

IBM Storage Insights

Security Guide



Note:

Before using this information and the product it supports, read the information in [“Legal notices” on page 51](#).

This edition applies to the current version of IBM Storage Insights (product number 5725-U02) and to all subsequent versions until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2017, 2022.**

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

Contents

- About this guide..... v**
 - Who should read this guide.....v

- Chapter 1. Summary..... 1**

- Chapter 2. What is the data collector.....3**

- Chapter 3. How is the metadata protected.....5**

- Chapter 4. What types of metadata are collected.....9**

- Chapter 5. How long is the metadata kept..... 11**

- Chapter 6. Who can access the metadata..... 13**
 - Metadata access controls and authorization..... 13
 - Metadata access for resolving issues.....13
 - IBM Support access for troubleshooting your tickets..... 14
 - Metadata access for quality improvements..... 14
 - Data backup and restore.....14
 - Requesting the deletion of personal information..... 15

- Chapter 7. Asset, capacity, and configuration metadata.....17**
 - Block storage system metadata.....18
 - Block volumes metadata..... 19
 - Block pools metadata..... 20
 - I/O groups, nodes, and ports metadata..... 21
 - Disks and managed disks metadata.....23
 - Block device adapters metadata..... 24
 - Block host adapters metadata..... 24
 - Enclosures metadata..... 24
 - File storage system metadata..... 25
 - File system metadata..... 25
 - Fileset metadata..... 26
 - File shares metadata..... 26
 - File system pools metadata.....26
 - Network shared disks metadata..... 27
 - File nodes metadata..... 27
 - Object storage systems metadata.....27
 - Switches metadata..... 29
 - Fabrics metadata..... 30
 - Groups metadata..... 30
 - Hosts metadata.....32
 - Virtual machines metadata..... 33

- Chapter 8. Performance metadata..... 35**
 - Performance metadata for storage systems that run IBM Spectrum® Virtualize..... 36
 - Performance metadata for DS8000..... 38
 - Performance metadata for XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R..... 40

| | |
|---|-----------|
| Performance metadata for IBM Spectrum Scale..... | 41 |
| Performance metadata for Dell EMC Unity, VMAX, and VNX..... | 42 |
| Performance metadata for Hitachi VSP storage systems..... | 45 |
| Performance metadata for NetApp storage systems..... | 45 |
| Performance metadata for Pure Storage systems..... | 47 |
| Performance metadata for switches..... | 47 |
| Performance metrics for hosts..... | 48 |
| Performance metrics for virtual machines..... | 49 |
| Legal notices..... | 51 |
| Privacy policy considerations | 52 |
| Trademarks..... | 53 |
| Index..... | 55 |

About this guide

In IBM Storage Insights Pro and IBM Storage Insights, detecting and resolving issues in a storage environment has never been easier. It combines cognitive storage management capabilities with a simplified yet robust IBM support experience to help you spend less time troubleshooting storage problems and more time planning for your future storage needs.

Who should read this guide

This publication is intended for administrators or IT professionals who deploy IBM Storage Insights Pro or IBM Storage Insights and want to learn more about security and data collection.

Administrators should be familiar with the following topics:

- General procedures for installing software on Microsoft Windows, AIX®, and Linux®.
- Storage area network (SAN) concepts.
- Storage resources and management concepts.

Chapter 1. Summary

The concerns that customers might have about deploying a data collector on-premises and processing and storing metadata off-premises are discussed.

IBM Storage Insights Pro and IBM Storage Insights are cloud service offerings that use a light-weight application that is called the data collector to collect and send asset, configuration, capacity, and performance metadata for analysis to an IBM Cloud data center and for presentation in the GUI.

Important: The security policies for collecting, sending, accessing, protecting, and storing metadata for IBM Storage Insights Pro and IBM Storage Insights are identical.

The key differences between both cloud service offerings lie in the exclusive features that IBM Storage Insights Pro provides to its subscribers, such as capacity planning analysis, reclamation analysis, and tiering analysis, and in the access to the metadata that is presented in the GUI for the cloud service offerings. In IBM Storage Insights Pro, subscribers have access to all of the metadata in the GUI, whereas in IBM Storage Insights, users have access to key capacity and performance metadata in the GUI, and IBM Support has read-only access to the set of metadata that they need to troubleshoot and close support tickets.

To discuss the security concerns that customers might have, the following questions are answered:

- What is the data collector?
- How is the metadata protected?
- What types of metadata are collected?
- How long is the metadata kept?
- Who can access the metadata that is collected?

Lists of the asset, capacity, and configuration metadata and the performance metadata that is collected and stored about your storage systems are also provided.

Related reference

Asset, capacity, and configuration metadata


The data collector collects and stores asset, capacity, and configuration metadata for block, file, and object storage systems and their resources. A list of the supported storage systems is provided.

Performance metadata

The data collector collects and stores performance metadata for IBM block storage systems and non-IBM block storage systems and it collects and stores file system and node performance metadata for IBM Spectrum Scale storage systems.

Chapter 2. What is the data collector

The data collector is the application that collects and delivers the metadata that is analyzed and presented in the GUI.



The data collector is a light-weight application that is installed on a server in your data center. It sends the metadata that is collected about your storage systems, such as asset, configuration, capacity, and performance metadata, from your data center to your instance of IBM Storage Insights Pro or IBM Storage Insights, which is in an IBM Cloud data center.

Important: Outbound metadata is sent by the data collector to a single, unique address, which is the IBM host name and port of your instance. This means that when you configure your firewall to send the metadata, you open a single path to a well-defined and secure endpoint.

In a matter of minutes, you can install the data collector and when you add the storage systems that you want to monitor, you get the capacity and performance insights that you need to monitor your data center. Because the metadata that IBM Support needs to investigate and close tickets is also collected, you can also upload logs automatically when you create or update tickets and IBM Support can access and investigate the metadata to resolve any issues that you might have.

Credentials for connecting to storage systems: To add and collect metadata from the storage systems that you want to monitor, you must provide the storage system's credentials. Depending on the type of storage system that you add for monitoring, you can provide the name and password of a user with privileges to collect the metadata, or an SSH user and SSH key. The credentials that are provided are encrypted before they are stored in the database for the instance, and the database is also encrypted. In addition, most storage systems support the creation of users with read-only roles, who can't make any changes to the configuration of the storage system.

Supported operating systems: Data collectors can be installed on servers or virtual machine that run AIX, Linux, or Windows (64-bit systems only). On the server or virtual machine, you must provide at least 1 GB of RAM and 3 GB of disk space.

Security certification: IBM Storage Insights, based on regular audits, has [ISO/IEC 27001 Information Security Management certification](#). Annually, the following audits are conducted: two KPI audits, one external Veritas ISO27001, 27017, and 27018 audit, and one IBM internal audit for each ISO2700x.

Key security characteristics

To ensure that metadata is collected securely, the data collector has the following characteristics:

In-built security

Communication with other entities, such as storage systems in the local data center and the IBM Storage Insights service in the IBM Cloud data center are initiated solely by the data collector. The data collector does not provide any remote APIs that might be used to interact with the data collector.

One-way communication

The data collector sends metadata out of your network to your instance of IBM Storage Insights Pro or IBM Storage Insights. Communication is outbound only; the data collector can't receive data from the internet or any other entity in your network. Here's how the one-way communication works:

1. The data collector sends out a request for work.
2. IBM Storage Insights responds with a data collection request.
3. The data collector communicates with the storage resource or starts a log collection.

Secure transmission

All communication between the data collector and IBM Storage Insights Pro or IBM Storage Insights in the IBM Cloud data center uses encryption based on HTTPS.

The communication that the data collector initiates with the server where it is installed, and the communication between the server and IBM Storage Insights Pro GUI or IBM Storage Insights GUI. HTTPS connections use certificates issued by Cloudflare, Inc. (issuer common name "Cloudflare Inc ECC CA-3") and use TLS 1.2 and TLS 1.3 with 256-byte keys.


Tip: Because HTTPS connections are used, the data collector can run on any computer that can access the internet over an outbound TCP connection to port 443. Port 443 is the standard port for HTTPS connections.

Chapter 3. How is the metadata protected

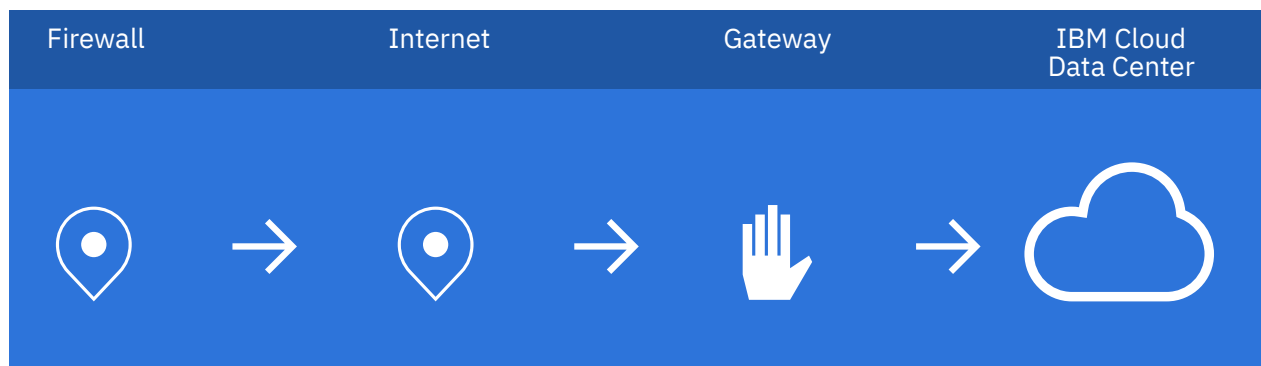
End-to-end protection is provided for the metadata that is collected, delivered, and stored for your IBM Storage Insights Pro or IBM Storage Insights service in the IBM Cloud data center. This protection includes meeting the requirements of the General Data Protection Regulation (GDPR).

