

Regain Control of your Environment with IBM Storage Insights

Hasan Abdullah Hashmi

Gaston Rius

Megan Gilge



 **Cloud**

Storage



International Technical Support Organization

Regain Control of your Environment with IBM Storage Insights

August 2015

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

First Edition (August 2015)

This edition applies to IBM Storage Insights.

This document was created or updated on March 4, 2017.

© Copyright International Business Machines Corporation 2015. 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	v
Trademarks	vi
IBM Redbooks promotions	vii
Preface	ix
Authors	ix
Now you can become a published author, too!	x
Comments welcome	xi
Stay connected to IBM Redbooks	xi
Chapter 1. Overview	1
1.1 Introduction to the IBM Spectrum Control family	2
1.2 Introduction to IBM Spectrum Control Storage Insights	3
1.2.1 Overview	3
1.2.2 Licensing	4
1.2.3 Data collection	5
1.2.4 Security	6
1.2.5 Support	7
1.2.6 IBM Storage Insights compared to Virtual Storage Center	8
Chapter 2. Getting started	11
2.1 Enrollment	12
2.2 Deployment	12
2.2.1 Hardware, software, and other requirements	12
2.2.2 Data collector deployment	12
2.3 Configuration overview	13
2.3.1 Supported storage systems	13
2.3.2 Storage resource configuration	13
2.3.3 Configuring the Servers view	17
2.3.4 Application and Department views	18
2.4 Insights views	20
Chapter 3. Scenarios	21
3.1 Improving capacity planning	22
3.2 Increasing storage utilization	25
3.3 Planning data placement	29
3.4 Performance reporting	31
3.5 Simplifying reporting	34
3.6 Using agentless servers	36
3.7 Application groups and departments	37
Related publications	41
Online resources	41
Help from IBM	41

Notices

This information was developed for products and services offered in the U.S.A.

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, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM 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 supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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 the names and addresses used by an actual business enterprise 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.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	IBM FlashSystem®	Redbooks®
Bluemix™	IBM SmartCloud®	Redpaper™
Cloudant®	IBM Spectrum Accelerate™	Redbooks (logo)  ®
Cognos®	IBM Spectrum Archive™	Storwize®
DB2®	IBM Spectrum Control™	System i®
DS8000®	IBM Spectrum Protect™	Tivoli®
FlashSystem™	IBM Spectrum Scale™	XIV®
Global Technology Services®	IBM Spectrum Storage™	
IBM®	IBM Spectrum Virtualize™	

The following terms are trademarks of other companies:

SoftLayer, and SoftLayer device are trademarks or registered trademarks of SoftLayer, Inc., an IBM Company.

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

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 up-to-the-minute Redbooks news and announcements
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



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

This IBM® Redpaper™ publication introduces you to the new, cloud-based IBM Spectrum Control™ Storage Insights (*IBM Storage Insights*, for short) offering, which is designed for small and medium businesses and organizations who need to quickly understand what is happening in their storage environment without implementing complex tools. IBM Storage Insights can be set up in less than 5 minutes and provides actionable insights about your storage in less than 30 minutes.

IBM Storage Insights is an off-premise software as a service (SaaS) offering that is offered through IBM Marketplace. This simple, graphical tool has built-in reports to help you rapidly understand what is happening in your environment and provides recommendations about how you can maximize the benefits of your storage and improve your decision-making process.

This publication is designed to help storage administrators learn about benefits, features, and key implementation scenarios.

The retention period for daily performance data was updated in February 2017.

IBM Marketplace links were added in March 2017.

Authors

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



Hasan Abdullah Hashmi is an IT Architect in Karachi, Pakistan. He obtained his Bachelor's degree in Computer Studies before joining IBM. He started his career at IBM and later completed his Master's in business. He worked for the Infrastructure Service as an IBM Systems, Storage, and Storage Software Specialist for 6 years. He has been performing the role of a Pre-Sales Infrastructure Architect for the last three years supporting the sales team to achieve valuable business targets. He is certified by The Open Group Architect Foundation for TOGAF-9 level and by IBM as an Experienced IT Architect. His current focus is to develop solutions on IBM Cloud and software-defined environments to help clients show IBM value in these domains.



Gaston Rius is a Senior IT Specialist from Buenos Aires, Argentina with 16 years of IT experience. Before joining IBM Global Technology Services® (GTS) in 2004, he worked as a Data Center and UNIX Support IT Specialist for three years at a major telecommunication company and then two years as Pre-Sales Support and Implementator for an IBM Business Partner. His current role in IBM is Storage Subject Matter Expert (SME) and his focus is in the field storage solutions design, implementation, and client support to multiple outsourced clients worldwide. His areas of expertise include a wide array of storage hardware and software products from multiple vendors, storage virtualization, storage area networks, storage tiering, remote copy services, IBM Tivoli® Storage Productivity Center, IBM SmartCloud® Virtual Storage Center, and Brocade Network Advisor. He is also interested in patents and cloud.



Megan Gilge is a Project Leader in the IBM International Technical Support Organization. Before joining the ITSO four years ago, she was an Information Developer in the IBM Semiconductor Solutions and IBM System i® areas.

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

Karen Orlando
International Technical Support Organization, Poughkeepsie Center

Brian De Guia
Thuan Nguyen
Socheat Sou
Markus Standau
Scott Trevett
IBM Systems

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 papers to be as helpful as possible. Send us your comments about this paper 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:

<https://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>



Overview

IBM Spectrum Control™ Storage Insights (*IBM Storage Insights*, for short) is an off-premise, cloud-based software as a service (SaaS) offering. This chapter provides an introduction to IBM Storage Insights as part of the IBM Spectrum Family of products and as a SaaS offering. It then describes the key features of IBM Storage Insights.

Tip: IBM Storage Insights is easy to configure and to use. It might take less time to sign up for a no-charge trial and try it than to read this paper.

For more information about a trial, see IBM Marketplace:

<https://www.ibm.com/marketplace/search/us/en-us?terms=IBM+Spectrum+Control+Storage+Insights>

This chapter contains the following topics:

- ▶ Introduction to the IBM Spectrum Control family
- ▶ Introduction to IBM Spectrum Control Storage Insights
- ▶ Introduction to IBM Spectrum Control Storage Insights

1.1 Introduction to the IBM Spectrum Control family

In the past few years, companies have been facing the exponentially increasing demand for new storage capacity, either driven by traditional workloads such as enterprise resource planning (ERP), human resources (HR), and customer relationship management (CRM), or new ones as mobile, social media, big data, and analytics. This tendency continues to grow as companies are becoming eager to take advantage of the insights that are provided by this new data to make more informed and better business decisions.

Daily, 2.5 Exabytes (2.5 billion GB) of data is created. About 90% of the worldwide data been created in the last two years and the trend keeps increasing.

With costs per GB declining, storage has become a commodity within companies that usually would try to meet these needs by simply procuring more storage capacity, often allowing them to postpone or overlook tasks, such as adopting an adequate Information Lifecycle Management (ILM) strategy. This postponement leads to challenges as storage management support staffing costs increases and the existing Business Continuity Recovery Services (BCRS) strategies must be redesigned a result of larger and more complex storage environments.

The IBM response to these new market challenges is delivered through the IBM Spectrum Storage™ family of software-defined offerings.

Spectrum Storage enhances the speed and efficiency of your storage and simplifies migration to new workloads by:

- ▶ Simplifying and integrating storage management and data protection across traditional and new applications
- ▶ Delivering elastic scalability with high performance for analytics, big data, social, and mobile
- ▶ Unifying silos to deliver data without borders with built-in hybrid cloud support
- ▶ Optimizing data economics with intelligent data tiering from flash to tape and cloud
- ▶ Building on open architectures that support industry standards, including OpenStack and Hadoop

The IBM Spectrum Family consists of the following products:

- ▶ *IBM Spectrum Accelerate*™ offers grid-scale block storage with rapid deployment that helps speed delivery of data across an enterprise and adds extreme flexibility to cloud deployments
- ▶ *IBM Spectrum Scale*™ is flash accelerated, industrial strength, highly scalable software-defined storage that enables global shared access to data with extreme scalability and agility for cloud and analytics
- ▶ *IBM Spectrum Virtualize*™ software is at the heart of IBM SAN Volume Controller and 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 a simpler, more responsive, scalable, and cost efficient IT infrastructure
- ▶ *IBM Spectrum Protect*™ enables reliable, efficient data protection and resiliency for software-defined, virtual, physical, and cloud environments. *IBM Spectrum Archive*™ enables you to automatically move infrequently accessed data from disk to tape to lower costs while retaining ease of use and without the need for proprietary tape applications.

- ▶ *IBM Spectrum Control* provides analytics-driven data management as well as efficient infrastructure management for virtualized, cloud, and software-defined storage to simplify and automate storage provisioning, capacity management, availability monitoring, and reporting. IBM Spectrum Control offerings can be deployed:
 - From the cloud. *IBM Storage Insights* supports IBM storage appliances and IBM software-defined storage installations, with no hardware to install and only one data collector to deploy, when one or more IBM storage systems are installed. It provides reporting and recommendations to improve utilization and performance.

Data collector: The data collector is not distributed as a virtual application or appliance. It is a small application that is downloaded as a compressed file. The file is unpacked and a script is run to launch the data collector. The data collector can be installed on a VM client or on a server.

- On-premises. *IBM Virtual Storage Center (VSC)*, a bundle of Spectrum Control and Spectrum Virtualize, provides a richer experience with an active administration console, automated storage provisioning, and tier optimization. Virtual Storage Center includes a storage service catalog, drivers for OpenStack and VMware, and application aware snapshot management.

To learn more about the IBM Spectrum Control Family, see the following link:

<http://www.ibm.com/systems/storage/spectrum>

1.2 Introduction to IBM Spectrum Control Storage Insights

IBM Spectrum Control Storage Insights is a SaaS offering with its core running over IBM SoftLayer. It combines IBM analytics leadership and a rich history of storage management expertise with a cloud delivery model, enabling you to take control of your storage environment.

The continuous delivery model is used seamlessly to deploy all application updates.

SaaS offering licensing is implemented through different subscription levels.

For more information about IBM Marketplace offerings and IBM Spectrum Control Storage Insights, see the following links:

- ▶ IBM Marketplace main portal
<https://www.ibm.com/us-en/marketplace>
- ▶ IBM Spectrum Control Storage Insights
<https://www.ibm.com/us-en/marketplace/analytics-driven-data-management>

1.2.1 Overview

The solution is oriented to small and medium businesses who want to avoid the expensive and time-consuming deployment of an on-premises solution and enable less experienced staff to manage storage environments more efficiently by delivering different insights as simple as looking at the many available dashboards. Large organizations can also deploy IBM Storage Insights to gain visibility of small storage environments, even if they are already using on-premises solutions like IBM Virtual Storage Center (VSC) to manage their core storage systems.

Some of the many benefits and features that IBM Storage Insights delivers are:

- ▶ Improve capacity planning. Take the guess work out of capacity planning with increased visibility into historical data growth rates and available capacity.
- ▶ Increase storage utilization. Delay future purchases by identifying and reclaiming provisioned, but unused, storage.
- ▶ Provide tiering recommendations to optimize data placement. Reduce costs by moving data to the most cost-effective storage tier with customizable analytics-driven data management.
- ▶ Enhance performance monitoring. Eliminate service disruptions due to lack of available storage.
- ▶ Simplify reporting. Deliver a consumer-oriented view of storage by reporting on storage by departments or applications.

Figure 1-1 shows the main benefits and features of IBM Storage Insights.

