Cytiva data bridge

Privacy and Security Manual

cytiva.com

Table of Contents

1 Introduction

About this manual

This manual describes the privacy and security considerations of the use of Cytiva data bridge.

Purpose of this manual

This manual describes the expected intended use of data bridge, the privacy and security capabilities included, and how these capabilities are configured.

Scope of this manual

This manual is valid for all data/information shared via data bridge.

Introduction to privacy and security

This manual assumes that the reader understands the concepts of privacy and security. Security protects both system and information from risks to confidentiality, integrity, and availability. Security and privacy work together to help reduce risk to an acceptable level. In healthcare, the privacy, security, and safety must be balanced, relating to the intended use of the product.

The customer is encouraged to use risk management procedures to assess and prioritize privacy, security, and safety risks. Using the risk management, the customer can determine how to best leverage the capabilities provided within the product.

Product description

Data bridge delivers an infrastructure where data can be shared with customers through a cloud database solution. The initial release focuses on raw materials data although the platform is designed to be scalable to incorporate additional data types in the future.

Contact information

For specific privacy and security inquiries, use the contact form found at *cytiva.com/contact*.

Abbreviations

The following terms and abbreviations are used in this manual:

Term/Abbreviation	Definition
2FA	Two-Factor Authentication

Term/Abbreviation	Definition
AWS	Amazon Web Services
DAC	discretionary access control
DRP	disaster recovery plan
DWH	Data Warehouse
ECOA	Electronic Certificate of Analysis. Type of eData. Contains signed off results of QA laboratory analysis for the finished product batch, in the form of a list of measurements and their values.
eData	Electronic data associated with a finished product lot, made up of several data types. eData for a specific lot should generally be treated as an atomic bundle of data, not as collection of independent records.
ETL	Extract, Transform, Load – data movement process starting by extracting from source system, transforming to target form, and loading the transformed data into the target system.
HMI	human-machine interface
IdP	Identity Provider
LGEN	Lot Genealogy. Type of eData. Parent-child tree of subassemblies and raw material batches that went into finished lot manufacturing, down to raw material level. Includes raw material batch provenance (batch ID, vendor, date of purchase, and date of manufacturing)
PLC	programmable logic controller
RBAC	role-based access control
SaaS	Software as a Service
SFDC	Salesforce.com
SOR	System of Record (also known as Source of Truth)

2 Privacy and security environment

Privacy and security in the environment

Data bridge has been designed for an intended use with the following expectations of privacy and security protection:

Data bridge is a single data repository for non-restricted, shareable eData available for Cytiva manufactured products. The design primarily uses Snowflake (SaaS) for both backend and front-end capabilities. The backend is Snowflake (a cloud-based SaaS database) that provides the database engine, access controls, and compute capabilities. The front-end is Snowsight, a web-based SQL client environment, that provides full flexibility to explore and to query the data, with certain data visualization capabilities included.

Snowflake (and specifically, Snowsight and Snowflake API) is, generally, a public-facing system, and as such there are no default restrictions on connectivity over public Internet (however all connections require authentication). All connections to Snowflake use the HTTPS protocol that are SSL-encrypted in flight. On the AWS side, resources inside Cytiva VPC are not public facing and reside on a private subnet. All S3 buckets have public access disabled.

The Snowflake account for data bridge uses AWS as the backend, so communication between Snowflake and AWS happens over the AWS backbone network, not over the public Internet. Connections between a Snowflake account and S3 buckets integrated as external stages (e.g., data bridge xml document storage location) automatically go over a Snowflake-maintained S3 Gateway, if the bucket and Snowflake account are in the same AWS region. Connections between Snowflake and Lambda functions integrated as external functions go over Snowflake PrivateLink and Private API Gateway (configured as part of the infrastructure setup for the data bridge AWS account).