Metadata collection, delivery, and storage in the cloud

To transform the metadata into insights and present them in IBM Storage Insights Pro or IBM Storage Insights, the data collector forwards metadata packages for analysis and storage to the IBM Cloud data center (located in Washington, D.C.).



To keep the metadata package safe on its journey to the cloud, the data collector uses Hypertext Transfer Protocol Secure (HTTPS), which encrypts the metadata and sends the metadata package through a secure channel to the IBM Cloud data center.



At the gateway, or reverse proxy gateway, the metadata package gets instructions to deliver the package to your IBM Storage Insights Pro or IBM Storage Insights service. Only data collectors that are associated with your service can collect and deliver metadata about your storage environment.

When the metadata package is delivered, the metadata is decrypted, analyzed, and stored.

From your data center to the internet

HTTPS connections are used to compress and encrypt the metadata that is collected about your storage systems and sent to the IBM Cloud data center.

As part of the onboarding process, you're provided with a host name and port number for your IBM Storage Insights Pro or IBM Storage Insights service. To secure the outbound communication between the data collector and IBM Storage Insights Pro or IBM Storage Insights, a Secure Sockets Layer (SSL) certificate is used. HTTPS connections use certificates issued by Cloudflare, Inc. (issuer common name "Cloudflare Inc ECC CA-3") and use TLS 1.2 and TLS 1.3 with 256-byte keys.

To send the metadata, complete the following tasks for your firewall:

- Update your firewall rules to allow outbound communication on the default HTTPS port 443 using the Transmission Control Protocol (TCP). The User Datagram Protocol (UDP) is not supported.
- Update your firewall rules to allow outbound communication to the following network endpoint: <https://insights.ibm.com>. If you use a proxy server with a separate firewall, ensure that you also update its rules.
- If a firewall exists between the data collector and the storage systems that it monitors, configure the firewall to allow SNMP traffic.

At the IBM Cloud data center

IBM Storage Insights Pro and IBM Storage Insights are hosted in IBM Cloud data centers, which comply with high physical, technical, and organizational security standards.

Key security

Each instance of IBM Storage Insights uses a local keystore that is dedicated to that instance and is password protected. The password for the keystore is generated randomly when the instance is created. The certificate in the keystore is unique to each instance and the keystore password is encrypted. (The encryption doesn't include hardware encryption.) The master password is kept encrypted in the service payload configuration in a secure location in IBM Cloud®.

There is only one external customer key, which is the public key that is certified by DigiCert. As part of the TLS Handshake and certificate exchange, the client (Web Browser) uses the signed certificate to verify that it is communicating with the IBM Storage Insights Pro or IBM Storage Insights gateway in IBM Cloud and that communications are not tampered with. For internal traffic, each customer's instance of IBM Storage Insights Pro or IBM Storage Insights has a unique key, which is protected with a unique, encrypted password, and which is self-signed by IBM to validate that the communication is between the customer and the customer's instance.

Key rotation: A new master key is created and added to the keystore when the instance is created and when the instance is upgraded. Instances are upgraded at least once every three months, which results in an implicit key rotation of not less than 90 days. The public key that is certified by DigiCert is updated every 2 years.

This results in end-to-end privacy and encryption for each instance of IBM Storage Insights Pro and IBM Storage Insights.

Physical protection

The data centers are rigorously controlled and onsite security is provided round the clock. Access to server-rooms is limited to certified employees and security controls are vetted by third-party auditors.

See <https://www.ibm.com/cloud-computing/bluemix/data-centers> and <https://www.ibm.com/cloud/security>.

Technical security

Each instance of IBM Storage Insights Pro or IBM Storage Insights is isolated from other instances of IBM Storage Insights Pro or IBM Storage Insights at the compute and at the storage layer.

Note: Dedicated SAN disks are used for each instance.

Each instance also uses its own database, which stores only data for that instance. The database is encrypted separately using native IBM DB2® encryption with a strong cryptographic algorithm, AES 256-bit encryption. An SSL certificate is used. For the encrypted DB2 database and its encrypted backups that are stored in IBM Cloud, the certificate is self-signed by IBM.

See https://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.sec.doc/doc/c0061758.html.

On a day-to-day basis, the following security software and services are used:

- Mantes Endpoint Protection to protect against malware

- IBM Bigfix® to comply with security and regulatory requirements
- IBM Security QRadar® SIEM to store and monitor system and application logs

Organizational security

Access to the infrastructure and instances for IBM Storage Insights Pro and IBM Storage Insights, is controlled:

- By restricting access to the members of the DevOps team and cloud service infrastructure teams who qualify as privileged users.
- By conducting regular system health and vulnerability scans at the source code level and on the running instances.
- By conducting regular penetration tests. External companies conduct the penetration tests.

GDPR: IBM Storage Insights meets the requirements of the EU General Data Protection Regulation (GDPR). Additional information related to IBM's privacy policy can be found at <https://www.ibm.com/privacy/us/en/>.

Chapter 4. What types of metadata are collected

Asset, capacity, configuration, and performance metadata are collected and stored for the storage devices that are monitored. Diagnostic data is also collected into log packages and added to support tickets.

Metadata about the configuration and operations of storage systems, switches, fabrics, and hosts is collected. The actual data that is stored on devices is not accessed or collected.

The metadata that is collected is used for the following tasks:

- To provide and improve services
- To analyze and get insights into storage usage and performance
- To generate charts and present data in IBM Storage Insights
- To upload logs automatically when support tickets are created or updated
- To enable IBM Support to investigate and close the issues that you might encounter

The types of the metadata that are collected and stored are as follows:

Asset metadata

Name, model, firmware, and type of a device.

Configuration metadata

Name and number of the resources that are associated with the devices such as the number of disks, pools, volumes, and ports.

Capacity metadata

Storage usage values such as provisioned capacity, unmapped capacity, pool capacity, used capacity, and raw capacity.

Performance metadata

Performance metrics, such as read and write data rates, I/O rates, and response times for storage systems and switches.

Relationships metadata

In IBM Storage Insights Pro, metadata is collected to represent general groups, applications, and departments that are created to show storage usage and performance metadata.

Data collector metadata

To monitor the performance of the data collector and resolve data collection issues, the name of the server that hosts the data collector, the status of the data collector, and when the data collector was last run is collected.

To get the metadata, the information that is used to connect to devices is also collected and stored. The information is stored in the database that was created for your IBM Storage Insights service. Passwords are encrypted before they are stored in the database.

IBM Support ticket and diagnostic log packages

When the IBM Storage Insights data collector collects a diagnostic log package from a device, it transfers the package to IBM Support by using HTTPS.

When users create tickets in IBM Storage Insights, they provide a name, an email address, and a phone number so that IBM Support can contact them. IBM Storage Insights collects and uploads the diagnostic data for IBM block storage systems to IBM Enhanced Customer Data Repository (ECuRep).

What is ECuRep: ECuRep is an IBM strategic worldwide Post Sales Technical Support solution for client diagnostic data transmission, storing, and analysis.

If an IBM block storage system is configured to use Blue Diamond Enhanced Secure Support, IBM Storage Insights collects and uploads the diagnostic data that is collected for the storage system to the Blue Diamond environment.

What is Blue Diamond: Blue Diamond Enhanced Secure Support is an enhancement to standard IBM remote software and hardware support. It adds extra layers of security and allows you to use a secure, dedicated portal to upload diagnostic data to IBM® Support.

Chapter 5. How long is the metadata kept

Information is provided about the retention periods for the metadata that is collected to provide storage services and to improve storage services.

When you add storage systems, asset, configuration, capacity, and performance metadata is collected about the storage systems. The metadata is collected and retained to provide and improve the analytical and monitoring services that IBM Storage Insights offers. For example, the metadata is analyzed to present key capacity and performance metrics and for detecting and investigating capacity and performance trends. The metadata is also used to show capacity savings, predict capacity shortfalls, and to provide reclamation and tiering recommendations.

As the metadata is collected, the aggregation level of the metadata changes. For asset, configuration, and capacity metadata, over a 24-month period, the aggregation levels of the metadata change from daily, to weekly, to monthly based on the age of the metadata. For performance metadata, over a 52-week period, the aggregation levels change from sample, to hourly, to daily based on the age of the performance metadata. In effect, a more granular view of new metadata is provided and a less granular view of aged metadata is provided.