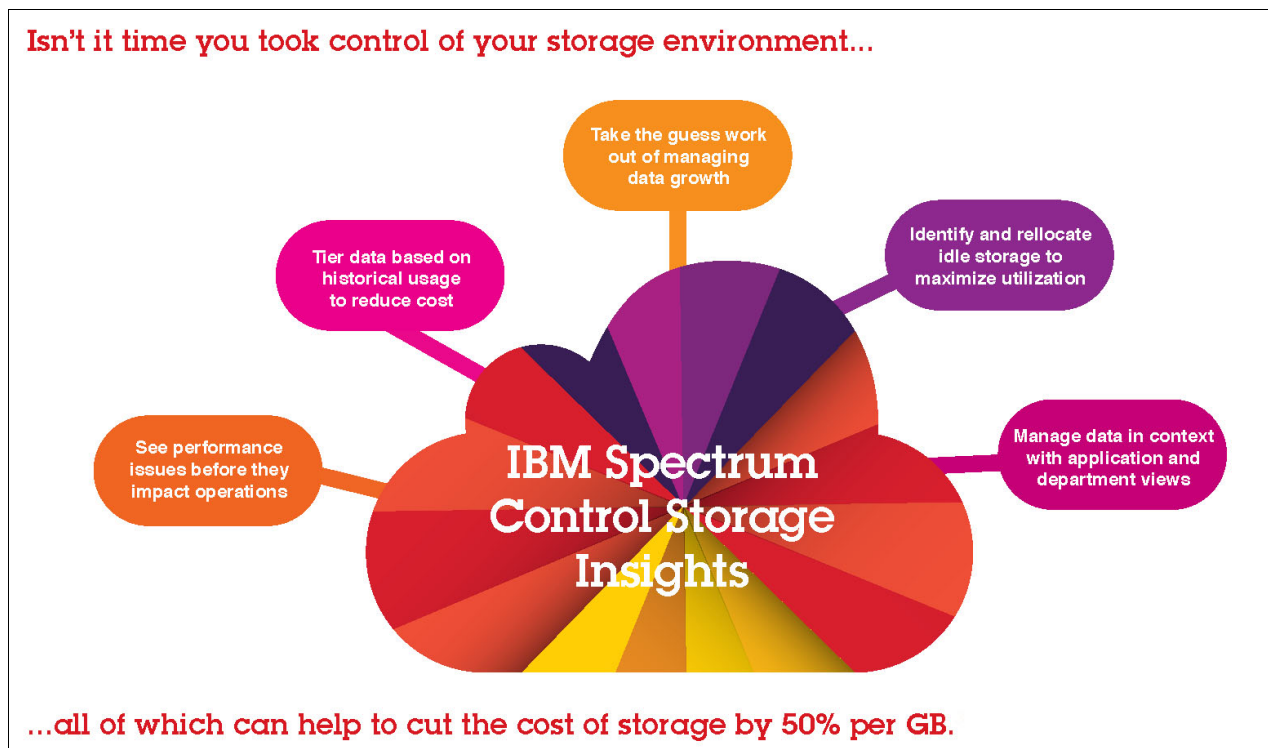


Figure 1-1 IBM Storage Insights features and benefits¹

1.2.2 Licensing

IBM Storage Insights licensing is based on a subscription model that is based on the total capacity managed.

To try IBM Storage Insights, sign into the IBM Marketplace website with your IBM ID and request a no-charge 30-day trial.

¹ Based on IBM's own experiences through internal implementation.

For more information about enrollment, subscriptions, and trials, see the following link:

<https://www.ibm.com/marketplace/search/us/en-us?terms=IBM+Spectrum+Control+Storage+Insights>

1.2.3 Data collection

IBM Storage Insights can be deployed in less than 5 minutes, starts to collect storage systems data immediately, and delivers the first set of actionable insights in less than 30 minutes of runtime.

To start deployment, you log in to your IBM Storage Insights instance and download the data collector software (available for the Microsoft Windows, Linux, and IBM AIX® operating systems). This data collector is for one-way transmission. It makes requests to upload metadata, but IBM Storage Insights does not make any requests to retrieve data.

After it is deployed, the data collector is ready to start to retrieve configuration and performance-related information from the storage systems and transmit it over the Internet to the IBM Storage Insights instance through a secure communication channel, as shown in Figure 1-2.

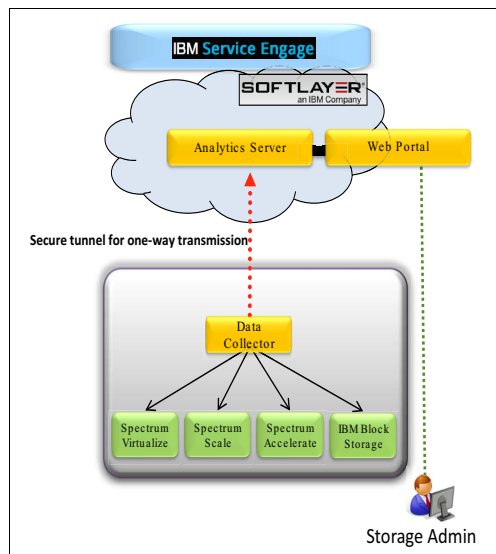


Figure 1-2 IBM Storage Insights components

This connection is for one-way transmission and IBM Storage Insights does not make any requests to retrieve data. The only information that is sent to IBM Storage Insights is the storage system *metadata*. This information includes:

- ▶ Information about the configuration of the Storage system (name, firmware, capacity, and so on).
- ▶ Information about the subresources of the storage system, such as volumes, pools, nodes, ports, and disks. This information includes the name, configuration, and capacity information for each subresource.
- ▶ Performance information for various levels of the storage system internal hierarchy.

Note: IBM Storage Insights cannot access the actual application data that is stored on the storage systems.

Data retention periods

IBM Storage Insights implements historical data retention periods for the data collected from the storage systems as well as for the internal tool logging. These periods vary depending on the granularity of the data to be retained.

The capacity data retention periods are:

- ▶ Daily: 12 weeks
- ▶ Weekly: 24 weeks
- ▶ Monthly: 24 months

The performance data retention periods are:

- ▶ Sample: 2 weeks
- ▶ Hourly: 4 weeks
- ▶ Daily: 12 months

1.2.4 Security

No data from your storage systems is sent to IBM Storage Insights. Multiple security techniques are in place to ensure security for the metadata that is sent. Storage device passwords are encrypted before they are stored on the IBM Storage Insights instance.

IBM Storage Insights cannot modify your environment: IBM Storage Insights has no capabilities to make any changes to the monitored environment. It collects only the storage systems metadata about subcomponents, configuration, and performance statistics.

This data collector is for one-way transmission. It makes requests to upload metadata, but IBM Storage Insights does not make any requests to retrieve data.

Sensitive information such as user names and passwords (or SSH certificates) are provided by IBM Storage Insights to each storage system upon each data collection iteration, and this information is transmitted over a secure communication channel that is established by the data collector.

You can change the connection credentials for each storage system at any time by selecting **Resources** → **Storage Systems** from the top menu bar in IBM Storage Insights, as shown in Figure 1-3. Select the wanted storage system, right-click, then select **Connections** → **Update credentials**. After you change credentials, you can test the connectivity to the storage device by selecting **Connections** → **Test Connection**.

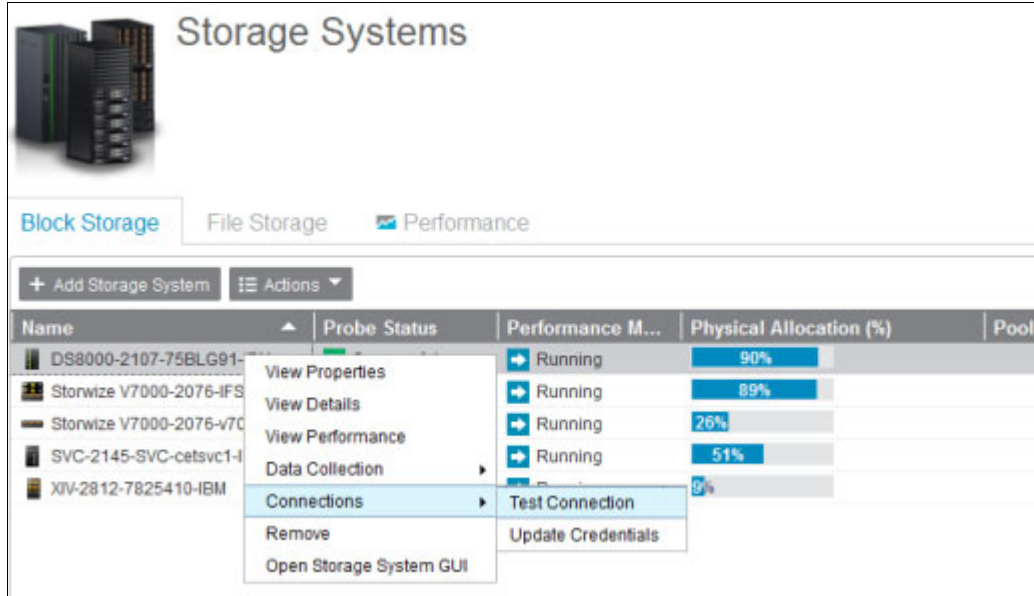


Figure 1-3 Changing and testing the storage system's credentials

For more information about security for IBM Spectrum Control Storage Insights, see the *IBM Spectrum Control Storage Insights Security Information* white paper:

<https://ibm.ent.box.com/v/storage-insights-security-wp>

1.2.5 Support

For support documentation, you can access the support website via IBM Storage Insights. As shown in Figure 1-4, look for the question mark (?) icon on the upper right side of the panel, hover over the icon, and select **Support** when the help options are presented.

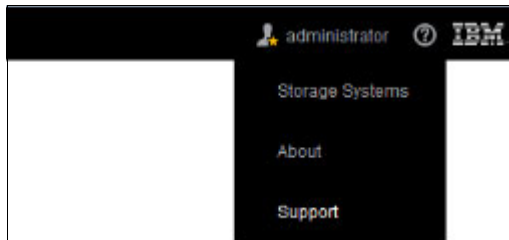


Figure 1-4 IBM Storage Insights help options

For more information about support, see the following websites:

- ▶ IBM Marketplace

<https://www.ibm.com/marketplace/search/us/en-us?terms=IBM+Spectrum+Control+Storage+Insights>

- ▶ IBM Storage Insights forum

<https://www.ibm.com/developerworks/community/forums/html/forum?id=d6c4eeb1-4ada-4807-a333-2b92fc93ecca>

1.2.6 IBM Storage Insights compared to Virtual Storage Center

IBM Storage Insights and Virtual Storage Center (VSC) are both storage offerings in the IBM Spectrum Storage family. One of the most distinctive differences between them is the deployment method: While IBM Storage Insights is deployed on the cloud, VSC is deployed on-premises.

Another aspect of the deployment to be considered is its complexity: The IBM Storage Insights implementation is fast and simple. It requires only the data collector to be installed on a VM (or real small physical server) and network access to the storage devices and to the Internet, whereas VSC requires dedicated physical server hardware to be deployed (hardware specifications vary from case to case and should be thoroughly sized during the planning phase to properly support each environment).

Code version currency is another major differentiator. VSC requires the user to check for periodical updates and maintain the version changes by upgrading of all of its many components according to vendor-specific instructions and guidance, which can be complex for components such as IBM DB2® or JazzSM. However, IBM Storage Insights only requires the user to deploy a small data collector. All code updates for IBM Storage Insights are done by IBM. The user never needs to install updates.

IBM Storage Insights and VSC licensing are both based on managed capacity. However, VSC requires the license to be procured and paid in full up-front (as for most standard software product licenses). IBM Storage Insights is delivered as a software as a service (SaaS) offering and the service access is entitled through a subscription-based model. For more information, see 1.2.2, “Licensing” on page 4.

Even though both IBM Storage Insights and VSC require standard-user access level on each of the storage systems that they connect to, IBM Storage Insights *does not provide any functionality to change anything in your storage system's configuration*. It only collects information and delivers multiple insights to the user based on IBM proprietary analytics.

In IBM Storage Insights, the Reclamation Insights view provides a simple graphical user interface approach to gain visibility (or insight) over provisioned but unused storage capacity on the environment. Efficiency of storage utilization is one of the most important topics for all organizations. When you can optimize utilization, you can reuse unused capacity and postpone purchasing more storage. This view enables you to quickly identify candidate volumes to be reclaimed and it can be accessed as quickly as two clicks away from any view in the tool.

With the same simplicity as reclamation, IBM Storage Insights also delivers tier planning recommendations to help you achieve proper data placement according to customizable criteria (either by volume I/O or I/O density) in addition to a comprehensive Capacity dashboard where you can quickly go through various aspects of the environment, such as Capacity Over Time (current usage and forecasting), Current Space by Tier or Device, Replication (protection level) of the volumes, and Thin Provisioning Capacity Savings.

Table 1-1 lists some of the main characteristics and features of both VSC and IBM Storage Insights and describes how these features are delivered. Use this comparison to consider specific product capabilities. It is not a definitive list of product features.

New functions: New functions are frequently released for both VSC and IBM Storage Insights. Always check the most recent product documentation for a full list of functions.