Main functionalities:

- Store eData both in relational database format (for flexibility of use), and in industry-standard data exchange document format (for maintaining atomicity of data bridge.
- Develop data interfaces between data bridge data repository, Bioprod, and LIMS (for reading SOR data into data bridge). Perform scheduled, automatic ETL for all data bridge relevant data from SORs.
- Perform scheduled, automatic ETL for all data bridge relevant data from SORs.
- Share data bridge data with Cytiva customers, while retaining full control over eligibility to receive and continued access to data bridge by specific product and data bridge type (such as eCoA, Lot Genealogy).
- Give Cytiva customers the ability to run free-form queries on top of and plug their own analysis tools to fetch data from the data bridge repository.
- Give Cytiva customers the ability to quickly get document versions of data bridge.

- Streamline and enable future automation of customer onboarding and access management while maintaining control over critical parts of the process.

3 Authentication, authorization, and audit logging

About this chapter

Data bridge includes a broad assortment of capabilities to enable privacy and security. This chapter describes the ability and use of these privacy and security capabilities.

In this chapter

3.1 Access controls

Introduction

The access control on data bridge is used to help control access to customer information on the system. Access control includes user account creation, assigning the privileges, and other features.

Identity provisioning

The provisioning of user accounts requires the steps of account creation, maintenance, and removal of the account when it is no longer needed. A user account is created to be used by a specific individual. This user account is associated with access rights, and is recorded in system security log files.

Salesforce IdP

Cytiva managed Salesforce is used as an identity provider that serves as the shared identity pool and handles user registration and authentication (through SAML 2.0). Authentication through Salesforce IdP, managed by Cytiva, is mandatory to access data bridge, as no other authentication mechanism is enabled for normal user accounts.

Predefined user roles

Predefined roles exist for both users and administrators:

Account	Description
USER_ECOA_READER	Allows reading ECOA data only (including downloading of ECOA XML files).
USER_ALL_READER	Allows reading ECOA, LGEN, and any future data sets (including downloading ECOA XML files, and ECOA+LGEN XML files).
SYSADMIN role	As object owner, grants permissions on objects (such as USAGE and SELECT) to roles covering access for a data bridge type in specific customer schema.
USERADMIN	Owner of all user accounts.

Security Controls

The data bridge application includes a variety of configurable security controls that include:

- Unique user identifiers (user IDs).
- Password complexity and length requirements and controls.

- Controls to revoke access or enable notification after a number of consecutive failed login attempts.
- Two-Factor Authentication or OAuth for access to the services.
- Utilize SSL certificates to secure site URL access.
- Controls to terminate a user session after a period of inactivity.
- Configurable access controls, including to enable or disable accounts.
- User passwords are stored using a salted hash format for authentication to such services.
- Passwords are not transmitted unencrypted.
- Passwords are not logged.
- No defined passwords are set.
- OAuth tokens are encrypted and not transmitted unencrypted.
- Access logs are stored in a secured centralized host to prevent tampering.
- Client-server communication logs are maintained temporarily to facilitate debugging and system monitoring.

Security Policies and Procedures

Vulnerability Scanning

The data bridge application maintains security policies and procedures that include the following administrative and technical safeguards:

The application has gone through vulnerabilities scanning. The purpose of the activity was to identify the security vulnerabilities in a product by evaluating against various cybersecurity techniques. In the scope of the scanning was the Snowflake platform and AWS Console. The penetration test team conducted an internal vulnerability assessment and penetration testing of the *eData v1.0* web application.

The test was performed in accordance with OWASP and covering most of the test cases relevant to web penetration testing and custom test cases based on environment setup. The purpose of this assessment was to verify the effectiveness of the security controls put in place to secure critical or sensitive information.

User authentication

The user authentication step verifies that the user attempting to access the system is indeed the user associated with the specific account.

All users from the customer's side connecting to data bridge are managed through the shared identity pool, shared with other customer facing Cytiva applications, such as the Customer Portal (*cytiva.com*). Administrative users are only managed in Snowflake.

Assigning access rights

Assigning access rights is the administrative process for connecting permissions with user accounts.

