage increases, and there is a risk of running out of space.

For more information about this topic, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/1SJMzQA>

Set alerts to mitigate the risk of overprovisioning

For the following parameters, you can configure alerts to mitigate the risk of overprovisioning:

- ▶ On the pool level:
 - Virtual Allocation %
 - Pool Shortfall %
 - Available Repository Space for DS8000 track space-efficient (TSE) Pools
 - Unallocated Volume Space
 - Unallocatable Volume Space
 - Virtual Capacity Limit
- ▶ On the volume level: Unallocatable Space

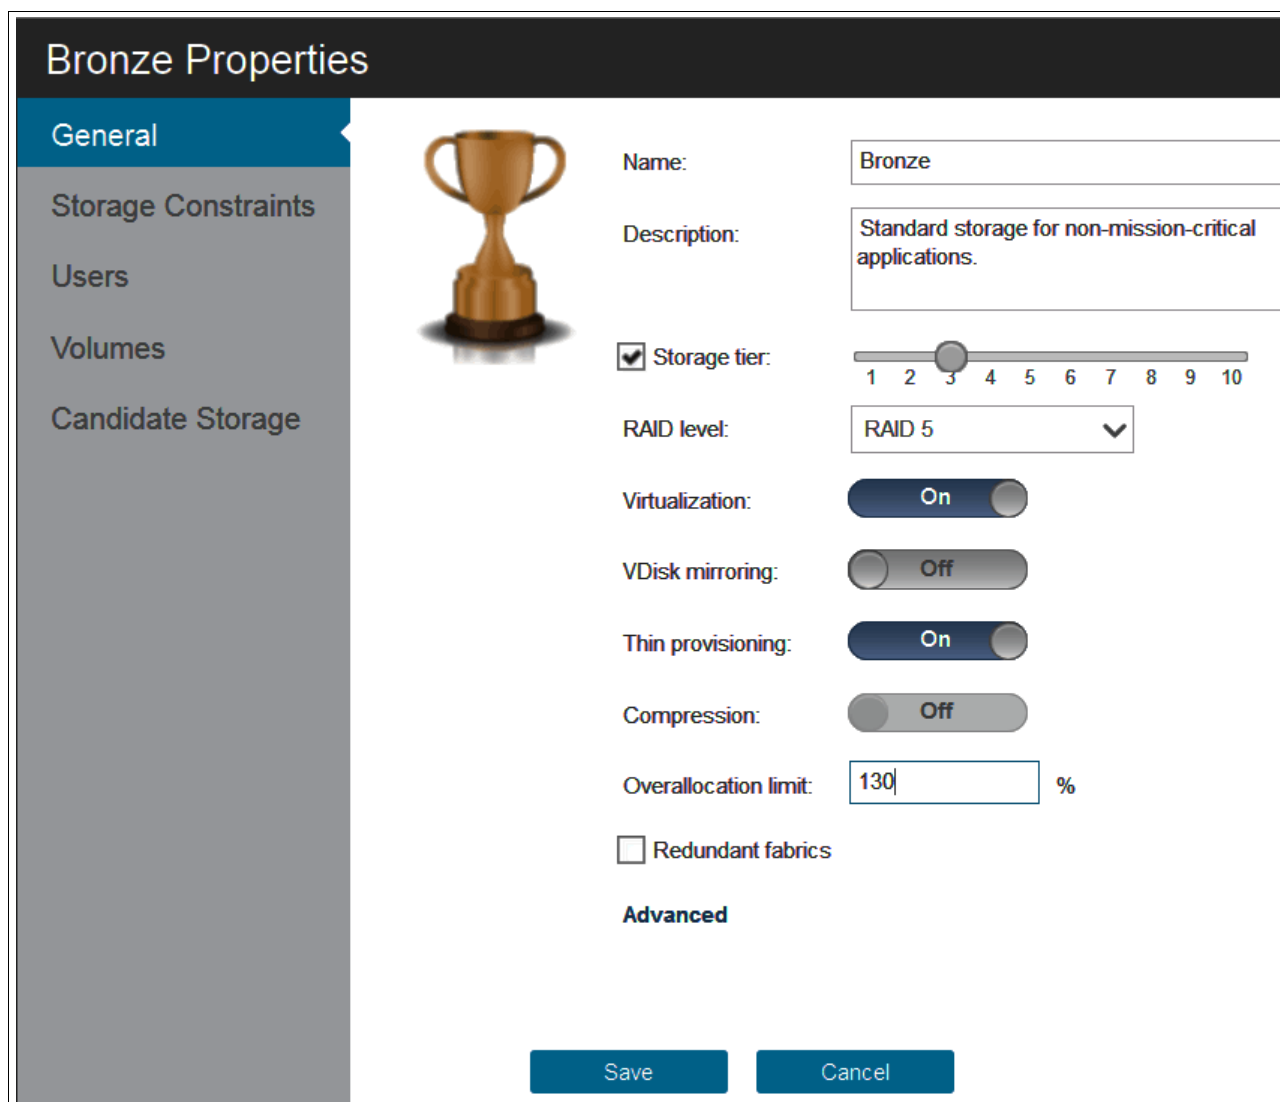
Available Repository Space is defined as the unallocated storage space in the repository for TSE thin-provisioned Volumes. This value is available only for DS8000 thin-provisioned pools.

Virtual Capacity Limit is defined as the maximum amount of virtual storage space that is available to allocate to volumes in the storage pools that are associated with the storage system. You can use this alert to be notified about space-efficient volumes.

For more information, see Chapter 6, “Advanced alert management” on page 177.

Using Spectrum Control service classes to control overprovisioning

You can control the level of overprovisioning with Spectrum Control when you use the service class attribute *Overallocation Limit %* to control your overprovisioning, as shown in Figure 5-12.



Bronze Properties


General

Storage Constraints

Users

Volumes

Candidate Storage



Name: Bronze

Description: Standard storage for non-mission-critical applications.

☒ **Storage tier:** 1 2 3 4 5 6 7 8 9 10

RAID level: RAID 5

Virtualization: On

VDisk mirroring: Off

Thin provisioning: On

Compression: Off

Overallocation limit: 130 %

☐ **Redundant fabrics**

Advanced

Save **Cancel**

Figure 5-12 Service class property *Overallocation Limit %*

If an overallocation limit is specified, Spectrum Control refines the list of candidate storage of a specified service class to include only the storage pools that have an overallocation percentage that does not exceed the overallocation limit.

The overall allocation percentage for a pool is determined by using the following measurements, as shown in Figure 5-13, which shows the Pools window with the Virtual Allocation percentage and the Soft Space parameter:

- ▶ For an XIV system storage pool, the Soft Space is divided by the Capacity.
- ▶ For other storage system pools, use the Virtual Allocation.

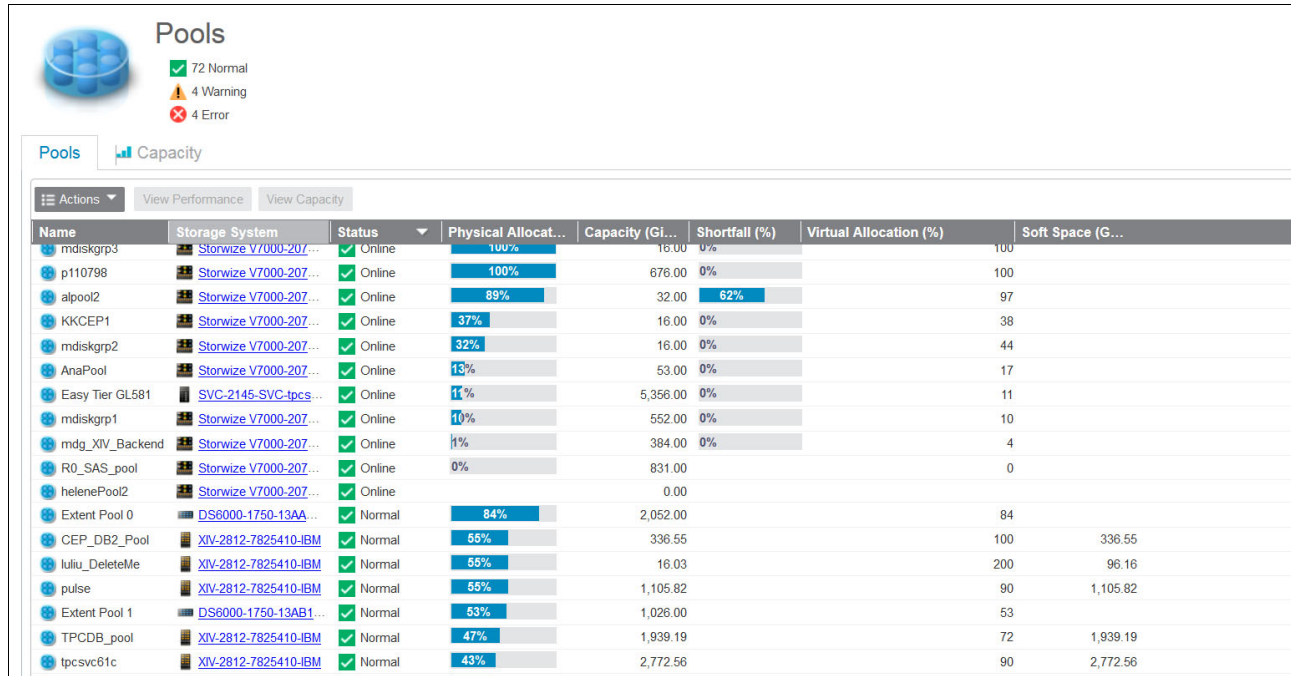


Figure 5-13 Virtual Allocation percentage and the Soft Space parameter

This parameter helps you control, for each service class, the amount that Spectrum Control overprovisions your physical capacity.

For more information about using service classes with Spectrum Control, see *IBM SmartCloud Virtual Storage Center*, SG24-8239.

5.1.4 Creating input for your chargeback report

Spectrum Control offers several options that help you better control the consumption of storage and get the necessary data for charging consumers based on their usage.

Spectrum Control is not meant to provide invoicing capabilities, so it cannot provide a chargeback report. But, Spectrum Control is aware of all the relevant data:

- ▶ Capacities
- ▶ Entity attributes
- ▶ Host mappings
- ▶ Spectrum Control applications
- ▶ User-defined properties

Therefore, Spectrum Control can be used to generate input for your chargeback report. Because input for a chargeback report is not yet considered a chargeback report, call the input report a *showback report*.

Considerations

Customer feedback shows that each customer has his own requirements that must be included in a showback report. Customer requirements depend on whether they are a cost or a profit center, or some of them might want to show their management an analysis of the IT costs because of each consumer.

To create a showback or a chargeback report, think about how you are going to meter, rate, group, and bill the costs. This section lists considerations for each of these categories. Depending on your environment, there are additional items. Here are the considerations:

- ▶ Metering
- ▶ Rating
- ▶ Grouping
- ▶ Example of a showback report instead of billing

Metering

The following prerequisites are required because the capacity influences the charges and functions like encryption (see “Rating”) influence the charges. Some of the information is stored with the volume, and other information is stored with the pool or even with the storage system. Data must be available when the report is generated.

- ▶ Devices must be configured in Spectrum Control and probes must run regularly.
- ▶ You might want to use custom tags, for example, for each pool in Spectrum Control to define costs per capacity or to define the usage of the pool.
- ▶ You can specify Spectrum Control tiers as needed.
- ▶ You can specify and use service classes to provision volumes with Spectrum Control.
- ▶ You can configure Spectrum Control applications and departments, as described in 5.2, “Organizing capacity data around applications and business units” on page 145.

Rating

You must define which entity on which you want to base your rating:

- ▶ Volume Capacity, which includes the following functions:
 - Easy Tier
 - Thin provisioning
 - Compression
 - Encryption
 - RAID level
- ▶ Spectrum Control: Tier
- ▶ Spectrum Control: Service classes

You also must determine when and how to meter the following items:

- ▶ The cutoff date, for example, the first day of a month is easier to establish, but needs definitions for a month with 28, 29, 30, and 31 days.
- ▶ Average usage over a specified time frame.
- ▶ Different rates may be applied if your data center is managed as a cost center versus being managed as a profit center.
- ▶ You might want to apply special rates for specific customers, for example, the more capacity they buy, the cheaper the costs per GB, or for customers within or external to your organization.

Grouping

Grouping can be done by using one of the following methods:

- ▶ Using the Spectrum Control Grouping function, which is described in 5.2, “Organizing capacity data around applications and business units” on page 145.

When you use the Spectrum Control Grouping function, consider the following items:

- Usage of naming conventions.
 - Only application groups or other departments can be added to departments.
 - Avoid double counting: Costs for shared resources can be automatically split between the groups.
- ▶ Naming conventions only.

Example of a showback report

For an example script that generates a showback report, see Appendix B, “Reporting scripts” on page 243.

Note: The script is also available on IBM developerWorks in the IBM Tivoli Storage Productivity Center wiki under Custom Reporting, found at:

<http://ibm.co/1NViY0>

That website is a good place to share custom reports.

The showback report that is shown in Appendix B, “Reporting scripts” on page 243 is provided as an example only and is highly dependent upon your environment (there is *no guarantee* that it will work for your environment and your requirements), but you might want to use this script to get started and modify it for your own needs. SQL skills are required.

The example is written by using native SQL and can be rendered by using the following options:

- ▶ Use native SQL in a batch report with the option of automating. This action generates a CSV file that can be used for further processing.
- ▶ IBM Cognos BI is required.
- ▶ Use Microsoft Excel. ODBC connectivity is required.

Note: If you have not installed the Cognos BI Reporting tool yet, you can do so at any time. For installation guidelines, see the following website:

<http://ibm.co/1SJmZQA>

Here are the limitations of the example script to create a showback report:

- ▶ Works only with block storage.
- ▶ Only tested on IBM systems.
- ▶ Does not include IBM FlashCopy®, Metro Mirror, Global Mirror, or Global Mirror with Change Volume relationships.
- ▶ Shows capacities (no historical data) so it can be used on a fixed date, for example, to run on the first day of every month. (See the information about the cutoff date in “Rating” on page 127).

- ▶ User-defined properties, such as custom tags, are string fields. To use them for calculations, they must be converted into a numerical value.
- ▶ This report is based on volumes and lists host connections (and not servers or hypervisors). Volumes that are assigned to the same hypervisor or server but to different HBAs on the same server or hypervisor are assigned to different host connections.

Note: If you use this report as an import to create other reports, make sure that you avoid double counting. As a preferred practice, create verification reports.

This Accounting report (one row per volume) provides the following data:

- ▶ VOLUMENAME
- ▶ VOLID
- ▶ HOSTLIST
- ▶ CAPACITY
- ▶ ALLOCATED_SPACE
- ▶ USED_SPACE
- ▶ POOLNAME
- ▶ POOL_CUSTOM_TAG1
- ▶ POOL_CUSTOM_TAG2
- ▶ POOL_CUSTOM_TAG3
- ▶ STORAGE_SYSTEM
- ▶ THINPROVISIONED
- ▶ ENCRYPTABLE
- ▶ ENCRYPTED
- ▶ MIRROR_COUNT
- ▶ STATUS
- ▶ RAID_LEVEL
- ▶ IS_COMPRESSED
- ▶ SPECTRUMCONTROL_TIER
- ▶ POOL_EASYTIER
- ▶ POOL_EASYTIER_STATUS
- ▶ STORAGESUBSYSTEM_CUSTOM_TAG1
- ▶ STORAGESUBSYSTEM_CUSTOM_TAG2
- ▶ STORAGESUBSYSTEM_CUSTOM_TAG3
- ▶ APPLIST
- ▶ VOL_TIER_CAPACITY
- ▶ CURRENT_DATE

Figure 5-14 shows some of the fields that you receive from the script.

VOLUMENAME	HOSTLIST	NUMHOSTS	CAPACITY	ALLOCATED_SPACE	USED_SPACE	POOLNAME
*BladeTestLUN5	BLADE_host218_Indigo-1,BLADE_host.	3	4	0	0	TPC_Pool2
*TPC_VOL0	BLADE223_LS21	1	30	19	19	TPC_Pool1
*ThinProv_Vol1	tpc232	1	20	0	0	TPC_Pool2
Bookstore_Wholesale_Vol1	tpc232	1	10	10	10	TPC_Pool2
*BladeTestLUN2	BLADE_host218_Indigo-1,BLADE_host.	3	4	0	0	TPC_Pool2
mhctest_0402_4	WIN-D2VMKKUITNA	1	1	1	1	TPC_Pool2
TPC_vDisk2_DS8K	tpc	1	250	250	250	TPC_Pool1
mhctest0402_3	WIN-2I0P2DE19BJ	1	1	1	1	TPC_Pool2
V_150408_082440	tpc	1	1	1	1	TPC_Pool2
*BladeTestLUN3	BLADE_host218_Indigo-1,BLADE_host.	3	4	0	0	TPC_Pool2
TPC_Vol_ZE1	tpc232	1	20	20	20	TPC_Pool2
*BladeTestLUN0	BLADE_host218_Indigo-1,BLADE_host.	3	4	0	0	TPC_Pool2
TPC_vDisk_DS8K	tpc	1	200	200	200	TPC_Pool2

Figure 5-14 Example of a showback report output

Figure 5-15 also shows some additional columns and fields that are received from the example showback report output.

VOLUMENAME	POOLNAME	POOL_CUSTOM_TAG1	THINPROVISIONED	UNT	MIRROR_CO	IS_COMPRESS	SPECTRUMCON	ED	TROL_TIER	POOL_EASYTIER
*BladeTestLUN5	TPC_Pool2	50.00	True		1	No				2 auto
*TPC_VOL0	TPC_Pool1	34	True		2	No				2 auto
*ThinProv_Vol1	TPC_Pool2	50.00	True		1	No				2 auto
Bookstore_Wholesale_Vol1	TPC_Pool2	50.00	False		1	No				2 auto
*BladeTestLUN2	TPC_Pool2	50.00	True		1	No				2 auto
mhtest_0402_4	TPC_Pool2	50.00	False		1	No				2 auto
TPC_vDisk2_DS8K	TPC_Pool1	34	False		1	No				2 auto
mhtest0402_3	TPC_Pool2	50.00	False		1	No				2 auto
V_150408_082440	TPC_Pool2	50.00	False		1	No				2 auto
*BladeTestLUN3	TPC_Pool2	50.00	True		1	No				2 auto
TPC_Vol_ZE1	TPC_Pool2	50.00	False		1	No				2 auto
*BladeTestLUN0	TPC_Pool2	50.00	True		1	No				2 auto
TPC_vDisk_DS8K	TPC_Pool2	50.00	False		1	No				2 auto

Figure 5-15 Additional columns of the example showback report output

The script's result shows volumes that are mapped to a host (assuming that only volumes with host connection are used for an import to a charge back report). If you want to see volumes with no host connections, replace the "inner join" text that is shown in Figure 5-16, which shows how to change the script to also get volumes with no host connections, with the "left outer join" in "t1".

```

with t1 as (
SELECT cast(AL1.name as VARCHAR(70)) as valid, Strip(AL2.UNKNOWN_HOST_NAME) as hostname
FROM TPCREPORT.STORAGEVOLUME AL1 inner join TPCREPORT.STORAGEVOLUME_PATH AL2 on (AL1.STORAGE_VOLUME_ID=AL2.STORAGE_VOLUME_ID)
and (AL1.BACKEND_DISK_ID is NULL) and (AL1.DETECTABLE = 'True')

group by AL1.name, AL2.UNKNOWN_HOST_NAME ),

```

Figure 5-16 Code snippet of the SQL query

The script shows the volumes and their VDisk mirrors. To get only the primary volumes, remove the two dashes in front of AL6.IS_PRIMARY=1 at the end of the temporary table t4 (Figure 5-17, which shows how to change the script to get only primary volumes).

```

, AL1.RAID_LEVEL
, AL1.IS_COMPRESSED
, AL4.POOL_TIER_ID as TPC_TIER
, AL1.BLOCK_SIZE
FROM TPCREPORT.STORAGEVOLUME AL1 LEFT OUTER JOIN TPC.T_RES_VDISK_COPY AL6 ON (AL1.STORAGE_VOLUME_ID=AL6.SVID)
and (AL1.DETECTABLE = 'True')and (AL1.BACKEND_DISK_ID is NULL)
left outer join TPCREPORT.STAGESUBSYSTEM AL3 on AL1.SUBSYSTEM_ID=AL3.SUBSYSTEM_ID
left outer join TPCREPORT.STORAGEPOOL AL2 on AL6.POOL_ID=AL2.POOL_ID
left outer join TPC.T_RES_STORAGE_POOL AL4 on AL2.POOL_ID=AL4.POOL_ID
-- and AL6.IS_PRIMARY=1
),

```

Figure 5-17 Another SQL code snippet

Then, you might want to see the Mirror_Count. Therefore, remove the dashes, as shown in Figure 5-18 on page 131.

```

select
t4.Volumename
, t4.volid
, case
when tnum.numservers = 0 then 'No Hosts mapped'
else substr (t3.hostname,1,length(t3.hostname)- 1)
end as Hostlist
, tnum.numservers as NUMHosts
, case
-- when (t4.StorageSubsystemname1 like 'SVC%') or (t4.StorageSubsystemname1 like 'Storwize%') then t4.MIRROR_COUNT
-- else 1
-- end as MIRROR_COUNT
, t4.Capacity
, t4.Allocated_Space
, t4.Used_Space
, t4.Pool_ID

```

Figure 5-18 A third SQL code snippet

Using native SQL in a batch report with the option of automating

The following section describes the steps to use native SQL in a batch report to automate the script to run on a regular basis.

Complete the following steps:

1. Obtain the showbackexport.sql and runshowbackexport.bat scripts, which can be found in Appendix B, “Reporting scripts” on page 243.
2. If you want to use these scripts outside of your Spectrum Control Server, then you must configure your database in the IBM DB2 client first, as described in “Instructions for creating a rollup report” on page 138.
3. Store the files in a location without a blank in the path, for example, to C:\showback.
4. Update the connection, path details, and the credentials in runshowbackexport.bat.
5. Update the path in runshowbackexport.bat, if necessary.
6. Now you can run or schedule runshowbackexport.bat to obtain a CSV file as an import for your chargeback on a regular basis.

Reference: The showbackreport.sql script uses the DB2 **EXPORT** command. For more information about this command, go to the following website, where you find information about how to specify the delimiter and to format data:

<http://ibm.co/1PDayRA>

Using reports in IBM Cognos Business Intelligence

For basic instructions about how to work with Cognos Business Intelligence (BI) Reports within Spectrum Control, see *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

To use reports IBM Cognos BI, complete the following steps:

1. Obtain the showback_Cognos_xml.txt file, which is found in “Sample script for Cognos” on page 253.
2. Copy the whole script into your clipboard.
3. Start Cognos Report Studio.
4. Choose any package, for example, Spectrum Control Capacity and Relationship Package.
5. Create a blank report.

6. Click **Tools** → **Open Report from Clipboard**, as shown in Figure 5-19.

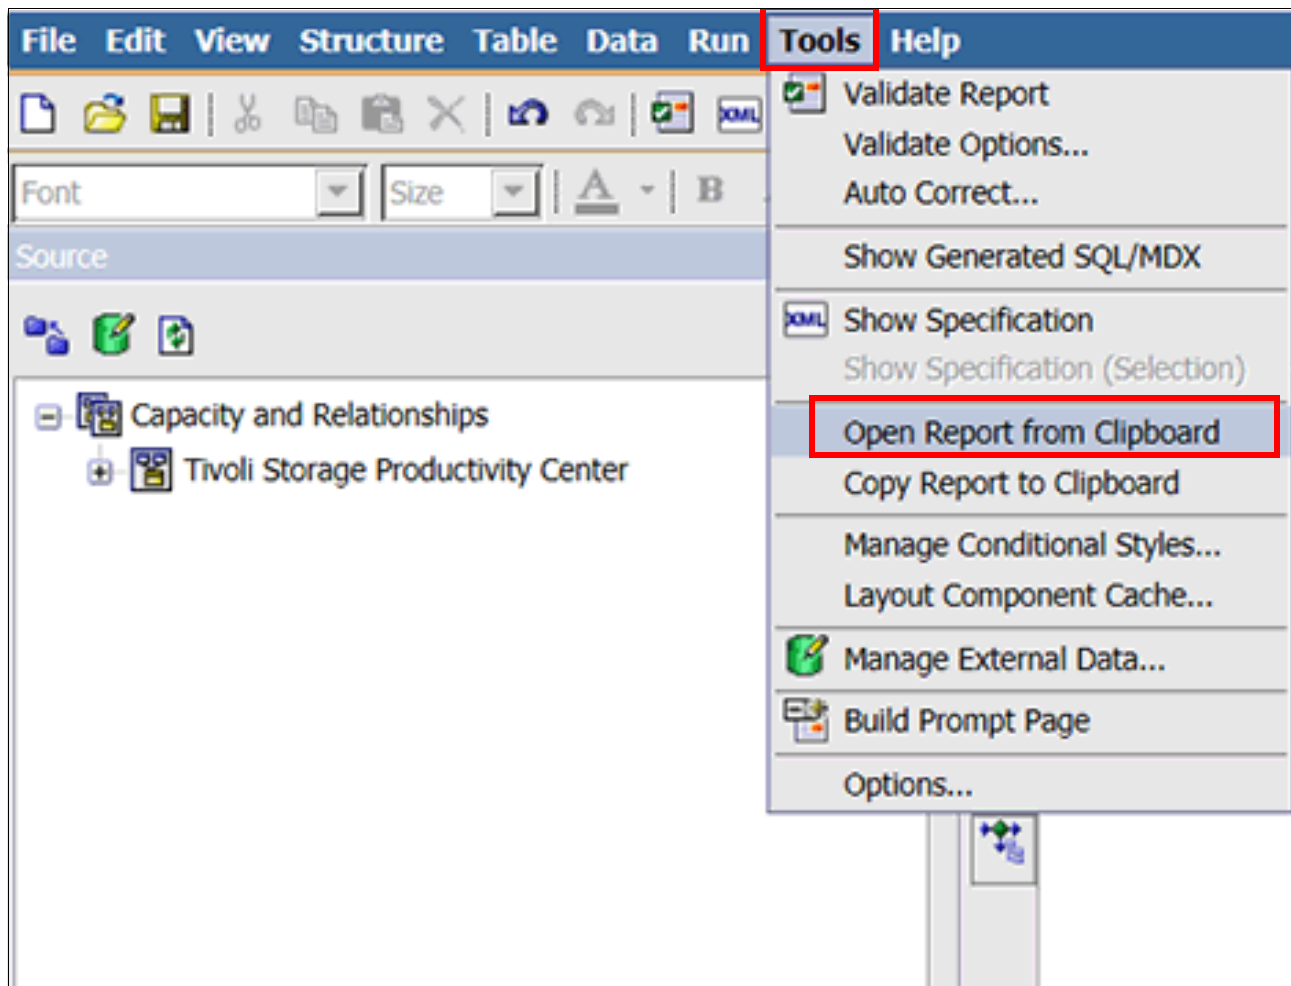


Figure 5-19 Open Report from Clipboard function

7. Save, run, or schedule your report.

Note: Using IBM Cognos Report Studio gives you the opportunity to schedule the report. You also can have the report output automatically sent to specified recipients or saved to the file system.

If you are using Spectrum Control, or a Tivoli Storage Productivity Center version earlier than Version 5.2.8, you might need to import the showback.sql file, which can be found in “Sample script with native SQL statement” on page 244, into an SQL Query. To do so, see Figure 5-31 on page 142 and steps 3a on page 142 to 3l on page 145 in “Instructions for creating a rollup report” on page 138.

Using Microsoft Excel

Databases are excellent for storing information, but sometimes you want to use or analyze the data in Excel. You can use the Data Connection wizard to create a dynamic connection between a DB2 database and your Excel workbook. To accomplish this task, complete the following steps.

Note: The following description is based on Microsoft Excel 2013 and might be slightly different for other versions of Microsoft Excel.

1. Obtain the showback.sql script, which is found in “Sample script with native SQL statement” on page 244.
2. Configure the Spectrum Control database TPCDB in the DB2 client, as described in “Instructions for creating a rollup report” on page 138
3. Set up your ODBC connection to the Spectrum Control database TPCDB. For Microsoft 64-bit systems, see the following website:
<http://www.ibm.com/support/docview.wss?uid=swg21384435>
4. Open a blank workbook in Microsoft Excel.
5. Click the **Data** tab, and then click **From Other Sources** → **From Data Connection Wizard**, as shown in Figure 5-20.

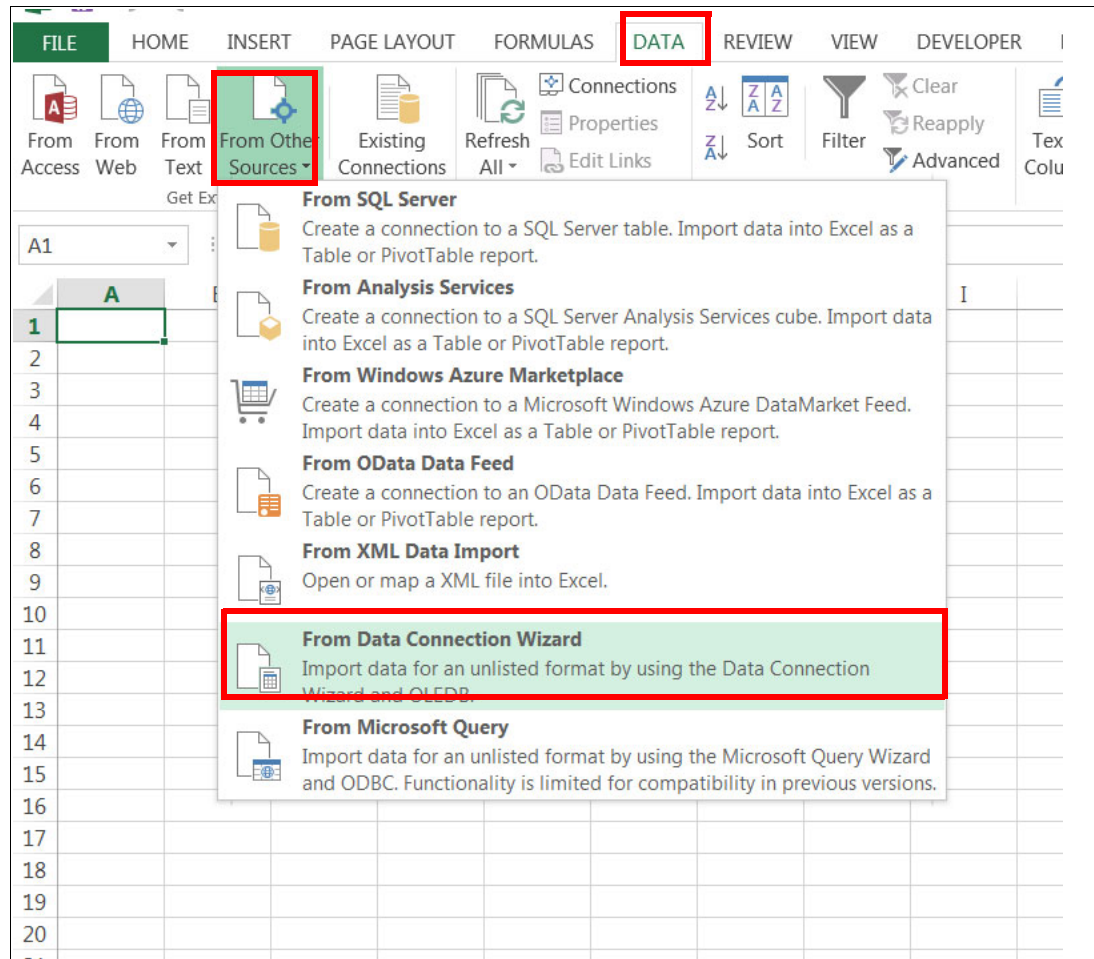


Figure 5-20 How to configure a DB2 database in Microsoft Excel

6. Click **Other/Advanced** and then **Next**.
7. In the next window, select **IBM OLE DB Provider for DB2 - DB2COPY1**.

8. Click the **Connection** tab of the Data Link Properties window. Select Existing data source and select TPCDB171 from the drop-down menu. Enter the correct credentials, and click **Test Connection**, as shown in Figure 5-21.

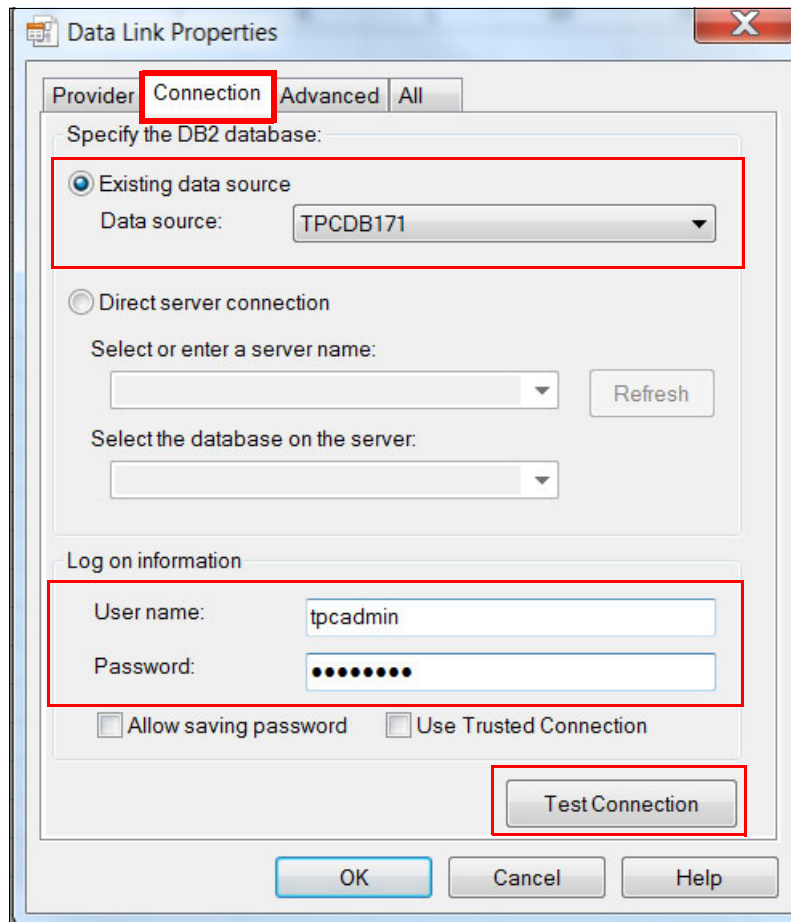


Figure 5-21 Specify the Spectrum Control TPCDB in Microsoft Excel

9. Clear **Connect to a specific table** in the Data Connection wizard, as shown in Figure 5-22 on page 135. Click **Next**.

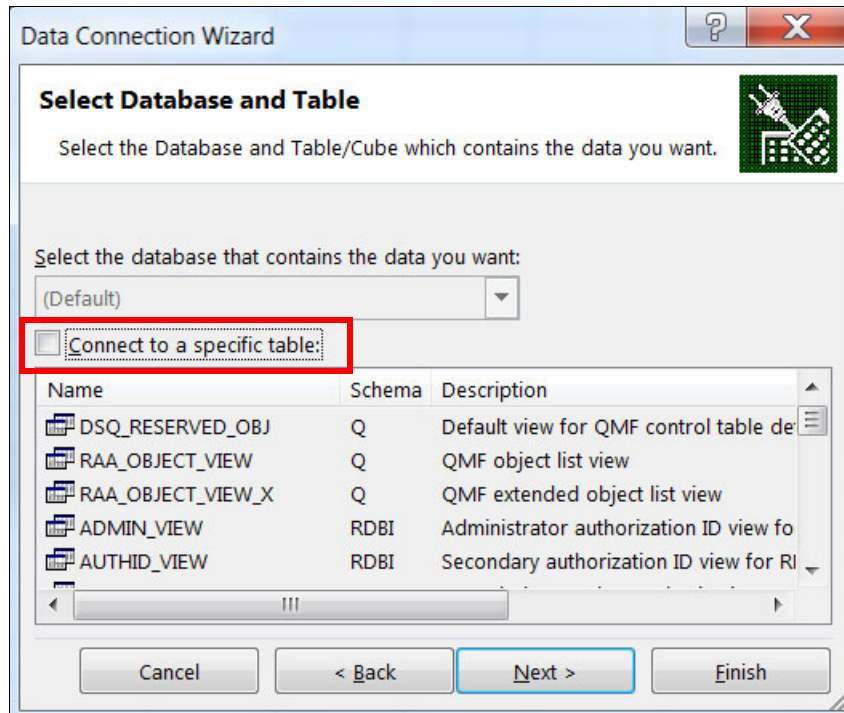


Figure 5-22 Select Database and Table in Microsoft Excel

10. Specify the Data Connection file name, as shown in Figure 5-23. Click **Finish**.

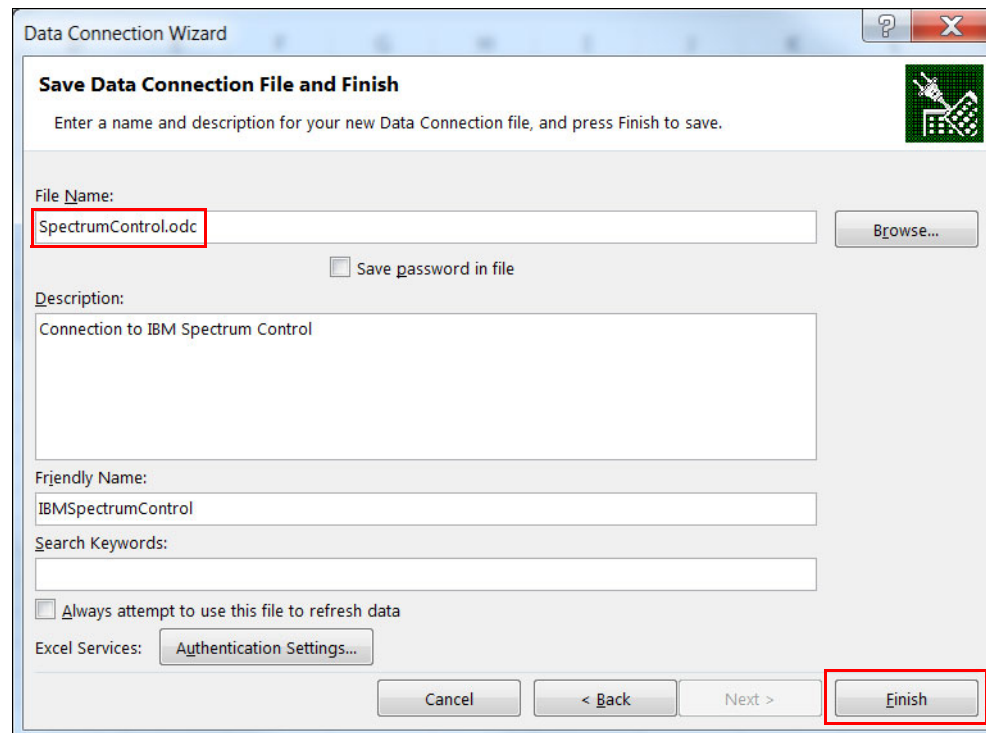


Figure 5-23 An example for specifying the Data Connection file name

11. If prompted, enter the credentials again.

12. Without selecting a table, click **OK** in the Select Table window.

13. In the Import Data window, click **Properties**, as shown in Figure 5-24.

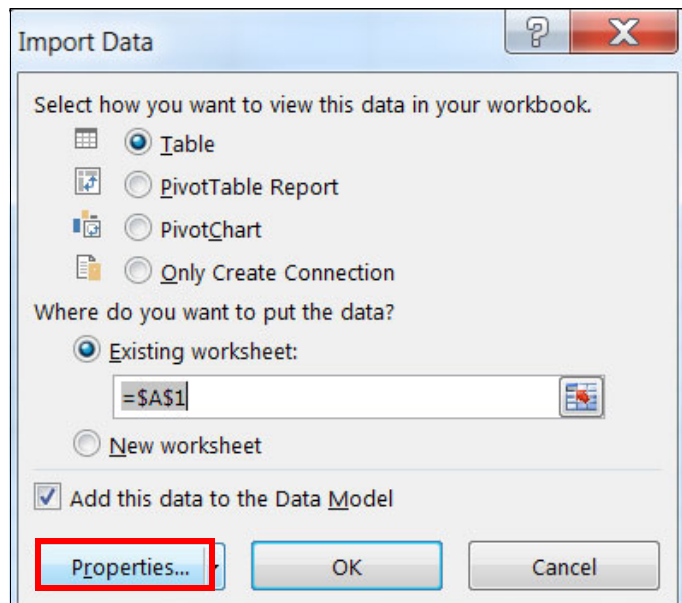


Figure 5-24 Import Data window

14. In the Connection Properties window (Figure 5-25 on page 137), click the **Definition** tab, and type SQL in to the Command type field. Copy and paste the content of the showback.sql file in to the Comment field.

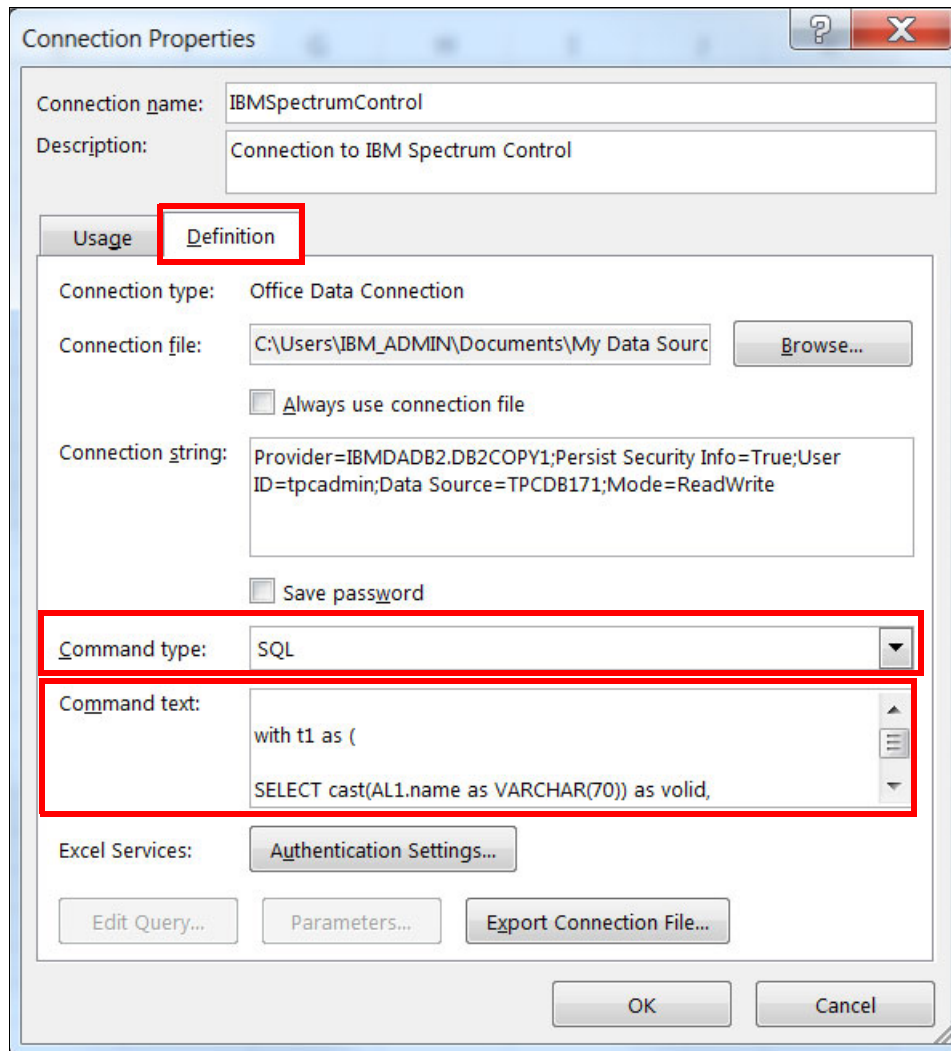


Figure 5-25 Specify the SQL statement in Microsoft Excel

15. Click **OK** twice. If prompted, add your credentials again.

The showback report is now shown in Microsoft Excel. You might want to save this report. The next time that you open the report, click the **Data** tab, and click **Refresh**.

5.1.5 Rollup reports across different Spectrum Control servers

This section describes the scenario where you manage more than one Spectrum Control server and want to get one report with data from multiple Spectrum Control servers.

This task is possible by using the Report Studio of the Cognos Business Intelligence (*Cognos BI*) Reporting tool. The Cognos BI reporting tool is available as an optional component of Spectrum Control. If you have not installed the Cognos BI Reporting tool yet, you can do so at any time. For installation guidelines, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1S4nLQM>

In Spectrum Control, Cognos BI uses the local database (TPCDB) as a data source for creating reports. The configuration of the local TPCDB database (and in DB2 and Cognos) is done during installation. The basic idea is to use the Spectrum Control server that has Cognos installed to use the local TPCDB and TPCDBs from other Spectrum Control servers as data sources. TPCDBs from other Spectrum Control servers must be configured manually to be accessed by DB2 and Cognos BI.

Depending on how you configure remote databases (TPCDBs) to Spectrum Control Cognos BI, you can use them differently:

- ▶ Configure remote databases so that users can choose which database they want to run the report against. In this case, they are prompted to choose the database when running the report. With this configuration, predefined reports or custom reports that use the Spectrum Control Packages can be used to run against local and remote databases. With this configuration, users can select only one database per report.
- ▶ Configure remote databases so that you can define a database per query. You can use this configuration to create and use a union query to combine multiple queries that can run against different databases in the same report. The advantage is that one report can contain data from multiple databases, so this method can be used to create a rollup report. The disadvantage is that you cannot use the predefined reports or the Spectrum Control Packages. Therefore, you must use native SQL for creating queries against multiple databases.

For more information about these setups, see the following website:

<http://ibm.co/1MBEwzd>

Instructions for creating a rollup report

This section describes how to configure a second Spectrum Control database and run a report across two Spectrum Control databases. You must repeat these steps to configure more Spectrum Control databases.

1. Configure your remote database.

On the Spectrum Control server where you want to run the report, you must add the remote TPCDB database to the system database catalog and assign an alias to the database. For example, you can catalog the remote TPCDB database as `tpcdb02`.

For more information about how to configure your database, see the following websites:

– <http://ibm.co/10HCYtb>

– <http://ibm.co/1LcYPkB>

Example 5-1 shows the commands that are used to configure your database.

Example 5-1 Configure your database by using the db2cmd command-line interface

```
db2 catalog tcpip node tpc02 remote myserver server 50000
db2 catalog database tpcdb as tpcdb02 at node tpc02
```

If you run **db2 list node directory**, you should see your newly configured database `tpcdb02`,

2. Next, you must configure your database in Cognos BI:

- a. In the Spectrum Control web-based GUI, click **Cognos** → **View predefined reports**.
- b. Click **Launch** → **IBM Cognos Administration**, click the **Configuration** tab, and click **Data Source Connections** (make sure that the path shows **Directory** → **Cognos**, as shown in Figure 5-26 on page 139).

- c. Click the New Data Source icon, as shown in Figure 5-26.