Table 1-1 Comparison of IBM Storage Insights and VSC features

Feature	IBM Storage Insights	Virtual Storage Center 5.2 (VSC)
Deployment	Minutes - On the cloud	Hours - On the premises
Third-party vendor support	No	Yes (through CIMOM agents)
Licensing	Periodic subscription (based on managed capacity)	Up front, one time only (based on managed capacity)
Customized Alerting	No	Yes
Element Manager Support	Yes	Yes
Historical data retention	Non-customizable. 24 months for Capacity and 13 months for Performance.	Customizable. Up to 48 months for Capacity and 39 months for Performance.
Product code currency	SaaS offering; code changes are delivered by IBM over the cloud	Multiple components that are upgraded by the user (IBM Tivoli Storage Productivity Center, DB2, JazzSM)
Capacity reclamation analysis	Yes (reclamation candidates based on predefined criteria)	No
Maximum quantity of storage devices per installation	Not limited	Not limited. Proper sizing of server hardware (processor, RAM, internal capacity) based on the environment to manage
Reporting	Pre-defined views	Pre-defined views and custom reporting (either with conventional Tivoli Storage Productivity Center GUI or with IBM Cognos®)
Smallest performance data sample granularity	5 minutes	1 minute
Host-level data collection	Limited information through agentless method	Detailed information through SRA deployment. Limited through agentless method
Filtered view by applications and departments	Yes	Yes
Hypervisor (ESX) view support	No	Yes
Performance monitoring	Yes, real-time and historical	Yes, real-time and historical
Thin provisioning and compression support	Yes	Yes
SAN switch support	No	Yes (no direct support, integration via SAN switch vendor agent)

Feature	IBM Storage Insights	Virtual Storage Center 5.2 (VSC)
Tiering support	Tiering Planning; easy to set up. Recommends volume migrations between tiers	Tiering optimization; complex to set up. Recommends and implements volume migrations between tiers

For the most current information about features, capabilities, and supported systems, see the following links:

IBM Storage Insights

<http://www.ibm.com/systems/storage/spectrum/insights/index.html>

Virtual Storage Center (VSC)

<http://www.ibm.com/software/products/en/virtual-storage-center>



Getting started

This chapter briefly describes the deployment and configuration activities to configure the IBM Storage Insights instance. It also provides guidance on the main sections of the tool (insights) and their usage.

It contains the following information:

- ▶ Enrollment
- ▶ Deployment
- ▶ Configuration overview
- ▶ Insights views

2.1 Enrollment

You enroll for IBM Storage Insights through IBM Marketplace. An IBM ID is required. You can also request a 30-day, no-charge trial.

For the IBM Marketplace product catalog, see the following link:

<https://www.ibm.com/marketplace/search/us/en-us?terms=IBM+Spectrum+Control+Storage+Insights>

IBM Storage Insights licensing is based on the total capacity that is managed and can be easily purchased, changed, or renewed.

For more information, see the IBM Spectrum Control Storage Insights FAQ:

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

2.2 Deployment

Only a single software component, the *data collector*, must be deployed to begin using IBM Storage Insights. The data collector acts as a proxy agent that communicates to the storage devices, gathers information that pertains to configuration and performance, and sends it to the IBM Storage Insights instance on the cloud.

2.2.1 Hardware, software, and other requirements

The data collector is a lightweight software component and requires either a physical server or virtual machine with at least 1 GB of RAM and 300 - 500 MB of disk capacity.

The Microsoft Windows, Linux, and AIX operating systems are supported.

The data collector also requires Ethernet network access to the storage systems for communication in addition to Internet connectivity to upload the gathered storage subsystem metadata to IBM Storage Insights.

2.2.2 Data collector deployment

To deploy the data collector in your environment, you must do the following steps:

1. Log in to your IBM Storage Insights instance and go to **Configuration** → **Data Collector**.
2. Select the operating system on which you want to deploy the agent and download it locally on the server.
3. Extract the data collector, test your network connection to the address for your IBM Storage Insights instance on port 9550, and run the data collector installation script.

When the data collector is up and running, its status appears in the **Configuration** → **Data Collector** view in IBM Storage Insights.

If the data collector stops working or the communication with IBM Storage Insights is interrupted, support detects this problem and an alert is sent to the email address that is used to subscribe to the IBM Storage Insights service. You can contact support to request a change to the email address for data collector alerts.

Internet access: If the server or virtual machine to be used to deploy the data collector does not come with direct access to the Internet, the data collector can be also configured to communicate with IBM Storage Insights through an Internet proxy server.

For more information about the data collector, such as deployment and configuring the agent to use an Internet proxy server, see the Installing the data collector information in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/tpch_saa_s_t_data_collector_welcome_wizard.html?lang=en

2.3 Configuration overview

This section briefly introduces the process to add and configure the storage systems in IBM Storage Insights in order to enable the tool to start collecting data and provide all of its available insights.

2.3.1 Supported storage systems

At the time of writing, the following storage systems are supported by IBM Storage Insights:

- ▶ IBM DS8000®
- ▶ IBM Spectrum Scale
- ▶ IBM FlashSystem™ V840 and FlashSystem V9000
- ▶ IBM Scale Out Network Attached Storage (SONAS)
- ▶ IBM SAN Volume Controller
- ▶ IBM Storwize V3500, Storwize V3700, Storwize V5000, and Storwize V7000
- ▶ IBM Storwize V7000 Unified
- ▶ IBM XIV® system

See the Adding storage systems topic in the IBM Knowledge Center for a list of supported storage systems:

http://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/tpch_saa_s_t_container_adding_ssystems.html?lang=en

2.3.2 Storage resource configuration

The following tasks are required to add and configure storage devices for IBM Storage Insights:

- ▶ “Adding a storage system” on page 14
- ▶ “Scheduling the probe and performance data collection” on page 15
- ▶ “Configuring pools tiers and defining their tiering thresholds” on page 16

Adding a storage system

Do the following steps to add a storage system:

1. Select **Resources** → **Storage Systems**.
2. In the Storage Systems view, click **Add Storage System**. A configuration wizard appears, as shown in Figure 2-1.

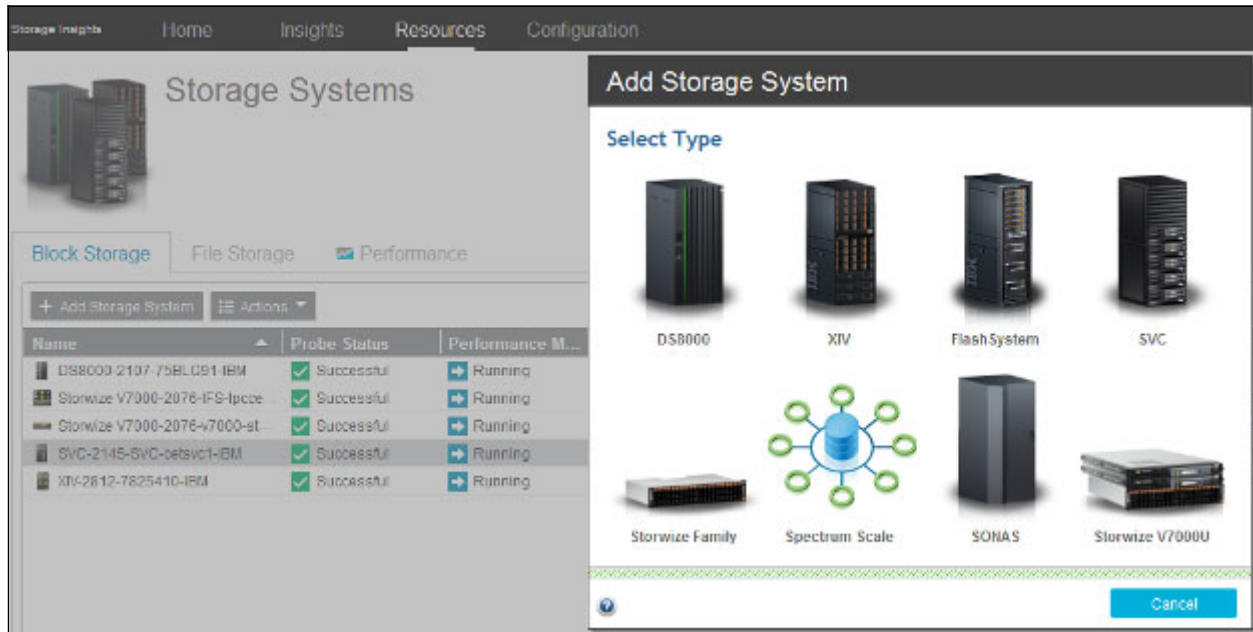


Figure 2-1 Adding a storage subsystem

The wizard displays a list of currently supported storage systems types and models.

3. To select one, click the corresponding picture.

When you are done, a new configuration pane like Figure 2-2 is displayed.

4. Provide the IP address of the storage system and the corresponding login credentials and click **Next**.

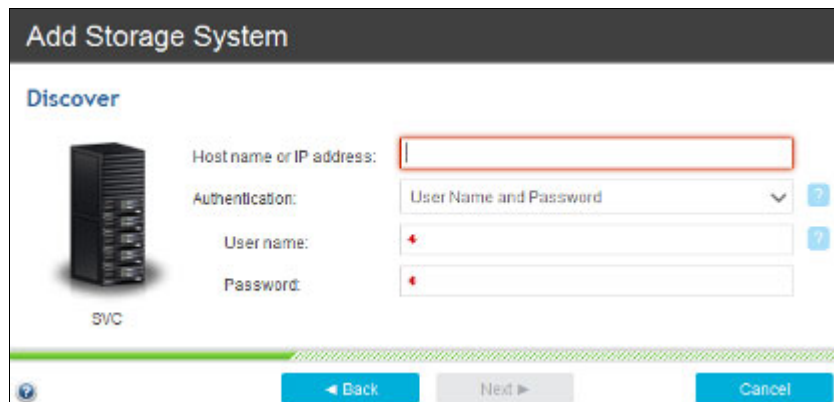


Figure 2-2 Providing storage subsystems connection information

IBM Storage Insights tests the connection to the storage system. If it connects successfully, it adds the system to the managed environment and displays the Data Collection Schedule configuration panel (see Figure 2-3), where you configure the storage system probe schedule, enable the performance monitor, and configure the granularity of performance samples.

5. Specify the settings that you want and click **Save** when you are done.

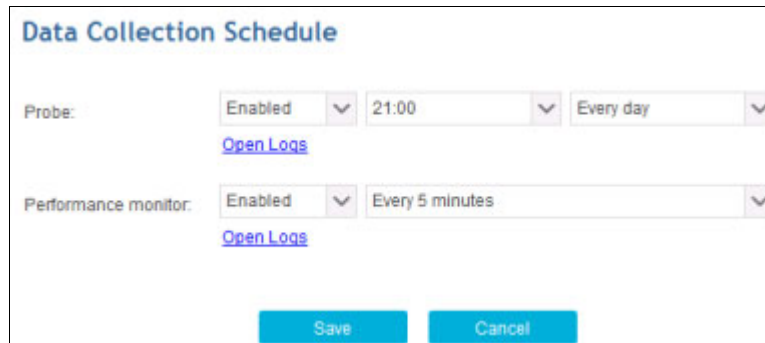


Figure 2-3 Configuring configuration probe and performance monitor schedules

6. Immediately after you add the storage subsystem, the configuration probe job is automatically run and the performance monitor starts to collect performance data (if enabled during the storage system setup). Information for the different insights views is available to the user after 30 minutes of data capture.

DS8000 and SSLv3 protocol: By default, IBM Storage Insights supports DS8000 systems whose firmware levels do not support the SSLv3 protocol. If you are running a DS8000 software version that uses SSLv3, see the IBM Spectrum Control Storage Insights documentation in the IBM Knowledge Center for details about how to change the communication method before you add the DS8000 system:

http://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/tpch_saas_t_configuring_ds8k_sslv3.html?lang=en

For detailed information about adding storage subsystems and specific notes for each model and type, see the Adding storage systems topic in the IBM Spectrum Control Storage Insights documentation in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/tpch_saa_s_t_container_adding_ssystems.html?lang=en

Scheduling the probe and performance data collection

Scheduling for the configuration probe and performance data collection is usually configured when the storage system is added to the managed environment.

To make more modifications (such as disabling either component or modifying its schedule), select **Resources** → **Storage Systems** from the menu, right-click the wanted storage system, and select **Data Collection** → **Schedule** to open the Data Collection Schedule pane (Figure 2-3).

For more information about working with probe configuration and performance data collection, see the Modifying data collections topic in the IBM Spectrum Control Storage Insights documentation in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/tpch_saa_s_t_configuring_container.html?lang=en

Configuring pools tiers and defining their tiering thresholds

To view insights for tiering recommendations in the Tier Planning Insight view, IBM Storage Insights requires a minimal configuration to define the different levels of tiers that are associated to each one of the storage system pools and to set the performance thresholds that are observed by the tool to present these recommendations.

Do the following steps to set the tier for each storage pool:

1. Select **Resources** → **Pools** on the top menu bar (Figure 2-4).
2. On the main pool view, select one or multiple pools, right-click the selection, and select **Set Tier** to define the level to be assigned to your pool selection.