Snowflake user accounts

Each user that should have access to data bridge needs a Snowflake account to be provisioned, with an identity attribute identical to the one returned by Salesforce IdP. That provisioning can be part of a manual process, or eventually automated through integration with any SCIM compliant identity and access management system.

Primary user ID (LOGIN_NAME) in Snowflake is equivalent to a user email address, to match how user identity is defined in the identity provider.

All AWS interfaces used by the design (mainly S3 and API gateway for Lambda) require authentication. The authentication is due to Snowflake accounts for data bridge being AWS-based. All authentication and authorization are based around IAM roles defined in the data bridge AWS account, set to trust Snowflake's IAM principles retrieved from Snowflake as part of setting up the integration. There is no need to provision IAM users and manage secret keys.

User access management

All user access management (granting and revoking of roles to customer user accounts) is performed by Cytiva administrators. While it would be technically possible to grant some users from the customer side permission to grant access roles scoped to a single customer to other users (making them *customer administrators*), this is highly not recommended, as there is no way to guardrail certain roles and certain users together. In Snowflake, if a role has permission to grant another role, it can grant this role to any security principal - in this context, for example, to a user account from another customer company.

Self-service permission management for data bridge customers requires an external IAM solution with finer-grained permission scoping to be feasible.

Administrative access accounts

Only relevant for the Cytiva support team, administrative accounts are created directly in Snowflake and have one or more built-in administrative roles assigned (e.g., ACCOUNTADMIN, SYSADMIN, SECURITYADMIN, or USERADMIN). Those accounts do not use the same Salesforce IdP for authentication (for example, it must be possible to use the accounts to fix issues with identity provider integration), but instead use passwords and 2FA codes managed directly in Snowflake.

Authorization for all user accounts follows the DAC (discretionary access control) and RBAC (role-based access control) access control models as implemented in Snowflake. All database objects in all schemas are owned by the **SYSADMIN** role. SYSADMIN role, as object owner, grants permissions on objects (such as USAGE and SELECT) to roles covering access for a data bridge type in specific customer schema.

Customer automation accounts

For customer accounts, the only possible way to authenticate is by going through Salesforce IdP. The authentication is shared with the Cytiva Customer Portal, so if a customer is already logged on to the Customer Portal and authorized to use data bridge, the customer is automatically logged in. Authentication uses SAML2 protocol to exchange verified identities. Both Snowflake and Salesforce are public-facing SaaS systems so exchange of SAML assertions occurs over the public internet.

A customer identity sent as SAML assertion subject (sub) must match LOGIN_NAME in Snowflake. This identity value is the same as the user's email address in their respective corporate domain.

Cytiva administrative accounts

Accounts for Cytiva administrators are native Snowflake authentication mechanisms. Basing administrative access on security integration like SSO risks locking administrators out from the system if anything goes wrong with the integration, including:

- Password authentication
- Second factor authentication

Application authentication

For (Customer automation accounts) automated data ingestion on the customer side, SSO authentication is not suitable because of the requirement of interactive login at each connection. If the customer requests the possibility to connect to their eData in a fully automated way, a separate automation account can be provisioned. This account would have a different authentication mechanism (OAuth or key pair, user + password as a last resort, but this is not recommended).

This account is not able to log in to Snowsight if any access mechanism other than user + password is used and is limited to using ODBC or the native SQL API.

User account lockout

If a user login fails after five consecutive attempts, the user is locked out of their account for a period of time (currently 15 minutes). Once the time elapses, the system automatically clears the lock, and the user can attempt to log in again.

To unlock the user before the time has elapsed, you can reset the timer by using the *ALTER USER* command.

For more information, see *https://docs.snowflake.com/en/user-guide/admin-user-management*.

Authorization – customer accounts

Authorization for all user accounts follows the DAC and RBAC access control models as implemented in Snowflake.