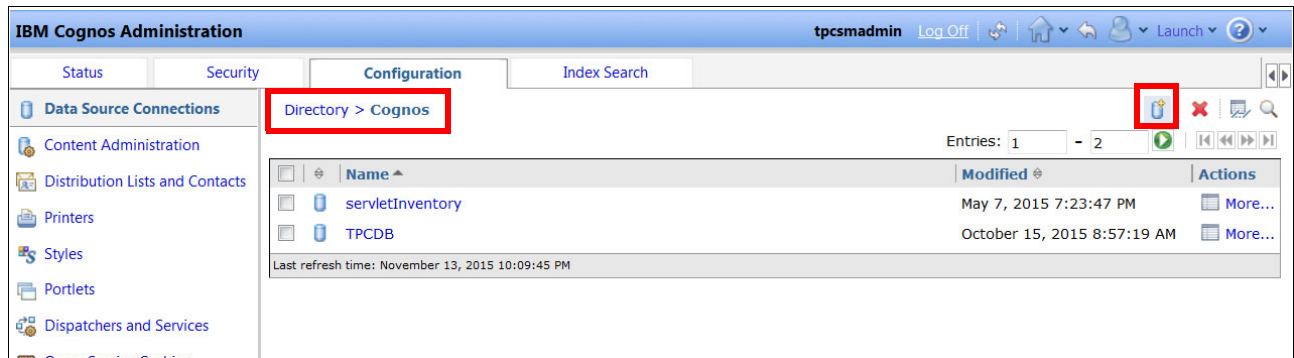


Figure 5-26 Add a data source in Cognos administration

- d. Enter the alias that you created in step 1 on page 138 in to the Name field, and then click **Next**.
- e. From the Type list, select **IBM DB2**, as shown Figure 5-27.
- f. Leave the **Configure JDBC connection** check box clear, as shown in Figure 5-27. Click **Next**.

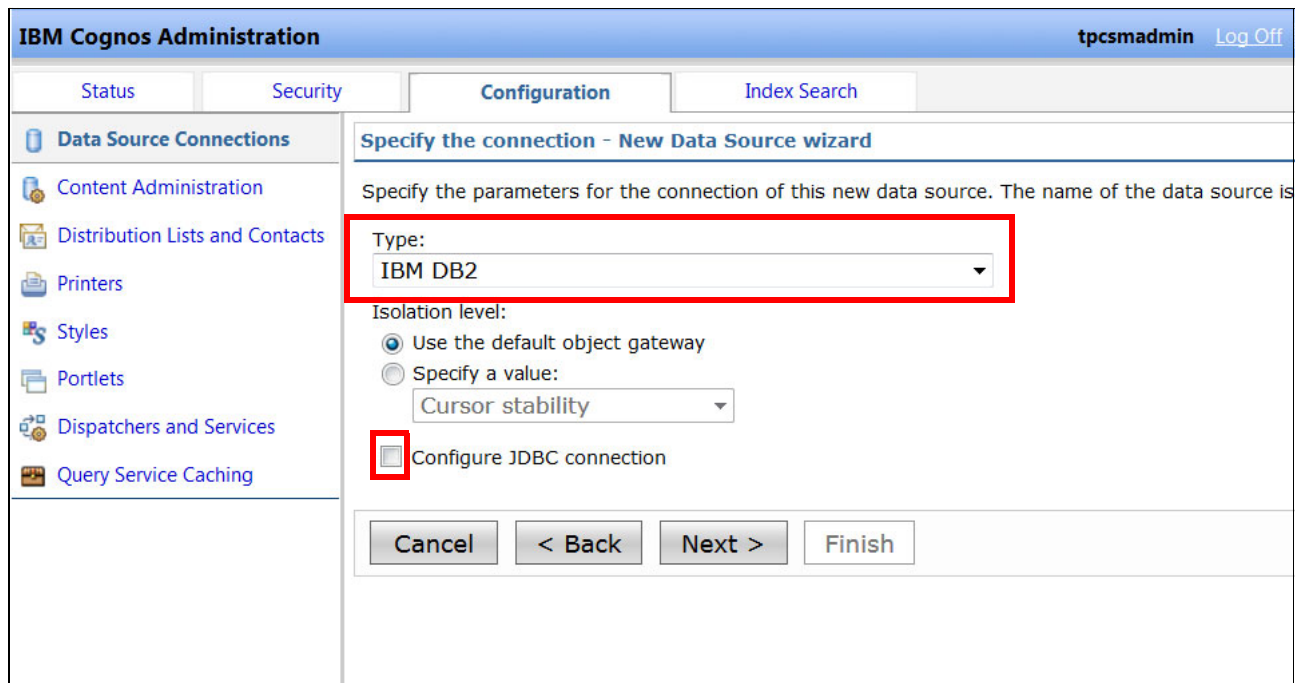


Figure 5-27 New Data Source wizard to specify a DB2 connection

- g. Enter the alias that you created in step 1 on page 138 in to the DB2 database name field.

- h. Select **Signons**, select the **Password** check box, and then enter a User ID and Password for the remote database, as shown in Figure 5-28.

The screenshot shows the IBM Cognos Administration interface. The top navigation bar includes 'Status', 'Security', 'Configuration' (selected), and 'Index Search'. The left sidebar lists various administration tasks, with 'Data Source Connections' selected. The main content area is divided into sections: 'Trusted context' (with an unchecked checkbox), 'Timeouts' (with fields for 'Connect time' and 'Reply time', both set to 0), 'Signon' (with radio buttons for 'No authentication', 'An external namespace: VMMProvider (Active)', and 'Signons' (selected)), and 'Testing' (with a 'Test the connection...' button). The 'Signons' section is highlighted with a red box, showing the 'Password' checkbox checked, the 'Create a signon that the Everyone group can use:' checkbox checked, and the 'User ID' field containing 'tpcadmin'. The 'Password' and 'Confirm password' fields are masked with dots. The 'Test the connection...' button in the 'Testing' section is also highlighted with a red box.

Figure 5-28 Provide credentials for a new data source connection

- i. Click **Test the connection**, and then click **Test**. Verify that the status of the connection is Succeeded, as shown in Figure 5-29 on page 141.

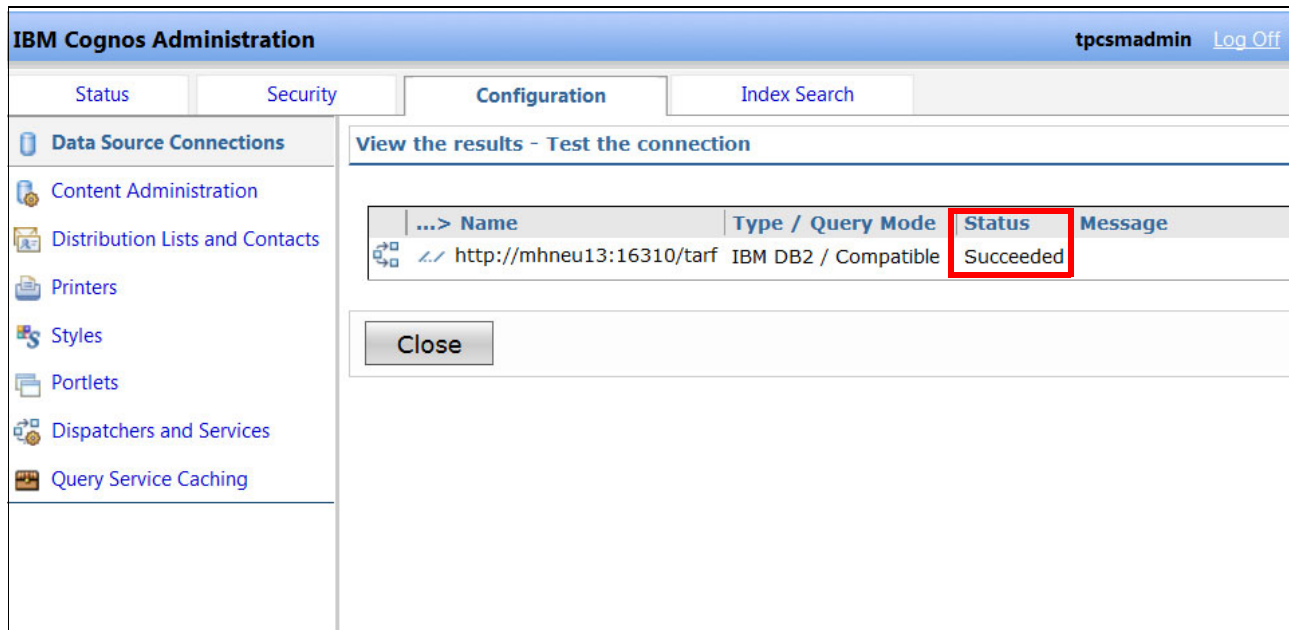


Figure 5-29 Result of the connection test

- j. Click **Close**. Click **Next**, and then click **Next** again. Click **Finish**.
- k. Check that TPCDB02, shown in Figure 5-30, now shows in the window that opens when you click **IBM Cognos Administration**, → **Configuration** tab → **Data Source Connections**. (Ensure that the path shows **Directory** → **Cognos**.)

Figure 5-30 shows the newly configured remote TPCDB02 (compare to Figure 5-26 on page 139).

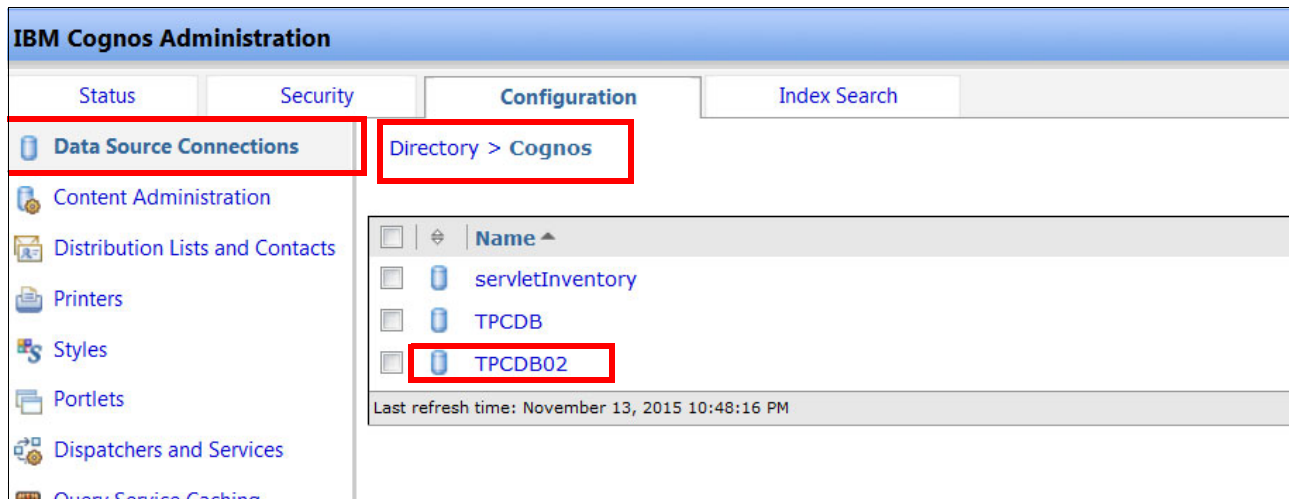


Figure 5-30 Data Source Connections

3. Create your Cognos BI Report from multiple Spectrum Control instances by using Report Studio.

For basic instructions about how to work with Cognos BI Reports within Spectrum Control see *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

Complete the following steps:

- a. Start Cognos Report Studio.
- b. Choose any package, for example, Spectrum Control Capacity and Relationship Package, and create a blank report.
- c. Open Query Explorer, as shown in Figure 5-31.
- d. Drag the SQL Item from the Toolbox to the right of your Query1.
- e. Select the SQL item and set the SQL Syntax Property to **Pass-Through**. Click the three dots next to the Data source field and select the local TPCDB as the data source, as shown in Figure 5-31.

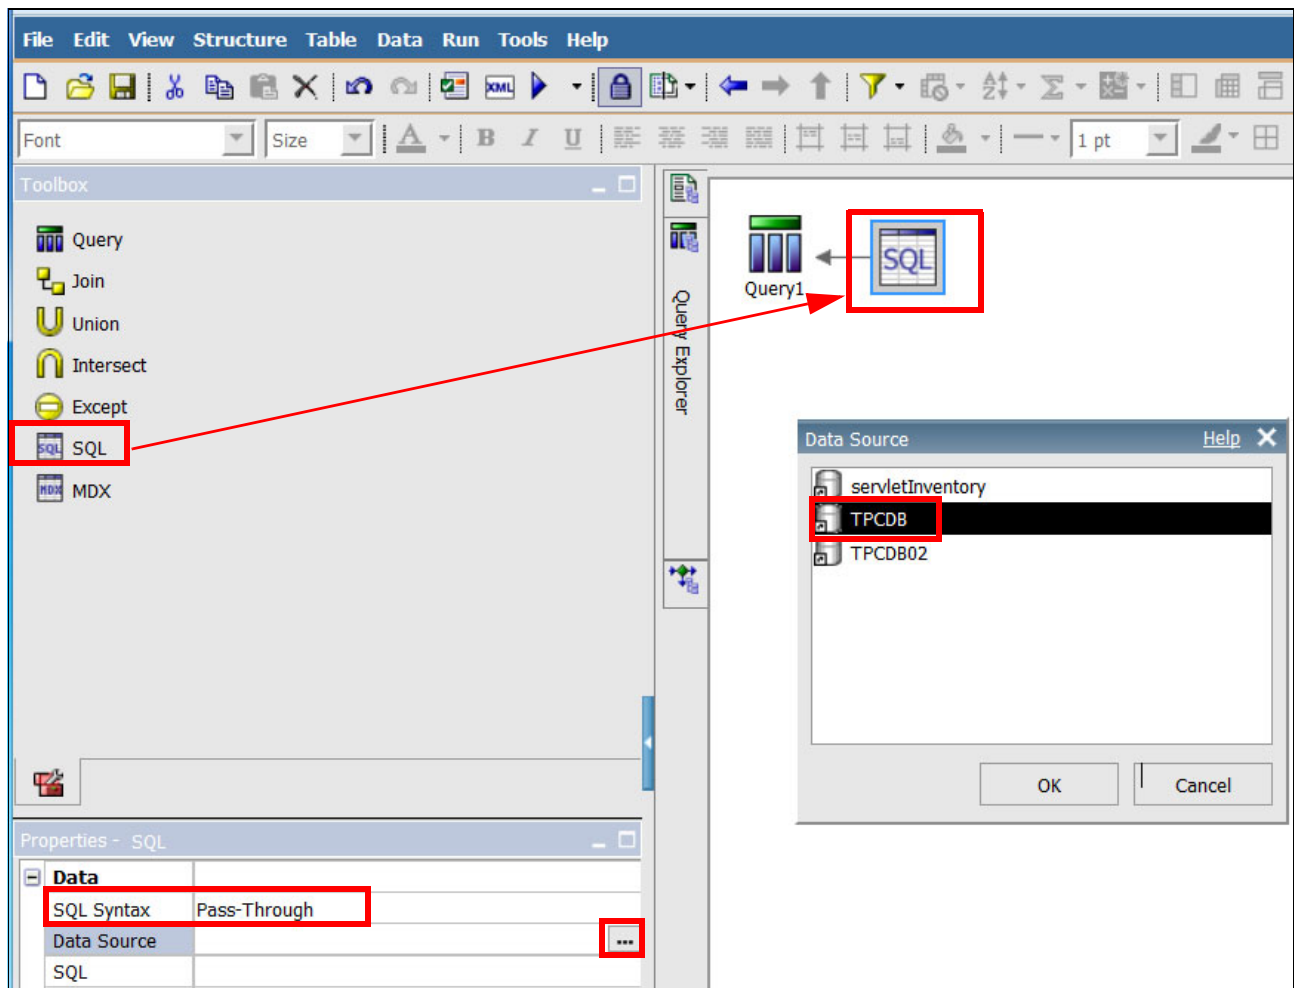


Figure 5-31 SQL properties

- f. Click **OK**.
- g. Double-click the SQL item. Insert the query statement that is shown in Example 5-2 on page 143, which gives you a list of storage virtualizer pools and their capacity, in which “SC1” refers to your local Spectrum Control server.

Example 5-2 SQL query

```
SELECT
  'SC1' as SCserver
, AL2.DISPLAY_NAME as StorageSubsystemName
, AL1.DISPLAY_NAME as StoragePoolName
, AL1.STORAGE_POOL_SPACE/ (1024*1024*1024) as Capacity
, AL1.REAL_SPACE/ (1024*1024*1024) as Allocated_Space
, AL1.STORAGE_POOL_AVAILABLE_SPACE/ (1024*1024*1024) as Available_Space
, AL1.STORAGE_POOL_CONSUMED_SPACE/ (1024*1024*1024) as Used_Space
FROM TPCREPORT.STORAGEPOOL AL1, TPCREPORT.STORAGESUBSYSTEM AL2
WHERE (AL2.SUBSYSTEM_ID=AL1.SUBSYSTEM_ID) AND (AL2.DISPLAY_NAME LIKE
'%SVC%')
with ur for read only
```

Note: Always specify with ur for read only in your SQL queries. Otherwise, your tables might get locked during the read operation, which might slow down the performance of your TPCDB.

For more information about this topic, see the following websites:

- ▶ <http://ibm.co/1QGS3uC>
- ▶ <http://ibm.co/1PKnYcM>

- Copy and paste Query1 so that you now have two queries. Query2 queries your second Spectrum Control server that you configured to be used remotely.
- Change the properties of the second SQL item to TPCDB02. Also, change “SC1” as SCserver to “SC2” as SCserver, as shown in Figure 5-32. (“SC2” is your remote Spectrum Control server.)

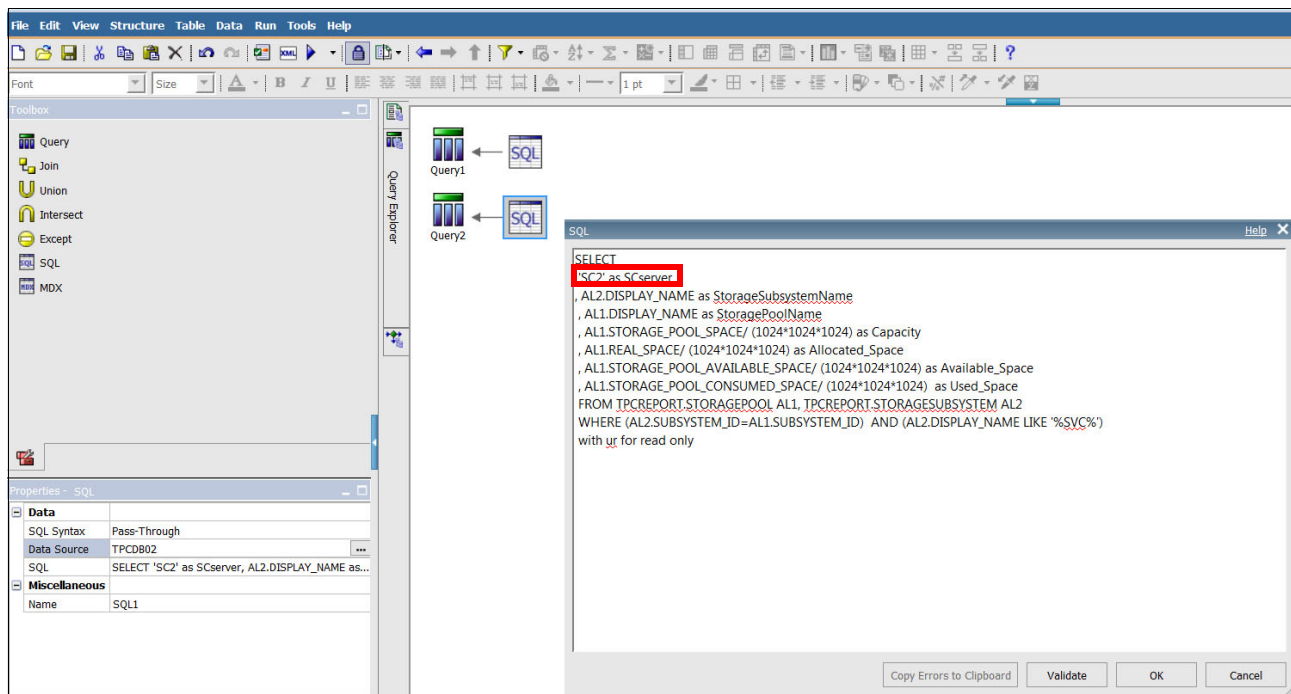


Figure 5-32 Rollup report query

- j. Drag a **Union** item from the Toolbox to the Query Explorer pane and add both of your Queries, as shown Figure 5-33.

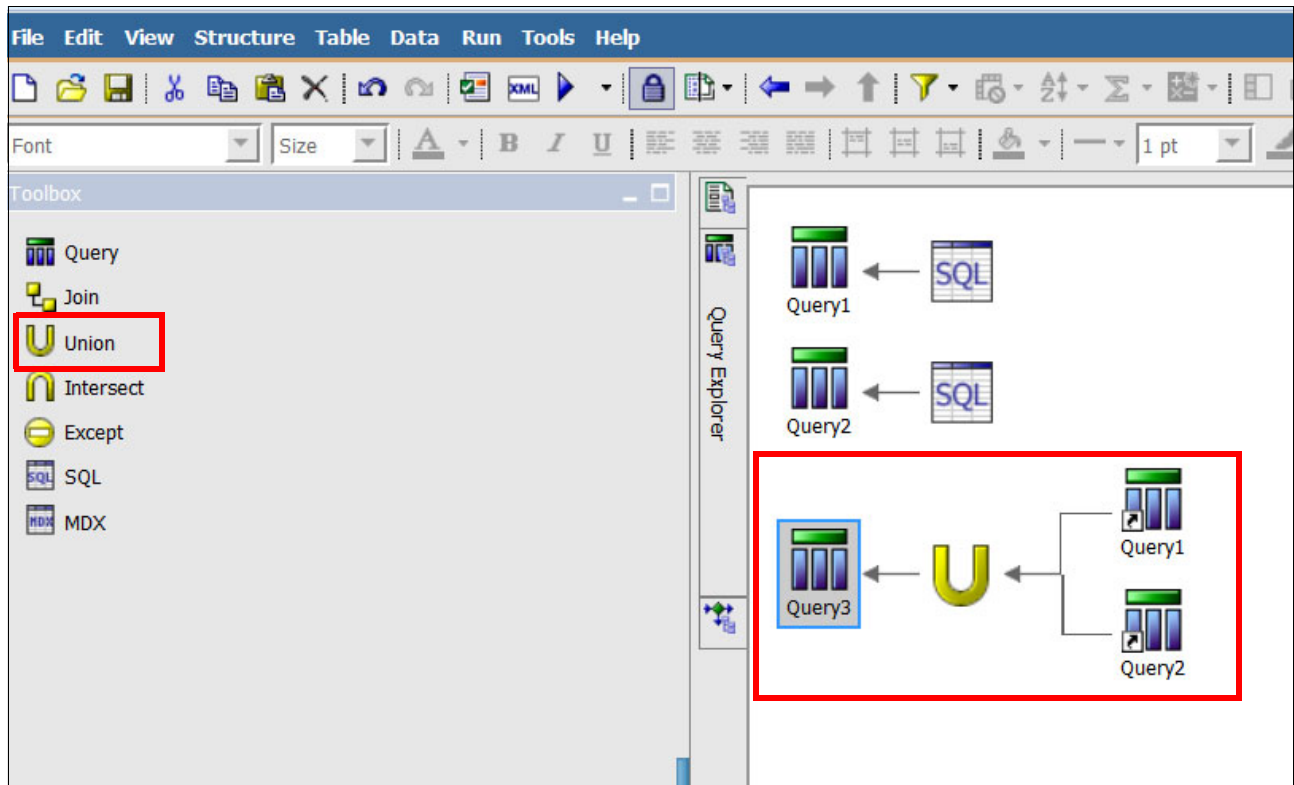


Figure 5-33 Union statement

- k. Open Query3 and drag all data items from the Source tab to the Data items tab of Query3, as shown in Figure 5-34.

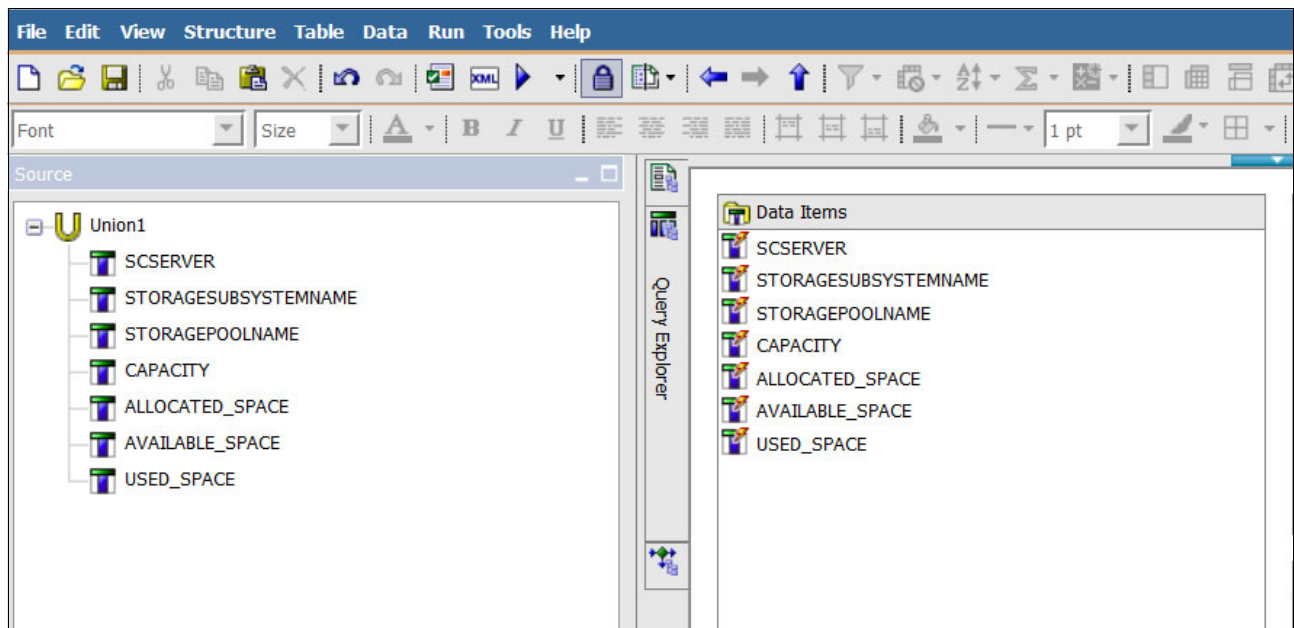


Figure 5-34 Add all data items to Query3

- l. Open **Page1** and drag Query3 to your page.
- m. Group the SCServer and the StorageSubsystemName columns.
- n. Run your Report.

You should get a report that shows your storage virtualizers in both Spectrum Control instances with their pools and the capacity metrics, as shown in Figure 5-35.

SCSERVER	STORAGESUBSYSTEMNAME	STORAGEPOOLNAME	CAPACITY	ALLOCATED_SPACE	AVAILABLE_SPACE	USED_SPACE
SC1	SVC Cust	Easy_Tier_pool_Site2	8,192	162	8,030	162
		EasyTier_pool	5,120	167	4,953	167
		EMC_Easy_Tier_Pool	400	34	365	35
		Flash 820 Flash only	6,144	182	5,962	182
		Flash Only 840	8,192	200	7,992	200
		MigrationPool_1024	16	0	16	0
		Quorum	19	0	19	0
		V7k_ctr11_SAS	5,119	410	4,709	410
		V7k_ctr22	8,191	410	7,781	410
SC2	SVC-2145-svc_27_28-IBM	TPC_Pool1	398	383	33	365
		TPC_Pool2	399	399	87	312
	SVC-2145-SVC_Cluster_80_81-IBM	Pool_1_DS8k15_Mirro	3,500	3,500	0	3,500
		Pool_2_DS8k03_Mirro	3,749	3,499	249	3,499
		Pool_3_Ds8k15_noMirro	3,499.5	3,501	0	3,499
	SVC-2145-SVC_Cluster-IBM	DS3400_pool	19	16	3	16
		EMC_pool	524	129	395	129
		MigrationPool_1024	18	0	18	0
		V7000_pool	1,024	346	684	340

Figure 5-35 Rollup report across two Spectrum Control servers

For more information about the Spectrum Control exposed views, see the following website:

<http://www.ibm.com/support/docview.wss?uid=swg27023813>

Note: If you use Local OS authentication for Cognos, as described in *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204, you must add the users to the db2users group to create reports with native SQL.

5.2 Organizing capacity data around applications and business units

When you justify costs for existing or new storage systems, have answers to the following questions:

- ▶ Which applications use how much and what kind of capacity?
- ▶ Who owns which application?

Spectrum Control can help you get answers to these questions through the *Grouping* function that is part of Spectrum Control. With the Grouping function, you can build a business-oriented hierarchy in Spectrum Control. This hierarchy can be as simple or as complex as you require, as shown in Figure 5-36.

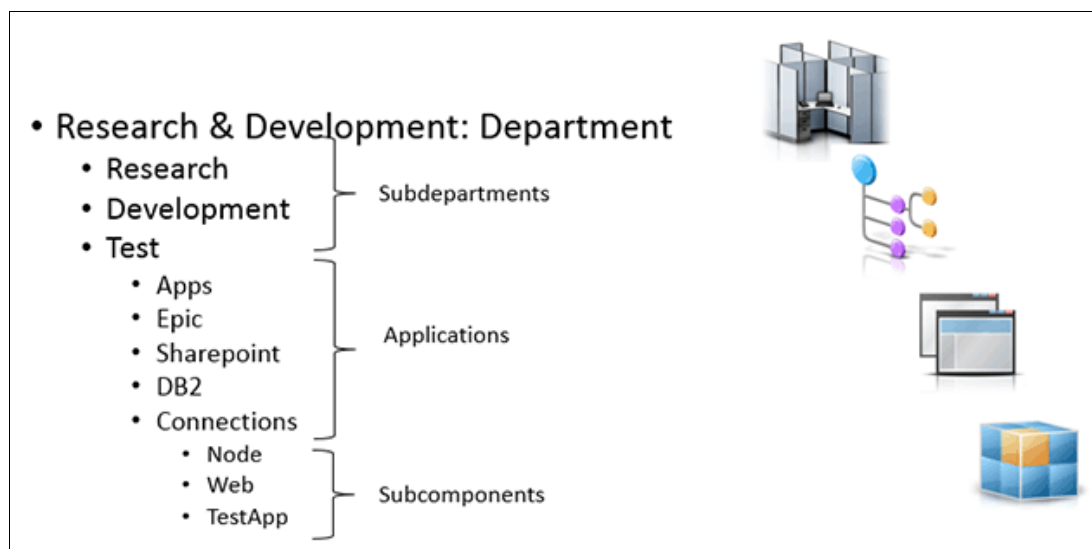


Figure 5-36 An example of a business-oriented hierarchy in Spectrum Control

Note: A subcomponent is an application that is a child of another application.

In the example shown in Figure 5-36, resources such as data stores, exports, file systems, file sets, servers, VMs, volumes, and volume groups were added to the subcomponents Node, Testapp, and Web, and to applications Apps, Epic, Sharepoint, and DB2. The applications Node, TestApp, and Web were then added to the application called Connections. The applications Connections, DB2, Sharepoint, Epic, and Apps were then added to the Department Test. The departments Test, Development, and Research were added to the Main-Department Research and Development.

There are three ways to add resources to applications:

1. Use filters to assign resources to the applications that are used in step 6 on page 150. For more information, see the Spectrum Control IBM Knowledge Center, found at:
<http://ibm.co/117WeDc>
2. Add resources manually to the applications that are used in step 14 on page 154. For more information, see the Spectrum Control IBM Knowledge Center, found at:
<http://ibm.co/1SIAFnD>
3. Add resources by using the command-line interface (CLI). You can do so in two ways:
 - a. Add resources individually.
 - b. Add resources by using bulk assignment.

For more information, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1MpzurM>

Note: Not all resources can be added by using all methods. For example, virtual machines (VMs) that are not listed as a server can be added only by using the CLI.

When you use the CLI, you can add all resource types and automate your bulk assignment of eligible storage resources to applications.

There are more functions that are available than are shown in this section. For more information, see the Spectrum Control IBM Knowledge Center.

The Application and Department function can be accessed on the Groups tab in the navigation bar.

Figure 5-37 shows the Application and Department function.

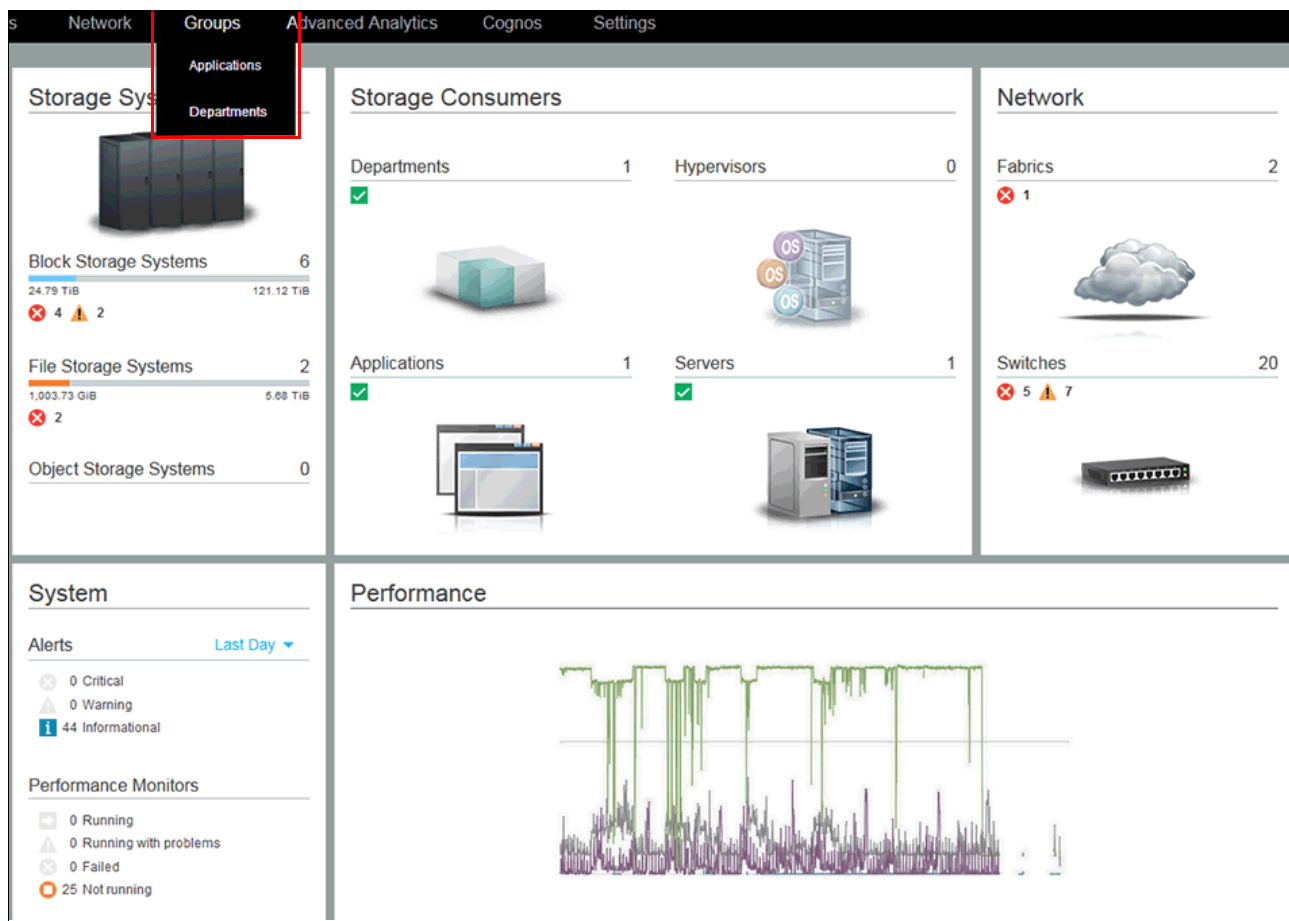


Figure 5-37 Application and Department function

In the example that is shown in Figure 5-37 on page 147, you can start creating your first application Connection by using the subcomponents Node and Web. Complete the following steps:

1. From the Groups tab, click **Application**.
2. Click **Create Application**, enter Connection as the application name, and click **Create**.

Type and Subtype are not mandatory. To keep this example simple, we do not specify Type and Subtype.

Figure 5-38 shows how to create an application.

The screenshot shows the 'Create Application' wizard interface. At the top, there is a '+ Create Application' button and an 'Actions' dropdown menu. Below this is a table with columns: Name, Status, Block Capacity (GiB), and File Capacity. The main area is titled 'Create Application' and contains a form with the following fields: 'Name' (set to 'Connection'), 'Description' (empty), 'Type' (dropdown menu with 'Enter a type' text), and 'Subtype' (dropdown menu with 'Enter a subtype' text). At the bottom right, there are 'Create' and 'Cancel' buttons. A red box highlights the 'Create Application' button at the top left and the 'Create' button at the bottom right. Another red box highlights the 'Name' field and its value 'Connection'.

Figure 5-38 Create Application wizard

3. Now, you have several options to either assign resources to the Connection application or specify subcomponents. In this example, we choose to create the subcomponent *Node*. Click **Create subcomponents that each have their own set of storage resources** and then click **Close**, as shown in Figure 5-39 on page 149, which shows the methods that are available after creating the Connection application.

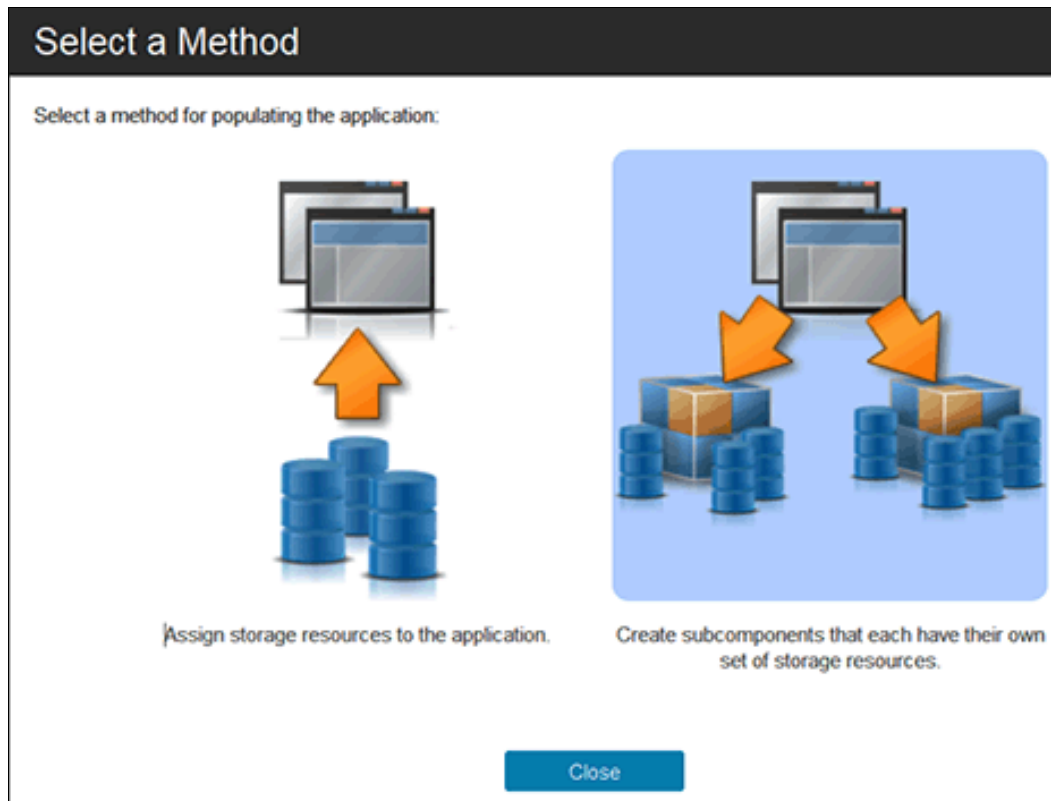


Figure 5-39 Methods to either add storage resources or add subcomponents to the application

4. In the Create Subcomponent window, enter Node in to the Name field, and click **Create**, as shown in Figure 5-40. This action creates the Node subcomponent.

Figure 5-40 Subcomponent window

When you look at the Connections application, you see the Node subcomponent, but all the Capacity values are still at 0 GiB because no storage resources are added yet.

Figure 5-41 shows the Connection application with the Node subcomponent.

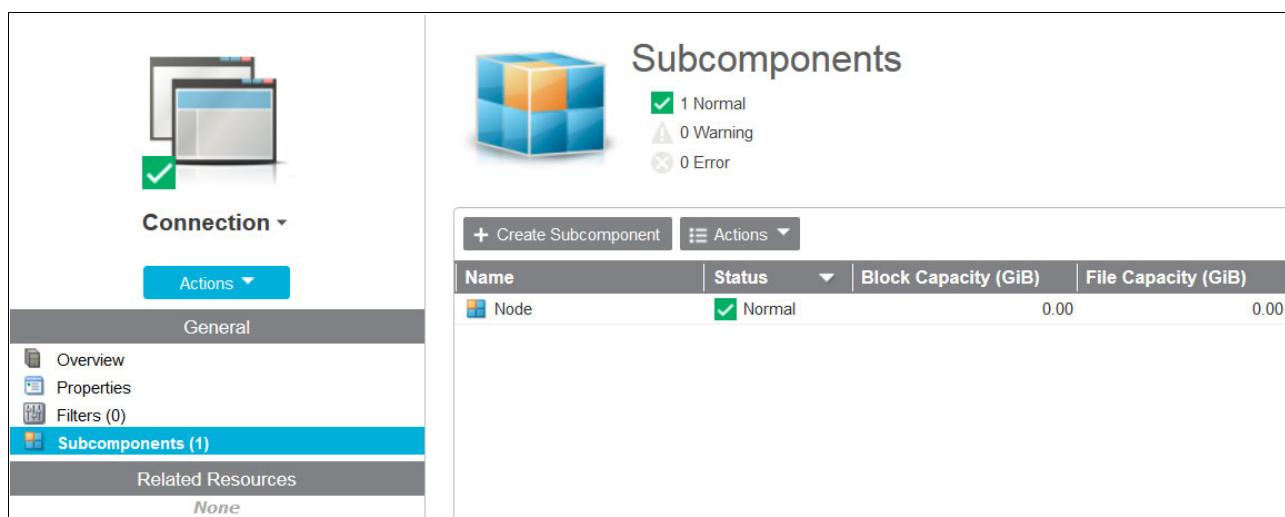


Figure 5-41 All capacity values are at 0 GiB because no storage resources are added yet

5. Double-click **Node** to add storage resources.
6. Click **Filters**, and then click **Create Filter**, as shown in Figure 5-42.

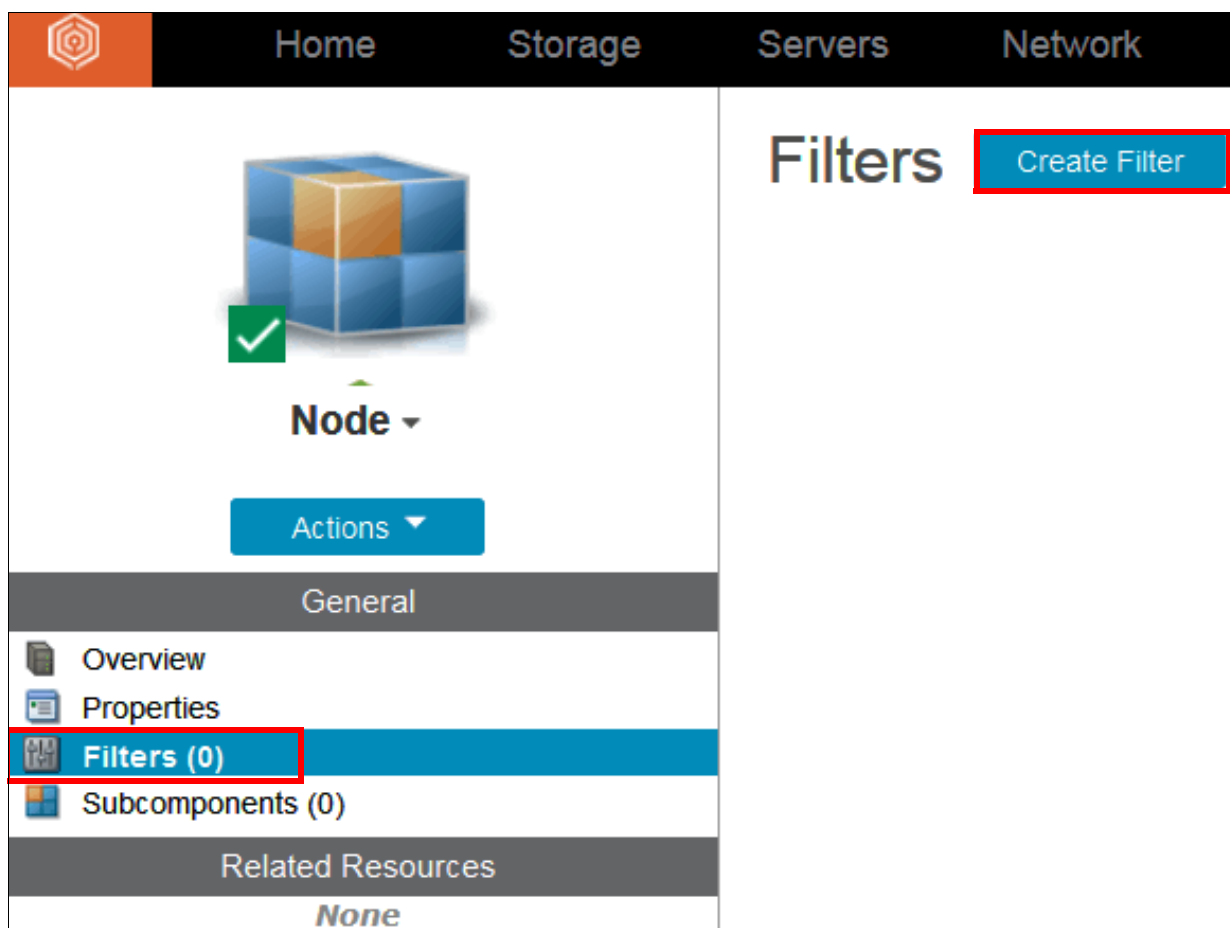



Figure 5-42 Example of creating a filter


7. Click **A selection of volumes that match a pattern**, as shown in Figure 5-43.
8. Define the Name Patterns for the Volumes. You can use the wildcards “*” for multiple characters or “?” for one character. Because of naming conventions in our example, company, all volumes that belong the Node subcomponent start with TPC. Therefore, enter TPC*, as shown in Figure 5-43, which shows the method of adding storage resources to an application or subcomponent by using the filter method.

Create Filter


Select a Resource Type




All storage that belongs to a set of servers or hypervisors



A selection of volumes that match a pattern



A selection of shares that match a pattern.



A selection of filesets that match a pattern.

Define Name Patterns

Volumes:


TPC*

☐ Belonging to:

☐ Server/Hypervisor:

☐ Storage System:

☐ Pool:



Preview

Save

Cancel

Figure 5-43 Add storage resources to an application or subcomponent by using the filter method

- Click **Preview** to see the volumes that match this pattern. Figure 5-44 shows all volumes that match the naming pattern.

Filter Matches

Actions

Filter...

Name	Server or Hypervisor	Storage System	Pool
TPC_vDisk2_DS8K	9.155.113.62	SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_vDisk_DS8K	9.155.113.62	SVC-2145-svc_27_28-IBM	TPC Pool2
TPC_VOL0		SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_VOL1		SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_VOL2		SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_VOL3		SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_Vol_ZE1	9.155.90.222 Demo	SVC-2145-svc_27_28-IBM	TPC Pool2
TPCMZAN-BAK	x3650_m3_17_VMWare	SVC-2145-SVC Cluster 80 8...	Pool 3 Ds8k15 noMirro
TPCMZAN-Prod	x3650_m3_17_VMWare	SVC-2145-SVC Cluster 80 8...	Pool 1 DS8k15 Mirro
tpcr_5A00		DS8000-2107-75ZA571-IBM	ITSO EasyTier
tpcr_5A00		DS8000-2107-75TV181-IBM	fb_r5_0
tpcr_5A01		DS8000-2107-75ZA571-IBM	ITSO EasyTier
tpcr_5A01		DS8000-2107-75TV181-IBM	fb_r5_0
tpcr_5A02		DS8000-2107-75ZA571-IBM	ITSO EasyTier
tpcr_5A02		DS8000-2107-75TV181-IBM	fb_r5_0
tpcr_5A03		DS8000-2107-75ZA571-IBM	ITSO EasyTier
tpcr_5A03		DS8000-2107-75TV181-IBM	fb_r5_0

Showing 137 items | Selected 0 items

Refreshed a few moments ago

Save

Close


Figure 5-44 Shows all volumes that match the naming pattern.