The following table lists the aggregation levels for asset, configuration, and capacity metadata based on the age of the data that is collected:

| Aggregation level | Metadata age |
|-------------------|--------------|
| Daily | 12 weeks |
| Weekly | 24 weeks |
| Monthly | 24 months |

The following table lists the aggregation levels for performance metadata based on the age of the data that is collected:

| Aggregation level | Metadata age |
|-------------------|--------------|
| Sample | 2 weeks |
| Hourly | 4 weeks |
| Daily | 52 weeks |

Based on the collection date, metadata is retained for up to two years.

Note: If you subscribe to IBM Storage Insights Pro and cancel your subscription, you'll still be able to use IBM Storage Insights. The metadata from IBM Storage Insights Pro is retained.

How long are diagnostic data packages kept

Typically, diagnostic data is automatically deleted from IBM Enhanced Customer Data Repository (ECuRep) 30 days after the ticket is closed. For information about the retention of data in ECuRep, see the [IBM terms of use for Exchanging diagnostic data with IBM](#).

Blue Diamond Enhanced Secure Support uses a secure, dedicated portal for diagnostic data packages. For more information about diagnostic data and Blue Diamond, contact the Blue Diamond team at the [Blue Diamond registration page](#).

Related tasks

Requesting the deletion of personal information

To delete the minimal personal information that was stored to provide you with monitoring and support services for your storage systems, you can submit a request to IBM Support.

Chapter 6. Who can access the metadata

Information is provided about access to the metadata that is collected and stored.

Access to the metadata that is collected and stored for your instance of IBM Storage Insights Pro or IBM Storage Insights is restricted:

- To the DevOps and cloud service infrastructure teams who are responsible for the maintenance and day-to-day operation of your instance
- To IBM Support for investigating and closing support tickets and for downloading support logs to investigate issues

Note: IBM Support has read-only access to the metadata that is collected about all your monitored block storage systems and their internal resources.

To access the metadata in the IBM Cloud network and ensure that the connection is secure, DevOps and cloud service infrastructure teams use a secure virtual private network (VPN) connection. Access to instances is only permitted from privileged user workstations, which must meet the strict security controls of IBM Security policies for production servers.

The access to metadata for DevOps and cloud service infrastructure teams is restricted:

- To the infrastructure for the cloud service
- To the operating system
- To add-on services such as agents
- To middleware components

Metadata access controls and authorization

Access controls and authorization checks are enforced for SaaS infrastructure components and services.

An approval process is used to authorize access to the following infrastructural elements and services:

- The network
- The operating system
- The middleware components
- The application
- Administrative services

When managing the changes to a production environment, the following change management processes are adhered to:

- Changes to the production environment must be recorded and must be approved by the change advisory board
- All support activities must be tracked in the IBM Support Portal for cloud services
- All operational and maintenance activities must be tracked by the internal ticketing system

Metadata access for resolving issues

To investigate and resolve issues, access is required to metadata and instances of IBM Storage Insights Pro or IBM Storage Insights.

To find the causes of issues, investigations are undertaken that might require access to the metadata that is collected and stored, or access to infrastructural elements, or both. For example, the DevOps team or IBM Support, might need to monitor instances of the application to determine the cause of interruptions in service, or to investigate interruptions in the collection of metadata. To resolve such issues, it might be necessary:

- To analyze the configuration of the instance
- To analyze log files
- To analyze the metadata that was collected

To thoroughly investigate some issues, it might also be necessary to package the metadata and transfer it to a secure IBM system so that the development team can complete the investigation.

IBM Support access for troubleshooting your tickets

To investigate hardware and software tickets, IBM Support has read-only access to the asset, configuration, capacity, and performance metadata that is collected for IBM storage systems and their internal storage resources.

The metadata might not provide enough information to close the ticket, so IBM Support might need to collect a log package from your storage systems. In this case, IBM Support can attach the log package to an open ticket and submit the log package to IBM Enhanced Customer Data Repository (ECuRep). Depending on the data governance requirements of a client, the diagnostic data package might be uploaded to the Blue Diamond Enhanced Secure Support environment instead of ECuRep.

Permit IBM Support to collect log packages: To save time when IBM® Support troubleshoots your ticket, you can permit IBM® Support to collect and upload log packages remotely without contacting you. To set this permission, click **Configuration > Settings**, and then click **Edit** in the **IBM Support Log Permissions** section. You can set this permission for each storage system.

This is the procedure for uploading the log packages to tickets:

1. The data collector submits a request to the storage system to create a log package.
2. The data collector uses Hypertext Transfer Protocol Secure (HTTPS), which encrypts the metadata, and sends the log package through a secure channel to IBM Storage Insights.
3. IBM Storage Insights sends the log package to ECuREP or to Blue Diamond Enhanced Secure Support environment.

Metadata access for quality improvements

Anonymized metadata is used to improve the quality of service and to enhance the product offering.

A subset of the metadata from all of the instances is aggregated and condensed for further analysis. The data that is used is anonymized:

- It does not include instance-specific metadata
- It does not include customer-specific metadata such as IP addresses

For example, the aggregated metadata contains such information as the number of different types of storage systems or the number of different firmware levels for the storage systems that are monitored. The aggregated metadata might contain GUI and usage metrics, but it doesn't contain the names, the serial numbers, or the IP addresses of the storage systems.

Data backup and restore

To restore instances, regular backups of the data are made automatically.

Backups are made daily, which means that recovery point objective (RPO) is one day, and the recovery time objective (RTO) is between 1.5 and 2 days.

Backups are stored both locally, in the same data center, and remotely. The latest backup of the instance is stored in a remote data center, whereas the five previous backups are stored in the local data center.

Requesting the deletion of personal information

To delete the minimal personal information that was stored to provide you with monitoring and support services for your storage systems, you can submit a request to IBM Support.

If you cancel your subscription for IBM Storage Insights Pro or decide that you no longer want to monitor your storage environment with IBM Storage Insights, you can request that the minimal personal information is deleted.

1. Go to [IBM Support](#).
2. Sign in.
3. Click **Go to my cases**.
4. Create a new case and request the deletion of your personal information.

Chapter 7. Asset, capacity, and configuration metadata

The data collector collects and stores asset, capacity, and configuration metadata for block, file, and object storage systems and their resources. A list of the supported storage systems is provided.

Asset, capacity, and configuration, metadata is collected and stored for the following IBM block storage systems and non-IBM block storage systems when they are added for monitoring.

Table 3. Asset, capacity, and configuration metadata by storage system

| Storage System | IBM Storage Insights Pro | IBM Storage Insights |
|--------------------------------------|---------------------------------|-----------------------------|
| DS8000® | Yes | Yes |
| Dell EMC Unity | Yes | No |
| Dell EMC VMAX | Yes | No |
| Dell EMC VNX, VNXe | Yes | No |
| FlashSystem 5000 | Yes | Yes |
| FlashSystem 5100 | Yes | Yes |
| FlashSystem 7200 | Yes | Yes |
| FlashSystem 9100 | Yes | Yes |
| FlashSystem 9200 | Yes | Yes |
| FlashSystem 900 | Yes | Yes |
| FlashSystem A9000 | Yes | Yes |
| FlashSystem A9000R | Yes | Yes |
| FlashSystem V9000 | Yes | Yes |
| Hitachi VSP | Yes | No |
| IBM Spectrum Accelerate | Yes | Yes |
| NetApp ONTAP 9 | Yes | No |
| Pure FlashArray//M and FlashArray//X | Yes | No |
| SAN Volume Controller | Yes | Yes |
| Storwize® V3500 | Yes | Yes |
| Storwize V3700 | Yes | Yes |
| Storwize V5000 | Yes | Yes |
| Storwize V7000 | Yes | Yes |
| Storwize V7000 Unified | Yes | Yes |
| XIV® | Yes | Yes |

In IBM Storage Insights Pro, asset, capacity, and configuration metadata is collected for the following file, object, and software-defined storage systems when they are added for monitoring:

- Dell EMC Unity
- Dell EMC VNX

- Dell EMC VNXe
- IBM Cloud Object Storage
- IBM Spectrum Scale
- NetApp ONTAP 9
- Storwize V7000 Unified

Block storage system metadata

View the list of the asset, configuration, and capacity metadata for block storage systems.

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored:

- Available Capacity (GiB) (Previously known as Available Pool Space)
- Available Written Capacity (GiB) (Previously known as Effective Available Capacity)
- Compressed
- Compression Savings (%)
- Custom Tag 1, 2, 3
- Deduplication Savings (%)
- Data Collection
- Data Reduction (GiB)
- Data Reduction Ratio
- Disks
- FC Ports
- Firmware
- FlashCopy®
- IP Address
- IP Ports
- Location
- Managed Disks
- Mapped Capacity (GiB) (Previously known as Assigned Volume Space)
- Model
- Name
- Overhead Capacity (GiB)
- Pools
- Ports
- Raw Capacity (GB)
- Read Cache (GiB)
- Remote Relationships
- Reserved Volume Capacity (GiB) (Previously known as Unused Space)
- Safeguarded Capacity (GiB)
- Serial Number
- Shortfall (%)
- Time Zone
- Total Provisioned (%) (Previously known as Provisioned Capacity (%))
- Total Provisioned (GiB) (Previously known as Provisioned Capacity (GiB))
- Total Savings (%) (Previously known as Total Capacity Savings (%))
- Total Savings Ratio
- Turbo Performance
- Type
- Unmapped Capacity (GiB) (Previously known as Unassigned Volume Space)
- Usable Capacity (GiB) (Previously known as Capacity(GiB))

Used Capacity (%) (Previously known as Physical Allocation)
Used Capacity (GiB) (Previously known as Used Pool Space (GiB))
Used Written Capacity (%) (Previously known as Effective Used Capacity)
Used Written Capacity (GiB) (Previously known as Effective Used Capacity)
VDisk Mirrors
Vendor
Volumes
Write Cache (GiB)
Written Capacity Limit (GiB) (Previously known as Effective Capacity)

Block volumes metadata

View the list of the asset, capacity, and configuration metadata for volumes.

