



CONNECT Data Services Service Description

Contents

CONNECT Data Services	4
Document Purpose and Audience.....	4
About CONNECT data services.....	4
Service Overview	8
Usage	9
Service Limitations.....	10
Regional Cloud Availability.....	10
Hardware and Software Requirements.....	10
Data Ownership.....	11
Data Privacy	11
Decommission of the Service.....	11
High Availability, Business Continuity, and Data Protection.....	12
Service Level Commitment.....	12
Additional Services	13

CONNECT Data Services

Last revision: 6/19/2024

Document Purpose and Audience

Document Purpose

This document describes CONNECT data services, including key features and limitations, as well as the operational parameters.

This document must be read in conjunction with the CONNECT Services service description, which describes the common services available for all functional digital services on CONNECT. Any additions or exceptions to the common services are described in this document.

Audience

The audience of this document are IT departments and business decision makers who are investigating whether to leverage AVEVA cloud offers in their own IT landscape.

About CONNECT data services

CONNECT data services provides a cloud-native platform for aggregating, storing, enriching, accessing, and analyzing real-time operations data from historians, edge devices, and more. CONNECT data services makes it easy to aggregate and store data that resides in your process control networks and on devices outside of your corporate network, such as remote or urban assets.

CONNECT data services provides:

- data scientists and business analysts with a central repository to query large, enriched datasets without disrupting the performance of historians used by local operators.
- engineers and field operations with monitoring capabilities, including live trends of real-time data and device health status, for ad hoc investigations into assets.
- engineers and management with a single place to aggregate production, performance, and other manufacturing information from multiple AVEVA MES sites. The collected data provides content for enterprise visualization and production performance reporting.
- developers with a data platform that simplifies programmatically accessing operations data for custom applications.
- data stewards with secure mechanisms for sharing access to data with users outside of your corporate network.

Key Features

- Native data collection from AVEVA PI Servers (including asset framework), AVEVA Historian, Edge Data Stores, AVEVA Adapters, and AVEVA MES.
- Rules can be configured to automatically enrich streams information with metadata.
- Flexibility to ingress data from other data sources with custom-developed applications that use our Open Message Format (OMF) specification.
- A database built for data streams that follow a sequence (such as time, but also other indices like drilling depth) and can be associated with other data streams and static metadata information into an Asset, to convey the operational context of the data.
- A database built for storing different types of industrial events (for example: downtime, excursions, batches, work orders, equipment use, material consumptions, etc.) and referenced data to those events and assets.
- A rich GraphQL API for traversing event and reference data relationships.
- Rules can be configured to automatically create assets from information provided by stream naming patterns.
- The ability to interactively shape data streams and assets with context into data views for easy use by non-expert data users, machine learning algorithms, and business intelligence tools, with specific connectivity to Microsoft Power BI that is out of the box.
- Choices in how to consume data, either within CONNECT data services via Trends, or externally via PowerBI or AI and ML platforms, 3rd party analytics tools, or other cloud applications via REST APIs.
- The ability to monitor assets in real time, view their status and current values, investigate issues with trending, and share trends with colleagues.
- Quick and secure community data sharing mechanism for companies to manage, control, and ensure transparency around the data shared with their external business partners, broadening who can find insights, deliver value from operational data, value-added service enablement, and enabling cross-company collaboration.
- The ability to query change data from internal sources and external partners, which allows the most recent data to be available to users and applications.
- The ability to configure and manage data collection software on a remote device, while monitoring the health status of the software and device.
- Agnostic industrial platform with open REST APIs and message format specification that simplify developing custom applications.

