

IBM Storage Insights

Security Guide



Note:

Before using this information and the product it supports, read the information in [“Legal notices” on page 19](#).

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, 2023.**

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. Security overview.....	1
Chapter 2. What security measures are built in.....	3
Chapter 3. What is the data collector.....	5
Chapter 4. How is the metadata protected.....	7
Chapter 5. What types of metadata are collected.....	11
Chapter 6. How long is the metadata kept.....	13
Chapter 7. Who can access the metadata.....	15
Metadata access controls and authorization.....	15
Metadata access for resolving issues.....	15
IBM Support access for troubleshooting your tickets.....	16
Metadata access for quality improvements.....	16
Data backup and restore.....	16
Requesting the deletion of personal information.....	16
Legal notices.....	19
Privacy policy considerations	20
Trademarks.....	20
Index.....	23

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. Security overview

Learn about the security measures related to deploying a data collector on-premises, processing and storing metadata off-premises, and session timeouts.

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 securely and efficiently send configuration, capacity, performance, and status 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.
- After you log in to the service, the security of your web browser session is important. To protect your session, you're automatically logged out after 2 hours 30 minutes of inactivity. For more security during extended use, the duration of an active login session is limited to approximately 8 hours. When you are logged out, you can log in again and pick up right where you left off.

The timeout durations for a session are set by default and can't be changed.

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, non-subscribers have access to specific capacity and performance metadata only. IBM Support also has read-only access to the set of metadata that they need to troubleshoot and close support tickets.

Tip: In the security documentation, the name IBM Storage Insights is used to refer to both IBM Storage Insights and IBM Storage Insights Pro unless a notable difference exists between the offerings.

To address the security concerns that you might have, the following questions are answered:

- What security measures are built-in?
- 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.

Chapter 2. What security measures are built in to IBM Storage Insights

Key security measures are built in to IBM Storage Insights to help ensure that it's a secure part of your organization.

Security and Privacy by Design (SPbD) at IBM is an agile set of focused security and privacy practices, including threat models, privacy assessments, security testing, and vulnerability management. SPbD@IBM is aligned with the United States National Institute of Standards and Technology (NIST's) [Secure Software Development Framework \(SSDF\)](#), which drive processes that are required across all business units.

Because IBM Storage Insights is a cloud-based service, the security of the connection between it and your storage environment is paramount. The IBM Storage Insights team used SPbD to build in security measures at the start and continues to carry it up through every aspect of the service.

In summary, security wasn't something that was tacked on after the service was developed, but was and is baked into the design and DNA of IBM Storage Insights:


- ISO/IEC 27001/27017/27018/27701 ISM certified
- Communication is one way, encrypted and compressed
- Metadata at rest is AES 256-bit encrypted
- Metadata streamed to IBM Cloud® is 128-bit encrypted
- Only metadata about your storage is collected
- Personal, identity, and application data are never accessed
- HIPAA / Blue Diamond ready
- Dedicated vulnerability tracking and threat response team (IBM PSIRT) *
- EU-US Privacy Shield and Swiss-US Privacy Shield Framework
- Meets the requirements of GDPR

*** About the IBM Product Security Incident Response Team (PSIRT):** This global team manages the receipt, investigation, and internal coordination of security vulnerability information related to IBM Storage Insights. [IBM PSIRT](#) is the centralized process through which IBM customers, security researchers, industry groups, government organizations, or vendors report potential IBM security vulnerabilities. IBM is committed to responding to new threats and risks. [IBM's Secure Engineering practices](#) were designed so that IBM can act in a timely fashion to a reported security vulnerability affecting IBM Storage Insights.

[Trusting](#) in the security of IBM Storage Insights is an important factor when organizations consider deploying the service within their environments. Understanding more about the security measures that IBM builds in can help address your concerns and gain the trust that you need to use it with peace of mind.

Chapter 3. 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 data collectors to the well-defined and secure network endpoint <https://insights.ibm.com:443>. Update your firewall rules to allow outbound communication to <https://insights.ibm.com> and to the HTTPS port 443 using the Transmission Control Protocol (TCP).