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for block storage volumes.

Auto Expand
Available Capacity (%) (Previously known as Unallocated Space)
Capacity (GiB)
Compression Savings (%)
Copies
Copy ID
Copy Relationship
Easy Tier®
Enterprise HDD Capacity (GiB)
Enterprise HDD Capacity (%)
Fast Write State
Formatted
Grain Size (KiB)
Hosts
I/O Group
ID
Last Data Collection
Mirror Role
Name
Nearline HDD Capacity (GiB)
Nearline HDD Capacity (%)
Node
Pool
Provisioned Capacity (%) (Previously known as Virtual Allocation)
RAID Level
Reserved Volume Capacity (GiB)
Safeguarded
Safeguarded Capacity (GiB)
Safeguarded Location
SCM Capacity (GiB)
SCM Capacity (%)
Shortfall (%)
SSD Capacity (GiB)
Storage System
Storage Virtualizer
Tier 0 Flash Capacity (GiB)

- Tier 0 Flash Capacity (%)
- Tier 1 Flash Capacity (GiB)
- Tier 1 Flash Capacity (%)
- Tier 2 Flash Capacity (GiB)
- Tier 2 Flash Capacity (%)
- Thin Provisioned
- Tier Distribution (%)
- Unique ID
- Used Capacity (%)
- Used Capacity (%) (Previously known as Physical Allocation)
- Used Capacity (GiB) (Previously known as Allocated Space)
- Virtualizer Disk Type
- Virtualizer Disk
- Warning Level (%)
- Written Capacity (GiB) (Previously known as Written Space)

Block pools metadata

View the list of the asset, capacity, and configuration metadata for block storage pools.

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for storage pools:

- Activity
- Available Capacity (GiB) (Previously known as Available Pool Space)
- Available Repository Capacity (GiB)
- Available Soft Capacity (GiB)
- Available Written Capacity (GiB) (Previously known as Effective Available Capacity)
- Back-end Storage Disk Type
- Back-end Storage Disks
- Back-end Storage RAID Level
- Back-end Storage System Type
- Compression Savings (%)
- Custom Tag 1, 2, 3
- Deduplication Savings (%)
- Easy Tier
- Encryption
- Encryption Group
- Enterprise HDD Available Capacity (GiB)
- Enterprise HDD Capacity (GiB)
- Extent Size (MiB)
- Format
- Last Data Collection
- LSS or LCU
- Managed Disks
- Mapped Capacity (GiB) (Previously known as Assigned Volume Space)
- Name
- Nearline HDD Available Capacity (GiB)
- Nearline HDD Capacity (GiB)
- Overhead Capacity (GiB)
- Overprovisioned Capacity (GiB) (Previously known as Unallocatable Volume Space)
- Owner Name
- RAID Level

- Rank Group
- Repository Capacity (GiB)
- Reserved Volume Capacity (GiB) (Previously known as Unused Space)
- Safeguarded Capacity (GiB)
- SCM Capacity (GiB)
- SCM Available Capacity (GiB)
- Shortfall (%)
- Soft Capacity (GiB)
- Solid State
- SSD Available Space (GiB)
- SSD Capacity (GiB)
- Status
- Storage System
- Tier
- Tier 0 Flash Capacity (GiB)
- Tier 0 Flash Available Capacity (GiB)
- Tier 1 Flash Capacity (GiB)
- Tier 1 Flash Available Capacity (GiB)
- Tier 2 Flash Capacity (GiB)
- Tier 2 Flash Available Capacity (GiB)
- Tier Distribution (%)
- Total Provisioned (%) (Previously known as Provisioned Capacity (%))
- Total Provisioned (GiB) (Previously known as Provisioned Capacity (GiB))
- Total Savings (%) (Previously known as Total Capacity Savings (%))
- Unmapped Capacity (GiB) (Previously known as Unassigned Volume Space)
- Unused Volume Capacity (%) (Effective Previously known as Unallocated Volume Space)
- Usable Capacity (GiB) (Previously known as Capacity(GiB))
- Used Capacity (%) (Previously known as Physical Allocation)
- Used Capacity (GiB) (Previously known as Allocated Space)
- Used Written Capacity (%) (Previously known as Effective Used Capacity)
- Used Written Capacity (GiB) (Previously known as Effective Used Capacity)
- Volumes
- Written Capacity Limit (GiB) (Previously known as Effective Capacity)
- Zero Capacity

I/O groups, nodes, and ports metadata

View the lists of the asset, capacity, and configuration metadata for I/O groups, nodes, Fibre Channel (FC) ports, and IP ports.

I/O groups

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for I/O groups:

- Compression
- Enclosure
- Name
- Total FlashCopy Memory (MiB)
- Total Mirroring Memory (MiB)
- Total Remote Copy Memory (MiB)
- Used Mirroring Memory (MiB)
- Used Remote Copy Memory (MiB)

Volumes

Nodes

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for nodes:

- Compression
- Configuration node
- Enclosures
- FC ports
- I/O group
- IP address
- IP ports
- Model
- Name
- Panel name
- Serial number
- Spare nodes
- WWN

FC Ports

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for FC ports:

- Acknowledged
- Connected NPIV Ports
- Connected Port
- Connected Resource
- Connected WWPN
- FC Port ID
- Frame
- Host Adapter
- I/O Enclosure
- Location
- Name
- Protocol
- Speed (Gbps)
- Status
- WWPN

IP Ports

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for IP ports:

- Acknowledged
- Duplex
- Failover
- Gateway
- Host Attach
- IP Address
- IQN
- Management

- MTU
- Name
- Node
- Remote Copy Relationship
- Speed (Gbps)
- Status
- Storage Attach
- Storage System
- Subnet

Disks and managed disks metadata

View the lists of the asset, capacity, and configuration metadata for disks and managed disks.

Disks

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for disks:

- Capacity (GB)
- Capacity (GiB)
- Class
- Firmware
- Hardware
- Name
- Serial number
- Slot
- Spare
- Speed (RPM)
- Status
- Vendor

Managed Disks

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for managed disks:

- Active Quorum
- Available Capacity (GiB)
- Back-end Storage System
- Capacity (GiB)
- Class
- Easy Tier
- Easy Tier Load
- Mode
- Name
- Pool
- RAID Level
- Storage System
- Volumes

Block device adapters metadata

View the list of the asset, configuration, and capacity metadata for device adapters.

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for device adapters.

- Name
- I/O Enclosure
- Frame
- Device Adapter Pair
- Location
- Serial number

Block host adapters metadata

View the list of the asset, configuration, and capacity metadata for host adapters.

Depending on the type of storage system, and the features that the storage system supports, some or all of the following metadata is collected and stored for host adapters.

- Name
- I/O Enclosure
- Frame
- Type
- Associated ports
- Speed
- Speed
- Location
- Serial number

Enclosures metadata

View the list of the configuration metadata for block storage system enclosures.

Metadata about enclosures is collected for the following block storage systems:

- FlashSystem 9100
- FlashSystem V9000
- FlashSystem 900
- SAN Volume Controller
- IBM Storwize family

Depending on the type of storage system, some or all of the following metadata is collected and stored for enclosures:

- Canisters
- Canister status
- Disks
- Disk slots
- I/O groups
- Machine Type Model
- Name
- Nodes
- Power[®] supplies
- Power supply status
- Site

Status

File storage system metadata

View the list of the asset, configuration, and capacity metadata for file storage systems.

Depending on the type of file storage system, and the features that the file storage system supports, some or all of the following metadata is collected and stored:

Restriction: Applies only to IBM Storage Insights Pro.

- Available File System Capacity (GiB)
- Cluster
- Custom Tag 1, 2, 3
- Data Collection
- Disks
- External Pool Used Capacity (GiB)
- File System Capacity (%)
- File System Capacity (GiB) (Previously known as Total File System Capacity)
- IP Address
- Location
- Model
- Name
- Raw Capacity (GB)
- Serial Number
- Type
- Used File System Capacity (GiB)
- Vendor
- Version

File system metadata

View the list of the asset, configuration, and capacity metadata for file systems.

Depending on the type of file storage system, and the features that the file storage system supports, some or all of the following metadata is collected and stored:

Restriction: Applies only to IBM Storage Insights Pro.

- Available Inodes
- Available Capacity (%)
- Capacity (%)
- Capacity (GiB) (Previously known as Total Capacity)
- Cluster
- Custom Tag 1, 2, 3
- File System Type
- Filesets
- Last Probe Time
- Maximum Inodes
- Name
- NSDs
- Path
- Physical Capacity (GiB)
- Pools
- Storage System
- Used Inodes

Used Inodes (%)
Used Capacity (GiB)

Fileset metadata

View the list of the asset, configuration, and capacity metadata for filesets.