Tip: If we had specified “TPC_” in Figure 5-43 on page 151, we would not see the “tpcr” volumes in Figure 5-44.


- Because only volumes on SVC-2145-svc_27-28-IBM belong to the Node subcomponent, click **Close** to redefine the filter.
- Select the **Belonging to** check box, select **Storage System**, and enter SVC-2145-svc_ in to the Storage System field to get only volumes with names that start with TPC that are on SVC-2145-svc_27-28-IBM, as shown in Figure 5-45 on page 153, which shows an example for a filter on volumes.

Create Filter


Select a Resource Type




All storage that belongs to a set of servers or hypervisors



A selection of volumes that match a pattern



A selection of shares that match a pattern.



A selection of filesets that match a pattern.

Define Name Patterns

Volumes: TPC*

☒ Belonging to:

☐ Server/Hypervisor:
 ☐ Storage System: SVC-2145-svc_27*
 ☐ Pool:

Preview

Save

Cancel

Figure 5-45 Example of how a filter can be specified

Tip: Filters can be restricted, for example, all volumes that are mapped to a given server or contained in a given pool.

12. Click **Preview** again and check whether the filter is correct now.

Figure 5-46 shows the volume selection after redefining the filter.

Filter Matches

Actions

Filter...

Name	Server or Hypervisor	Storage System	Pool
TPC_vDisk2_DS8K	9.155.113.62	SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_vDisk_DS8K	9.155.113.62	SVC-2145-svc_27_28-IBM	TPC Pool2
TPC_VOL0		SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_VOL1		SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_VOL2		SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_VOL3		SVC-2145-svc_27_28-IBM	TPC Pool1
TPC_Vol_ZE1	9.155.90.222 Demo	SVC-2145-svc_27_28-IBM	TPC Pool2

Showing 7 items | Selected 0 items

Refreshed a few moments ago

Save

Close

Figure 5-46 Show only volumes with names that start with TPC and are on SVC-2145-svc_27_28-IBM

- Click **Save**. When you provision a volume with a name that starts with “TPC on SVC-2145-svc_27_28”, it is added to the Node subcomponent automatically.

Note: The filter function allows resources of a given type matching a certain naming pattern to be associated. Filters are dynamic and are continuously evaluated. Using the filter function keeps the application self-maintaining.

- If the naming conventions were not implemented everywhere, you must add a server explicitly through the web-based GUI. Agentless servers can also be added to an application. Right-click the agentless server **9.155.90.222 Demo** and select **Add to application**, as shown in Figure 5-47 on page 155, which shows how to add storage resources explicitly through the web-based GUI.

Servers Alerts Tasks Discovered Servers

+ Add Server

Actions

View Properties

View Details

View Data Path

Provision Storage

Analyze Tiering

Add to Application

Modify Ports

Remove

Name	Status	Probe Status	Agent State	OS Type
WIN-2I0P2DE19BJ	Normal	Successful	Up	Windows
WIN-D2VMKKUITNA	Normal	Successful	Up	Windows
ws-vcenter.mainz.de.ibm.com	Normal	Successful	Up	Windows
x3550-m2-03.mainz.de.ibm.com	Normal	Failed	Up	Windows
9.155.113.248		N/A		Linux
9.155.116.119		N/A		Windows
9.155.116.142		N/A		Other
9.155.116.143		N/A		Other
9.155.116.153		N/A		Other
9.155.50.201		N/A		AIX
9.155.87.155		N/A		Windows
9.155.90.222_Demo		N/A		Windows
p8-e870-01-lpar1	Agentless	N/A		Other
Test_5005076801105D7E	Agentless	N/A		Other
x3650_m3_17_VMWare	Agentless	N/A		Other
9.155.90.216	Deploying	Never Probed	Deploying	Unknown
9.155.90.223	Deploying	Never Probed	Deploying	Unknown

Showing 28 items | Selected 1 item

Figure 5-47 Add a data store to an application

15. A list of existing applications and subcomponents opens. In our example, we must create the web application. To do so, click **Create Application**, enter Web in to the Name field, and click **Create**.

Figure 5-48 shows how to create an application while explicitly adding storage resources.

The screenshot shows a software interface with a dark header bar containing the text "Add to Application". Below the header is a toolbar with a "+ Create Application" button and an "Actions" dropdown menu. A table lists three items: "Accounting" with hierarchy "/Accounting", "all_types" with hierarchy "/all_types", and "Book Sales DB" with hierarchy "/Book Sales DB". Each item has a checkbox to its left. Below the table is a "Create Application" sub-dialog. This sub-dialog has a dark header with the text "Create Application". On the left is an icon of a laptop. To the right of the icon are four labels: "Name:", "Description:", "Type:", and "Subtype:". The "Name:" field contains the text "Web". The "Description:" field is empty. The "Type:" and "Subtype:" fields are dropdown menus with the placeholder text "Enter a type" and "Enter a subtype" respectively. At the bottom of the sub-dialog are two buttons: "Create" and "Cancel".

	Name	Hierarchy	Type
<input type="checkbox"/>	Accounting	/Accounting	
<input type="checkbox"/>	all_types	/all_types	
<input type="checkbox"/>	Book Sales DB	/Book Sales DB	

Create Application

Name: Web

Description:

Type: Enter a type

Subtype: Enter a subtype

Create Cancel

Figure 5-48 Create the Web application

16. The Web application must be added to the Connection application. To do so, click **Groups**, click **Applications**, right-click the Web application, select **Add to Application**, select the Connections application check box, and click **Save**.
17. When you view the Subcomponents application under the Connections application, you see the two subcomponents with their capacity values, as shown in Figure 5-49 on page 157.

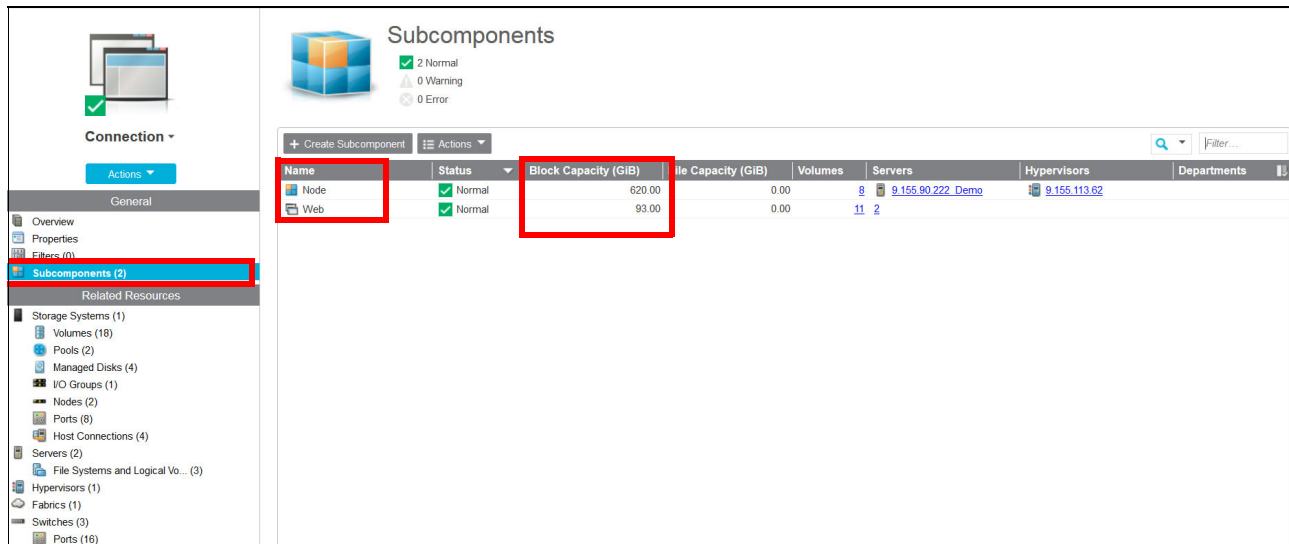


Figure 5-49 Connection application showing all the subcomponents and their resources

Related resources are explained in 5.3, “Getting an overview of the resources that are associated with your applications” on page 164.

18. Looking at the application’s level view in Figure 5-50, you see the aggregated capacity values of the two subcomponents of the Connection application.




Figure 5-50 Aggregated capacity values of the Connection application

To add a VM to an application, you must use the **tpctool** command.

References: For more information about how to use the **tpctool** command, see the Spectrum Control IBM Knowledge Center at the following websites:

- For information about command modes, go to:
<http://ibm.co/10C7TW0>
- For information about tpctool commands, go to:
<http://ibm.co/1NuuQdF>

19. Start a command window (CMD) and from the <install>/cli directory, run the following command to get the tpctool prompt:

```
tpctool -user tpcadmin -pwd password
```

20. Because you must use **tpctool** to add the VM to the application, you also must create the TestApp application by using tpctool by running the following command:

```
mkappgroup -name TestApp
```

Figure 5-51 shows how to create the TestApp application.

```
C:\Windows\System32\cmd.exe
tpctool> mkappgroup -name TestApp
Name      ID              Status
TestApp   /Application/9352510 SUCCESS
```

Figure 5-51 The mkappgroup command and the associated ID

21. To get the ID of the VM, run the following command:

```
lsvm -l
```

Figure 5-52 shows the **lsvm -l** command.

```
tpctool> lsvm -l
```

ID	Name	Hypervisor Name	OS Type	Host Name	IP Address
ws-vcenter+9.155.113.62	ws-vcenter	9.155.113.62	"Microsoft Windows Server 2008 R2 (64-bit)"	ws-vcenter.mainz.de.ibm.com	9.155.115.170
ws-tpc+9.155.113.62	ws-tpc	9.155.113.62	"Microsoft Windows Server 2008 R2 (64-bit)"	vmworlagewin150.mainz.de.ibm.com	9.155.115.171
p6-570-lab-2v3.mainz.de.ibm.com	p6-570-lab-2v3	-	AIX	p6-570-lab-2v3.mainz.de.ibm.com	9.155.157.28
p7-770-04v1.mainz.de.ibm.com	p7-770-04v1	-	AIX	p7-770-04v1.mainz.de.ibm.com	9.155.116.215

Figure 5-52 VMs and their parameters

22. To confirm the ID of your TestApp application, run the following command:

```
lsappgroup -l
```

The **lsappgroup** command lists all applications and their parameters, as shown in Figure 5-53.

```
tpctool> lsappgroup -l
```

Name	ID	Parents	Type	Subtype	Status	Description	UDP1	UDP2	UDP3	Paths
all_types	/Application/1879616	-	-	-	NORMAL	-	-	-	-	/all_types
Accounting	/Application/4578583	-	-	-	NORMAL	-	-	-	-	/Accounting
"Book Sales DB"	/Application/178148	-	-	-	NORMAL	-	-	-	-	"/Book Sales DB"
Connection	/Application/9311950	-	-	-	NORMAL	-	-	-	-	/Connection
Customer1	/Application/3366090	-	Abrechnung	-	NORMAL	-	-	-	-	/Customer1
datastoretest	/Application/904489	-	-	-	NORMAL	-	-	-	-	/datastoretest
eco123	/Application/8321086	-	-	-	WARNING	-	-	-	-	/eco123
ESX	/Application/151002	-	ESX	UM	NORMAL	-	-	-	-	/ESX
Flashtest	/Application/348784	-	-	-	NORMAL	-	-	-	-	/Flashtest
LAB_TSM	/Application/162622	TSM	LAB	TSM	NORMAL	-	-	-	-	/TSM/LAB_TSM
mhi	/Application/9311410	-	-	-	NORMAL	-	-	-	-	/mhi
NIU	/Application/9291535	-	-	-	NORMAL	-	-	-	-	/NIU
Node	/Application/9311951	Connection	-	-	NORMAL	-	-	-	-	/Connection/Node
"Online Transactions"	/Application/178147	"Book Sales DB"	-	-	NORMAL	-	-	-	-	"/Book Sales DB/Online Transactions"
"Peters App"	/Application/861652	-	"Tiered Volumes T1"	-	NORMAL	"Test app"	-	-	-	"/Peters App"
residencytest	/Application/9309739	-	-	-	NORMAL	-	-	-	-	/residencytest
"Retail Transactions"	/Application/178146	"Book Sales DB"	-	-	NORMAL	-	-	-	-	"/Book Sales DB/Retail Transactions"
TestApp	/Application/9352510	-	-	-	NORMAL	-	-	-	-	/TestApp

Figure 5-53 The lsappgroup command shows the application IDs

23. To add the ws-vcenter VM to the TestApp application, run the following command:

```
modifyappgroup -add -id /Application/9352510 -type vm -key ws-vcenter+9.155.113.62
```

Figure 5-54 shows how to add a VM to an application.

```
tpctool> modifyappgroup -add -id /Application/9352510 -type vm -key ws-vcenter+9.155.113.62
```

Group Name	Member Key	Member Type	Status
TestApp	ws-vcenter+9.155.113.62	Virtual Machine	SUCCESS

Figure 5-54 Add the ws-vcenter VM to the TestApp application

24. To check the application members of the TestApp application, run the following command:

```
lsappgroupmembers -id /Application/9352510
```

Figure 5-55 shows the members of the TestApp application.

```
tpctool> lsappgroupmembers -id /Application/9352510
Member Key      Member Name      Member Type
=====
ws-vcenter+9.155.113.62 ws-vcenter.mainz.de.ibm.com "Virtual Machine"
```

Figure 5-55 The ws-vcenter VM as a member of the TestApp application

25. Now, you must add the TestApp application as a subcomponent to the Connection application. You also want to add a few volumes that do not follow any naming conventions to the Connection application. In this scenario, a colleague informed you that he does not need the WIN-2I0P2DE19BJ server in his Web application anymore, and that he marked this server as “notused”. This task can be done in one step by running the **modifyappgroupvia** command.

Reference: For more information about the usage of custom tags, see *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

26. To get the volume IDs, you must determine the device GUID on which the volumes are by running the following command:

```
lsdev -l
```

Figure 5-56 shows the output of the **lsdev -l** command.

```
tpctool> lsdev -l
GUID      Name      Label
=====
2107.75ZA571+0 DS8000-2107-75ZA571-IBM DS8000-21
2107.75TU181+0 DS8000-2107-75TU181-IBM DS8000-21
0000020061014FA2+0 SUC-2145-SUC_Cluster-IBM SUC-2145-
00000200A0C1454C+0 Storwize U7000-2076-U7000 Unified 2-IBM Storwize
ONTAP:0135081557+0 N3400-01a Nseries_
```

Figure 5-56 Devices and their parameters

27. Now, you can create a list with the volumes of this DS8000 storage system. To pipe the output to the file location, open another command line, go to the <install>/cli directory, and run the following command:

```
tpctool -user tpcadmin -pwd password lsvol -l -dev 2107.75ZA571+0 >>
C:\temp\vol.txt
```

28. Open the vol.txt file and find the labels for the volumes that you want to add to the Connection application.

Figure 5-57 shows an example for the **lsvol -l** output.

ID	Label	Size	Format	FlashCopy	Real Us
IBM.2107-75ZA571-0111+1+2107.75ZA571+0	"ITSO_ESXi_vol (ID:0111)"	50.00	FB	-	50.
IBM.2107-75ZA571-2d40+1+2107.75ZA571+0	"PACI_2D40 (ID:2d40)"	23.79	CKD	-	23.
IBM.2107-75ZA571-2d3f+1+2107.75ZA571+0	"PACI_2D3F (ID:2d3f)"	7.93	CKD	-	7.
IBM.2107-75ZA571-2d3e+1+2107.75ZA571+0	"PACI_2D3E (ID:2d3e)"	7.93	CKD	-	7.
IBM.2107-75ZA571-2d3d+1+2107.75ZA571+0	"PACI_2D3D (ID:2d3d)"	7.93	CKD	-	7.
IBM.2107-75ZA571-2d3c+1+2107.75ZA571+0	"PACI_2D3C (ID:2d3c)"	7.93	CKD	-	7.
IBM.2107-75ZA571-2d3b+1+2107.75ZA571+0	"PACI_2D3B (ID:2d3b)"	7.93	CKD	-	7.
IBM.2107-75ZA571-2d3a+1+2107.75ZA571+0	"PACI_2D3A (ID:2d3a)"	7.93	CKD	-	7.
IBM.2107-75ZA571-2d39+1+2107.75ZA571+0	"PACI_2D39 (ID:2d39)"	7.93	CKD	-	7.

Figure 5-57 List of volumes for the specified device

29. Using the `inputfile.txt` input file and the `modifyappgroupviafile` command, you can add and remove multiple devices to or from multiple applications in one step. To do so, run the following command:

```
modifyappgroupviafile -inputfile c:\temp\inputfile.txt
```

Figure 5-58 shows the `modifyappgroupviafile` command.

```
tpctool>
tpctool>
tpctool> modifyappgroupviafile -inputfile c:\temp\inputfile.txt
tpctool>
```

Figure 5-58 Example of the `modifyappgroupviafile` command

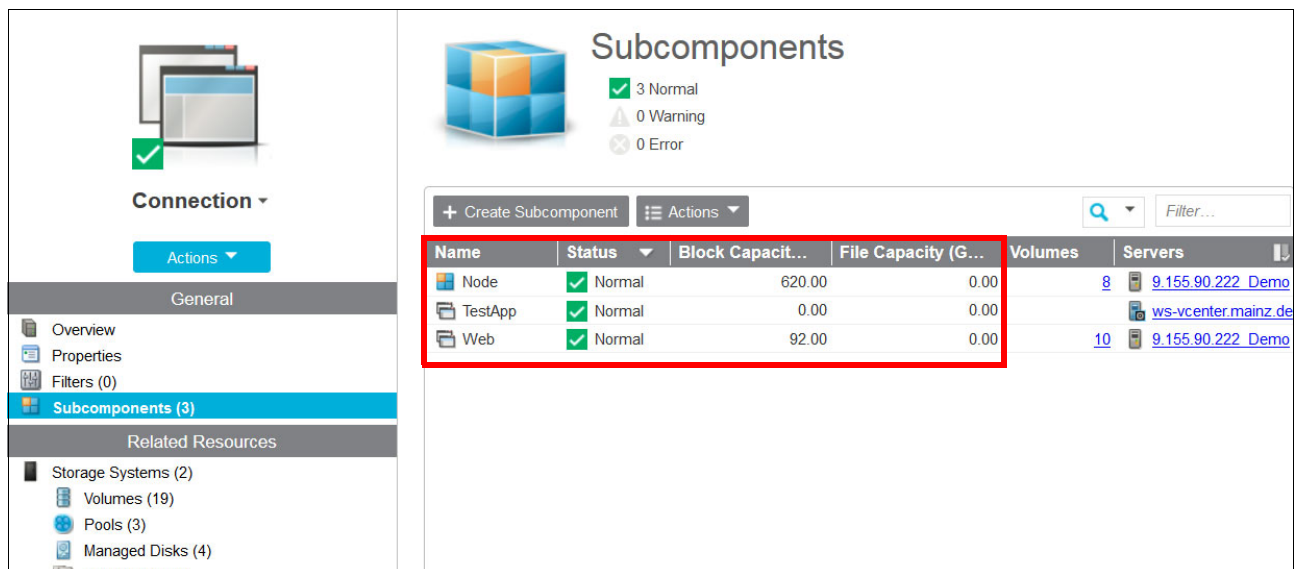
Figure 5-59 shows the content of the `inputfile.txt` file.

```
-id /Application/9311950
  -add -type appgroup -memberid /Application/9352510
  -add -type volume -devicename DS8000-2107-752A571-IBM -name KD2F1C ,"CKD_2F8F (ID:2f8f)"
-id /Application/9315330
  -remove -type server| -tagkey udp3 -tagvalue notused
```

Figure 5-59 Example of an input file for the `modifyappgroupviafile` command

Now, the Connection application lists three subcomponents and their capacity values.

Figure 5-60 shows the subcomponents of the Connection application.



Name	Status	Block Capacit...	File Capacity (G...	Volumes	Servers
Node	✓ Normal	620.00	0.00	8	9.155.90.222 Demo
TestApp	✓ Normal	0.00	0.00		ws-vcen...mainz.de
Web	✓ Normal	92.00	0.00	10	9.155.90.222 Demo

Figure 5-60 Subcomponents and the related resources for the Connection application

The capacity values of Connection application changed. Figure 5-61 on page 161 shows the capacity values for the Connection application.

+ Create Application		Actions		Conn		Reset
Name	Status	Block Capacity (GiB)	File Capacity (GiB)	Volumes	Servers	Hyp
Connection	✓ Normal	697.29	0.00	19 2		9.11

Figure 5-61 The parameters of the Connection

30. The DB2, Sharepoint, Epic, and Apps applications can be created in a similar fashion. To create the Test department, select the Connection and Apps applications and click **Add To → Department**, as shown in Figure 5-62.

Applications

✓ 19 Normal
⚠ 1 Warning
✖ 0 Error

+ Create Application

≡ Actions ▼

Name	Status	Block Capacity (GiB)	File Capacity (GiB)	Vol
eco123	⚠ Warning	426.00	0.00	
Accounting	✓ Normal	0.00	0.00	
Apps	✓ Normal	450.00	0.00	
Book Sales DB	✓ Normal	92.00	0.00	
Connection	✓ Normal	697.29	0.00	
Customer1	✓ Normal	5.00	0.00	
datastoretest	✓ Normal	0.00	0.00	
DB2	✓ Normal	1,138.00	0.00	

Add to...

Application

Remove

Department

Figure 5-62 Add multiple applications to a department

31. Click **Create Department**, enter Test into the Name field, and click **Create**, as shown in Figure 5-63.

Figure 5-63 Create a department

Similarly, you can add the DB2 and Epic applications to the Research department and the Sharepoint application to the Development department.

32. Now, the Development, Research, and Test departments must be added to the main Research & Development department. To do so, click **Groups** → **Departments**, select the Development, Research and Test departments, click **Add to** → **Department** → **Create Department**, enter Name Research & Development, and click **Create**.

The view of the department's condition and capacity consumptions are rolled up to the department level. This table can be exported to CSV, PDF, or HTML formats. Figure 5-64 shows the accumulated capacity data of the main department.

Name	Status	Block Capacity (GiB)	File Capacity (GiB)	Applications	Subdepartments	Type	Subtype
Research & Development	Normal	8,512.29	8,997.27				3

Figure 5-64 Accumulated capacity values of the Research & Development department

33. Click the Research & Development department to get to the condition and capacity consumptions of each subdepartment. Figure 5-65 shows the subdepartments of the Research & Development department.

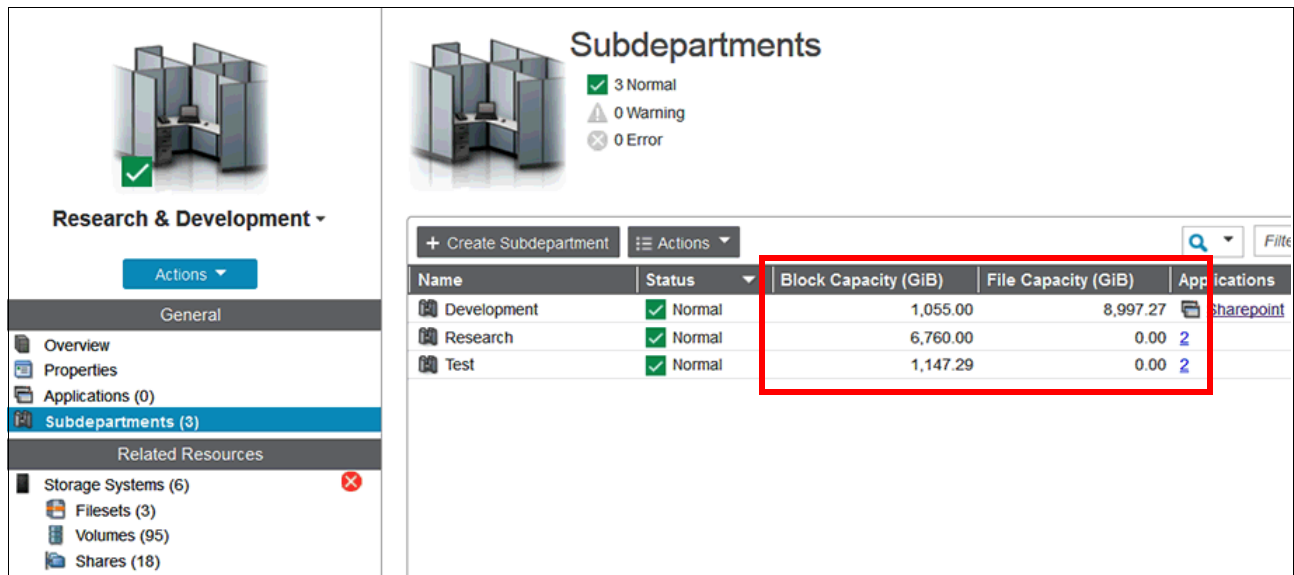


Figure 5-65 Capacity breakdown per subdepartment

5.3 Getting an overview of the resources that are associated with your applications

In the overview window of each of the departments or applications, you also can determine the application's or department's historical capacity consumption and the capacity usage per Tier.

Figure 5-66 shows an overview of the Research & Development department.



Figure 5-66 Overview with historical capacity usage and capacity usage by tier

The Spectrum Control Grouping function provides application and department insights, such as:

- ▶ Volume activity over time
- ▶ Network activity over time

You can also quickly identify resources that are related to the application or department and their condition.

5.4 Drilling into resource details for problem determination and performance analysis

In high-priority situations, for example, when there is a poor response time for client application servers, Spectrum Control assists with performance troubleshooting. This section describes how Spectrum Control can help you with performance troubleshooting.

Complete the following steps:

1. Open the server's Overview window, as shown in Figure 5-67.

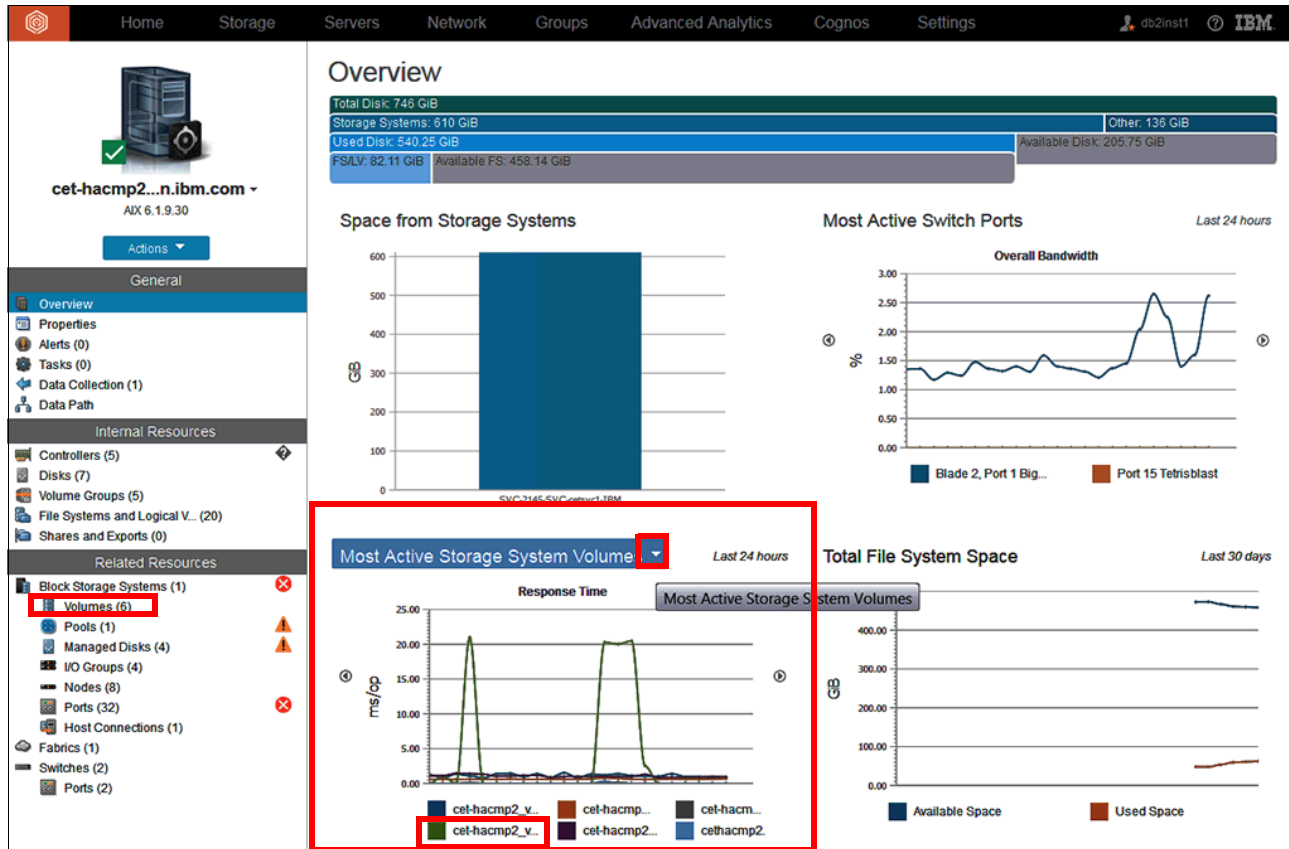


Figure 5-67 Overview window

Note: Regardless whether you deployed a storage resource agent to see the server's Overview window, you also can see the server's Overview window from *agentless servers*. For more information about agentless servers, see *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

To get help on how to create the CSV file that can be used as an import file for creating agentless servers, download a script from the ones found at the following website:

<http://ibm.co/1QEBTTi>

These scripts use SQL queries to extract the host connection information from the Spectrum Control database to a CSV file, which then can be manually imported by using the web-based GUI Agentless Server definition process.

2. You should be able to see the Most Active Storage System Volumes entry in one quadrant (you can use the drop-down menu to select this chart), as shown in Figure 5-67 on page 165. Determine the volumes that have high response times.

In our example, the cet-hacmp2_vol_tcrdblog volume shows high response times, as shown Figure 5-68.

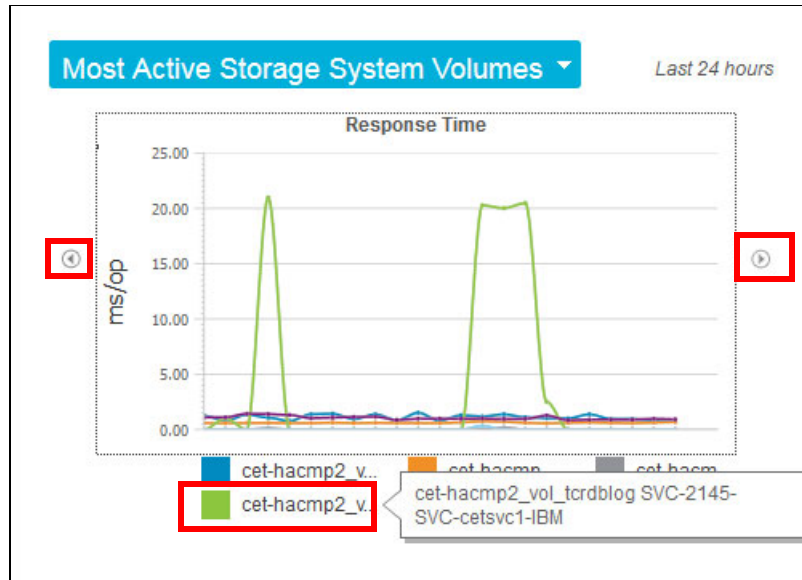


Figure 5-68 The cet-hacmp2_vol_tcrdblog volume shows high response times

You can use the arrows to view other parameters, such as Volume Utilization, I/O Rate, Data Rate, and Read Cache Hits.

3. In the left navigation pane under Related Resources, click **Volumes** and then click the **Performance** tab, as shown in Figure 5-69 on page 167, which shows the Volumes Performance for a given server.

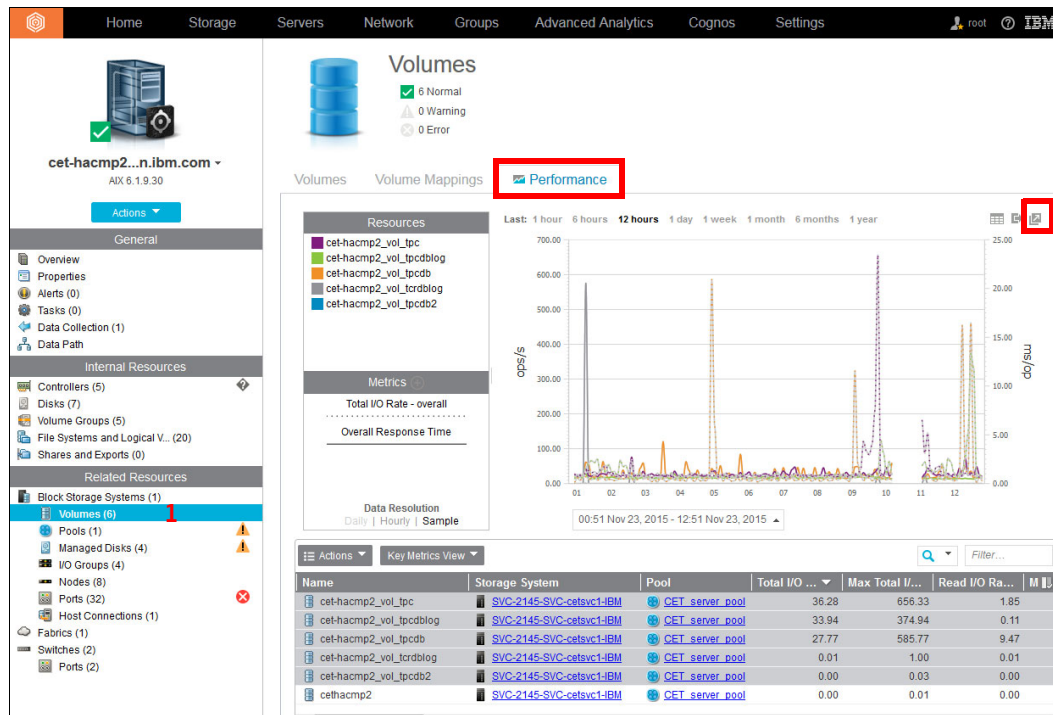


Figure 5-69 Volume Performance for a given server

- Click the **Open in new window** icon to open the graph in a new window, as shown in Figure 5-69.

5. In Figure 5-70, which shows volume performance for a given volume and a given time frame, if you cannot see the pool column, click the header row and select the **Pool** check box.



Figure 5-70 Volume performance data for a given volume and a given time frame

6. Select or clear the volumes that you want see in the chart.
7. Adjust the time frame for which the performance data is shown, either relatively or absolutely, to check whether the problem occurs on a regular basis.
8. Select the Data Resolution: Sample, Hourly, or Daily

Note: If you choose a long time frame, data resolution switches automatically to a higher aggregation level because only a certain number of dots can be displayed in the graph.

9. Select or clear the performance metrics in which you are interested. For example, select **Read Response Time** and **Write Response Time** to distinguish if the reads or the writes are causing the problem.

Figure 5-71 shows how to select Volume Performance metrics.

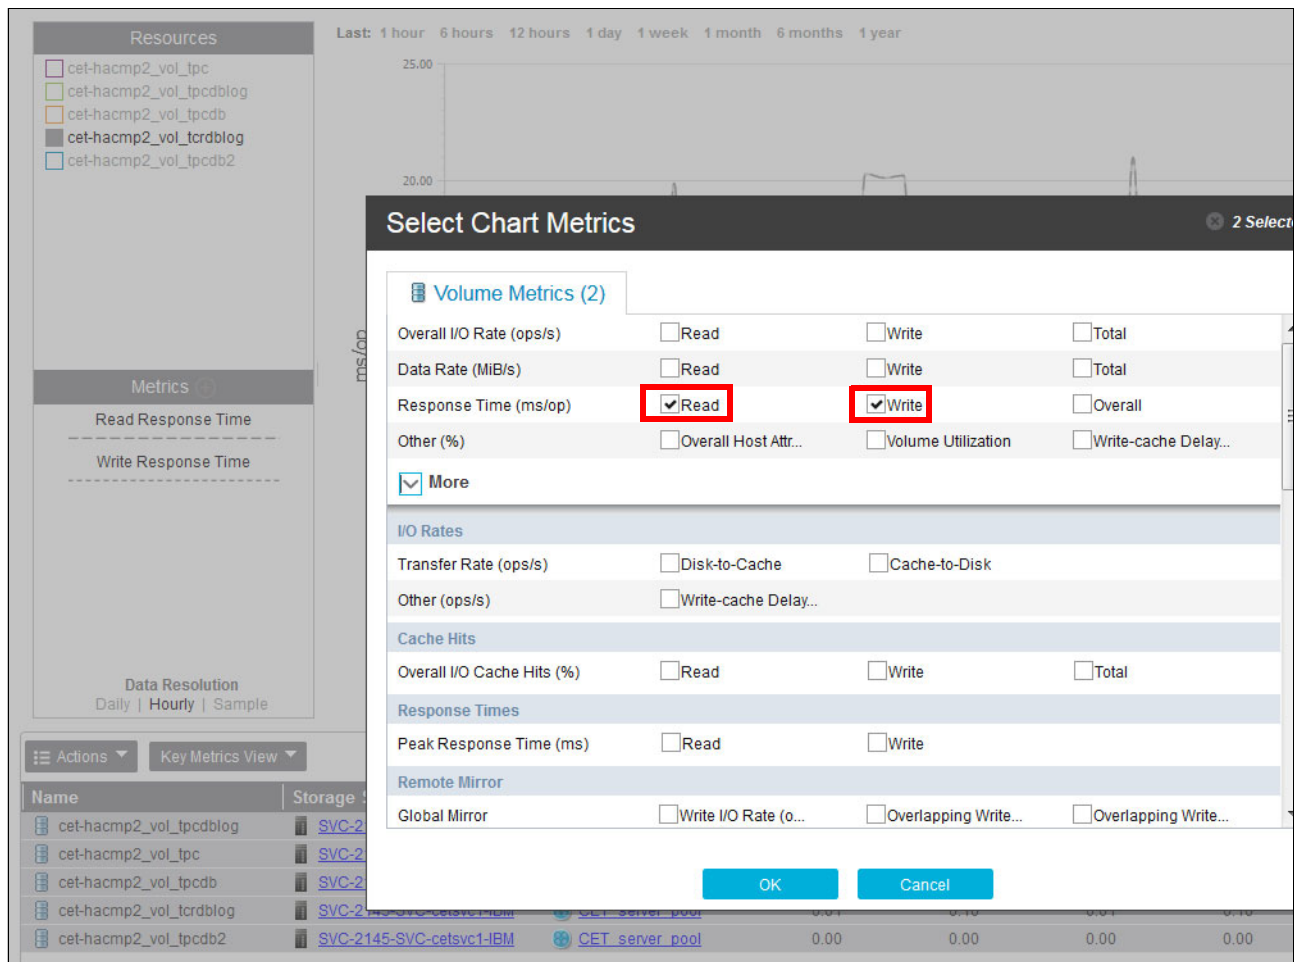


Figure 5-71 Select Volume Performance metrics

10.Click **Pools**, as shown in Figure 5-72, which shows how to select the Pool Performance chart for a given pool. Select the pool in which the volume is in and click **View Performance** to check whether the whole pool is showing high response times.

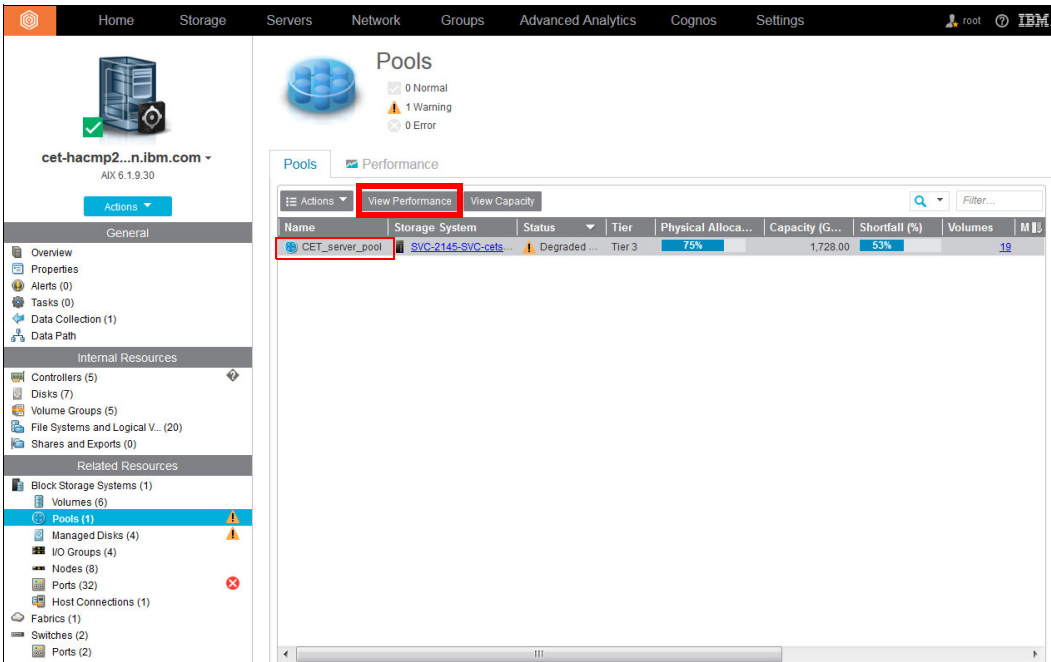


Figure 5-72 Open the Pool performance chart for a given pool

11.To show the same time frame that is shown in the Volume performance chart, click the Synchronize Time icon in the Volume performance chart, as shown in Figure 5-73, which shows the Synchronize Time icon.



Figure 5-73 Synchronize Time icon

Note: Using the Synchronize Time icon synchronizes the time range across all the performance views that are displayed in *separate* browser windows.

Use this action when you change the time range in a performance view and want to apply the same time range to the other performance views.

12. Depending on the storage subsystem on which your pool is located, you can select Disk and Volume Performance metrics so that you can determine whether the problem is caused by the back end, as shown in Figure 5-74, which shows the performance metrics for Pool performance for a given pool on a storage virtualizer.

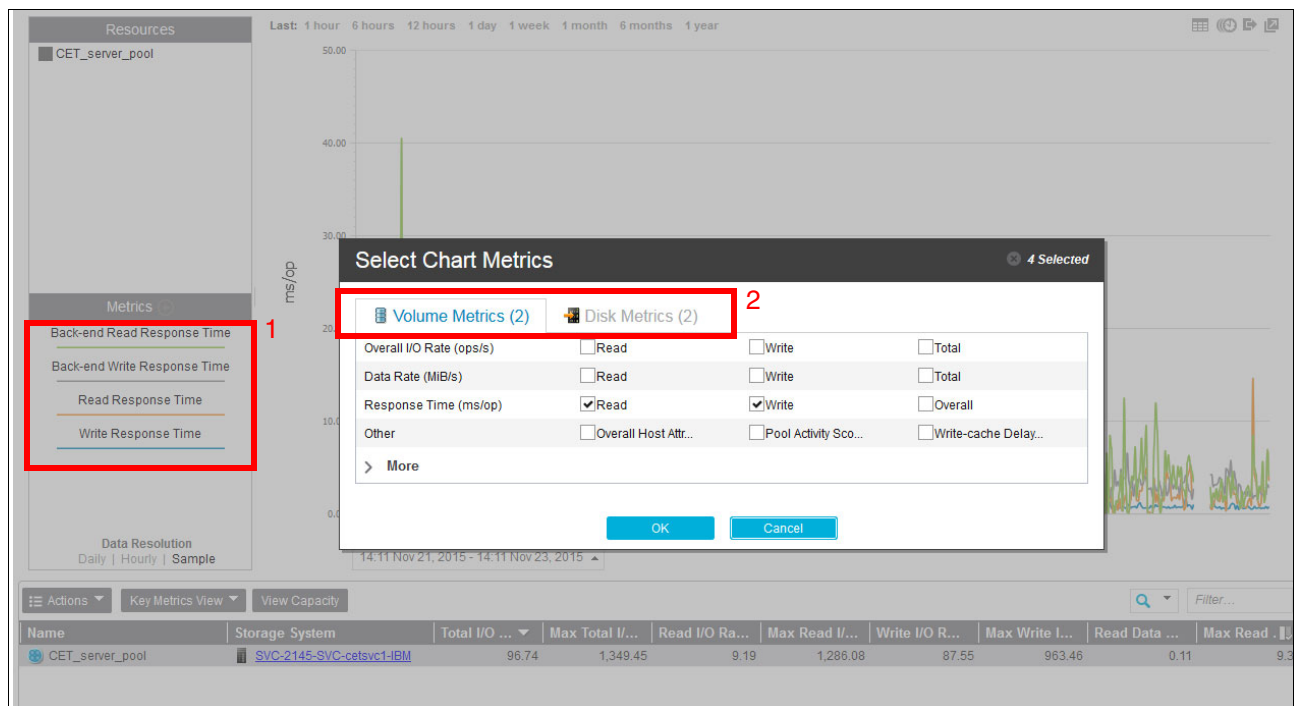


Figure 5-74 Performance metrics for pool performance for a pool on a storage virtualizer

13. From here, you can drill down, for example, to see the performance of all volumes in that pool to see whether they have the same issue. Right-click the pool and select **Volume Performance**, as shown in Figure 5-75, which shows the drill-down possibilities from a storage virtualizer pool.