- DAC – database object privilege management:

- All database objects in all schemas are owned by the SYSADMIN role.
- All roles in the account are owned by the USERADMIN role.
- SYSADMIN role, as object owner, grants permissions on objects (such as USAGE and SELECT) to roles covering access for a data bridge type in specific customer schema (e.g., for customer X and customer Y):
  - CUST_X_ECOA_READER
  - CUST_X_LGEN_READER
  - CUST_Y_ECOA_READER
  - CUST_Y_LGEN_READER
- Those DB object access roles are created as part of the customer onboarding process.

- RBAC – management of permissions for users:
  - Access roles described above are granted in turn to roles relating to functional user types. Functional roles can also be granted to each other, forming a privilege hierarchy:
    - CUST_X_ECOA READER granted to CUST_X_USER_ECOA_READER
    - CUST_X_USER_ECOA_READER granted to CUST_X_USER_ALL_READER
    - CUST_X_LG_READER granted to CUST_X_USER_ALL_READER

    Currently, 2 functional roles for customers are included in the design:
    - USER_ECOA_READER – allows reading ECOA data only (including down-loading of ECOA xml files).
    - USER_ALL_READER – allows reading ECOA, LGEN, and any future data sets (including downloading ECOA xml files and ECOA+LGEN xml files).
  - Functional roles are also created as part of the customer onboarding process.
  - Functional roles are granted directly to user accounts depending on applicable user persona (analyst, data scientist, etc.). However, mapping between such personas and Snowflake roles has to be maintained externally (for example, in an external IAM tool), or through a manually managed dictionary table in Snowflake. There is no possibility to directly extend Snowflake role attributes, and the only “free” attribute for custom descriptions is COMMENT.
- Access policies to fine tune the security model.

3.2 Audit logging and accountability controls

Introduction

Privacy and security information logging and control provide accountability through security surveillance, auditable records, and reporting.

Logging

Logging (both for error logging and audit logging) is handled as follows:

- Snowflake ACCOUNT_USAGE views that include LOGIN_HISTORY, QUERY_HISTORY, and other views. This logging mechanism (_HISTORY tables) is provided out of the box by Snowflake and remains available for querying and analysis for 365 days.

For long-term Snowflake log storage, the following process runs on a daily to monthly basis:

- Data from selected history views, including ACCESS_HISTORY, LOGIN_HISTORY, QUERY_HISTORY, and TASK_HISTORY for the period since the last log archival is selected and inserted into a temporary table.
- Temporary table is unloaded (copied to a file format) to a logging stage, connected with S3 log bucket.
- Temporary table is wiped, and a timestamp of the last log entry archived for each history view is persisted for the next log archival task run.

Snowflake history views are read-only for all Snowflake users, including account administrators, which prevents log tampering.

AWS management events are automatically logged to CloudWatch and CloudTrail, according to organization level settings in the BTG AWS organization. Logging of events on document storage S3 buckets needs to be configured, with retention period set to match Snowflake.

Attempts at tampering with CloudWatch and CloudTrail logs are restricted and monitored on the level of the Cytiva AWS organization.

Log storage location (S3 bucket) has a life cycle policy defined to move older logs to a cheaper tier, and automatically delete logs older than defined log storage period.

Log file content and management

Log file content and management is operated by Snowflake. For more information, see *https://docs.snowflake.com/en/developer-guide/logging-tracing/logging-tracing-overview*.

4 Patient privacy consent management

Patient privacy

Data bridge does not handle (create, transfer, or store) patient data, therefore the patient privacy consent is not applicable to data bridge.

5 Information protection

About this chapter

This chapter describes privacy and security operations, and contains guidelines for the preparation of a secure environment for data bridge.

Defense in depth

Security operations are best implemented as part of an overall "defense in depth" information assurance strategy. This strategy is used throughout an information technology system that addresses personnel, physical security, and technology. The layered approach of defense in depth limits the risk that the failure of a single security safeguard allows the system to be compromised.

In this chapter

5.1 Network security

Wireless network security

Radio signals are used in a wireless network communication, therefore wireless devices require special security consideration. Effective techniques and tools exist for improving the security of wireless communication. This section describes the characteristics for wireless connections for data bridge.