Depending on the type of file storage system, and the features that the file storage system supports, some or all of the following metadata is collected and stored.

Restriction: Applies only to IBM Storage Insights Pro.

Cache Role
Comment
File System
Home System Name
Name
Path
State
Used Capacity (GiB)

File shares metadata

View the list of the asset, configuration, and capacity metadata for file shares.

Depending on the type of file storage system, and the features that the file storage system supports, some or all of the following metadata is collected and stored.

Restriction: Applies only to IBM Storage Insights Pro.

Cluster
Discovered Time
Name
Path
Protocols
Shared Servers
State
Storage System
Ticket

File system pools metadata

View the list of the asset, configuration, and capacity metadata for file system pools.

Depending on the type of file storage system, and the features that the file storage system supports, some or all of the following metadata is collected and stored:

Restriction: Applies only to IBM Storage Insights Pro.

Available Capacity (GiB)
Capacity (%)
Capacity (GiB) (Previously known as Total Capacity)
Cluster
External Used Capacity (GiB)
File System
Inactive Used Capacity (%)
Inactive Used Capacity (GiB)
Name
Storage System

Used Capacity (GiB)

Network shared disks metadata

View the list of the asset, configuration, and capacity metadata for network shared disks.

Depending on the type of file storage system, and the features that the file storage system supports, some or all of the following metadata is collected and stored about network shared disks:

Restriction: Applies only to IBM Storage Insights Pro.

- Available Capacity (GiB)
- Capacity (%) (Previously known as Disk Capacity)
- Capacity (GiB) (Previously known as Total Disk Capacity)
- Cluster
- Correlated Storage Volume
- Custom Tag 1, 2, 3
- Failure Group
- File System
- ID
- Name
- Pool
- Probe Time
- Storage System
- Type
- Used Capacity (GiB)

File nodes metadata

View the list of the asset, configuration, and capacity metadata for file nodes.

Depending on the type of file storage system, and the features that the file storage system supports, some or all of the following metadata is collected and stored about file nodes.

Restriction: Applies only to IBM Storage Insights Pro.

- Cache Gateway Node
- Cluster
- IP Address
- Name
- Role
- Serial number
- Storage System
- Version

Object storage systems metadata

View the lists of the asset, configuration, and capacity metadata for object storage systems and their internal resources.

Some or all of the following metadata is collected and stored for object storage systems and their internal storage resources such as:

- Access pools
- Accesser[®] nodes
- Storage pools
- Slicestor[®] nodes

- Sites
- Vaults

Restriction: Applies only to IBM Storage Insights Pro.

Storage systems metadata

Access pools
Available Capacity (GiB)
Capacity (%), Capacity (GiB) (Previously known as Total Capacity)
Containers
Custom Tag 1, 2, 3
Data Collection
IP Address
Location
Name
Objects
Type
Used Capacity (GiB)
Vaults
Vendor
Version

Access pools metadata

Mirrors
Name
Protocol
Vaults
Site

Accesses nodes metadata

Access Pool
IP Address
Model
Name
Serial Number
Site
Software Version

Storage pools metadata

Capacity (%), Capacity (GiB) (Previously known as Total Capacity)
Name
Sets
Site
Storage Pools
Used Capacity (GiB)
Vaults

Slicestor nodes metadata

Available Capacity (GiB)

Capacity (%), Capacity (GiB) (Previously known as Total Capacity)
Drives
IP Address
Model
Name
Serial Number
Set ID
Site
Software Version
Storage Pool
Used Capacity (GiB)

Sites metadata

Access Nodes
Accessibility
Available Capacity (GiB)
Capacity (%), Capacity (GiB) (Previously known as Total Capacity)
Name
Slicestor Nodes
Used Capacity (GiB)
Vaults

Vaults metadata

Access Pools
Accessibility, Capacity Quota (%)
Creation Date
Hard Quota (GiB)
IDA
Mirror
Name
Sites
Soft Quota (GiB)
Storage Pool
Storage System
Used Capacity (GiB)

Switches metadata

View the list of the asset, configuration, and status metadata for switches.

Depending on the type of switch, some or all of the following metadata is collected and stored for switches.

Acknowledged
Chassis
Condition
Connected Ports
Custom Tag 1, 2, and 3
Data Collection
Domain ID
Fabric
Firmware

- IP Address
- Last Successful Monitor
- Last Successful Probe
- Links
- Location
- Mode
- Model
- Name
- Performance Monitor Status
- Probe Status
- Ports
- Principal Switch of Fabric
- Serial Number
- System UUID
- Vendor
- WWN

Fabrics metadata

View the list of the asset, configuration, and status metadata for fabrics.

Depending on the type of fabric, some or all of the following metadata is collected and stored for fabrics.

- Acknowledged
- Active Zone Set
- Condition
- Connected Ports
- Custom tag 1, 2, and 3
- Fabric Type
- Links
- Location
- Name
- NPV Switches
- Principal Switch of Fabric
- Switches
- Switch Ports
- System UUID
- WWN

Groups metadata

View the lists of the asset, configuration, and capacity metadata for general groups, applications, and departments.

Metadata for the resources in general groups, applications, and departments is also collected and stored. That metadata includes filter definitions for applications.

Restriction: Applies only to IBM Storage Insights Pro.

General groups metadata

- Name
- Subgroups

Applications metadata

- Block Capacity (GiB)
- Custom Tag 1, 2, 3
- Departments
- File Capacity (GiB)
- Name
- Object Capacity (GiB)
- Hosts
- Subtype
- Type
- Vaults
- Volumes

Accesses nodes metadata

- Access Pool
- IP Address
- Model
- Name
- Serial Number
- Site
- Software Version

Storage pools metadata

- Capacity (%), Capacity (GiB) (Previously known as Total Capacity)
- Name
- Sets
- Site
- Storage Pools
- Used Capacity (GiB)
- Vaults

Slicestor nodes metadata

- Available Capacity (GiB)
- Capacity (%), Capacity (GiB) (Previously known as Total Capacity)
- Drives
- IP Address
- Model
- Name
- Serial Number
- Set ID
- Site
- Software Version
- Storage Pool
- Used Capacity (GiB)

Sites metadata

- Access Nodes
- Accessibility
- Available Capacity (GiB)

Capacity (%), Capacity (GiB) (Previously known as Total Capacity)
Name
Slicestor Nodes
Used Capacity (GiB)
Vaults

Vaults metadata

Access Pools
Accessibility, Capacity Quota (%)
Creation Date
Hard Quota (GiB)
IDA
Mirror
Name
Sites
Soft Quota (GiB)
Storage Pool
Storage System
Used Capacity (GiB)

Hosts metadata

View the list of the asset, capacity, and configuration metadata for hosts.

Metadata is collected about hosts, their drives, HBAs, datastores, VMDKs and the related storage resources, such as the volumes that they use.

The following metadata is collected and stored for hosts:

Cluster
Domain Name
Drive Capacity (GiB)
Flash Copy Provisioned Capacity (GiB)
Flash Copy Used Capacity (GiB)
IP Address
Location
Model
Name
OS Type
Primary Provisioned Capacity (GiB)
Primary Used Capacity (GiB)
Remote Mirror Provisioned Capacity (GiB)
Remote Mirror Used Capacity (GiB)
SAN Capacity (GiB)
Serial Number
Used SAN Capacity (GiB)
vCenter
VDisk Mirror Provisioned Capacity (GiB)
VDisk Mirror Used Capacity (GiB)
Vendor

The following metadata is collected and stored for HBAs:

Associated Drives

HBA WWN
Name
Type
IQN

The following metadata is collected and stored for drives:

Available Drive Capacity (GiB)
Capacity (GiB)
Capacity (GB)
Name
Paths
Serial Number
Vendor

The following metadata is collected and stored for datastores:

Available Datastore Capacity (GiB)
Available File System Capacity (GiB)
Cluster
Datastore Capacity (GiB)
Datastore Capacity (GB)
File System
File System Capacity (GiB)
Name
Paths
Serial Number
Swap Space
Type
Vendor
Used Capacity (%)
Used Datastore Capacity (GiB)
Used File System Capacity (GiB)

The following metadata is collected and stored for VMDKs:

Datastore
Drives
Paths
Size (GiB)
Used Capacity (GiB)
Virtual Machine
Virtual Machine Drive

Virtual machines metadata

View the list of the asset, capacity, and configuration metadata for virtual machines.

Metadata is collected about virtual machines, their drives, HBAs and the related storage resources.

The following metadata is collected and stored for virtual machines:

Cluster
Datastores
Host
IP Address
Location
Name

OS Type

Chapter 8. Performance metadata

The data collector collects and stores performance metadata for IBM block storage systems and non-IBM block storage systems and it collects and stores file system and node performance metadata for IBM Spectrum Scale storage systems.