Name	Storage System	Physical Allocation (%)	Capacity (GiB)	Shortfall (%)	Tier
DS8K_2541-poolC9	SVC-2145-SVC	77%	20.00	0%	Tier 2
DS8k_2541_PM-pool	SVC-2145-SVC		39.73	0%	Tier 2
DS8K_A271-poolC5	SVC-2145-SVC		40.00	0%	Tier 2
DS8K_A271-PM-pool	SVC-2145-SVC-cetsvc1-IBM	0%	2.56	0%	Tier 2
DS8k_A271-PM-pool	SVC-2145-SVC-cetsvc1-IBM		0.94	0%	Tier 2
DS8K_L581-poolC10	SVC-2145-SVC-cetsvc1-IBM	2%	100.00	0%	Tier 2
DS8K_LG91-poolC7	SVC-2145-SVC-cetsvc1-IBM	0%	11.00	0%	Tier 2
DS8K_LG91_L3-poolC7	SVC-2145-SVC-cetsvc1-IBM	0%	9.00	0%	Tier 2
DS8k_LG91-PM-pool	SVC-2145-SVC-cetsvc1-IBM		10.00	0%	Tier 2
parent_pool	SVC-2145-SVC-cetsvc1-IBM		3.97	0%	Tier 3
PM_pool	SVC-2145-SVC-cetsvc1-IBM		40.00	0%	Tier 3
TPC-R_pool	SVC-2145-SVC-cetsvc1-IBM		480.00	57%	Tier 3
v7k_ford1-poolC4	SVC-2145-SVC-cetsvc1-IBM	0%	20.00	0%	Tier 3
XIV_5410-PM-pool	SVC-2145-SVC-cetsvc1-IBM	64%	48.05	84%	Tier 3

Figure 2-4 Defining the tier level for each storage pool

Tier performance thresholds are user customizable.

3. Select **Configuration** → **Tiers** from the top menu bar to display the tiers view.
4. Click **Define Tier Thresholds**.

- On the Define Tier Thresholds pane (Figure 2-5), select the criteria to be used on each tier level (either *volume total I/O rate* (read and write I/O operations per second) or *volume I/O density* (volume total I/O operations divided by volume size in Gigabytes) and specify the wanted threshold value.

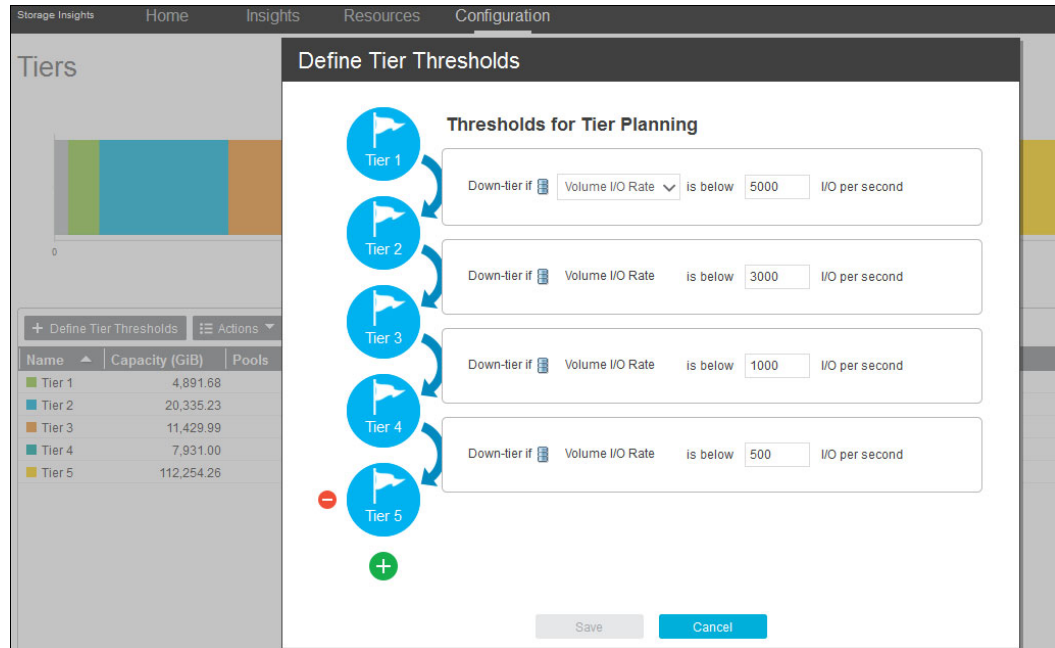


Figure 2-5 Defining down-tiering thresholds

When you define performance characteristics for each tier of storage, IBM Storage Insights can make recommendations for potential volumes migrations between tiers (retiering) to help you create more efficient data placement.

Note: Tiers must start at Tier 1 and be consecutive.

For more information about tiering, see the Creating tiers topic in the IBM Spectrum Control Storage Insights documentation in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSQR88/com.ibm.spectrum.si.doc/tpch_saa_s_t_container_configuring_tiers.html?lang=en

2.3.3 Configuring the Servers view

You can use IBM Storage Insights to collect information about the server and to identify storage subsystem subcomponents (controllers, pools, volumes), which is presented from a server-centric view.

The **Resources** → **Servers** page shows the list of servers that are defined on the IBM Storage Insights instance. You can add, remove, and configure the servers that are in this list.

Servers can be added by clicking **Add Server**, where you can choose to add them manually (one at a time) or to provide a list that is formatted as a .csv comma delimited file (the field format to be used is provided by IBM Storage Insights in a wizard).

When you select any of the servers that are listed in the Servers view, you can drill down into that particular server. As shown in Figure 2-6, you can explore the server view and gain visibility over its related storage systems' subcomponents (volumes, pools, ports) and critical server aspects, such as volume performance and utilization by using dashboards.

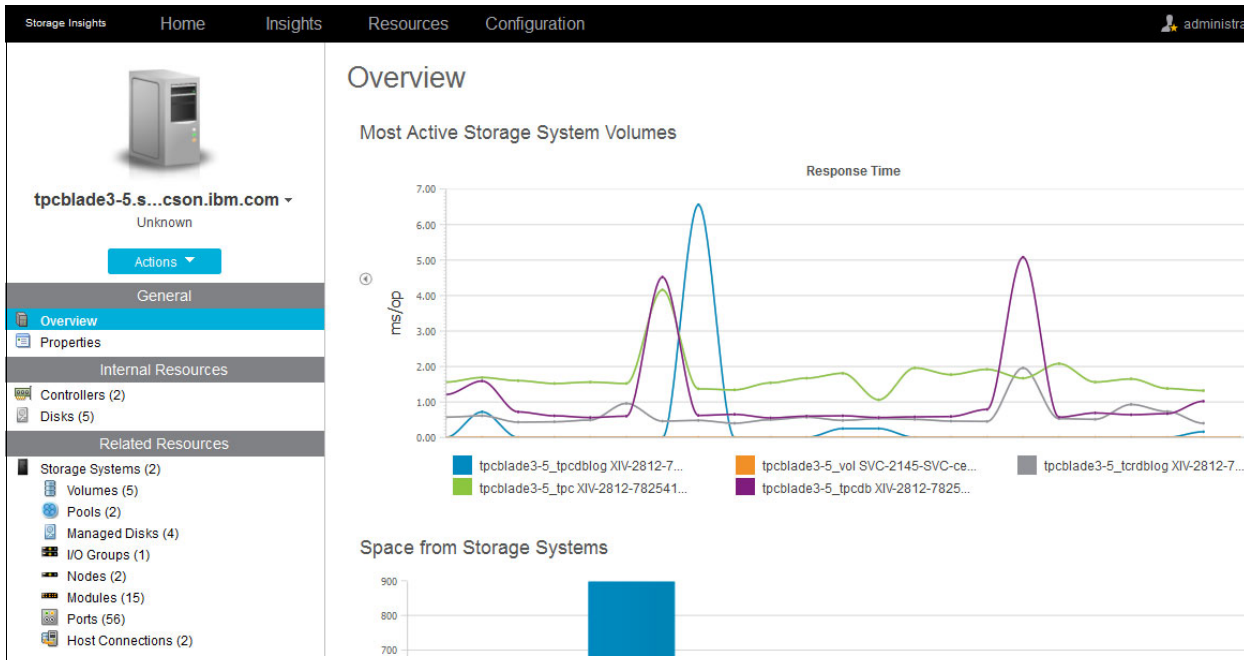


Figure 2-6 Servers view in IBM Storage Insights

Server management capabilities: IBM Storage Insights cannot modify anything about your servers' configuration or the storage subcomponents that are associated with them. The server in IBM Storage Insights is implemented as a logical object in the IBM Storage Insights component hierarchy (storage systems subcomponents, servers, applications, departments, and so on).

For more information about working with servers, see the Adding Servers topic in the IBM Spectrum Control Storage Insights documentation in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/tpch_saa_s_t_container_adding_servers.html?lang=en

2.3.4 Application and Department views

You can use IBM Storage Insights to model your organization's applications by creating applications and applications subcomponents (optional), defining their hierarchies, and associating them to the corresponding storage resources that they use with filters.

After the applications are created, you can use IBM Storage Insights to monitor storage resources aspects, such as capacity and performance from an application perspective. This view provides the type of visibility that the application's user has. Figure 2-7 on page 19 shows an example of the Department view.

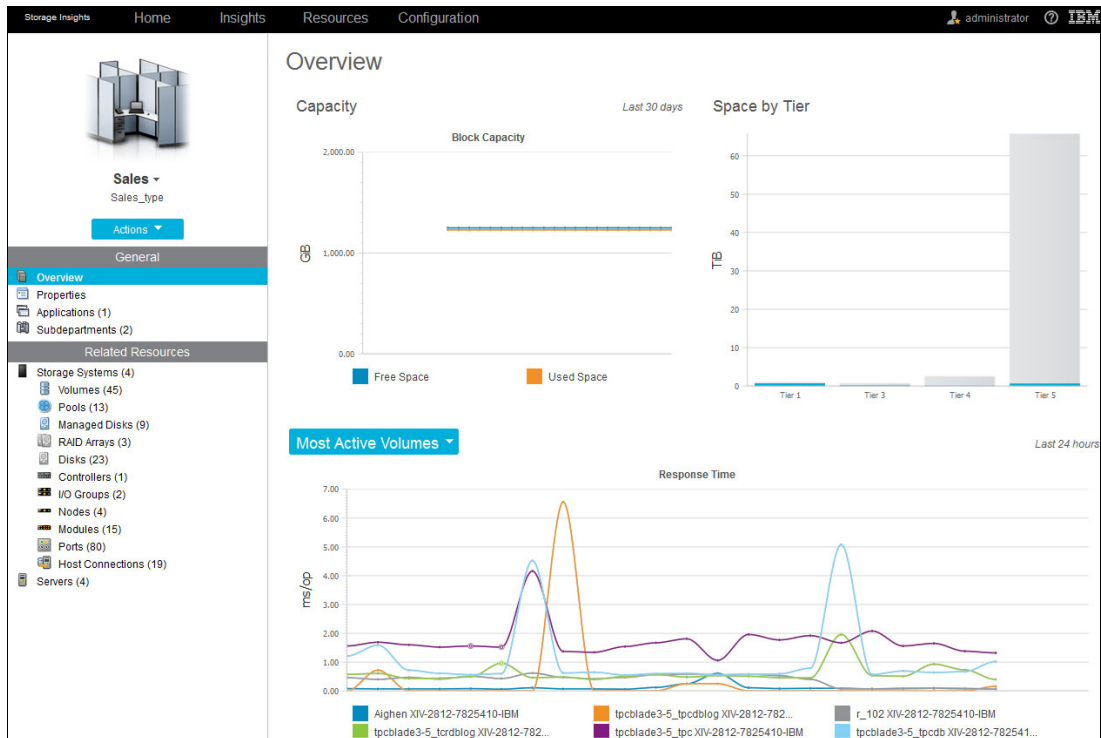


Figure 2-7 Department view in IBM Storage Insights

You can use a similar process to create departments, subdepartments (optional), and their hierarchies, and to associate the applications to their corresponding departments. Figure 2-8 shows an example of organizational departmental modeling.




Figure 2-8 Departmental hierarchy

For more information about working with applications and departments, see the following topics in the IBM Spectrum Control Storage Insights documentation in the IBM Knowledge Center:

- ▶ Creating applications

http://www.ibm.com/support/knowledgecenter/SSQR8/com.ibm.spectrum.si.doc/tpch_saas_t_container_creating_applications.html?lang=en

- ▶ Creating departments

http://www.ibm.com/support/knowledgecenter/SSQR8/com.ibm.spectrum.si.doc/tpch_saas_t_container_creating_departments.html?lang=en

2.4 Insights views

As briefly described in 1.2.1, “Overview” on page 3, you can receive benefits from IBM Storage Insights in less than 30 minutes after deployment.