Note: For information about which resources to which you can drill down, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1Yqm0WZ>

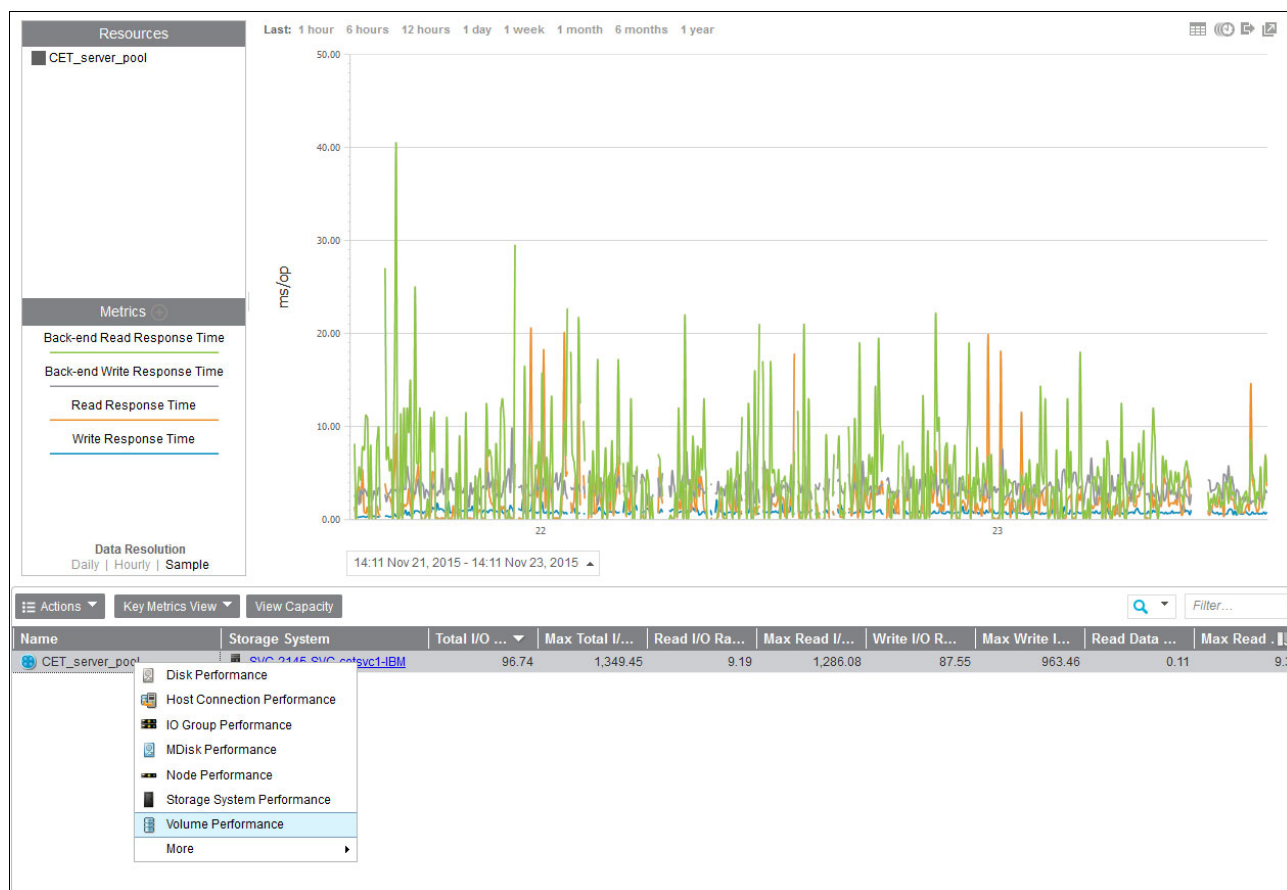


Figure 5-75 Drill-down possibilities from a storage virtualizer pool

Note: By default, the five most used volumes regarding total I/O are shown in the performance graph. But, you can select or clear the volumes for which you need the performance graph. For more information, see *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204.

14. Select the Resources volumes for which you want to see the performance data and the Metrics that are you are interested in, as shown in Figure 5-75, which shows the volume performance for a given pool.

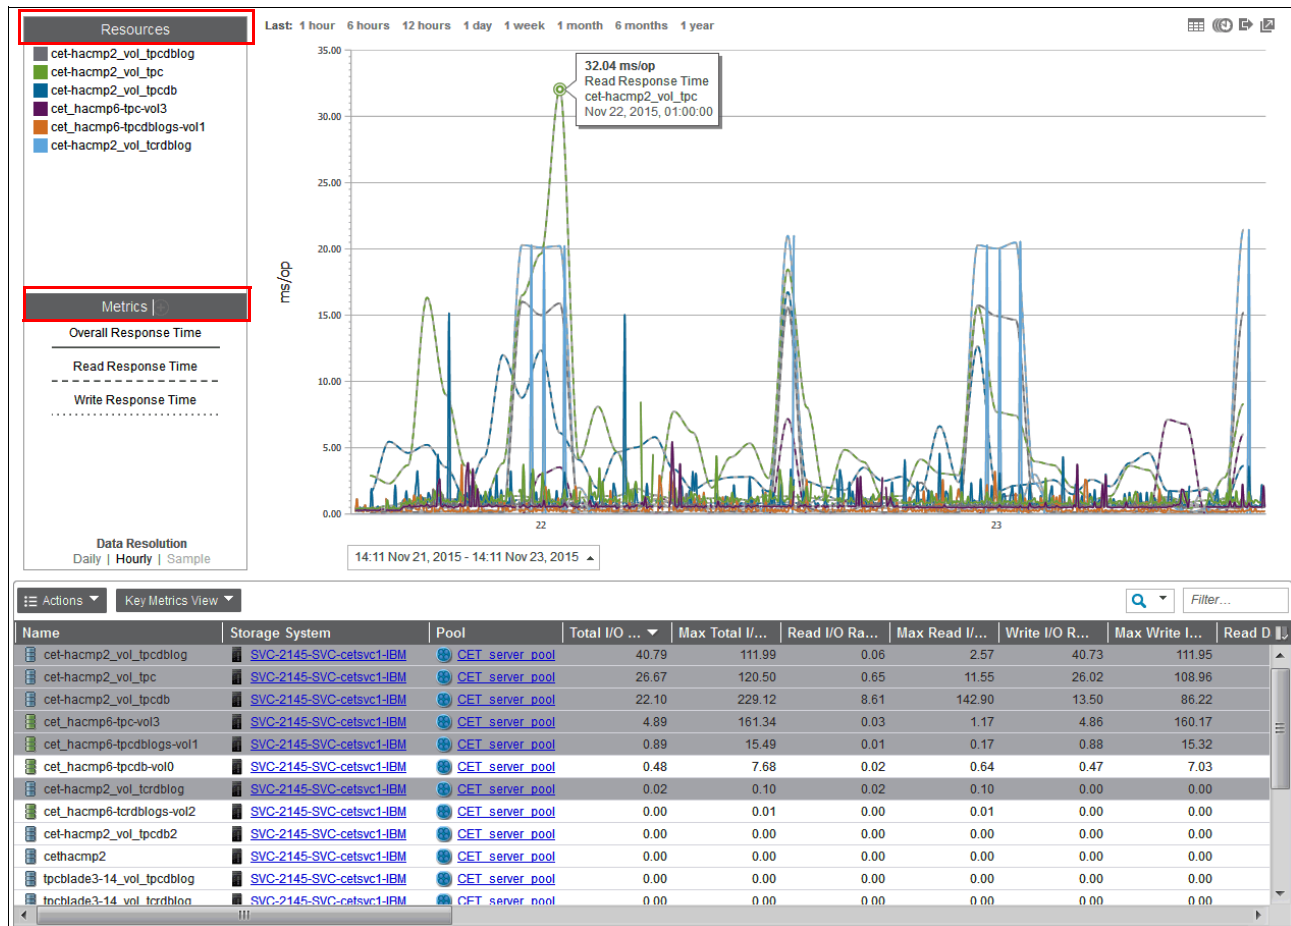


Figure 5-76 Volume performance for a given pool

15. You can easily correlate the graph to the volume by hovering your cursor over the datapoints, as shown in Figure 5-76.

16. Drill down to the pool's MDisk Performance, as shown in Figure 5-75 on page 172, to see whether one of the disks is providing too long response times.

17. Drill down to Disk Performance Metrics to check whether the problem is caused by any disk as shown in Figure 5-77, which shows how to drill down from MDisk performance to Disk performance.

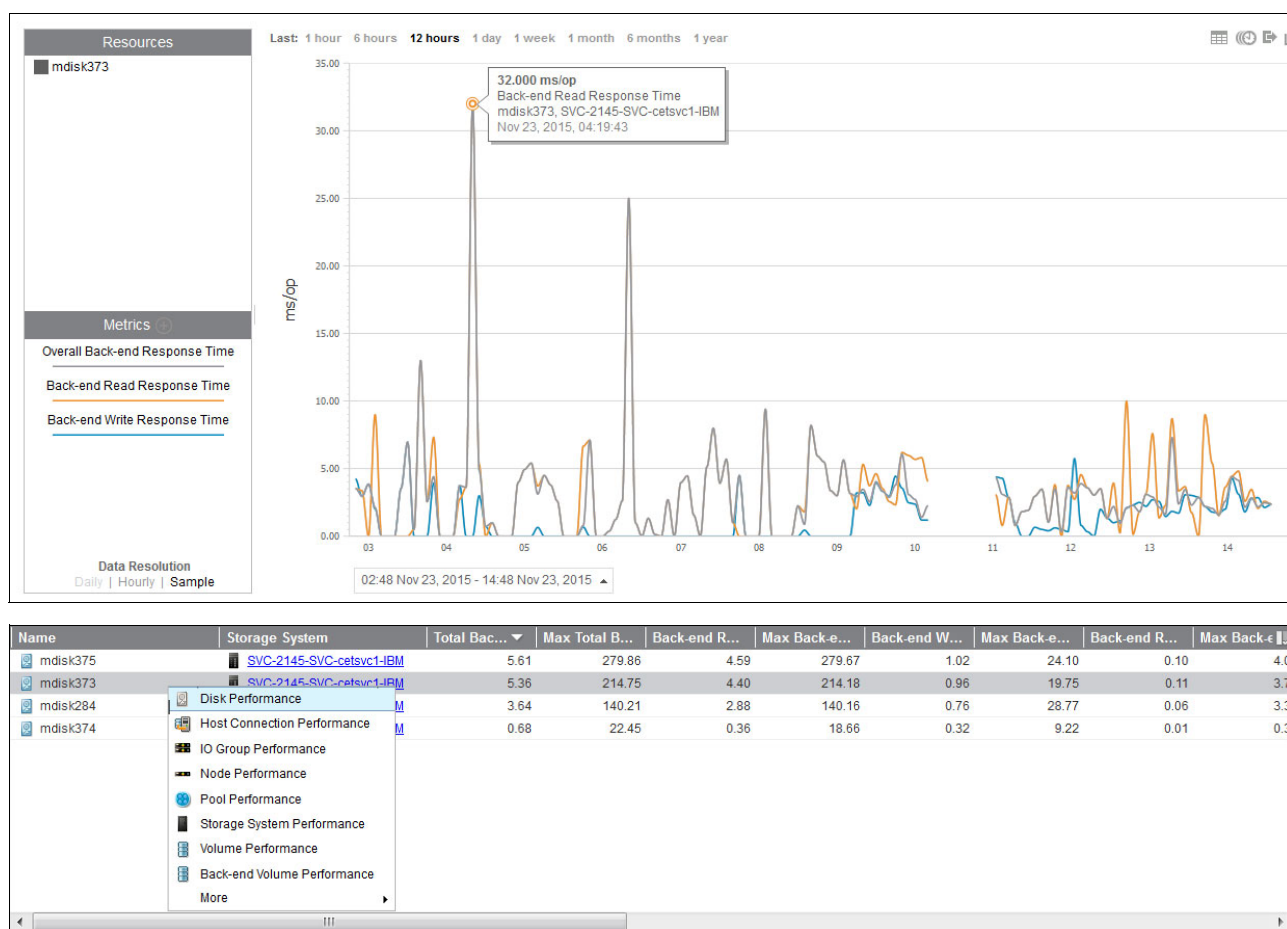


Figure 5-77 Drill down from MDisk performance to Disk performance

18. If you are troubleshooting a volume of a storage virtualizer pool, be sure to check the node balance of your storage system to make sure that everything is balanced. To do so, go to the Pool Performance chart and drill down to the Node Performance chart, as shown in Figure 5-75 on page 172, which shows that not all nodes of the storage virtualizer are used equally.



Figure 5-78 Not all nodes of the storage virtualizer are used equally

Note: For more information about available performance metrics in Spectrum Control, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1XafEZA>



Advanced alert management

With IBM Spectrum Control (Spectrum Control) Standard and Advanced Edition, the ability to define alerts and be notified of capacity changes, configuration changes, and performance conditions on resources is enhanced.

This chapter describes and provides practical scenarios that use this extensive set of new alert conditions to better monitor and identify more of the critical situations in your environment. These scenarios are focused on adding value from a storage administrator perspective. Top-level entities (storage system, switch, and server) are covered.

For more information about alerting in Spectrum Control, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1UJizmy>

6.1 Alerting overview

Determining when and how you are alerted to conditions and violations of performance thresholds within a storage environment is important to helping you maintain and administer storage resources.

With Spectrum Control, you can now apply an extensive set of alert conditions to detect problems. You can define alerts for almost all attributes of a resource, including attributes for status, configuration, capacity, and performance. Use this extensive set of alert conditions to better monitor your environment without needing to look at all the metrics manually. Instead, look once and then define thresholds according to your environment. From time to time, you should review your alerts and thresholds because business requirements change.

In Spectrum Control, there are over 50 conditions on which you can receive alerts, including conditions such as Used Capacity, Total Volume Capacity, and Available Pool Space. For performance alerts, you can now receive alert on any of the hundreds of metrics that are collected by Spectrum Control.

You can now configure multiple alerts for the same condition with different severity, suppression, and notification settings.

6.2 Alerts

The Spectrum Control Alert Server expands the ability of Spectrum Control to alert at the storage component level, providing added flexibility for capacity and performance alerts and attribute-based configuration alerts. The alert server is an independent component server that is built on WebSphere Liberty Core.

For more information about WebSphere Liberty Core, see WebSphere Application Server IBM Knowledge Center, found at:

http://www.ibm.com/support/knowledgecenter/SSEQTP/mapfiles/product_welcome_was.html

The Alert Server is packaged with all Spectrum Control offerings and is automatically installed. It coexists with both the Device and Data Servers with the same level of service controls and logging.

If you use a custom start/stop script, you must include the Alert Server as well.

On the Spectrum Control Dashboard, you see how many alerts of which category occurred in the last hour, last day, or last week.

Figure 6-1 on page 179 shows the Alert Overview on the Dashboard.

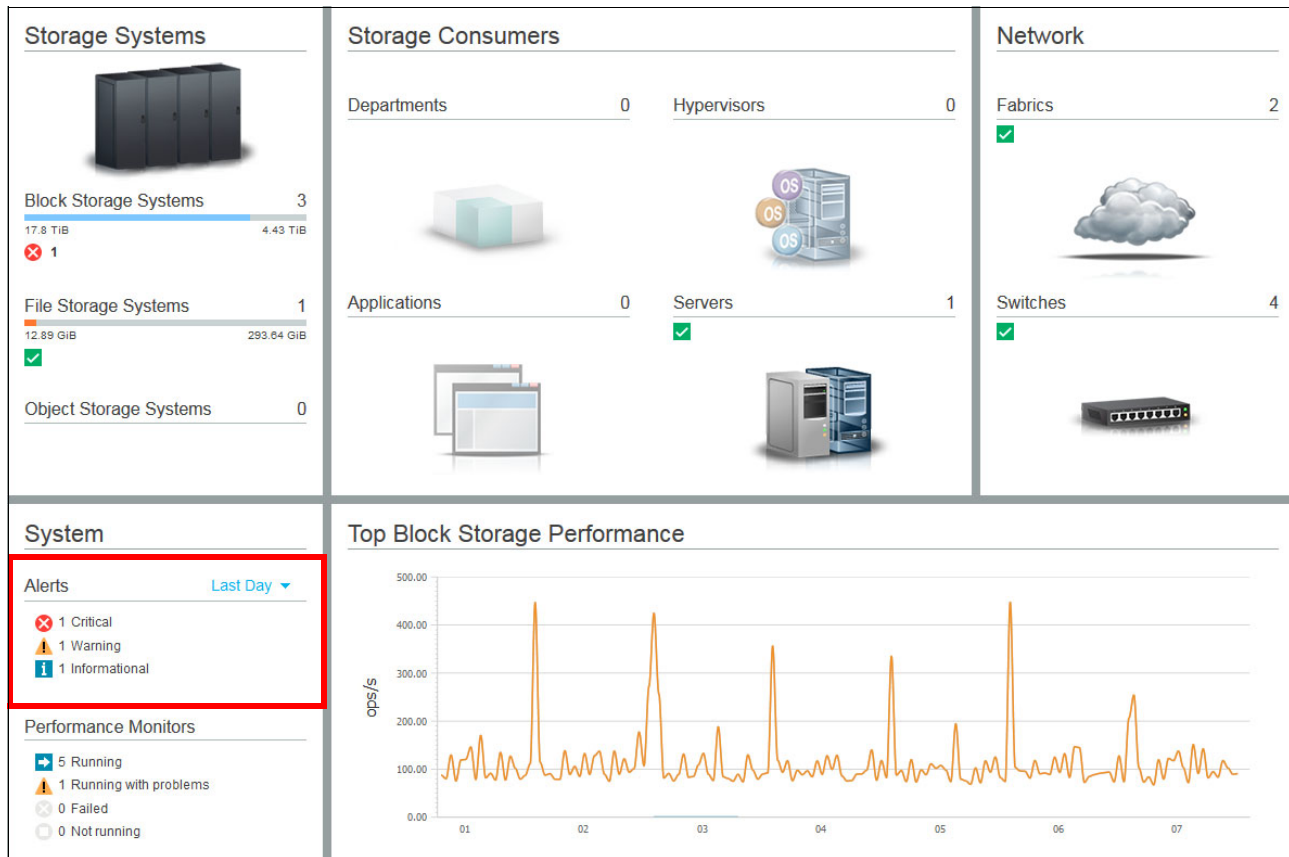


Figure 6-1 Alert Overview on the Dashboard

6.2.1 Features

The following list shows the main features that are available starting with Spectrum Control V5.2.8:

- ▶ Expands the total number of storage resources that can be tracked and alerted.
- ▶ Classifies alerts in three categories for simplified configuration: Informational, warning, and error.
- ▶ Centralizes alerts and threshold violations under a single tab, where you have the possibility to acknowledge them.
- ▶ Three different Alert or Threshold Violation Suppression settings are available:
 - Only alert once until problem clears.
 - Only generate alerts every x minutes, hours, or days.
 - Do not alert until condition has been violated for more than x minutes, hours, or days.
- ▶ Alert notification supports email, Omnibus, SNMP, and Windows/UNIX Event Logs.
- ▶ Multiple thresholds can be set for the same metric for different severities, which can be combined with specifying different email recipients or different scripts.

6.2.2 Triggering conditions for alerts

You can set up Spectrum Control so that it examines the data of your resources for the conditions and performance thresholds that you specify.

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. Some attribute and performance conditions require that you enter values for triggering alerts. In general, the following types of conditions can trigger alerts:

- ▶ An attribute of a resource changed.
- ▶ A performance metric for a resource fell outside a specified range.
- ▶ The storage infrastructure changed.
- ▶ Data cannot be collected for a resource.

6.2.3 Reporting

You can use reporting functions to view overview and detailed information about your storage, for example:

- ▶ Spectrum Control web-based GUI
- ▶ Native SQL
- ▶ Cognos BI

For more information, see Chapter 5, “Service-oriented storage management using IBM Spectrum Control” on page 113.

6.3 Alerting and event processing

With Spectrum Control, you can define alerts for almost all attributes of a resource, including attributes for status, configuration, capacity, and performance. For performance alerts, you can now alert on any of the hundreds of metrics that are collected by Spectrum Control.

Alerts and Threshold violations are defined on the device level, either from the resource list page or from the resource page.

Figure 6-2 shows the Edit Alert Definitions menu on the resource list page.

The screenshot displays the 'Block Storage Systems' section of a management interface. At the top, there's a navigation bar with 'Home', 'Storage', 'Servers', 'Network', and 'Groups'. Below this, a summary shows '2 Normal', '0 Warning', and '1 Error' status. The main table lists storage systems with columns for Name, Condition, Probe Status, and Performance. A context menu is open for the 'Storwize V7000-2076-V7000_21' system, showing options like 'View Properties', 'View Details', 'View Performance', 'Data Collection', 'Edit Alert Definitions' (highlighted), 'Edit Alert Notification Settings', 'Analyze Tiering', 'Add to Capacity Pool', 'Connections', 'Remove', and 'Open Storage System GUI'.

Name	Condition	Probe Status	Performance...
FlashSystem 840-9840-FlashSy...	✓ Normal	✓ Successful	❌ Disabled
SVC-2145-svc_27_28-IBM	✓ Normal	✓ Successful	➡ Running
Storwize V7000-2076-V7000_21	✓ Normal	✓ Successful	⚠ Running wit...

Figure 6-2 Edit Alert Definitions menu on the resource list page

Figure 6-3 shows the Alert Definitions menu on the resource page.

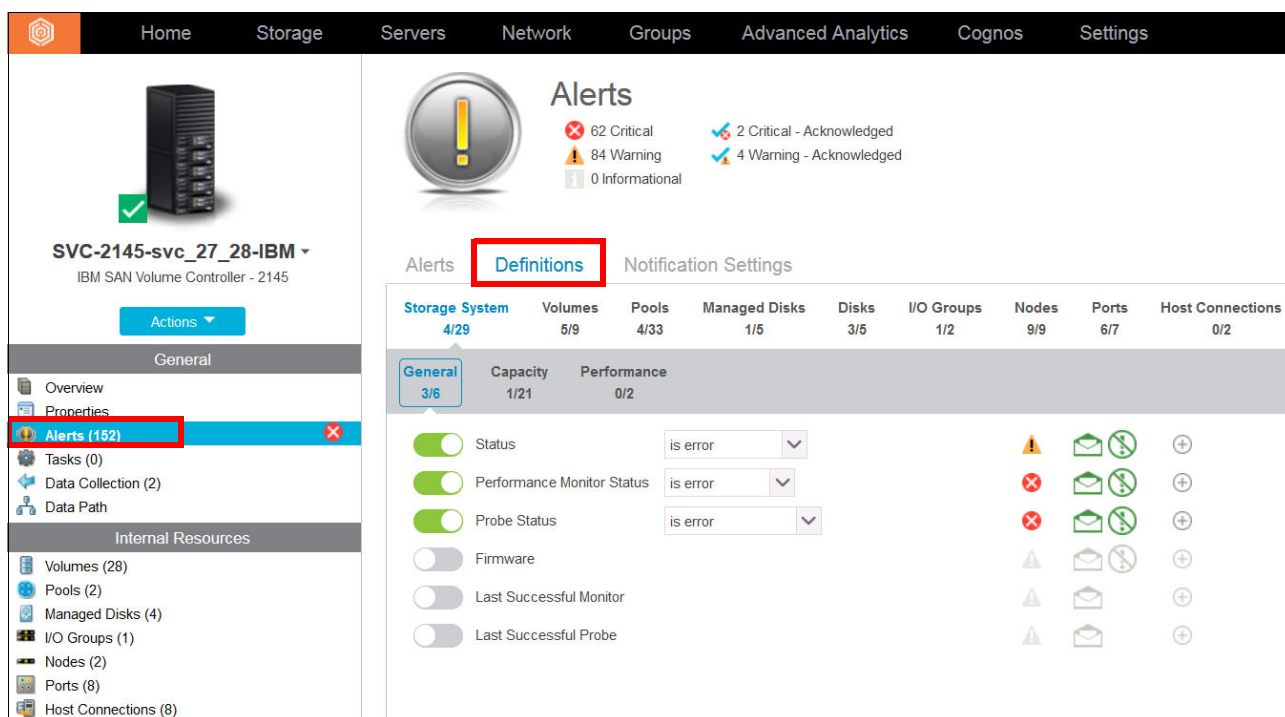


Figure 6-3 Edit Alert Definitions menu on the resource page

Reference: For more information about alert and event processing, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/1UJizmy>

For improvements of the alerting function in future versions of Spectrum Control, see the *What's new* topic in the Spectrum Control IBM Knowledge Center.

The following sections describe practical user scenarios from a storage administrator perspective that can be implemented in your business environments.

6.3.1 Scenario 1: Monitoring and being notified about storage utilization

In this scenario, a storage administrator wants to be notified if the file systems, which can be on a server, NAS, or hypervisor, are being underutilized for an extended period so that storage can be reclaimed to minimize storage costs. (This scenario can be expanded to go into detail about capacity threshold levels, the level of suppression, and who to notify).

To delve deeper into this scenario, assume that the storage administrator recently was notified that the request to purchase more storage in the coming year was denied. Given that the budget remains flat, the administrator wants to put a plan in place to identify whether any of the server storage is being underutilized.

Before making any major changes, the administrator wants to be notified if any of the file systems have extra capacity so the administrator can then take a deeper look into what storage might be a target for storage reclamation. To accomplish this task, complete the following steps:

1. From Spectrum Control, go to **Server's Overview**. Select **Alerts** in the left pane, as shown in Figure 6-4) and select the **Definitions** tab in the Definition pane.

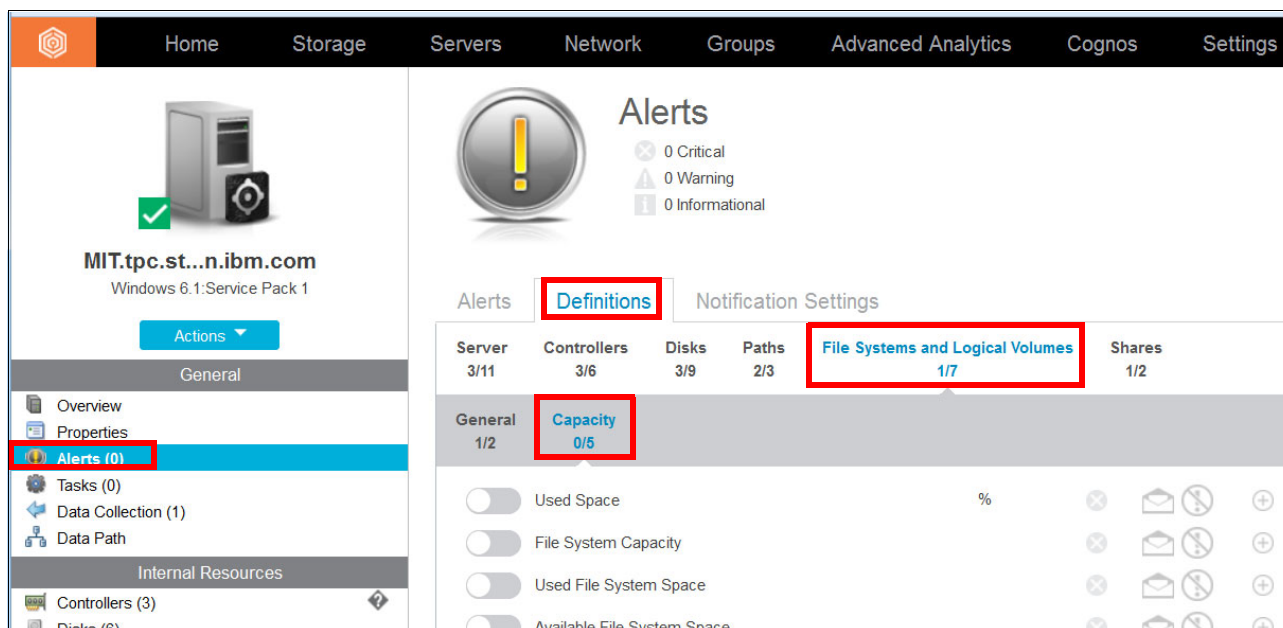


Figure 6-4 Alert definition window

2. Select the **File Systems and Logical Volumes** tab. On this tab, you notice that the alerts are categorized into General for Configuration alerts and Capacity for all Capacity-related alerts. Select the **Capacity** tab.

3. Enable **Used Space** and set it to be less than or equal to 25%, as shown in Figure 6-5, which shows an example of how to set a Used Space alert on a file systems or logical volume level. Now, you can track all file systems that are using less than 75% of the available capacity and might be good candidates for storage reclamation.

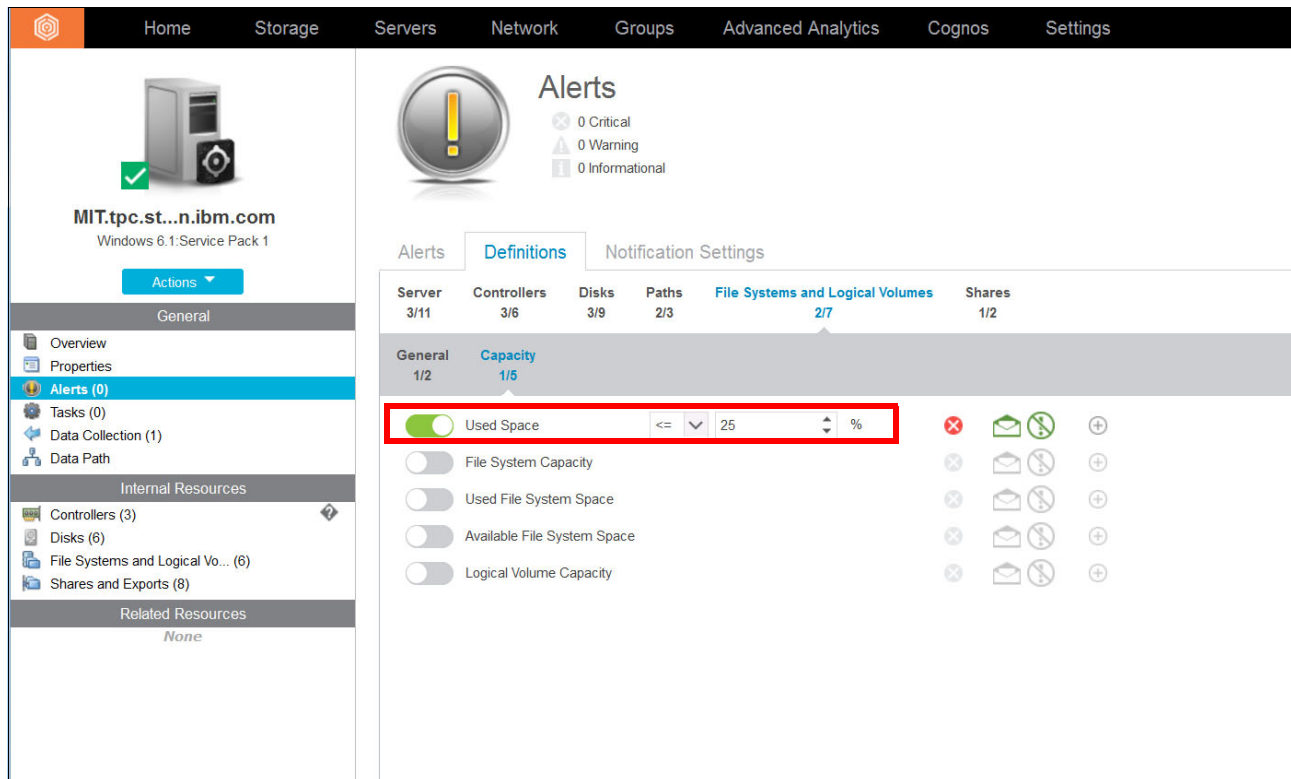


Figure 6-5 Example of how to set a Used Space alert on file systems or logical volume level

4. To notify the correct people, in Figure 6-6, which shows how to override notification settings for a special alert, click the envelope icon to set alert notification email addresses for this specific alert. Then, click **Override Notification settings**, select the **Email** check box, insert the email address, and click **Done**.

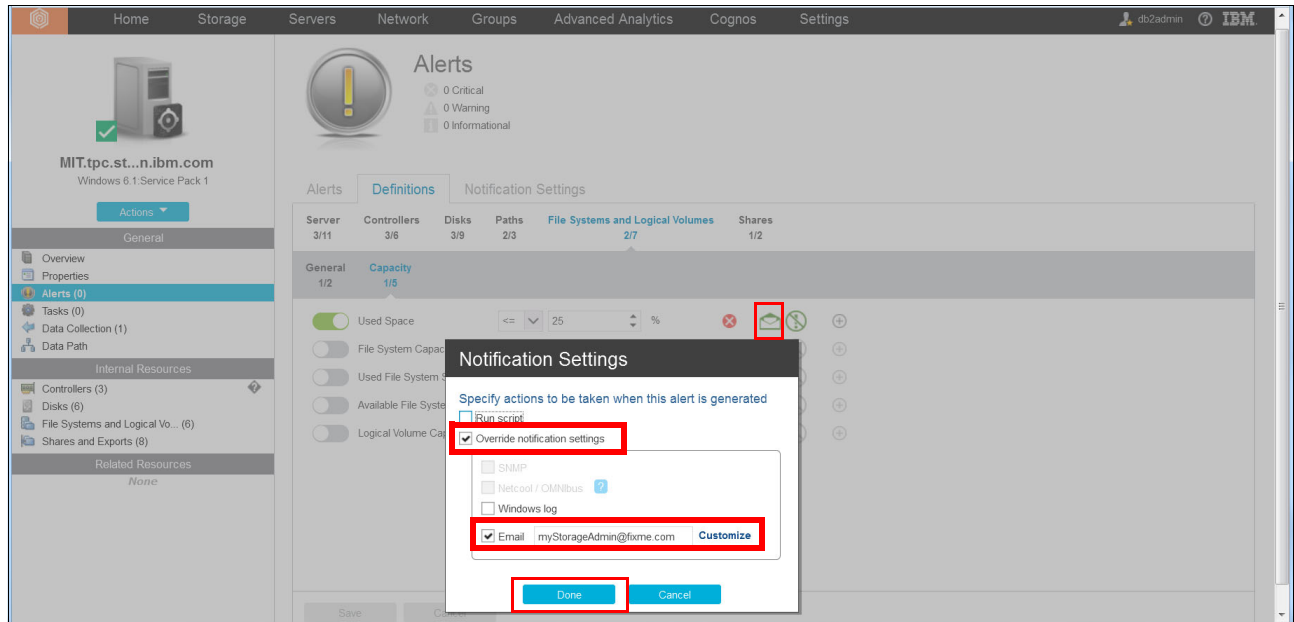


Figure 6-6 How to override notification settings for a special alert

5. If you know that the file system is being monitored daily and want to avoid a flood of emails, you can change the default suppression setting from **Only alert once until the problem clears** to **Only generate alerts every 5 days**. Click the struck through exclamation mark next to the envelop icon, as shown in Figure 6-6 on page 185. Then, click **Only generate alerts every**, select **5**, and select **days** as shown in Figure 6-7, which shows an example of alert suppression. With this suppression setting, you apply a fixed window of time where suppression occurs regardless of violation. This setting ensures that it does not matter whether the utilization within the five days is good or bad.

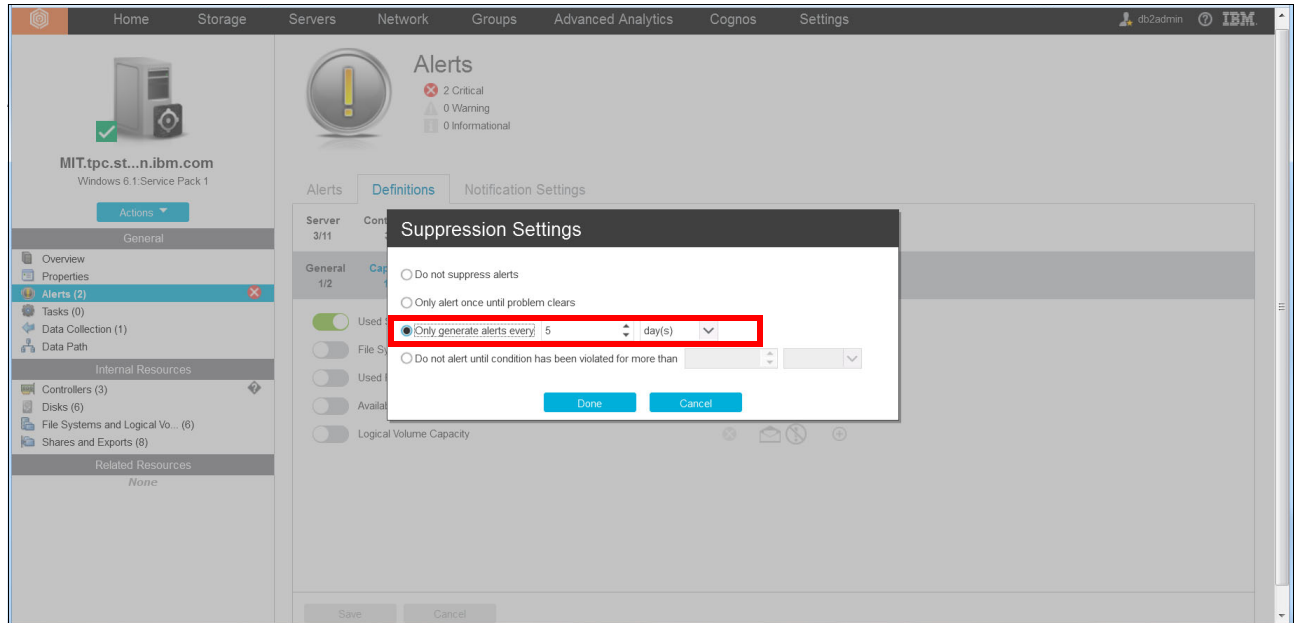


Figure 6-7 Example of alert suppression

Alternatively, you can choose **Do not alert until condition has been violated for more than 5 days**, as shown in Figure 6-8, which shows the Suppression Settings dialog box. In this case, the violation must occur consistently for more than 5 days to trigger an alert without any clearing of the condition. With this option, you will get an initial alert when the problem occurs.

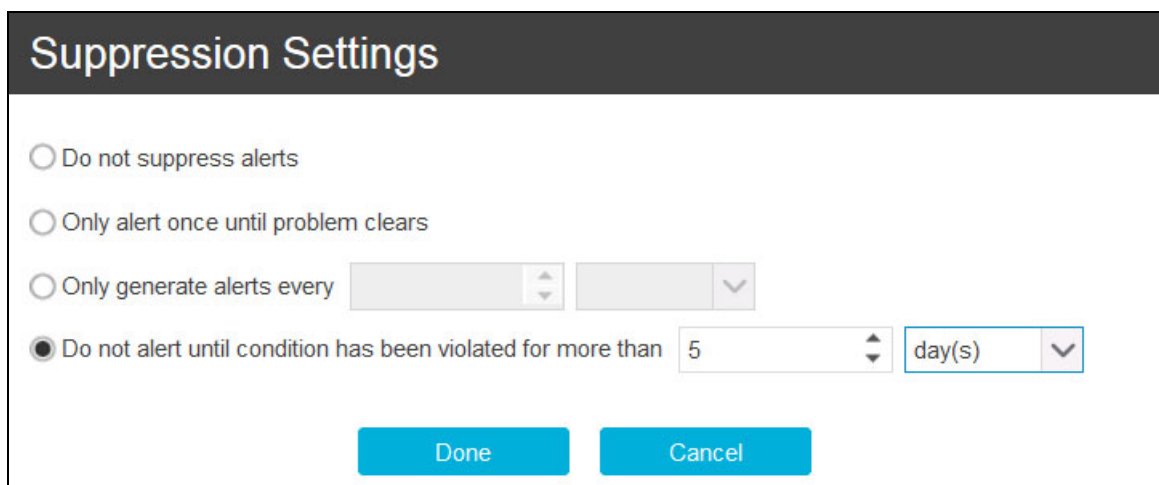


Figure 6-8 Suppression Settings dialog box

6.3.2 Scenario 2: Performance troubleshooting with advanced alert management

In this scenario, a storage administrator wants to track the performance of an IBM Storwize V7000 Unified storage system. Some of the clients are complaining about general performance degradation, so it is beneficial to set a few metrics on Spectrum Control to send notifications that are related to device performance.

To delve deeper into this scenario, the storage administrator wants to make sure all of the processors are being used to balance the subsystems workload and that none of the central processing units (CPUs) are taking on too much of the workload.

Also, the administrator wants to know whether the port bandwidth might be the problem, and specifically whether it is the sending or receiving of data that is causing a bottleneck on any ports. As with the CPUs, the administrator wants to know whether any ports are being underutilized so that the administrator can rebalance workload if necessary. The last thing that the administrator must consider is the backup window each night. For example, performance always spikes each night for about 3 hours while clients initiate their backups, so the administrator wants to *not* be notified of this performance slow down.

To resolve this scenario, complete the following steps:

1. From the Spectrum Control Dashboard, go to the Overview window for the IBM V7000 Unified storage system (Figure 6-9, which shows how to define the metrics on which you want to define thresholds). In the V7000 Overview window, open the Alerts Definitions pane, click **Ports**, and click the **Performance** tab. Click **Add Metric**, and select the performance alerts check boxes for **Overall Port Bandwidth Percentage**, **Port Send Bandwidth Percentage**, and **Port Receive Bandwidth Percentage**. Click **OK**.

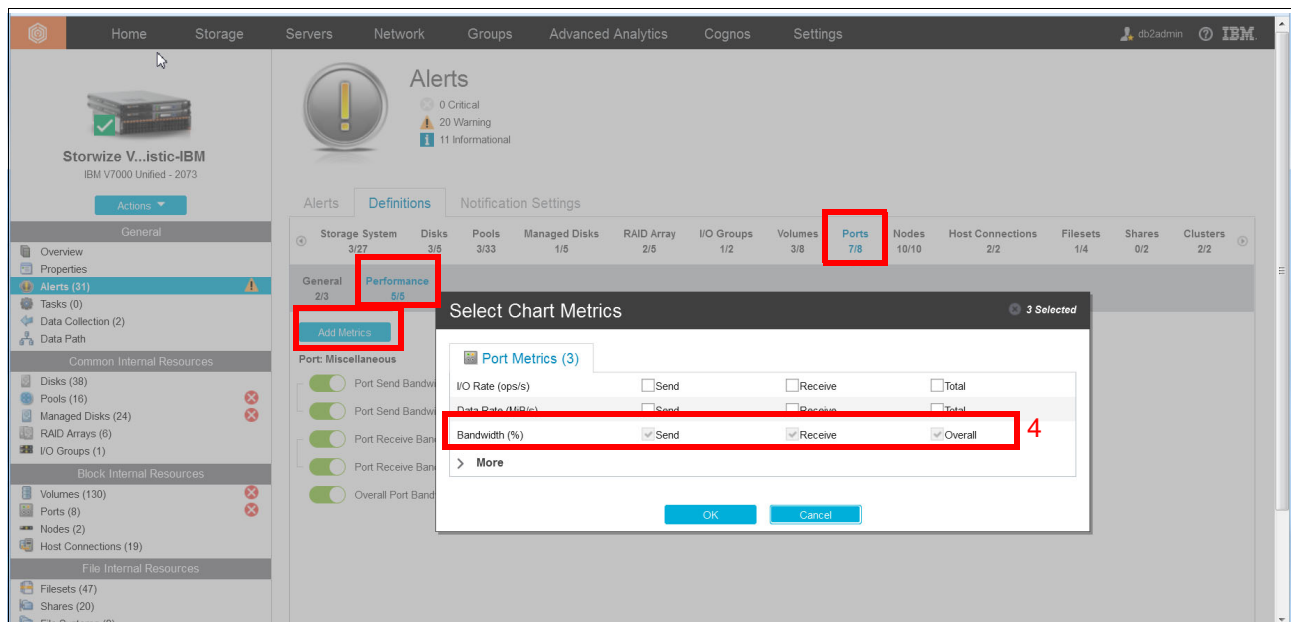


Figure 6-9 Define the metrics on which you want to define thresholds

- After enabling all metrics, set Overall Port Bandwidth Percentage to less than 15% (Figure 6-10, which shows how to set thresholds for individual performance metrics) to track whether any Subsystem Ports are being underutilized. The Overall Port Bandwidth Percentage alert can be defined with Informational severity level because this is not an immediate performance impact. To do so, click the drop-down arrow next to the red cross and select the yellow exclamation mark. For Port Send Bandwidth Percentage and Port Receive Bandwidth Percentage, define a Warning severity threshold at greater than 75%.

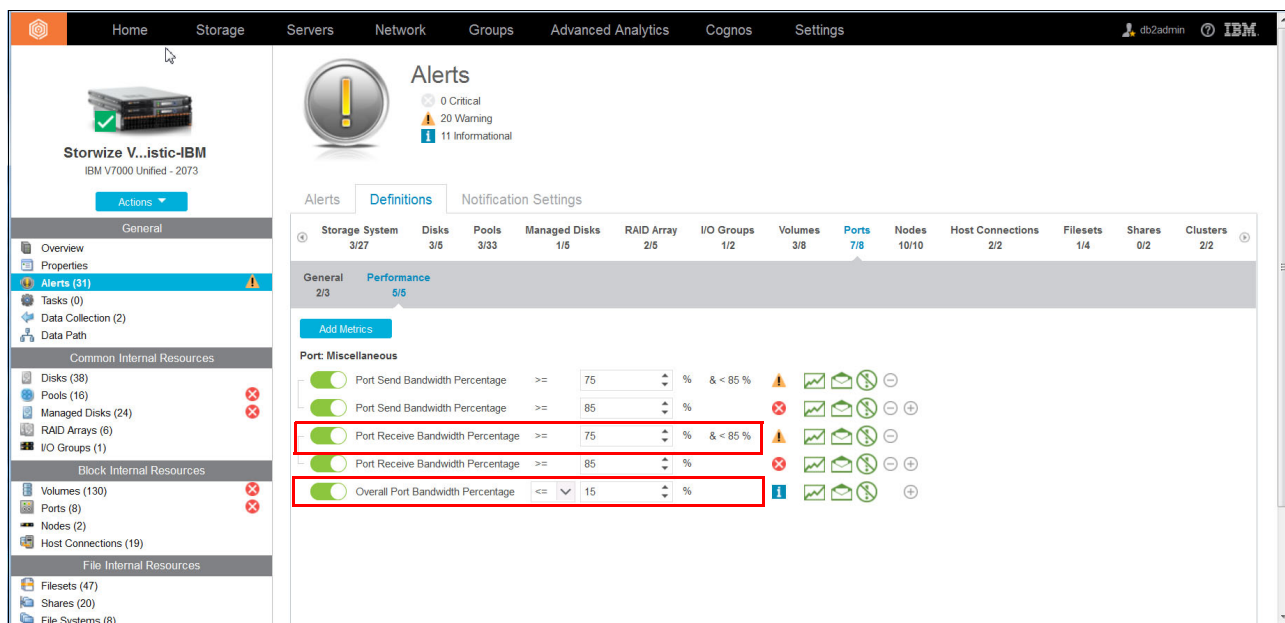


Figure 6-10 Set thresholds for individual performance metrics

- Set the Critical severity threshold at greater than 85%. To do this task, define the Warning threshold alert at greater than 75% as shown in Figure 6-11, which shows how to set a warning and a critical threshold on one metric, then click the (+) icon for this alert to define the Critical severity threshold of greater than 85%.

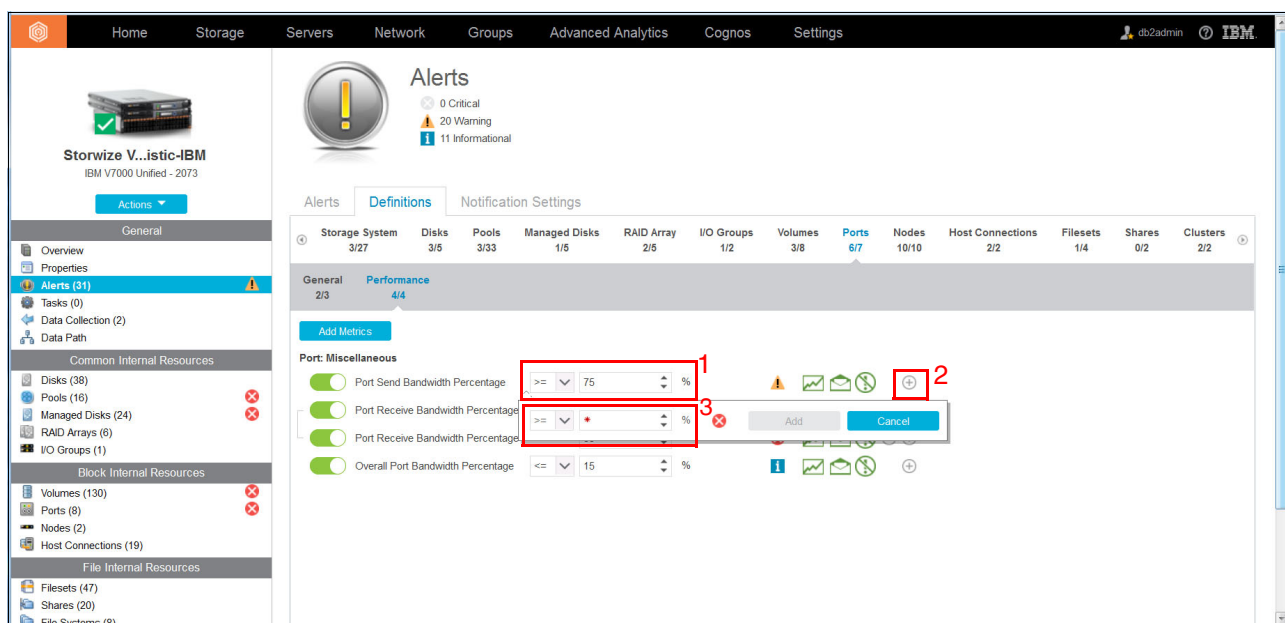


Figure 6-11 Set a warning and a critical threshold on the same metric

- Now, from the Port Performance tab, click the **Nodes Performance** tab (see Figure 6-12, which shows how to set performance thresholds for nodes, on the Alerts Definitions window). Then, use the same methods to define Node Performance Alerts for the subsystem. To identify idle CPU usage, add the System CPU Utilization metric from the Add Metrics dialog box and define the alert with less than or equal to 10%. Define this alert with Informational severity because it is not an immediate performance impact. To track excessive CPU utilization, define the System CPU Utilization with a Warning threshold of greater than 75%, then click the (+) icon on this alert to define the Critical threshold of greater than 90%.