Performance metadata is collected and stored for the following IBM block storage systems and non-IBM block storage systems when they are added for monitoring:

Table 4. Block performance metadata by storage system

| Storage System | IBM Storage Insights Pro | IBM Storage Insights |
|--------------------------------------|---------------------------------|-----------------------------|
| DS8000 | Yes | Yes |
| Dell EMC Unity | Yes | No |
| Dell EMC VMAX | Yes | No |
| Dell EMC VNX, VNXe | Yes | No |
| FlashSystem 5000 | Yes | Yes |
| FlashSystem 5100 | Yes | Yes |
| FlashSystem 7200 | Yes | Yes |
| FlashSystem 9100 | Yes | Yes |
| FlashSystem 9200 | Yes | Yes |
| FlashSystem 900 | Yes | Yes |
| FlashSystem A9000 | Yes | Yes |
| FlashSystem A9000R | Yes | Yes |
| FlashSystem V9000 | Yes | Yes |
| Hitachi VSP | Yes | No |
| IBM Spectrum Accelerate | Yes | Yes |
| NetApp ONTAP 9 | Yes | No |
| Pure FlashArray//M and FlashArray//X | Yes | No |
| SAN Volume Controller | Yes | Yes |
| Storwize V3500 | Yes | Yes |
| Storwize V3700 | Yes | Yes |
| Storwize V5000 | Yes | Yes |
| Storwize V7000 | Yes | Yes |
| Storwize V7000 Unified | Yes | Yes |
| XIV | Yes | Yes |

Node and file system performance metadata is also collected and stored for IBM Spectrum Scale storage systems.

Restriction: Applies to IBM Storage Insights Pro.

Performance metadata for storage systems that run IBM Spectrum® Virtualize

View the lists of the performance metadata for IBM SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, IBM Storwize storage systems, and for IBM FlashSystem® devices that run IBM Spectrum Virtualize.

Overview

The performance metadata is divided into the following categories:

- [“Volume metadata” on page 36](#)
- [“Disk metadata” on page 37](#)
- [“Pool metadata” on page 37](#)
- [“Port metadata” on page 37](#)
- [“Node metadata” on page 38](#)

Volume metadata

The following key metadata is collected for volume performance:

- Data Rate (Read, Total, Unmap, Write)
- Overall Host Attributed Response Time Percentage
- Overall I/O Rate (Read, Total, Unmap, Write)
- Pool Activity Score
- Response Time (Overall, Read, Unmap, Write)
- Volume Utilization
- Write Cache Delay Percentage

The following I/O rate metadata is collected for volume performance:

- Transfer Rate (Cache-to-Disk, Disk-to-Cache)
- Unaligned Unmap I/O Rate
- Write-Cache Delay I/O Rate

The following cache hit percentage metadata is collected for volume performance:

- Overall I/O Cache Hits (Read, Total, Write)

The following response time metadata is collected for volume performance:

- Peak Response Time (Read, Unmap, Write)

The following remote mirror metadata is collected for volume performance:

- Global Mirror (Overlapping Write I/O Rate, Overlapping Write Percentage, Secondary Write Lag, Write I/O Rate)

The following volume cache metadata is collected for volume performance:

- Cache Hits (Dirty Writes, Read, Total, Write)
- I/O Rate (Destage, Read, Total, Write)
- Response Time (Destage, Stage)
- Transfer Rates (Cache-to-Disk, Disk-to-Cache)
- Write Delay Percentage (Flush-through, Total Delay, Write-through)
- Write Delay Rate (Flush-through, Total Delay, Write-through)
- Data Rate (Fast-Write Writes)

The following volume copy cache metadata is collected for volume performance:

- Cache Hits (Dirty Writes, Read, Read-ahead, Total, Write)
- I/O Rate (Destage, Prestage, Stage, Total)
- Response Time (Destage, Prestage, Stage)
- Transfer Rates (Cache-to-Disk, Disk-to-Cache)
- Write Delay Percentage (Flush-through, Total Delay, Write-through)
- Write Delay Rate (Flush-through, Total Delay, Write-through)
- Data Rate (Fast-Write Writes)

The following compression metadata is collected for volume performance:

- Compressed Volumes (Data Rate, I/O Rate, Response Time)
- Uncompressed Volumes (Data Rate, I/O Rate, Response Time)

The following miscellaneous metadata is collected for volume performance:

- Cache to Host Transfer Response Time
- Non-Preferred Node Usage Percentage
- Transfer Size (Overall, Read, Write)
- Unaligned Write I/O Rate

The following legacy cache metadata is collected for volume performance:

- Dirty Write Percentage of Cache Hits
- Write-Cache I/O Rate (Flush-through, Overflow, Write-through)
- Write-Cache Percentage (Flush-through, Overflow, Write-through)

Disk metadata

The following key metadata is collected for disk performance:

- Data Rate (Read, Total, Write)
- I/O Rate (Read, Total, Write)
- Response Time (Overall, Read, Write)

The following response time metadata is collected for disk performance:

- Peak Back-end Response Time (Read, Write)
- Peak Back-end Queue Time (Read, Write)
- Queue Time (Overall, Read, Write)

The following miscellaneous metadata for disk performance:

- Cache Destage (In-Flight I/O, Target I/O)
- Transfer Size (Overall, Read, Write)

Pool metadata

The following key metadata is collected for pool performance:

- Max Write Cache Fullness
- Write Cache Fullness

Port metadata

The following key metadata is collected for port performance:

- Data Rate (Receive, Send, Total)

- I/O Rate (Receive, Send, Total)

The following I/O rate metadata is collected for port performance:

- Port-to-Disk I/O Rate (Receive, Send, Total)
- Port-to-Host I/O Rate (Receive, Send, Total)
- Port-to-Local Node I/O Rate (Receive, Send, Total)
- Port-to-Remote Node I/O Rate (Receive, Send, Total)

The following data rate metadata is collected for port performance:

- Port-to-Disk Data Rate (Receive, Send, Total)
- Port-to-Host Data Rate (Receive, Send, Total)
- Port-to-Local Node Data Rate (Receive, Send, Total)
- Port-to-Remote Node Data Rate (Receive, Send, Total)

The following response time metadata is collected for port performance:

- Port-to-Local Node Response Time (Overall, Receive, Send)
- Port-to-Remote Node Response Time (Overall, Receive, Send)

The following error rate metadata is collected for port performance:

- Frame Errors (CRC Error Rate)
- Link Errors (Invalid Link Transmission Rate, Invalid Transmission Word Rate, Link Failure Rate, Primitive Sequence Protocol Error Rate, Loss of Signal Rate, Loss of Sync Rate)
- Port Congestion Index
- Port Protocol Errors (Port Send Delay I/O Percentage, Port Send Delay Time, Zero Buffer Credit Percentage, Zero Buffer Credit Timer)

The following miscellaneous metadata is collected for port performance:

- Port-to-Local Node Queue Time (Overall, Receive, Send)
- Port-to-Remote Node Queue Time (Overall, Receive, Send)

Node metadata

The following metadata is collected for node performance:

- Cache Fullness (Max Read Cache Fullness, Max Write Cache Fullness, Read Cache Fullness, Write Cache Fullness)
- Compression CPU Utilization (Core 1 to Core 8)
- CPU Utilization (Compression CPU, System CPU)
- Garbage Collection (Data Movement Rate, Data Rewrite Rate, Extent Collection Rate, New Address Write Rate, Reclaimable Capacity, Recovered Capacity Rate)
- Node Utilization by Node
- System CPU Utilization (Core 1 to Core 8)

Performance metadata for DS8000

View the lists of performance metadata for DS8000 storage systems.

Overview

The performance metadata is divided into the following categories:

- [“Volume metadata” on page 39](#)
- [“Disk metadata” on page 39](#)

- [“Port metadata” on page 39](#)

Volume metadata

The following key metadata is collected for volume performance:

- Cache Holding Time
- Data Rate (Read, Total, Write)
- Overall I/O Rate (Read, Total, Write)
- Pool Activity Score
- Response Time (Overall, Read, Write)
- Volume Utilization
- Write-Cache Delay Percentage

The following I/O rate metadata is collected for volume performance:

- Average Transfer Rate (Cache-to-Disk, Disk-to-Cache)
- High Performance FICON[®] (Read, Total, Write)
- Normal I/O Rate (Read, Total, Write)
- PPRC Transfer Rate
- Record Mode Read I/O Rate
- Sequential I/O Rate (Read, Total, Write)
- Write-Cache Delay I/O Rate

The following cache hit metadata is collected for volume performance:

- Normal I/O Cache Hits (Read, Total, Write)
- Overall I/O Cache Hits (Read, Total, Write)
- Record Mode Read Cache Hit Percentage
- Sequential I/O Cache Hits (Read, Total, Write)

The following average miscellaneous metadata is collected for volume performance:

- Average Transfer Size (Overall, Read, Write)
- HPF I/O Percentage
- Sequential I/O Percentage

Disk metadata

The following key metadata is collected for disk performance:

- Data Rate (Read, Total, Write)
- Disk Utilization Percentage
- I/O Rate (Read, Total, Write)
- Response Time (Overall, Read, Write)

The following transfer size metadata is collected for disk performance:

- Average Transfer Size (Overall, Read, Write)

Port metadata

The following key metadata is collected for port performance:

- Bandwidth (Overall, Receive, Send)
- Data Rate (Receive, Send, Total)

- I/O Rate (Receive, Send, Write)
- Port Utilization (Overall, Receive, Send)
- Response Time (Overall, Receive, Send)

The following I/O rate metadata is collected for port performance:

- FCP I/O Rate (Receive, Send, Total)
- FICON I/O Rate (Receive, Send, Total)

The following data rate metadata is collected for port performance:

- FCP Data Rate (Receive, Send, Total)
- FICON Data Rate (Receive, Send, Total)

The following response time metadata is collected for port performance:

- FCP Response Time (Overall, Send, Total)
- FICON Response Time (Overall, Send, Total)

The following error rate metadata is collected for port performance:

- Frame Errors (CRC Errors, Error Frame, Invalid Relative Offset Rate)
- Link Errors (Invalid Link Transmission, Invalid Transmission Words, Link Failures, Primitive Sequence Protocol Error Rate, Sequence Timeouts, Signal Loss, Sync Loss)
- Overload Errors (Extreme I/O Concurrency Percentage, I/O Busy Percentage, I/O Busy Rate, I/O Overrun Percentage, I/O Overrun Rate)
- Port Protocol Errors (Credit Recovery Link Resets, Duplicate Frames, Link Reset Received, Link Reset Transmitted, Out of Order ACK, Out of Order Data)
- Zero Receive Buffer Credit Percentage
- Zero Send Buffer Credit Percentage

The following remote mirror metadata is collected for port performance:

- PPRC Data Rate (Receive, Send, Total)
- PPRC I/O Rate (Receive, Send, Total)
- PPRC Response Time (Receive, Send)

The following transfer size metadata is collected for port performance:

- Average Transfer Size (Overall, Receive, Send)

Performance metadata for XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R

View the lists of performance metadata for XIV systems, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R.