These benefits are unique to the Insights views. The main goal of the Insights views is to make it simple to obtain information that might otherwise be complex to obtain by using any other tools or traditional methods. The collected storage systems’ information is processed by using proprietary IBM analytics, and insights are delivered to the user by means of dynamic dashboards and simple views that provide quick visibility into storage environment health and efficiency-improving recommendations, which can help your organization make more informed and better decisions.

The following Insights views are available to users:

- ▶ Capacity Insight

Detailed view of device capacity utilization and historical growth, thin provisioning and compression savings, capacity growth forecast based on empirical data such as historical growth rates and available capacity, among several other offered insights. At the time of writing, forecast is not available for a tier, a subsystem, or a pool. See IBM Knowledge Center for the latest list of available functions:

https://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/tpch_saas_welcome.html

- ▶ Performance Insight

Integrated view of all the storage systems performance with drill-down capabilities to simplify and accelerate the performance-related troubleshooting.

- ▶ Reclamation Insight

Identification of provisioned but unused capacity. This view helps organizations improve the utilization across the environment, postpone more capacity purchases, and repurpose reclaimed capacity.

- ▶ Tier Planning Insight

Optimize data placement by providing retiering recommendations based on actual empirical performance requirements of the volumes.

For more information about each Insights view, see the Gaining insights topic in the IBM Spectrum Control Storage Insights documentation in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/tpch_saas_c_container_gaining_insights.html?lang=en

Chapter 3, “Scenarios” describes use cases for these Insights views.



Scenarios

This chapter describes industry usage scenarios for IBM Storage Insights. When you configure IBM Storage Insights as described in Chapter 2, “Getting started” on page 11, you can use these scenarios to maximize benefit from IBM Storage Insights.

This chapter describes the following typical scenarios for using IBM Storage Insights:

- ▶ Improving capacity planning
- ▶ Increasing storage utilization
- ▶ Planning data placement
- ▶ Performance reporting
- ▶ Simplifying reporting
- ▶ Using agentless servers
- ▶ Application groups and departments

3.1 Improving capacity planning

IBM Storage Insights helps you improve and automate your capacity planning. It collects data for analysis via its data collector and provides results based on its analytics.

The historical data is kept for a period which IBM Storage Insights uses to provide predictive analysis for future capacity and performance demands. For more information, see “Data retention periods” on page 6.

Note: IBM Storage Insights capacity planning can provide projections up to two years for block-level and file-level storage subsystems.

The data collector gathers and transfers data to the IBM Storage Insights secure portal for analysis and provides valuable reports. The following IBM Storage Insights reports are useful for capacity planning:

- ▶ Block capacity projection
- ▶ File capacity projection
- ▶ Reclaimable capacity guidelines

The preferred usage of capacity planning reports is explained in the following scenario.

Company A purchased IBM storage and an IBM Storage Insights subscription. A few months later, when the business and storage capacity demands grew, the growth raised a concern about sizing review for the higher IT management. The CIO of the company wants the storage team to predict current and future capacity utilization for 3 - 24 months. The company must not run out of storage capacity. Consider what the storage team can do to cater to this critical business need. The CIO wants to avoid the following things:

- ▶ Continuing to add storage blindly
- ▶ Needless spending on storage
- ▶ Wasting money to buy unneeded storage
- ▶ Performance issues to become a constraint for the business

In a typical, routine scenario, the storage team cannot do much. For action to be possible, they must take manual readings of the data for several months, use a scripting technique, and spend several laborious hours extracting this information, or purchase an on-premises monitoring tool.

IBM Storage Insights can answer all the challenges in a matter of minutes. After registration, the service is up and running in less than 5 minutes with actionable insights in less than 30 minutes.

IBM Storage Insights reports not only give the current statistics but also provide valuable insights into the future growth projections. It can also help address the following concerns:

- ▶ What is the current block and file level capacity utilization?
- ▶ How much storage has been used in the last 3 months (file and block level)?
- ▶ What will be the future storage utilization in the next 3 months (can be up to 24 months)?
- ▶ Will storage space be filled soon?
- ▶ What is the data growth rate? Is it steady growth or does it have some abnormal behavior?

Storage capacity projection is the analysis that an organization carries out to assess the storage growth and utilization trends in its environment.

IBM Storage Insights provides several built-in reports that might prove useful for the CIO's capacity projection and planning requirements.

The standard dashboard for the storage infrastructure has different views for capacity, performance, and tier planning. Both block-level and file-level storage trend reporting can be done with IBM Storage Insights. Projections with IBM Storage Insights predictive data analysis can be seen after only a few hours of data collection. The longer the data is collected, the more accurate the predictions IBM Storage Insights can make.

Figure 3-1 shows how the dashboard might look when future projections are disabled.

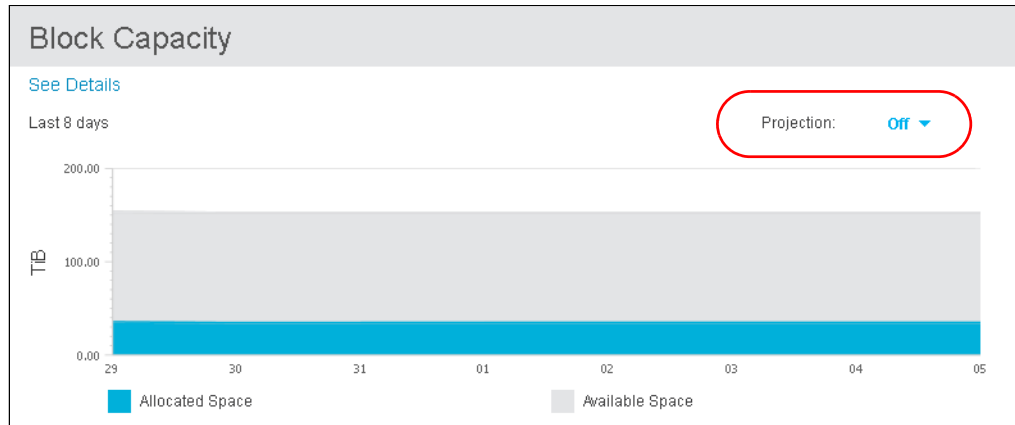


Figure 3-1 Future projection is disabled

Here are some reports for different periods of future projections under block and file capacity:

- ▶ Figure 3-2 shows block capacity projections for 3 months.

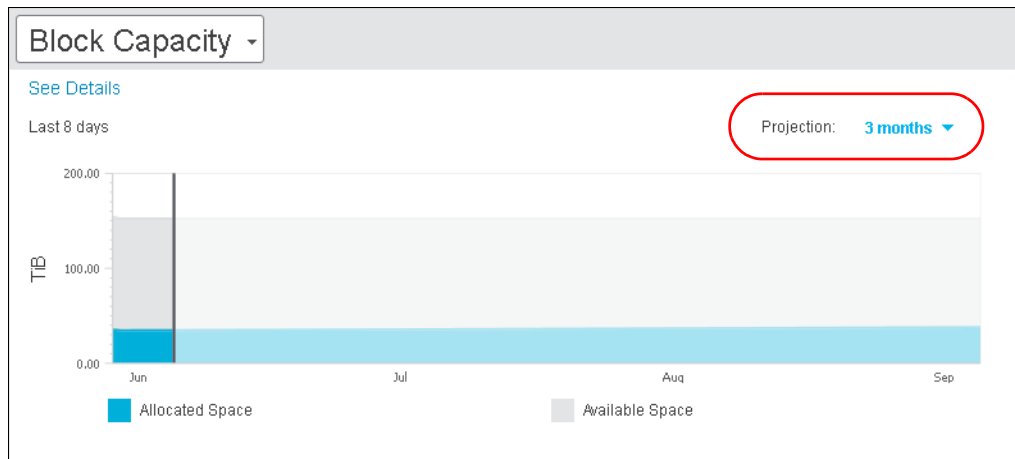


Figure 3-2 Three-month predictive analysis for block capacity

- ▶ Figure 3-3 shows block capacity projections for up to one year.

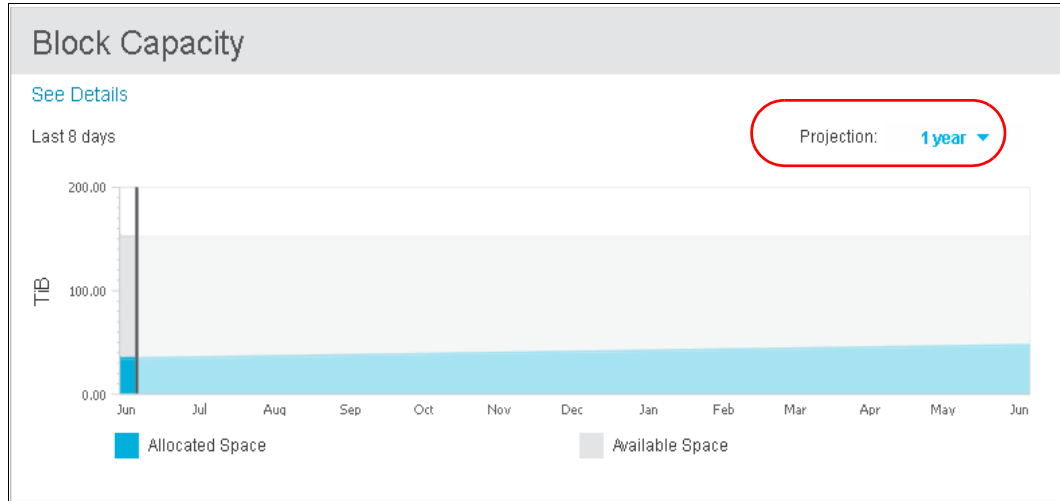


Figure 3-3 One-year block capacity predictive analysis

- ▶ Figure 3-4 shows a block capacity projection for two years.

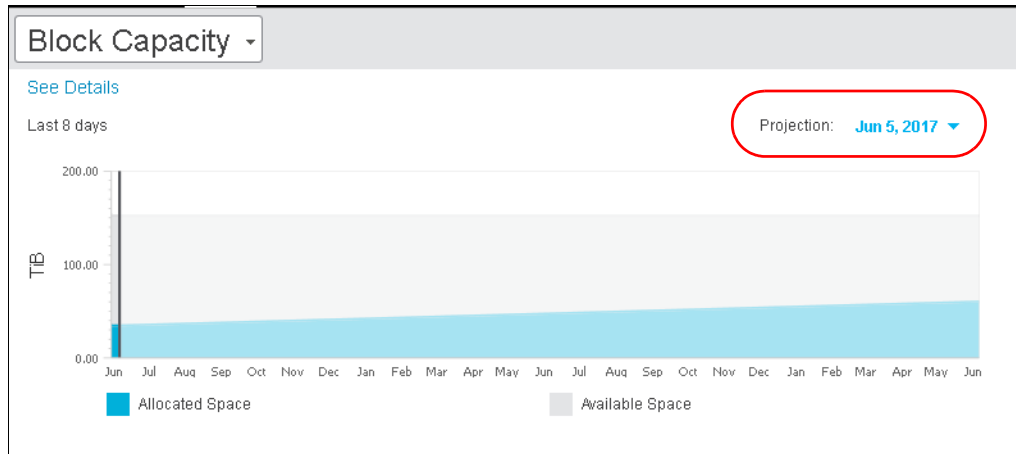


Figure 3-4 Two-year block capacity predictive analysis

The same predictive analysis technique can be used to extract reports for file type capacities:

- ▶ Figure 3-5 on page 25 shows file capacity projections for three months.

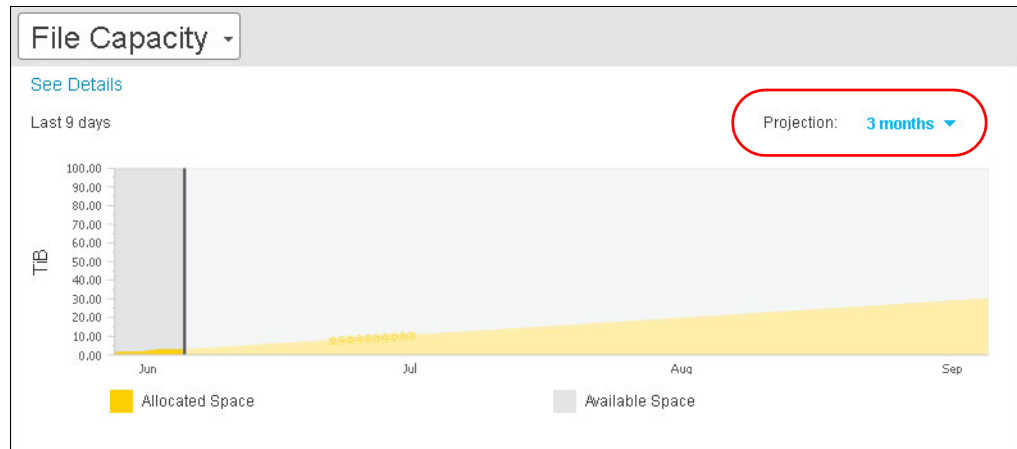


Figure 3-5 File capacity projections for 3 months

- Figure 3-6 shows file capacity projections for up to one year.