The screenshot shows the IBM Storwize V7000 Unified Alerts Definitions window. The 'Nodes' tab is selected in the top navigation bar, labeled '1'. The 'Add Metrics' button is highlighted with a red box and labeled '2'. The 'CPU Utilization' section contains three rows of metrics, each with a red box around the threshold and severity settings. The first row is 'System CPU Utilization' with a threshold of '75' and severity 'Warning', labeled '4'. The second row is 'System CPU Utilization' with a threshold of '90' and severity 'Warning', labeled '4'. The third row is 'System CPU Utilization' with a threshold of '10' and severity 'Informational', labeled '3'.

Figure 6-12 Set Performance Thresholds for Nodes

5. To account for the known backup window, you must set suppression on each alert you want to avoid during this window. To do this task, click the suppression icon on each alert. The default value is to alert only once until the problem is cleared. Because the backup window lasts for 3 hours each night, you must do the following steps:
 - a. Select the suppression radio button **Do not alert until condition has been violated more than 3 hours**, as shown in Figure 6-13, which shows an example for alert suppression that can be used for a backup window. This suppression setting is okay during the backup window, but it also means that during the rest of the day you get alerts only if the corresponding thresholds also are continuously violated for 3 hours.

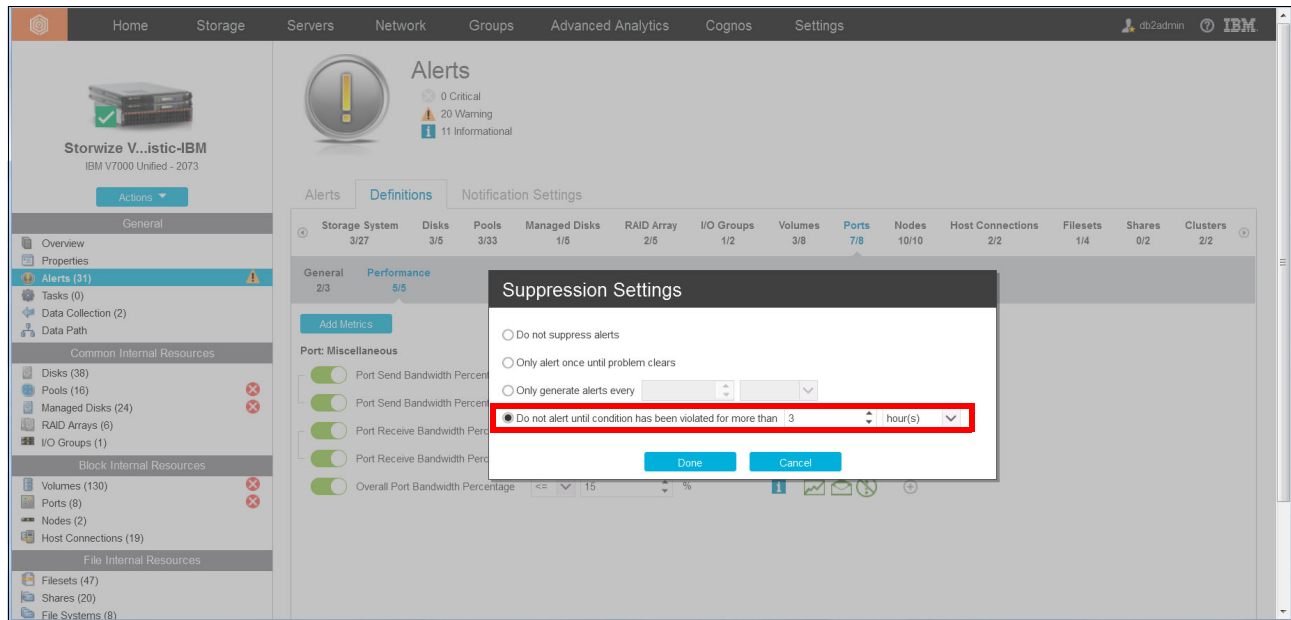


Figure 6-13 Example for alert suppression that can be used for a backup window

- b. What you need is a blackout window for the backup time frame, but this is not available yet. The alternative suppression setting is **Only generate alerts every three hours**. With this option, you receive the alert notification only once.

Figure 6-14 shows another example for alert suppression that can be used for a backup window.

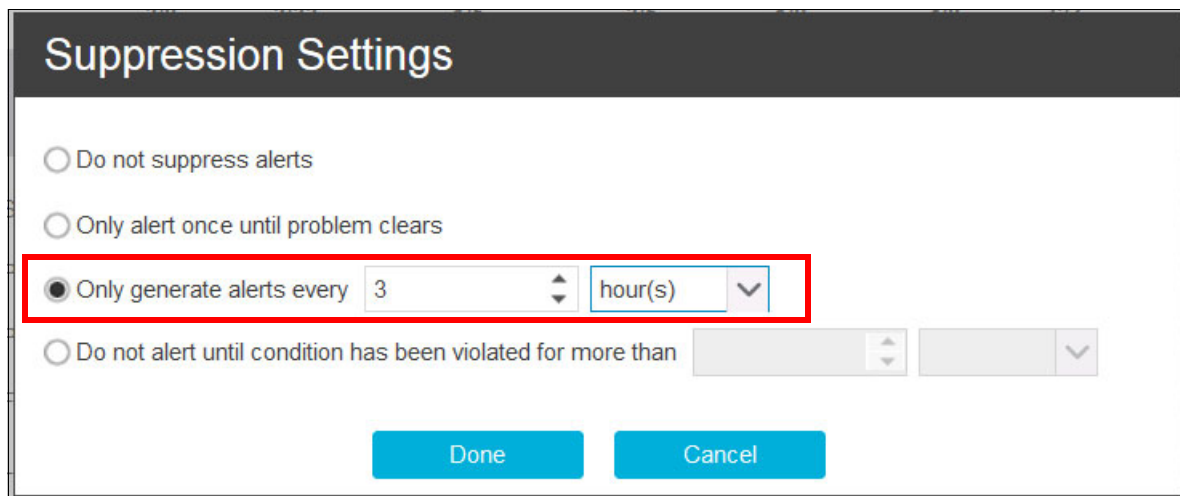


Figure 6-14 Another example for alert suppression that can be used for a backup window

6.3.3 Scenario 3: Closely monitoring new application usage with advanced alert management

In this scenario, a storage administrator recently received an email that a department that the administrator supports is upgrading one of their medical imaging software applications. Clients can use this new application to scan medical images with a greater level of detail, but it also requires the size of the application request to almost triple. Previously, the application requests were about 4 KB and they had about 5000 I/O requests per second. The department is using an IBM Storwize V7000 Unified storage system for storage and managed to stay under the maximum transfer rate.

To delve deeper into this scenario, the storage administrator, as the department implements its new application, wants to closely monitor the I/O transfer size on the back-end storage that is assigned to make sure that it is sufficient. The existing performance monitoring of the Storwize V7000 Unified storage system experienced some intermittent gaps, so the administrator wants to make sure that they can quickly gather diagnostic tests on the Spectrum Control server so that support can be quickly engaged if needed.

To track the performance of the Storwize V7000 Unified storage system, you must run the performance monitors. If there is a problem with the performance monitors, a script is initiated to open a ticket in the company's ticket interface.

Complete the following steps:

1. Go to the Block Storage for the IBM V7000 Unified Overview window and click the **Alert Definitions** tab. Click the **Managed Disks Performance Alerts** tab. Click **Add Metrics** and select the check boxes for **Total Back-end I/O Rate**, **Overall Back-end Transfer Size**, and **Back-end Write Transfer Size**.

Figure 6-15 shows an example of how to select performance thresholds for MDisks.

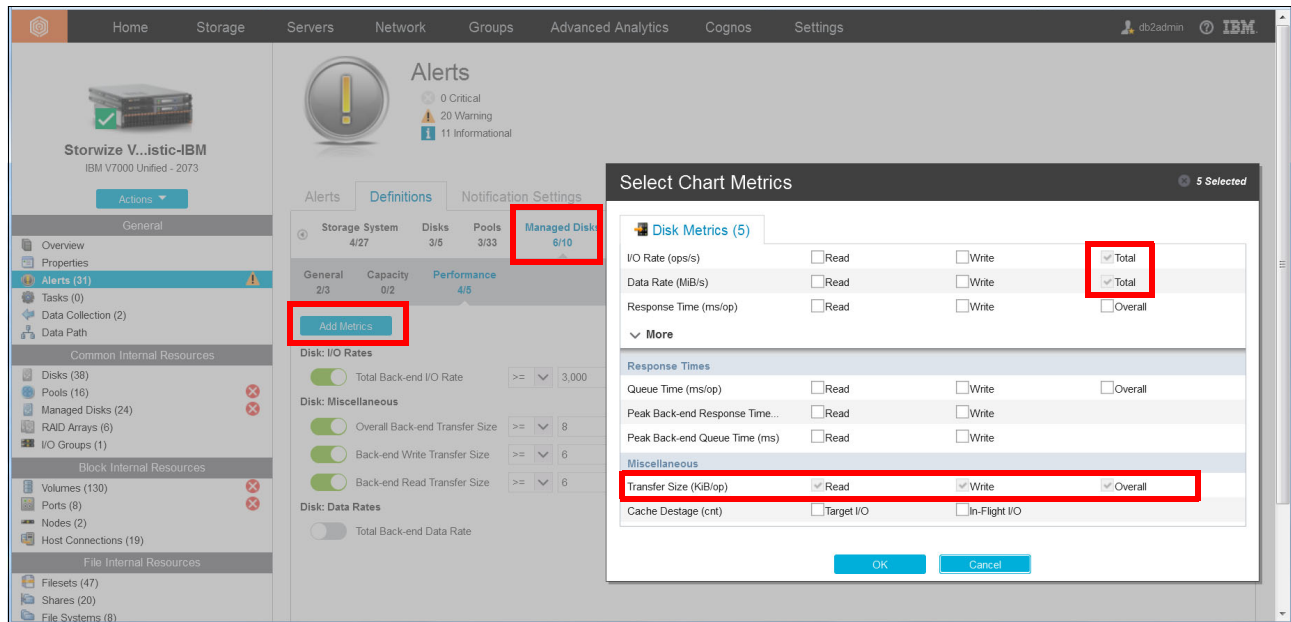


Figure 6-15 Example of how to select performance thresholds for MDisks

- After enabling all metrics, set the **Overall Back-end Transfer Size** alert to greater than 8 KiB/op, as shown in Figure 6-16, which shows an example on setting performance thresholds for MDisks. Then, define the **Back-end Write Transfer Size** and **Back-end Read Transfer Size** to greater than 6KiB/op. To track the overall transfer rate, define **Total Back-end I/O Rate** by setting the alert to greater than 3,000 ops/s.

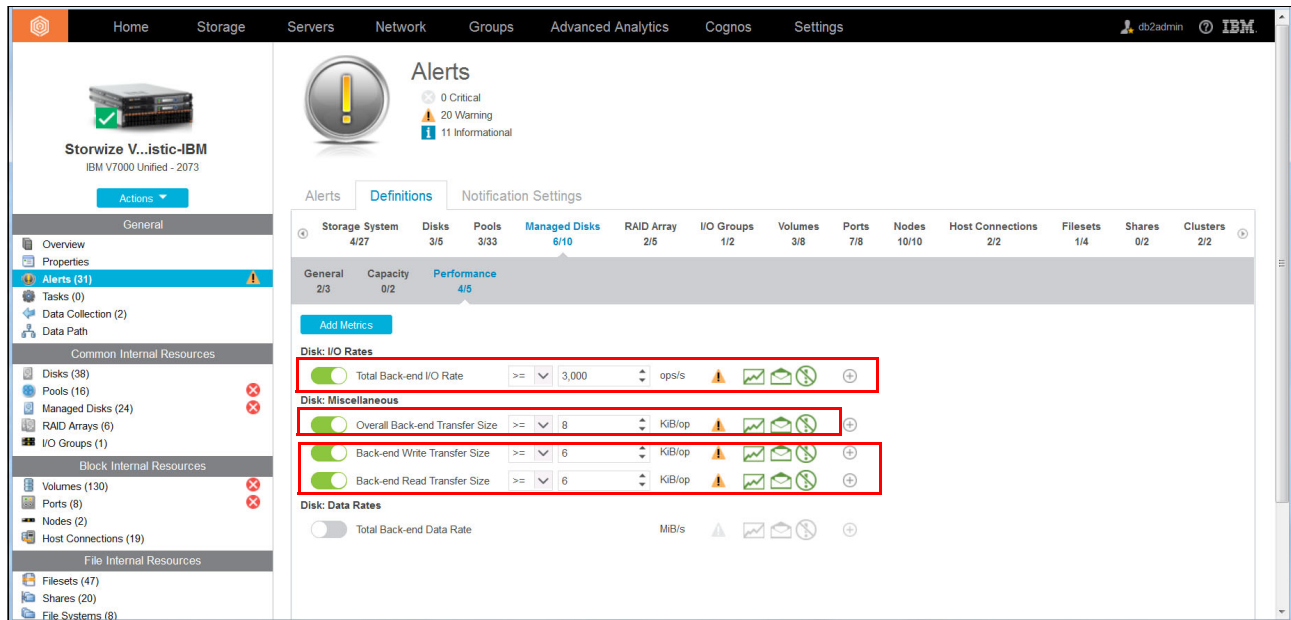


Figure 6-16 Example of setting performance thresholds for MDisks

- To ensure you that you can quickly engage support if Performance Monitors fail, you can define attribute alerts on the storage system, as shown Figure 6-17, which shows how to enable an alert on Performance Monitoring, by clicking the **General** tab. Enable the **Performance Monitor Status** alert setting the attribute to **is not normal**.

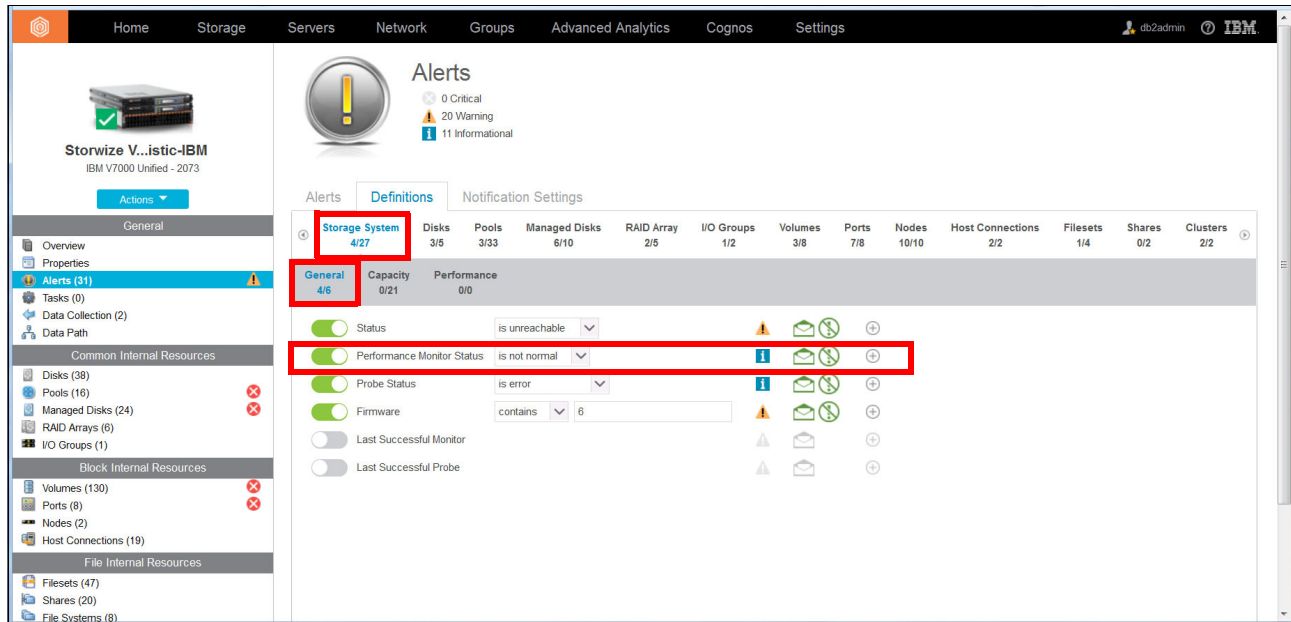


Figure 6-17 Enable an alert on Performance Monitoring

- Then, from the Notification Settings window, which is shown in Figure 6-6 on page 185, set **Run Script** to upload the Spectrum Control service.bat/sh script to run upon failure. Click the envelope icon, click **Run Script**, click **Select File**, and select **Upload Script**, as shown in Figure 6-18, which shows how to run a script upon alert notification. Click **Browse** and select **Script**. Define the storage resource agent by using the drop-down arrow next to Run script on Storage Resource agent and select the Storage Resource Agent. Click **Done**.

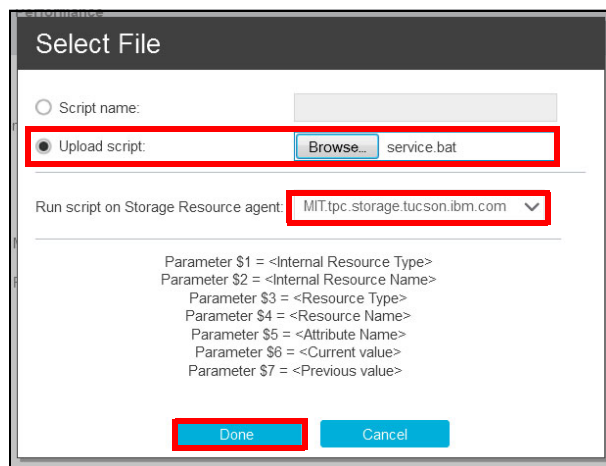


Figure 6-18 Run a script upon alert notification

Reference: For more information about running scripts as a triggered action for an alert condition, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1N6PFY4>



Managing IBM Spectrum Scale with IBM Spectrum Control

IBM Spectrum Scale (Spectrum Scale) is a software-defined, scalable, high-performance data and file management solution (formerly IBM General Parallel File System (IBM GPFS™)).

Note: IBM Spectrum Scale Version 4.1 is the first version of Spectrum Scale and is based on GPFS. For versions before Version 4.1, the name GPFS is maintained.

Spectrum Scale integrated with IBM Spectrum Control (Spectrum Control) delivers scalable capacity and performance to handle the demanding data analytics. Storage administrators can combine flash, disk, cloud, and tape storage into a unified system that is high-performing and has a lower cost than traditional approaches. Integrating Spectrum Control with Spectrum Scale with business enterprises provides intelligent analytics for managing storage and the capability to monitor multiple Spectrum Scale installations (and other storage).

Spectrum Scale is a storage infrastructure that can adapt to any demanding applications across the enterprise. By including Spectrum Scale in the software-defined infrastructure (SDI), organizations can streamline data workflows, help improve service, reduce costs, manage risk, and deliver business results while positioning the enterprise for future growth.

This chapter describes the Spectrum Scale enhancements and practical troubleshooting scenarios of Spectrum Scale by using the Spectrum Control web-based GUI from a system and storage administrator perspective. In addition, this chapter describes the benefits that Spectrum Control offers beyond the management capabilities of the Spectrum Scale CLI and GUI.

7.1 Historical overview of Spectrum Scale support in Spectrum Control

Starting with Spectrum Control Version 5.2.3 (formerly known as IBM Tivoli Storage Productivity Center), when initial support was added, it gradually improved with each release. This section describes what is improved.

Enhancements in Version 5.2.3

This is the first version when support was added for GPFS. You can use the web-based GUI to add GPFS clusters and GPFS Storage Server (GSS) / Elastic Storage Server (ESS) systems for monitoring. These systems can be added as file storage systems (similar to Storwize V7000 Unified or Scale Out Network Attached Storage). The initial data collection and alerting function was enabled in this release.

The Spectrum Scale cluster can contain one or more servers, or a combination of servers and ESS systems.

The initial firmware level of Spectrum Scale supported in Spectrum Control was V4.1 and V4.1 TL1 (technical level 1 fix pack), where the cluster nodes can run on Linux Red Hat Enterprise 6.

Enhancements in Version 5.2.4

In this release, the following enhancements were made available:

- ▶ Support for all Linux platforms for Spectrum Scale nodes
- ▶ Spectrum Scale node correlation to Spectrum Control managed servers (Storage Resource Agent and agentless)
- ▶ Spectrum Scale NSD correlation to SAN storage volumes
- ▶ Status updates for Spectrum Scale nodes and NSDs
- ▶ Spectrum Scale snapshots for file systems and file sets
- ▶ Files system inode overview chart
- ▶ Capacity bars showing total/user/free space, including Spectrum Scale snapshots used space
- ▶ Alert configuration and viewing alert logs

Enhancements in Version 5.2.5

In this release, a series of enhancements in terms of usability of the Spectrum Scale and ESS were made available:

- ▶ Display for the largest and fullest pools in the overview window of a cluster
- ▶ Spectrum Scale volume mapping display enhancements
- ▶ Correlation of Spectrum Scale clusters with related fabrics and switches
- ▶ Show of file systems mounted per Spectrum Scale node
- ▶ Show quotas from Spectrum Scale file systems
- ▶ Server details window for Spectrum Scale NSDs
- ▶ Show the metadata and data space breakdown per Spectrum Scale pools
- ▶ Add inodes details
- ▶ Display file set comments

Enhancements in version 5.2.6

In this release, support for Spectrum Scale V4.1.1 was introduced with some further usability enhancements:

- ▶ New charts in the overview window showing file sets closest to running out of free inodes:
 - One view showing file sets with the least number of free inodes
 - One view showing file sets with the most number of total inodes (max inodes)
- ▶ Adding a Spectrum Scale cluster to Spectrum Control by using a non-root user.
- ▶ Spectrum Scale probe is extended to collect information about remote mounted file systems (three new columns: owning cluster, owning storage subsystem and creation name).
- ▶ Update of GPFS/GSS to Spectrum Scale (rebranding).
- ▶ AIX cluster node support.

Enhancements in Version 5.2.7

In this release, further usability enhancements were introduced:

- ▶ With Spectrum Scale V4.1.1, the node protocols information can be reviewed on the Nodes chart. A new column that is named Protocols was added, which contains comma-separated values of the enabled protocols, for example, NFS, Object, and SMB.
- ▶ Status failure is enhanced to show whether the cluster export service is unavailable on any given node, and the reason of the failure (suspended, network down, or failed).

Enhancements in Version 5.2.8

In this release, a series of high-value enhancements were introduced:

- ▶ Monitoring the performance at the cluster, node and file system level (with Spectrum Scale Version 4.x1):
 - Analyze the performance of multiple clusters with the performance troubleshooting charts to determine which clusters have the heaviest I/O loads, highest data rates, and longest response times.
 - View which node that mounts various file systems have the heaviest I/O loads and the highest CPU usage or memory usage.
 - View which file system on a cluster has the heaviest I/O loads and longest response time.
- ▶ Monitoring object storage systems:
 - View which GPFS file systems are associated with an object storage system, end to end troubleshooting by viewing the relationship between Spectrum Scale and the SAN-attached storage, and determining which object container is stored and on a file system and correlate it with back-end storage.
 - View the capacity and space usage of specific object accounts and containers.
 - Monitor object count and space quotas for containers and plan for capacity.

7.2 Troubleshooting scenarios

This section describes practical troubleshooting scenarios of Spectrum Scale by using the Spectrum Control web-based GUI. These general scenarios provide resolutions to a set of questions that an administrator of a Spectrum Scale typically must handle:

- ▶ Which clusters or nodes have health problems?
- ▶ Which clusters are running out of free space?
- ▶ Which file systems or pools are running out of capacity?
- ▶ Which file systems are mounted on which nodes?
- ▶ How much space is occupied by snapshots, and how obsolete are some of the snapshots?
- ▶ Which quotas are close to being exceeded or already are exceeded?
- ▶ Which file sets are close to running out of free inodes?
- ▶ Are all the volumes backing network-attached storage (NSDs) performing correctly?
- ▶ Which NSDs are at risk of becoming unavailable, or are unavailable?
- ▶ Is the performance of the file systems and cluster nodes satisfactory?
- ▶ Are all nodes that are fulfilling critical roles in the cluster up and running?
- ▶ Which object accounts are managing the most container and violating quotas?
- ▶ Which object container is violating the space quota?
- ▶ Send a notification when a node goes offline or file systems fill up beyond a threshold.

7.2.1 Adding Spectrum Scale files and objects to Spectrum Control

With Spectrum Control V5.2.8, there are many user interface changes. The dashboard is reorganized around the concept of *storage providers* and *storage consumers (applications and departments)*. There also is a performance chart for *Top Block Storage Performance*, which is introduced (Figure 7-1).

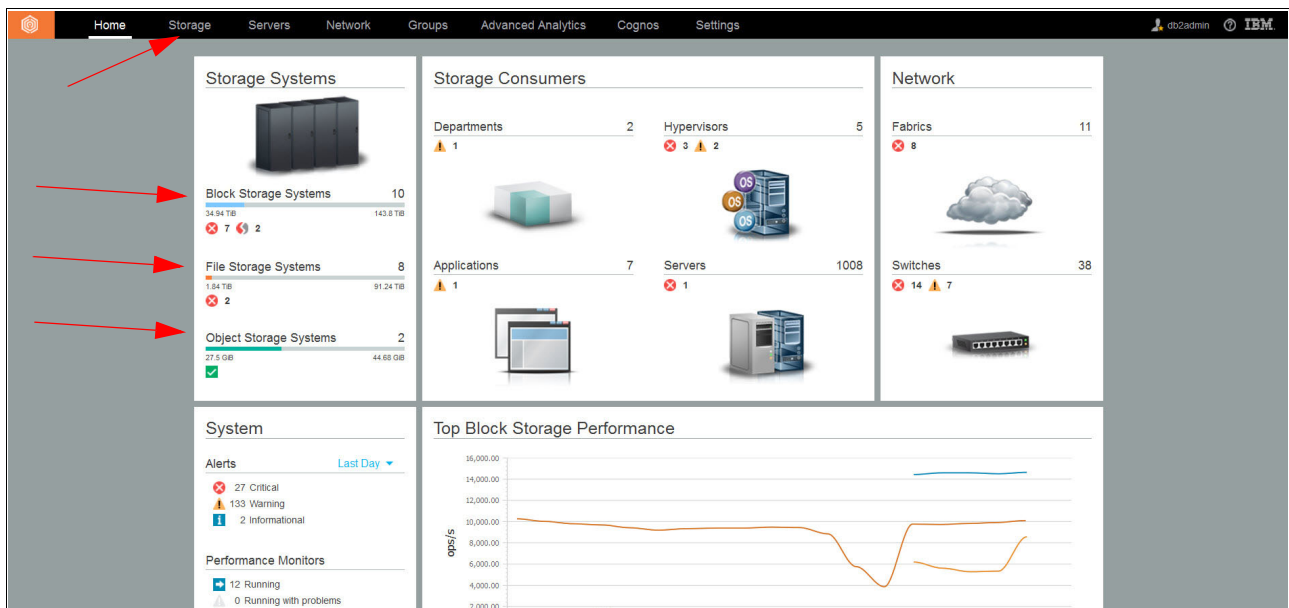


Figure 7-1 Spectrum Control dashboard

The storage providers are split into three parts, mainly because of the addition of object storage support for Spectrum Scale to the existing block and file storage. The storage providers (Storage systems) section now contains the following systems:

- ▶ Block storage systems
- ▶ File storage systems
- ▶ Object storage system

The same sections can be found under the menu that is shown in Figure 7-2.

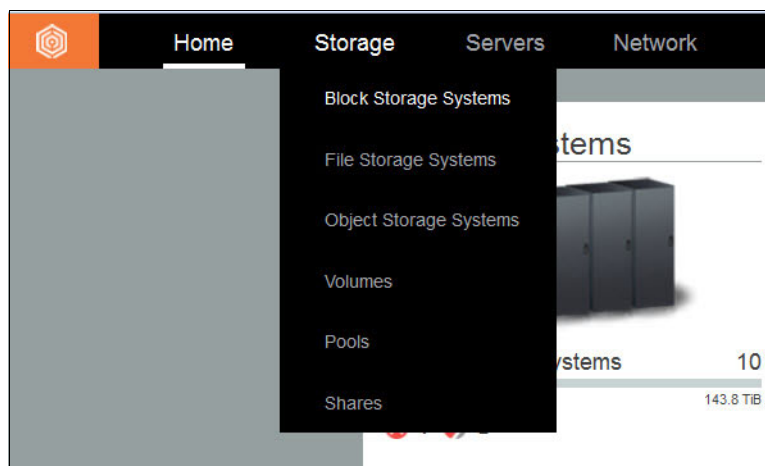


Figure 7-2 Block/File/Object storage systems

To have Spectrum Scale managed by Spectrum Control, click **File**, click the **Object Storage System** direct link, or use the top navigation menu.

If you want to add the Spectrum Scale file personality only, click **File Storage Systems**, click **+ Add Storage System**, and select **Spectrum Scale**. Other options are Scale Out Network Attached Storage, Storwize V7000 Unified, and NetApp (Figure 7-3).

Note: Within the File Personality Add Device wizard, it is possible to add the object-related credentials for Spectrum Scale.

Figure 7-3 shows an example of adding the Spectrum Scale file personality.

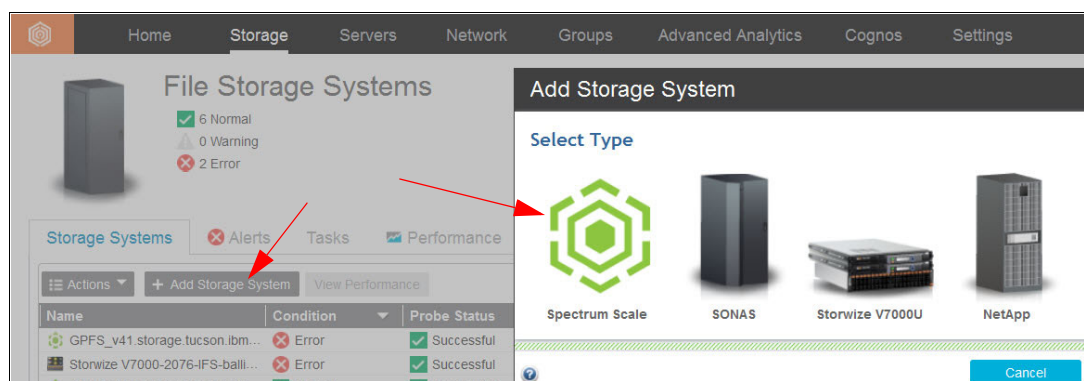


Figure 7-3 Adding Spectrum Scale file personality

If you want to add Spectrum Scale Object personality, click **Object Storage Systems**, click **+ Add Storage System**. The window where you enter the IP and Hostname / User Name / Password opens, and the **Specify authentication credentials for object storage** option is already selected. The Object personality requires an extra set of credentials, which are related to *Keystone*, which is the OpenStack identity service. For more information about the Keystone role, see the following website:

<http://ibm.co/1l0nK5t>

For examples about configuring OpenStack access to monitor the object storage system, and about user requirements, see the following website:

<http://ibm.co/1UJnxjg>

For Object Storage Systems, Spectrum Scale is the only type that is supported by Spectrum Control (Figure 7-4).

Note: It is not possible to add the Spectrum Scale as an Object storage system only without specifying the file side credentials.

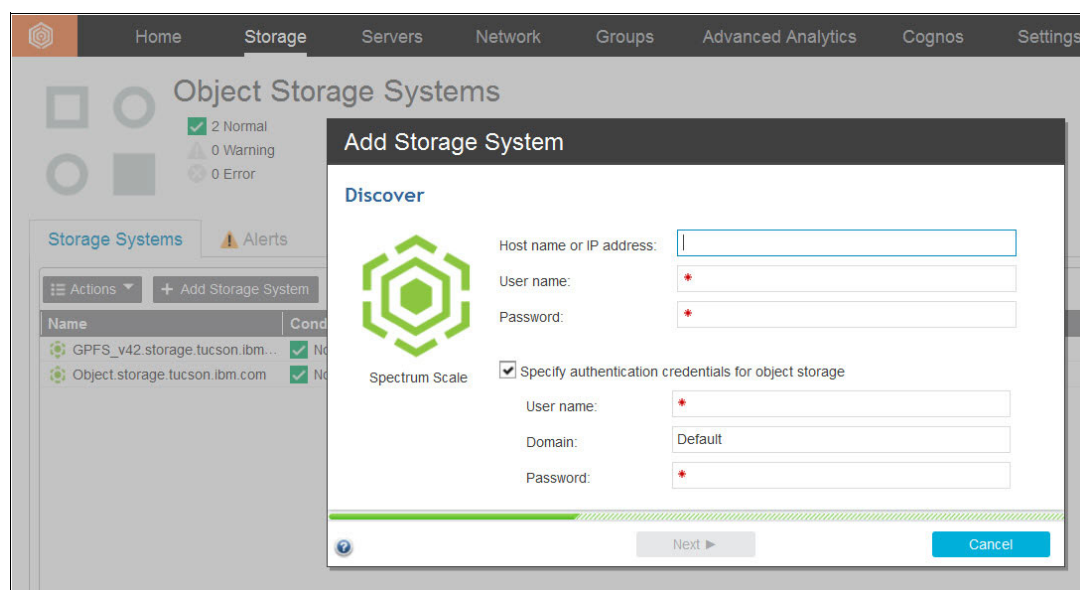


Figure 7-4 Add Spectrum Scale object personality

This section does not cover how to add or collect asset, status, and performance data of Spectrum Scale because these tasks are similar to the other storage device types. One of the differences with Spectrum Scale is when it is added without root privileges. For more information about adding Spectrum Scale without using root privileges, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1Pn0o0p>

Performance monitoring is available only on Spectrum Scale if the “Performance Monitoring tool” (also called ZiMon) is installed, configured, and running on the Spectrum Scale cluster, as of Spectrum Control V5.2.8. More details are available in the IBM Knowledge Center, found at:

<http://ibm.co/1l0o1VT>

Note: Performance monitoring of Spectrum Scale is supported only in Spectrum Control, starting with Version 4.1.1.

7.2.2 Usage scenarios for troubleshooting Spectrum Scale

This section adds details to the list of previously asked questions in 7.2, “Troubleshooting scenarios” on page 198. It showcases the advantages of using the Spectrum Control web-based GUI as an alternative to using the Spectrum Control command-line interface (CLI) or GUI, which are available for Spectrum Scale Version 4.2.

A storage administrator often must quickly handle a reported problem that is reported by Spectrum Control alerting (see Chapter 6, “Advanced alert management” on page 177) or by an application owner. This problem consumes Spectrum Scale resources, so determining the root cause of a problem within a short period is paramount in today’s dynamic environments.

One of the first advantages of using Spectrum Control is that a storage administrator can monitor many Spectrum Scale clusters and, when logging in to the dashboard, make an immediate assessment of the managed clusters (Figure 7-5 on page 202). The following sections describe items that the administrator should investigate.

Which clusters or nodes have health problem

Starting with Spectrum Control V5.2.8, the way the status is reported changed. The *condition* metric was introduced, which represents the propagated status, and the *status* metric represents the device status. For more information, see the *Identifying problems on resources* topic in the IBM Knowledge Center at the following website:

<http://ibm.co/1SJHV1H>

From the Storage System pane of the dashboard, you can easily tell that there are two file storage systems (out of eight) where the condition is “Error” (see Figure 7-5). Clicking the **File Storage System** link opens the details window of the file storage systems. One of the clusters is the one in the “Error” condition (Figure 7-6).

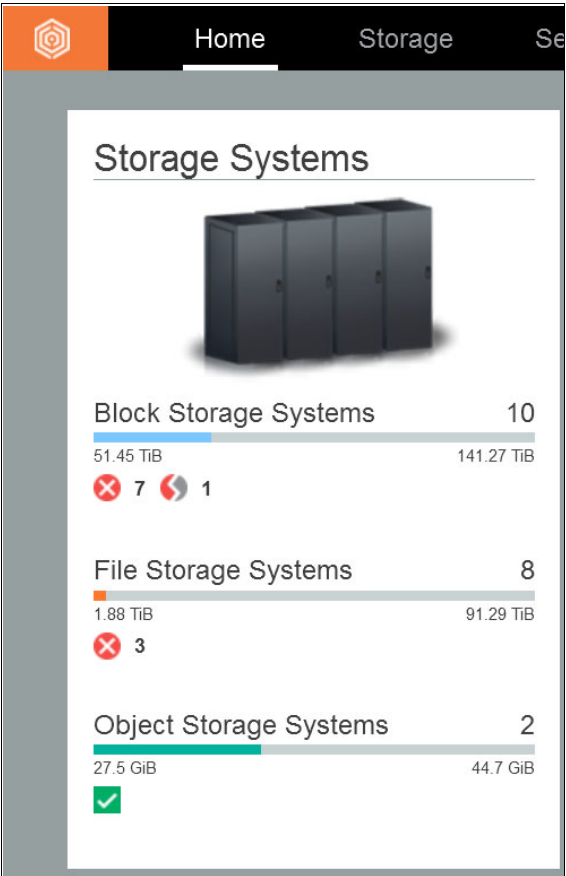


Figure 7-5 Spectrum Control dashboard pane of Storage Systems

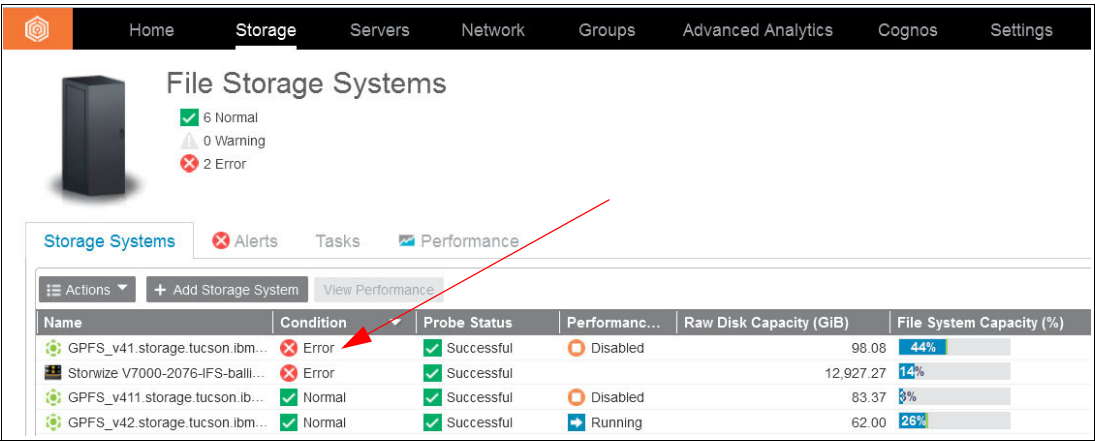


Figure 7-6 Spectrum Scale cluster condition

This condition is a propagated value (which is the same case for the dashboard). So, there is a possibility that one of the internal components of the cluster has an issue. From the File Storage Systems window, you can drill into the GPFS_v41.storage.tucson.ibm.com cluster by either double-clicking the item from the table or by right-clicking and selecting **View Details** from the drop-down menu (Figure 7-7).

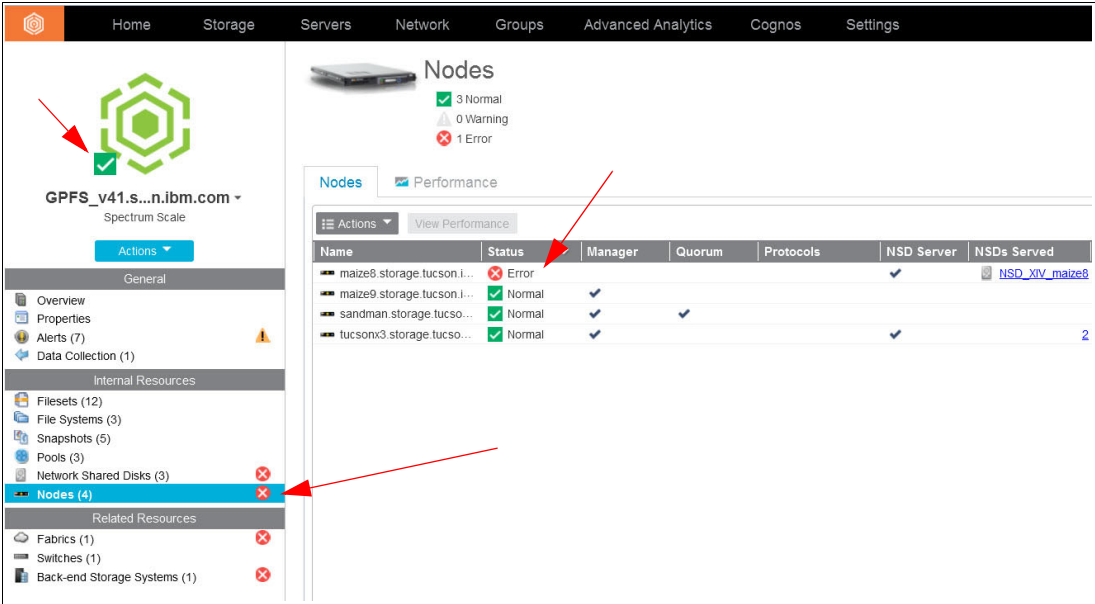


Figure 7-7 Spectrum Scale cluster nodes overlay window

As shown in Figure 7-7, the detail window of GPFS_v41.storage.tucson.ibm.com opens, with the Nodes overlay selected. In the left pane, you see a red struck circle, indicating that one of the Nodes statuses is in “Error”, and as seen in the right side table, maize8.storage.tucson.ibm.com node status is “Error”. This status shows the device status of maize8.storage.tucson.ibm.com node, which must be investigated further by logging in to one of the Spectrum Scale manager nodes.

The device status of the cluster itself is healthy, which is reported from the cluster, and despite one of the nodes being in “Error” status, the cluster continues to function because the node is not a critical node (Manager or Quorum).

Example 7-1 shows the cluster node error status.

Example 7-1 Investigate cluster node error status

```
[root@maize9 ~]# mm1scluster
```

GPFS cluster information

=====

```
GPFS cluster name:      GPFS_v41.storage.tucson.ibm.com
GPFS cluster id:        9485674298140078727
GPFS UID domain:        GPFS_v41.storage.tucson.ibm.com
Remote shell command:   /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:        CCR
```

GPFS cluster configuration servers:

```
Primary server:  sandman.storage.tucson.ibm.com (not in use)
Secondary server: (none)
```

Node Designation	Daemon node name	IP address	Admin node name
8	sandman.storage.tucson.ibm.com	9.11.97.44	sandman.storage.tucson.ibm.com
quorum-manager			
15	maize8.storage.tucson.ibm.com	9.11.92.11	maize8.storage.tucson.ibm.com
20	tucsonx3.storage.tucson.ibm.com	9.11.97.51	tucsonx3.storage.tucson.ibm.com
21	maize9.storage.tucson.ibm.com	9.11.92.75	maize9.storage.tucson.ibm.com
			manager
			manager

```
[root@maize9 ~]# ssh root@maize8.storage.tucson.ibm.com
```

```
ssh: connect to host maize8.storage.tucson.ibm.com port 22: Connection refused
```

As shown in Example 7-1, the **mm1scluster** command ran correctly, but **maize8.storage.tucson.ibm.com** is not reachable by SSH. This is the likely cause of why the node has an “Error” status.

Note: For the scenario in Example 7-1 to succeed, a connection to SSH is necessary for communication.

Now, look again at the Spectrum Scale GUI to see what the health of the cluster or node is.

Note: The Spectrum Scale GUI is only available starting with Version 4.2.

For more information about installing the Spectrum Scale GUI, see the IBM Knowledge Center, found at:

<http://ibm.co/1VZTPH0>

Figure 7-8 on page 205 shows the Spectrum Control V4.2 GUI dashboard.

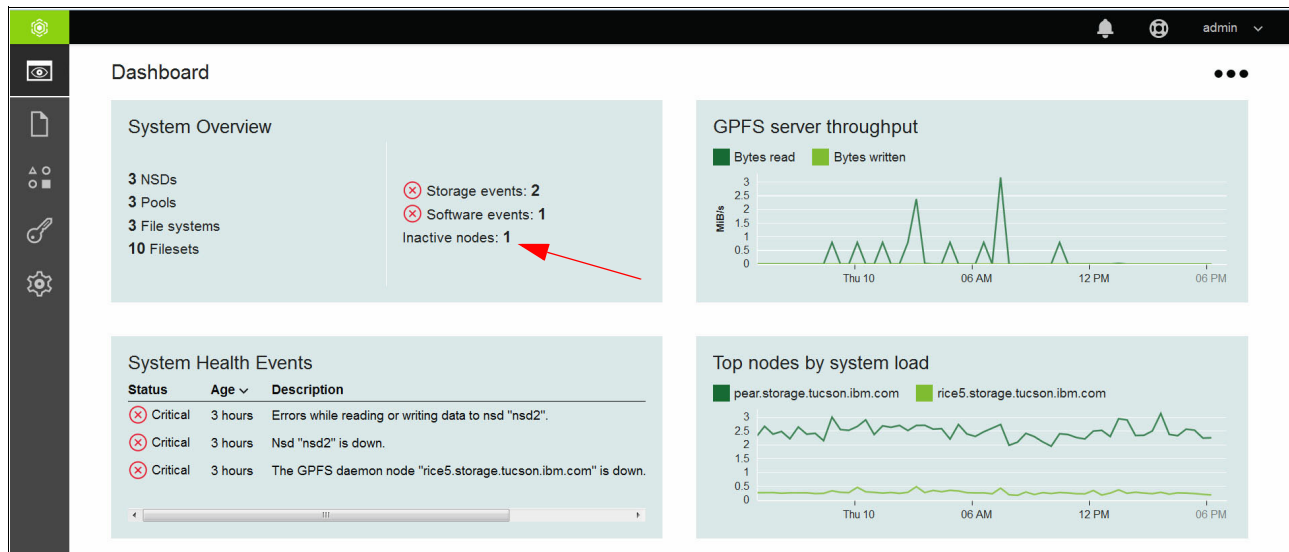


Figure 7-8 Spectrum Scale V4.2 GUI dashboard

In the Spectrum Scale GUI interface dashboard section, you see that one of the nodes of the cluster is inactive (Figure 7-8 (Spectrum Control V4.2)). There is no dedicated pane for nodes as there is in Spectrum Control (Figure 7-7 on page 203 (Spectrum Control V4.1)) in the Events window. However, you can see the Critical error about the GPFS daemon on node "rice5.storage.tucson.ibm.com" being down. On this same window, you can run a fix procedure to start the daemon (Figure 7-9 (Spectrum Control V4.2)).