Overview

The performance metadata is divided into the following categories:

- [“Volume metadata” on page 40](#)
- [“Port metadata” on page 41](#)

Volume metadata

The following key metadata is collected for volume performance:

- Data Rate (Read, Total, Write)

- Overall I/O Rate (Read, Total, Write)
- Response Time (Overall, Read, Write)

The following cache hit metadata is collected for volume performance:

- Data Cache Hits (Overall, Read, Write)
- Overall I/O Cache Hits (Read, Total, Write)
- SSD Read Cache Hits (I/O, Data)

The following response time metadata is collected for volume performance:

- Cache Hit Response Time (Overall, Read, Write)
- Cache Miss Response Time (Overall, Read, Write)
- Response Time by Transfer Size (Small, Medium, Large, Very Large)
- SSD Read Cache Hit Response Time

The following miscellaneous metadata is collected for volume performance:

- Average Transfer Size (Overall, Read, Write)
- Data Transfer Size (Small, Medium, Large, Very Large)
- I/O Transfer Size (Small, Medium, Large, Very Large)
- Pool Activity Score
- Volume Utilization

Port metadata

The following I/O rate metadata is collected for port performance:

- I/O Rate (Receive, Send, Total)

The following data rate metadata is collected for port performance:

- Data Rate (Receive, Send, Total)

The following response time metadata is collected for port performance:

- Response Time (Receive, Send, Total)

The following miscellaneous metadata is collected for port performance:

- Bandwidth (Overall, Receive, Send)

Performance metadata for IBM Spectrum Scale

View the lists of performance metadata for IBM Spectrum Scale storage systems.

Overview

The performance metadata that is collected and stored is divided into these categories:

- [“Node metadata” on page 41](#)
- [“File system metadata” on page 42](#)

Node metadata

The following key metadata is collected for IBM Spectrum Scale cluster node performance:

- CPU Utilization (User, System, Total)
- Memory Used (Cache and Buffer, Total)
- I/O Rate (Read, Total, Write)

File system metadata

File system metadata is collected for these resources:

- File systems
- IBM Spectrum Scale storage systems

Tip: For a storage system, the metadata consists of summary values for all the file systems on the storage system.

The following file system performance metadata is collected:

- Data Rate (Read, Total, Write)
- I/O Rate (Read, Total, Write)
- Maximum Data Rate (Read, Total, Write)
- Maximum I/O Rate (Read, Total, Write)
- Response Time (Overall, Read, Write)

Performance metadata for Dell EMC Unity, VMAX, and VNX

View the lists of performance metadata for Unity, VMAX, and VNX storage systems.

Overview

Performance metadata is collected and stored for the following storage systems:

- [“Unity storage systems” on page 42](#)
- [“VMAX storage systems” on page 43](#)
- [“VNX storage systems” on page 44](#)

Restriction: For Dell EMC, performance metadata is available only for block storage systems.

Unity storage systems

The following performance metadata is collected and stored for Unity resources:

- [“Volume metadata for Unity” on page 42](#)
- [“Disk metadata for Unity” on page 43](#)
- [“Port metadata for Unity” on page 43](#)
- [“Node metadata for Unity” on page 43](#)
- [“File system metadata for Unity” on page 43](#)

Volume metadata for Unity

The following key metadata is collected for volume performance on Unity storage systems:

- Overall I/O Rate (Read, Write, Total)
- Transfer Size (Read, Write, Overall)
- Data Rate (Total)

The following cache hit metadata is collected for volume performance on Unity storage systems:

- Overall I/O Cache Hits (Read, Write, Total)

The following data rate metadata is collected for volume performance on Unity storage systems:

- Data Rate (Read, Write)

The following response time metadata is collected for volume performance on Unity storage systems:

- Response Time (Overall)

The following miscellaneous metadata is collected for volume performance on Unity storage systems:

- Transfer Size (Read, Write)

Disk metadata for Unity

The following key metadata is collected for disk performance on Unity storage systems:

- Overall I/O Rate (Read, Write, Total)
- Response Time (Overall)
- Data Rate (Read, Write, Total)
- Transfer Size (Read, Write, Overall)

Port metadata for Unity

The following key metadata is collected for port performance on Unity storage systems:

- Overall I/O Rate (Total)
- Data Rate (Send, Receive, Total)
- Transfer Size (Send, Receive, Overall)

The following I/O rate metadata is collected for port performance on Unity storage systems:

- I/O Rate (Send, Receive)

The following data rate metadata is collected for port performance on Unity storage systems:

- Data Rate (Send, Receive)

Node metadata for Unity

The following key metadata is collected for node performance on Unity storage systems:

- CPU Utilization (System CPU)

File system metadata for Unity

The following key metadata is collected for file system performance on Unity storage systems:

- Overall I/O Rate (Read, Write, Total)
- Data Rate (Read, Write, Total)
- Transfer Size (Read, Write, Overall)

VMAX storage systems

The following performance metadata is collected and stored for VMAX storage systems:

- [“Volume metadata for VMAX” on page 43](#)
- [“Disk metadata for VMAX” on page 44](#)
- [“Port metadata for VMAX” on page 44](#)

Volume metadata for VMAX

The following key metadata is collected for volume performance on VMAX storage systems:

- Data Rate (Read, Total, Write)
- Normal I/O Rate (Read, Total, Write)
- Overall I/O Rate (Read, Total, Write)
- Response Time (Overall, Read, Write)

- Sequential I/O Rate (Read, Total, Write)
- Volume Utilization

The following cache hit metadata is collected for volume performance on VMAX storage systems:

- Normal I/O Cache Hits (Read, Total, Write)
- Overall I/O Cache Hits (Read, Total, Write)
- Sequential I/O Cache Hits (Read, Total, Write)

The following transfer size metadata is collected for volume performance on VMAX storage systems:

- Average Transfer Size (Overall, Read, Write)

Disk metadata for VMAX

The following key metadata is collected for disk performance on VMAX block storage systems:

- Data Rate (Read, Total, Write)
- Disk Utilization Percentage
- I/O Rate (Read, Total, Write)
- Response Time (Overall)

The following transfer size metadata is collected for disk performance on VMAX block storage systems:

- Transfer Size (Overall, Read, Write)

Port metadata for VMAX

The following key metadata is collected for port performance on VMAX storage systems:

- Data Rate (Total)
- I/O Rate (Total)
- Response Time (Overall)
- Transfer Size (Overall)

VNX storage systems

The following performance metadata is collected and stored for VNX storage systems:

- [Volume metadata](#)
- [Disk metadata](#)
- [Port metadata](#)

Volume metadata for VNX

The following key metadata is collected for volume performance on VNX storage systems:

- Data Rate (Read, Total, Write)
- Overall I/O Rate (Read, Total, Write)

The following transfer size metadata is collected for volume performance on VNX storage systems:

- Average Transfer Size (Overall, Read, Write)

Disk metadata for VNX

The following key metadata is collected for disk performance on VNX storage systems:

- Data Rate (Read, Total, Write)
- I/O Rate (Read, Total, Write)

The following transfer size metadata is collected for disk performance on VNX storage systems:

- Transfer Size (Overall, Read, Write)

Port metadata for VNX

The following key metadata is collected for port performance on VNX storage systems:

- Data Rate (Total)
- I/O Rate (Total)
- Transfer Size (Overall)

Performance metadata for Hitachi VSP storage systems

View the lists of performance metadata for Hitachi VSP F and G Series storage systems/>>.