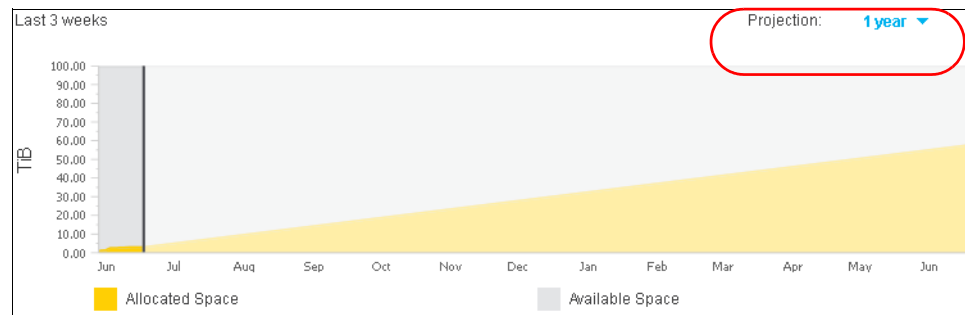


Figure 3-6 File capacity projections for up to one year

With these reports extracted in minutes, the team can now respond to the CIO about the capacity utilization and trends. This response helps the CIO make decisions about when storage is needed and how much storage is required.

3.2 Increasing storage utilization

This scenario explains how you can use the IBM Storage Insights interface to list unused capacity in your storage infrastructure. When you know which spaces are underutilized, you can decide whether to reclaim space.

Unused capacity is storage that is idle in your environment without being used by any application. When IBM Storage Insights runs its assessment, it displays volumes that come under the reclaimable capacity criteria as “reclaimable space” to help you improve storage utilization.

IBM Storage Insights analyzes the data and looks for zero workload volumes. Zero workload volumes are volumes that have no I/O on them for at least two weeks.

IBM Storage Insights also lists volumes that are orphaned. Orphaned volumes are volumes that are not mapped to any host and that have no workload running on them.

Note: Volumes that are included in a replication relationship are not shown as reclaimable.

After it identifies zero workload and orphaned volumes, IBM Storage Insights makes recommendations for reclaimable space. The IBM Storage Insights storage reclamation module can help answer the following questions.

- ▶ Which and how much storage is sitting idle in my infrastructure?
- ▶ Are there any orphaned volumes in my storage infrastructure?
- ▶ Is there any reclaimable capacity in my storage environment?
- ▶ How much unused storage capacity can be reclaimed?

Note: IBM Storage Insights reviews the last two weeks of data to provide recommendations about reclaimable capacity.

Consider this scenario. Company B has storage with huge amounts of space that is assigned to production, testing, and development environments. The company is quickly running out of usable storage space and is unable to keep track of thousands of volumes that are being allocated and expiring during short periods of time.

With IBM Storage Insights in place, the storage team can quickly and easily determine which volumes are unused and how much capacity can be reclaimed. As shown in Figure 3-7, it gives an overview of the reclaimable capacity and provides a pie chart of reclaimable versus non-reclaimable storage capacity.

Figure 3-7 shows a pie chart with recommendations for reclaimable versus non-reclaimable space.

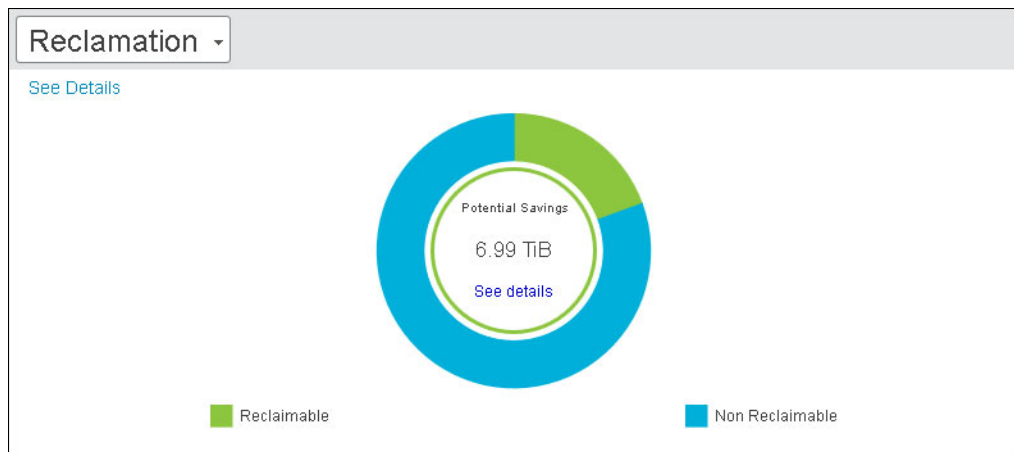


Figure 3-7 Recommendations for reclaimable versus non-reclaimable space

There are a few other views that show reclaimable capacity. The view in Figure 3-8 is broken down by storage subsystem. If you have multiple storage subsystems in your environment, IBM Storage Insights can help you identify which storage has more reclaimable capacity according to the predefined algorithms.




Figure 3-8 Reclaimable capacity that is separated by storage subsystem

Note: IBM Storage Insights can start providing results with only a few minutes of the data but it is a leading practice to have 2 weeks of historical data for better accuracy.

Reclamation reports also provide views with useful recommendation on saving expensive storage capacity. As shown in Figure 3-9, a breakdown by tier of reclamation recommendations clearly shows which tier is the most unproductive and has higher reclaimable capacity.

In this scenario, you can see that Tier 2 is the most unproductive. By reclaiming unused capacity from it, you can save approximately 3 TB. If you reclaim unused storage from the entire storage subsystem, you can reclaim approximately 7 TB, which is significant.

Figure 3-9 shows a breakdown of unused storage capacity by tier.




Figure 3-9 Unused storage capacity breakdown by tier

With the IBM Storage Insights reclamation views, you can easily determine which part of storage is sitting idle and is not being used. This view is helpful when you want to recover unused storage capacity from different departments and applications.

Note: If you know that certain volumes will remain unused and you do not want them to appear in the reclamation assessment, IBM Storage Insights allows you to ignore volumes that the IBM Storage Insights reclamation algorithm determines to be unused. These volumes will be excluded from analysis for all future assessments for reclamation.

Figure 3-10 shows how to exclude volumes from reclamation analysis.

Name	Tier	Time of Last IO	Pool	Storage System	Capacity (...)	Allocated Space (GiB)
SaaS-smcho-vol01		observed	CET_XIV_pool	XIV-2812-7825410-IBM	961.58	0.00
SaaS-smcho-vol02		observed	CET_XIV_pool	XIV-2812-7825410-IBM	961.58	0.00
SaaS-110718-Vol11		observed	SaaS-Pool01	XIV-2812-7825410-IBM	560.92	0.00
epic47492-kk5		observed	Thin_pool	XIV-2812-7825410-IBM	560.92	1.25
epic47492-kk6	Tier 5	None observed	Thin_pool	XIV-2812-7825410-IBM	560.92	4.50

Figure 3-10 Excluding volumes from analysis

Another method is to extract a reclamation report on your system. Then, you can use the listing to manually remove idle volumes from your storage subsystem.

As shown in Figure 3-11, several types of reports (including CSV, HTML, and PDF) can be exported from IBM Storage Insights for manual review and processing.

Tip: When you export a list of unused volumes from IBM Storage Insights, consider creating a script to delete them.

Figure 3-11 shows reclaimable capacity report export options.

Name	Tier	Time of Last IO	Pool	Storage System	Capacity (...)	Allocated Space (GiB)	Hosts	Applications
tpcblade3-14_tpcdb2		observed	v7ku_block_pool1	Storwize V7000-2078-IFS-tpc...	1,024.00	316.26		
SaaS-smcho-vol01		observed	CET_XIV_pool	XIV-2812-7825410-IBM	961.58	0.00		
SaaS-smcho-vol02	Tier 5	None observed	CET_XIV_pool	XIV-2812-7825410-IBM	961.58	0.00		
SaaS-110718-Vol11	Tier 5	None observed	SaaS-Pool01	XIV-2812-7825410-IBM	560.92	0.00		
epic47492-kk5	Tier 5	None observed	Thin_pool	XIV-2812-7825410-IBM	560.92	1.25		
SaaS-110718-Vol10	Tier 5	None observed	SaaS-Pool01	XIV-2812-7825410-IBM	560.92	0.00		

Figure 3-11 Options to export reclaimable capacity reports

Figure 3-12 shows a download pop-up window for list of reclaimable volumes.

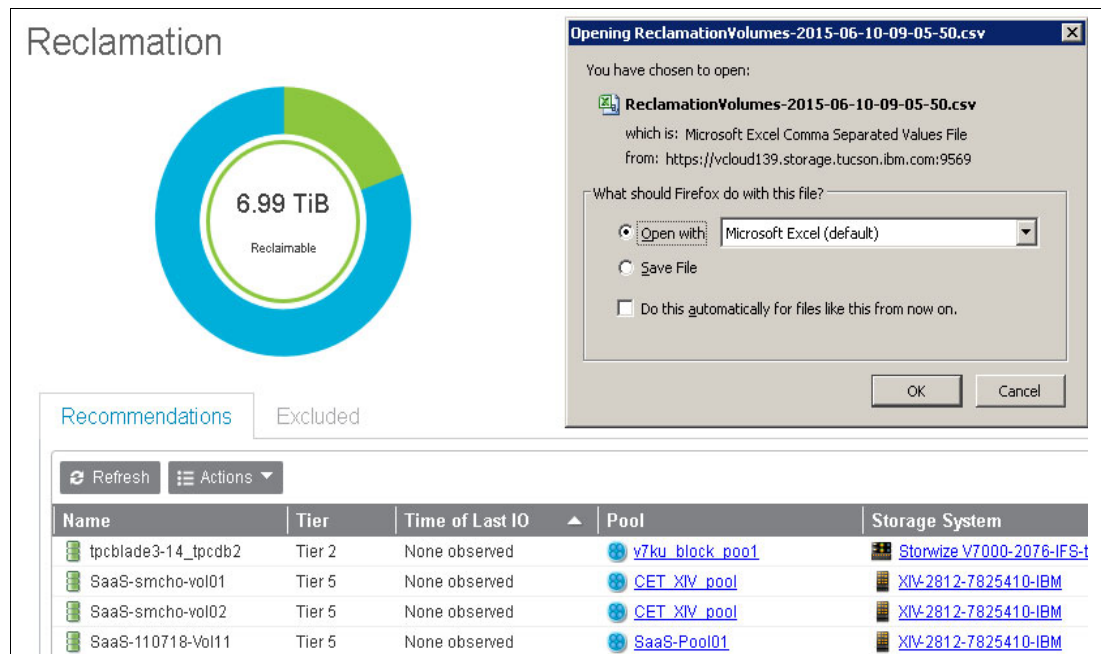


Figure 3-12 Extracting a list of reclaimable volumes from IBM Storage Insights

Also note that the volumes that are mapped in a relationship do not show up as reclaimable. This configuration helps you avoid considering target replication volumes that usually have low I/O as unused.

3.3 Planning data placement

In this scenario, you have many storage subsystems with many volumes. Most changes are not tracked and after a few years, you lose track of capacity allocation.

You want to know that you are using the expensive storage pools for the most performance-demanding workloads.

It is a leading practice to place critical data on high-performing storage pools and noncritical data on low-performing and inexpensive ones. Recommendations are generated by IBM Storage Insights for tiering the data according to the thresholds that you define.

By using the IBM Storage Insights interface, you can plan to optimize your data placement by defining performance tiers in your environment. Here is an example of how to maximize the benefits of this feature.

First, you need to define tiers in the IBM Storage Insights interface. To do this task, you must define the performance thresholds for each tier. You can define up to a maximum of nine tier levels in your entire storage infrastructure.

For more preferred practices and more detailed information about defining tiers, see “Configuring pools tiers and defining their tiering thresholds” on page 16.

As shown in Figure 3-13, you can use the tier definition pane to define tiers along with their thresholds in IOPS or I/O density rates.

