




CONNEXION is a SaaS or on-premise platform that gives you the ability to simply and efficiently deploy a new generation of third-party smart energy and operational applications such as analytics, fault detection and diagnostics, auto demand response, predictive forecasting and others.

Over the past decade, enterprise data solutions have revolutionized how organizations manage finances, supply chains, sales and labor, but little if anything has addressed how they use and value operational and energy data in relation to managing their core business. It provides a cost effective, enterprise intelligence solution that enhances an enterprise's profitability and improves their ability to achieve corporate energy, sustainability and operational goals. CONNEXION was developed specifically to:

- ✓ Provide an energy and operational data warehouse and analytics engine for buildings
- ✓ Provide a hardware normalization layer that streamlines the integration into facilities
- ✓ Enable the integration of third-party application solutions
- ✓ Enable the injection of energy data into the corporate business intelligence layer
- ✓ Create an open architecture that facilitates the creation of custom applications
- ✓ Engage the key stakeholders across the facility, energy, and business domains
- ✓ Create a scalable, secure solution with flexible deployment models

CONNEXION Advantages

- ✓ Fully compatible with existing building systems infrastructure
- ✓ Retain ownership of your data
- ✓ Open source data structure (*Haystack*)
- ✓ Support for both current and future applications
- ✓ Integrate corporate IT management systems



Technology

CONNEXION takes an innovative approach by creating the first scalable and secure data management platform. Using CONNEXION, disparate technologies can be integrated, heterogeneous networks can be bridged, legacy and best-in-class technologies can be seamlessly integrated, and stakeholder-centric applications can be rapidly deployed.

The CONNEXION Application Server is a complete Haystack compliant server (<http://project-haystack.org>) built on top of two big data, Apache licensed open source databases:

- ✓ An Apache Cassandra NoSQL fault tolerant, triple redundant database for long term, time series data storage and archiving.
- ✓ A MongoDB NoSQL fault tolerant database for long term, metadata storage.

CONNEXION Software provides a complete suite of services and APIs:

- ✓ Haystack compliant API (*Haystack API*)
- ✓ JSON compliant API (*CX API*)
- ✓ Data ingestion and synchronization supporting Haystack and CSV
- ✓ Data export supporting Haystack, CSV, Excel, Zinc and JSON
- ✓ Weather data ingestion
- ✓ Utility data ingestion
- ✓ Data quality management (*missing, failed and data error*)
- ✓ Reporting engine
- ✓ Alerting engine
- ✓ Project creation and management
- ✓ User creation and management
- ✓ Database management
- ✓ Interface support for other analytic dashboards



Key Features

Flexible Data Model

The Dynamo and BigTable concepts built into Cassandra allow for complex data structures that would be difficult to model in traditional, relational databases. The model works with a wide array of data modeling use cases (*i.e. Netflix Instant Queue streaming, user preferences, application logging, etc.*). Similarly, CMMS and building system data can become complex. This requires unique modeling with the flexibility to change the model dynamically while still ingesting data.

No Single Point of Failure

In a cluster, all nodes are created equal. Data is distributed across the cluster and each node is capable of handling read/write requests. When configured with the proper data replication strategy, individual node failures can be resolved with no downtime.

Linear Scalability

By default, all data stored in a cluster is distributed across the entire cluster. As a result of adding new nodes to a cluster, data becomes more sparsely distributed across the nodes, thus reducing the load that each node is under.

CONNEXION's database schema is based on the Project Haystack schema (www.project-haystack.org). Project Haystack is an open source initiative to streamline working with IoT data. It standardizes the semantic data models and web services with the goal of making it easier to unlock value from the vast quantity of data being generated by smart devices from our homes, buildings, factories, and cities. This tag-based schema (<http://project-haystack.org/tag>) is very compatible with both Cassandra and MongoDB's tag (*key-value pair*) based architecture.

CONNEXION Visualization & Reporting

CONNEXION also provides a complete, open sourced dashboard framework. The dashboard provides user friendly custom views and analytics as a standard feature of CONNEXION. Customers can select more sophisticated dashboard displays from the portfolio of Visualization and Reporting Applications.