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 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. For more information about the requirements for data collectors, see the following topics:

- <https://www.ibm.com/docs/en/storage-insights?topic=collectors-installing-data-windows>
- <https://www.ibm.com/docs/en/storage-insights?topic=collectors-installing-data-aix-linux>

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:

Built-in 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.

Data collectors use prepackaged commands and code from IBM Storage Insights to run pre-defined operations only. Remote code loading is not possible.

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.

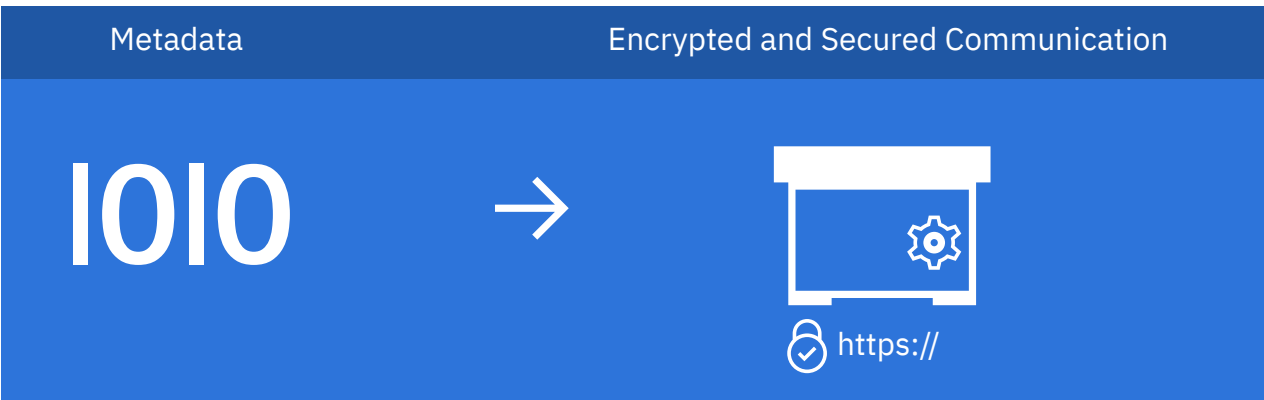
Tip: Because HTTPS connections are used, the data collector can run on any computer that can access the internet over an outbound TCP connection to port 443. Port 443 is the standard port for HTTPS connections.

Chapter 4. How is the metadata protected

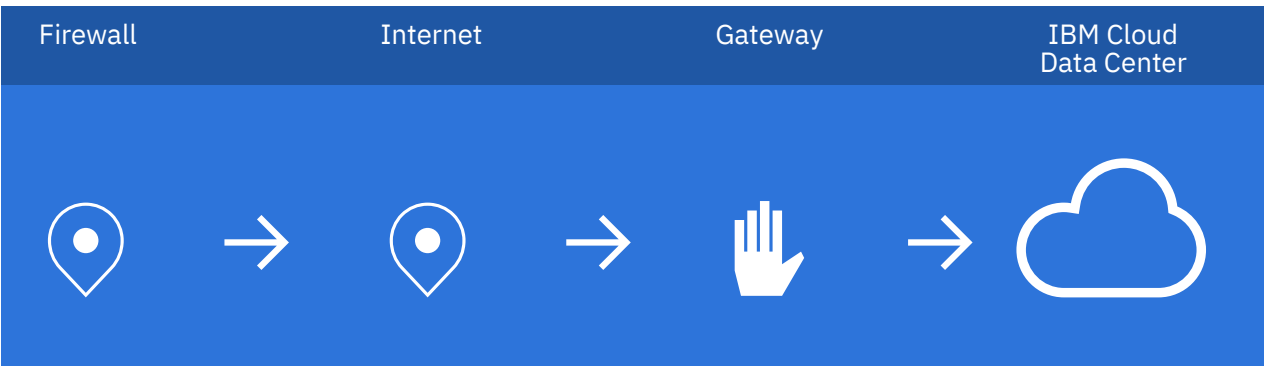
End-to end protection is provided for the metadata that is collected, delivered, and stored for your IBM Storage Insights service in the IBM Cloud data center. This protection includes meeting the requirements of the General Data Protection Regulation (GDPR).