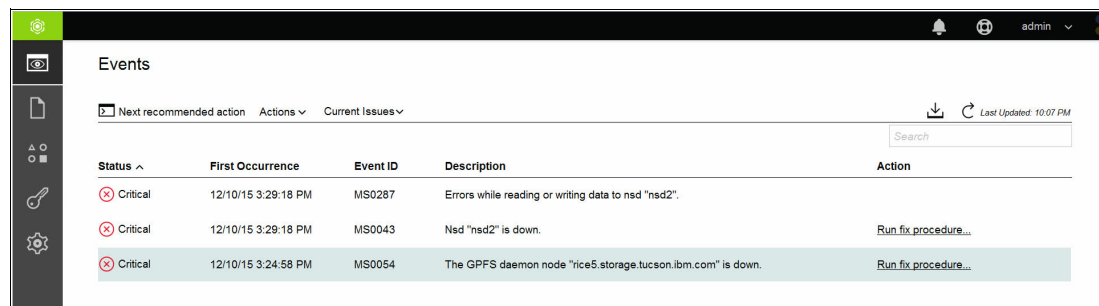


Figure 7-9 Spectrum Scale GUI - Events window with Run fix procedure

One of fundamental differences between the two GUIs is that you can use Spectrum Scale to change the configuration/run actions on the cluster, and Spectrum Control shows the cluster entities in a consistent manner with other storage devices that Spectrum Control manages.

Which clusters are running out of free space

Spectrum Control can manage many Spectrum Scale clusters. By looking at the File Storage System window, a storage administrator can immediately determine which cluster is running out of free space, as shown in Figure 7-10.

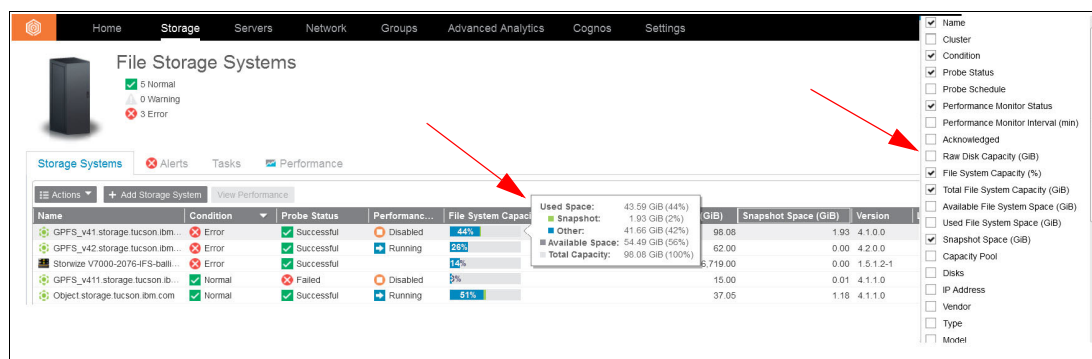


Figure 7-10 Spectrum Scale file system capacity

By using the File System Capacity stacked bar, you can easily grasp the usage level. Hovering your cursor over the bar shows a breakdown of the Used Space into Available Space and Total Capacity. This same information is available throughout the various columns of data from this table, which can be enabled by right-clicking the table header.

Which file systems or pools are running out of capacity

The File Storage System level File System Capacity stacked bar is an aggregated value of the File System Capacity(%) from the details window, which is shown in Figure 7-10, but by looking into the details of each File System or Storage Pool, you can determine which file system or pool is running out of capacity.

Figure 7-11 shows Spectrum Scale File Systems details window.

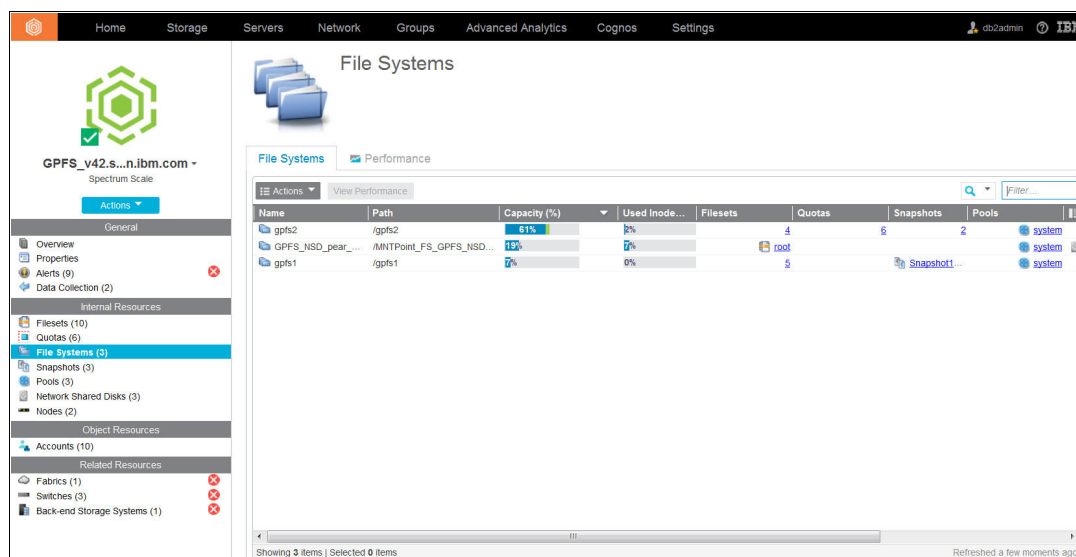


Figure 7-11 Spectrum Scale file systems capacities

Figure 7-11 shows that one of the file systems (gpfs2) is at 61% capacity usage, and this is likely the file system that will run out of capacity and needs attention.

Although there are no historical capacity charts for file storage system-based storage devices in Spectrum Control V5.2.8, you can access this information by opening the Overview window and reviewing the Total File System Space usage (Figure 7-12).

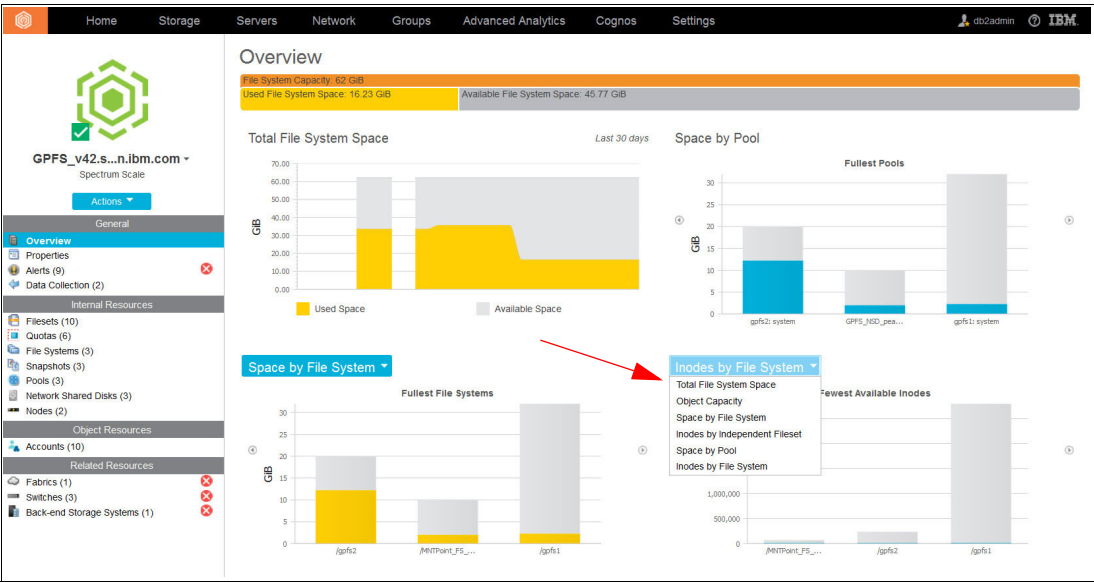


Figure 7-12 Spectrum Scale Overview window

The Overview window is customizable (this is a common function in the web-based GUI) and different pods with valuable information can be selected (Figure 7-12).

Switching to the Pools window (Figure 7-13), you see that there is a *system* pool in which all the file systems are placed.

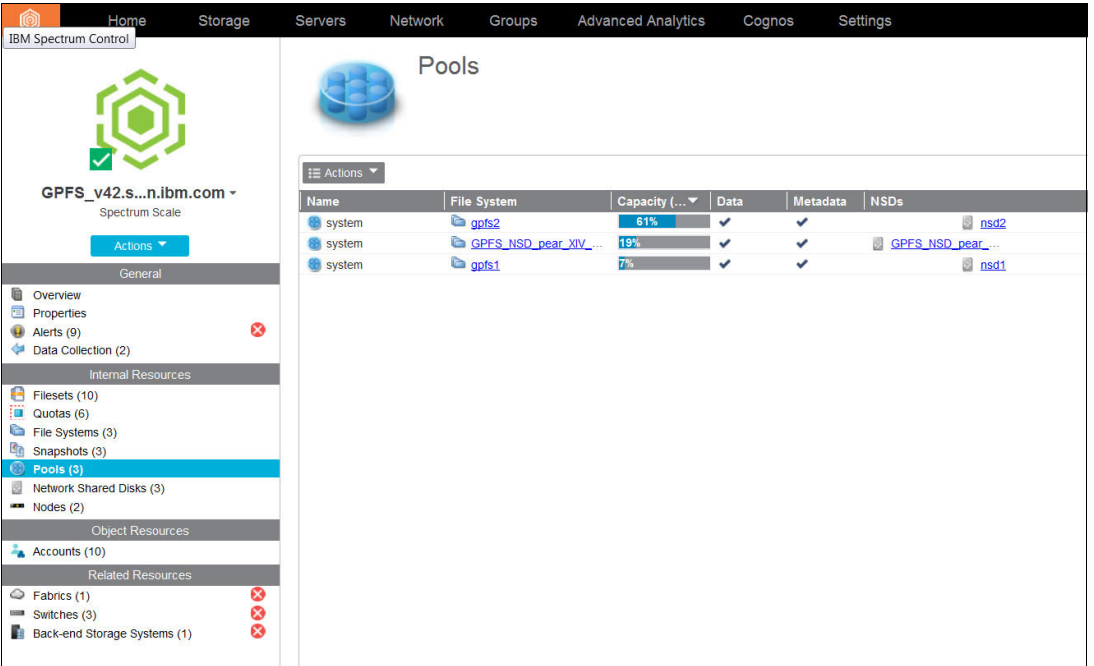


Figure 7-13 Spectrum Scale Pools window

Spectrum Scale has two types of storage pools:

- ▶ Internal storage pools
- ▶ External storage pools

The internal storage pools are the pools that are managed by Spectrum Scale, and the external storage pools are managed by external applications such as IBM Spectrum Protect (formerly Tivoli Storage Manager).

Note: Spectrum Control supports only internal storage pools.

Using the Spectrum Scale CLI, you can view the same pools information by running the `mm1spool` command (Example 7-2).

Example 7-2 Spectrum Scale listing of file system and storage pool

```
[root@rice5 ~]# mm1spool gpfs2
Storage pools in file system at '/gpfs2':
Name      Id  BlkSize Data Meta Total Data in (KB)  Free Data in (KB)  Total Meta in (KB)  Free Meta in (KB)
system    0   256 KB yes  yes  20971520      8257792 ( 39%)    20971520      8270848 ( 39%)
```

However, it is not possible to list information about multiple storage pools at the same time; you can list information only for a specific file system.

The Spectrum Scale web-based GUI is organized slightly differently, that is, around capacity values. When you click **Monitoring** → **Capacity**, you can see the various entities (File System, Pools, Filesets, Use, and so on) to which capacity values can apply (Figure 7-14), although there is no dedicated window as in Spectrum Control for Pools, where some of the related details can be seen from the window itself, whether the pools contains Data, Metadata, and which NSD is being mounted.

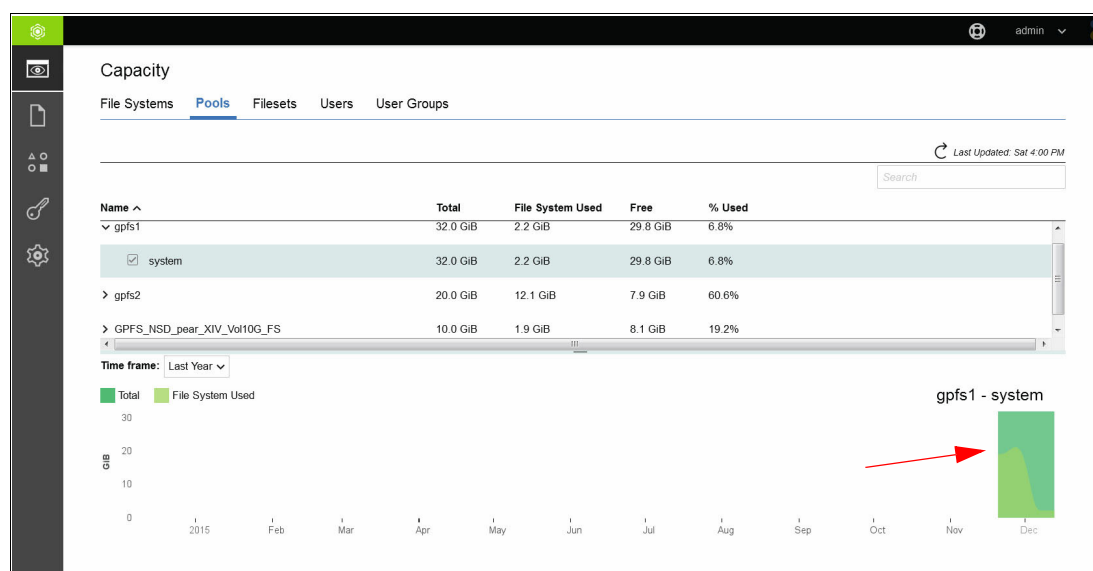


Figure 7-14 Spectrum Scale capacity values for Pools

The Spectrum Scale GUI has the advantage of historical charting on the File System/Pools level, as shown in Figure 7-14. Historical charting is available in Spectrum Control V5.2.8 only for block-based devices.

Which file systems are mounted on which nodes

In “Which file systems or pools are running out of capacity” on page 206, you saw that the Spectrum Control web-based GUI has the overall advantage of showing related resources. For example, a storage administrator might be interested in which file systems are mounted on which nodes. You can use the Spectrum Control Nodes window to answer this question (Figure 7-15).

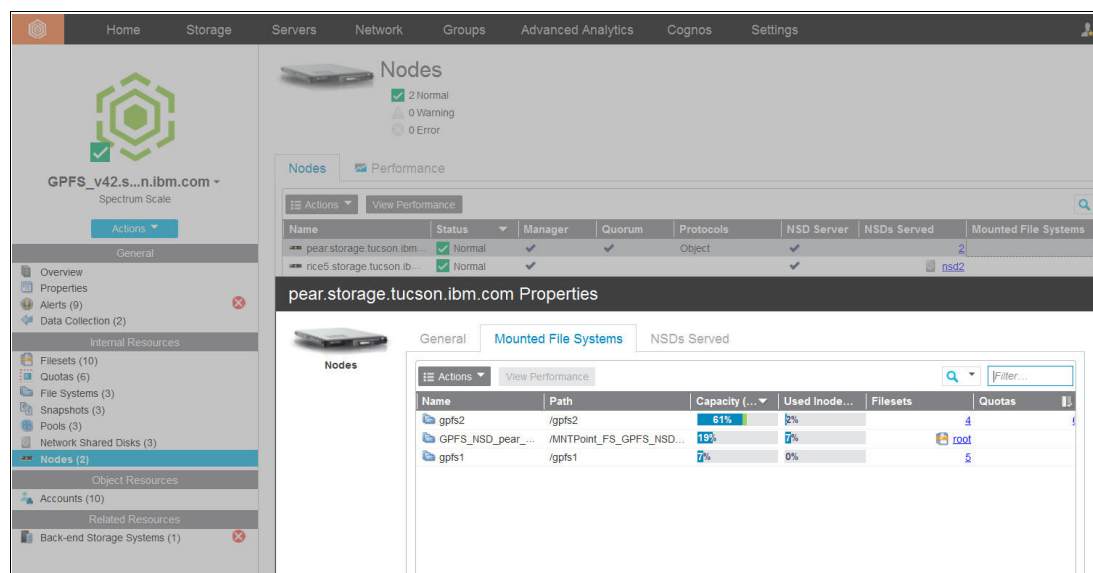


Figure 7-15 Spectrum Control Nodes - mounted file systems

Using the Spectrum Control CLI or operating system commands, you can determine the file systems that are mounted per node. You can also run the `df` command, as shown in Example 7-3.

Example 7-3 Spectrum Scale file systems mounted by node

```
[root@rice5 ~]# df -h
```

File system	Size	Used	Avail	Use%	Mounted on
/dev/mapper/volGroup--root	67G	17G	51G	24%	/
devtmpfs	7.7G	0	7.7G	0%	/dev
tmpfs	7.8G	64K	7.8G	1%	/dev/shm
tmpfs	7.8G	7.7G	27M	100%	/run
tmpfs	7.8G	0	7.8G	0%	/sys/fs/cgroup
/dev/sda2	497M	279M	219M	57%	/boot
/dev/GPFS_NSD_pear_XIV_Vol10G_FS	10G	2.0G	8.1G	20%	
/MNTPoint_FS_GPFS_NSD_pear_XIV_Vol10G					
/dev/gpfs1	32G	2.2G	30G	7%	/gpfs1
/dev/gpfs2	20G	13G	7.9G	61%	/gpfs2

If you use the Spectrum Scale `mmdf` or `mmdsh` commands, you must set the working collective (WCOLL) environment variable. With this command, you can see more details about a specific file system mount, as shown in Example 7-4.

Example 7-4 Spectrum Scale `mmdf` command

[root@rice5 ~]# mmdf gpfs1 --block-size=auto						
disk name	disk size	failure group	holds metadata	holds data	free in full blocks	free in fragments

Disks in storage pool: system (Maximum disk size allowed is 286 GB)						
nsd1	32G	1	Yes	Yes	29.81G (93%)	5.234M (0%)

(pool total)	32G				29.81G (93%)	5.234M (0%)
=====						
(total)	32G				29.81G (93%)	5.234M (0%)
=====						
Inode Information						

Total number of used inodes in all Inode spaces:				5863		
Total number of free inodes in all Inode spaces:				242905		
Total number of allocated inodes in all Inode spaces:				248768		
Total of Maximum number of inodes in all Inode spaces:				2796224		

The Spectrum Scale GUI cannot show which file system is mounted per node.

How much space is occupied by snapshots and how obsolete are some of the snapshots

Spectrum Scale can create snapshots at the file system, file set, and file level to preserve the content at a single point. The snapshots at the entire file system level are known as *global snapshots*. These snapshots are read-only and can cause a significant impact in terms of storage space.

Some of the questions a storage administrator might have is “How much space is occupied by snapshots?” and “How obsolete are some of the snapshots?”. Spectrum Control has several windows that can answer these questions.

The first window that can help you answer these questions is the overall File Storage System window, which is found under the **File System Capacity (%)** stacked bar chart, as shown in Figure 7-16 on page 211.

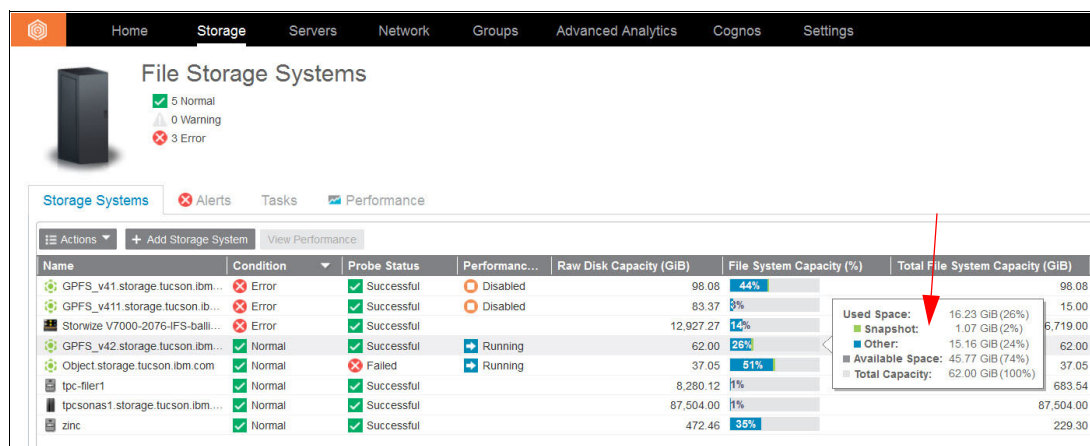


Figure 7-16 Spectrum Control file storage system level snapshot used space

This is an aggregated value. For more information about which file systems have snapshots, the storage administrator must go to the details of the managed Spectrum Scale server, as shown in Figure 7-11 on page 206.

By viewing the details window of the managed Spectrum Scale snapshots view, the storage administrator can tell how obsolete the snapshots are either by looking at the snapshots' name (which contain @date as trailing characters) or by enabling the "CREATION_TIME" column, as shown in Figure 7-17.

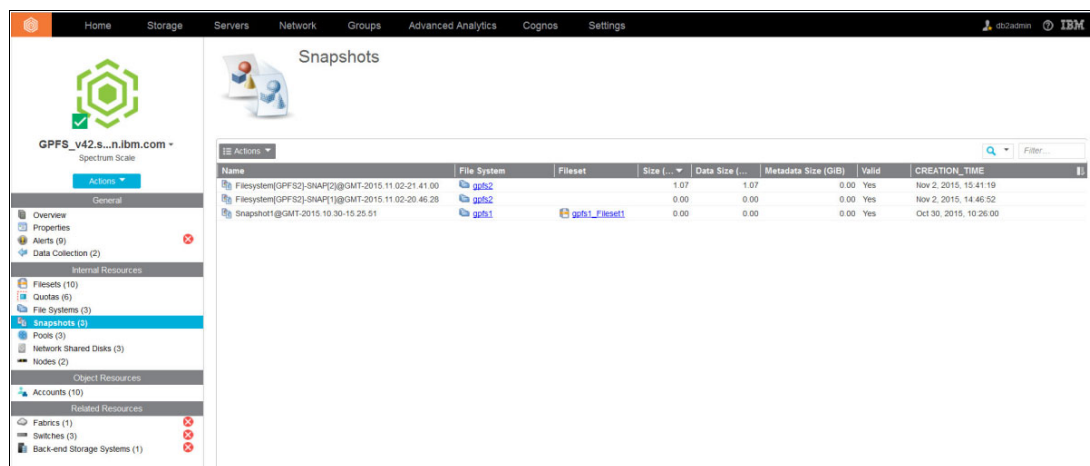
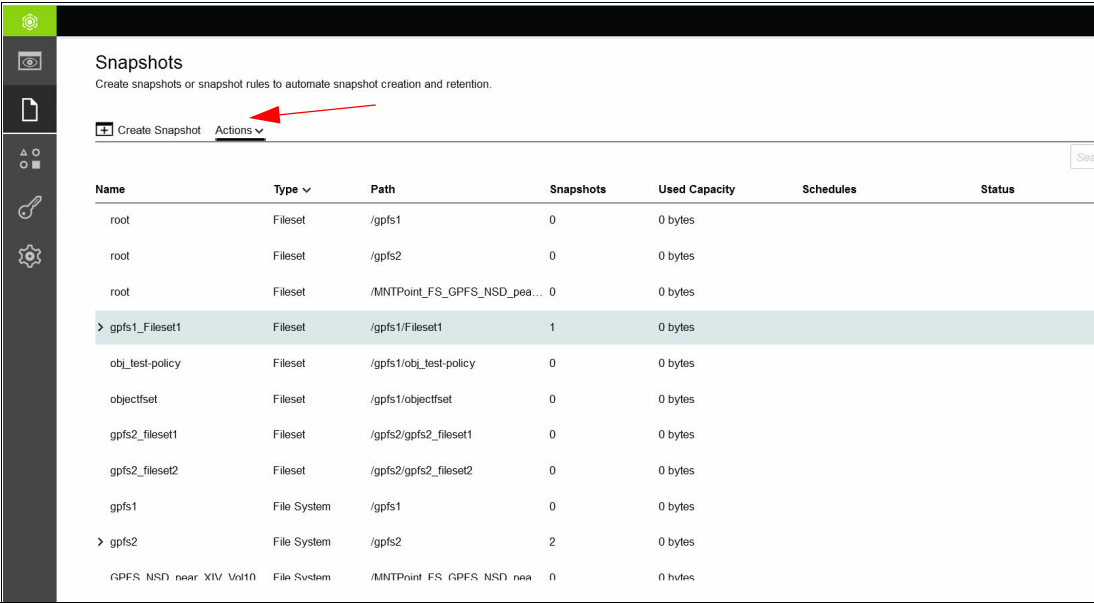


Figure 7-17 Spectrum Control showing snapshots of Spectrum Scale

In contrast to Spectrum Control, the Spectrum Scale GUI has more capabilities and shows a richer set of information about snapshots. The snapshots table can be accessed by clicking **Files** → **Snapshots**, as shown in Figure 7-18.



Snapshots
Create snapshots or snapshot rules to automate snapshot creation and retention.

Create Snapshot Actions

Name	Type	Path	Snapshots	Used Capacity	Schedules	Status
root	Fileset	/gpfs1	0	0 bytes		
root	Fileset	/gpfs2	0	0 bytes		
root	Fileset	/MNTPoint_FS_GPFS_NSD_pea...	0	0 bytes		
> gpfs1_Fileset1	Fileset	/gpfs1/Fileset1	1	0 bytes		
obj_test-policy	Fileset	/gpfs1/obj_test-policy	0	0 bytes		
objectfset	Fileset	/gpfs1/objectfset	0	0 bytes		
gpfs2_fileset1	Fileset	/gpfs2/gpfs2_fileset1	0	0 bytes		
gpfs2_fileset2	Fileset	/gpfs2/gpfs2_fileset2	0	0 bytes		
gpfs1	File System	/gpfs1	0	0 bytes		
> gpfs2	File System	/gpfs2	2	0 bytes		
GPFS_NSD near XIV Vol10	File System	/MNTPoint_FS_GPFS_NSD nea	0	0 bytes		

Figure 7-18 Spectrum Scale web-based GUI snapshots table

Here, you can do snapshots at the file system, file set, and file level. Spectrum Control can show only file system and file set level snapshots. In terms of actions, the Spectrum Scale GUI can create snapshots and edit existing ones (Figure 7-18).

You can use the Spectrum Scale CLI to list details about one snapshot at a time, as shown in Example 7-5.

Example 7-5 Spectrum Scale mmls snapshot

```
[root@rice5 ~]# mmlssnapshot gpfs1
Snapshots in file system gpfs1:
Directory          SnapId   Status  Created                Fileset
Snapshot1@GMT-2015.10.30-15.25.51 1      Valid   Fri Oct 30 11:26:00 2015  gpfs1_Fileset1
```

- Which quotas are close to being exceeded or already are exceeded**
- Spectrum Scale allows definitions of quotas that can be applied to the following components:
- ▶ Individual users
 - ▶ Groups of users
 - ▶ Individual file sets

The quotas can be enabled either at the creation time of the file system or afterward by using the **mmchfs** command. You can enable quotas by using the **mmquotaon** command, display information about a quota by using the **mmlsquota** command, and edit quotas by using the **mmedquota** command with your default editor. You can create reports on quotas by using the **mmrepquota** command, as shown in Example 7-6 on page 213.

Example 7-6 Spectrum Scale mmrepquota command

```
[root@rice5 ~]# mmrepquota -u -v gpfs2
*** Report for USR quotas on gpfs2
```

Name	type	Block Limits					File Limits					entryType
		KB	quota	limit	in_doubt	grace	files	quota	limit	in_doubt	grace	
root	USR	4477696	0	0	0	none	8	0	0	0	none	default on
testuser	USR	3004160	4194304	8388608	0	none	3	0	0	0	none	e
testuser2	USR	3004032	10485760	15728640	0	none	3	0	0	0	none	d_fsys

Note: The `mmsetquota` command can be used to set explicit quotas, for example, for specific users/groups.

For more information and examples, see the Spectrum Scale IBM Knowledge Center, found at:

<http://ibm.co/1Qedh1j>

For more information about quotas, see *IBM Spectrum Scale (formerly GPFS)*, SG24-8254.

Spectrum Control also can show the quotas of Spectrum Scale through the **Quotas** view, as shown in Figure 7-19.

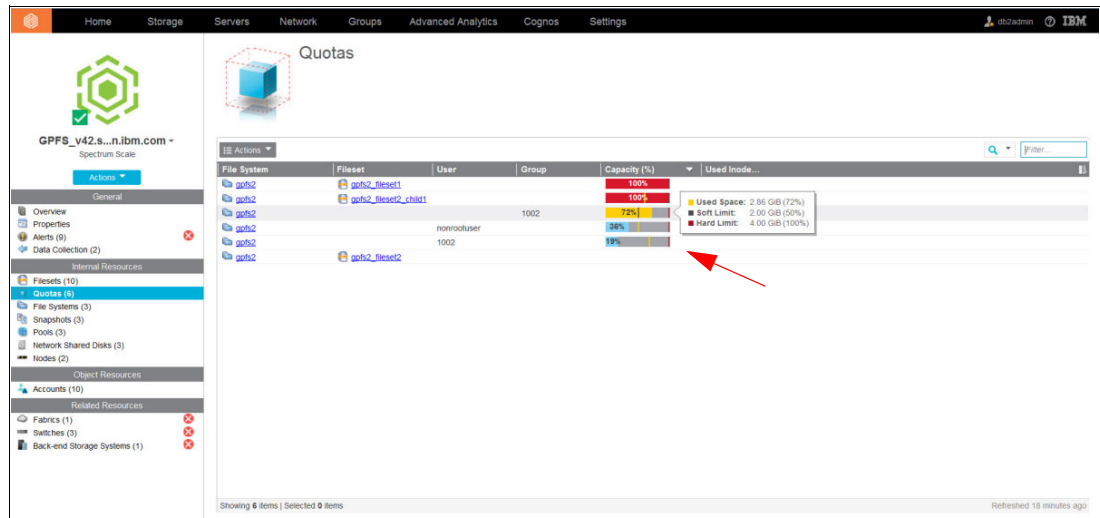


Figure 7-19 Spectrum Control quotas

A storage administrator can quickly determine whether the soft or hard limit is violated. Figure 7-19 shows the lines in the stacked bar chart, which denote the soft/hard limits on each quota. Sorting is possible on this column (by default, it is sorted on which quota limits are reached) and the capacity value can exceed 100% if quota limit checking is not activated.

In the Spectrum Scale web-based GUI, you can create or edit a quota. You can see whether a capacity value can exceed 100% by viewing the status of that quota, which is either “Enforced” or “Disabled”, and by viewing the small icon next to the Used columns for Soft % used and Hard % Used (Figure 7-20). The Quotas can be found by clicking **Files** → **Quotas**.

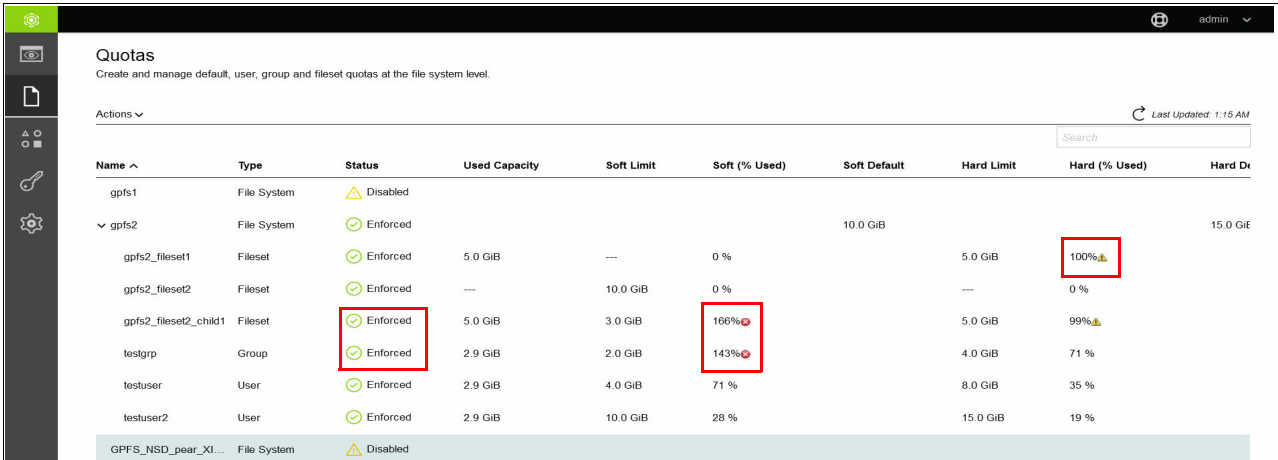


Figure 7-20 Spectrum Scale web-based GUI quotas

Which file sets are close to running out of free inodes

Spectrum Scale uses file sets that are defined as a subtree of a file system. In many respects, a file set behaves like an independent file system. You can use file sets to administer partitions of a file system at a granular level.

Spectrum Control can show details about the file sets, and a storage administrator can easily determine which file set is at risk of running out of inodes, as shown in Figure 7-21, by looking at the stacked bar chart, which is common through Spectrum Control table for User Inodes (%).

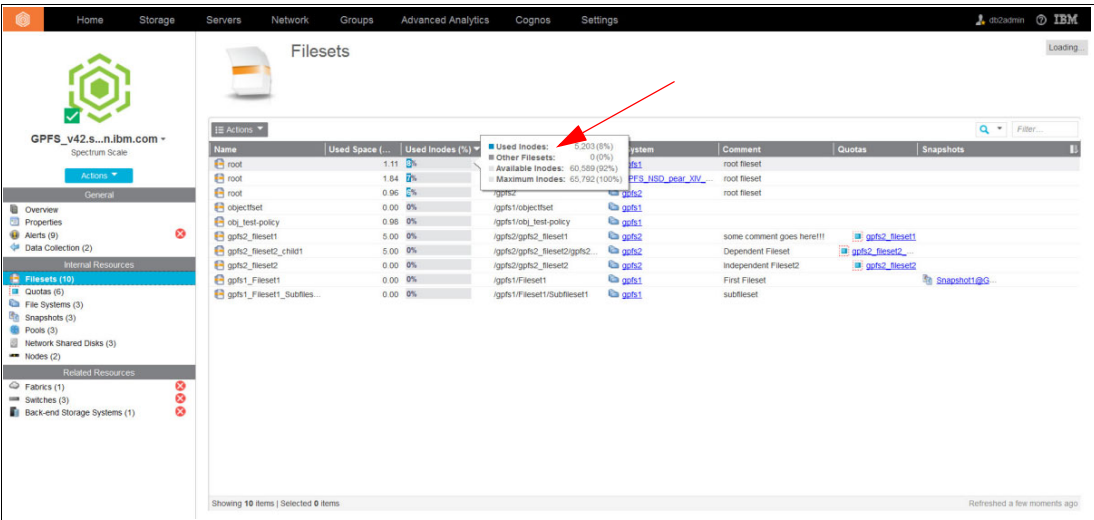


Figure 7-21 Spectrum Control file sets

The inodes are important when working with Spectrum Scale. An inode is a UNIX operating system data structure that contains important information about the files in the file system, which is known as metadata. Every time a file or directory is created on the file system, a new inode is allocated. If at the creation time of the file system the bytes/inode ratio is not correct, there is a risk of running out of inodes. A larger byte per ratio means that fewer inodes are created, but the typical impact of 1% of the size of the file system is higher. A smaller byte per ratio that is combined with a high number of files on the file system can lead to exhaustion of the inodes number.

Spectrum Scale also has the concept of a dependent file set, which has its inode space allocated from an independent file set, which means it many dependent file sets, which also can lead to depletion of the inodes. The advantages of using dependent file sets are that at snapshot time, a snapshot can be created for an independent file set that includes all the dependent file sets.

The Spectrum Control Overview window has a menu that you can use to determine which independent file sets have the fewest available inodes, as shown in Figure 7-22.

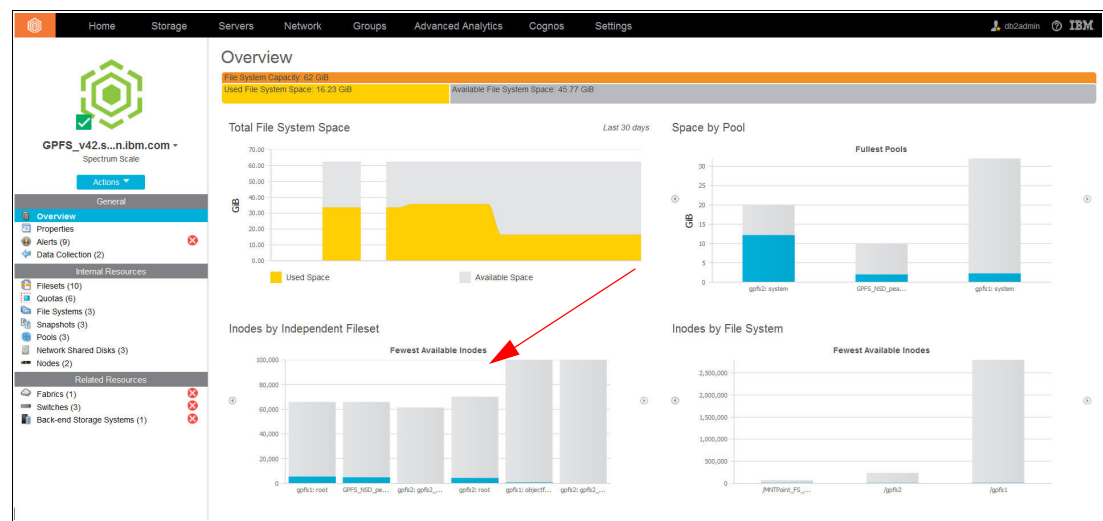


Figure 7-22 Spectrum Control Overview window - inodes by independent file set

Using the Spectrum Scale web-based GUI, you can see the same information presented by Spectrum Control in Figure 7-21 on page 214. The only difference is that you can use the Spectrum Scale web-based GUI to create and modify the file sets. You can access the file sets by using the left menu pane and clicking **Files** → **Filesets**, as shown in Figure 7-23.

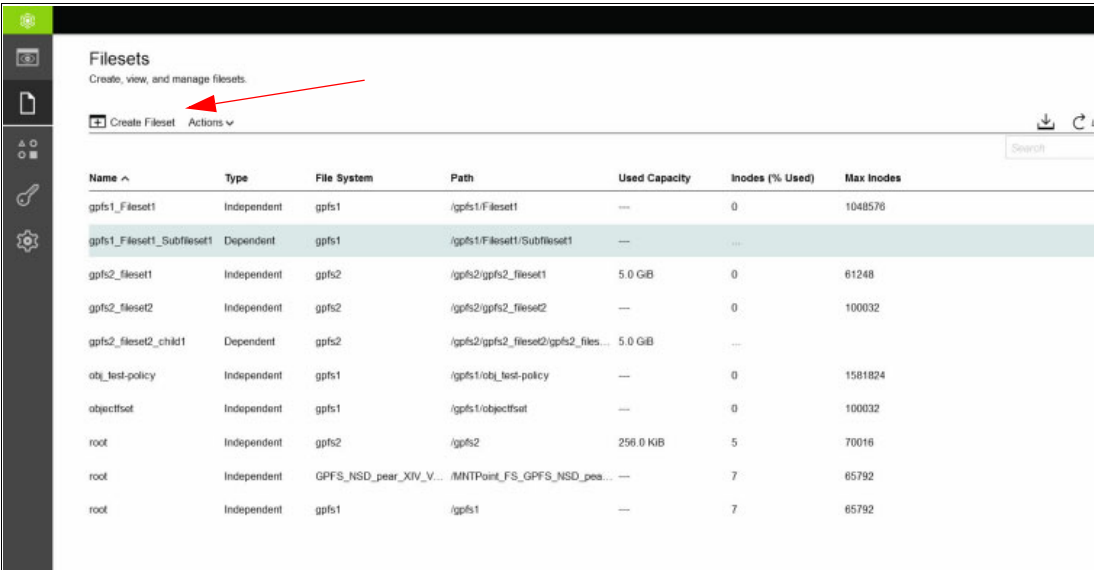


Figure 7-23 Spectrum Scale GUI file sets

By using the Spectrum Scale CLI, you can run the **mmfsfileset** command (Example 7-7) to show a specific file set that is related to a file system; to show the inode usage, run the **mmdf** command.

Example 7-7 Spectrum Scale mmfsfileset command

```
[root@rice5 ~]# mmfsfileset gpfs1 -d --block-size auto
Collecting fileset usage information ...
Filesets in file system 'gpfs1':
Name                               Status  Path
Data
root                               Linked  /gpfs1
1.112G
gpfs1_Fileset1                     Linked  /gpfs1/Fileset1
256K
gpfs1_Fileset1_Subfileset1         Linked  /gpfs1/Fileset1/Subfileset1
256K
objectfset                         Linked  /gpfs1/objectfset
2.016M
obj_test-policy                    Linked  /gpfs1/obj_test-policy
1007M
```

Are the volumes that are backing up the network shared disks performing correctly

Figure 7-24 on page 217 shows the scenario when a performance problem must to be investigated and the Spectrum Scale topology has the nodes that are directly attached with SAN storage.

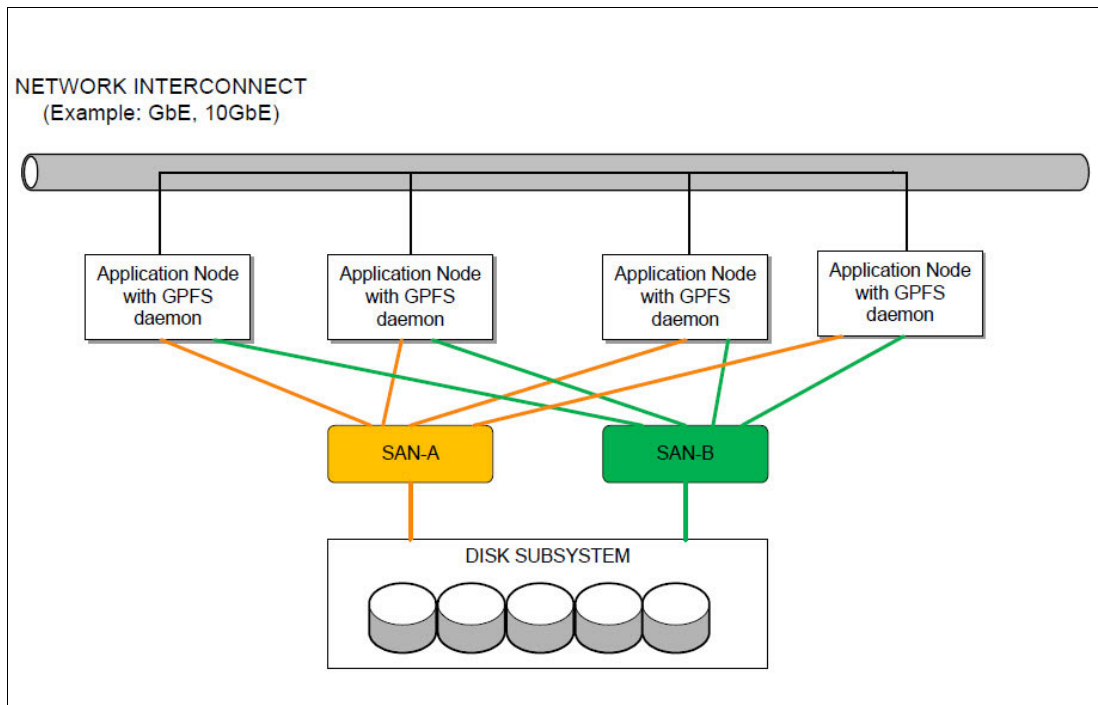


Figure 7-24 Spectrum Scale with direct attached storage (SAN)

Spectrum Scale can show details about the network shared disks (NSDs). These NSDs provide a block-level interface over TCP/IP network. The mounted file system looks the same to the users and applications (Spectrum Control handles the I/O request transparently), as shown in Figure 7-25.

The screenshot shows the IBM Spectrum Control interface. The left sidebar contains navigation options: Overview, Properties, Alerts (9), Data Collection (2), Internal Resources (Filesets (10), Quotas (6), File Systems (3), Snapshots (3), Pools (0), Network Shared Disks (3), Nodes (2)), and Related Resources (Fabrics (1), Switches (3), Back-end Storage Systems (1)). The main panel displays 'Network Shared Disks' with a status summary (3 Normal, 0 Warning, 0 Error) and a table of NSDs.

Name	Status	Disk Capa...	Type	Pool	File System	NSD Servers	Mounted Storage Volume
nsd2	Normal	61%	Data, Metadata	system	gofs2	peer.storage.1	
GPFS_NSD_pear_XIV_Vor10G	Normal	10%	Data, Metadata	system	GPFS_NSD_pear_XIV_Vor10G_FS	peer.storage.tu	GPFS_NSD_pear_XIV_Vor10G
nsd1	Normal	7%	Data, Metadata	system	gofs1	peer.storage.tu	peer_XIV_paths_v1

Figure 7-25 Spectrum Control network shared disks

Spectrum Control is useful in Spectrum Scale environments because it has visibility into the SAN storage. The performance of the block storage backing the NSD volumes can be investigated through the related resources area (Figure 7-26) for cases where the back-end storage is monitored by Spectrum Control.

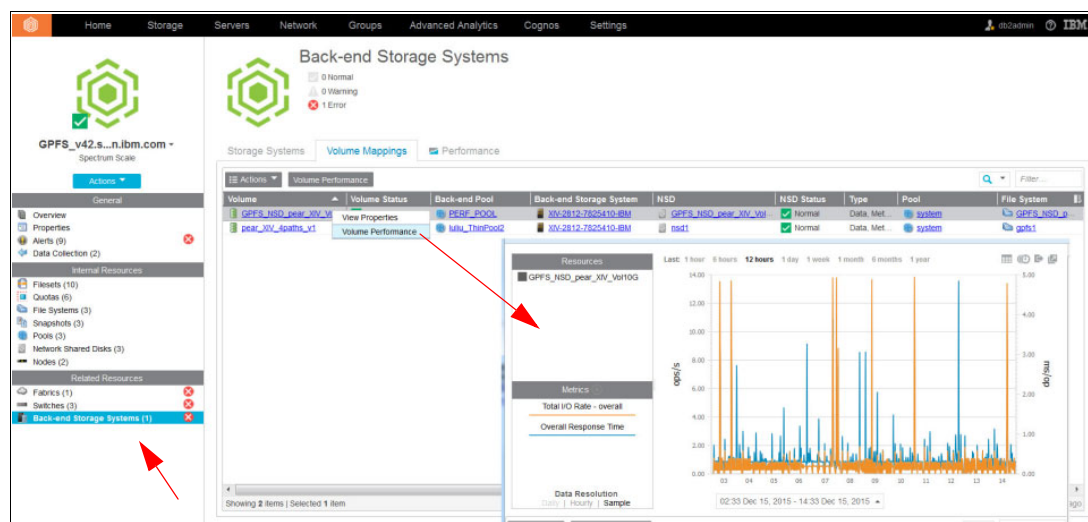


Figure 7-26 Spectrum Control related resources - backing volumes

Spectrum Control automatically can correlate the monitored storage devices with the LUNs that are assigned to be used by Spectrum Scale. To do so, use the web-based GUI and go to Related Resources (arrow 1 in Figure 7-26). Here, you can see details about the fabrics and switches through which the SAN storage is connected.

The performance charts can be displayed for either the Volume Mappings or for the related storage systems by using the Performance tab. These performance charts display the typical metrics that you can find for the block storage devices, which by default are “Overall I/O Rate” and “Response Time”.

The Storage System tab along with the “Condition” status can help you answer the next question a storage administrator wants to know: Which NSDs are at risk of becoming unavailable? Whenever the backing storage condition is in the “Error” state, this can indicate that the storage device is at risk of becoming unavailable and requires attention.

In addition to the block-level performance, Spectrum Control can collect performance information for the file side of Spectrum Scale, which is described in “Is the performance of the file systems and cluster nodes satisfactory” on page 219.

Note: The Spectrum Scale GUI cannot report on the SAN-attached storage performance.