Key Benefits

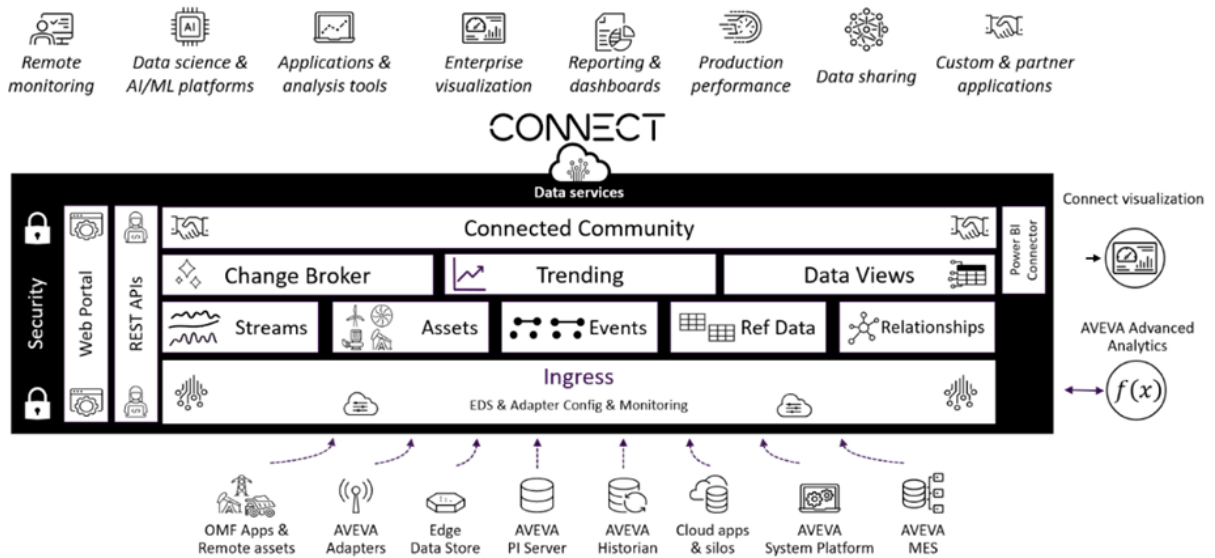
- **Purpose-built for OT, ready for the enterprise:** Native integration to AVEVA PI System, AVEVA Historian, and AVEVA MES means that CONNECT data services provides the easiest path for getting industrial operations data into the cloud in a way that protects the control network and is optimized for OT data, at scale. Data is contextualized, and curated operational data is immediately available to enterprise business applications, IT applications, analytics, and data science tools.

- **Save time on collecting data:** Leverage a portfolio of data collection products that connect to a large catalogue of protocols and do not require you to write any code.
- **Contextualized operational data:** Easily organize your streaming operational data as assets that provide additional context to operational data. Make your operational data more usable and valuable by grouping multiple related data streams together and complementing them with static information such as vendor, location, site, etc. Capture industrial events that provide time context to operational data and reference data that provides context to applications and users.
- **Enable machine learning, data science experiments, and BI reporting without impacting your process control network:** Allow users to query large curated operational datasets in context without needing your process control network or historian to support movement of large volumes of data.
- **Standard, scalable and transparent data sharing that puts you in control of your data:** Easily share operational data with your external ecosystem of trusted business partners, analytics providers, and digital service providers in a secure and standard way. No more managing multiple VPNs, FTPs, emailing large spreadsheets, flying on site to get access to data, or having a new IT conversation every time you need to share data externally with a new business partner.
- **Reduced security risks with sharing data externally:** CONNECT data services is a true multi-tenant system that naturally separates authentication, users, and data into separate tenants. As a result, there is better protection of intellectual property, and your business partners don't need to install hardware and software within your site to get access to your data. It's a more standard and secure way of sharing operational data.
- **One industrial platform for accessing operational data:** With secure cross tenant data sharing, CONNECT data services provides one platform that makes it easy for partners and digital solutions providers to access industrial operational data across many industries without having to install and manage additional hardware and software at their customer sites.
- **Easy integration into existing IT cloud applications and investments:** CONNECT data services has an open modern REST API that makes it easy to interact with the data and integrate with custom applications, IT applications, and other cloud services.
- **Remote operations monitoring at scale:** Monitor and manage your enterprise's remote assets from a centralized and secure location. Provide your field operations staff and subject matter experts with access to data from anywhere in the world while minimizing truck rolls and improving your sustainability goals.
- **Ready-to-use industrial cloud service:** No need to architect, assemble, and manage micro-services in the cloud; use self-service data management built for industrial operations data that is enterprise and IT-ready.
- **Nothing for you to manage or maintain:** Since CONNECT data services is a Software as a Service (SaaS) offering that is developed, operated, and maintained by AVEVA, there is no infrastructure or system upgrades for you to manage or maintain. You can focus on getting value from your operational data that is owned by you and AVEVA does not have access to.

Architecture

CONNECT data services is a multi-tenant cloud architecture with network protection, including mechanisms to mitigate denial-of-service attacks. A combination of logical and physical segmentation architecture enforces isolation between CONNECT data services tenants. CONNECT data services organizes tenants on multiple clustered instances. CONNECT data services is highly reliable and available with automatic failover capabilities.