Metadata collection, delivery, and storage in the cloud

To transform the metadata into insights and present them in IBM Storage Insights, 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 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.

After you sign up, you're provided with a host name and port number for your IBM Storage Insights service. To secure the outbound communication between the data collector and IBM Storage Insights at the well-defined and secure network endpoint `https://insights.ibm.com:443`, 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.

Tip: The security of your web browser session is also important. To protect your session, you're automatically logged out after 2 hours 30 minutes of inactivity. For more security during extended use, the duration of an active login session is limited to approximately 8 hours. When you are logged out, you can log in again and pick up right where you left off.

At the IBM Cloud data center

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 gateway in IBM Cloud and that communications are not tampered with. For internal traffic, each customer's instance of 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.

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

IBM Storage Insights is built with a multi-tenant SaaS architecture. Multiple SaaS instances, or tenants, are hosted from a single multi-tenant application that spans the resources of many shared servers and services. Even though any two tenants might share common resources, each tenant does not see the data of other tenants; let alone even knows others exist.

In this multi-tenant SaaS architecture, IBM Storage Insights uses a virtualization technology called "containers". If you are familiar with Docker, containers is the technology behind it. The resulting container consists of just the application and a very small overhead for dependencies. The application within the container is comprised of multiple, independent micro-services based on a functional area. For example, there is one micro-service for the web server and another to process performance data. A collection of all the containers for the various micro-service applications make up the entire multi-tenant IBM Storage Insights server.

To keep track of all the IBM Storage Insights containers, Kubernetes is used as the container management tool. Kubernetes organizes containers into pods that are deployed on nodes in the cluster. Each IBM Storage Insights tenant is containerized within a Kubernetes cluster, which enables scalability, high-availability, and disaster tolerance. The Kubernetes cluster uses enterprise class IBM Cloud security, providing optimal communication and lower front-end latency to IBM Storage Insights containers and services. Additionally, back-end storage and SAN resources utilize the same enterprise class IBM Cloud security.

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

- Crowdstrike EDR and Crowdstrike Prevent to protect against malware
- IBM SOS® to comply with security and regulatory requirements
- IBM Security QRadar® SIEM to store and monitor system and application logs

For more information about IBM Cloud's compliance and certifications, see <https://cloud.ibm.com/docs/overview?topic=overview-security>.

Database security

IBM Storage Insights uses IBM Cloud databases built on Apache Cassandra. It's designed to power real-time applications with high availability and massive scalability. With its NoSQL workloads, a smooth and secured experience is natively integrated into the IBM Cloud. Cassandra database protects against unauthorized access, provides data resiliency, is SOC/ISO certified, and GDPR/HIPAA/PCI DSS compliant.

For more information about Cassandra's compliance and certifications, see <https://cloud.ibm.com/docs/databases-for-cassandra?topic=databases-for-cassandra-security-compliance>.

Organizational security

Access to the infrastructure and instances for 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 5. What types of metadata are collected

Metadata is the information that IBM Storage Insights collects about your storage devices and environment.

Metadata about your storage devices can include, but is not limited to the following information:

- Inventory and configuration metadata such as name, model, firmware, type, and more
- Inventory and configuration metadata for internal components such as volumes, pools, disks, ports, and more
- Capacity metrics such as capacity, usable capacity, used capacity, compression ratios, and more
- Performance metrics such as read and write data rates, I/O rates, response times, and more
- Diagnostic data, system failure logs, maintenance levels, and more support-related information

IBM Storage Insights analyzes this metadata to help you identify problems with your storage before they impact your business. Performance bottlenecks, capacity usage and shortages, loss of connectivity or access to devices, and configuration issues are just a few of the things that metadata can spotlight. To get 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.