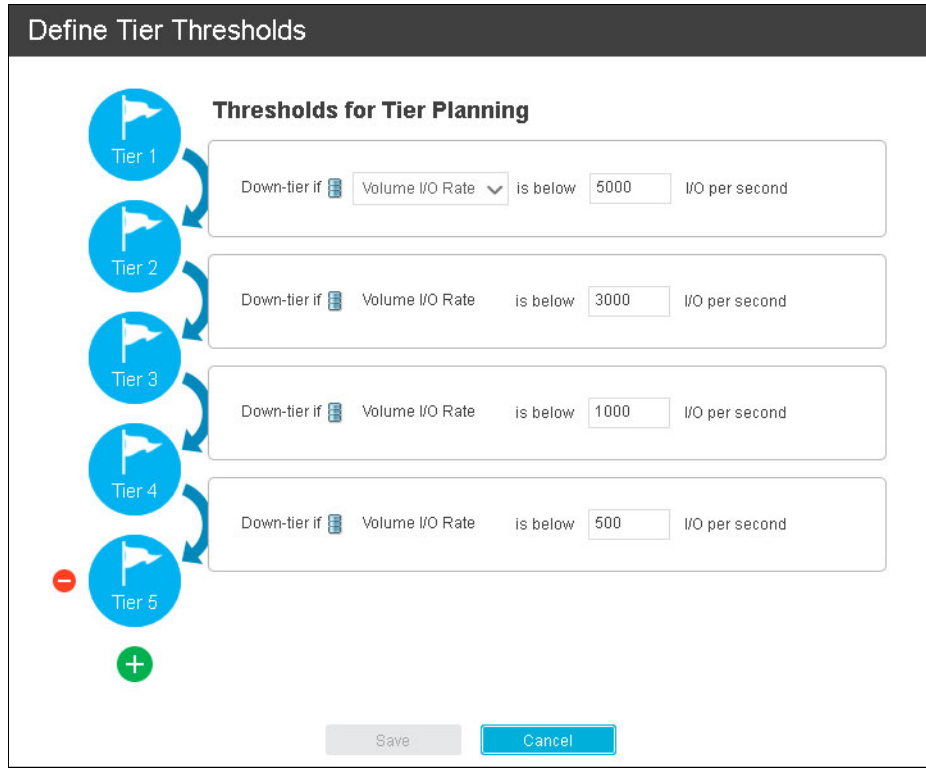


Figure 3-13 Tier definition pane

Tier levels are performance thresholds that are defined to demarcate the performance levels that are expected from each level. So, whenever you assign a pool to a tier level, IBM Storage Insights runs its analysis to give recommendations about optimization of data placement.

For example, consider that you have an environment with the following three tiers defined:

- ▶ Tier 1 is 5,000 IOPS and above
- ▶ Tier 2 is 2,500 - 5,000 IOPS
- ▶ Tier 3 is 2,500 IOPS and less

You want to get the recommendations for your storage infrastructure. IBM Storage Insights automatically recommends the volume placement according to the criteria so you can retier the data to optimize data placement.

Tip: Tier planning can be used to identify whether to promote or demote the volume to specified tiers.

As shown in Figure 3-14, when you define tiers, they appear with color coding.

Name	Capacity (GiB)	Pools	Threshold Type	Range
Tier 1	4,885.68	7	I/O Rate	> 5,000 IOPS
Tier 2	20,335.23	17	I/O Rate	3,000-5,000 IOPS
Tier 3	11,429.99	17	I/O Rate	1,000-3,000 IOPS
Tier 4	7,931.00	13	I/O Rate	500-1,000 IOPS
Tier 5	100,239.71	48	I/O Rate	< 500 IOPS

Figure 3-14 Tier thresholds defined

With the tiers defined, you can view the recommendations from the IBM Storage Insights interface. As shown in Figure 3-15, you can see current versus recommended comparison for allocation according to the tier definition.

Figure 3-15 shows current versus recommended tier planning.




Figure 3-15 Current versus recommended allocations

Using these tiering recommendations, you can optimize your data placement that is an optimal combination for cost and performance. A more detailed tier planning view (Figure 3-16) can be used for more specific data placement recommendations.

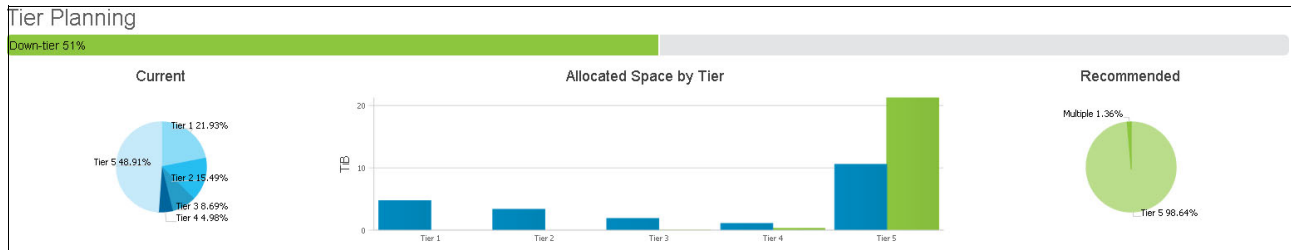


Figure 3-16 Detailed tier planning

3.4 Performance reporting

In any organization, a storage subsystem’s capability to sustain heavy workloads is the key performance indicator (KPI) of its potential. If a storage subsystem’s performance cannot be

gauged, this lack of information can become a major gap in understanding its efficiency to manage the workload that is being run on it.

IBM Storage Insights provides several options for performance reporting.

Consider the example of company C, who is facing severe performance issues in its storage environment and is not sure how to locate the issue. It wants answers to the following questions:

- ▶ Are there performance issues on the storage subsystem?
- ▶ Which volumes have the most performance-demanding workload running on them?
- ▶ Since when the storage performance is being maxed out?
- ▶ Is there a higher tier in the environment that can cater to this workload?

All these queries can be responded to with the help of IBM Storage Insights. After the basic setup is done, the company can have access to all the performance data and reporting, including a high-level overview of the entire storage infrastructure performance. See Figure 3-17.




Figure 3-17 High-level overview of the storage performance

The company can also have a detailed performance view that is broken down by storage subsystem. This view can be used to identify which storage subsystem is under more stress and, which has the highest utilization.

Figure 3-18 shows a detailed performance view by storage system.

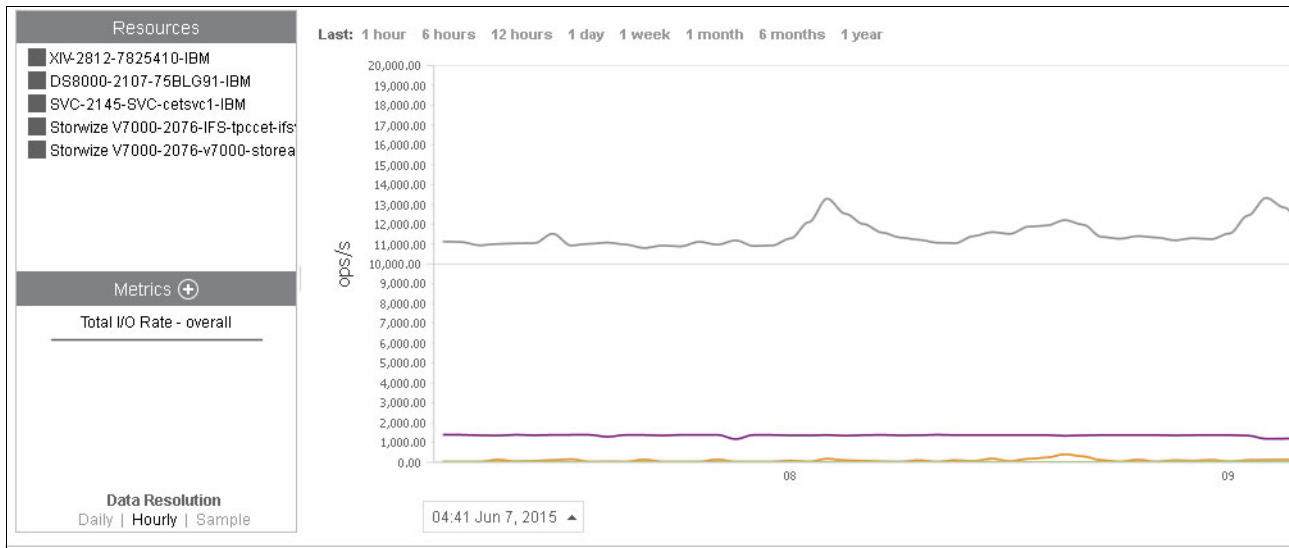


Figure 3-18 Detailed performance view by storage system

You can extract a report for top volume utilization and determine which volume is getting the most performance-demanding workload. As shown in Figure 3-19, there are five volumes with the highest utilization in the environment. This view can help you see the volumes that are performing the most workload in the environment. This view can help you find the performance issues in your environment.

Figure 3-19 shows top volume utilization.




Figure 3-19 Top volume utilization

Another way to find the performance issues is a listing by server. As shown in Figure 3-20, it shows volumes that are used by tpccblade6. This report can be used to determine whether there is any performance issue with the storage subsystem or any of its volumes.

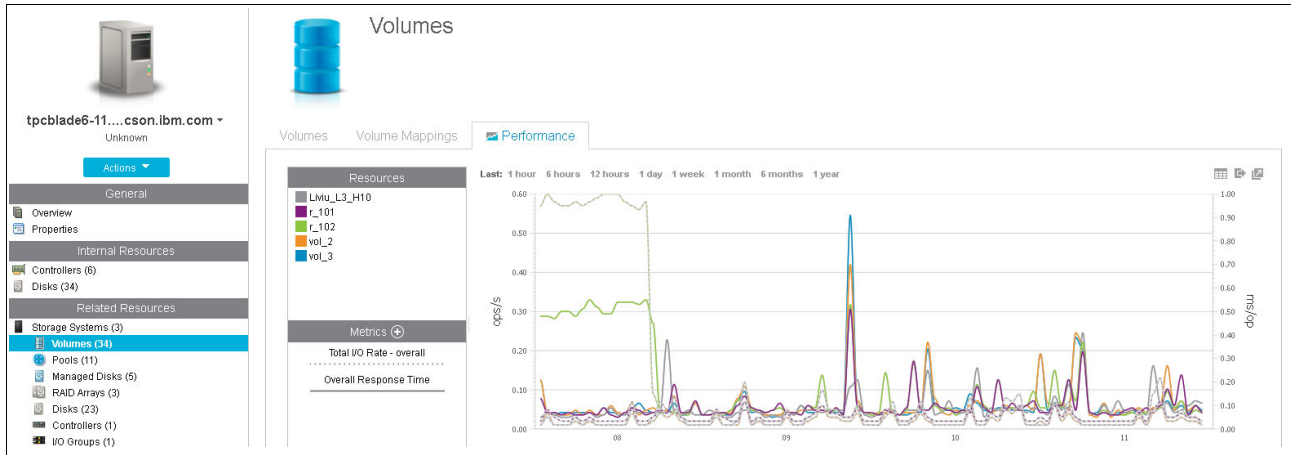


Figure 3-20 Top volume utilization for a selected server

These reports can help you identify the key areas in Company C's storage infrastructure, which can provide useful insights about your storage environment. This analysis can be useful when you work with other teams of servers and application domains find the issue.

3.5 Simplifying reporting

Reporting is an essential part of any application, but tool reporting must be simple and resilient enough to fulfill the information requirements and also be easy to understand.

IBM Storage Insights provides simple and useful reporting features. It provides the benefit of live action troubleshooting rather than slow and painful conventional reporting. You can simply log on to the IBM Storage Insights interface and can view predefined graphical and pictorial representation about the health of your storage infrastructure immediately.

IBM Storage Insights can help you address everyday queries for your storage subsystem in a matter of minutes, which traditionally required much manual data collection and hours of laborious efforts by the storage administration team.

Consider a scenario where the IT manager wants the following information:

- ▶ What is the current data growth rate?
- ▶ How much capacity is being saved with thin provisioned volumes?
- ▶ Which departments are using the most storage capacity?

These questions are few of the many queries that can be easily addressed by IBM Storage Insights.

As shown in Figure 3-21, you can see the projections for up to three months of data growth. This report is predefined and is available in the IBM Storage Insights web interface for quick reporting.

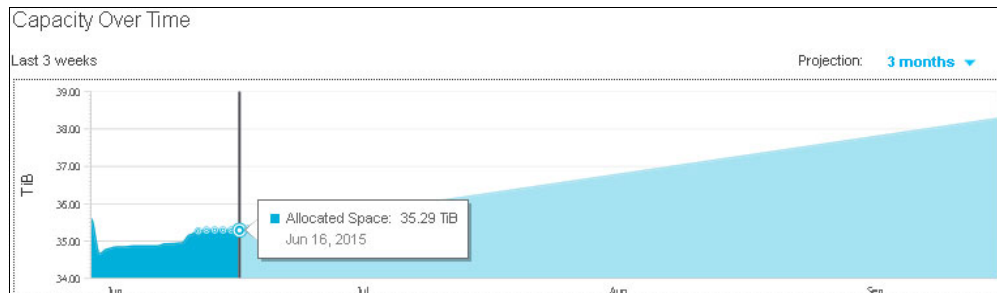


Figure 3-21 Projected capacity growth over time

As shown in Figure 3-22, the storage infrastructure benefits from using thin-provisioned volumes. This graph shows a saving of almost 97 TB. This information can be an important tool to help convince stakeholders to use thin provisioning.