CONNECT data services has native connectivity to PI Servers (via the PI to Data Hub agent), to AVEVA Historian, Edge Data Store (EDS), AVEVA Adapters, and AVEVA MES (via the Events to CONNECT agent). This native connectivity provides data collection capabilities to hundreds of data source protocols. For data source protocols where there are no compatible AVEVA products, CONNECT data services also accepts OMF messages (a message specification) from any client that has a secure and valid connection, and appropriate authority. For viewing, analyzing, or sharing data outside of CONNECT data services, REST APIs are available that can be queried with a secure and authorized connection(s). All communication channels into and out of CONNECT data services are encrypted.



Service Overview

CONNECT data services is a multi-tenant, cloud-native offering built on top of components from the Microsoft Azure platform and natively integrated with CONNECT, which provides the identity management and authentication mechanism for CONNECT data services. Client applications, on the other hand, are directly authenticated via CONNECT data services. Whether it's granting permissions for CONNECT users and groups or client applications, this is managed via role-based access control within CONNECT data services.

- **Tenancy**

- A CONNECT account can be tied to only one CONNECT data services tenant. Within a CONNECT data services tenant, data-related resources are encapsulated within a namespace, which is located in a specific geographical region.
- The namespace's region is determined by the region of the CONNECT folder, in which CONNECT data services is enabled. This means selection of a region for the CONNECT folder determines what region CONNECT data services namespace will be deployed in when the service is turned "on" in that CONNECT folder. If the CONNECT folder's region is not supported by CONNECT data services, the namespace can be deployed in a different region already supported by CONNECT data services.

- **User Management**

- The integration with CONNECT extends to user management. All users and user groups available for role assignments in CONNECT data services originate from CONNECT.
- Within CONNECT data services, multiple roles are supported. Some roles have pre-defined permissions while other roles are created and customized by the customer. The permissions you assign to customized roles can change over time. Within CONNECT data services, CONNECT users and groups can be assigned one or more roles, and you can change these assignments over time as role permissions are edited and roles are added to and removed from CONNECT data services.
- For authenticating clients that don't interact with CONNECT or the CONNECT data services user portal (for example custom web applications), CONNECT data services provides several ways to securely authenticate a client, including issuing client IDs and secrets, and supporting authorization code flow.

- **Community Data Sharing**

- Community data sharing allows one tenant (such as an Industrial Operator company) to create a community and invite other tenants (such as their service providers and trusted business partners) to join the community. The Tenant Administrator and Community Administrator roles can join a tenant to a community. Each tenant manages their own users in the community.

- Each tenant can share CONNECT data services resources (such as streams) into the community. All members of the community, regardless of their geographic location, have Read access to the shared resource information and underlying data. Customers are responsible for controlling what data they decide to share with their external trusted business partners. If there are data residency regulatory requirements, CONNECT data services gives customers the control to configure their namespaces to opt-out of cross-region data sharing. They can also control the preferred region for a community.
- Customers can create different communities for different business partners (tenants) they want to share operational data with.



For more information, see the [CONNECT data services documentation](#).

Usage

CONNECT accounts that subscribe to CONNECT data services will be charged for its use based on the overall consumption of data from the CONNECT data services tenant. Specifically, the number of Total Streams Accessed for the CONNECT data services tenant is counted and reported to CONNECT each day at midnight Coordinated Universal Time (UTC).

The **Total Streams Accessed** count is the number of unique data streams that have been accessed across the tenant's various namespaces and communities during a calendar day as measured in Coordinated Universal Time (UTC).

- If the same stream is accessed multiple times from a tenant's namespace in one day, it is counted as one (1) Total Streams Accessed.
- If the same shared stream is accessed multiple times from different communities in one day, it is counted as one (1) Total Streams Accessed per community that it was accessed from.
- If the same shared stream is accessed multiple times from the same community in one day, it is counted as one (1) Total Streams Accessed.

- Shared streams that are accessed from the community by another CONNECT data services tenant that you have securely provided access to do not count toward your CONNECT data services tenant’s Total Streams Accessed metric.

Service Limitations

CONNECT data services has the following limitations:

- The default maximum for number of namespaces per CONNECT data services tenant is limited to five (5). To increase this maximum for a particular tenant, customers must communicate this request to AVEVA.
- The CONNECT data services portal supports only the English language.
- AVEVA retains the right to suspend inactive subscriptions to CONNECT data services that have been procured by the customer or end user under a no-fee subscription to CONNECT data services.