Wireless network security is not applicable to data bridge.

5.2 Data storage and encryption

Data encryption

Data is encrypted at rest using vendor managed keys (Snowflake default encryption, SSE-S3).

Data integrity capabilities

Data bridge has capabilities to make sure that the data is not accidentally or maliciously modified.

De-identification capabilities

Data bridge is not a medical device and does not handle (create, transfer, or store) patient data. Therefore data bridge does not contain de-identification (anonymization and pseudonymization) capabilities.

5.3 External connections

Security controls provided by the cloud provider

The data bridge is hosted on the AWS cloud. This section provides an overview of how AWS fulfills the security, privacy, compliance, and risk management requirements of the product.

More information about security provided by AWS is available from the AWS Security Website, including AWS's overview of security processes.

6 System protection

Introduction

This chapter describes the guidelines for how to configure and maintain the product in a way that continuously protects privacy and security.

Protection from malicious attacks

The computing environment is increasingly hostile, and threats continue to grow from denial of service attacks and malicious software, including computer viruses, worms, Trojan horses, and other malware. Vigilant defense on many levels is required to keep the systems free from intrusion by malicious software. The protective features are enabled as part of the third-party hosting service.

Continuous Data Protection (CDP) encompasses a comprehensive set of features that help protect data stored in Snowflake against human error, malicious acts, and software failure. At every stage within the data life cycle, Snowflake enables your data to be accessible and recoverable in the event of accidental or intentional modification, removal, or corruption.

More information can be found in the formal snowflake documentation under Continuous Data Protection: *Continuous Data Protection | Snowflake Documentation*

Server and workstation security

Server and workstation security are operated by Snowflake.

7 Remote access

Remote connection

Remote connection to the product is not applicable.

8 Personal information collected by the product

Personal information

Data bridge is not a medical device and does not handle (create, transfer, or store) patient data. Data bridge does not collect personal information.

Data bridge does require collecting limited personal information as a requirement for an account in Salesforce.

For more information on customer privacy rights and how Cytiva processes personal data, see *Cytiva Privacy Policy*.

9 Disaster recovery considerations

Disaster recovery plan (DRP)

Cytiva data bridge is based on the Snowflake platform that is provided by a third party. As such, Cytiva depends on the Snowflake platform to maintain adequate performance and accessibility to the customer data.

The Snowflake platform has robust processes for dealing with disaster recovery. One of those processes is time travel that enables accessing historical data at any point within a defined period. It serves as a powerful tool for performing the following tasks:

- Restoring data-related objects (tables, schemas, and databases) that might have been accidentally or intentionally deleted.
- Duplicating and backing up data from key points in the past.
- Analyzing data usage/manipulation over specified periods of time.

On data bridge, Snowflake time travel is currently set up to store this data for 90 days, which is the maximum possible period.

For more information, refer to Snowflake documentation about business continuity and data recovery: *Understanding & Using Time Travel | Snowflake Documentation*.

Although Snowflake is using multiple built-in features preventing data loss and securing data recovery, it is recommended that the customer risk assesses the use of data bridge and implements appropriate procedures to deal with any outages.

10 Additional privacy and security considerations

Additional risks

Data bridge has been designed with privacy and security functionality integrated into the core design. However, there exist privacy and security residual risks that must be mitigated when data bridge is integrated into the work environment. This section describes some risks that should be imported into the risk assessment of the deployment of data bridge for proper mitigation.

Risk mitigation

To securely use data bridge, make sure that:

- Accounts are not shared.
- Accounts are deleted (or requested to be deleted) when an associate leaves their company.

11 Product security supplemental documents

Software Bill of Materials (SBOM)

SBOM is available for data bridge upon request. Contact the sales representative for a copy of SBOM.

Page intentionally left blank

Give feedback on this document

Visit cytiva.com/techdocfeedback or scan the QR code.

cytiva.com

For local office contact information, visit cytiva.com/contact

29747242 AA V:1 04/2024