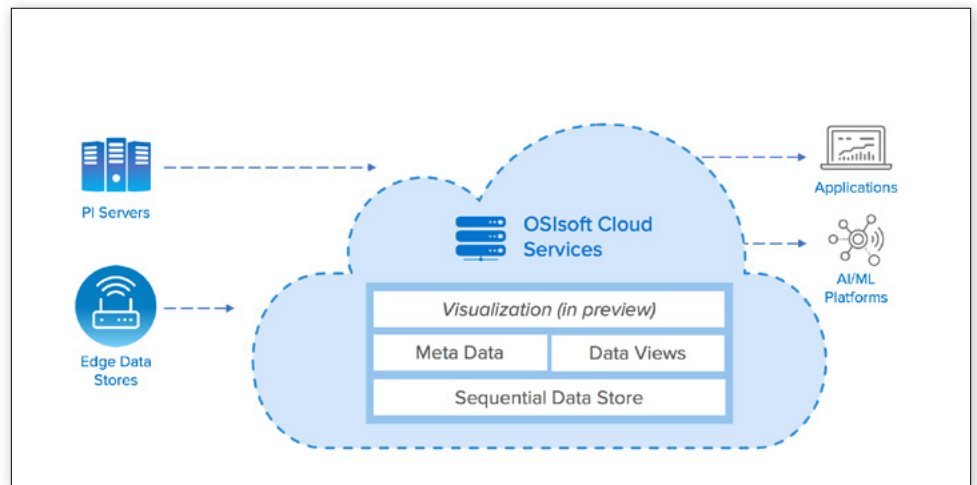


AT A GLANCE

OSIsoft Cloud Services (OCS) is a database platform as a service (dbPaaS) designed for real-time operations data. With OCS, you get turnkey connectivity to data created within the primary control network and at the edge of your industrial network, while you also gain the ability to add contextual information to your time-based operations data. OCS is managed and operated by OSIsoft, so you avoid the time and complexity of database maintenance.

BENEFITS

- Create a single source of operations data for all queries
- Run ad hoc scenarios with no impact to critical operations
- Connect operations data to the cloud for easy access across your organization and with partners
- Offload database maintenance and updating to OSIsoft experts
- Pay based on the data consumed



OSIsoft Cloud Services (OCS) is a cloud-native, multi-tenant data management solution built for accessing and sharing real-time operations data.

Flexible, Secure, and Easy Access to Operations Data

Increasingly, industrial operations data is a critical foundation for achieving business insight, reducing costs, improving processes, and driving new revenue opportunities. As more eyes consume this data, the value of it grows. But sharing operations data can be a challenge. The sources of the data uses proprietary protocols. The data resides across many geographical sites. The data is transmitted within various networks (e.g. process control, corporate, or cellular). The same data is needed by multiple applications, systems, and users in different parts of the organization. And these are just some of the challenges.

OCS provides users inside and outside your organization with flexible, secure, and easy access to relevant operations data. From data collection to data access, to data delivery through a REST API, OCS provides users with an easy, no-code way to prepare and share access to data. To authenticate a user, Azure Active Directory, Microsoft accounts, and Google accounts are supported. To control what is accessed, permissions can be granted to users and roles for one or more data streams. To manage where data is stored and accessed, OCS namespaces (a group of data-related features) can be deployed in the US West, West Europe, or both. In distributed environments, OCS natively connects to data collected on the edge using OSIsoft's Edge Data Store and data collected within the process control network using PI Server.

Since consumers of the data may be unfamiliar with the intricacies of industrial operations, OCS makes it easy to create self-documenting rules that applies common language to data streams. For example, an operator can use special characters in the data stream name to parse its first segment in the name to indicate the equipment type it represents. Or, a process engineer can configure a view of operations data for a data scientist that is in row-column format with column headers labeled with metadata. As a result:

- Developers don't spend time connecting their app to multiple databases or managing those databases.
- Data scientists can query large datasets on-demand without impacting the systems used by operations.
- IT does not need to manage custom code, VPNs, or external accounts.
- Operations spend less time on preparing data and making it accessible to other users inside and outside their company.

Simplifies Using Cloud with Operations Data by Providing:

Out-of-The-Box Connectivity and Aggregation

A plethora of protocols and legacy equipment exist on the operations floor, and connectivity into the control network is not a given. Connecting a SCADA or DCS system to the cloud will likely require writing or, at the minimum, compiling code. Today, with the PI System's 500+ connectivity options, you can connect to data in your process control network without having to deal with or manage code. And, to get this data to the cloud, the PI Server natively integrates to OCS, which can scale to ingest all of your PI Servers' data. OCS also integrates with EDS, so no code is needed to collect data from remote and rugged environments.

The Use of Microservices without the Assembly

To fully leverage the scale and elasticity of cloud computing for industrial operations, you need an end-to-end data management solution for operations data. One way to achieve this is to assemble cloud platform components (or microservices). Components for ingestion, analytics, storage (hot, warm, and cold), visualization, and orchestration all need to be assembled, configured, and tuned to serve your solution. This effort has a steep learning curve and also requires ongoing maintenance to support growth. With OCS, this assembly and management of microservices is done by OSIsoft, so that you can focus more time on analyzing data streams.

A New Sequential Data Store Designed for Operations Data

While usually indexed by time, operations data can contain other dimensions that occur in sequence (e.g. drilling depth, product lot number, etc.). The sequential data store in OCS is built from the ground up to allow for data streams to have a configurable primary index, as well as secondary indices that enable compound indices. This lets you quickly retrieve data across a variety of parameters for richer analysis. At the same time, users have all the capabilities typically expected of time-series databases, like native acceptance of out-of-order data and future data, and fast retrieval of high-resolution data.

A Straightforward Consumption Model

When paying for services from large cloud vendors, a variety of metrics are used to track usage. Storage volume, ingestion rates, and hourly query rate capacity, just to name a few. It often takes complicated spreadsheets and a lot of assumptions to estimate your usage and costs. With OCS, consumption is determined by a single metric: monthly average of your data streams accessed per day.

FIND OUT MORE

Interested in creating an OSIsoft Cloud Services account?

Go to cloud.osisoft.com

“

...we noticed significant performance increases because OCS is optimized for this large amount of time-series data where that's not necessarily the case for MongoDB.

— John Rogers
Software Engineer,
DERNetSoft

”