NSDs are also at the risk of becoming unavailable whenever the mounting servers (nodes) lose connectivity with the NSD. A node typically serves multiple NSDs. In the Nodes window, the user can easily determine whether any of the NSDs are at risk by viewing the NSDs Served column. When you open the property notebook, you see the list of NSDs Served along with the “Connectivity” column, which shows “Connected” if the status of the connectivity is good (Figure 7-27 on page 219).

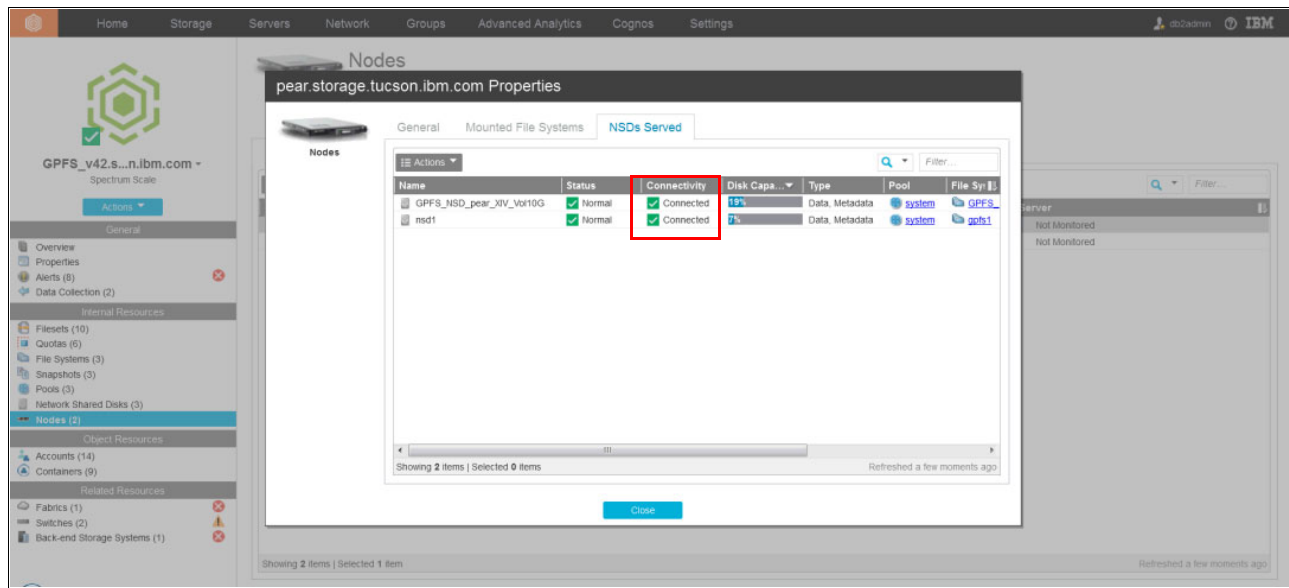


Figure 7-27 Status of NSDs served along with the Connectivity column

Is the performance of the file systems and cluster nodes satisfactory

As stated in “Enhancements in Version 5.2.8” on page 197, starting with Spectrum Control V5.2.8, monitoring support for performance is available for Spectrum Scale. The performance chart is available for file systems and nodes.

Using these performance charts, a storage administrator can easily determine whether things are out of balance, or compare the performance across multiple clusters. Using this information, an administrator can understand which of the clusters are doing most of the workloads, which helps them plan.

Figure 7-28 shows an example of a Spectrum Control performance chart.

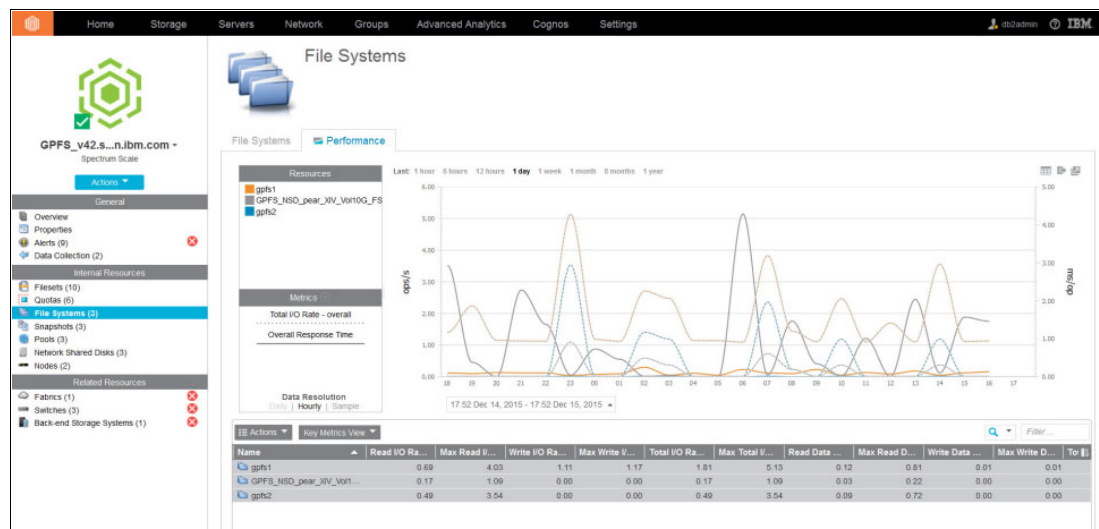


Figure 7-28 Spectrum Control file stem performance chart

Figure 7-28 on page 219 shows the performance of the three file systems that are managed by the GPFS_v42 cluster. The metrics that are supported for file systems can be shown for Read/Write or Total/Overall and are as follows:

- ▶ Overall I/O Rate (ops/s)
- ▶ Data Rate (MiB/s)
- ▶ Response Time (ms/op)

Figure 7-29 shows the performance of the two nodes that make up the GPFS_v42 cluster. Here are the metrics that are supported for the nodes:

- ▶ System CPU Utilization (%), User CPU Utilization (%), and Total CPU Utilization (%)
- ▶ Memory Utilization (%), Cache and Buffer Utilization (%), and Total Memory Utilization (%)

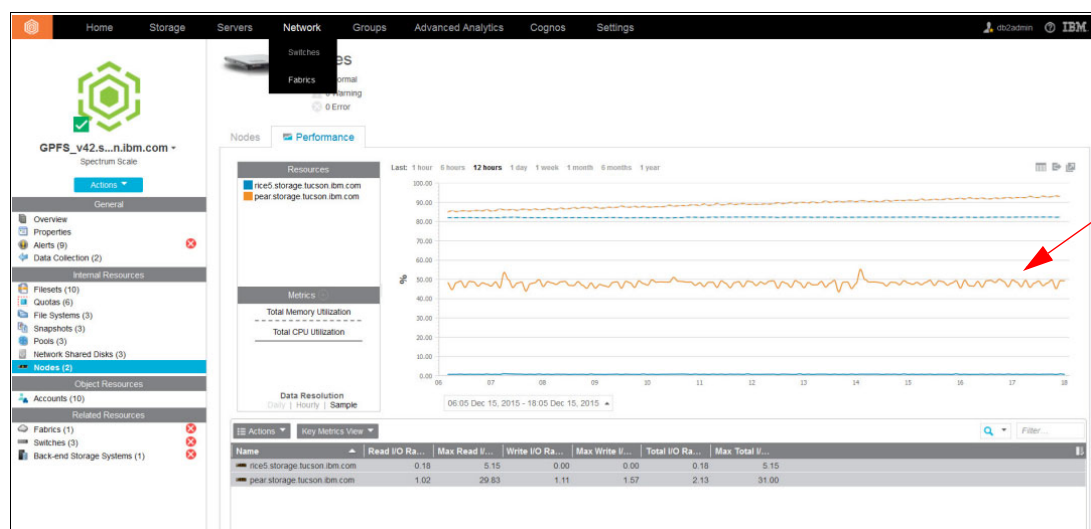


Figure 7-29 Spectrum Control Nodes performance chart

In Figure 7-29, the CPU utilization for rice5 cluster node hovered around 50% and the pear cluster node is at 2-3% in the last 12 hours. This usage is similar to a longer period in the past. This metric suggests that either the rice5 node hardware is potentially becoming obsolete and performance is marginal, or the workload is not balanced across the cluster nodes and some configuration adjustment is necessary. Another possibility is that some user process on rice5 is consuming more processor than the system processes. This can be a sign of an application of which the storage administrator is unaware, but if you use the charts that are offered by Spectrum Control, this consumption is easy to detect.

The Spectrum Scale GUI offers performance charts, which can be viewed in the following ways:

- ▶ Viewing them on the dashboard
- ▶ Clicking **Monitoring** → **Performance**

On the dashboard, there is a chart with two views showing the *throughput* and the *top node by system load* (Figure 7-30 on page 221).

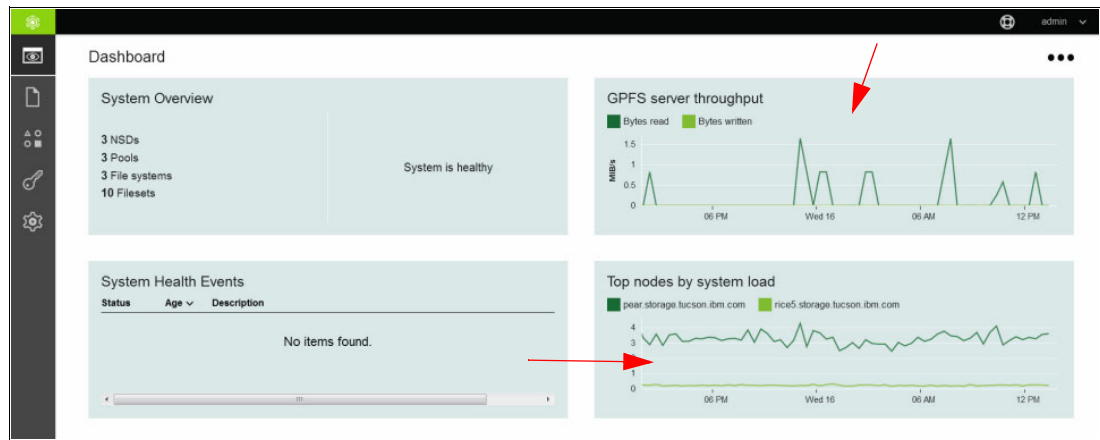


Figure 7-30 Spectrum Scale GUI performance charts

In the Monitoring → Performance window, the number of views is configurable. You can show either one or two views with various resources, the aggregation level, and metrics, as shown in Figure 7-31.

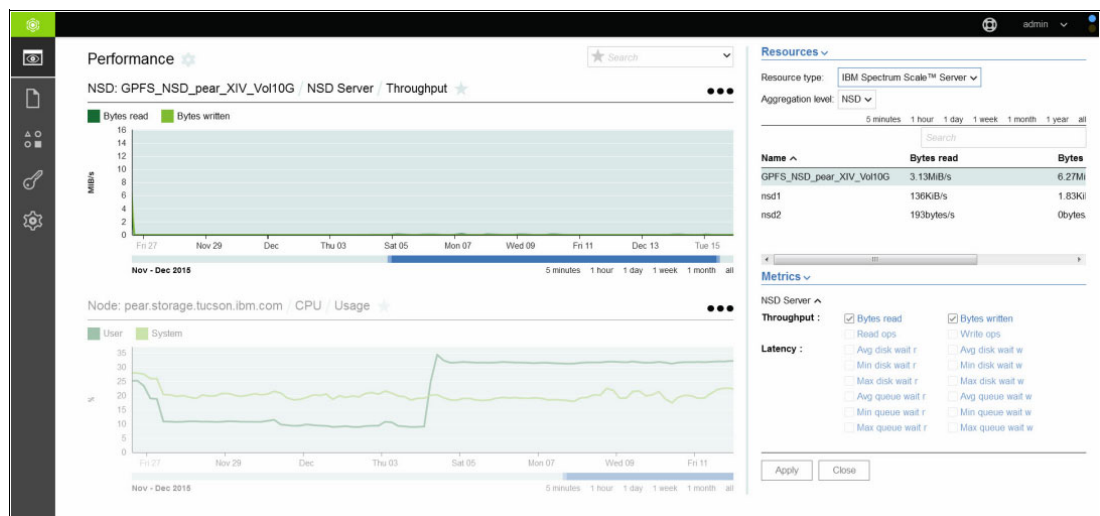


Figure 7-31 Spectrum Scale GUI performance monitoring charts

Resources types can be the following ones:

- Network (InfiniBand or Ethernet)

Aggregation level:

- Node
- Adapter per node
- Adapter
- Cluster

- System resources

Aggregation level:

- Node
- Cluster

► Spectrum Scale Server

Aggregation level:

- NSD
- File System
- Pool
- Cluster

► Spectrum Scale Client

Aggregation level:

- File System
- Cluster

The metrics that are available for each resource type are much richer than what Spectrum Control has available. You can find more information about the metrics in the Spectrum Scale IBM Knowledge Center, found at:

<http://ibm.co/1VZUc5m>

For more information about performance monitoring by using the Spectrum Scale GUI, see the Spectrum Scale IBM Knowledge Center, found at:

<http://ibm.co/1K9fjLN>

Are all nodes that fulfill critical roles in the cluster running

A node in the Spectrum Scale cluster is considered to be any server on which the product is installed with direct storage access or network access to another node. Depending on the type of access, each node can have different roles. Some of the nodes are essential to the normal functioning of the cluster (cluster manager, quorum node, file system manager, or meta-node), and other nodes, such as the application node, have a smaller role, but as a storage administrator, you want to make sure that all nodes fulfilling critical roles have a good health status.

The Spectrum Control Nodes window is shown in Figure 7-32.

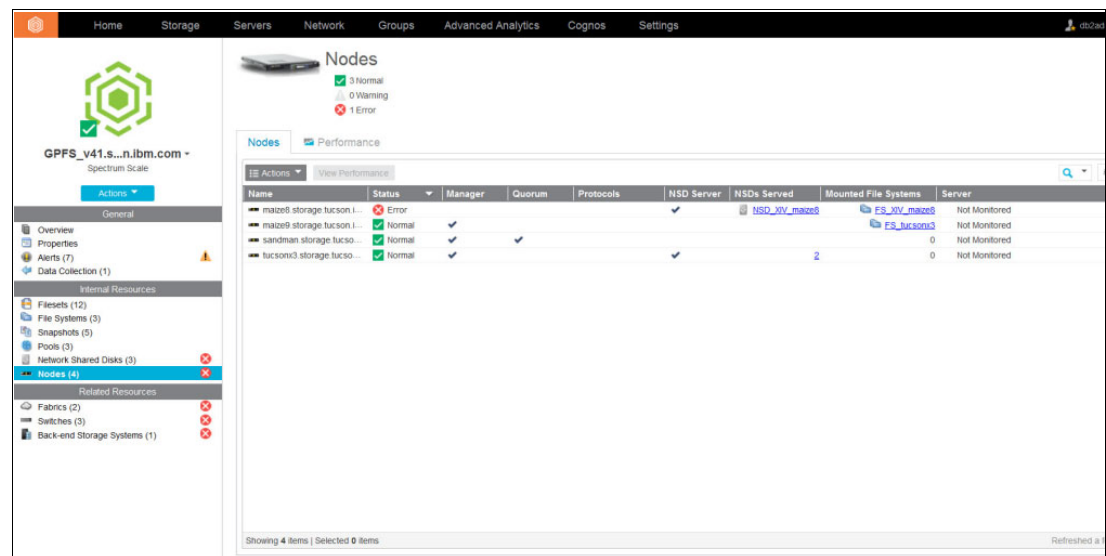


Figure 7-32 Spectrum Control nodes window

Figure 7-32 on page 222 shows the maize8 node in “Error” status, but all the rest of the nodes fulfilling critical roles (Manager, Quorum) are in “Normal” status, so the Spectrum Scale cluster continues to function correctly.

Which object accounts manage the most containers and are violating quotas

Starting with Spectrum Control V5.2.8, the Spectrum Scale object personality is supported and you have visibility into object accounts and containers. One task a storage administrator must deal with is which accounts are managing the most containers or which quotas are being violated.

Note: Object/Swift has its own set of quotas, that is, these are not GPFS quotas.

For more information about installing Object Storage (Swift) manually and quotas, see the IBM Cloud Manager with OpenStack 4.3.0 IBM Knowledge Center, found at:

<http://ibm.co/1UbV19R>

In Figure 7-33, you can sort the table on the Containers column and concurrently see Space Quota (%), which is in violation of the administrator account.

You can use Spectrum Scale GUI to display Accounts-related information by going to the left pane and clicking **Object** → **Accounts**. The Object menu shows users and roles that are not being shown as views in Spectrum Control. Spectrum Scale can create or edit the users and roles.

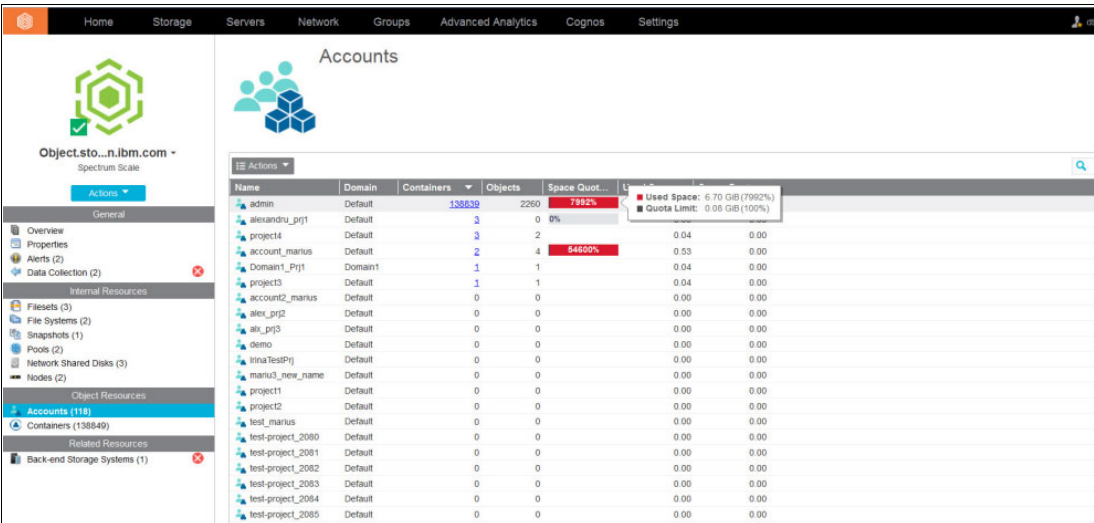


Figure 7-33 Spectrum Control Object accounts

Which object containers are violating space quotas

After a storage administrator determines which accounts are violating the space quota, the next question is which containers are violating space quotas or object counts? Spectrum Control can display information about containers, as shown in Figure 7-34.

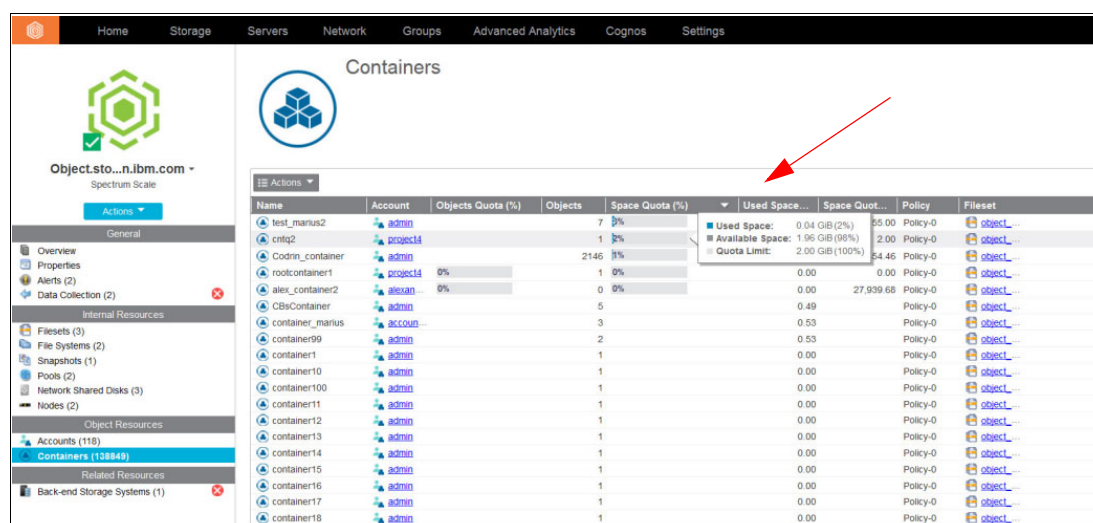


Figure 7-34 Spectrum Control displays information about containers

The Spectrum Scale GUI does not have a view around containers, but you can see them by using the Accounts view.

Being notified when a node goes offline or a file system violates a threshold

The alerting capabilities were improved in Spectrum Control V5.2.8. You can define alerts and be notified when configuration or capacity violations occur.

Note: As of Spectrum Control V5.2.8, threshold alerts are not supported for Spectrum Scale performance.

For more information about these alert improvements, see the IBM Knowledge Center, found at:

<http://ibm.co/1SJHV1H>

This section describes how to define a warning severity alert for when a Spectrum Scale node goes offline or the file system fills beyond 40%. Complete the following steps:

1. Open the Alerts definition window, where you can set alerts for a series of metrics. Click **NAS Node** and then the **General** tab, as shown in Figure 7-35 on page 225.

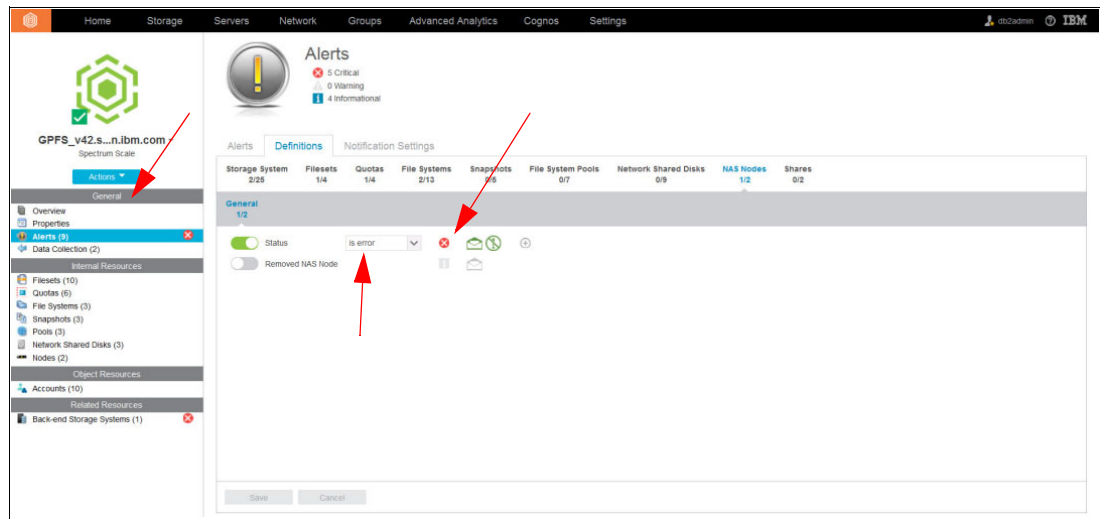


Figure 7-35 Spectrum Scale alerts definition for nodes

2. Select **is error** for the Status, and for severity, select the “error” icon, as shown in Figure 7-35.
3. Click **File Systems** and then the **Capacity** tab and look for the Used Space % metric. Change the alert to Enabled status by using the switch. Enter “>=” and “40%”, and for the severity, set the “warning” icon for an early alert. Select + to double the metric, enter “>=” and “80%”, and select the “error” icon for severity. Click **Save**, as shown in Figure 7-36.

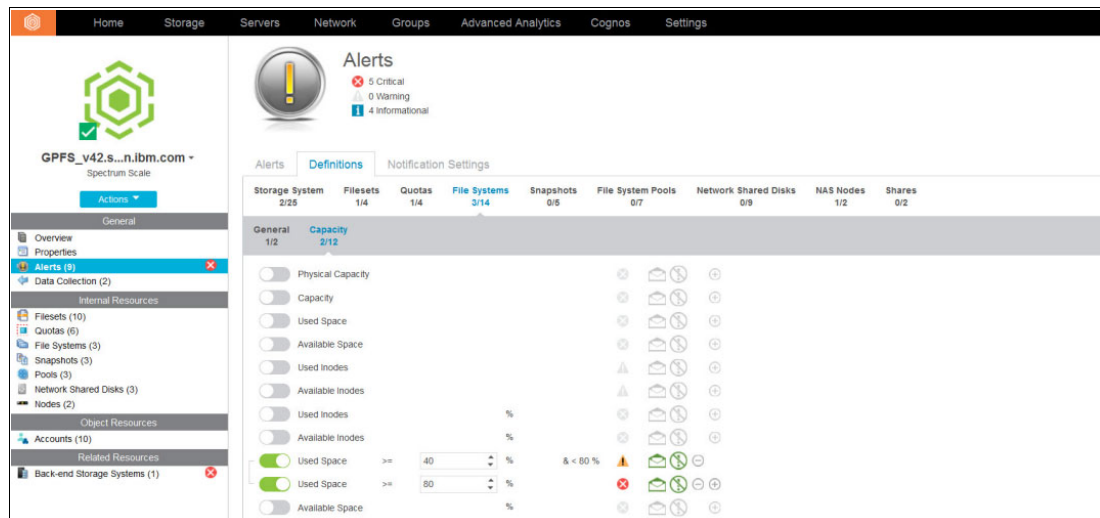


Figure 7-36 Spectrum Control alert definition for file systems

When you run the next probe, check whether the alert is being triggered for any of your nodes or file systems, as shown in Figure 7-37.

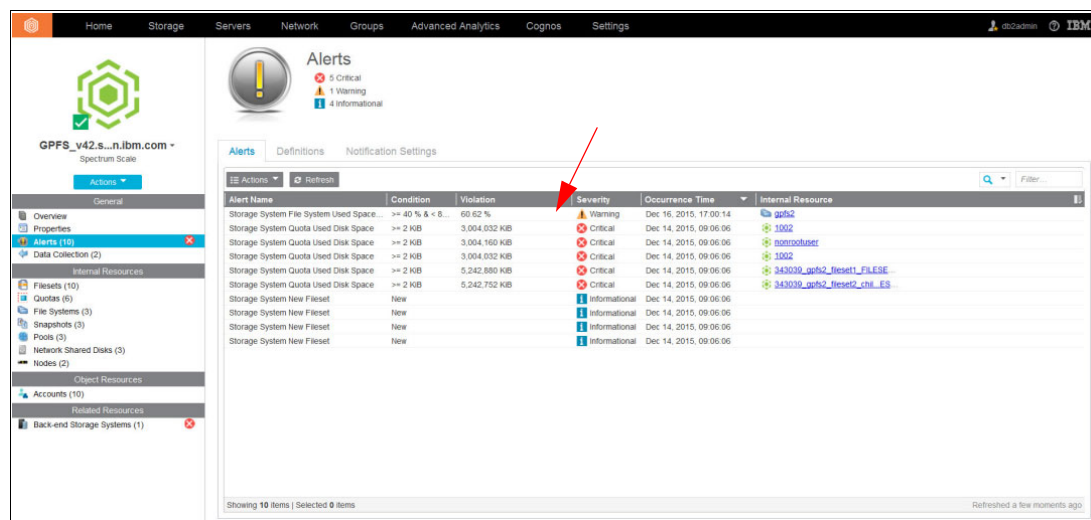


Figure 7-37 Spectrum Control Alert for File System violation triggered

As shown in Figure 7-37, the alert for the File System Used Space Percent for the “>= 40%” setting was triggered and the severity of the alert is “warning”.

Spectrum Scale GUI has a limited event monitoring capability. It is not possible to define any alerts and triggering conditions, but you can use it to run fix procedures directly from the GUI, as shown in Figure 7-9 on page 205.

7.3 Advantages of using the Spectrum Control web-based GUI to manage Spectrum Scale clusters

Using the Spectrum Control web-based GUI to manage Spectrum Scale clusters has the following advantages:

- ▶ Monitoring of multiple Spectrum Scale clusters and keeping historical data for an extended period
- ▶ Comparing performance across all the managed Spectrum Scale clusters, allowing side-by-side comparison at various levels (node/file system) and taking advantage of all the charting features of Spectrum Control (time sync and table view)
- ▶ Correlating the SAN-attached storage and performance charting of the block level LUNs
- ▶ Showing the related resources, such as back-end storage systems, switches, and fabrics
- ▶ Comprehensive set of alerting capabilities
- ▶ Allows grouping of resources into applications and departments
- ▶ Health status information at the resource level

7.4 Advantages of using the Spectrum Scale GUI to manage Spectrum Scale

Using the Spectrum Scale GUI to manage an individual Spectrum Scale cluster has the following advantages:

- ▶ Monitoring only a single cluster at a time, but has a richer set of functions and details
- ▶ Monitoring performance on a larger set of metrics, including the networking layer for InfiniBand or Ethernet
- ▶ Visualizing the internal details and changing the existent configuration
- ▶ Display of events and running fix procedures in place
- ▶ Display of topological information



IBM Spectrum Control Advanced Edition overview

IBM Spectrum Control Advanced Edition (Spectrum Control Advanced) provides all of the functions of Spectrum Control Standard Edition, and all of the advanced analytics functions. This powerful solution enables organizations to optimize provisioning, capacity, availability, and reporting. With Spectrum Control Advanced V5.2.8, updates to the new cloud configuration and provisioning functions are available. Using the Provisioning Storage wizard, you can request the amount of storage that you need while conforming to the requirements of the placement of volumes and shares.

This chapter provides a brief overview if you are not already familiar with the functions that were previously exclusively available only with IBM Virtual Storage Center (VSC). The topics in this chapter complement the Spectrum Control IBM Knowledge Center, which is referenced for product details throughout this chapter. You can find it at the following website:

<http://ibm.co/1N39nno>

For a comparison of Spectrum Control Advanced features compared to Spectrum Control Standard Edition features, see 1.6.3, “Spectrum Control Standard Edition features” on page 7 and 1.6.4, “Spectrum Control Advanced Edition comparison to the Standard Edition” on page 11.

The following publications contain details and use cases that demonstrate Spectrum Control Advanced analytics functions:

- ▶ *IBM SmartCloud Virtual Storage Center*, SG24-8239
- ▶ *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204

8.1 Introduction

Spectrum Control Advanced provides advanced storage analytics for storage optimization and provisioning capabilities. Those capabilities are used in cloud-based and software-defined storage environments where the customer's workload requirements must be met in real time.

Spectrum Control Advanced uses real performance metrics and advanced analytics to make recommendations to optimize storage pools and volumes by redistributing workloads across storage environment. By using real performance metrics, it enables optimization decisions to be made based on actual usage patterns, rather than on predictions.

Spectrum Control Advanced includes all of the features of Spectrum Control Standard Edition, and the following features:

- ▶ Tiered storage optimization with intelligent analytics
- ▶ A service catalog with policy-based provisioning
- ▶ Self-service provisioning with restricted use logins
- ▶ Application-aware, hardware-assisted snapshot management from IBM Spectrum Protect Snapshot

To take advantage of the optimization and provisioning capabilities that are available in Spectrum Control Advanced, some configuration is required. This configuration is called *cloud configuration*, and it specifies storage tiers, service classes, and capacity pools that are used to transform your environment to an efficient, self-optimized environment.

For information about key concepts of the following advanced topics, see the Spectrum Control IBM Knowledge Centers:

- ▶ Storage optimization:
<http://ibm.co/1VZUKIx>
- ▶ Cloud configuration:
<http://ibm.co/10KR0ho>
- ▶ Provisioning storage
<http://ibm.co/10XmE8v>

For an overview of product licensing and Spectrum Advanced Edition features, see the *IBM Spectrum Control Advanced Edition and IBM Spectrum Control Advanced Select Edition* topic, found on the Spectrum Control website:

<http://ibm.co/1P33zf0>

For information about how to get started quickly with a typical installation of Spectrum Control Advanced, see the Spectrum Control IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.8/com.ibm.spectrum.sc.pdf/qs_trtscae.pdf

8.2 Cloud configuration

With the Spectrum Control cloud configuration function, you can organize your storage environment by categorizing the resources in tiers and defining service classes with specific qualities of service (QoS). This function is a requirement in optimizing your storage environment. It is also highly preferred in a traditional storage provisioning context. To provision storage, you specify only the storage capacity and storage quality that is required. By using the service class, volumes are created with the required characteristics.

To optimize storage performance, Spectrum Control analyzes tiering and, according to the recommendations that are generated, volumes are redistributed across each tier and can be moved to higher or lower tiers. Cloud configuration involves the following steps:

1. Assign the tier to storage pools, which is described at the following website:
<http://ibm.co/10mzibE>
2. Refine the customization by assigning tags to storage resource, which is described at the following website:
<http://ibm.co/1SJr65P>
3. Define and configure capacity pools, which is described at the following website:
<http://ibm.co/1PnPPvG>
4. Define and configure service classes, which is described at the following website:
<http://ibm.co/10dVoPD>

8.2.1 Tiering

Storage tiering is a step toward defining service classes and provisioning volumes that require a certain tier level. It also allows you to optimize the placement of volumes by using Spectrum Control optimization tasks, which redistribute volumes within the tier or move the volumes to higher or lower tiers. Spectrum Control uses 10 tiering levels of storage pools: tier 1 represents the best performing tier, and tier 10 the least performing tier. You can use as many levels as you want, based on the storage systems and disk technologies in your environment.

For example, you might categorize the following storage tiers:

- ▶ Tier 1. Leave this level available for future use.
- ▶ Tier 2. Flash disk pools and solid-state drive (SSD) disk pools.
- ▶ Tier 3. Leave this level available for future use.
- ▶ Tier 4. Fibre Channel and SAS disk pools.
- ▶ Tier 5. Leave this level available for future use.
- ▶ Tier 6. Hybrid pool.
- ▶ Tier 7. Leave this level available for future use.
- ▶ Tier 8. NL-SAS disk pool.
- ▶ Tier 9. Leave this level available for future use.
- ▶ Tier 10. SATA disk pool.

Figure 8-1 shows an example of how to set tiers 1-10.

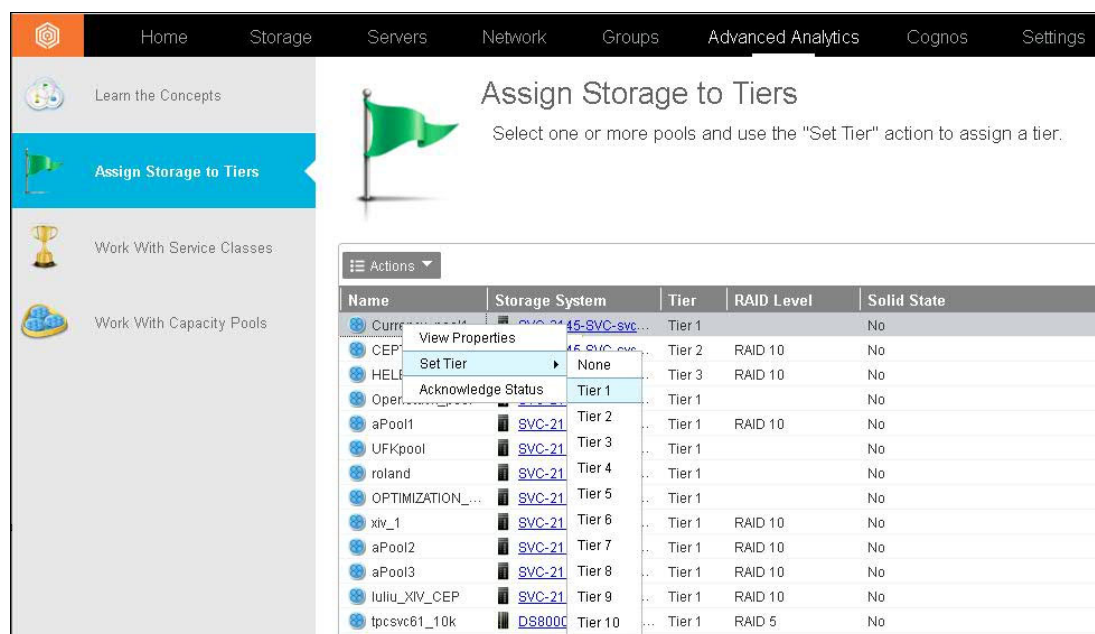


Figure 8-1 Example of assigning tiers 1-10 to storage

For more information about tiering, see the Spectrum Control IBM Knowledge Center at the following website:

<http://ibm.co/10mzibE>

8.2.2 Service classes

A *service class* is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular QoS, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, storage tier, volume mirroring, multipathing policy, and whether storage resources must be able to encrypt or thin provision volumes. Service classes simplify provisioning requests by representing a level or type of storage quality. When you request storage, only the required capacity and service class must be specified. Before you can provision storage, you must create service classes that describe the capabilities and characteristics of the storage that you want to be able to provision.

Service classes can later be modified or deleted as the needs of your installation change. Spectrum Control accounts for all attributes of the service class, which is specified in a provisioning request, and also current storage resource utilization (space, performance, and status) to identify the most appropriate resource for the new volume or share. Service classes are a key point in mapping business requirements (capacity, accessibility, performance, and availability) in infrastructure capabilities (media type, disk technologies, RAID Levels, encryption, compression, and thin provisioning). Two types of service classes are available:

- **Block-storage service classes:**


This type describes attributes and capabilities of block storage resources. When you provision volumes, you specify the requirements by using a block-storage service class.

Based on the requirements of the service class, Spectrum Control identifies a storage pool for the volume. Spectrum Control provides three predefined block-storage service classes.

Although the classes are predefined, you might want to change them or adjust the tiers, RAID levels, and so on:

- Gold: Defined for mission-critical applications (highest-performing storage).
- Silver: Defined for applications in production (high-performing storage).
- Bronze: Defined for non-mission-critical applications (standard storage).
- File-storage service classes:
 - Normal Isolation (normal isolation file storage): The file system from which the NAS share is provisioned can contain other NAS shares. Shared storage is allowed.
 - Enhanced Isolation (enhanced isolation file storage): The file system from which the NAS share is provisioned cannot contain other NAS shares. Dedicated storage is required.

Figure 8-2 shows the default service classes and a custom service class.



Name	Type	Used Space (%)	Total Capacity (GiB)	Available Space (GiB)	Unavailable Space (GiB)	Description
Silver	Block	0%	0.00	0.00	0.00	High-performing storage for applications in production.
Gold	Block	0%	0.00	0.00	0.00	Highest-performing storage for mission-critical applications.
D8K_SC	Block	0%	2,821.00	2,821.00	0.00	Service class for only D8K.
Bronze	Block	0%	0.00	0.00	0.00	Standard storage for non-mission-critical applications.
Normalisolation	File	0%	0.00	0.00	0.00	Normal isolated file storage.
Enhancedisolation	File	0%	0.00	0.00	0.00	Enhanced isolated file storage.

Figure 8-2 Default service classes and example of a custom service class

For more information about service classes, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/10mzGqr>

8.2.3 Capacity pools

Capacity pools are logical groups of storage resources and are optional for the more advanced user. They include these items:

- Storage systems
- Storage pools
- File systems of file storage systems
- IBM Scale Out Network Attached Storage Network Shared Disks (NSDs)

Capacity pools can group storage resources in any way that serves the needs of your environment or business. By configuring capacity pools, you can track the used and available space for block and file storage on any set of storage resources. You can use capacity pools to define a set of storage resources from which provisioning requests must be satisfied. When you request storage by using the Provision Storage wizard or the vSphere Web Client extension, you can specify a capacity pool, and the provisioning request is restricted to the resources in this capacity pool. You can also associate a service class with capacity pools. If a service class is associated with capacity pools, provisioning requests for the service class must be specified and be constrained to one of the associated capacity pools. Because you can also grant to non-administrative users permission to provision storage by using the service class, associating capacity pools with the service class restricts the users to a specific set of resources.

Figure 8-3 shows the capacity pool wizard.

Create Capacity Pool

Name: Resource tags:

Description:

Capacity: Used Block 7,618.00 GiB (47.14%) 8,541.00 GiB Total 16,159.00 GiB

Figure 8-3 Capacity pool wizard

For more information about configuring capacity pools, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/1PnPPvG>

8.3 Storage provisioning

You can provision storage resources to servers, hypervisors, or a cluster. You can provision storage resources, such as storage volumes or NAS file shares, to one or more servers, one or more hypervisors, or one cluster. To provision the storage resources, you must specify only the storage capacity and storage quality that is represented by using the service class definition. Spectrum Control identifies the storage resources that satisfy the requirements of the specified service class from a pool of available resources.

By default, the pool of available resources includes all block-storage or file-storage resources that are known to Spectrum Control. If you specify a capacity pool, new storage volumes and NAS file shares are provisioned only from storage resources in that capacity pool. From the storage resources that can provide the required service class, Spectrum Control identifies the best storage placement. The determination of the best storage placement is based on storage system free space. For provisioning storage volumes, performance data is also considered. Preference is first given to storage pools and systems that contain volumes for the selected server, hypervisor, or cluster. Preference is then given to systems that have available performance data.

For more information about storage provisioning, see the Spectrum Control IBM Knowledge Center:

<http://ibm.co/10dVDKD>

Also, for more information and use cases, see *IBM SmartCloud Virtual Storage Center*, SG24-8239.

8.4 Optimization

You can optimize resources in your storage environment. The optimization function is used to help improve performance of your storage environment and to better use storage resources. The optimization function has three main tasks:

- ▶ Balance pools
- ▶ Analyze tiering
- ▶ Transforming volumes

For more information about how to implement and use optimization, see the following links on the Spectrum Control IBM Knowledge Center:

- ▶ Optimizing storage tiering
<http://ibm.co/110qm37>
- ▶ Optimizing storage Pools
<http://ibm.co/1R1bdbC>
- ▶ Transforming and migrating volumes
<http://ibm.co/1ndqRbQ>

Note: For more information and practical use cases, see *IBM SmartCloud Virtual Storage Center*, SG24-8239.

8.4.1 Balancing pools

Balancing the workload of volumes across storage pools is done with the balance pools function, where the workload of volumes is balanced across pools on the same tier. The balance pools function is used to analyze the pool performance and, according to the recommendations for pool activity, volumes are moved from pools with high-activity values to pools with low-activity values in the same tier. The balance pools function helps you achieve a more balanced workload across the pools within same tier.

For more information about balancing pools, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1ZMLqNz>

8.4.2 Analyzing tiering

The analyze tiering function is used for optimizing the placement of volumes on storage tiers. The purpose of analyze tiering is to tier volumes based on the criteria that you set in tiering policies. For example, you can tier volumes based on the volume workload, on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated that you can implement to ensure that the volumes are placed on the tiers that best match the workload requirements. The analyze tiering function can help improve storage performance by moving volumes with heavy workloads to the tiers that meet the workload requirements of the volumes. It also can help you to reduce storage costs by moving volumes with low workloads from higher and more expensive tiers to lower and less expensive tiers.

For more information about how to use optimization, see the following links at the Spectrum Control IBM Knowledge Center:

- Analyze tiering by storage pools:

<http://ibm.co/1Ribg7t>

- Analyze tiering by volumes:

<http://ibm.co/1P3EzUS>

Figure 8-4 shows an example of analyzing tiering in the GUI.



Figure 8-4 Example of how to analyze tiering in the GUI

8.4.3 Volume transformation

The volume transformation function is used to convert volumes in storage pools to fully allocated volumes, compressed volumes, or thin-provisioned volumes. It can be also used to move volumes to other pools on the same storage system or to pools that are enabled for Easy Tier on the same storage system.

Here are some of the key advantages of using volume transformation from Spectrum Control Advanced:

- You can do transform volumes from the context of a server. For example, you can check the file system to volume mapping, then the performance, and then convert a volume, whereas from the Spectrum Scale web-based GUI, you cannot.
- You can more easily pause and resume the conversion with Spectrum Control Advanced.

Using Spectrum Control Advanced, you can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

Spectrum Control Advanced can set these characteristics for a volume:

- Migrate pool

The migrate pool option can move the selected volume on the recommended pool to one that you selected from the available pools.

- Thin provisioning

The thin provisioning option enables the storage to present the required capacity to the host while allocating only the used capacity in terms of space on the physical storage media. By using the thin provisioning option, you can convert fully allocated volumes to thin-provisioned volumes or thin-provisioned volumes to fully allocated volumes.

- Compression

The compression option is used to convert your fully allocated volume to a compressed volume, or to transform the compressed volume to a fully allocated volume.

- Easy Tier

Easy Tier is a performance function that automatically migrates or moves extents off a volume to, or from, one MDisk storage tier to, or from, another MDisk storage tier. Easy Tier monitors the host I/O activity and latency on the extents of all volumes with the Easy Tier function enabled in a multitier storage pool over a 24-hour period (heatmap creation). Then, it creates an extent migration plan based on this activity and then dynamically moves high activity or hot extents to a higher disk tier within the storage pool. It also moves extents whose activity has dropped or cooled from the high-tier MDisk back to a lower-tiered MDisk.

For more information about transforming volumes, see the Spectrum Control IBM Knowledge Center, found at:

<http://ibm.co/1ndqRbQ>

8.5 IBM Spectrum Protect Snapshot

IBM Spectrum Protect Snapshot (formerly known as IBM Tivoli FlashCopy Manager) uses advanced snapshot capabilities of disk subsystems to create application, database, and VM-aware point-in-time snapshots. These snapshots can be retained on disk only, or stored by using IBM Spectrum Protect to take advantage of the full range of advanced data protection and data reduction capabilities. As the explosion of data continues to increase, and the backup windows decrease, an application-aware snapshot operation consumes much less time than conventional tape backup and can be offloaded to a backup server.

Although many storage systems are now equipped with volume snapshot tools, these hardware-based snapshot technologies provide only “crash consistent” copies of data. Many business critical applications, including those that rely on a relational database, need an extra snapshot process to ensure that all parts of a data transaction are flushed from memory and committed to disk before the snapshot. This process is necessary to ensure that you have a usable, consistent copy of the data.

IBM Spectrum Protect Snapshot helps deliver the highest levels of protection for mission-critical IBM DB2, SAP, Oracle, Microsoft Exchange, and Microsoft SQL Server applications that use integrated, application-aware snapshot backup and restore capabilities. This protection is achieved through the exploitation of advanced IBM storage hardware snapshot technology to create a high-performance, low-impact application data protection solution.

The snapshots that are captured by IBM Spectrum Protect Snapshot can be retained as backups on local disk. With optional integration with IBM Spectrum Protect, customers can use the full range of advanced data protection and data reduction capabilities, such as data deduplication, progressive incremental backup, hierarchical storage management, and centrally managed policy-based administration.

For more information about configuring and using IBM Spectrum Protect Snapshot, see *IBM SmartCloud Virtual Storage Center*, SG24-8239 and the *IBM Spectrum Control Advanced Edition and IBM Spectrum Control Advanced Select Edition* topic in the Spectrum Control Product Licenses IBM Knowledge Center, found at:

<http://ibm.co/1N39nno>



A

Managing security vulnerabilities

Media attention has drawn public focus to security vulnerabilities in both software and hardware products. An application such as IBM Spectrum Control (Spectrum Control) or any of its included components might have security vulnerabilities. If this happens, it is important to understand the process IBM uses to communicate information to the customer.

This appendix provides an overview of the IBM Product Security Response Team and how their processes relate to managing any possible security vulnerabilities in Spectrum Control and its related components. The following resources are described in this appendix:

- ▶ Monitor the IBM Support portal.
- ▶ Monitor the IBM Product Security Incident Response Team (PSIRT) blog.
- ▶ Subscribe to My Notifications.

Note: A description of these possibilities does not imply that a product is affected by any particular vulnerabilities. Always see published security bulletins.