The following performance metadata is collected and stored for Hitachi resources:

- [“Volume metadata for Hitachi” on page 45](#)
- [“Port metadata for Hitachi” on page 45](#)
- [“Node metadata for Hitachi” on page 45](#)

Volume metadata for Hitachi

The following key metadata is collected for volume performance on Hitachi storage systems:

- Overall I/O Rate (Read, Write, Total)
- Data Rate (Read, Write, Total)
- Response Time (Read, Write, Overall)
- Transfer Size (Read, Write, Overall)

The following cache hit metadata is collected for volume, pools, and storage systems performance on Hitachi storage systems:

- Overall I/O Cache Hits (Read, Write)

Port metadata for Hitachi

The following key metadata is collected for port performance on Hitachi storage systems:

- Overall I/O Rate (Total)
- Data Rate (Total)
- Response Time (Overall)
- Transfer Size (Overall)

Node metadata for Hitachi

The following key metadata is collected for node performance on Hitachi storage systems:

- CPU Utilization (System CPU)

Performance metadata for NetApp storage systems

View the lists of performance metadata for NetApp storage systems running ONTAP 9.

The following performance metadata is collected and stored for NetApp resources:

- [“Volume metadata for NetApp” on page 46](#)
- [“Disk metadata for NetApp” on page 46](#)

- [“Port metadata for NetApp” on page 46](#)
- [“Node metadata for NetApp” on page 46](#)
- [“File system metadata for NetApp” on page 47](#)

Volume metadata for NetApp

The following key metadata is collected for volume performance on NetApp storage systems:

- Overall I/O Rate (Read, Write, Total)
- Transfer Size (Overall)
- Data Rate (Total)

The following cache hit metadata is collected for volume performance on NetApp storage systems:

- Overall I/O Cache Hits (Read, Total)

The following data rate metadata is collected for volume performance on NetApp storage systems:

- Data Rate (Read, Write)

The following response time metadata is collected for volume performance on NetApp storage systems:

- Response Time (Read, Write, Overall)

The following miscellaneous metadata is collected for volume performance on NetApp storage systems:

- Transfer Size (Read, Write)

Disk metadata for NetApp

The following key metadata is collected for disk performance on NetApp storage systems:

- Overall I/O Rate (Read, Write, Total)
- Response Time (Overall)
- Data Rate (Read, Write, Total)
- Transfer Size (Read, Write, Overall)

The following response time metadata is collected for disk performance on NetApp storage systems:

- Response Time (Read, Write)

Port metadata for NetApp

The following key metadata is collected for port performance on NetApp storage systems:

- Overall I/O Rate (Total)
- Response Time (Send, Receive, Overall)
- Data Rate (Total)
- Transfer Size (Send, Receive, Overall)

The following I/O rate metadata is collected for port performance on NetApp storage systems:

- I/O Rate (Send, Receive)

The following data rate metadata is collected for port performance on NetApp storage systems:

- Data Rate (Send, Receive)

Node metadata for NetApp

The following key metadata is collected for node performance on NetApp storage systems:

- CPU Utilization (System CPU)

File system metadata for NetApp

The following key metadata is collected for file system performance on NetApp storage systems:

- Overall I/O Rate (Read, Write, Total)
- Response Time (Overall)
- Data Rate (Read, Write, Total)
- Transfer Size (Read, Write, Overall)

The following response time metadata is collected for volume performance on NetApp storage systems:

- Response Time (Read, Write)

Performance metadata for Pure Storage systems

View the list of performance metadata for Pure FlashArray//M and FlashArray//X.

Volume metadata for Pure

The following key metadata is collected for volume performance on Pure FlashArray//M and FlashArray//X.

- Overall I/O Rate (Read, Write, Total)
- Data Rate (Read, Write, Total)
- Response Time (Read, Write, Overall)

Performance metadata for switches

View the lists of the performance metadata for switches.

Overview

The performance metadata is divided into the following categories:

- [“Key metadata” on page 47](#)
- [“I/O rates metadata” on page 47](#)
- [“Error rates metadata” on page 48](#)
- [“Port Protocol Errors metadata” on page 48](#)
- [“Link Errors metadata” on page 48](#)
- [“Miscellaneous metadata” on page 48](#)

Key metadata

The following key metadata is collected for switch performance:

- Data Rate (Receive, Sent, Total)
- Bandwidth (Overall, Receive, Send)
- Other (Total Port Error Rate)

I/O rates metadata

The following I/O rate metadata is collected for switch performance:

- Frame Rate (Receive, Send, Total)

Error rates metadata

The following error rate metadata is collected for switch performance:

- Bad EOF CRC Error Rate
- CRC Errors
- Discarded Class 3 Frames
- Error Frames
- F_BSY Frames
- F-RJT Frames
- Long Frames
- Short Frames

Port Protocol Errors metadata

The following port protocol error metadata is collected for switch performance:

- Class 3 Receive Timeout Frame Rate
- Class 3 Send Timeout Frame Rate
- Credit Recovery Link Reset Rate
- Discarded Frames
- Link Reset Received
- Link Reset Transmitted
- Port Congestion Index
- Zero Buffer Credit Percentage
- Zero Buffer Credit Rate

Link Errors metadata

The following link error metadata is collected for switch performance:

- Encoding Disparity
- Invalid Link Transmission Rate
- Invalid Transmission Words
- Link Failures
- Primitive Sequence Protocol Errors
- Signal Loss
- Synch Loss

Miscellaneous metadata

The following miscellaneous metadata is collected for switch performance:

- Frame Size (Overall, Receive, Send)
- Other (Link Quality Percentage)

Performance metrics for hosts

View the list of performance metadata for hosts.

Overview

The performance metadata is divided into the following categories:

- [“Drive metadata” on page 49](#)
- [“Host metadata” on page 49](#)

Drive metadata

The following key metadata is collected for drive performance:

- Overall I/O Rate (Read, Write, Total)
- Data Rate (Read, Write, Total)
- Response Time (Read, Write, Total)

Host metadata

The following miscellaneous metadata is collected for host performance:

- Total CPU Utilization (%)
- Total Memory Utilization (%)

Performance metrics for virtual machines

View the list of performance metadata for virtual machines.

Overview

The performance metadata is divided into the following categories:

- [“Drive metadata” on page 49](#)
- [“Node metadata” on page 49](#)
- [“Virtual Machine metadata” on page 49](#)

Drive metadata

The following key metadata is collected for drive performance:

- Overall I/O Rate (Read, Write, Total)
- Data Rate (Read, Write, Total)
- Response Time (Read, Write, Total)

Node metadata

The following miscellaneous metadata is collected for node performance:

- Total CPU Utilization (%)
- Total Memory Utilization (%)

Virtual Machine metadata

The following miscellaneous metadata is collected for virtual machine performance:

- Host VM CPU Ready Utilization (%)
- Host VM Memory Ballooned Utilization (%)

Legal notices

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
224A/101
11400 Burnet Road
Austin, TX 78758
U.S.A*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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. 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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at <https://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <https://www.ibm.com/privacy/details> in the sections entitled “Cookies, Web Beacons and Other Technologies,” and “IBM Cloud services”.

Trademarks

IBM, the IBM logo, and [ibm.com](https://www.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 <https://www.ibm.com/legal/us/en/copytrade.shtml>.

Intel, Intel logo, Intel Xeon, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java[™] and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

Red Hat[®] is a registered trademark of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Index

A

access controls [13](#)
Access, IBM Support [14](#)
aggregation of asset and capacity metadata [11](#)
aggregation of performance metadata [11](#)
asset metadata [9](#)
authorization [13](#)

B

backups [14](#)

C

cloud service infrastructure teams access to metadata [13](#)
configuration metadata [9](#)
connection details data [9](#)

D

data collector [3](#)
datastores metadata [32](#)
Dell EMC
 view performance metrics [42](#)
DevOps access to metadata [13](#)
diagnostic data packages [9](#)
drives metadata [32](#), [33](#)

E

ECuRep [9](#), [14](#)

F

firewall [3](#)

G

gdpr [5](#)

H

HBAs metadata [32](#), [33](#)
Hitachi VSP
 view performance metrics [45](#)
hosts metadata [32](#)

I

IBM Spectrum Scale
 file systems, view performance metrics [41](#)
 nodes, view performance metrics [41](#)
IBM Support [14](#)
IBM Support access to metadata [13](#)

IEPD [9](#)

M

metadata access [13](#)
metadata access authorization [13](#)
metadata access controls [13](#)
metadata aggregation levels [11](#)
metadata backups [14](#)
metadata collected for storage systems [17](#)
metadata collection [5](#)
metadata delivery [5](#)
metadata encryption [5](#)
metadata for quality improvements [14](#)
metadata retention periods [11](#)
metadata storage [5](#)
metadata types [9](#)
minimum installation requirements [3](#)

N

NetApp ONTAP [9](#)
 view performance metrics [45](#)

O

organizational security [5](#)

P

performance metadata [9](#)
performance metrics
 Dell EMC [42](#)
 Hitachi VSP [45](#)
 IBM Spectrum Scale [41](#)
 NetApp ONTAP [9](#) [45](#)
 Pure [47](#)
personal data,
 delete [15](#)
personal information,
 delete [15](#)
physical protection [5](#)
PI,
 delete [15](#)
Pure
 view performance metrics [47](#)

Q

quality of service improvements [14](#)

R

resolve issues [13](#)

S

security certification [3](#)
service interruptions [13](#)
support tickets [9](#)
Support, access [14](#)
Support, Blue Diamond [14](#)
supported operating systems [3](#)
supported storage systems [17](#)

T

technical security [5](#)
ticket data [9](#)
trademarks [53](#)

V

virtual machines metadata [33](#)
VMDKs metadata [32](#)



Product Number: 5725-U02