Regional Cloud Availability

CONNECT data services is accessed via the public Internet using HTTPS/TLS (a secure transport mechanism). The web applications can be accessed via any supported web browser.

CONNECT data services is available for deployment in the following Microsoft Azure public cloud regions:

- Americas - US West - California
- Asia-Pacific - Australia East - New South Wales
- Europe North - Ireland

Hardware and Software Requirements

CONNECT data services is executed through application streaming technology therefore, client hardware requirements are minimal. Client software requirements are given below.

Component	Minimum/Recommended
Web browser	HTML5 compatible browser, including the latest versions of Google Chrome, Mozilla Firefox, and Microsoft Edge.

Data Ownership

With CONNECT data services, AVEVA has designed capabilities that demonstrate value at scale and complement the way customers collect, use, manage, and share their operational data. The data a customer collects and generates in CONNECT data services is exclusively the customer's responsibility to manage and share. It is simple for a customer to manage the selection of data and leverage communities to share the data with external partners, service providers, analytics providers, and other vendors.

Simply stated, CONNECT data services is built with this core principle in mind: **Our customers own their data**. Customers exclusively own the data uploaded, collected, and created as a computational result through features and functions in CONNECT data services. CONNECT data services communities help you maintain control of and share your data with external trusted business partners, rather than make copies of your data in external partners' systems and tenants.

Data Privacy

AVEVA is committed to the privacy of your data and does not examine, manipulate, or share customer data. Our primary goal is to provide you with a secure platform you can operate, while giving you the freedom and confidence to do so without our examination or intervention. Customers maintain their data privacy as we are data neutral and data unaware.

It must be noted that while AVEVA provides software services to customers that eases control of and streamlines providing access to their data, AVEVA does not enable or enforce the data sharing policies agreed to between corporate entities.

Specific AVEVA DevOps roles have controlled access to backend servers and logs to view some metadata in CONNECT data services as required to operate the service. AVEVA enforces strict personnel surety, procedural, and technical controls over privileged DevOps access to customer data in production systems.

Decommission of the Service

Upon request and confirmation from the customer to decommission CONNECT data services, AVEVA will follow a process for the decommissioning and destruction of data to include the deletion of all files and data held within the service:

- CONNECT data services instances and the data stored within will be deleted.
- A backup of the data store may be provided (for an additional fee) upon request from the customer within **60 days** of the request. If the customer's subscription to CONNECT data services has expired or has been terminated, the request for a backup of the data store must be delivered to AVEVA within 60 days of the expiration or termination.

High Availability, Business Continuity, and Data Protection

To ensure high availability, business continuity, and data protection, CONNECT data services follows the timeline given below.

- **Database Storage**

Different backup mechanisms are used depending upon the nature of data stored:

- For sequential and asset data, customized Azure Blob Storage mechanisms use blob snapshots for incremental backups that occur every **24 hours**. Full backups of this data occur every **seven days**.
- For configuration data (for example data views, asset rules, data collection configurations), Azure SQL Point in Time Restore is used.

All backup data is stored in the same cloud region as the cloud service. All data is replicated across multiple data centers within the same region.

All backup data is retained for **90 days**. After this retention period, AVEVA will delete the data. This retention policy is not practiced for free trials and evaluation licenses.

- **Disaster Recovery**

In the event of a service failure, AVEVA initiates a recovery process in accordance with RPO and RTO objectives detailed below.

Cloud Service	Recovery Point Objective (RPO)
CONNECT data services	24 hours

Cloud Service	Recovery Time Objective (RTO)
CONNECT data services	24 hours

- **Redundancy**

CONNECT data services is a highly reliable and available multi-tenant cloud architecture that implements Microsoft Azure Service Fabric automatic failover capabilities. CONNECT data services redundancy mechanisms provided by Microsoft Azure Service Fabric are regularly exercised. For example, CONNECT data services clusters exercise failover during automatic weekly updates.

Service Level Commitment

AVEVA Cloud Services are governed by the AVEVA General Terms and Conditions.

The AVEVA Cloud Service Level Commitment is a supporting document that describes the service level commitment for all available AVEVA Cloud Services.

Both documents are available on the AVEVA web site at <https://www.aveva.com/en/legal>.

Additional Services

AVEVA offers an extensive collection of Customer Success Accelerators, well-defined, outcome-based services that are designed to ensure you realize the maximum benefit from your investment in our software through all the lifecycle stages of your software application.

For more details, visit the Customer Success Accelerators site at <https://www.aveva.com/en/support/customer-first/success-accelerators/>.