IBM Product Security Response Team overview

The PSIRT manages the receipt, investigation, and internal coordination of security vulnerability information that is related to IBM offerings. This team coordinates with IBM product and solutions teams to investigate and identify the appropriate response plan. Their processes allow for communication of confirmed vulnerabilities with customers.

Communicating about security vulnerabilities

IBM does not publicly disclose or confirm security vulnerabilities until IBM conducts an analysis of the product and issues fixes or mitigations. Additionally, IBM does not intend to provide vulnerability details that can enable someone to craft an exploit. Communication occurs primarily through the security bulletins that are published on the IBM Support portal, but targeted and discrete communication might also occur. After they are published, the security bulletins for a product provide information to allow the user to identify the following items:

- ▶ Public security vulnerability details
- ▶ Affected versions of the product
- ▶ Workarounds or mitigations
- ▶ Remediations or fixes

Before publishing a security bulletin for a product, information about whether a product is affected by it is not available. In some situations, with much public attention, a security bulletin might be published indicating that a product is not affected by a particular vulnerability, but in most cases that is not done.

Finding the security bulletins

The security bulletins are published on the <http://www.ibm.com> website. You can choose to be notified when new security bulletins are published, or you can go to <http://www.ibm.com> and review content from the perspective of security for all products or the perspective of all support content for a product. However you prefer to find the needed security bulletins, becoming familiar with the following websites helps.

IBM Product Security Response Team blog

The IBM Product Security Response Team blog is a good place to start. It consolidates the security bulletins from all IBM products into a common location and provides links to other related security resources. To check out the blog, open a web browser and go to the following website:

<http://www.ibm.com/blogs/PSIRT>

IBM Support portal

Additionally, each product's support portal provides a section that is called *Flashes, alerts and bulletins* where you can browse the security bulletins for that product. For Spectrum Control, the support portal can be found at the following website:

http://www.ibm.com/support/entry/portal/product/tivoli/tivoli_storage_productivity_center

For other IBM products, open the IBM Support portal and select your product at the following website:

<http://www.ibm.com/support/entry/portal/support>

The flashes, alerts, and bulletins that are listed on the support portal can be filtered by version or operating system to help you locate the appropriate content.

My Notifications

You can also choose to be notified immediately of security bulletins and flashes by subscribing to My Notifications for selected products. Sign up for notifications by going to the following website:

<http://www.ibm.com/software/support/einfo.html>

Resolving security vulnerabilities for Spectrum Control

If Spectrum Control is affected by a security vulnerability or a set of vulnerabilities, a security bulletin is published after fixes or a workaround to mitigate the issue are available. The security bulletin covers all of the actions that must be taken to resolve the vulnerabilities. These actions cover Spectrum Control and any affected internal components. The Spectrum Control bulletins might mention but will not go into detail for other products that integrate with it. Depending on the product version, the vulnerability, and the components of Spectrum Control that you are using, a separate set of actions to follow after applying a fix might be required.

Because each vulnerability is unique, review the security bulletin thoroughly to identify whether you are affected and what actions to take.

Vulnerabilities for other products

Security bulletins for Spectrum Control cover vulnerabilities for it plus any internal components, but what about other software products that are bundled with Spectrum Control or integrated with it? Here are a few examples of products you might be using and must maintain:

- ▶ IBM DB2
- ▶ Jazz for Service Management
- ▶ IBM Tivoli Common Monitoring
- ▶ IBM WebSphere Application Server (used with Jazz for Service Management)
- ▶ Copy Services Manager
- ▶ IBM Spectrum Protect Snapshot
- ▶ IBM Tivoli Monitoring

The answer is the same for all of them and is similar to what you do for Spectrum Control: For those products, you should reference the IBM Product Response Team blog or the individual product support portals for any published security bulletins. Each product team is responsible for publishing their own security bulletins. Review any published security bulletins to identify affected versions and follow the remediation or fix information. If these products are affected, they might provide a maintenance fix. It is acceptable to apply a fix pack to these products to resolve security vulnerabilities or other problems. Upgrading to a new version or release is not typically supported by Spectrum Control and might result in compatibility issues. Even with fix pack updates, read the security bulletins carefully for any compatibility concerns. For example, if a common protocol is disabled, both products must accommodate a new protocol.

Vulnerabilities with the operating systems

Some security vulnerabilities are found at the operating system level. To confirm and resolve any security vulnerabilities for them, see the operating system vendor websites and their security information.

Reporting a security vulnerability

Customers and other entitled users of Spectrum Control, or any of the bundled products, should contact IBM technical support to report security issues. If the IBM technical support team determines that a reported issue is a security vulnerability, it informs IBM PSIRT and that sets the collaboration process in motion.

Other individuals or groups can report security vulnerabilities by using information that is available at the following website:

<http://www.ibm.com/security/secure-engineering/report.html>



Reporting scripts

This appendix provides three different versions of an example script to create a showback report.

For considerations regarding a showback report and for more information about how to use and modify the script, see 5.1.4, “Creating input for your chargeback report” on page 126.

This appendix contains the following reporting scripts:

- ▶ Sample script with native SQL statement
- ▶ runshowbackexport.bat
- ▶ showbackexport.sql
- ▶ Sample script for Cognos

Overview

The following sections show three different versions of an example script that creates a showback report:

- ▶ In “Sample script with native SQL statement” on page 244, you find the native SQL statement that is used in all versions. This native SQL statement can be used to run natively and in applications that support DB2 SQL, for example, Microsoft Excel 2013.
- ▶ If you want to run the showback report regularly as a batch report, see “runshowbackexport.bat” on page 249 and “showbackexport.sql” on page 249.
- ▶ You can use “Sample script for Cognos” on page 253 to run the showback report in Cognos.

Sample script with native SQL statement

Example B-1 shows the showback.sql script.

Example B-1 showback.sql example

with t1 as (

```
SELECT cast(AL1.name as VARCHAR(70)) as volid, Strip(AL2.UNKNOWN_HOST_NAME) as
hostname
FROM TPCREPORT.STORAGEVOLUME AL1 inner join TPCREPORT.STORAGEVOLUME_PATH AL2 on
(AL1.STORAGE_VOLUME_ID=AL2.STORAGE_VOLUME_ID)
and (AL1.BACKEND_DISK_ID is NULL) and (AL1.DETECTABLE = 'True')
```

group by AL1.name, AL2.UNKNOWN_HOST_NAME),

```
t2 as (
SELECT e.volid,e.hostname,
XMLELEMENT(
  NAME "volid", XMLELEMENT(
    NAME "volid", e.volid
  ),
  XMLELEMENT(
    NAME "hostname", e.hostname
  )
)
AS "Result"
FROM t1 E
),
```

```
t3 as (
SELECT volid, rtrim(
replace (
replace (
  XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME "hostname", t2.hostname)
    ORDER BY t2.hostname
  )
  as varchar (5000)),
```



```

        '<hostname>',''), '</hostname>',''),';')
    AS hostname
  FROM t2
  group by volid
),

tnum as (
select count(hostname) as numservers, volid from t1  group by volid
),

t4 as (
SELECT
AL1.STORAGE_VOLUME_ID
, cast(AL1.name as VARCHAR(70)) as volid
, AL1.ELEMENT_NAME as Volumenname
, AL1.SPACE/ (1024*1024*1024) as Capacity
, AL1.REAL_SPACE / (1024*1024*1024) as Allocated_Space
, AL1.CONSUMED_SPACE/ (1024*1024*1024) as Used_Space
, AL2.POOL_ID
, AL2.Display_NAME as Poolname
, AL2.UDP1
, AL2.UDP2
, AL2.UDP3
, AL3.DISPLAY_NAME as Storagesubsystemname1
, AL3.UDP1 as StorageSubsystem_Custom_Tag1
, AL3.UDP2 as StorageSubsystem_Custom_Tag2
, AL3.UDP3 as StorageSubsystem_Custom_Tag3
, case
  when AL3.User_Provided_Name is not null then AL3.User_Provided_Name
  else
    case
      when AL3.TYPE = 'FASTT' then substr(AL3.DISPLAY_NAME,1, length
(AL3.DISPLAY_NAME)-37)
      when AL3.TYPE = 'DS4000' then
substr(AL3.DISPLAY_NAME,8,posstr(substr(AL3.DISPLAY_NAME,8,
length(AL3.DISPLAY_NAME)),'-')-1)
      when AL3.TYPE = 'DS5000' then
substr(AL3.DISPLAY_NAME,8,posstr(substr(AL3.DISPLAY_NAME,8,
length(AL3.DISPLAY_NAME)),'-')-1)
      else AL3.DISPLAY_NAME
    end
  end as Storagesubsystemname
, AL1.SPACE_EFFICIENT_VOLUME
, AL1.ENCRYPTABLE
, AL1.ENCRYPTED
, AL1.MIRROR_COUNT
, AL1.OPERATIONAL_STATUS
, AL1.DETECTABLE
, AL1.RAID_LEVEL
, AL1.IS_COMPRESSED
, AL4.POOL_TIER_ID as TPC_TIER
, AL1.BLOCK_SIZE

```

```

FROM TPCREPORT.STORAGEVOLUME AL1 LEFT OUTER JOIN TPC.T_RES_VDISK_COPY AL6 ON
(AL1.STORAGE_VOLUME_ID=AL6.SVID)
and (AL1.DETECTABLE = 'True')and (AL1.BACKEND_DISK_ID is NULL)
left outer join TPCREPORT.STORAGESUBSYSTEM AL3 on
AL1.SUBSYSTEM_ID=AL3.SUBSYSTEM_ID
left outer join TPCREPORT.STORAGEPOOL AL2 on AL6.POOL_ID=AL2.POOL_ID
left outer join TPC.T_RES_STORAGE_POOL AL4 on AL2.POOL_ID=AL4.POOL_ID
-- and AL6.IS_PRIMARY =1
),

```

```

tt as (
select
AL2.POOL_ID
,AL5.EASY_TIER as Pool_EasyTier
,AL5.EASY_TIER_STATUS as Pool_EasyTier_Status
from TPCREPORT.STORAGEPOOL AL2
left outer join TPCREPORT.StorageSubsystem_Tiering_Info AL5 on
AL5.Entity_ID=AL2.POOL_ID
group by AL2.POOL_ID, AL5.EASY_TIER, AL5.EASY_TIER_Status
),

```

```

t5 as (
SELECT
AL2.name as appgroup
, cast(AL3.name as VARCHAR(70)) as volid
FROM TPC.T_GROUP2RES AL1, TPC.T_GROUP AL2, TPCREPORT.STORAGEVOLUME AL3
WHERE (AL2.GROUP_ID=AL1.GROUP_ID AND AL1.RESOURCE_ID=AL3.STORAGE_VOLUME_ID) AND
(AL1.RES_TYPE=194 AND AL2.FLAGS=6)
group by AL3.name, AL2.name ),

```

```

t6 as (
SELECT a.volid, a.appgroup,
XMLELEMENT(
NAME "volid", XMLELEMENT(
NAME "volid", a.volid
),
XMLELEMENT(
NAME "appgroup", a.appgroup
)
)
AS "Appresult"
FROM t5 A
),

```

```

t7 as (
SELECT volid, rtrim(
replace (
replace (
XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME "appgroup", t6.appgroup)
ORDER BY t6.appgroup
)
as varchar (500)),
'<appgroup>', '' ), '</appgroup>', '' ), '' );

```

```

AS appgroup
FROM t6
group by volid
),

t8 as (
SELECT
  cast(AL1.name as VARCHAR(70)) as volid
  , AL2.TIER || ': ' || (cast(AL2.TIER_SPACE *512/(1024*1024*1024) as
varchar(10))) as Tier_Capacity
  FROM TPCREPORT.STORAGEVOLUME AL1, TPCREPORT.STORAGESUBSYSTEM_TIERING_INFO AL2
  WHERE (AL1.STORAGE_VOLUME_ID=AL2.ENTITY_ID)
  group by AL1.name,AL2.TIER, AL2.TIER_SPACE
),

t9 as (
SELECT e.volid,e.Tier_Capacity,
XMLELEMENT(
  NAME "volid", XMLELEMENT(
    NAME "volid", e.volid
  ),
  XMLELEMENT(
    NAME "Tier_Capacity", e.Tier_Capacity
  )
)
AS "Result"
FROM t8 E
),

t10 as (
SELECT volid, rtrim(
replace (
replace (
  XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME "Tier_Capacity", t9.Tier_Capacity)
    ORDER BY t9.Tier_Capacity
  )
  as varchar (500)),
  '<Tier_Capacity>','), '</Tier_Capacity>','),';')
AS Tier_Capacity
FROM t9
group by volid
)

select
t4.Volumename
, t4.volid
,case
  when tnum.numservers = 0 then 'No Hosts mapped'
  else substr (t3.hostname,1,length(t3.hostname)- 1)
  end as Hostlist
, tnum.numservers as NUMHosts
-- , case

```

```

-- when (t4.StorageSubsystemname1 like 'SVC%') or (t4.StorageSubsystemname1 like
'Storwize%') then t4.MIRROR_COUNT
-- else 1
-- end as MIRROR_COUNT
, t4.Capacity
, t4.Allocated_Space
, t4.Used_Space
, t4.Pool_ID
, t4.Poolname
, t4.UDP1 as Pool_Custom_Tag1
, t4.UDP2 as Pool_Custom_Tag2
, t4.UDP3 as Pool_Custom_Tag3
, t4.StorageSubsystemname as Storage_System
, t4.SPACE_EFFICIENT_VOLUME as ThinProvisioned
, t4.ENCRYPTABLE
, t4.ENCRYPTED
, t4.OPERATIONAL_STATUS as Status
, case when t4.RAID_LEVEL = -1 then 'unavailable' else cast(t4.RAID_LEVEL as
varchar(10)) end as Raid_level
-- , t4.RAID_LEVEL
, t4.IS_COMPRESSED
, t4.TPC_TIER as SpectrumControl_Tier
-- , t4.BLOCK_SIZE
, tt.Pool_EasyTier
, tt.Pool_EasyTier_Status
, t4.StorageSubsystem_Custom_Tag1
, t4.StorageSubsystem_Custom_Tag2
, t4.StorageSubsystem_Custom_Tag3
, substr (t10.Tier_Capacity,1,length(t10.Tier_Capacity)- 1) as Vol_Tier_Capacity
, substr (t7.appgroup,1,length(t7.appgroup)- 1) as Applist
, current_date as Current_date
from t3 inner join t4 on (t3.volid = t4.volid)
inner join tnum on (t3.volid = tnum.volid)
left outer join t7 on (t4.volid = t7.volid) left outer join t10 on (t4.volid
= t10.volid)
left outer join tt on (t4.pool_id = tt.pool_id)
with ur for read only

```

Sample script to run as a batch report

If you want to run the showback report regularly as a batch report, use the scripts that are shown in “runshowbackexport.bat” on page 249 and “showbackexport.sql” on page 249.

runshowbackexport.bat

Example B-2 shows the .bat file that is used to run “showbackexport.sql” on page 249.

Example B-2 The .bat file to run runshowbackexport

```
if "%DB2CLP%" == "" db2cmd /c /i /w %0 %* & goto :eof
db2 connect to tpcdb user db2admin using pwd
db2 -tvf db2 -tvf
C:\showback\showbackexport.sql
```

showbackexport.sql

Example B-3 on page 249 shows the showbackexport.sql script, which exports the results to a batch report.

Example B-3 Export the showback.sql results to a CSV file

```
EXPORT TO "C:\showback\showback.csv" OF DEL MODIFIED BY COLDEL; MESSAGES
"C:\showback\showback.log"
```

```
with t1 as (
```

```
SELECT cast(AL1.name as VARCHAR(70)) as volid, Strip(AL2.UNKNOWN_HOST_NAME) as
hostname
FROM TPCREPORT.STORAGEVOLUME AL1 inner join TPCREPORT.STORAGEVOLUME_PATH AL2 on
(AL1.STORAGE_VOLUME_ID=AL2.STORAGE_VOLUME_ID)
and (AL1.BACKEND_DISK_ID is NULL) and (AL1.DETECTABLE = 'True')
```

```
group by AL1.name, AL2.UNKNOWN_HOST_NAME ),
```

```
t2 as (
SELECT e.volid,e.hostname,
XMLELEMENT(
NAME "volid", XMLELEMENT(
NAME "volid", e.volid
),
XMLELEMENT(
NAME "hostname", e.hostname
)
)
AS "Result"
FROM t1 E
),
```

```
t3 as (
SELECT volid, rtrim(
replace (
replace (
XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME "hostname", t2.hostname)
ORDER BY t2.hostname
)
)
```

```

        as varchar (5000)),
        '<hostname>',''), '</hostname>',''),';')
    AS hostname
FROM t2
group by volid
),

tnum as (
select count(hostname) as numservers, volid from t1  group by volid
),

t4 as (
SELECT
AL1.STORAGE_VOLUME_ID
, cast(AL1.name as VARCHAR(70)) as volid
, AL1.ELEMENT_NAME as Volumenname
, AL1.SPACE/ (1024*1024*1024) as Capacity
, AL1.REAL_SPACE / (1024*1024*1024) as Allocated_Space
, AL1.CONSUMED_SPACE/ (1024*1024*1024) as Used_Space
, AL2.POOL_ID
, AL2.Display_NAME as Poolname
, AL2.UDP1
, AL2.UDP2
, AL2.UDP3
, AL3.DISPLAY_NAME as Storagesubsystemname1
, AL3.UDP1 as StorageSubsystem_Custom_Tag1
, AL3.UDP2 as StorageSubsystem_Custom_Tag2
, AL3.UDP3 as StorageSubsystem_Custom_Tag3
, case
when AL3.User_Provided_Name is not null then AL3.User_Provided_Name
else
case
when AL3.TYPE = 'FASTT' then substr(AL3.DISPLAY_NAME,1, length
(AL3.DISPLAY_NAME)-37)
when AL3.TYPE = 'DS4000' then
substr(AL3.DISPLAY_NAME,8,posstr(substr(AL3.DISPLAY_NAME,8,
length(AL3.DISPLAY_NAME)),'-')-1)
when AL3.TYPE = 'DS5000' then
substr(AL3.DISPLAY_NAME,8,posstr(substr(AL3.DISPLAY_NAME,8,
length(AL3.DISPLAY_NAME)),'-')-1)
else AL3.DISPLAY_NAME
end
end as Storagesubsystemname
, AL1.SPACE_EFFICIENT_VOLUME
, AL1.ENCRYPTABLE
, AL1.ENCRYPTED
, AL1.MIRROR_COUNT
, AL1.OPERATIONAL_STATUS
, AL1.DETECTABLE
, AL1.RAID_LEVEL
, AL1.IS_COMPRESSED
, AL4.POOL_TIER_ID as TPC_TIER
, AL1.BLOCK_SIZE

```

```

FROM TPCREPORT.STORAGEVOLUME AL1 LEFT OUTER JOIN TPC.T_RES_VDISK_COPY AL6 ON
(AL1.STORAGE_VOLUME_ID=AL6.SVID)
and (AL1.DETECTABLE = 'True')and (AL1.BACKEND_DISK_ID is NULL)
left outer join TPCREPORT.STORAGESUBSYSTEM AL3 on
AL1.SUBSYSTEM_ID=AL3.SUBSYSTEM_ID
left outer join TPCREPORT.STORAGEPOOL AL2 on AL6.POOL_ID=AL2.POOL_ID
left outer join TPC.T_RES_STORAGE_POOL AL4 on AL2.POOL_ID=AL4.POOL_ID
-- and AL6.IS_PRIMARY =1
),

```

```

tt as (
select
AL2.POOL_ID
,AL5.EASY_TIER as Pool_EasyTier
,AL5.EASY_TIER_STATUS as Pool_EasyTier_Status
from TPCREPORT.STORAGEPOOL AL2
left outer join TPCREPORT.StorageSubsystem_Tiering_Info AL5 on
AL5.Entity_ID=AL2.POOL_ID
group by AL2.POOL_ID, AL5.EASY_TIER, AL5.EASY_TIER_Status
),

```

```

t5 as (
SELECT
AL2.name as appgroup
, cast(AL3.name as VARCHAR(70)) as volid
FROM TPC.T_GROUP2RES AL1, TPC.T_GROUP AL2, TPCREPORT.STORAGEVOLUME AL3
WHERE (AL2.GROUP_ID=AL1.GROUP_ID AND AL1.RESOURCE_ID=AL3.STORAGE_VOLUME_ID) AND
(AL1.RES_TYPE=194 AND AL2.FLAGS=6)
group by AL3.name, AL2.name ),

```

```

t6 as (
SELECT a.volid, a.appgroup,
XMLELEMENT(
NAME "volid", XMLELEMENT(
NAME "volid", a.volid
),
XMLELEMENT(
NAME "appgroup", a.appgroup
)
)
AS "Appresult"
FROM t5 A
),

```

```

t7 as (
SELECT volid, rtrim(
replace (
replace (
XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME "appgroup", t6.appgroup)
ORDER BY t6.appgroup
)
as varchar (500)),
'<appgroup>', '' ), '</appgroup>', '' ), '' );

```

```

AS appgroup
FROM t6
group by volid
),

t8 as (
SELECT
  cast(AL1.name as VARCHAR(70)) as volid
  , AL2.TIER || ': ' || (cast(AL2.TIER_SPACE *512/(1024*1024*1024) as
varchar(10))) as Tier_Capacity
  FROM TPCREPORT.STORAGEVOLUME AL1, TPCREPORT.STORAGESUBSYSTEM_TIERING_INFO AL2
  WHERE (AL1.STORAGE_VOLUME_ID=AL2.ENTITY_ID)
  group by AL1.name,AL2.TIER, AL2.TIER_SPACE
),

t9 as (
SELECT e.volid,e.Tier_Capacity,
XMLELEMENT(
  NAME "volid", XMLELEMENT(
    NAME "volid", e.volid
  ),
  XMLELEMENT(
    NAME "Tier_Capacity", e.Tier_Capacity
  )
)
AS "Result"
FROM t8 E
),

t10 as (
SELECT volid, rtrim(
replace (
replace (
  XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME "Tier_Capacity", t9.Tier_Capacity)
    ORDER BY t9.Tier_Capacity
  )
  as varchar (500)),
  '<Tier_Capacity>',''), '</Tier_Capacity>',''),';')
AS Tier_Capacity
FROM t9
group by volid
)

select
t4.Volumename
, t4.volid
,case
  when tnum.numservers = 0 then 'No Hosts mapped'
  else substr (t3.hostname,1,length(t3.hostname)- 1)
  end as Hostlist
, tnum.numservers as NUMHosts
-- , case

```



```

-- when (t4.StorageSubsystemname1 like 'SVC%') or (t4.StorageSubsystemname1 like
'Storwize%') then t4.MIRROR_COUNT
-- else 1
-- end as MIRROR_COUNT
, t4.Capacity
, t4.Allocated_Space
, t4.Used_Space
, t4.Pool_ID
, t4.Poolname
, t4.UDP1 as Pool_Custom_Tag1
, t4.UDP2 as Pool_Custom_Tag2
, t4.UDP3 as Pool_Custom_Tag3
, t4.StorageSubsystemname as Storage_System
, t4.SPACE_EFFICIENT_VOLUME as ThinProvisioned
, t4.ENCRYPTABLE
, t4.ENCRYPTED
, t4.OPERATIONAL_STATUS as Status
, case when t4.RAID_LEVEL = -1 then 'unavailable' else cast(t4.RAID_LEVEL as
varchar(10)) end as Raid_level
-- , t4.RAID_LEVEL
, t4.IS_COMPRESSED
, t4.TPC_TIER as SpectrumControl_Tier
-- , t4.BLOCK_SIZE
, tt.Pool_EasyTier
, tt.Pool_EasyTier_Status
, t4.StorageSubsystem_Custom_Tag1
, t4.StorageSubsystem_Custom_Tag2
, t4.StorageSubsystem_Custom_Tag3
, substr (t10.Tier_Capacity,1,length(t10.Tier_Capacity)- 1) as Vol_Tier_Capacity
, substr (t7.appgroup,1,length(t7.appgroup)- 1) as Applist
, current_date as Current_date
from t3 inner join t4 on (t3.volid = t4.volid)
inner join tnum on (t3.volid = tnum.volid)
left outer join t7 on (t4.volid = t7.volid) left outer join t10 on (t4.volid
= t10.volid)
left outer join tt on (t4.pool_id = tt.pool_id)
with ur for read only

;

```

Sample script for Cognos

Example B-4 shows how you can use the showback_Cognos_xml.txt script to run the generate a showback report in Cognos.

Example B-4 Cognos XML showback script example

```

<report xmlns="http://developer.cognos.com/schemas/report/9.0/" useStyleVersion="10"
expressionLocale="en-gb">
    <modelPath>/content/folder[@name='IBM Spectrum Control
Packages']/package[@name='Capacity and
Relationships']/model[@name='2015-12-03T23:00:55.249Z']/</modelPath>
    <drillBehavior/>

```

```

        <layouts>
        <layout>
        <reportPages>
        <page name="Page1"><style><defaultStyles><defaultStyle
refStyle="pg"/></defaultStyles></style>
        <pageBody><style><defaultStyles><defaultStyle
refStyle="pb"/></defaultStyles></style>
        <contents>

        <list horizontalPagination="true" refQuery="Query2"
name="List1">

        <noDataHandler>
        <contents>
        <block>
        <contents>
        <textItem>
        <dataSource>
        <staticValue>No Data Available</staticValue>
        </dataSource>
        <style>
        <CSS value="padding:10px 18px;"/>
        </style>
        </textItem>
        </contents>
        </block>
        </contents>
        </noDataHandler>
        <style>
        <CSS value="border-collapse:collapse"/>
        <defaultStyles>
        <defaultStyle refStyle="ls"/>
        </defaultStyles>
        </style>
        <listColumns><listColumn><listColumnTitle><style><defaultStyles><defaultStyle
refStyle="lt"/></defaultStyles></style><contents><textItem><dataSource><dataItemLabel
refDataItem="VOLUMENAME"/></dataSource></textItem></contents></listColumnTitle><listColumnB
ody><style><defaultStyles><defaultStyle
refStyle="lm"/></defaultStyles></style><contents><textItem><dataSource><dataItemValue
refDataItem="VOLUMENAME"/></dataSource></textItem></contents></listColumnBody></listColumn>
<listColumn><listColumnTitle><style><defaultStyles><defaultStyle
refStyle="lt"/></defaultStyles></style><contents><textItem><dataSource><dataItemLabel
refDataItem="VOLID"/></dataSource></textItem></contents></listColumnTitle><listColumnBody><
style><defaultStyles><defaultStyle
refStyle="lm"/></defaultStyles></style><contents><textItem><dataSource><dataItemValue
refDataItem="VOLID"/></dataSource></textItem></contents></listColumnBody></listColumn><list
Column><listColumnTitle><style><defaultStyles><defaultStyle
refStyle="lt"/></defaultStyles></style><contents><textItem><dataSource><dataItemLabel
refDataItem="HOSTLIST"/></dataSource></textItem></contents></listColumnTitle><listColumnBod
y><style><defaultStyles><defaultStyle
refStyle="lm"/></defaultStyles></style><contents><textItem><dataSource><dataItemValue
refDataItem="HOSTLIST"/></dataSource></textItem></contents></listColumnBody></listColumn><list
Column><listColumnTitle><style><defaultStyles><defaultStyle
refStyle="lt"/></defaultStyles></style><contents><textItem><dataSource><dataItemLabel
refDataItem="NUMHOSTS"/></dataSource></textItem></contents></listColumnTitle><listColumnBod
y><style><defaultStyles><defaultStyle
refStyle="lm"/></defaultStyles></style><contents><textItem><dataSource><dataItemValue
refDataItem="NUMHOSTS"/></dataSource></textItem></contents></listColumnBody></listColumn><list

```


256 IBM Spectrum Family: IBM Spectrum Control Standard Edition

```

Body></listColumn><listColumn><listColumnTitle><style><defaultStyles><defaultStyle
refStyle="lt"/></defaultStyles></style><contents><textItem><dataSource><dataItemLabel
refDataItem="VOL_TIER_CAPACITY"/></dataSource></textItem></contents></listColumnTitle><list
ColumnBody><style><defaultStyles><defaultStyle
refStyle="lm"/></defaultStyles></style><contents><textItem><dataSource><dataItemValue
refDataItem="VOL_TIER_CAPACITY"/></dataSource></textItem></contents></listColumnBody></list
Column><listColumn><listColumnTitle><style><defaultStyles><defaultStyle
refStyle="lt"/></defaultStyles></style><contents><textItem><dataSource><dataItemLabel
refDataItem="APPLIST"/></dataSource></textItem></contents></listColumnTitle><listColumnBody
><style><defaultStyles><defaultStyle
refStyle="lm"/></defaultStyles></style><contents><textItem><dataSource><dataItemValue
refDataItem="APPLIST"/></dataSource></textItem></contents></listColumnBody></listColumn><li
stColumn><listColumnTitle><style><defaultStyles><defaultStyle
refStyle="lt"/></defaultStyles></style><contents><textItem><dataSource><dataItemLabel
refDataItem="CURRENT_DATE"/></dataSource></textItem></contents></listColumnTitle><listColumn
nBody><style><defaultStyles><defaultStyle
refStyle="lm"/></defaultStyles></style><contents><textItem><dataSource><dataItemValue
refDataItem="CURRENT_DATE"/></dataSource></textItem></contents></listColumnBody></listColumn
n></listColumns></list></contents>
</pageBody>
<pageHeader>
<contents>
<block><style><defaultStyles><defaultStyle
refStyle="ta"/></defaultStyles></style>
<contents>
<textItem><style><defaultStyles><defaultStyle
refStyle="tt"/></defaultStyles></style>
<dataSource>
<staticValue/>
</dataSource>
</textItem>
</contents>
</block>
</contents>
<style>
<defaultStyles>
<defaultStyle refStyle="ph"/>
</defaultStyles>
<CSS value="padding-bottom:10px"/>
</style>
</pageHeader>
<pageFooter>
<contents>
<table>
<tableRows>
<tableRow>
<tableCells>
<tableCell>
<contents>
<date>
<style>
<dataFormat>
<dateFormat/>
</dataFormat>
</style>
</date>
</contents>
<style>
<CSS
value="vertical-align:top;text-align:left;width:25%"/>

```

```

        </style>
      </tableCell>
    <tableCell>
      <contents>
        <pageNumber/>
      </contents>
    <style>
      <CSS
value="vertical-align:top;text-align:center;width:50%"/>
    </style>
  </tableCell>
  <tableCell>
    <contents>
      <time>
        <style>
          <dataFormat>
            <timeFormat/>
          </dataFormat>
        </style>
      </time>
    </contents>
    <style>
      <CSS
value="vertical-align:top;text-align:right;width:25%"/>
    </style>
  </tableCell>
</tableCells>
</tableRow>
</tableRows>
<style>
  <defaultStyles>
    <defaultStyle refStyle="tb"/>
  </defaultStyles>
  <CSS value="border-collapse:collapse;width:100%"/>
</style>
</table>
</contents>
<style>
  <defaultStyles>
    <defaultStyle refStyle="pf"/>
  </defaultStyles>
  <CSS value="padding-top:10px"/>
</style>
</pageFooter>
</page>
</reportPages>
</layout>
</layouts>
  <XMLAttributes><XMLAttribute name="RS_CreateExtendedDataItems" value="true"
output="no"/><XMLAttribute name="listSeparator" value="," output="no"/><XMLAttribute
name="RS_modelModificationTime" value="2015-12-03T23:00:55.343Z"
output="no"/></XMLAttributes><queries><query name="Query2">
  <source>

    <sqlQuery name="SQL2" dataSource="TPCDB" type="passThrough">
    <sqlText>with t1 as (

SELECT cast(AL1.name as VARCHAR(70)) as valid, Strip(AL2.UNKNOWN_HOST_NAME) as hostname
FROM TPCREPORT.STORAGEVOLUME AL1 inner join TPCREPORT.STORAGEVOLUME_PATH AL2 on
(AL1.STORAGE_VOLUME_ID=AL2.STORAGE_VOLUME_ID)

```

```
and (AL1.BACKEND_DISK_ID is NULL) and (AL1.DETECTABLE = &apos;True&apos;);
```

```
group by AL1.name, AL2.UNKNOWN_HOST_NAME ),
```

```
t2 as (
SELECT e.volid,e.hostname,
XMLELEMENT(
NAME &quot;volid&quot;;, XMLELEMENT(
NAME &quot;volid&quot;;, e.volid
),
XMLELEMENT(
NAME &quot;hostname&quot;;, e.hostname
)
)
AS &quot;Result&quot;
FROM t1 E
),
```

```
t3 as (
SELECT volid, rtrim(
replace (
replace (
XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME &quot;hostname&quot;;, t2.hostname)
ORDER BY t2.hostname
)
as varchar (5000)),
&apos;&lt;hostname&gt;&apos;;&apos;&apos;)),
&apos;&lt;/hostname&gt;&apos;;&apos;&apos;,&apos;&apos;,&apos;&apos;,&apos;&apos;,&apos;&apos;,&apos;&apos;))
AS hostname
FROM t2
group by volid
),
```

```
tnum as (
select count(hostname) as numservers, volid from t1 group by volid
),
```

```
t4 as (
SELECT
AL1.STORAGE_VOLUME_ID
, cast(AL1.name as VARCHAR(70)) as volid
, AL1.ELEMENT_NAME as Volumenname
, AL1.SPACE/ (1024*1024*1024) as Capacity
, AL1.REAL_SPACE / (1024*1024*1024) as Allocated_Space
, AL1.CONSUMED_SPACE/ (1024*1024*1024) as Used_Space
, AL2.POOL_ID
, AL2.Display_NAME as Poolname
, AL2.UDP1
, AL2.UDP2
, AL2.UDP3
, AL3.DISPLAY_NAME as Storagesubsystemname1
, AL3.UDP1 as StorageSubsystem_Custom_Tag1
, AL3.UDP2 as StorageSubsystem_Custom_Tag2
, AL3.UDP3 as StorageSubsystem_Custom_Tag3
, case
when AL3.User_Provided_Name is not null then AL3.User_Provided_Name
```

```

else
  case
    when AL3.TYPE = 'FAST' then substr(AL3.DISPLAY_NAME,1, length
(AL3.DISPLAY_NAME)-37)
    when AL3.TYPE = 'DS4000' then
substr(AL3.DISPLAY_NAME,8,posstr(substr(AL3.DISPLAY_NAME,8,
length(AL3.DISPLAY_NAME)),'-')-1)
    when AL3.TYPE = 'DS5000' then
substr(AL3.DISPLAY_NAME,8,posstr(substr(AL3.DISPLAY_NAME,8,
length(AL3.DISPLAY_NAME)),'-')-1)
    else AL3.DISPLAY_NAME
  end
end as Storagesubsystemname
, AL1.SPACE_EFFICIENT_VOLUME
, AL1.ENCRYPTABLE
, AL1.ENCRYPTED
, AL1.MIRROR_COUNT
, AL1.OPERATIONAL_STATUS
, AL1.DETECTABLE
, AL1.RAID_LEVEL
, AL1.IS_COMPRESSED
, AL4.POOL_TIER_ID as TPC_TIER
, AL1.BLOCK_SIZE
FROM TPCREPORT.STORAGEVOLUME AL1 LEFT OUTER JOIN TPC.T_RES_VDISK_COPY AL6 ON
(AL1.STORAGE_VOLUME_ID=AL6.SVID)
and (AL1.DETECTABLE = 'True')and (AL1.BACKEND_DISK_ID is NULL)
left outer join TPCREPORT.STORAGESUBSYSTEM AL3 on AL1.SUBSYSTEM_ID=AL3.SUBSYSTEM_ID
left outer join TPCREPORT.STORAGEPOOL AL2 on AL6.POOL_ID=AL2.POOL_ID
left outer join TPC.T_RES_STORAGE_POOL AL4 on AL2.POOL_ID=AL4.POOL_ID
-- and AL6.IS_PRIMARY =1
),

tt as (
select
AL2.POOL_ID
,AL5.EASY_TIER as Pool_EasyTier
,AL5.EASY_TIER_STATUS as Pool_EasyTier_Status
from TPCREPORT.STORAGEPOOL AL2
left outer join TPCREPORT.Storagesubsystem_Tiering_Info AL5 on AL5.Entity_ID=AL2.POOL_ID
group by AL2.POOL_ID, AL5.EASY_TIER, AL5.EASY_TIER_Status
),

t5 as (
SELECT
AL2.name as appgroup
, cast(AL3.name as VARCHAR(70)) as volid
FROM TPC.T_GROUP2RES AL1, TPC.T_GROUP AL2, TPCREPORT.STORAGEVOLUME AL3
WHERE (AL2.GROUP_ID=AL1.GROUP_ID AND AL1.RESOURCE_ID=AL3.STORAGE_VOLUME_ID) AND
(AL1.RES_TYPE=194 AND AL2.FLAGS=6)
group by AL3.name, AL2.name ),

t6 as (
SELECT a.volid, a.appgroup,
XMLEMENT(
NAME 'volid', XMLEMENT(
NAME 'volid', a.volid
),
XMLEMENT(

```



```

)
)
AS 'Appresult';
FROM t5 A
),

t7 as (
SELECT valid, rtrim(
replace (
replace (
XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME 'appgroup', t6.appgroup)
ORDER BY t6.appgroup
)
as varchar (500)),
'<appgroup>','</appgroup>'),
'</appgroup>','</appgroup>'),
AS appgroup
FROM t6
group by valid
),

t8 as (
SELECT
cast(AL1.name as VARCHAR(70)) as valid
, AL2.TIER || ' ' || (cast(AL2.TIER_SPACE *512/(1024*1024*1024) as
varchar(10))) as Tier_Capacity
FROM TPCREPORT.STORAGEVOLUME AL1, TPCREPORT.STORAGESUBSYSTEM_TIERING_INFO AL2
WHERE (AL1.STORAGE_VOLUME_ID=AL2.ENTITY_ID)
group by AL1.name,AL2.TIER, AL2.TIER_SPACE
),

t9 as (
SELECT e.valid,e.Tier_Capacity,
XMLELEMENT(
NAME 'valid', XMLELEMENT(
NAME 'valid', e.valid
),
XMLELEMENT(
NAME 'Tier_Capacity', e.Tier_Capacity
)
)
AS 'Result';
FROM t8 E
),

t10 as (
SELECT valid, rtrim(
replace (
replace (
XMLSERIALIZE(CONTENT XMLAGG(XMLELEMENT(NAME 'Tier_Capacity', t9.Tier_Capacity)
ORDER BY t9.Tier_Capacity
)
as varchar (500)),
'<Tier_Capacity>','</Tier_Capacity>'),
'</Tier_Capacity>','</Tier_Capacity>'),
AS Tier_Capacity
FROM t9

```

```

group by volid
)

select
t4.Volumename
, t4.volid
, case
  when tnum.numservers = 0 then 'No Hosts mapped';
  else substr (t3.hostname,1,length(t3.hostname)- 1)
  end as Hostlist
, tnum.numservers as NUMHosts
-- , case
-- when (t4.StorageSubsystemname1 like 'SVC') or (t4.StorageSubsystemname1 like
'&apos;Storwize%&apos;') then t4.MIRROR_COUNT
-- else 1
-- end as MIRROR_COUNT
, t4.Capacity
, t4.Allocated_Space
, t4.Used_Space
, t4.Pool_ID
, t4.Poolname
, t4.UDP1 as Pool_Custom_Tag1
, t4.UDP2 as Pool_Custom_Tag2
, t4.UDP3 as Pool_Custom_Tag3
, t4.StorageSubsystemname as Storage_System
, t4.SPACE_EFFICIENT_VOLUME as ThinProvisioned
, t4.ENCRYPTABLE
, t4.ENCRYPTED
, t4.OPERATIONAL_STATUS as Status
, case when t4.RAID_LEVEL = -1 then 'unavailable'; else cast(t4.RAID_LEVEL as
varchar(10)) end as Raid_level
-- , t4.RAID_LEVEL
, t4.IS_COMPRESSED
, t4.TPC_TIER as SpectrumControl_Tier
-- , t4.BLOCK_SIZE
, tt.Pool_EasyTier
, tt.Pool_EasyTier_Status
, t4.StorageSubsystem_Custom_Tag1
, t4.StorageSubsystem_Custom_Tag2
, t4.StorageSubsystem_Custom_Tag3
, substr (t10.Tier_Capacity,1,length(t10.Tier_Capacity)- 1) as Vol_Tier_Capacity
, substr (t7.appgroup,1,length(t7.appgroup)- 1) as Applist
, current_date as Current_date
from t3 inner join t4 on (t3.volid = t4.volid)
inner join tnum on (t3.volid = tnum.volid)
left outer join t7 on (t4.volid = t7.volid) left outer join t10 on (t4.volid =
t10.volid)
left outer join tt on (t4.pool_id = tt.pool_id)
with ur for read only</sqlText>
<mdProjectedItems><mdProjectedItem name="VOLUMENAME"/><mdProjectedItem
name="VOLID"/><mdProjectedItem name="HOSTLIST"/><mdProjectedItem
name="NUMHOSTS"/><mdProjectedItem name="CAPACITY"/><mdProjectedItem
name="ALLOCATED_SPACE"/><mdProjectedItem name="USED_SPACE"/><mdProjectedItem
name="POOL_ID"/><mdProjectedItem name="POOLNAME"/><mdProjectedItem
name="POOL_CUSTOM_TAG1"/><mdProjectedItem name="POOL_CUSTOM_TAG2"/><mdProjectedItem
name="POOL_CUSTOM_TAG3"/><mdProjectedItem name="STORAGE_SYSTEM"/><mdProjectedItem
name="THINPROVISIONED"/><mdProjectedItem name="ENCRYPTABLE"/><mdProjectedItem
name="ENCRYPTED"/><mdProjectedItem name="STATUS"/><mdProjectedItem
name="RAID_LEVEL"/><mdProjectedItem name="IS_COMPRESSED"/><mdProjectedItem

```

```

name="SPECTRUMCONTROL_TIER"/><mdProjectedItem name="POOL_EASYTIER"/><mdProjectedItem
name="POOL_EASYTIER_STATUS"/><mdProjectedItem
name="STORAGESUBSYSTEM_CUSTOM_TAG1"/><mdProjectedItem
name="STORAGESUBSYSTEM_CUSTOM_TAG2"/><mdProjectedItem
name="STORAGESUBSYSTEM_CUSTOM_TAG3"/><mdProjectedItem
name="VOL_TIER_CAPACITY"/><mdProjectedItem name="APPLIST"/><mdProjectedItem
name="CURRENT_DATE"/></mdProjectedItems></sqlQuery></source>
    <selection><dataItem
name="VOLUMENAME"><expression>[SQL2].[VOLUMENAME]</expression></dataItem><dataItem
name="VOLID"><expression>[SQL2].[VOLID]</expression></dataItem><dataItem
name="HOSTLIST"><expression>[SQL2].[HOSTLIST]</expression></dataItem><dataItem
name="NUMHOSTS"><expression>[SQL2].[NUMHOSTS]</expression></dataItem><dataItem
name="CAPACITY"><expression>[SQL2].[CAPACITY]</expression></dataItem><dataItem
name="ALLOCATED_SPACE"><expression>[SQL2].[ALLOCATED_SPACE]</expression></dataItem><dataItem
name="USED_SPACE"><expression>[SQL2].[USED_SPACE]</expression></dataItem><dataItem
name="POOLNAME"><expression>[SQL2].[POOLNAME]</expression></dataItem><dataItem
name="POOL_CUSTOM_TAG1"><expression>[SQL2].[POOL_CUSTOM_TAG1]</expression></dataItem><dataItem
name="POOL_CUSTOM_TAG2"><expression>[SQL2].[POOL_CUSTOM_TAG2]</expression></dataItem><dataItem
name="POOL_CUSTOM_TAG3"><expression>[SQL2].[POOL_CUSTOM_TAG3]</expression></dataItem><dataItem
name="STORAGE_SYSTEM"><expression>[SQL2].[STORAGE_SYSTEM]</expression></dataItem><dataItem
name="THINPROVISIONED"><expression>[SQL2].[THINPROVISIONED]</expression></dataItem><dataItem
name="ENCRYPTABLE"><expression>[SQL2].[ENCRYPTABLE]</expression></dataItem><dataItem
name="ENCRYPTED"><expression>[SQL2].[ENCRYPTED]</expression></dataItem><dataItem
name="STATUS"><expression>[SQL2].[STATUS]</expression></dataItem><dataItem
name="RAID_LEVEL"><expression>[SQL2].[RAID_LEVEL]</expression></dataItem><dataItem
name="IS_COMPRESSED"><expression>[SQL2].[IS_COMPRESSED]</expression></dataItem><dataItem
name="SPECTRUMCONTROL_TIER"><expression>[SQL2].[SPECTRUMCONTROL_TIER]</expression></dataItem><dataItem
name="POOL_EASYTIER"><expression>[SQL2].[POOL_EASYTIER]</expression></dataItem><dataItem
name="POOL_EASYTIER_STATUS"><expression>[SQL2].[POOL_EASYTIER_STATUS]</expression></dataItem><dataItem
name="STORAGESUBSYSTEM_CUSTOM_TAG1"><expression>[SQL2].[STORAGESUBSYSTEM_CUSTOM_TAG1]</expression></dataItem><dataItem
name="STORAGESUBSYSTEM_CUSTOM_TAG2"><expression>[SQL2].[STORAGESUBSYSTEM_CUSTOM_TAG2]</expression></dataItem><dataItem
name="STORAGESUBSYSTEM_CUSTOM_TAG3"><expression>[SQL2].[STORAGESUBSYSTEM_CUSTOM_TAG3]</expression></dataItem><dataItem
name="VOL_TIER_CAPACITY"><expression>[SQL2].[VOL_TIER_CAPACITY]</expression></dataItem><dataItem
name="APPLIST"><expression>[SQL2].[APPLIST]</expression></dataItem><dataItem
name="CURRENT_DATE"><expression>[SQL2].[CURRENT_DATE]</expression></dataItem></selection>
    </query></queries></report>

```

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *IBM SmartCloud Virtual Storage Center*, SG24-8239
- ▶ *IBM Software-Defined Storage Guide*, REDP-5121
- ▶ *IBM Spectrum Scale (formerly GPFS)*, SG24-8254
- ▶ *IBM Tivoli Storage Productivity Center Beyond the Basics*, SG24-8236
- ▶ *IBM Tivoli Storage Productivity Center V5.2 Release Guide*, SG24-8204
- ▶ *Implementing the IBM Storwize V7000 and IBM Spectrum Virtualize V7.6*, SG24-7938
- ▶ *Introduction Guide to the IBM Elastic Storage Server*, REDP-5253
- ▶ *Regain Control of your Environment with IBM Storage Insights*, REDP-5231

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- ▶ *IBM Spectrum Control Version 5.2.8 Administrator's Guide*, SC27-6588
- ▶ *IBM Spectrum Control Version 5.2.8 Installation Guide* SC27-6584
- ▶ *IBM Spectrum Control Version 5.2.8 User's Guide*, SC27-6588

Online resources

These websites are also relevant as further information sources:

- ▶ IBM Spectrum Control IBM Knowledge Center:
<http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>
- ▶ IBM Spectrum Scale IBM Knowledge Center
http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome

- ▶ *IBM Spectrum Control Standard Edition Version 5.2.8 Quick Start Guide*, found at:
http://www.ibm.com/support/knowledgecenter/api/content/nl/en-us/SS5R93_5.2.8/com.ibm.spectrum.sc.pdf/qstrtscse.pdf
- ▶ For more information about IBM Storage Insights, see the following websites:
 - <http://www.ibm.com/systems/storage/spectrum/insights>
 - <http://www.ibm.com/marketplace/cloud/analytics-driven-data-management/us/en-us>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM Spectrum Family: IBM Spectrum Control Standard Edition

SG24-8321-00
ISBN 0738441481



(0.5" spine)
0.475" <-> 0.873"
250 <-> 459 pages



SG24-8321-00

ISBN 0738441481

Printed in U.S.A.

Get connected