Important:

- Use of IBM Storage Insights and the collection and use of metadata is governed by the [IBM Cloud Service agreement](#) and the [IBM Storage Insights Service Description](#).
- The data that is stored on your storage devices is never viewed or accessed by IBM Storage Insights.

IBM Support ticket and diagnostic log packages

When you create tickets in IBM Storage Insights, you provide a name, an email address, and a phone number so that IBM Support can contact you. IBM Storage Insights also collects and uploads the diagnostic data for IBM block storage systems to IBM Enhanced Customer Data Repository (ECuRep) or Blue Diamond Enhanced Secure Support, depending on your configuration.

What is ECuRep

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

When the diagnostic log package is collected from a device, IBM Storage Insights transfers it to IBM Support and ECuRep. To secure the transmission of that data, multiple methods are used, such as HTTPS protocol. For more information, see <https://www.ibm.com/support/pages/enhanced-customer-data-repository-ecurep-send-data-https>.

About encryption: When diagnostic data is transmitted, that data is encrypted. For information about the data encryption that is used for ECuRep, see the following links:

- <https://www.ibm.com/support/pages/ecurep-encryption-information-0>
- <https://www.ibm.com/support/pages/node/6259449>

What is Blue Diamond Enhanced Secure Support

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.

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.

About encryption: In Blue Diamond environments, data at rest is stored on encrypted storage.

Chapter 6. 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.

As metadata about monitored devices is collected, the aggregation level of that metadata changes. For configuration, status, 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:

Table 1. Asset, configuration, and capacity metadata	
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:

Table 2. Performance metadata	
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 7. Who can access the metadata

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

Access to metadata is carefully controlled and governed by the [IBM Cloud Service Agreement](#) and the [IBM Storage Insights Service Description](#).

Key teams can access metadata. IBM Support, Development, DevOps, and cloud infrastructure teams have a level of access that's needed to help ensure that your day-to-day storage operations run smoothly. The wider IBM Storage Insights team has limited access to improve your product experience and help resolve any issues that you might encounter.

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.

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 the related IBM Storage Insights service.

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 or collect the existing log packages.
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.

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

For information about privacy policy considerations, see IBM's Privacy Policy at <https://www.ibm.com/privacy>.

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 [15](#)
Access, IBM Support [16](#)
aggregation of asset and capacity metadata [13](#)
aggregation of performance metadata [13](#)
asset metadata [11](#)
authorization [15](#)

B

backups [16](#)

C

cloud service infrastructure teams access to metadata [15](#)
configuration metadata [11](#)
connection details data [11](#)

D

data collector [5](#)
data encryption [3](#)
DevOps access to metadata [15](#)
diagnostic data packages [11](#)

E

ECuRep [11](#), [16](#)

F

firewall [5](#)

G

gdpr [7](#)
GDPR [3](#)

I

IBM Support [16](#)
IBM Support access to metadata [15](#)
IEPD [11](#)
ISO/IEC [3](#)

M

metadata access [15](#)
metadata access authorization [15](#)
metadata access controls [15](#)
metadata aggregation levels [13](#)
metadata backups [16](#)
metadata collection [7](#)
metadata delivery [7](#)

metadata encryption [7](#)
metadata for quality improvements [16](#)
metadata retention periods [13](#)
metadata storage [7](#)
metadata types [11](#)
minimum installation requirements [5](#)

O

organizational security [7](#)

P

performance metadata [11](#)
personal data,
 delete [16](#)
personal information,
 delete [16](#)
physical protection [7](#)
PI,
 delete [16](#)

Q

quality of service improvements [16](#)

R

resolve issues [15](#)

S

security [3](#)
security certification [5](#)
security measures [3](#)
service interruptions [15](#)
support tickets [11](#)
Support, access [16](#)
Support, Blue Diamond [16](#)
supported operating systems [5](#)

T

technical security [7](#)
ticket data [11](#)
trademarks [20](#)



Product Number: 5725-U02