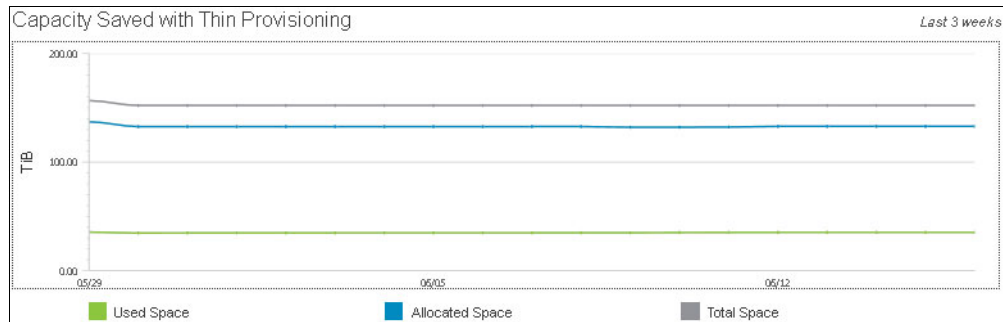


Figure 3-22 Capacity benefit of thin-provisioned volumes

The reports can also be used to provide more information that every IT manager wants to see. If you cannot tell which department is using how much of your storage resources, you probably will not be able to justify your value to the organization. This reporting feature helps you not only determine storage utilization for each department but also show it to them.

As shown in Figure 3-23, you can see capacity from different high-performance storage subsystems that are assigned to the application (in this case, the Sales portal). This way, you can see how much storage is being used by which application and whether to promote the workload to a higher tier.

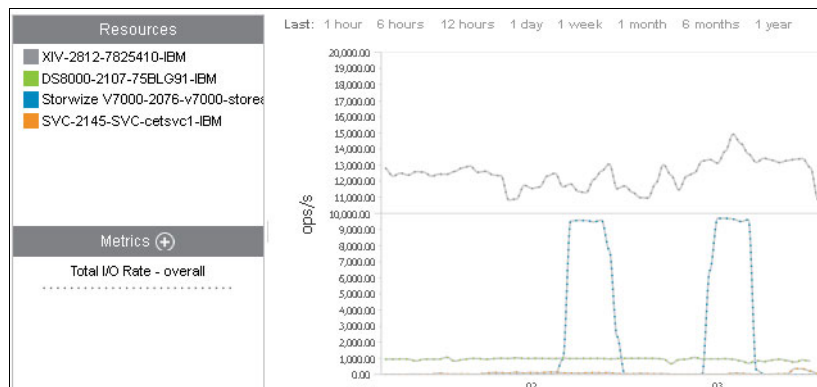


Figure 3-23 Application use for high-performance storage subsystems

3.6 Using agentless servers

The IBM Storage Insights environment is intended to be easy to use and flexible for the user.

Most reporting products for storage subsystems require an agent to be deployed on hosts to provide useful reports. Agents are small pieces of code that can run on a physical or virtual server and help you probe useful information and statistics from the managed endpoint or host.

No matter how efficient the agents are, they add overhead on the hosts, which can negatively affect system performance.

IBM Storage Insights gathers important server statistics without putting any extra workload on the host. The information that is provided by the agentless end points to IBM Storage Insights can be used to actively troubleshoot storage-related capacity and performance issues.

Figure 3-24 shows agentless server statistics.




Figure 3-24 Agentless server statistics view

You can use the agentless server feature to view details about each added host and troubleshoot issues that pertain to specific servers and the storage pools that are assigned to them.

Figure 3-25 shows the agentless server view by pools.




Figure 3-25 Agentless server view by pools

You can also get performance statistics to find out whether there is any performance bottleneck on the host ports. Figure 3-26 clearly shows throughput of each Fibre Channel port, which can provide an overview of the health of the port throughput.

Figure 3-26 shows the agentless server view by ports.




Figure 3-26 Agentless server view by ports

3.7 Application groups and departments

Most organizations are composed of departments, each with a role that needs to be fulfilled in line with the company's vision and goals. The IT department is also part of the team and has its own role and value. The role of IT is to support and provide technological values that are tightly integrated with the rest of the business to ensure smooth business.

IBM Storage Insights provide you the opportunity to show value to your in-house customers that can be your application or database teams, be another department or even another company. By modeling the applications in your organization, you create applications and application subcomponents, organize them into hierarchies, and associate them with the storage resources that they use.

By defining applications in IBM Storage Insights, you can monitor the space consumption and the performance statistics of the storage infrastructure that are critical for your environment. You can also create applications and map them to the departments or subdepartments in your organization. This way, you can know how much resource is being allocated to each department.

Consider the following example.

When you want to monitor the space utilization of a particular department, you can associate the application with the department and can obtain the consumption of storage resources that are being utilized.

You can also use this function to chargeback or show value to each department by using the information that is received from creating and mapping applications to each department.

As shown in Figure 3-27, department “Sales” is defined in IBM Storage Insights and you can view how much capacity and performance workload is being run by the application. This information can be used to better understand the department’s application performance and capacity requirements and project growth trends. This information helps you plan future purchases for storage infrastructure for this application.

Figure 3-27 shows sales application and department mapping.

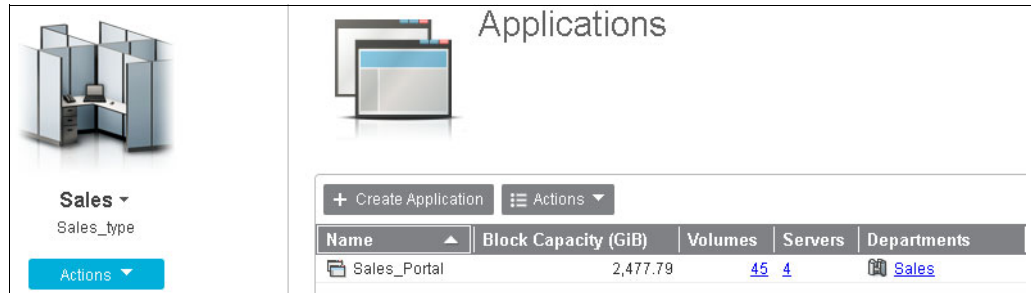


Figure 3-27 Sales application and department mapping

You can help address storage questions from each department that pertain to resource utilization, including the following questions:

- ▶ How much capacity is my department using?
- ▶ How many servers and volumes are used by my application?
- ▶ What are the top volumes that are used by my department?

These questions can be answered by using the IBM Storage Insights application and department reporting functions. As shown in Figure 3-28, each volume of the “Sales” application can be seen with different response times and I/O. This information can help the IT team to make the “Sales” department aware of congestion and slow application responses and explain what must be done to improve performance.

Figure 3-28 shows top volume I/O and response times for the Sales department.

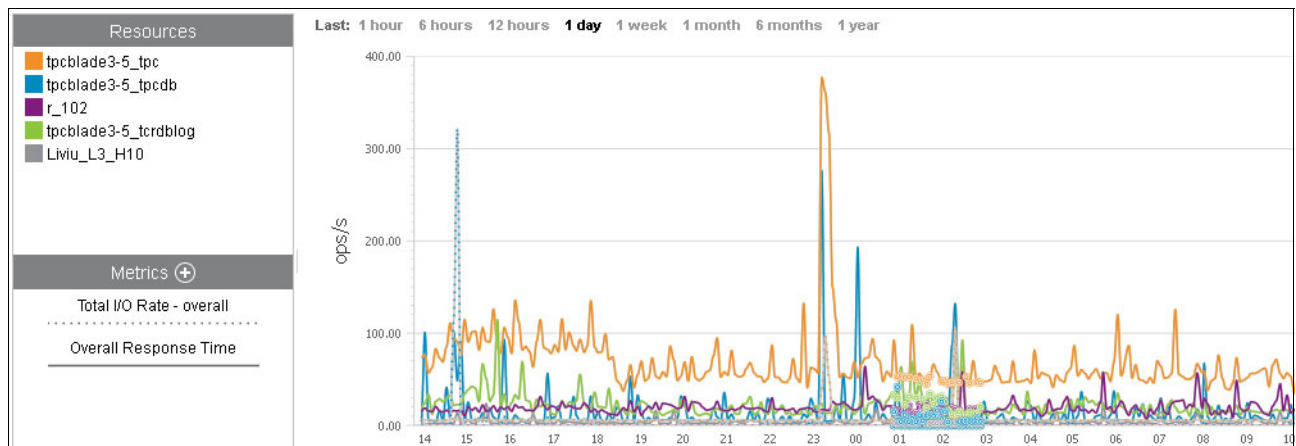


Figure 3-28 Top volume I/O and response times for the Sales department

In other views, you can see how much capacity is being allocated to each application or department. Figure 3-29 shows the pools that comprise the “Sales Portal” application that is mapped to the Sales department. It shows physical storage allocation from each storage subsystem for the Sales department.

Name	Storage System	Physical Allocation (%)	Capacity (GiB)
CET_FB_P4	DS8000-2107-75BLG91-IBM	77%	1,480.00
Irina_XIV_pool	Storwize V7000-2076-v7000-st...	55%	480.00
mdiskgrp0	Storwize V7000-2076-v7000-st...	35%	3,620.00
ThickChild05	Storwize V7000-2076-v7000-st...	34%	50.00
CET_server_pool	SVC-2145-SVC-cetsvc1-IBM	63%	1,727.00
CET_XIV_pool	XIV-2812-7825410-IBM	15%	28,783.36
CURRENCY_POOL	XIV-2812-7825410-IBM	1%	7,772.79
garbage_pool	XIV-2812-7825410-IBM	0%	4,583.54
Iuliu_RegularPool1	XIV-2812-7825410-IBM	0%	160.26
PERF_POOL	XIV-2812-7825410-IBM	0%	11,186.40
pool_xiv	XIV-2812-7825410-IBM	2%	5,609.23
TPC_TEST	XIV-2812-7825410-IBM	12%	14,391.68
TPC_TEST_	XIV-2812-7825410-IBM	3%	1,121.85

Figure 3-29 Physical storage subsystem allocation for the Sales department

IBM Storage Insights helps you filter your report according to your needs. As can be seen in Figure 3-30, the Sales department can be filtered and viewed by pools, controllers, servers, and so on. These filters are a few of the many filters that you can use to simplify the final report according to your requirements.

Figure 3-30 shows the Sales department detailed views.

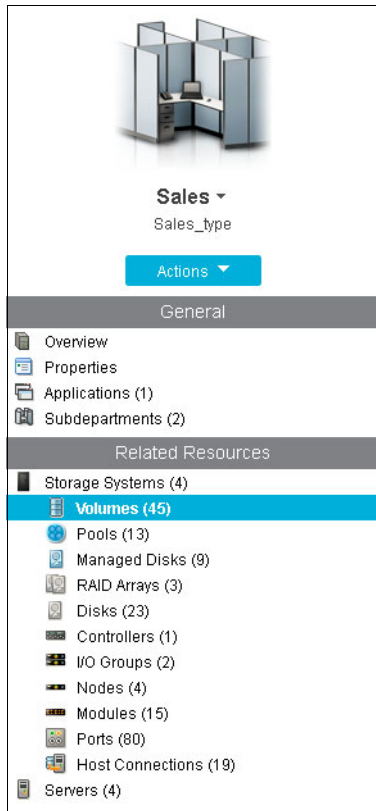


Figure 3-30 Sales department detailed views

These views can help you understand what resources that each department and its applications are using. You can then make an action plan to meet performance and capacity growth requirements based on IBM Storage Insights recommendations and reports.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

Online resources

These websites are also relevant as further information sources:

- ▶ IBM Marketplace
<https://www.ibm.com/marketplace/search/us/en-us?terms=IBM+Spectrum+Control+Storage+Insights>
- ▶ IBM Spectrum Control Storage Insights documentation
http://www.ibm.com/support/knowledgecenter/SSQR88/com.ibm.spectrum.si.doc/tpch_saas_welcome.html?lang=en
- ▶ IBM Spectrum Control Storage Insights product page
<http://www.ibm.com/systems/storage/spectrum/insights>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5231-00

ISBN 0738454370

Printed in U.S.A.

Get connected

